



C L O S E D D O O R S E S S I O N S

2 0 2 6

W H I T E P A P E R

From Capacity to Consequence

*Decision-Making, Legitimacy and Collective Action in an Age of
Increasing Complexity*

CDS₂₆

CONTENTS

Executive Summary	2
Foreword	6
What Are The Closed Door Sessions?	9
Methodological Note	12
The CDS 26 Framework	16
The Nature of the Inquiry	19
Session One · Attention, Participation and Community	28
Session Two · Capital, Impact and Implementation	38
Session Three · Technology, Governance and Social Cohesion	49
Cross-Session Findings	61
The Questions We Did Not Finish	71
What Questions Do We Ask Next?	81

Executive Summary

The Closed Door Sessions 2026 (CDS 26) were convened to create a form of inquiry that ordinary institutional settings often struggle to sustain. The purpose was not to produce consensus, issue recommendations, or summarise a sequence of discussions. It was to examine how complex public problems are understood when people from different disciplines, sectors and institutional positions are placed in a room designed for candid, cross-disciplinary inquiry. The White Paper therefore reconstructs the movement of thought within the room. It treats the participants' statements, disagreements, hesitations, reframings and unresolved tensions as evidence. It asks not only what was said, but what the behaviour of the inquiry revealed about the nature of the problems under discussion.

The central proposition emerging from the White Paper is that contemporary societies are not primarily constrained by lack of awareness, capital, expertise or technological capability. They are increasingly constrained by the weakening or insufficiency of the structures that convert those capacities into legitimate, accountable and socially useful action. The inquiry did not begin from this proposition. It emerged through the repeated transformation of the stated questions. What appeared to be a question about attention became a question about civic agency. What appeared to be a question about capital became a question about accountability and public benefit. What appeared to be a question about technology became a question about human judgement, legitimacy and governance.

This pattern matters because it suggests that many contemporary public problems are misdiagnosed when treated as shortages of information, money or innovation. The Sessions repeatedly found that many of the relevant capacities already exist in abundance. Societies possess unprecedented quantities of information, but information does not reliably become shared judgement or sustained action. Capital is available, but capital does not become public benefit simply because it is labelled philanthropic, catalytic, ESG-aligned or impact-oriented. Technological capability is expanding rapidly, but capability does not become progress unless it remains governed by human purposes, trusted data, moral constraint and legitimate authority. The difficulty lies in conversion: the passage from capacity to consequence.

This finding should not be mistaken for a claim of institutional collapse. The evidence does not support that conclusion. Participants identified many possible mediators: education, community, leadership, trusted messengers, patient capital, local legitimacy, measurement, implementation, data provenance, human-in-the-loop design, child safety and public-private collaboration. These mechanisms remain important. What the evidence supports is more precise: these mediating structures are fragmented, contested, unevenly trusted or insufficiently developed for the scale and speed of the problems under discussion. The White Paper therefore does not argue that societies lack solutions altogether. It argues that they lack sufficiently legitimate and accountable means of converting capacities into shared public purpose.

One of the strongest surviving findings is that legitimacy can no longer be assumed from institutional category alone. Governments retain formal authority, but may be slow, mistrusted, politically volatile or technically behind. Private capital can move quickly, but is not automatically legitimate simply because it is capable. Technology companies possess expertise and infrastructure, but their incentives do not necessarily align with public good. Communities may carry trust, but often lack scale. International institutions may convene, but may struggle to enforce. Participants did not converge on a single answer to this legitimacy problem. What the evidence supports is that legitimacy increasingly has to be earned, assembled and tested through practice, rather than inherited automatically from position.

A second surviving finding is that measurement is both indispensable and dangerous. Across the inquiry, participants returned to the question of how public benefit is known, verified and trusted. Measurement appeared as a condition of accountability in relation to civic engagement, impact capital, ESG, public health, social cohesion, AI and data systems. Yet the room also recognised that measurement can distort the very thing it seeks to validate. Attention metrics may not measure civic judgement. ESG reporting may become paperwork rather than consequence. AI outputs may appear authoritative without being verifiable. Impact metrics may reduce complex social goods to narrow indicators. The inquiry did not resolve what should be measured, or by whom. It clarified that measurement is not a technical afterthought, but a site of moral, political and institutional judgement.

A third finding is that complex problems defeat single-domain expertise without making expertise irrelevant. The room advanced when participants brought specific knowledge from their own fields and then allowed that knowledge to be tested against others. Civic questions became technological, psychological and political. Capital questions became security, development and legitimacy questions. Technology questions became educational, moral, economic and geopolitical questions. The evidence supports the view that expertise is necessary but insufficient when problems move across systems. The value of interdisciplinary inquiry lies not in replacing expertise with generalism, but in exposing the limits of any one explanatory field.

A fourth finding is that crisis remains one of the most reliable mechanisms for alignment, but an inadequate one. Participants repeatedly observed that societies often act when risk becomes immediate, visible or undeniable. Crisis can concentrate attention, mobilise capital and accelerate governance. Yet crisis-led action is late, costly and uneven. The inquiry therefore left open one of its most important questions: whether societies can build forms of foresight, trust, measurement and institutional legitimacy capable of acting before crisis supplies the urgency.

A fifth finding is that the behaviour of the room itself was evidence of the problem it diagnosed. The participants did not simply disagree and converge in predictable ways. They repeatedly moved from first-order questions to deeper questions. They found that proposed solutions in one register often created difficulties in another. Emotional storytelling may move people, but may risk manipulation. Private capital may accelerate delivery, but may raise legitimacy concerns. Measurement may support accountability, but may distort value. Technology may extend access, but may weaken human judgement. These tensions did not indicate failure of the inquiry. They indicated that the questions had reached the level at which values, authority, evidence and consequences intersect.

The White Paper distinguishes between what participants believed, what the evidence supports and what can reasonably be inferred. Participants believed many different things: that education matters, that community matters, that private capital can help where public systems fail, that ESG is inadequate, that technology can be both tool and threat, that child safety may offer a basis for global agreement, that human judgement must be preserved. The evidence does not establish any one of these as a complete answer. It supports the more limited inference that public-good action now depends on mediating structures that are themselves contested. What can reasonably be inferred is that modern societies possess considerable capacities, but increasingly struggle to connect those capacities to legitimacy, accountability and shared human purpose.

The principal uncertainties are therefore not peripheral. They are constitutive of the inquiry. Who has legitimacy to act on behalf of the public good when authority is distributed across governments, markets, communities, investors, companies and international bodies? How does awareness become responsible action without relying on crisis, coercion or manipulation? What should be measured when public benefit is claimed? Can moral constraint survive competitive systems that reward speed, scale, attention, profit or strategic advantage? Can shared facts be rebuilt, or must coordination now be designed for conditions of persistent epistemic fragmentation? How can human judgement be preserved when systems increasingly automate interpretation, allocation and decision-making? How can local legitimacy be reconciled with global scale? What forms of capital can legitimately act where public systems fail? What kinds of institutions are needed to think across boundaries before action is claimed?

These questions are not presented as rhetorical prompts or deferred recommendations. They are the frontier of inquiry revealed by CDS 26. They matter because each governs several others. Legitimacy shapes who may act, measure and govern. Measurement shapes what institutions and markets pursue. Shared facts shape whether disagreement can become deliberation. Human judgement shapes whether technology remains accountable. Local legitimacy shapes whether global action can be trusted. Foresight shapes whether societies can act before crisis forces alignment.

The contribution of CDS 26 lies in making these relationships visible. Its value was not that it produced agreement, but that it created conditions in which disagreement could become analytically useful. The inquiry showed that sophisticated participants from different fields can identify deeper problems when they are not required to defend institutional positions, deliver immediate solutions or perform consensus. It also showed that unresolved questions may be more valuable than premature answers when the problem itself has not yet been adequately named.

This White Paper should therefore be read as a reconstruction of an inquiry rather than a report of conclusions. It does not claim to settle the questions it raises. It does not offer a programme of action. It seeks to clarify the conditions under which action becomes difficult, and to identify the questions that future inquiry must take seriously. For diplomats, policymakers, economists, sociologists, historians, technologists, security specialists, institutional investors and philanthropic leaders, its significance lies in the proposition that contemporary public problems cannot be addressed by capacity alone. They require the deliberate rebuilding of the structures that allow knowledge, capital and technology to become responsible public action.

Foreword

The Closed Door Sessions (the CDS) were created to provide a form of inquiry that has become increasingly difficult to sustain within ordinary institutional life. Many of the problems now facing governments, markets, families, communities and international systems do not belong to a single discipline. They cannot be understood adequately through policy, economics, technology, philanthropy, psychology, diplomacy or security studies alone. They move across those fields, altering shape as they do so. A question that appears to concern public trust may also concern platform design, education, capital incentives and geopolitical vulnerability. A question that appears to concern investment may also concern legitimacy, measurement, local need and public authority. A question that appears to concern technology may also concern human judgement, moral constraint, social cohesion and the limits of governance.

The CDS exists because complex problems require conditions in which those relationships can be examined without being prematurely simplified. Such conditions are rare. Existing institutions are often required to produce decisions, positions, strategies, communiqués or public narratives before the underlying problem has been properly understood. Governments must respond to urgency. Markets must allocate capital. Media must create legibility. Philanthropy must demonstrate impact. Companies must manage risk and reputation. Academic disciplines must preserve methodological boundaries. Each of these functions is necessary, but none naturally creates space for open-ended, cross-disciplinary inquiry in which uncertainty can be held long enough to become analytically useful.

The purpose of The CDS is not to create consensus. Consensus can be valuable where action is required, but it can also conceal unresolved disagreement, flatten complexity or turn inquiry into performance. The more difficult and often more productive task is to understand why disagreement persists. Disagreement may reveal a difference of evidence, a difference of values, a difference of time horizon, a difference of institutional position, or a difference in what participants believe the problem actually is. In this sense, disagreement is not an obstacle to the work of the CDS. It is one of its primary sources of evidence.

The Closed Door Sessions were therefore designed not as panels, presentations or negotiations, but as structured inquiries. Participants were invited not to represent sectors formally, nor to defend predetermined positions, but to think aloud with others whose expertise, incentives and assumptions differed from their own. The value of such a format lies in what happens when an initial question does not remain stable. A discussion may begin in one domain and end in another. A problem may be reframed. A proposed answer may expose a deeper difficulty. A point of convergence may reveal the limits of agreement. A moment of uncertainty may prove more important than a confident conclusion.

This White Paper from CDS 26 reconstructs that inquiry. It does not offer a transcript, a summary of proceedings, or a set of recommendations disguised as findings. Nor does it treat CDS 26 as an event to be reported. Its purpose is analytical: to follow the movement of thought within the room and to identify what became clearer because the inquiry was shared. The document therefore attends not only to what participants said, but to how the conversation behaved: where it converged, where it resisted closure, where it escalated, where participants hesitated, where a question changed form, and where competing explanations remained plausible.

For this reason, the behaviour of the room is treated as evidence. In conventional reports, uncertainty is often tidied away and disagreement reduced to a range of views. Here, uncertainty and disagreement are part of the material. They reveal the structure of the problem. If a room of experienced participants repeatedly moves away from a stated question towards a deeper one, that movement matters. If experts from different fields agree on a diagnosis but diverge over remedies, that divergence matters. If a proposed solution in one discipline creates a difficulty in another, that tension matters. These are not imperfections in the inquiry. They are indications of complexity.

The White Paper also proceeds from the belief that questions can be as important as answers. This is not an argument against action. It is an argument against premature certainty. In public life, badly formed questions can produce inadequate institutions, misleading metrics, misplaced capital and fragile consensus. Better questions do not delay responsibility; they clarify what responsibility requires. They help distinguish what is known from what is assumed, what is contested from what is unresolved, and what can be acted upon from what still requires deeper inquiry.

Better questions do not delay responsibility; they clarify what responsibility requires.

The reader should therefore approach this document not as a conclusion, but as a reconstruction of the intellectual work participants made possible. Its value lies in the discipline of tracing how complex problems were examined by people who did not share a single field, institution or vocabulary, but who were willing to test their assumptions in the presence of others. The ambition is not to settle the questions raised by CDS 26. It is to make them more exact, more useful and more difficult to dismiss.

What Are The Closed Door Sessions?

The Closed Door Sessions are a private, structured forum for cross-sector inquiry. They were created to bring together individuals whose work, decisions and institutional positions sit within different parts of the contemporary global system: capital, technology, science, diplomacy, crisis response, health, sustainability, infrastructure, culture and social behaviour. The purpose of the Sessions is not to produce a communiqué, declaration, policy position or collective commitment. Their purpose is to create the conditions for disciplined, candid and intellectually serious discussion among participants who would not ordinarily examine the same questions from the same table.

The CDS is distinct from a conference, panel discussion, summit, negotiation or policy forum. It is not organised around speeches, public performance, institutional representation or advocacy. Participants are not invited to present prepared positions or to defend sectoral interests. The format is instead designed around inquiry: a common set of questions is placed before small, mixed tables, allowing participants to examine the same problem from different disciplinary, professional and institutional perspectives. The emphasis is on thought, not performance; on the quality of exchange, not the production of immediate agreement.

The Closed Door Sessions 2026 convened forty participants and ten moderators across five tables. Each table comprised eight participants and two moderators. This structure was intended to create a setting small enough for sustained participation, but broad enough to bring different forms of expertise into contact with one another. Each table considered the same three Sessions, with common questions put to all tables. Each Session was allocated thirty-five minutes of discussion, with moderators permitted to bid for additional five-minute extensions where they judged that a question, tension or line of inquiry remained unresolved or particularly productive. The use of multiple tables allowed the inquiry to proceed in parallel. Rather than producing a single conversation shaped by a small number of dominant voices, the format created five independent sites of discussion, each testing the same questions under different interpersonal and disciplinary conditions. Between Sessions, participants rotated seats, and often tables, ensuring that each round of discussion brought different combinations of participants and moderators into conversation with one another.

The Sessions were structured around three domains of inquiry. The first considered attention, influence and civic motivation. The second considered capital, accountability and public benefit. The third considered technology, population and the future of cohesion. These domains were not treated as specialist silos. Participants were drawn from different sectors precisely because the questions were expected to exceed the boundaries of any one field. Capital was not considered only by capital holders; technology was not considered only by technologists; civic motivation was not considered only by political or cultural actors. Each table was composed to allow different forms of knowledge, responsibility and experience to interact.

The cross-sector composition was central to the design. Contemporary public questions increasingly involve interdependent systems. Issues of motivation, trust, capital, technology, health, climate, infrastructure, security and social cohesion cannot be understood adequately through one institutional lens alone. For this reason, the Sessions brought together participants from sovereign, philanthropic, scientific, cultural, financial, diplomatic, technological, humanitarian and intergenerational contexts. The purpose was not to create representation in a formal sense, but to ensure that each question was examined from more than one position of experience and responsibility.

The Chatham House Rule governed the Sessions. This was essential to the format. It allowed participants to speak candidly without attribution and without the constraints that often accompany public forums, recorded panels or institutional negotiations. The White Paper that follows is therefore anonymised. It does not attribute comments to individual participants, nor does it present the Sessions as a record of personal positions. The purpose is to reconstruct the inquiry itself: where questions led, where perspectives converged, where tensions remained, and what forms of evidence and judgement were brought into the room.

Moderators played a substantive role in maintaining the structure of the inquiry. Each table had two moderators, whose task was to hold the question, encourage participation, preserve focus and help distinguish between assertion, evidence, disagreement and unresolved tension. The role of the moderator was not to steer participants towards consensus, nor to impose conclusions. It was to support the conditions in which a serious inquiry could unfold within a limited period of time.

Following each Session, moderators guided a short table debrief. These debriefs were not designed to produce votes, resolutions or unified table positions. Their purpose was to identify the most salient points of convergence, disagreement, uncertainty and further inquiry. They allowed each table to pause, clarify what had been tested, and name the questions or tensions that remained. The debriefs also formed part of the evidential basis for the anonymised synthesis that followed.

The format was designed in this way because complex questions often change shape when examined from more than one discipline or institutional perspective. A single plenary conversation can easily become performative or hierarchical. A conventional panel often privileges presentation over inquiry. A negotiation requires position-taking. A policy forum may seek actionable outputs before the underlying problem has been adequately understood. The CDS was designed to avoid these pressures. By using small tables, common questions, parallel inquiry, cross-sector composition, moderated discussion and protected anonymity, the Sessions created a setting in which participants could test assumptions, disagree, reframe problems and examine uncertainty without being required to reach a collective position.

