



ESCAPING THE MULTIPOLAR TRAP IN GLOBAL CLIMATE NEGOTIATIONS

A DELIBERATIVE NEGOTIATION TECHNOLOGY AND SIMPOL-BASED SIMULATION

Ernest Thiessen, John Bunzl and Leland Beaumont

About the Authors



ERNEST M. THIESSEN

Ernest M. Thiessen is a Canadian engineer and peace-technology innovator and the inventor of the Smartsettle negotiation and collaboration systems. He holds a PhD in Civil and Environmental Engineering from Cornell University and has worked for over three decades at the intersection of decision science, conflict resolution and public-interest negotiation. His recent work focuses on the application of multi-stakeholder consensus optimisation to global challenges including climate change and post-conflict reconstruction.



JOHN M. BUNZL

John M. Bunzl is the co-founder and trustee of Simpol (Simultaneous Policy), an international citizens' movement dedicated to solving global collective-action problems through coordinated, simultaneous implementation of public policy across nations. He is the author of *The Simpol Solution* (with Nick Duffell), and has spent more than two decades advancing democratic strategies for overcoming multipolar traps in climate and other global governance issues. See <https://simpol.org>



LELAND R. BEAUMONT

Leland R. Beaumont is an electrical engineer and a systems thinker focused on civic innovation, wisdom studies, and large-scale social coordination. He holds Bachelor's and Master's degrees in Electrical Engineering and spent the early part of his career at Bell Laboratories and Lucent Technologies. Since retiring from industry in 2001, he has applied engineering and systems-analysis methods to the study of human behavior, emotional competency, wisdom, and democratic capacity. He is the creator of the Emotional Competency website, The Wise Path, the Wise Living Toolkit, and has contributed to applied wisdom education through Wikiversity. His recent work emphasizes coalition-building strategies and simultaneous policy implementation to address global collective-action problems. He lives in Middletown, New Jersey.

Cover image: PPstock / shutterstock.com

The views expressed in this report are those of the author(s) alone. They do not necessarily reflect the views of the Toda Peace Institute. An online edition of this and related reports and policy briefs can be downloaded on our website: toda.org/policy-briefs-and-resources.html



Toda Peace Institute
Samon Eleven Bldg. 5F,
3-1 Samon-cho, Shinjuku-ku
Tokyo 160-0017, Japan

Tel. +81-3-3356-5481
Fax. +81-3-3356-5482
Email: contact@toda.org

© 2026 Toda Peace Institute

Abstract

Global climate negotiations have produced ambitious targets but limited coordinated policy implementation. A central barrier is the multipolar collective-action problem, in which states hesitate to act for fear of economic disadvantage. This report examines how deliberative technologies can restructure climate negotiation architecture to enable multi-issue, mutually beneficial agreements that can be simultaneously implemented without undermining relative competitiveness. Using a Smartsettle Infinity simulation of an alternative global climate negotiation architecture, the report demonstrates how private preference modelling, structured trade-offs, and optimisation algorithms can generate Pareto-superior policy packages at the global scale. The simulation illustrates how coordinated policy design may overcome strategic deadlock while preserving national sovereignty. Rather than proposing specific climate policies, the contribution lies in illustrating a **deliberative decision-support architecture** capable of identifying coordinated, politically viable outcome packages under realistic strategic constraints. The findings point toward new pathways for coupling deliberative negotiation technology with citizen-driven political mobilisation to strengthen global climate governance. The analysis is intended for public agencies, philanthropic funders, multilateral institutions, and research organisations with the capacity to support experimentation alongside formal climate negotiations.

1. Introduction

Over the past three decades, the risks posed by climate change have become increasingly clear, even as international diplomatic efforts have expanded in scope and ambition. Even the Summaries for Policymakers issued by the IPCC are negotiated line by line with governments, underscoring how structural negotiation dynamics shape not only policy commitments but also the communication of scientific findings. While Conference of the Parties (COP) language has evolved over time, including contested terms such as ‘phase out fossil fuels,’ these rhetorical shifts have not translated into coordinated implementation, reinforcing the limits of target-based and ratchet-style mechanisms. Meanwhile aggregate global greenhouse gas emissions continue to rise, and the gap between nationally determined pledges and trajectories consistent with climate stabilisation remains substantial.