The White Paper is therefore not a report of conclusions reached by the participants. It is an anonymised reconstruction of the inquiry conducted through the Closed Door Sessions. Its source material includes the common questions, the table discussions, the moderator observations, the table debriefs and the subsequent synthesis of what the parallel inquiry revealed. The value of the format lies in the disciplined comparison of independent conversations held under shared conditions. This allows the White Paper to examine not only what was discussed, but how the same questions behaved when placed before different combinations of expertise, judgement and experience.

Methodological Note

This White Paper was constructed from a structured, anonymised inquiry conducted under the Chatham House Rule. The Rule was adopted to allow participants to speak candidly without attribution, institutional constraint or public positioning. Accordingly, the analysis that follows does not identify individual speakers, attribute views to named participants, or present comments as institutional positions. Participant contributions are treated as evidence of the inquiry, not as statements of organisational policy or personal commitment.

The evidence base for the White Paper consists of four principal forms of material: the common questions placed before each table; the transcripts of table discussions; moderator observations and written notes; and the short structured debriefs conducted after each Session. These materials were considered together. No single comment, table or intervention was treated as sufficient evidence for a finding. The reconstruction instead examined how questions developed across the room, where independent discussions converged, where they diverged, where assumptions were challenged, and where uncertainty or disagreement remained unresolved.

The 2026 Sessions were organised through five independent tables considering the same questions in parallel. This structure produced a different form of evidence from a single plenary discussion. In a single conversation, the direction of inquiry can be shaped heavily by one sequence of interventions, by the authority of particular speakers, or by the social dynamics of the room as a whole. By contrast, multiple tables considering the same questions simultaneously allow repeated patterns to be distinguished from isolated observations. Where similar tensions, reframings or unresolved questions appeared across separate tables, those recurrences were treated as analytically significant. Where different tables approached the same question in different ways, those divergences were also treated as evidence.

Each table was moderated. The role of the moderators was to maintain the structure of the inquiry, preserve focus on the common questions, encourage participation, and support the distinction between assertion, example, disagreement and unresolved tension. Moderators were not asked to produce consensus, steer conclusions or represent the table's views. Their observations and notes were included in the evidence base because they recorded not only substantive points made by participants, but also the behaviour of the discussion: hesitation, escalation, convergence, contradiction, reframing and areas where the table did not reach closure.

Following each Session, moderators guided a short table debrief. These debriefs were not votes, conclusions or collective statements. They were structured pauses designed to identify what had become salient within the discussion: where perspectives appeared to align, where disagreement remained, which questions had changed form, and which uncertainties required further attention. The debriefs therefore served as interpretive evidence, not as formal findings. They helped clarify the movement of the inquiry without converting that movement into consensus. Following each Session, participants rotated seats for the subsequent discussion, and often changed tables entirely. This ensured that each Session brought different combinations of participants and moderators into conversation, reducing the influence of fixed group dynamics and allowing questions to be tested across a wider range of interpersonal and disciplinary contexts.

The White Paper is not a transcript. A transcript preserves sequence and wording; it does not by itself distinguish between a passing observation, a repeated concern, a contested proposition or a structurally significant reframing. The purpose of this document is not to reproduce the discussions in chronological order, but to reconstruct the inquiry they generated. This required comparing contributions across tables, testing candidate interpretations against contradictory evidence, distinguishing participant belief from supported inference, and identifying where the evidence justified an analytical claim.

The White Paper is also not a consensus document. The Sessions were not designed to produce an agreed position, declaration, recommendation or collective view. No finding in this document should be read as the opinion of all participants, or as the settled view of any table. Where the analysis identifies convergence, it refers to patterns that appeared across the evidence base, not to unanimity. Where it identifies disagreement, that disagreement is treated not as a failure of the format but as part of the evidence. In several instances, disagreement clarified the structure of a problem more effectively than agreement would have done.

The reconstruction methodology used in this White Paper proceeded through comparison, challenge and inference. First, the evidence from each table was examined in relation to the common questions. Second, repeated propositions, tensions and reframings were identified. Third, those propositions were tested against contradictory evidence, competing explanations and unresolved uncertainties. Fourth, claims were retained only where they were supported by patterns across the evidence rather than by isolated statements. This method allowed the analysis to distinguish between what participants believed, what the evidence supports, and what can reasonably be inferred.

Participant views are treated as first-order evidence. They show how individuals in the room understood a question, framed a problem, offered examples, challenged assumptions or proposed explanations. Moderator observations are treated as second-order evidence. They show how the discussion moved, where attention intensified, where disagreement persisted, and where participants appeared to reframe the original question. Reconstructed analysis is the White Paper's interpretation of these materials. Reasonable inference refers to conclusions that can be drawn from recurring patterns in the evidence, while remaining conscious of uncertainty and limitation.

The behaviour of the room is treated as evidence because the inquiry was not limited to substantive claims. How participants thought together was part of what the Sessions were designed to reveal. Convergence indicated that different tables or participants had encountered a similar difficulty, though it did not automatically establish truth. Disagreement indicated that competing explanations remained plausible, though it did not automatically indicate failure. Hesitation, reframing and unresolved tension were particularly important because they often marked the point at which a stated question proved too narrow, an assumption became unstable, or a proposed answer created a difficulty in another domain.

The methodology has limitations. The Sessions brought together a selected group of participants; they were not designed to be statistically representative. The evidence reflects the composition of the room, the questions posed, the time available, the dynamics of each table and the judgement of moderators and analysts. Some perspectives may have been more strongly represented than others. Some participants may have shaped discussion more forcefully than others. The use of transcripts and notes allows careful reconstruction, but no reconstruction can capture every nuance of tone, emphasis or interpersonal context. The White Paper therefore does not claim to establish general laws or definitive conclusions.

The conclusions that can reasonably be drawn are interpretive rather than statistical. The White Paper can identify patterns of inquiry, recurring tensions, supported propositions, competing explanations and questions that remained unresolved. It can show where independent tables considering common questions encountered similar difficulties, and where they diverged. It can distinguish between claims that were asserted, claims that were challenged, and claims that survived comparison across the evidence base. It cannot claim that all participants agreed, that the findings are universally generalisable, or that the questions examined have been resolved.

The value of the methodology lies in disciplined reconstruction. By combining common questions, independent tables, moderated discussion, table debriefs, transcripts and moderator observations, the White Paper examines both the content of the Sessions and the movement of the inquiry itself. Its claims should therefore be read as the product of structured interpretation under conditions of anonymity, parallel discussion and cross-sector comparison.

The CDS 26 Framework

CDS 26 inquiry was structured around three Sessions. All five tables considered the same questions in the same sequence. Each table contained participants from multiple sectors, disciplines and institutional perspectives. The sector lenses were used to bring different forms of knowledge to bear on the same question, rather than to divide the inquiry into separate sector-specific discussions.

SESSION ONE Attention, Participation and Community

ORIGINAL QUESTIONS AS POSED

What truly motivates people to care, act, and engage in the civic and social life of their nations?

How can influence, political, cultural, or financial, become a tool for inclusion rather than division?

In societies shaped by large-scale institutions (states, platforms, corporations) versus those driven by grassroots agency, which models sustain long-term civic motivation?

PURPOSE OF THE SESSION

Session One examined attention, influence and civic motivation. It considered how people come to care about public or social questions, how participation is formed or weakened, and how influence may operate either as a source of inclusion or division.

SECTOR LENSES USED

Technology, AI & Frontier Intelligence	The algorithmic architecture shaping attention, influence, and perception.
Global Capital, Sovereign Wealth & Family Offices	How influence is funded, incentivised, or strategically directed.
Healthcare, Longevity & Human Reinvention	Mental health, cognition, and behavioural drivers of engagement.
Diplomacy, Crisis Response & Education for Global Stability	How legitimacy, trust, and civic identity are formed.
Sustainability, Climate Transition & Regenerative Economies	Mobilisation, activism, and collective action dynamics.
Infrastructure, Energy, Mobility & New Supply Chains	Systems and platforms that distribute attention at scale.
Social Behaviour, Consumer Economies & Global Commerce	Identity, status signalling, and social narratives shaping participation.

SESSION TWO Capital, Impact and Implementation

ORIGINAL QUESTIONS AS POSED

Can private wealth and entrepreneurial capital succeed where public policy often fails, in restoring trust and cohesion?

What models of accountability can ensure that “impact” is measured with rigour rather than rhetoric?

Do Western ESG frameworks genuinely align with the dynastic and legacy-driven governance models of major family offices, or are we seeing a divergence between institutional regulation and the long-term stewardship values of private capital?

PURPOSE OF THE SESSION

Session Two examined the relationship between capital, accountability and public benefit. It considered how capital is deployed, how impact is defined and assessed, and how different models of governance, stewardship and implementation shape the credibility of public-good claims.

SECTOR LENSES USED

Technology, AI & Frontier Intelligence	Verification, monitoring, transparency, and accountability systems.
Global Capital, Sovereign Wealth & Family Offices	Stewardship, deployment, and long-term capital governance.
Healthcare, Longevity & Human Reinvention	Population economics, longevity, and health system sustainability.
Diplomacy, Crisis Response & Education for Global Stability	Legitimacy, reporting standards, and institutional trust.
Sustainability, Climate Transition & Regenerative Economies	Impact measurement, regeneration, and transition frameworks.
Infrastructure, Energy, Mobility & New Supply Chains	Physical resilience, transition finance, and systemic constraints.
Social Behaviour, Consumer Economies & Global Commerce	Consumption, brand signalling, and value formation.

SESSION THREE Technology, Governance and Social Cohesion

ORIGINAL QUESTIONS AS POSED

Can technological progress genuinely rebuild social bonds, or does it risk deepening isolation and polarisation?

How should nations confront the demographic crises of ageing populations and declining fertility, alongside regions still grappling with youth unemployment and rapid growth?

As nations balance welfare, innovation, and regulation, which models best foster long-term cohesion, those prioritising protection and stability, or those optimising for rapid innovation and growth?

PURPOSE OF THE SESSION

Session Three examined the relationship between technological progress, governance, population change and social cohesion. It considered how technological, demographic, economic and regulatory pressures shape the conditions under which societies maintain or weaken social bonds.

SECTOR LENSES USED

Technology, AI & Frontier Intelligence	Automation, AI governance, and the redefinition of labour.
Global Capital, Sovereign Wealth & Family Offices	Who funds the future, and under what conditions.
Healthcare, Longevity & Human Reinvention	Fertility decline, longevity, and workforce sustainability.
Diplomacy, Crisis Response & Education for Global Stability	Models of governance, mediation, and regulation.
Sustainability, Climate Transition & Regenerative Economies	Intergenerational trade-offs and planetary limits.
Infrastructure, Energy, Mobility & New Supply Chains	Food, energy, water, and physical resilience.
Social Behaviour, Consumer Economies & Global Commerce	Cohesion versus polarisation in social and economic life.

The Nature of the Inquiry

The Closed Door Sessions 2026 were designed around three stated questions, and yet the behaviour of the room suggests that the stated questions functioned less as topics than as instruments of discovery. Participants did not simply answer them. They tested them, resisted them, redefined their terms and moved beyond them. This movement is itself evidence. It shows that when complex problems are placed before a cross-sector group, the first formulation of the problem is often revealed to be provisional. The room repeatedly discovered that what appeared to be a question about attention, capital or technology was in fact a question about agency, legitimacy, accountability and governance.

The most important feature of the inquiry was therefore not consensus, but transformation. In each Session, participants began from a surface problem and moved towards a deeper one. Attention became the problem of converting awareness into judgement and action. Capital became the problem of converting financial capacity into legitimate public benefit. Technology became the problem of converting capability into humanly governed progress. This pattern does not prove that all complex problems are the same. It does show that the room repeatedly found insufficiency in first-order explanations. Participants entered through sectoral language and left with systemic questions.

One possible interpretation of this behaviour is that the room was simply composed of people inclined to think systemically: investors, diplomats, technologists, security specialists, health leaders, philanthropists and institutional actors are accustomed to seeing second-order consequences. This interpretation has some force. The composition of the room shaped the inquiry. A different room might have spent longer on operational delivery, lived experience, regulatory detail or political feasibility. The evidence does not permit the conclusion that CDS 26 discovered the definitive nature of modern complexity. It does, however, support the more limited conclusion that, among this particular cross-sector group, sectoral explanations repeatedly failed to satisfy the evidence generated by discussion.

A second interpretation is that participants were avoiding specificity. When faced with difficult questions, sophisticated rooms can move upwards into abstraction. A discussion of social media becomes a discussion of shared reality; a discussion of impact investing becomes a discussion of legitimacy; a discussion of AI becomes a discussion of human agency. This risk must be acknowledged. Abstraction can be a way of evading decision, responsibility or measurement. Yet the evidence does not support the view that the room merely retreated into generality. Participants repeatedly grounded abstract questions in concrete examples: vaccine hesitancy, hostile narratives, food prices, hunger, water scarcity, defence investment, rural hospitals, weather data, healthcare algorithms, AI hallucination, child safety and philanthropic capital. The movement towards systemic explanation was not detached from evidence; it was produced by the inadequacy of single-domain explanations to account for that evidence.

A third interpretation is that the inquiry revealed the increasing interdependence of modern challenges. This has stronger explanatory power. In Session One, attention could not be separated from platforms, family, community, education, institutional trust, emotional communication and democratic leadership. In Session Two, capital could not be separated from public policy, ESG, security, local needs, implementation, measurement, philanthropy, infrastructure and geopolitics. In Session Three, technology could not be separated from human judgement, data integrity, national security, basic needs, labour value, child safety and global governance. Participants did not collapse disciplinary boundaries because they ignored disciplinary difference. They crossed boundaries because the objects under discussion no longer remained inside them.

This collapse of disciplinary boundaries was one of the most consistent behaviours of the room. Diplomatic questions became technological questions; technological questions became educational questions; capital questions became security questions; civic questions became psychological questions; humanitarian questions became economic questions. The behaviour of the room therefore supports the inference that complex problems increasingly present themselves as failures of relation between systems rather than failures within a single system. Expertise remained necessary, but insufficient. The doctor, investor, diplomat, technologist, philanthropist, security specialist or community leader could each identify part of the problem, and yet the inquiry advanced when participants tested the limits of their own explanatory field against the others.

The relationship between expertise and complexity was therefore ambivalent. Expertise provided substance, precision and examples. It allowed participants to identify mechanisms that a general conversation would have missed: how vaccine communication failed, how capital structures interact with subsidy, how ESG can mismeasure impact, how AI hallucination complicates trust, how blockchain might support data provenance, how food insecurity becomes conflict risk. Yet expertise also revealed its own insufficiency. No disciplinary perspective could close the inquiry. The evidence supports the view that, under conditions of complexity, expertise is most valuable when it is willing to be destabilised by neighbouring forms of expertise.

Disagreement was central to that destabilisation. It did not usually function as obstruction. It often advanced the inquiry by preventing premature closure. In Session One, social media was described both as corrosive and democratising. That disagreement prevented the room from adopting a simple anti-platform explanation. In Session Two, ESG was criticised as paperwork but also defended, implicitly or explicitly, as a framework capable of evolution if connected to security, climate and livelihoods. That disagreement prevented the room from rejecting all non-financial accountability. In Session Three, technology was treated as both a threat and a tool. That disagreement prevented the inquiry from becoming either technophobic or accelerationist.

The behaviour of the room suggests that disagreement advanced inquiry when it forced a distinction. Participants moved forward when disagreement clarified that a concept contained multiple meanings: community as local or digital; trust as affective or evidentiary; impact as intention or outcome; technology as tool or infrastructure; leadership as inspiration or authority; measurement as accountability or distortion. Disagreement failed to converge when the underlying standard of judgement remained unsettled. For example, participants could debate whether ESG works, but without a shared account of what ESG is meant to measure, convergence remained limited. They could debate AI governance, but without shared agreement over who holds legitimate authority, the question remained open.

The room's use of quotations as inflection points shows how conceptual movement occurred. The statement "we can't have a common conversation because we are not operating from the same starting point" reframed disagreement itself. It distinguished ordinary pluralism from epistemic fragmentation. The phrase "we were taking a knife to a gunfight" reframed public communication as an asymmetry between evidence and emotional mobilisation. "We can all measure, but we are all measuring different things" transformed measurement from a technical solution into a legitimacy problem. "Technology is a tool" prevented technological determinism, while "how do you create a mechanism where morality is part of the competition framework?" moved the inquiry from ethical aspiration to system design. These phrases mattered because they changed what the room understood the problem to be.

Uncertainty also played a constructive role. The room did not treat uncertainty only as lack of knowledge. It often functioned as evidence that the problem had outgrown existing categories. Participants were uncertain whether community remains local or has become digital; whether emotional communication can be ethical; whether private capital can rebuild trust; whether ESG can be repaired; whether impact can be measured without distortion; whether AI can be governed by existing institutions; whether child safety can provide a basis for global agreement. These uncertainties persisted not because participants failed to think sufficiently, but because competing answers remained plausible and each answer carried costs.