Although climate change is inherently multi-dimensional, negotiations have typically treated emissions reduction as a single-issue problem, detached from linked fiscal, technological, and enforcement instruments. Framed around voluntary targets and implemented sequentially, fears of competitive disadvantage and incentives to defect persist. While national commitments are increasingly documented (e.g., by the Climate Action Tracker [1]), aggregate outcomes continue to diverge sharply from trajectories consistent with climate stabilisation, reflecting a persistent implementation and coordination gap. Figure 1 summarizes scientific consensus as reflected in IPCC assessment reports, illustrating the divergence between emissions trajectories and stabilization pathways. This persistent disconnect between knowledge, commitment, and coordinated implementation points to a deeper structural governance challenge, not merely a deficit of scientific understanding or political goodwill.

[1] https://www.decisionchange.net/_files/ugd/9464b0_a7e8cb725b374c8e8f8a205154e091e0.pdf

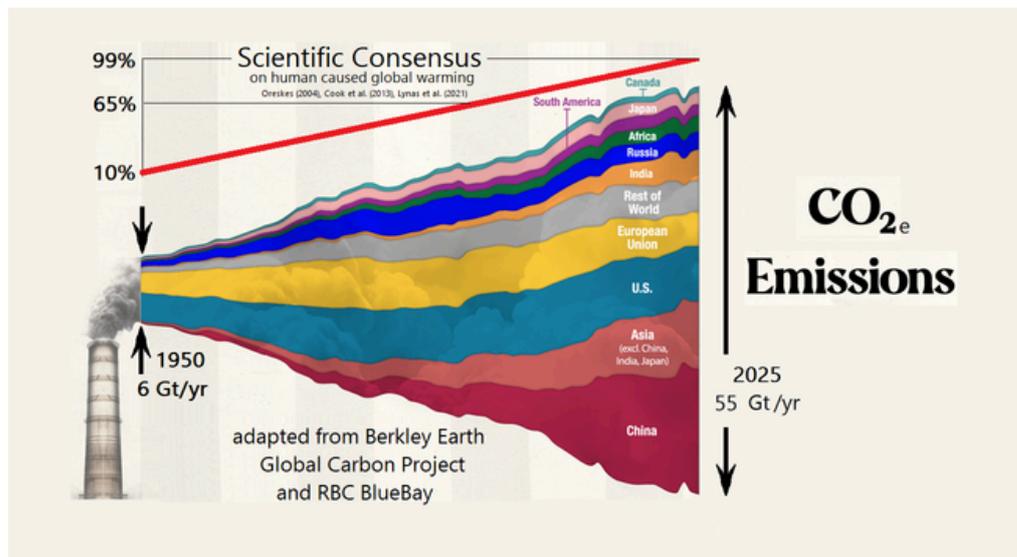


Figure 1. While scientific consensus on anthropogenic global warming has increased from roughly 10% in the 1950s to near-universal agreement today, global CO₂ emissions have continued to rise.

The COP process under the United Nations Framework Convention on Climate Change (UNFCCC), which began in the early 1990s, represents one of the most ambitious multilateral negotiation efforts in history. Nearly all nations participate. Yet COP outcomes remain dominated by aspirational targets, voluntary pledges, and uneven implementation [UNFCCC 2015]. Enforcement mechanisms are weak, coordination remains fragmented, and fears of economic competitiveness continue to constrain ambition (Victor 2011; Keohane and Victor 2016). The result is what may be described as a *multipolar trap*: a collective-action dilemma in which every major emitter hesitates to act decisively unless others do so first. This coordination failure is a well-documented feature of international environmental treaty-making under conditions of strategic interdependence (Schelling 1960; Barrett 2003).

Several interlocking structural dynamics help explain why COP processes have struggled to produce decisive results. First, although nearly all nations participate, each ultimately acts under domestic political and economic constraints. Because no government can be certain that others will fully implement ambitious measures on comparable timelines, coordinated action repeatedly breaks down.

Second, climate diplomacy has historically emphasised targets rather than detailed, synchronised policy instruments. Emissions goals are announced, but the fiscal, regulatory, and industrial measures required for aligned implementation are rarely negotiated in an integrated way. This uncertainty sustains fears of economic disadvantage and carbon leakage. While these dynamics resemble familiar Nash-equilibrium and free-rider problems (Keohane and Victor 2016, see also Section 2.2), the multipolar trap is not merely a static equilibrium outcome. It arises from the interaction of many sovereign actors, asymmetric costs, and the absence of credible mechanisms for simultaneous commitment and implementation. Under these conditions, even Pareto-superior agreements remain unreachable—not because cooperation is irrational in principle, but because no actor can safely move first.

Third, negotiations suffer from a persistent lack of trade-off mechanisms. High-emitting or economically vulnerable states are often asked to assume large adjustment costs without parallel compensatory instruments. The absence of linked financial, technological, or fiscal tools capable of redistributing burdens—such as coordinated climate finance or global transaction mechanisms—severely constrains the agreement space. Fourth, verification and enforcement remain weak: monitoring is fragmented, penalties are politically contested, and compliance incentives are insufficiently aligned with national self-interest. Finally, while civil society organisations have proven highly effective at agenda-setting and public mobilisation, they possess limited institutional leverage to ensure that governments both conclude robust agreements and subsequently implement them (Abbott and Snidal 2000; Victor 2011). Without mechanisms that translate popular concern into sustained, synchronised electoral accountability across borders, climate governance remains vulnerable to short-term political cycles and concentrated economic interests.

Together, these factors suggest that the central challenge of climate governance is structural rather than ideological. Even highly motivated actors may be unable to reach stable cooperative outcomes if the negotiation framework itself does not permit credible trade-offs, synchronised commitments, and enforceable coordination.

This report examines whether *deliberative negotiation technologies*—defined here as systems that preserve stakeholder agency while enabling structured, computable trade-offs—can help address these structural weaknesses when combined with simultaneous implementation strategies and citizen-driven electoral coordination. The analysis draws on a structured simulation of coordinated climate negotiations, documented in a companion video (Thiessen, Bunzl and Beaumont 2026), using **Smartsettle Infinity**—a deliberative optimisation platform designed to support complex, multi-party decision-making. The platform’s core concepts trace back to Ernest Thiessen’s doctoral research at Cornell University on visualising complex trade-offs and facilitating fair and efficient collaborative solutions to multi-stakeholder decision problems (Thiessen, Loucks, and Stedinger 1998; Keeney and Raiffa 1993; Belton and Stewart 2002).

Using a Smartsettle Infinity simulation of coordinated climate negotiations, this report explores how multi-issue optimisation, private preference modelling, and fairness-based algorithms can support structured negotiation among sovereign actors with divergent interests. Rather than facilitating dialogue or persuasion, the system functions as **negotiation decision-support infrastructure**: it enables parties to articulate preferences privately, evaluate trade-offs across linked issues, and identify outcome packages that are jointly superior to available alternatives.

The simulation does not model actual diplomatic behaviour, bargaining tactics, or political decision-making by states. Instead, it tests whether a structured, multi-issue, simultaneous negotiation framework can *identify* coalition-proof outcomes under plausible strategic constraints. Accordingly, the purpose is not to advocate any particular climate policy outcome *ex ante*, but to demonstrate a negotiation-enabled decision-support architecture capable of producing concrete, jointly agreed policy packages under realistic strategic constraints. Questions of enforcement and compliance are addressed later through mechanisms of simultaneous implementation, domestic legal authorisation, and coordinated electoral accountability.

Numerous institutional reforms have been attempted within the COP process, including market mechanisms, transparency frameworks, loss-and-damage facilities, and voluntary climate finance arrangements. While each has addressed important sub-problems, none has altered the underlying incentive structure governing national decision-making. The novelty of this work lies not in climate modelling or policy design, but in demonstrating that coordinated global agreements are structurally feasible when negotiations are re-architected. The simulation thus functions as a proof of coordinated feasibility under multipolar conditions, rather than an idealised equilibrium or a normative prescription.

2. Theoretical foundations of multi-party climate negotiation

2.1 COLLECTIVE-ACTION THEORY AND THE MULTIPOLAR TRAP

International climate negotiations exhibit key features of a classic collective-action problem. Each country may hesitate to reduce emissions if it believes others will not do their fair share, since acting alone can seem costly. When many countries think this way, the result is worse for everyone: global emissions remain high and climate risks increase. What makes sense for individual countries in the short-term leads to worse outcomes for the world as a whole.

At an institutional level, these dynamics stand in contrast to the design principles identified in the commons governance literature, particularly those advanced by Elinor Ostrom (Ostrom 1990; Ostrom et al. 2002), who demonstrated that durable cooperation depends on clearly defined rules, credible commitment mechanisms, and locally legitimate enforcement structures. Climate governance differs in scale but not in principle: without institutional mechanisms that translate shared interests into enforceable commitments, cooperation remains fragile. While game-theoretic extensions such as correlated equilibria relax some coordination constraints, they do not resolve the institutional problem of credible simultaneous commitment among sovereign actors.

2.2 NASH EQUILIBRIUM AND ITS LIMITS IN CLIMATE GOVERNANCE

A **Nash equilibrium** is a situation in which **each participant is choosing a strategy that is optimal given the strategies of others**, and therefore **no one has an incentive to change their behaviour on their own** (Nash 1950).