The persistence of unresolved tension is therefore one of the most important features of the inquiry. The room did not resolve the tension between top-down leadership and bottom-up community, between public authority and private capability, between local legitimacy and global scale, between profitability and public goods, between speed and governance, between technological openness and security, between emotional persuasion and manipulation, between democratic consent and long-term responsibility. These were not loose ends at the margins. They were structurally central. The inquiry repeatedly reached points where a solution in one register created a problem in another.

This is particularly visible in the relationship between evidence and judgement. Participants valued evidence, but they did not treat evidence as self-acting. In Session One, evidence failed if it did not become emotionally and socially intelligible. In Session Two, evidence of impact required measurement, and yet measurement itself required judgement about what counts. In Session Three, data required verification, but verification required governance and legitimacy. The evidence supports the inference that complex inquiry requires judgement not after evidence, but within evidence: in selecting what to measure, deciding whom to trust, interpreting uncertainty, and determining when sufficient proof exists for action.

The room converged most readily when participants could identify a shared failure of mediation. There was convergence around the insufficiency of information alone, the inadequacy of capital without accountability, and the danger of technological capability without governance. There was also convergence when examples revealed consequences across domains: hunger becoming conflict, misinformation becoming civic fragmentation, AI hallucination becoming decision risk, ESG becoming bureaucratic distortion, crisis becoming a mechanism for alignment. Convergence occurred when participants could see that the same phenomenon had implications for several systems at once.

Convergence failed when the inquiry required assigning authority or defining standards. Participants could agree that trust is needed, but not who should be trusted. They could agree that measurement matters, but not what should be measured. They could agree that technology needs governance, but not who should govern it. They could agree that community matters, but not what form of community now carries obligation. They could agree that long-term action is necessary, but not how democratic societies can sustain it. The evidence suggests that convergence around diagnosis was easier than convergence around institutional design.

This difference between diagnostic convergence and prescriptive divergence is central to understanding the room. Participants frequently agreed that a problem existed, but diverged when the conversation turned to remedies. This behaviour should not be interpreted as failure. In complex systems, remedies redistribute authority, risk, cost and legitimacy. Agreement that AI requires governance does not settle whether governance should be public, private, technical, international or hybrid. Agreement that impact requires measurement does not settle whether measurement should prioritise livelihoods, health, conflict reduction, financial return or avoided harm. Agreement that civic agency is weakened does not settle whether the remedy is education, regulation, community, leadership or platform reform. The divergence reflects the political nature of solutions.

The inquiry also showed that legitimacy, trust and governance were not separate themes but recurrent conditions for thought itself. Participants struggled to evaluate proposed solutions because each depended on an authority structure. Emotional storytelling might motivate action, but who decides the ethical boundary? Impact capital might scale solutions, but who validates benefit? AI governance might be technically possible, but who defines the moral operating system? Child safety might unite actors, but who sets the rules? The room repeatedly discovered that modern complexity is not only technical or financial. It is constitutional in the broader sense: it concerns the allocation of authority under conditions of interdependence.

The role of moderators was significant because moderation made the inquiry observable. Moderator interventions did not merely keep time or distribute participation. They often clarified when the room had drifted from the stated question into a deeper one, and they preserved unresolved tensions rather than forcing consensus. Moderator observations later captured the same pattern: movement from sectoral positions to systemic concerns; uncertainty around measurement; tension between leadership and democracy; concern over whether AI or capital can be trusted; recognition that participants were measuring different things; awareness that some participants dominated certain lines of inquiry. These observations are evidence of room behaviour because they record not only what was said but how the inquiry moved.

The presence of dominant voices also requires analysis. In Session Three, moderator observations noted that one participant repeatedly redirected conversation towards a particular AI governance framework. This shaped the room's trajectory. The evidence therefore requires caution: the prominence of AI may reflect both the objective importance of the topic and the dynamics of participation. Yet even this dominance is analytically useful. Complex inquiries are not neutral systems of idea exchange; they are social processes shaped by confidence, expertise, status, urgency and rhetorical force. The room's task was not to eliminate those dynamics, but to make them productive through challenge, reframing and the entry of other perspectives.

The room also revealed a relationship between complexity and humility. Participants often spoke from expertise, but the inquiry advanced most when expertise encountered its limit. The statement "I don't know" or its functional equivalent appeared repeatedly in the structure of the discussion, even when not phrased directly. Participants could identify why something matters, but not how to govern it; why measurement is needed, but not what to measure; why community matters, but not which communities can carry obligation; why technology must be constrained, but not who can do so legitimately. This humility was not intellectual weakness. It was a more accurate response to problems in which evidence, authority and consequence do not align neatly.

The room's behaviour supports the interpretation that complex problems are not simply complicated problems with more variables. A complicated problem may require expertise, coordination and technical execution. A complex problem changes as actors respond to it; it crosses institutional boundaries; it involves competing values; it produces unintended consequences; it lacks a single legitimate authority; it cannot be solved by information alone. The inquiry repeatedly encountered this form of complexity. Attempts to solve one part of the problem altered the conditions elsewhere. More emotional communication may improve action but risk manipulation. More private capital may accelerate delivery but raise legitimacy concerns. More technological capability may solve access problems but weaken human judgement. More measurement may improve accountability but distort value.

The strongest interpretation of the room's behaviour is therefore that the CDS inquiry exposed the conditions under which sector-specific expertise becomes insufficient and systemic judgement becomes necessary. This interpretation has greater explanatory power than the alternatives. The room did move from sectoral concerns towards systemic concerns, but that alone does not explain why. The room did repeatedly discover that the stated problem was not the actual problem, but that alone does not explain the direction of discovery. The room did reveal the limits of disciplinary expertise, but it also depended on expertise. The room did demonstrate interdependence, but interdependence alone does not explain the persistence of legitimacy, trust and governance questions. The most complete interpretation is that participants repeatedly encountered failures of conversion between systems, and therefore had to think across the boundaries that normally separate civic life, capital, technology, security and governance.

This interpretation also explains why the room did not produce easy closure. If the underlying difficulty is conversion between systems, then no single disciplinary solution can resolve it. A sociological solution must confront political authority. A political solution must confront economic incentives. An economic solution must confront measurement and legitimacy. A technological solution must confront human judgement and governance. A diplomatic solution must confront sovereignty and bad-faith actors. A security solution must confront civil liberties and moral constraint. The room's unresolved tensions were therefore not signs that participants failed to finish the work. They were evidence of the work's real object.

What participants believed was plural and sometimes contradictory. They believed in community, leadership, capital, technology, public policy, education, measurement, moral courage, child safety, trusted messengers and better governance. The evidence does not support the conclusion that any one of these is sufficient. What can reasonably be inferred is that the inquiry revealed the need for spaces and institutions capable of holding multiple forms of knowledge together long enough for deeper propositions to emerge. The room's value lay not in eliminating disagreement, but in allowing disagreement to reveal the inadequacy of first-order answers.

*We can't have a common conversation because we are not
operating from the same starting point.*

The nature of the inquiry was therefore diagnostic rather than programmatic. It did not issue instructions. It clarified the conditions under which instructions would be premature. It showed that before societies can act coherently on attention, capital or technology, they must understand the mediating structures that convert capacity into responsibility. The room itself briefly performed such mediation: participants from different disciplines entered a shared process, challenged one another's assumptions, introduced evidence from their own fields, reframed problems under pressure, and left several questions unresolved because the evidence did not justify closure. That behaviour is not incidental to the White Paper. It is one of its findings.

Attention, Participation and Community

THE ORIGINAL QUESTION

Session One began from the question of what motivates people to care, act and engage in civic and social life. The question appeared, at first, to invite a discussion about attention: how public-good issues attract notice, why some causes break through while others remain neglected, and what forms of message or messenger can move people from indifference to concern. Yet across the room, the inquiry did not remain within the category of attention. The participants repeatedly found that attention was not, by itself, the scarce good. Information was abundant; awareness was often already present; exposure to crisis was continuous. The more difficult problem was that attention no longer reliably became judgement, and judgement no longer reliably became action.

The strongest proposition to emerge from Session One is therefore that contemporary societies have developed increasingly powerful mechanisms for capturing attention and distributing information, while becoming weaker at converting that attention into shared judgement and sustained collective action. The fracture was not simply between knowledge and ignorance, or between engagement and apathy. It was between awareness and agency. Participants returned to this fracture through different languages: algorithmic distortion, weakened community, institutional distrust, short-termism, emotional manipulation, declining shared facts, loss of human interaction, absence of courageous leadership, and the difficulty of making abstract public-good problems personally meaningful. These were not separate conclusions so much as different attempts to name the same broken civic conversion mechanism.

The inquiry opened with a problem of scale. The formal question referred to civic and social life within the nation, but several participants resisted the assumption that the nation remains the primary unit through which civic motivation forms. Some located responsibility in family, neighbourhood, faith, local community or value system; others described digital communities that are transnational, interest-based and detached from geography. This early resistance mattered because it showed that the room did not take the classical civic subject for granted. The citizen was no longer assumed to be someone primarily addressed by the nation-state, informed by common institutions, and moved into action through shared public narratives. The participants' first difficulty was therefore not how to motivate citizens within a stable civic order, but whether the social containers that once produced civic motivation still exist in recognisable form.

That shift was visible in the repeated return to community. Some participants believed community remained the most plausible site for rebuilding action because it is close enough to be felt and durable enough to carry generational obligation. The formulation of "self-referential altruism" captured this with particular precision: people care more readily about those to whom they can perceive some relation, whether through family, neighbours, friends, locality or descendants. What the evidence supports, however, is not the simple conclusion that community is the answer. The room also recognised that community itself has changed. Younger people may experience community online, not through embodied proximity. Digital communities can reduce isolation for those who cannot find belonging locally, but they can also intensify loneliness, fragment reality and produce affiliation without reciprocal obligation. The proposition that community can repair civic agency therefore survived only in qualified form: community matters because it supplies relational meaning, but it is no longer clear which forms of community can sustain responsibility.

The inquiry then moved from community to information. Across the tables, participants repeatedly encountered the difficulty of acting together when people do not share a factual starting point. The sentence "we can't have a common conversation because we are not operating from the same starting point" marked an inflection in the inquiry. It was not merely an observation about misinformation. It identified the precondition for deliberation. Disagreement can be politically productive where participants accept common facts, common procedures or common institutions through which disagreement can be tested. The evidence from the room suggests that participants were less concerned with disagreement itself than with the loss of the shared ground on which disagreement can become judgement.

There was no full consensus on the cause of this loss. Some participants attributed it to social media platforms, commercial algorithms and echo chambers. Others resisted the suggestion that older media systems were neutral, noting that mainstream media had also been shaped by money, power and bias. This disagreement is important because it prevents the surviving proposition from collapsing into a narrow anti-technology argument. The evidence does not support the claim that digital systems alone broke civic life. It supports a more limited and stronger claim: digital systems have intensified an existing vulnerability by making influence easier, verification weaker and emotional capture more profitable.

This concern deepened when participants turned from information to judgement. Several tables identified the risk that attention systems do not simply show people content; they train habits of response. The problem was described in relation to young people, AI companionship, social media dependency, loss of critical thinking and the substitution of reaction for reflection. One participant's reference to large numbers of young people having an "AI friend" was not significant because the statistic itself was verified in the room. Its significance lay in the reaction it produced: participants treated it as evidence of anxiety about the replacement of human relation by technological simulation. The room's behaviour around this point, its hesitation and concern, indicated that the participants were not only discussing media literacy. They were asking whether the relational and cognitive capacities needed for civic agency are being weakened at the stage of formation.

We were taking a knife to a gunfight.

The Covid vaccine example provided the clearest evidence that the room had moved beyond information deficits. Participants recognised that institutions had often responded to vaccine hesitancy with data when the resistance was emotional, social and narrative. The phrase "we were taking a knife to a gunfight" marked a conceptual breakthrough. It identified an asymmetry between evidence-led institutions and actors who understand emotional transmission. The further example of a patient asking, "Is it too late for me to get the vaccine?" showed how an abstract public-health claim can become behaviourally meaningful when translated into human consequence. The evidential value of this quotation lies in the transition it records: the room recognised that truth does not act on its own. Truth requires form, carrier, emotional intelligibility and social trust.

Yet the proposed remedy created its own unease. If constructive actors must use emotional force to compete with misinformation, demagoguery or platform-driven outrage, then civic repair may require the same tools that damaged civic judgement. Some participants suggested that pro-social communication must become more compelling, personal and emotionally resonant. Others worried that answering loud voices with louder voices risks reproducing the pendulum of performative conflict. This disagreement did not weaken the central proposition. It strengthened it, because it showed that the room had identified not a messaging problem but an ethical problem of influence. The question was no longer how to make people listen, but how to move people without manipulating them.

The same structure appeared in the discussion of leadership. Some participants believed that long-term civic action requires courageous leaders able to withstand unpopularity. The statement “this is not a popularity contest” functioned as an inflection point because it named the democratic cost of long-term responsibility. Decisions made for peace, public health, climate resilience, technological restraint or defence preparedness may be necessary and still unpopular. The example of Juan Manuel Santos in Colombia illustrated this tension: moral courage did not produce universal approval; it produced division, backlash and historical ambiguity. What participants believed varied. Some placed hope in leaders, institutions or elders with moral authority. Others placed more trust in local community, decentralised governance or bottom-up action. What the evidence supports is that neither top-down leadership nor bottom-up community was sufficient by itself. Each was an attempted bridge across the same gap between awareness and action.

This is not a popularity contest.

The discussion of non-traditional messengers extended this point. Participants looked to athletes, the Pope, cultural figures, entertainers and even influencers as possible carriers of public-good messages. This was not a superficial turn to celebrity. It reflected an implicit recognition that formal authority no longer guarantees transmission. If younger publics receive information through TikTok, sport, entertainment, religion, music or peer networks, then civic communication must travel through those channels or fail to reach the people it seeks to move. The evidence supports the inference that messenger legitimacy has become plural, segmented and context-dependent. It does not support the simpler claim that traditional institutions are obsolete. Participants still returned to stabilising institutions, education, family and government protection. The stronger conclusion is that no single authority now monopolises trust.

This explains why the room repeatedly returned to education. Education appeared not as a policy sector but as a precondition for agency. Participants linked it to critical thinking, self-accountability, value formation and the capacity to navigate information environments shaped by commercial incentives and algorithmic amplification. The concern was not only that people may believe falsehoods. It was that people may lose the habit of testing claims, comparing sources, sustaining attention, entering disagreement and distinguishing emotional salience from truth. The evidence supports the view that education was understood as civic infrastructure. It was one of the few interventions that could plausibly operate before attention is captured and before misinformation hardens into identity.

The room also tested a material interpretation of civic motivation. Several interventions suggested that people act when issues touch their welfare, children, health, food, water, safety, costs or immediate environment. This interpretation diverged from explanations centred on platforms or institutions. It suggests that the conversion from awareness to action depends on material relevance. People may care about hunger, climate or public health in the abstract, but they act when the issue becomes proximate, embodied or economically legible. This does not contradict the central proposition. It clarifies one of its mechanisms. The conversion mechanism fails when public-good problems remain abstract, distant or morally general; it functions better when those problems become personal, relational and materially intelligible.

There was, however, a contradictory strand in the evidence. One participant suggested that comfort may itself be the ultimate distraction. This complicates the assumption that civic disengagement arises primarily from hardship, misinformation or exclusion. If comfort reduces urgency, then action cannot be explained only by need. Yet other participants noted that affluent societies are not necessarily passive; they may be full of rage, division, dissent and political mobilisation. The evidence therefore does not support a simple scarcity-or-comfort model. It supports a more complex account in which motivation arises from perceived stake, grievance, belonging, fear, aspiration, comparison and moral identification, while the direction of action depends on the structures that interpret those feelings.

The moderator observations are important because they show how the room behaved when confronted with these difficulties. Across tables, there was recurring convergence around the inadequacy of information alone. There was repeated energy around community, storytelling, critical thinking, trusted messengers and leadership. There was hesitation around emotional influence, because the tools of mobilisation can become tools of manipulation. There was disagreement over whether technology democratises or degrades information. There was unresolved tension between regulation and freedom, between top-down leadership and bottom-up community, between local belonging and digital identity, and between civic participation as desirable and civic participation as potentially destructive. These dynamics show that the participants were not moving toward a neat solution. They were circling a structural difficulty: the mechanisms that capture attention are often stronger than the mechanisms that form judgement, and the mechanisms that distribute information are often stronger than those that generate responsibility.