Imagine two people deciding whether to cooperate or act selfishly. If both cooperate, both do well. If one cooperates and the other acts selfishly, the selfish one benefits more. If both act selfishly, both do poorly.

If each person expects the other to act selfishly, then acting selfishly becomes the safest choice for each, even though everyone ends up worse off (a structure analogous to the Prisoner's Dilemma). That outcome is a Nash equilibrium: no one can improve their situation by changing their choice alone. This doesn't mean that Nash equilibria are inherently undesirable; in many settings they do yield efficient outcomes. The difficulty arises when equilibrium coincides with collectively inferior results.

In climate negotiations, this often manifests as countries settling for weak commitments that no one strongly opposes, even though all would benefit from deeper cooperation. Escaping such equilibria requires the creation of alternative **cooperative solution concepts** that restructure the payoff landscape. This is not simply a matter of moral persuasion; it is a matter of altering the incentive geometry itself.

2.3 PARETO OPTIMALITY AND COOPERATIVE GAIN

These limitations motivate the shift from equilibrium concepts to **Pareto optimality** and cooperative gain. Pareto optimality characterizes outcomes in which no stakeholder can be made better off without making at least one other stakeholder worse off (Keeney and Raiffa 1993). In simple terms, a Pareto-optimal outcome represents a point where all mutually beneficial improvements have been exhausted: any further change would require someone to accept a loss. Importantly, Pareto optimality says nothing about fairness or equity—it merely identifies the boundary of feasible mutual improvement.

The novelty lies not in the existence of Pareto-optimal outcomes in principle, but in their structural unreachability under target-based, sequential negotiations—and in their attainability once negotiations are re-architected to allow multi-issue trade-offs and simultaneous commitment. Accordingly, for applied governance the more relevant and intuitive concept is mutual improvement: outcomes in which all parties are better off relative to a realistic baseline or status quo. For example, in a climate negotiation, an agreement that reduces

emissions while also improving economic outcomes for all participating regions—through technology sharing, financial transfers, or risk reduction—would be Pareto-superior to inaction. By contrast, agreements that visibly impose net losses on specific actors, even if they advance global goals, are inherently fragile. This distinction is crucial for political durability: agreements that clearly harm some parties tend to unravel over time as those parties withdraw support or seek to renegotiate.

In multi-dimensional negotiations, stakeholders evaluate outcomes across several incommensurable dimensions rather than a single linear utility scale. Such problems, involving many issues and many parties, create opportunities to realize substantial cooperative gains precisely because losses on one dimension can be offset by gains on another. When negotiations are structured to allow such trade-offs, participants can move from inefficient outcomes—where cooperation is underutilized—to the **Pareto efficiency frontier**, which represents the set of outcomes where total joint gains are maximized. Whether a trade-off is sufficient is determined by each stakeholder's privately specified preference function, which defines acceptable compensation across all dimensions. Figure 2 illustrates this concept for a simplified three-party problem.[2] Moving from the origin toward the efficiency frontier corresponds to discovering combinations of commitments and compensations that leave all parties better off than their starting positions.

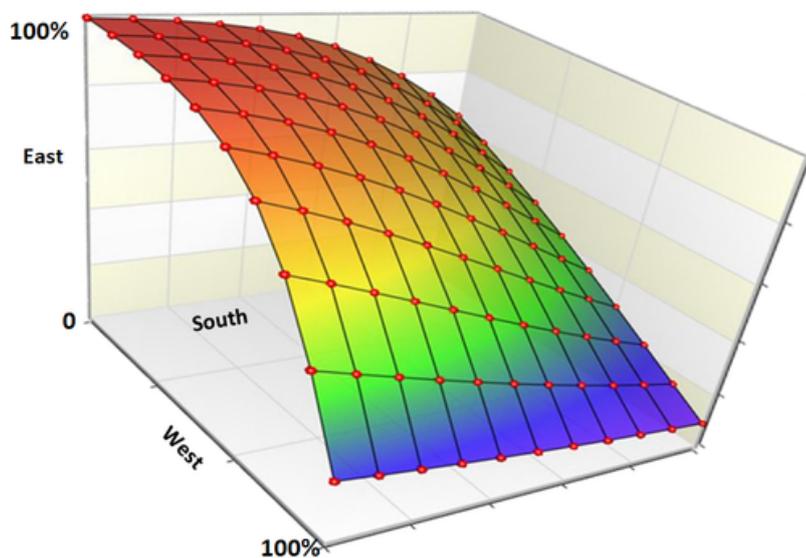


Figure 2. *Efficiency Frontier in a three-party problem. Gains for each individual party are represented by their own axis, while total joint gain is implicit in the surface rather than represented as a separate axis. All parties would be better off by moving from the origin (labelled 0) to the Efficiency Frontier (coloured surface).*