WHAT EMERGED

What emerged was not a theory of attention, but a theory of civic conversion. The strongest surviving proposition has greater explanatory power than the competing alternatives because it accounts for the widest range of evidence across the room. The algorithmic explanation explains echo chambers, misinformation and emotional capture, but it cannot account by itself for the recurrence of family, community, leadership, education, material welfare and moral courage. The institutional explanation explains distrust of politics, media and expert authority, but it does not account fully for the way people may still trust athletes, parents, religious figures, local communities or digital peers. The psychological explanation explains why fear, comparison, aspiration and belonging matter, but it does not account for the institutional and technological environments that shape those responses. The leadership explanation explains the need for long-term moral courage, but not why leadership struggles to reach publics whose attention is fragmented across multiple social worlds. The community explanation explains relational motivation, but not the collapse of shared facts or the ethical problem of influence.

The proposition of a broken civic conversion mechanism is stronger because it incorporates these explanations without reducing the evidence to any one of them. Algorithmic distortion, institutional distrust, weakened community, short-termism, emotional resonance, material relevance and leadership failure all become manifestations of a deeper condition: awareness no longer reliably travels through shared institutions, trusted relationships and stable narratives into judgement and action. The room repeatedly returned to this condition whether participants were discussing vaccines, climate, hunger, defence, AI, young people, social media, positive narratives, local community or courageous leadership.

What participants believed was not uniform. Some believed the answer lay in stronger regulation of platforms and technology. Some believed the answer lay in education and critical thinking. Some believed it lay in rebuilding community. Some argued for better storytelling and emotionally intelligent communication. Some emphasised leadership and institutions. Others looked to trusted messengers outside politics. The evidence supports none of these as a sufficient explanation or remedy. It supports the view that each addresses a different point in the same broken chain: attention must be captured, information must be trusted, meaning must be formed, judgement must be exercised, and action must be socially sustained.

What can reasonably be inferred is that Session One revealed a crisis of civic agency under conditions of informational abundance. The room did not suggest that people know too little. It suggested that knowledge is no longer reliably organised into responsibility. People may be more informed and less engaged; more connected and lonelier; more exposed to crisis and less able to act; more expressive and less deliberative; more mobilised and less oriented towards shared purpose. The consequence is that public-good issues may become visible without becoming governable. They circulate as content, crisis, grievance, identity or performance, but fail to become sustained collective action.

This has implications for institutions. Institutions that assume evidence will convert into behaviour may continue to misread the problem. Public agencies, philanthropies, companies and civic organisations may produce accurate information, but if they do not attend to trust, messenger, emotional resonance, material relevance and community context, they may not produce action. Conversely, actors skilled in emotional mobilisation may produce action without truth, responsibility or constructive purpose. The room therefore identified a double danger: evidence without agency on one side, mobilisation without judgement on the other.

The implications for democracy are equally difficult. Democratic systems require publics who can receive information, test claims, tolerate disagreement, recognise common stakes and act beyond immediate self-interest. The evidence from Session One suggests that each of these capacities is under pressure. Yet the room also resisted anti-democratic conclusions. Participants recognised the dangers of strong leaders who can act without consent, and the dangers of publics who no longer trust democratic institutions. The question was not whether democracy should be replaced by technocratic or authoritarian long-termism. It was whether democratic societies can rebuild the conditions under which long-term responsibility becomes politically possible.

The implications for community are unresolved but central. If community remains the most plausible site of obligation, then civic renewal may need to begin at smaller scales than national rhetoric usually assumes. Yet if community is increasingly digital, segmented and self-selected, then it may reinforce fragmentation rather than repair it. The evidence supports a distinction between community as belonging and community as responsibility. Digital groups may provide identity and companionship; local or relational communities may provide mutual obligation; national communities may provide institutional scale. Session One did not resolve how these forms relate, but it showed that civic action depends on some form of belonging capable of carrying obligation across time.

The implications for communication are also ethically unsettled. The room recognised that emotionally resonant stories can move people where data cannot. It also recognised that emotional force can be used to mislead. This tension matters because any attempt to repair civic conversion will require communication that is more than informational but less than manipulative. The room did not define that boundary. It did, however, make clear that civic communication cannot return to a model in which facts are simply published and action follows. Meaning must be constructed, but the construction of meaning must itself remain accountable to truth.

- 01 The central unresolved question is whether the broken conversion mechanism can be repaired, or whether societies must learn to govern under conditions where shared reality remains permanently weakened. Participants repeatedly proposed repair mechanisms: education, family formation, platform responsibility, regulation, technical intervention, community building, moral leadership, stabilising institutions and trusted messengers. The evidence does not yet show whether these can operate at sufficient scale or speed.
- 02 A second unresolved question concerns ethical influence. If public-good actors must communicate emotionally in order to compete, what distinguishes civic persuasion from manipulation? The vaccine example showed that emotional narrative may be necessary; the hesitation around louder pro-social messaging showed that necessity does not settle legitimacy. The chapter's evidence leaves open whether ethical influence can be designed, taught or institutionally governed.
- 03 A third unresolved question concerns the meaning of community. Participants agreed that community matters, but disagreed over whether community is local, familial, digital, national, religious, value-based or generational. The evidence suggests that civic agency depends on community, but does not identify which forms of community can sustain durable public responsibility rather than fragmented identity.
- 04 A fourth unresolved question concerns leadership. The room recognised the need for long-term moral courage, but also the democratic costs of unpopular action. Can democratic leaders take decisions that require long horizons when attention systems reward short-term reaction and electoral systems punish discomfort? Can non-political leaders compensate for political weakness, or does that simply displace accountability into unelected forms of influence?

05

A fifth unresolved question concerns education. Participants repeatedly returned to critical thinking, but did not resolve what kind of education is needed: media literacy, civic education, moral formation, family-based value transmission, digital self-accountability, or training in deliberation and disagreement. Nor did they resolve whether education can act quickly enough when attention systems shape cognition continuously.

06

A sixth unresolved question concerns responsibility. If social media platforms, AI systems and commercial algorithms shape civic judgement, who is accountable for the consequences: companies, governments, schools, families, users, investors or civil society? The room's repeated movement across these possible actors suggests that responsibility is dispersed, but dispersed responsibility can easily become no responsibility.

07

The final unresolved question concerns measurement. Participants implicitly asked what should count as success: attention, awareness, engagement, judgement, trust, action, community strength, avoided harm, reduced polarisation, institutional legitimacy or long-term behavioural change. Until this is clarified, interventions may continue to optimise for the wrong thing. Session One suggests that the relevant measure is not whether people notice a public-good problem, but whether they become capable of acting upon it responsibly with others.

Capital, Impact and Implementation

THE ORIGINAL QUESTION

Session Two began from the question of whether private wealth and entrepreneurial capital can succeed where public policy often fails in restoring trust and cohesion, and what models of accountability can ensure that impact is measured with rigour rather than rhetoric. The question assumed a distinction between private capital, public policy, entrepreneurship and accountability. Across the room, however, that distinction became increasingly unstable. Participants did not simply debate whether private capital is more agile than government, or whether public policy is too slow. They repeatedly encountered a deeper difficulty: societies do not primarily lack capital, but lack trusted mechanisms for converting capital into legitimate, measurable and durable public benefit.

The surviving proposition from Session Two is that capital now faces an accountability gap. Private wealth, public funds, philanthropy, entrepreneurial capital and institutional investment may all be capable of addressing public-good problems, but none is inherently legitimate by source alone. Capital becomes socially useful only when it is tied to credible purpose, appropriate time horizons, local relevance, rigorous measurement, implementation capacity and public trust. Without those conditions, capital risks becoming rhetoric, compliance, extraction, reputational display or a substitute for public responsibility.

This proposition has greater explanatory power than the simpler claim that private capital should replace failed public policy. The evidence across the room repeatedly resisted that conclusion. Participants recognised that public systems are often slow, bureaucratic and politically constrained, but they also recognised that private capital can be mistrusted, short-term, self-interested, geographically misaligned or insufficiently accountable. The room's inquiry therefore moved beyond the binary of public versus private and towards the conditions under which any capital can legitimately claim to produce impact.

RECONSTRUCTING THE INQUIRY

The first movement in the inquiry was away from source and towards function. Some participants believed that private capital can move faster than public policy, especially where government systems fail to meet urgent needs. Others emphasised the distinctive role of public frameworks, regulation, subsidy, procurement and legitimacy. What the evidence supports is neither a pro-private nor pro-public conclusion, but a more demanding one: the relevant question is not where capital comes from, but what problem it is solving, under what authority, with what evidence, and for whose benefit.

Fund the jockey, not the horse.

This was visible in the repeated distinction between capital as funding and capital as responsibility. One table framed effective capital through identifiable human purpose: capital becomes transformative when someone assumes responsibility for a concrete problem and remains accountable for outcomes. The phrase “fund the jockey not the horse” captured this belief in capable, purpose-driven agents. As an evidential inflection point, it shows how some participants understood capital as relational and judgement-based. Money alone does not produce change; it must be entrusted to people or institutions capable of execution. Yet the same formulation also exposes a limit. If systemic problems are reduced too far to exceptional individuals, then accountability may become personal charisma rather than institutional design.

The inquiry therefore turned repeatedly to measurement. Participants broadly agreed that impact cannot be asserted through intention or language. It must be demonstrated. Yet they also found that measurement is not a neutral technical exercise. One of the most important formulations, “we can all measure, but we are all measuring different things,” identified the measurement gap at the centre of the session. The problem is not only that some actors fail to measure impact. It is that different actors measure return, social benefit, risk, standard of living, avoided harm, health outcomes, trust, local resilience, conflict reduction or reputational compliance. Without shared evaluative architecture, measurement can produce the appearance of accountability without establishing legitimacy.

This difficulty explains the recurring frustration with ESG. Participants did not all reject ESG. Some argued that ESG, if properly understood, should include security, hunger, climate, livelihoods and defence, because these are conditions of social stability. Others regarded ESG as bureaucratic, politically vulnerable and conceptually confused. The phrase “ESG is paperwork” was an inflection point because it distinguished reporting from consequence. It did not prove that all ESG frameworks are useless; rather, it showed why ESG loses legitimacy when compliance becomes detached from the actual problem being solved. The evidence supports the narrower proposition that ESG and impact are too often conflated. ESG may describe process, risk, disclosure or compliance; impact must describe real-world change.

The same distinction appeared in the critique of “impact tourists.” Some participants observed that institutional capital entered impact investing when it was fashionable, then froze or withdrew when political, market or geopolitical conditions shifted. This exposed a tension between capital that seeks durable public benefit and capital that follows reputational or cyclical incentives. The room did not conclude that institutional capital is incapable of impact. It did, however, recognise that capital cannot be trusted if its commitment disappears under pressure. Durability became one of the implicit tests of legitimacy.

The inquiry then moved from measurement to time horizon. Several participants argued that meaningful public-good innovation requires patience: medical devices may take a decade to reach market; climate and infrastructure transitions require long-term stability; defence and resilience cannot be built only after crisis; local ecosystems require capital suited to their own development cycles. This evidence challenged conventional venture-capital assumptions about rapid scale, short return cycles and market capture. At the same time, participants did not reject profitability. The formulation “there is no sustainability without profitability, and no profitability without sustainability” marked an attempt to overcome a false opposition. It suggested that financial durability and public benefit must be mutually reinforcing, not sequential or decorative.

*There is no sustainability without profitability, and no profitability
without sustainability.*

The evidence supports this as a condition, not a complete answer. Some problems can and should be solved through profitable, cost-competitive technologies. Participants cited examples in climate, hydrogen, health and infrastructure where market viability may allow public-good solutions to survive political volatility. Yet other problems, particularly those involving basic needs, humanitarian crises, prevention or vulnerable communities, may not produce market-rate returns without public, philanthropic or blended support. The room therefore did not establish profitability as the sole criterion for legitimacy. It established that impact dependent indefinitely on unstable subsidy or political fashion remains fragile, while public goods that cannot be monetised still require accountable non-market capital.

This distinction became sharper when participants discussed Western policy volatility. Several contributions contrasted long-term industrial strategies in China or other Asian contexts with Western democratic policy cycles. The evidence showed concern that Western administrations can reverse climate, infrastructure or social-investment incentives abruptly, undermining technologies and capital formation that require decades. Some participants inferred from this that policy should not be the primary driver of transition. Others pointed to examples where stable policy had supported energy transition, or argued that public policy remains indispensable for scale. The evidence supports neither anti-policy nor policy-only interpretations. It supports the conclusion that capital allocation for public benefit requires institutional stability, whether that stability comes from public policy, long-term private capital, blended structures or durable regulatory frameworks.

The conversation's treatment of infrastructure extended this point. Participants repeatedly moved beyond startup language towards population-level systems: grids, health provision, cooling cities, water, energy, rural care, behavioural health and security. This shift mattered because it changed the scale of accountability. Startups can demonstrate product-market fit; infrastructure must demonstrate social fit, resilience, maintenance, access and legitimacy. The phrase "entangled capital" captured the structural consequence of this shift. At sufficient scale, capital is no longer simply private or public. It becomes entangled with governance, security, infrastructure, public trust and geopolitical power. This proposition reframed the session: the more systemic the capital, the greater the legitimacy burden.

The room's discussion of trust in wealth made this burden explicit. Participants recognised that private capital can meet needs where public policy fails, but also that concentrated wealth may generate suspicion or hostility. Trust in responsible private-sector actors may be higher than trust in politicians or media in some contexts, but trust in the very wealthy is not automatic. Communities may ask whether capital is aligned with their interests, whether it extracts value, whether it delivers what it promises, and whether it remains accountable after reputational benefit has been secured. This evidence undermines any simple claim that private capital restores trust by acting. It restores trust only when action is demonstrably aligned with need, outcome and responsibility.

This also means that concentrated capital does not only respond to public-good problems once they have been defined. It may influence which problems become visible, investable, measurable or institutionally actionable in the first place. This agenda-setting capacity is not inherently illegitimate; in some contexts, it may bring neglected problems into view or move resources where public systems have failed to act. Yet it does increase the burden of accountability, because the power to fund solutions can also become a power to shape the definition of need, the acceptable form of evidence and the criteria by which success is judged.

Locality complicated the question further. Some participants argued that capital often imposes metrics, time horizons and expectations from dominant financial geographies onto communities where different forms of growth, resilience or benefit may matter more. Local capital for local solutions was therefore not simply a slogan; it was a challenge to external evaluation. A technology, fund or infrastructure project may appear successful under one financial model while failing to address the lived problem in another context. The evidence supports the conclusion that impact measurement must be sensitive to place without becoming so localised that comparison, scale or accountability becomes impossible.

The strongest reframing came when participants linked basic needs to security. Hunger, water scarcity, climate instability, health access and livelihoods were not treated as separate humanitarian categories. They were described as drivers of migration, conflict, displacement and social instability. The phrase “you cannot be green if you are in the red” worked as an evidential inflection because it grounded impact in material survival. Environmental, technological or social ambitions cannot be politically or morally sustained where basic human needs are unmet. This shifted the inquiry from philanthropic generosity to security architecture. Investment in food, water, health and livelihoods became not only compassionate but structurally preventive.

You cannot be green if you are in the red.

This reframing also exposed the limits of Northern impact discourse. One participant insisted that “impact should be for everyone,” challenging frameworks that remain oriented around the concerns, categories or reputations of wealthy institutions in the Global North. If impact is to be credible, it must reach people in need, operate at scale, improve livelihoods, secure essential services and acknowledge interdependence across geographies. What happens in the Global South affects security, migration, markets and political stability elsewhere. The evidence supports the inference that impact cannot be legitimate if it remains geographically selective or reputationally convenient.

The inquiry's final movement concerned implementation. Participants repeatedly suggested that capital too often funds novelty rather than deployment. The complaint that capital should stop financing innovation alone and finance the implementation of innovations already developed was one of the session's clearest practical inflections. It exposed a misallocation of prestige and risk. New ideas attract attention, but public benefit often depends on scaling, maintenance, procurement, training, distribution, local adaptation and adoption. The room's behaviour around this point suggested convergence: participants from different sectors recognised that a solution unrealised at scale is not yet impact.

We can all measure, but we are all measuring different things.

WHAT EMERGED

What emerged from Session Two was not a rejection of private capital, public policy, ESG, philanthropy or entrepreneurship. It was a demand that each be judged by function, evidence and legitimacy. The strongest proposition survived because it explains the broadest range of evidence: the central difficulty is an accountability gap in capital. This gap appears when private capital acts faster than public systems but lacks democratic legitimacy; when public policy sets goals but cannot sustain them across political cycles; when ESG records compliance but not consequence; when impact investors claim purpose without defining risk, return and beneficiary; when capital funds innovation but not implementation; when local problems are assessed by external metrics; and when basic needs are treated as charity rather than security.

The competing interpretations each explain part of the evidence. The public-policy failure interpretation explains why participants looked to private and philanthropic capital. It does not explain why the room remained concerned about mistrust of wealth, measurement, local legitimacy and public responsibility. The capital-discipline interpretation explains the critique of venture timelines, subsidy dependence and short-term returns. It does not fully account for the need for policy, blended finance and non-market provision. The ESG legitimacy interpretation explains frustration with bureaucracy and compliance. It does not account for those who saw ESG as potentially relevant when expanded to security and resilience. The locality interpretation explains the need for context-specific capital, but not the need for scalable institutional structures. The security interpretation explains the connection between hunger, water, climate, migration and conflict, but not the full range of measurement and trust problems. The accountability-gap proposition contains all of these without reducing the evidence to any one of them.

What participants believed varied substantially. Some believed private capital is the only credible way to bridge public-policy gaps. Some believed public frameworks remain indispensable. Some placed trust in entrepreneurial founders and mission-driven investors. Some emphasised infrastructure, others prevention, others measurement, others local needs, others basic services, and others geopolitical risk. What the evidence supports is that capital becomes legitimate only when it can answer a sequence of questions: what problem is being solved, who defines that problem, what capital is appropriate, what time horizon is required, who benefits, who measures, who bears risk, who is accountable, and what happens after initial funding.

What can reasonably be inferred is that the room understood impact as a governance problem. This is not governance only in the formal sense of laws, boards, regulators or public agencies. It is governance in the broader sense of ordering capital towards public consequence. The inquiry repeatedly returned to the same difficulty: money can move, but movement is not impact; innovation can be financed, but financing is not implementation; measurement can be produced, but measurement is not legitimacy; intention can be sincere, but sincerity is not accountability.

The consequences are significant for institutional investors and family offices. If impact remains a label, it will be vulnerable to political backlash, reputational scepticism and internal confusion. If it becomes a discipline of problem definition, capital matching, measurement, implementation and durability, it may contribute to public trust. This requires investors to be explicit about their risk-return-impact balance, rather than concealing trade-offs beneath aspirational language. It also requires humility about what different forms of capital can and cannot do. Philanthropy can absorb risk and seed models; catalytic capital can de-risk; institutional capital can scale; public policy can legitimise and structure; local capital can contextualise; entrepreneurial capital can execute. Confusing these roles weakens accountability.

The consequences for public systems are equally important. Public policy cannot simply delegate public goods to private actors without addressing legitimacy. Where government fails to deliver, private capital may create hope, access and innovation; it may also create dependency, suspicion or unequal provision. The evidence suggests that public systems need capital partners, but also frameworks that prevent public responsibility from dissolving into private discretion. The question is not whether private capital should do more. It is how private capital can contribute without becoming the unaccountable infrastructure of public life.

The consequences for ESG are particularly severe. If ESG remains associated with paperwork, reversibility, box-ticking or political fashion, it will not command confidence. Yet the room's evidence does not support abandoning all attempts to align capital with public outcomes. It suggests that ESG must be distinguished from impact and, where useful, subordinated to clearer questions of problem, consequence and accountability. A defence investment may contribute to democratic security; a climate investment may fail if it does not reach those in material insecurity; a green technology may be irrelevant if it cannot be implemented affordably; a humanitarian intervention may have security value if it reduces displacement or conflict. Existing labels do not reliably capture these relationships.

The consequences for measurement remain unresolved but unavoidable. Participants repeatedly wanted rigorous evidence, but there was no agreement on what should be measured. Financial return is necessary but insufficient. Social return can be powerful but may be methodologically fragile. Avoided harm is difficult to prove. Lack of conflict may be an important indicator but hard to attribute. Local resilience matters but resists standardisation. The room therefore revealed a paradox: impact requires measurement to become accountable, but premature or simplistic measurement can distort impact by forcing complex public goods into inadequate metrics.

The consequences for implementation may be the most practical. Capital markets often reward invention, announcement and scale narratives more readily than deployment. Yet for communities, implementation is where impact becomes real. Water systems, health access, food security, rural clinics, climate adaptation and educational tools do not produce public benefit through concept alone. They require financing models that support installation, maintenance, training, uptake, local governance and long-term viability. Session Two therefore suggests that the frontier of impact is not only innovation. It is the disciplined conversion of innovation into functioning social infrastructure.

QUESTIONS LEFT UNRESOLVED

O1

The first unresolved question is who should define impact. Participants variously implied investors, communities, beneficiaries, governments, technical experts, entrepreneurs and independent evaluators. Each has a claim; none is sufficient alone. Investors can define risk and return, but not social legitimacy by themselves. Communities can identify need, but may lack technical or financial visibility. Governments can confer public authority, but may be slow or politicised. Experts can measure, but may narrow what counts. The room did not resolve how these forms of authority should be combined.

O2

A second unresolved question concerns measurement. The session exposed the need for rigorous proof, but did not establish whether impact should be measured through livelihoods, health, standard of living, safety, trust, avoided harm, conflict reduction, resilience, economic return or systemic risk. The difficulty is not only methodological. It is political and moral. What societies choose to measure determines what capital will pursue.

- 03 A third unresolved question concerns the boundary between public responsibility and private initiative. Private capital may be necessary where public systems fail, but if it becomes the main provider of public goods, who ensures fairness, access and accountability? Conversely, if public systems retain authority but cannot act at speed, how can urgent problems be addressed? The room identified the need for partnership but did not settle the terms of legitimacy.
- 04 A fourth unresolved question concerns profit. Participants broadly accepted that sustainable impact cannot depend indefinitely on subsidy, yet they also recognised that many essential needs may not generate market-rate returns. How should capital distinguish between problems that can become profitable, problems that require catalytic de-risking, and problems that must remain public or philanthropic obligations?
- 05 A fifth unresolved question concerns locality and scale. Local capital and local solutions may better match lived needs, but population-level problems often require large pools of institutional capital and infrastructure-scale coordination. The room did not resolve how to preserve local relevance while achieving sufficient scale.
- 06 A sixth unresolved question concerns ESG. Should ESG be repaired, clarified, replaced or separated from impact altogether? If ESG is too bureaucratic, it may lose credibility. If abandoned, capital may lose one of the few existing languages for non-financial accountability. The evidence points towards reform and differentiation, but not towards a settled framework.
- 07 A seventh unresolved question concerns implementation. Why does capital continue to favour new ideas over the scaling of existing solutions, and what financial structures would change that bias? Participants gestured towards infrastructure bonds, public-private partnerships, blended finance and catalytic capital, but the room did not establish which mechanisms can reliably move capital from demonstration to deployment.
- 08 The final unresolved question concerns trust. Can capital rebuild trust in institutions and social systems, or does concentrated capital inevitably deepen suspicion? The evidence suggests that trust depends on delivered outcomes, transparent purpose, appropriate measurement and long-term commitment. Yet trust cannot be engineered solely through performance metrics. It also depends on legitimacy, relationship and the perception that those deploying capital are not merely solving problems for others, but accepting responsibility within a shared social order.

Technology, Governance and Social Cohesion

THE ORIGINAL QUESTION

Session Three began from the question of whether technological progress can rebuild social bonds or whether it risks deepening isolation, polarisation and social fragmentation. The question also invited participants to consider demographic pressures, welfare innovation, regulation, long-term cohesion, and the competing claims of protection, stability, rapid innovation and growth. At first glance, the inquiry appeared to ask whether technology is socially cohesive or socially corrosive. Across the room, however, participants did not remain within that binary. They repeatedly found that technology itself was not the deepest object of inquiry. The deeper question was whether societies still possess the human, institutional, ethical and epistemic capacities needed to govern technologies that are advancing faster than shared rules, public understanding and political legitimacy.

The strongest proposition to emerge from Session Three is therefore that the room exposed a governance and human-agency gap. Societies are developing technologies, capital systems and infrastructures capable of reshaping human life faster than they can define the purposes, safeguards, responsibilities and common rules required to govern them. The same technology may widen or narrow inequality; extend care or lower barriers to harm; support truth or generate authoritative falsehood; strengthen human agency or substitute for it. The decisive issue is not technological progress in itself, but whether technology remains anchored in human judgement, trusted data, basic needs, moral constraint, institutional legitimacy and shared-risk governance.

This proposition has greater explanatory power than the simpler claim that technology is either good or bad for social cohesion. Participants did not converge around technological optimism or pessimism. They converged around conditionality. Technology was repeatedly described as a tool, an amplifier, an infrastructure, a market force, a security risk and a potential instrument of public benefit. The question became: under what conditions can such tools serve human purposes rather than reorganise human life around logics of speed, profit, dependency, strategic advantage or unverified authority?

RECONSTRUCTING THE INQUIRY

The inquiry began by unsettling technological determinism. Across the room, participants resisted the idea that technology has a single social direction. Some emphasised its capacity to extend healthcare, improve access, support humanitarian supply chains, strengthen weather forecasting, scale behavioural interventions or connect isolated people. Others emphasised misinformation, hallucinated outputs, social media polarisation, AI dependence, algorithmic exploitation, digital exclusion and the risk of agentic systems moving beyond human comprehension. What the evidence supports is not a verdict on technology, but a claim about governance: the same technical capacity can produce contrary social outcomes depending on the institutional, moral, economic and epistemic systems in which it is deployed.

The formulation “technology is a tool” served as an early inflection point because it prevented the inquiry from becoming a morality play about machines. A tool can be used well or badly, but that does not mean its effects are neutral. Tools embed assumptions, incentives, forms of access, power relations and dependencies. The room’s repeated return to AI, blockchain, social media, digital finance, health technologies, weather data and humanitarian logistics showed that technologies increasingly mediate not only communication, but care, knowledge, security, capital allocation and public trust. If technology is a tool, it is now a tool operating inside the central infrastructures of social life.

This raised the problem of legitimacy. Participants could identify beneficial uses of technology, particularly in mental health, healthcare access, humanitarian systems and public-service delivery, but repeatedly asked who should decide how such tools are governed. Governments may have formal authority but may lack speed, technical understanding or public trust. Companies may have expertise and infrastructure but are shaped by commercial incentives. Investors can allocate capital but do not automatically possess democratic legitimacy. International organisations may convene but may not enforce. Technical communities may understand systems but may not be politically accountable. The evidence therefore supports the conclusion that technological capability has outpaced legitimate authority.

The inquiry became more acute when participants considered AI in relation to human judgement. Several interventions suggested that AI may be valuable when it augments human capacity, improves measurement, supports diagnosis, enables prediction or scales scarce expertise. Yet participants also worried that younger generations may grow up assuming that AI “does the thinking.” The statement “we had to all learn to critically think growing up without AI” marked an important shift. It was not simply a complaint about education. It identified a developmental risk: if human judgement is not formed before technological dependency becomes habitual, societies may lose the very capacities required to audit, contest and govern machine outputs.

I don't believe what I cannot prove with extremely high probability.

This concern was reinforced by discussions of misinformation and hallucination. Participants did not treat false information as only a social media problem. In several exchanges, they considered the authority of machine-generated outputs in domains where mistakes have material consequences. Weather forecasting offered one of the most concrete examples. AI can produce persuasive maps and predictions, but if those outputs are wrong, farmers, humanitarian actors or public authorities may make consequential decisions on false premises. The problem is not that AI may err; human systems err as well. The problem is that AI outputs can appear authoritative while their data provenance, reasoning path or uncertainty remains opaque. The sentence “I don't believe what I cannot prove with extremely high probability” identified the truth-governance problem with particular force. Trust in technological systems increasingly requires verifiability, provenance and proof, not merely confidence in computation.

This led to the discussion of data integrity, blockchain, real-world assets, real data assets, satellite verification and immutable records. Some participants believed that technical systems could help solve problems of trust by making data harder to alter, easier to verify and more suitable for capital markets, regulation or insurance. Others questioned whether technical assurance can substitute for democratic accountability. The evidence supports neither a rejection nor an endorsement of these tools. It supports the narrower conclusion that technical verification can strengthen trust only if the systems producing it are themselves legitimate, comprehensible and accountable. A mathematically secure record may still raise political questions about who defines the data, who controls the infrastructure, who benefits from its use, and who can challenge its conclusions.

The same difficulty appeared in relation to public institutions. The observation that “government mindsets are seven years behind” functioned as a diagnosis of institutional lag. Participants were not simply criticising bureaucracy. They were identifying a temporal mismatch between technological acceleration and public governance. Regulation, procurement, education, democratic deliberation and international agreement often operate slowly. AI systems, platform incentives, digital finance and autonomous agents evolve quickly. The room’s concern was that the pace of technological change may not allow societies enough time to establish legitimate oversight before harms become embedded.

Government mindsets are seven years behind.

Yet participants did not agree that faster private innovation is the answer. Several exchanges exposed a paradox. Public systems may be too slow, but private systems may be too incentive-driven. Platforms optimise for engagement; companies pursue market share; investors reward growth; defence actors respond to adversarial pressure; AI firms race towards capability. One participant’s question, “how do you create a mechanism where morality is part of the competition framework?” became one of the central inflection points of the session. It shifted the inquiry from ethics as exhortation to ethics as structure. If moral constraint remains external to markets, it may be outcompeted by faster, less constrained systems. The room therefore moved towards the problem of how ethical behaviour can be made economically and strategically viable, not merely desirable.

This was the strongest contribution of the discussion around AI governance. Participants considered human-in-the-loop design, moral operating systems, legal compass frameworks, AI for good, smart contracts, open-source architectures and blockchain-based rules. These proposals differed technically, but they shared an underlying premise: human constraint must be built into systems before those systems become too autonomous, too profitable or too widely adopted to be meaningfully restrained. The concern was sharpened by the question of agentic AI. The image of AI agents communicating with one another in languages humans do not understand was an escalation point. The question “how are we going to make sure that we as humanity always have an upper hand over the agentic silicon lifeform?” moved the room from regulation of tools to preservation of human authority over autonomous systems.

*How are we going to make sure that we as humanity always
have an upper hand over the agentic silicon lifeform?*

The room did not settle whether this fear is imminent, speculative or already material. Moderator observations note that some exchanges were shaped by a dominant participant repeatedly returning to AI governance. This matters methodologically. The prominence of AI in the session should not be mistaken for unanimous prioritisation. Other participants might have taken the conversation further into demographics, welfare, ageing, basic needs, public-sector reform or social contracts if the conversational centre of gravity had not repeatedly returned to AI. Yet the dominance of AI was itself evidence of the room’s behaviour. AI became the site onto which wider anxieties about agency, truth, governance, capital, security and human value were projected.

The inquiry also moved beyond AI to the conditions of social cohesion. Some participants argued that technology cannot be discussed apart from water, food, energy, healthcare, sanitation and education. The statement that if people cannot eat or drink, the result becomes an economic and security issue because people “flee, fight, or die” grounded the conversation in material life. This was an important corrective to abstract technological speculation. It suggested that cohesion cannot be built through advanced tools alone. Societies that lack foundational infrastructure are not simply waiting for innovation; they are exposed to insecurity, displacement and conflict. Technology may assist, but it cannot replace the political and economic work of securing basic needs.

This evidence complicates any interpretation that treats AI governance as the whole of Session Three. Across the room, participants repeatedly linked technological futures to food systems, health equity, climate stress, water, humanitarian supply chains, demographic imbalance and public welfare. The deeper issue was not only how to govern AI, but how to ensure that technological progress remains aligned with human needs rather than detached from them. For many populations, the relevant technological breakthrough may be reliable water access, not artificial general intelligence. The room's attention to basic needs supports the inference that human agency is materially conditioned. People cannot exercise meaningful agency in technological futures if their immediate survival systems are unstable.

The discussion of cohesion therefore became more ambitious. One table introduced the distinction between “thin cohesion” and “deep cohesion.” Thin cohesion referred to partial forms of belonging, identity management or social order that may hold temporarily but do not address shared planetary risk. Deep cohesion would require recognition that societies are entangled through climate, resources, finance, technology, security and intergenerational vulnerability. This distinction matters because it reframed cohesion away from sentiment. The room was not asking how to make people feel more unified. It was asking whether shared risk can become the basis for a more durable social contract.

The value of being human is diminishing.

The labour-value question intensified this inquiry. Some participants raised the possibility that AI and automation may reduce the economic value of human labour. The claim that “the value of being human is diminishing” was not important because the room accepted it as fact. Its importance lay in the question it forced: if societies continue to value human beings primarily through labour productivity, what happens when technological systems reduce the need for human labour? This exposed a human-value problem beneath the governance problem. Technological cohesion cannot be achieved simply by regulating tools if the social contract remains organised around forms of work that technology is displacing or devaluing.

The room did not resolve this. Some participants focused on universal basic income, digital currencies, virtual reality, displacement and future labour markets. Others returned to education, basic needs, interpretive intelligence and human participation. What the evidence supports is that the labour question is part of the wider agency gap. Human beings require more than income replacement if technology changes their economic role. They require social worth, participation, judgement, belonging and responsibility. Without these, welfare innovation may preserve subsistence while failing to preserve agency.

Security repeatedly interrupted more optimistic accounts of technology. Participants involved in defence and diplomacy warned that bad actors, hostile states and adversarial systems cannot be ignored. The statement that “there are inherently bad people in this world who want to do irreparable harm” challenged any purely benevolent model of AI for good. In defence, cyber, digital finance and national security, ethical actors must operate in environments where others may not share their moral constraints. This produced a central disagreement. Some participants emphasised open, decentralised or humanistic governance; others argued that defensive capacity and strategic realism are necessary. The evidence supports a synthesis: ethical technological governance must be able to withstand adversarial use, or it will remain aspirational.

This is why the room repeatedly asked who carries responsibility. Is responsibility located in government, industry, investors, technical designers, international bodies, public institutions or civil society? The question became especially difficult because each actor possesses partial capacity and partial legitimacy. Governments can regulate but may lag; companies can build but may optimise for profit; investors can fund but may chase returns; international bodies can articulate norms but may struggle to enforce; technical systems can embed rules but cannot define legitimacy on their own. The inquiry did not produce an institutional answer. It revealed the inadequacy of existing responsibility structures for technologies whose effects are distributed, fast-moving and transnational.

The search for a universal entry point culminated in the discussion of children. The sentence “there isn’t a nation in the world that doesn’t care about their children” functioned as an inflection point because it identified a possible basis for global agreement that is neither purely technical nor purely ideological. Child safety offered a way to frame AI and digital governance around a near-universal human priority. This did not solve the governance problem, but it showed how participants tried to move from abstraction to political feasibility. In a world of divergent values, sovereignties and strategic interests, child protection may supply a limited but powerful foundation for alignment.

There isn’t a nation in the world that doesn’t care about their children.

WHAT EMERGED

What emerged from Session Three was a conditional theory of technological cohesion. Technology can support social bonds only where human purposes are defined, data can be trusted, basic needs are secured, judgement is preserved, responsibility is assigned and moral constraint is structurally embedded. Without those conditions, technology may intensify fragmentation, inequality, dependency, misinformation and strategic vulnerability. The strongest proposition survived because it explains the evidence across multiple domains: AI in care, digital exclusion, weather data, blockchain, humanitarian supply chains, defence, social media, labour displacement, education, basic needs, child safety and global governance all returned to the same underlying gap between capability and governability.

The competing interpretations each explain part of the session. The technological determinist interpretation captures the speed and power of AI, but it cannot explain the repeated insistence that technology is shaped by context, infrastructure and incentives. The governance interpretation captures institutional lag, but not the educational and psychological concern that human judgement itself may weaken. The infrastructure interpretation explains the return to water, food, energy and health, but not the questions about data provenance or agentic AI. The market interpretation explains why ethical systems may be outcompeted, but not the role of security threats or public legitimacy. The human-value interpretation explains the labour displacement concern, but not the full range of truth and accountability problems. The governance and human-agency proposition is stronger because it incorporates these partial explanations into a wider account of the conditions required for technology to serve human life.

What participants believed varied substantially. Some believed technical architecture can embed moral rules. Some believed blockchain or smart contracts may support accountability. Some believed AI for good can monitor harm and reduce abuse. Some believed governments are too slow and companies too powerful. Some believed market incentives must be altered. Some believed basic needs must come first. Some believed child safety may provide the basis for global consensus. Others remained more concerned with bad actors, national security, institutional lag or the erosion of critical thinking. What the evidence supports is not the sufficiency of any one of these remedies, but the inadequacy of treating technology as self-governing.

What can reasonably be inferred is that Session Three revealed a crisis of responsibility under conditions of technological acceleration. The room did not merely worry that technology might be misused. It worried that the systems capable of misuse are becoming too fast, too distributed, too opaque and too economically powerful for existing forms of governance to contain. This concern applied to governments, corporations, investors and users alike. It was not a rejection of technological progress. It was a recognition that progress without legitimate constraint may erode the human capacities and social conditions it claims to advance.

The consequences for education are direct. If critical thinking is not deliberately protected, AI may become a substitute for judgement before citizens have developed the ability to question it. Education cannot be limited to technical fluency. It must include epistemic discipline: how to test claims, understand uncertainty, recognise manipulation, compare evidence, distinguish authority from proof, and preserve the habit of independent judgement. Without this, societies may become more technologically capable while less able to deliberate about technology's uses.

The consequences for institutions are equally serious. Public trust may deteriorate if technological systems make consequential decisions without clear accountability. A healthcare algorithm that denies claims, a weather system that misleads farmers, a social platform that accelerates hate speech, a defence technology that changes thresholds for force, or an AI model that hallucinates authoritative information all raise the same question: who is answerable? Existing institutional forms distribute responsibility too thinly. The result may be a world in which harms occur through systems that no single actor fully controls or accepts responsibility for.

The consequences for markets are structural. If the most profitable or strategically advantageous technologies are less constrained than the most ethical ones, moral design will remain fragile. Investors and capital allocators therefore become part of the governance problem. The room's evidence suggests that capital cannot remain neutral where technological systems shape social reality, security and human agency. Capital allocation either reinforces technologies of trust or technologies of mistrust. The distinction may not be visible in conventional financial metrics, but it is central to long-term social cohesion.

The consequences for security are unavoidable. AI, digital finance, cyber systems and automated decision-making cannot be governed as if all actors share benign intentions. Democracies face a particular tension: they must preserve ethical constraints while competing against actors who may reject them. If they abandon moral limits, they undermine the values they claim to defend; if they ignore adversarial realities, they become vulnerable. This is one of the unresolved contradictions at the heart of technological governance.

The consequences for the social contract extend beyond regulation. If human labour value declines, if digital systems mediate access to services, if data becomes infrastructure, and if AI substitutes for judgement, societies will need to define human worth, participation and responsibility in terms broader than productivity. Cohesion cannot be sustained by technological efficiency alone. It requires people to experience themselves as agents within systems, not merely users, data points, recipients, risks or displaced workers.

QUESTIONS LEFT UNRESOLVED

01

The first unresolved question is who should govern advanced technology. Participants identified governments, companies, investors, technical experts, international organisations, civil society and hybrid bodies as possible actors, but no single actor combined legitimacy, competence, speed, enforcement power and public trust. The room recognised the responsibility gap without resolving institutional design.

02

A second unresolved question concerns human judgement. Can societies preserve critical thinking when AI becomes increasingly capable, convenient and authoritative? The participants agreed that education matters, but did not determine what form of education is sufficient, who should provide it, how early it must begin, or whether it can keep pace with adoption.

03

A third unresolved question concerns truth. Can trusted data infrastructure be made both technically reliable and democratically accountable? Blockchain, immutable records, satellite verification and real data assets may strengthen provenance, but they do not answer who defines the relevant data, who governs access, who audits the systems, or how citizens contest errors.

04

A fourth unresolved question concerns moral competition. Can ethical AI be made economically competitive, or will systems optimised for speed, profit and strategic advantage dominate? Participants repeatedly identified the problem but did not resolve how morality can be priced, rewarded or structurally embedded in markets.

- 05 A fifth unresolved question concerns agentic AI. Can human authority be embedded before autonomous systems become too complex or self-referential to govern? The room did not determine whether the risk is immediate or speculative, but the anxiety itself was evidence of declining confidence in traditional models of control.
- 06 A sixth unresolved question concerns basic needs. What level of water, food, health, energy, education and safety must be secured before advanced technological transformation can be socially legitimate? Participants recognised that technology cannot repair cohesion if foundational insecurity persists, but did not define the threshold at which technological progress becomes responsible rather than distracting.
- 07 A seventh unresolved question concerns human value. If labour is no longer the primary basis of economic worth for many people, what social, political and moral framework will sustain dignity, participation and responsibility? The room identified the risk but did not develop an alternative account of human value.
- 08 An eighth unresolved question concerns global governance. Can any transnational AI framework survive sovereignty, geopolitical competition, bad-faith actors and unequal technical capacity? Child safety emerged as a plausible entry point, but it remains uncertain whether a narrow consensus around children can extend to broader questions of labour, defence, data, markets and human agency.
- 09 The final unresolved question concerns whether societies can act before crisis. The Covid example suggested that risk can galvanise extraordinary collaboration, speed and institutional adaptation. Yet waiting for crisis means accepting preventable harm as the mechanism of coordination. Session Three left open whether technological responsibility can be built through foresight, or whether societies will continue to govern only after disruption has made inaction impossible.

Cross-Session Findings

The strongest finding to emerge across the Closed Door Sessions 2026 is that the three inquiries did not remain within their stated domains. Session One did not end as an inquiry into attention. Session Two did not end as an inquiry into capital. Session Three did not end as an inquiry into technology. In each case, participants moved from the existence of a capacity towards the failure of the structures required to convert that capacity into legitimate, accountable and socially useful action. The pattern was sufficiently consistent across the room to support the proposition that CDS 26 revealed a systems-conversion crisis: modern societies can generate awareness, capital and technological capability at scale, but are increasingly weak at converting those capacities into shared judgement, public benefit and humanly governed progress.

This proposition should not be accepted merely because it offers an elegant synthesis. It must be tested against the evidence. The three Sessions could be interpreted separately: Session One as a crisis of civic agency, Session Two as a crisis of capital accountability, and Session Three as a crisis of technological governance. Each of these interpretations is valid within its own field. Yet the behaviour of the room suggests that the participants were encountering the same structural difficulty in different forms. They repeatedly discovered that the surface resource was not sufficient. Information did not produce action. Capital did not produce impact. Capability did not produce progress. Between capacity and consequence lay missing or weakened mediating structures: trust, legitimacy, measurement, education, verification, accountability, community, moral constraint and long-term institutional stewardship.

In Session One, the room began with attention and arrived at agency. Participants did not conclude that people simply need to be informed. In many cases, they are already saturated with information. The problem lay in the broken passage from awareness to judgement and from judgement to action. The formulation “I know...” / “I will act upon...” captured this conversion gap with unusual economy. Its importance lies not in its phrasing, but in the conceptual distinction it records. Knowledge is not yet agency. Awareness is not yet responsibility. Exposure is not yet civic action. Participants repeatedly tested possible mediators: community, family, education, critical thinking, positive narratives, trusted messengers, courageous leadership, regulation and platform responsibility. The multiplicity of attempted repairs is itself evidence that no single institution now performs the conversion function reliably.

The same structure reappeared in Session Two. The room did not primarily discover a shortage of capital. Participants were surrounded by people and institutions capable of allocating capital, yet the question became whether capital could be trusted to produce durable public benefit. The inflection “we can all measure, but we are all measuring different things” exposed the problem. Accountability cannot be achieved by measurement alone if the objects of measurement are unaligned. Likewise, “ESG is paperwork” did not prove that ESG is useless; it identified the distinction between formal compliance and substantive consequence. The evidence supports the conclusion that capital requires conversion into legitimacy through problem definition, appropriate capital type, time horizon, local relevance, implementation and public accountability. Without these, capital may remain active while failing to become impact.

Session Three repeated the structure in technological form. The inquiry did not settle whether technology is cohesive or divisive. It showed that technology is conditional. The statement “technology is a tool” rejected technological determinism, but did not make technology neutral. Tools operating at the scale of AI, digital finance, social media, data infrastructure and defence systems are embedded in power, markets, cognition and governance. The unresolved problem was whether societies can convert capability into humanly governed progress before speed, profit, strategic competition or opacity make that conversion harder. “How do you create a mechanism where morality is part of the competition framework?” was therefore one of the strongest cross-session inflection points. It identified the difficulty of converting ethical intention into durable system design when markets reward capability faster than they reward restraint.

The systems-conversion proposition has greater explanatory power than a simple legitimacy-crisis interpretation. Legitimacy was present throughout: participants questioned political authority, institutional credibility, wealthy actors, platforms, technologists, governments and international bodies. Yet legitimacy alone does not explain the repeated concern with measurement, education, implementation, data provenance, time horizon and critical thinking. A legitimate actor may still fail if it cannot measure outcomes, implement at scale, verify data, sustain attention or govern technical systems. Legitimacy is therefore necessary but insufficient. It is one mediating structure within the wider conversion problem.

Nor is the evidence fully explained by an agency crisis. Session One and parts of Session Three certainly centred on agency: citizens may fail to act, young people may lose critical capacities, and societies may struggle to retain human authority over AI. Yet agency alone does not explain the failure of impact capital to produce accountable public benefit, or the difficulty of aligning ESG, philanthropy, venture capital, public policy and blended finance. Agency must be understood within institutions, markets and systems of measurement. The evidence does not support an individualised diagnosis. It supports a systemic one.

An incentive-crisis interpretation is also powerful but incomplete. Across the Sessions, incentives repeatedly appeared misaligned. Platforms reward outrage and emotional capture; capital markets reward speed, scale and return; technology markets reward capability, attention and strategic advantage; political systems reward short-term responsiveness over long-term stewardship. This explains much of the evidence. Yet incentive failure does not fully account for the loss of shared facts, the role of family and community, the need for trusted messengers, the difficulty of human judgement, or the search for moral legitimacy. Incentives describe what systems reward; they do not fully describe the social and institutional conditions through which people recognise, trust and act upon public purposes.

A governance-crisis interpretation also explains significant parts of the inquiry, especially Session Three. Governments were described as slow, mistrusted, politically volatile or technically behind. Participants questioned who should govern AI, data, platforms, capital and public goods. Yet governance alone risks narrowing the inquiry to formal authority and regulation. Session One showed that civic action depends on social and psychological mediators as much as formal governance. Session Two showed that capital legitimacy depends not only on regulation, but on measurement, local fit, implementation and trust. Governance is central, but it is not the whole condition.

Trust is perhaps the most tempting interpretation. Trust appeared everywhere: trust in facts, trust in messengers, trust in institutions, trust in capital, trust in data, trust in technology, trust in leaders. Yet trust alone is too static a category. The Sessions were not merely about trust being absent. They were about the absence of reliable processes through which trust can be earned, tested, measured, verified and converted into coordinated action. In Session One, trust had to become shared judgement. In Session Two, trust had to become accountable impact. In Session Three, trust had to become governed capability. The conversion proposition therefore explains trust as both a precondition and an outcome of functioning systems.

The cross-session evidence also requires attention to the room's behaviour. Participants repeatedly moved away from the surface wording of the questions. They resisted simple binaries. Attention became agency; public versus private capital became accountability; technology as good or bad became governance and human authority. This pattern was not incidental. It suggests that, when placed in a cross-sector environment, participants recognised that the stated problems were entry points into failures of mediation. The room's hesitations and disagreements were especially revealing. Participants hesitated over whether emotional influence can be ethical, whether ESG can be repaired, whether blockchain can govern trust, whether private capital can act legitimately, whether leadership can be courageous and democratic, whether AI governance is already too late, and whether crisis is now the only reliable mechanism for alignment. These unresolved tensions indicate that the inquiry had reached structural questions rather than merely operational ones.

The sociological evidence points to weakened relational infrastructures. Session One repeatedly returned to family, community, human interaction, local belonging and shared facts. Session Two returned to local capital, livelihoods, basic needs and the social legitimacy of wealth. Session Three returned to child safety, digital exclusion, human value, critical thinking and basic infrastructure. In each case, social cohesion depended on relationships that cannot be manufactured by information, capital or technology alone. What participants believed differed: some looked to local community, some to online community, some to leadership, some to global frameworks. What the evidence supports is that public action depends on relational structures capable of carrying obligation across difference and time.

The political-science evidence points to a crisis in the capacity of public authority to organise long-term responsibility. Democratic leaders face short-term incentives and public backlash; public policy is vulnerable to electoral reversal; governments may be technically behind; international agreements may be difficult to enforce; private actors may possess capacity without public legitimacy. The room did not endorse authoritarian efficiency, though participants did compare Western short-termism unfavourably with longer-term state strategies elsewhere. The stronger inference is that democratic societies face a particular conversion problem: they must translate long-term risks into legitimate action without abandoning pluralism, consent or accountability.

The economic evidence points to capital and market systems that do not automatically internalise public benefit. In Session One, attention markets reward negative or emotionally charged narratives. In Session Two, capital may reward speed, novelty, scale or compliance rather than durable impact. In Session Three, technology markets may reward capability and profit before moral constraint. The phrase “you cannot be green if you are in the red” grounded this economic dimension in material conditions. Public benefit cannot be sustained where economic survival is ignored. Yet the inverse is also present: economic systems that ignore public benefit may create insecurity, displacement and conflict. The evidence supports an economic interpretation in which public goods must be priced, measured or institutionally protected, but also shows that not all public goods can be reduced to market return.

The diplomatic evidence shows that the conversion crisis is not confined to domestic institutions. Digital sovereignty, geopolitical mistrust, GCC philanthropic retrenchment, China's long-term planning, Western policy volatility, AI governance and child safety all appeared as cross-border problems. The proposed use of child safety as a universal basis for AI governance was significant because it sought a minimal shared ground across cultures and states. The sentence "there isn't a nation in the world that doesn't care about their children" functioned as an attempt to identify a diplomatic conversion mechanism: a shared human priority capable of translating technological anxiety into global agreement. Whether that mechanism can bear the weight placed upon it remains unresolved.

The security evidence widened the definition of security itself. Participants repeatedly linked misinformation, social cohesion, water, food, health, migration, defence, AI, digital finance and bad actors. Security was not confined to military preparedness. It appeared wherever failed conversion could generate instability: awareness becoming polarisation rather than judgement; capital failing to secure basic needs; technology empowering adversaries or eroding human control. The evidence supports the inference that future security depends not only on deterrence or defence capability, but on the integrity of social, informational, financial and technological infrastructures.

The technological evidence is especially important because it both supports and tests the proposition. Technology appeared as a potential conversion tool: AI could measure impact, detect hate speech, support healthcare, improve weather prediction or verify data. Yet technology also appeared as a source of further conversion failure: hallucination, opacity, dependency, market competition, agentic autonomy and governance lag. "I don't believe what I cannot prove with extremely high probability" identified the need for verification; "government mindsets are seven years behind" identified institutional lag; "how are we going to make sure that we as humanity always have an upper hand over the agentic silicon lifeform?" identified the escalation from tool use to human authority. The evidence therefore does not support technological rejection. It supports the view that technology can repair conversion failures only if it is itself governed by trustworthy conversion structures.

The social-psychological evidence explains why information, capital and technology fail when treated as rational inputs. Participants repeatedly acknowledged that people act through fear, belonging, aspiration, comparison, trust, grief, hope, status, family and perceived stake. The Covid vaccine example, in which data failed where a human story might have moved behaviour, was central because it showed that cognition is socially and emotionally mediated. Similarly, capital is trusted when it appears aligned with human need, and technology is accepted when it appears to preserve agency rather than displace it. The evidence supports the inference that conversion is not only institutional; it is psychological and relational.

From a systems perspective, the three Sessions describe failures of feedback between interdependent systems. Information systems produce awareness, but civic systems do not convert it into judgement. Capital systems produce allocation, but social systems do not always receive it as legitimate benefit. Technological systems produce capability, but political and ethical systems do not govern it at the same pace. The result is not a single institutional breakdown, but a coordination failure across systems operating at different speeds, with different incentives and different forms of legitimacy. This explains why participants repeatedly proposed hybrid solutions: education and regulation, capital and accountability, technology and human judgement, local legitimacy and global scale.

The contradictory evidence should be taken seriously. Participants did not describe a world without remedies. They identified practical mechanisms in every Session: critical thinking, trusted messengers, community formation, patient capital, blended finance, local investment, implementation funding, human-in-the-loop design, data verification, smart contracts, child safety and AI governance. This evidence prevents the systems-conversion proposition from becoming a theory of collapse. The condition is not total failure, but fragile and uneven mediation. Societies still possess institutions and practices that can convert capacity into public benefit, but those institutions are under strain and do not operate reliably at the scale or speed required.

There is also a risk that the proposition over-integrates distinct problems. Civic agency, capital accountability and technological governance are not identical. Each has its own history, institutional structure and field of practice. The evidence would not support treating them as interchangeable. The proposition survives this challenge because it does not erase their differences. It identifies a shared structural pattern across them. In each Session, an abundant capacity was shown to require mediating structures that are contested, weakened or insufficiently developed. The pattern is common; the mechanisms are domain-specific.

Another challenge concerns power. Systems-conversion language can sound neutral, as though society simply lacks coordination. Yet some actors benefit from failed conversion. Platforms may benefit from attention without judgement. Capital may benefit from public-good language without accountability. Technological firms may benefit from capability before governance. Political actors may benefit from polarisation. Adversaries may benefit from mistrust. The evidence therefore requires that the systems-conversion proposition include conflict as well as coordination. The problem is not merely that systems fail to connect; it is that existing incentives may reward some actors for preventing connection.

This does not displace the central proposition, but it qualifies it. Conversion failure should not be understood only as the result of complexity, fragmentation or institutional weakness. In some cases, the gap between capacity and accountability may be advantageous to actors already positioned within the system. Attention without judgement can reward platforms and political communicators that benefit from emotional capture. Capital without clear accountability can allow public-good language to confer legitimacy before consequence is demonstrated. Technological capability before governance can advantage firms, investors or states that move faster than regulatory, ethical or civic scrutiny. Mistrust, fragmentation and delay may also serve adversarial actors. The evidence therefore suggests that failed conversion is not always accidental. It may also reflect interests that are protected by the absence of effective mediation.

The consequences, if the proposition is correct, are significant. Increasing awareness will not by itself produce civic agency. Increasing capital will not by itself produce impact. Increasing technological capability will not by itself produce human progress. Institutions that optimise for outputs alone, whether content, funding or innovation, will misdiagnose the problem. The relevant question becomes how those outputs are converted: what forms of trust, evidence, legitimacy, measurement, education, governance and implementation allow capacity to become public benefit.

For democratic societies, the implication is that long-term public action depends on rebuilding conversion mechanisms before crisis forces alignment. Participants repeatedly noted that crisis can galvanise action, as in Covid, defence shocks or acute security threats. Yet crisis-driven conversion is costly, reactive and often inequitable. If societies can only coordinate after harm becomes undeniable, they will repeatedly govern from emergency rather than foresight. The unresolved question is whether deliberative systems can generate sufficient urgency without requiring catastrophe as proof.

For institutional investors and capital allocators, the implication is that legitimacy must be treated as part of impact. Capital cannot claim public benefit through label, intention or allocation alone. It must show problem fit, beneficiary relevance, measurement discipline, implementation capacity and long-term responsibility. This does not mean all capital must become philanthropic. It means that claims of social benefit require conversion structures that make benefit observable, accountable and trusted.

For technologists and policymakers, the implication is that governance cannot be postponed until after capability matures. If moral constraint is not embedded before systems become profitable, autonomous or geopolitically necessary, it may become difficult to impose later. Yet governance cannot be purely restrictive, because societies also need technologies that solve basic problems and defend against real threats. The evidence supports a model of technological responsibility that is neither accelerationist nor prohibitive, but conditional on human judgement, verification, legitimacy and security realism.

For civil society, the implication is that community, education and trusted messengers remain central but insufficient. They are not nostalgic remedies; they are conversion infrastructures. They help translate information into meaning, meaning into judgement and judgement into action. Yet they must operate in environments shaped by platforms, markets and geopolitical forces. Civil society cannot repair conversion alone, but without it, formal systems lack the relational grounding necessary for legitimacy.

The questions that survived all three Sessions are therefore not merely practical. They are constitutional in the broadest sense. Who has the legitimacy to act on behalf of the public good? How does awareness become responsible action? How does capital become accountable impact? How does technology remain under human authority? What should be measured, and who decides? How can long-term responsibility survive short-term incentives? How can societies act before crisis forces alignment? These questions remained because none of the Sessions could resolve them within a single domain.

How do you create a mechanism where morality is part of the competition framework?

The deepest uncertainty is whether existing institutions can repair the systems-conversion crisis or whether new institutional forms are required. The CDS format itself offers limited but relevant evidence. For a short period, it created a structured environment in which people from different sectors could test assumptions, disagree under conditions of trust, and move from surface questions to underlying propositions. It did not solve the condition it diagnosed. It did, however, demonstrate that when knowledge, trust, disagreement and responsibility are deliberately held in the same room, conversion becomes at least temporarily possible. That observation should not be overstated. It is not a conclusion. It is evidence that the crisis identified across the Sessions is not only a matter of content, capital or technology, but of the conditions under which people and institutions are able to think and act together.

The Questions We Did Not Finish

The Closed Door Sessions did not end by producing a settled programme of action. They produced a more precise understanding of why action remains difficult even where awareness, capital and technological capability are present. The questions that survived the inquiry were not those that lacked discussion. They were the questions to which repeated discussion produced only partial, competing or unstable answers. They remained active because each Session returned to them in different forms: through civic agency, capital accountability, technological governance, legitimacy, trust, measurement, leadership, basic needs and human judgement.

OI

The first unfinished question is who has the legitimacy to act on behalf of the public good. This appeared first in the difficulty of civic motivation: if citizens no longer share facts, trusted institutions or stable communities, then public-good action cannot simply be addressed to “the nation” or assumed to pass through conventional authority. The question reappeared in the capital discussion, where private wealth, philanthropy, venture capital, public policy and blended finance were all treated as potentially useful but not automatically legitimate. It appeared again in the technology discussion, where participants repeatedly asked whether governments, companies, investors, international bodies, technical experts or civil society should govern AI, data and digital systems. The recurrence of the question across all three domains suggests that legitimacy is no longer attached securely to inherited institutional categories.

Proposed answers proved insufficient because each candidate source of legitimacy carried its own deficit. Governments possess public authority but may be slow, mistrusted, politically volatile or technically behind. Private capital can move quickly but may be suspected of self-interest, extraction or reputational motive. Technology companies possess capability but are shaped by market incentives. International bodies can convene but may lack enforcement power. Communities possess proximity and trust but may lack scale. The evidence supports the conclusion that legitimacy must now be assembled rather than assumed, but it does not resolve how such assembly should occur. Participants agreed that no single actor can carry the burden alone. They disagreed over whether legitimacy should be restored through public institutions, distributed through local community, earned by outcomes, embedded in technical systems or shared through hybrid governance.

O2

The second unfinished question is how awareness becomes responsible action. This question was explicit in Session One, but it did not remain there. The inflection “I know...” / “I will act upon...” identified the gap between information and agency. It reappeared in Session Two when participants asked how impact capital moves from intention to implementation, and in Session Three when the room asked whether technological knowledge, including AI outputs and data systems, produces judgement or dependency. Across the Sessions, participants recognised that knowing more does not necessarily produce acting better. Societies may be more informed, more connected and more technologically capable while remaining unable to coordinate sustained public action.

The proposed answers were all partial. Education was repeatedly invoked, but the room did not settle whether education means critical thinking, civic formation, media literacy, moral reasoning, family-based values or technical fluency. Storytelling was proposed, but the room remained uneasy about the boundary between ethical persuasion and manipulation. Community was proposed, but participants disagreed over whether community is local, digital, national, familial or value-based. Leadership was proposed, but courageous leadership may be polarising, unpopular or democratically fragile. What remains unknown is how to build a conversion pathway from attention to judgement and from judgement to action without relying either on manipulation, crisis or coercion.

03

The third unfinished question is what should be measured, and who decides. This question appeared most explicitly in the capital discussion, where the statement “we can all measure, but we are all measuring different things” identified a central accountability problem. It was also present in Session One, where participants implicitly questioned whether attention, awareness, engagement, trust, judgement or action should count as civic success. It reappeared in Session Three through data provenance, AI hallucination, blockchain, real data assets and the need to verify technological outputs. Measurement became the point at which trust, legitimacy and accountability met, but it also became the point at which disagreement became unavoidable.

Participants broadly agreed that rhetoric without measurement is inadequate. They also agreed that existing metrics often distort the thing they claim to assess. ESG was criticised as “paperwork” where compliance becomes detached from consequence. Social media engagement was treated as a misleading proxy for civic value. AI outputs were treated as insufficient unless verifiable. Yet no consensus emerged over the proper object of measurement. Should public-good systems measure attention, trust, action, avoided harm, livelihoods, health, conflict reduction, local resilience, safety, standard of living, data integrity or human judgement? Each answer privileges a different theory of value. The evidence supports the need for shared evaluative architecture, but leaves unresolved whether such architecture can be rigorous without becoming reductive.

04

The fourth unfinished question is how long-term responsibility can survive short-term systems. This question appeared in Session One through the conflict between short-term attention and long-term civic problems. It appeared in Session Two through unstable policy cycles, venture timelines, subsidy dependence and the need for patient capital. It appeared in Session Three through rapid technological acceleration, public-sector lag and market competition around AI. Across the inquiry, participants repeatedly identified a temporal mismatch: long-term problems require continuity, while the systems acting upon them reward immediacy.

The proposed answers remained unstable. Some participants looked to institutions as stabilising baselines. Others looked to community and generational interest. Others looked to long-term capital, infrastructure structures, public-private partnerships or non-Western state strategies. Some looked to moral leadership capable of withstanding unpopularity. Yet each answer met resistance. Institutions may be mistrusted or slow. Communities may be fragmented. Capital may flee under pressure. Democratic leaders must remain electable. Authoritarian long-termism raises political and moral dangers. The evidence supports the claim that time horizon is a structural problem, not a matter of will alone. What remains unresolved is how democratic, market-based and plural societies can create durable long-term action without suppressing contestation, consent or responsiveness.

05

The fifth unfinished question is whether trust can be restored, or whether systems must operate under conditions of permanent distrust. Trust appeared in every Session, but never as a simple remedy. In Session One, trust was needed for shared facts, messengers and civic judgement. In Session Two, trust was required for capital to be received as public benefit rather than self-interest. In Session Three, trust depended on data, human oversight, technical reliability and legitimate governance. The room repeatedly returned to trust because it is a precondition for conversion: awareness requires trusted interpretation; capital requires trusted accountability; technology requires trusted verification.

Yet the inquiry did not establish how trust is rebuilt. Some participants believed trust can be earned through outcomes. Others looked to family, community, local proximity or trusted messengers. Others looked to measurement, transparency, data provenance or moral operating systems. The problem is that each mechanism can also fail. Outcomes can be contested; communities can become echo chambers; messengers can be captured by incentives; measurement can become bureaucracy; technical verification can lack democratic legitimacy. The evidence supports the conclusion that trust must be earned through repeated, accountable practice, but it does not resolve which institutions are capable of sustaining that practice at scale.

06

The sixth unfinished question is whether public-good actors can compete without becoming what they oppose. This question first appeared in the attention discussion, where participants recognised that evidence-led institutions often lose to emotional, performative or manipulative actors. The phrase “we were taking a knife to a gunfight” marked the recognition that truthful communication may be structurally weaker than emotionally optimised communication. It reappeared in Session Two, where impact capital must compete with ordinary capital, ESG backlash, subsidy instability and market pressures. It appeared again in Session Three, where ethical AI may be outpaced by faster, more profitable or less constrained systems. Across all three Sessions, participants confronted the same problem: public-good action must compete inside systems whose incentives may be hostile to public-good behaviour.

The proposed answers were incomplete because they risked reproducing the problem. More emotionally powerful civic communication may become manipulation. More profitable sustainability may narrow attention to investable problems. More competitive AI may sacrifice moral constraint. More forceful security responses may undermine the values they defend. Participants agreed that public-good actors cannot ignore incentives, markets, narratives or strategic competition. They disagreed over how far such actors can adapt to those logics without being absorbed by them. What remains unresolved is whether ethical conduct can be made competitive without diluting the ethical constraint that gives it legitimacy.

07

The seventh unfinished question is how local legitimacy can be reconciled with global scale. This question appeared in Session One through community, family, locality and digital belonging. It appeared in Session Two through local capital, local solutions, Global South needs, infrastructure scale and impact measurement across contexts. It appeared in Session Three through digital sovereignty, global AI governance, child safety, basic needs and technological systems that cross borders. Participants repeatedly returned to the fact that legitimacy is often local, while the problems under discussion are often global.

The proposed answers remain incomplete. Local community can generate trust, but may lack scale. Global institutions can coordinate, but may lack proximity. Digital systems can cross borders, but may flatten cultural difference. Capital can scale infrastructure, but may impose external metrics. Universal frames such as child safety may create agreement, but may not extend to labour, data, defence, sovereignty or economic justice. The evidence supports the need for multi-level governance, but this phrase does not itself solve the problem. What remains unknown is how to build systems that are locally legitimate, globally coordinated and capable of acting before harm spreads across borders.

08

The eighth unfinished question is how societies should understand human value when technology, capital and institutions increasingly treat people as users, data points, labour inputs, beneficiaries, risks or consumers. This question appeared in Session One through concern about loss of human interaction, AI companionship and the weakening of relational formation. It appeared in Session Two through livelihoods, basic needs, survival-brain dynamics and the danger that impact frameworks measure systems rather than human flourishing. It appeared in Session Three most explicitly through the possibility that AI and automation may reduce the economic value of human labour. The claim that “the value of being human is diminishing” was not accepted as settled fact, but it forced the inquiry beyond access, efficiency and productivity.

Participants did not resolve this question because existing frameworks remained partial. Labour value is inadequate if technology reduces demand for human work. Consumer value is inadequate if it reduces people to market behaviour. Beneficiary status is inadequate if it removes agency. Data value is inadequate if it turns persons into extractable inputs. Civic value is difficult to sustain where shared facts and communities weaken. The evidence supports the need to define human value beyond productivity and consumption, but the Sessions did not produce a sufficient account of what replaces those measures. What remains unresolved is how institutions should value human judgement, care, belonging, dignity and responsibility in systems increasingly optimised for efficiency and scale.

09

The ninth unfinished question is whether crisis is now the only reliable mechanism for alignment. This question appeared repeatedly, sometimes explicitly and sometimes by implication. In Session One, participants asked how to make long-term issues visible before they become emergencies. In Session Two, capital often moved when risk became legible, measurable or urgent. In Session Three, Covid, defence, AI and security threats showed that crisis can collapse barriers, accelerate decision-making and force collaboration. The Covid vaccine example demonstrated that extraordinary alignment is possible when risk is immediate, shared and publicly undeniable.

The proposed answers sought alternatives to crisis: foresight, institutions, patient capital, education, community, child safety, measurement, shared risk and moral leadership. Yet participants repeatedly returned to the difficulty of creating urgency without emergency. Long-term climate, food, water, AI and social cohesion risks do not always generate action before visible harm. The evidence supports the claim that crisis remains one of the most powerful conversion mechanisms available to modern societies. What remains unresolved is whether societies can build anticipatory systems strong enough to act before crisis supplies the political, emotional and financial conditions for action.

10

The tenth unfinished question is whether shared facts can be restored, or whether collective action must now be built without them. This question was most explicit in Session One, where “we can’t have a common conversation because we are not operating from the same starting point” identified the collapse of shared factual ground. It reappeared in Session Two through disagreement over measurement and impact, because actors cannot coordinate if they do not accept common standards of evidence. It reappeared in Session Three through AI hallucination, data provenance and the need to verify claims with extremely high probability. Across the Sessions, shared facts were not treated as an intellectual luxury. They were treated as infrastructure for action.

The proposed answers differed. Some looked to education and critical thinking. Others looked to trusted messengers, technical verification, blockchain, real data assets, corporate responsibility, regulation or self-accountability. Each answer addresses part of the problem. None resolves whether shared facts can be restored in an environment where platforms, states, markets, AI systems and communities all produce competing realities. The evidence supports the conclusion that action requires some shared basis of trust and verification. It does not resolve whether this basis will be institutional, technical, communal, legal or plural.

II

The eleventh unfinished question is what role private actors should play when they are capable but not automatically legitimate. This question arose in all three Sessions. In Session One, participants looked beyond government and formal institutions to athletes, influencers, cultural leaders, religious figures, platforms and companies as possible civic messengers. In Session Two, private capital, family offices, philanthropies and investors were treated as essential but legitimacy-limited actors. In Session Three, companies, technologists and investors were among the actors most capable of shaping AI and digital infrastructure, but they were not treated as democratically authorised to govern the social consequences alone.

The unresolved issue is therefore not only one of legitimacy, but also one of interest. Private actors may bring capital, expertise, infrastructure, speed and implementation capacity that public systems cannot easily supply. At the same time, they may carry incentives to preserve influence, manage reputational risk, shape agendas, retain control over infrastructure, or define success in terms that align with their own institutional position. The question is not whether such interests disqualify private actors from public-good work, but how they should be acknowledged, governed and held in relation to public responsibility.

The evidence supports the conclusion that private actors are now indispensable to public-good action. The evidence also supports the conclusion that indispensability is not legitimacy. This is why the question survived. Participants could not simply exclude private actors, because they possess capital, capability, speed, infrastructure, influence and expertise. They could not simply empower them, because public-good action requires accountability, legitimacy and safeguards against self-interest. What remains unresolved is how private capacity should be brought into public responsibility without privatising authority.

12

The twelfth unfinished question is whether the mediating structures identified across the Sessions can operate at sufficient scale. Across the inquiry, participants proposed education, community, trusted messengers, measurement, patient capital, blended finance, local legitimacy, implementation funding, human-in-the-loop design, data verification, moral constraints, child safety and AI governance. These were not trivial proposals. They were evidence that the room did not regard the systems-conversion crisis as total or irreversible. Yet the repeated return to the same gaps suggests that existing mediators are fragmented, underpowered or domain-specific.

The unresolved issue is scale. Local community can build trust, but not necessarily govern AI. Measurement can improve impact capital, but not necessarily repair civic agency. Child safety may support AI governance, but not resolve digital sovereignty or labour displacement. Education may build critical thinking, but may not counter platform incentives quickly enough. Blended finance can support implementation, but may not repair democratic legitimacy. The evidence supports the need for mediating structures, but leaves unresolved whether they can be connected into a system capable of matching the scale of the problems discussed.

The final unfinished question concerns The CDS itself. The room's behaviour across the Sessions suggests that structured, cross-sector inquiry can temporarily reconnect knowledge, disagreement, trust and responsibility. Participants moved beyond surface questions, challenged assumptions, identified competing interpretations and, at times, acknowledged uncertainty rather than forcing premature agreement. This does not prove that CDS 26 solved the systems-conversion crisis it diagnosed. It does, however, raise a question that remained implicit throughout: what kinds of spaces, institutions or processes are now needed to allow societies to think and act across domains?

This question survived because the format itself became evidence. The room did not produce uniform answers, but it did produce a shared capacity to recognise deeper problems beneath sector-specific language. Historians, political theorists, economists, diplomats, security specialists, technologists and investors may interpret that modestly: as a useful deliberative exercise. They may also interpret it more structurally: as evidence that modern systems require intermediary spaces where authority is not assumed, expertise is not siloed, disagreement is protected and public purpose can be tested before action is claimed. What remains unresolved is whether such spaces can influence institutions, capital, technology and policy beyond the room in which the inquiry occurred.

What Questions Do We Ask Next?

The questions that should shape the next stage of inquiry are not simply the questions left unanswered. They are the questions whose answers would alter how the other questions are understood. Across the Closed Door Sessions, many issues remained unresolved: trust, legitimacy, measurement, AI governance, civic agency, private capital, public responsibility, shared facts, basic needs and human value. Yet not all unresolved questions carry equal explanatory weight. Some are derivative; others are foundational. The frontier of inquiry lies where a question affects several systems at once and where existing evidence is insufficient because the problem is not only empirical, but institutional, political, ethical and epistemic.

OI

The first question is how societies convert awareness into responsible collective action without relying on crisis, manipulation or coercion. This question has priority because it underlies civic life, capital allocation and technological governance alike. If people, institutions and markets cannot act until harm becomes visible, measurable or emotionally undeniable, then public-good systems remain structurally reactive. The Sessions repeatedly suggested that crisis can concentrate attention and force alignment, but they also showed that crisis-driven action is late, uneven and often costly. Before this question can be answered, future inquiry must distinguish between awareness, concern, judgement, commitment and sustained action. It must test the assumption that information can be converted into agency through better communication alone. The evidence currently missing is comparative: which forms of social, institutional or narrative mediation produce durable action before crisis, and under what conditions do they fail? This question requires further evidence, better measurement and political judgement, because it concerns both behavioural mechanisms and the legitimacy of mobilisation.

02

The second question is what forms of legitimacy are adequate for public-good action when authority is distributed across governments, private capital, technology companies, communities and international bodies. This question matters because many of the actors with capacity lack uncontested legitimacy, while many of the actors with legitimacy lack speed, capital or technical competence. The room repeatedly returned to this mismatch without resolving it. Governments, investors, companies, experts, local communities and international organisations each appeared necessary and insufficient. Before this question can be answered, inquiry must clarify whether legitimacy is being treated as democratic consent, competence, proximity, transparency, outcomes, legality, moral authority or some combination of these. The assumptions requiring testing are that legitimacy can be earned through results, that hybrid governance can combine legitimacy and capability, and that public trust can be transferred between actors. The missing evidence is institutional: which arrangements allow distributed authority to act without dissolving accountability? This question requires institutional experimentation and political judgement, not merely further discussion.

03

The third question is what should be measured when public benefit is claimed. This question should shape future inquiry because measurement is the point at which rhetoric, accountability and legitimacy meet. The Sessions repeatedly showed that measurement can clarify impact, but also distort it. Attention metrics do not measure civic judgement. ESG reporting does not necessarily measure social consequence. AI outputs do not establish truth without provenance and verification. Before this question can be answered, inquiry must identify the object of measurement: awareness, trust, action, avoided harm, health, livelihoods, security, resilience, dignity, human judgement, or institutional legitimacy. The assumption requiring testing is that complex public benefit can be reduced to stable metrics without losing its meaning. The missing evidence concerns both methodology and governance: who defines the metric, who audits it, who is accountable for error, and how measurement changes behaviour once introduced. This question requires better measurement, philosophical clarification and institutional design.

04

The fourth question is how moral constraint can survive competition. This question sits at the frontier because it appeared in every domain. Civic communication must compete with manipulation; impact capital must compete with higher-return alternatives; ethical AI must compete with faster, less constrained systems; democratic states must compete with adversaries who may not share their restraints. The question is not whether morality matters. Participants broadly accepted that it does. The unresolved issue is whether moral constraint can be made structurally durable in systems that reward speed, attention, profit or strategic advantage. Before this can be answered, inquiry must distinguish between moral aspiration, legal obligation, market incentive, technical design and strategic necessity. The assumptions requiring testing are that ethical systems can be made competitive without being diluted, and that markets or technologies can internalise moral limits without external governance. The missing evidence lies in cases where restraint has survived competitive pressure. This question requires institutional experimentation, economic analysis, security analysis and ethical clarification.

05

The fifth question is whether shared facts can be rebuilt, or whether coordination must now be designed for conditions of permanent epistemic fragmentation. This question matters because the Sessions repeatedly treated shared facts as infrastructure for action, but did not establish whether that infrastructure can be restored. Civic deliberation, impact measurement and AI governance all depend on some accepted basis of evidence. Yet platforms, political polarisation, institutional mistrust, competing metrics and machine-generated outputs all weaken that basis. Before this question can be answered, inquiry must distinguish factual agreement from trust in institutions, and data verification from shared interpretation. The assumption requiring testing is that better evidence or better technology can restore common ground. The missing evidence concerns environments where plural publics retain the ability to act together despite disagreement over sources, authorities or narratives. This question requires further evidence, better measurement, technological inquiry and political theory.

06

The sixth question is how human judgement should be preserved when systems increasingly automate interpretation, allocation and decision-making. This question should shape future inquiry because it links education, AI, capital, governance and democracy. The Sessions repeatedly returned to critical thinking, AI dependence, data verification, institutional lag and the risk of outsourcing judgement to systems whose outputs appear authoritative. Before this question can be answered, inquiry must define what kind of judgement is at stake: epistemic judgement, moral judgement, civic judgement, investment judgement, clinical judgement or strategic judgement. The assumption requiring testing is that human oversight remains meaningful once systems become sufficiently complex, fast or embedded. The missing evidence concerns where human-in-the-loop systems genuinely preserve authority, and where they merely provide procedural reassurance. This question requires technological inquiry, educational research, institutional experimentation and ethical clarification.

07

The seventh question is how local legitimacy can be reconciled with global scale. This question belongs at the frontier because the Sessions repeatedly showed that trust is often local while problems are transnational. Community may motivate action, but climate, AI, migration, capital flows, digital systems, defence and supply chains exceed local boundaries. Capital may scale solutions, but risk imposing external priorities. Global governance may coordinate action, but lack proximity and trust. Before this question can be answered, inquiry must distinguish between scale of problem, scale of decision, scale of legitimacy and scale of implementation. The assumptions requiring testing are that global frameworks can respect local difference, and that local legitimacy can be aggregated without being diluted. The missing evidence lies in institutional forms that preserve local accountability while acting across borders. This question requires diplomatic inquiry, institutional experimentation, sociology and political judgement.

08

The eighth question is what forms of capital can legitimately act where public systems fail. This question matters because private actors are increasingly indispensable to public-good outcomes, but indispensability does not confer authority. Across the inquiry, private capital, family offices, philanthropy, investors, companies and technologists appeared as capable actors whose role remained unresolved. Before this question can be answered, inquiry must distinguish between capital that discovers, de-risks, scales, subsidises, governs, owns or maintains public-good infrastructure. The assumption requiring testing is that capital can be made accountable through measurement and outcomes alone. The missing evidence concerns long-term cases where private capital has acted in public-good domains without undermining public responsibility, local legitimacy or democratic accountability. This question requires economic analysis, governance design, legal inquiry and political judgement.

09

The ninth question is how basic needs should be understood within discussions of technology, capital and social cohesion. This question has priority because the Sessions repeatedly returned from advanced systems to food, water, health, safety, energy, education and livelihoods. These were not peripheral humanitarian concerns; they appeared as conditions for agency, legitimacy and security. Before this question can be answered, inquiry must clarify whether basic needs are being treated as moral obligations, security infrastructure, investment categories, development indicators or foundations of democratic participation. The assumption requiring testing is that advanced technological or financial systems can generate cohesion while foundational insecurity persists. The missing evidence concerns how unmet basic needs interact with misinformation, migration, conflict, institutional trust and technological adoption. This question requires security studies, development economics, sociology and measurement design.

IO

The tenth question is how human value should be understood if labour, productivity and consumption become less adequate measures of social worth. This question did not dominate every discussion, but where it appeared it altered the terms of the inquiry. AI, automation, welfare innovation, impact measurement and civic agency all raise the question of whether people are being treated as workers, consumers, users, beneficiaries, data sources, risks or citizens. Before this question can be answered, inquiry must distinguish economic value from civic value, relational value, moral value and political agency. The assumptions requiring testing are that income replacement is sufficient where labour value declines, and that efficiency gains necessarily translate into human flourishing. The missing evidence concerns how societies preserve dignity, participation and responsibility when economic roles change. This question requires philosophical and ethical clarification, social psychology, political theory and institutional experimentation.

II

The eleventh question is what kinds of spaces or institutions allow complex problems to be thought across boundaries before action is claimed. This question matters because the room's behaviour itself became evidence. Participants advanced the inquiry when disciplinary expertise was brought into contact with other forms of expertise, and when disagreement was held long enough to reveal deeper problems. Convergence occurred around diagnosis more readily than prescription. Before this question can be answered, inquiry must distinguish deliberative space, advisory body, governance institution, investment forum, diplomatic process and civic assembly. The assumption requiring testing is that cross-sector dialogue can influence institutions rather than remain episodic. The missing evidence concerns the conditions under which inquiry changes subsequent decisions, capital allocation, governance design or public narratives. This question requires institutional experimentation and historical analysis.

The final question is whether societies can build foresight institutions strong enough to act before harm becomes undeniable. This question cuts across all the others. It concerns awareness before crisis, capital before collapse, technology before loss of control, security before conflict and legitimacy before institutional failure. The Sessions repeatedly showed that risk becomes actionable when it becomes visible, measurable, emotional or immediate. The frontier question is whether foresight can create sufficient authority without waiting for emergency. Before this question can be answered, inquiry must distinguish prediction, preparedness, prevention, precaution and pre-emption. The assumptions requiring testing are that societies can act on probabilistic evidence, and that publics will accept costs for harms that have not yet occurred. The missing evidence concerns political systems, capital structures and governance models that have sustained preventive action over long periods. This question requires further evidence, better measurement, institutional experimentation and political judgement.

These questions define the next frontier because each one governs several others. Legitimacy shapes who may measure, govern and act. Measurement shapes what capital, technology and public institutions pursue. Shared facts shape whether disagreement can become deliberation. Human judgement shapes whether AI, capital and public systems remain accountable. Local legitimacy shapes whether global action can be trusted. Basic needs shape whether agency is possible. Foresight shapes whether action can occur before crisis supplies the missing urgency. The next stage of inquiry therefore cannot begin from a single sector. It must begin from the points at which sectors fail to convert their own capacities into responsible shared action.



Closed Door Sessions

PUBLISHED ON SIMPLE

[Find the latest version here.](#)