2.4 MULTI-ATTRIBUTE UTILITY AND PREFERENCE MODELLING

Climate policy choices involve multiple dimensions: economic cost, energy security, employment, technological innovation, geopolitical influence, environmental outcomes, and social equity. Multi-attribute utility theory (MAUT), pioneered by scholars such as Raiffa and Keeney, provides formal methods for representing such multidimensional preferences in a mathematically consistent way (Keeney and Raiffa 1993; Raiffa 1982). Rather than collapsing complex trade-offs into a single metric, MAUT preserves the structure of stakeholder priorities across issues. This allows trade-offs—such as between emissions reductions and financial transfers—to be assessed transparently rather than implicitly. Decision-making then proceeds by first identifying Pareto-superior outcomes relative to baseline, and then selecting among them using fairness-based optimisation criteria that ensure no participant is disproportionately disadvantaged.

[2] For a more in-depth illustration of this concept, the reader is referred to the [companion video](#) (Thiessen, Ernest M., John Bunzl and Leland Beaumont 2026.)

2.5 CONFLICT STRUCTURE AND COLLABORATIVE OUTCOMES

Collaborative outcomes in high-stakes negotiations require simultaneous attention to both self-interest and mutual gain. In climate negotiations, adversarial strategies tend to produce deadlock, while purely accommodative approaches often yield symbolic agreements that unravel during implementation. Durable collaboration instead depends on structured mechanisms that make it possible to identify, evaluate, and balance divergent priorities across multiple issues.

These perspectives converge on a central insight: **negotiation failure is often structural rather than ideological**. Even when governments broadly agree on the science of climate change and express sincere commitments to mitigation, they may still be unable to cooperate effectively if the negotiation framework itself creates incentives to delay, defect, or free-ride. For example, when emissions reductions are negotiated primarily as unilateral or sequential commitments, countries face persistent fears of competitive disadvantage if others fail to act in parallel. In such settings, reluctance to commit may reflect rational strategic caution rather than lack of political will.

From this perspective, persistent failure does not stem from unwillingness to cooperate, but from the absence of decision structures that allow actors to make credible, synchronized commitments and to trade off costs and benefits across linked issues. Addressing climate change therefore requires not only ambition and consensus on goals, but negotiation architectures capable of transforming shared intent into stable, collectively implementable agreements. While optimised agreements improve decision quality and political feasibility, successful implementation still depends on domestic legal authority and coordinated political action.

3. Deliberative technologies as global decision infrastructure

3.1 FROM BARGAINING TO STRUCTURED OPTIMISATION

Deliberative negotiation technologies such as **Smartsettle Infinity** differ fundamentally from traditional bargaining platforms. Rather than facilitating dialogue, persuasion, or positional exchange, these systems support collaboration by enabling parties to express preferences in structured, computable form and to explore trade-offs across multiple issues simultaneously, with private preference inputs preventing sequential signalling or strategic posturing. In this context, *deliberation* does not refer to open discussion or facilitated conversation. Instead, deliberation is embedded in the formal articulation of stakeholder priorities, preserving agency, sovereignty, and pluralism while allowing systematic optimisation across those preferences without requiring persuasion, consensus-seeking dialogue, or rhetorical agreement.

Smartsettle Infinity is the culmination of several decades of research and development in decision analysis and negotiation support. It operates through a secure neutral platform in which multiple stakeholders privately enter structured preference functions over a shared set of issues and options. The system then evaluates potential agreements against these inputs to identify outcome sets that are *Pareto-superior* relative to baseline positions—solutions in which all participants are better off than under their realistic alternatives.

In conventional diplomatic bargaining, outcomes are typically driven by strategic posturing, asymmetric information, and sequential concessions. These dynamics often privilege power over efficiency and make mutually beneficial trade-offs difficult to identify or sustain. By contrast, deliberative optimisation platforms formalise the negotiation problem itself. Rather than competing over positions, participants jointly search an explicitly defined feasible solution space. This shifts negotiation from rhetorical compromise toward mathematically coherent optimisation, enabling agreements that are both substantively balanced and strategically stable.

3.2 SINGLE NEGOTIATING FRAMEWORKS AND MULTI-ISSUE LINKAGE

A central design element of this approach is the **Single Negotiating Framework (SNF)**—a shared, formal representation of all issues, options, constraints, and interdependencies to be negotiated. The SNF establishes the problem architecture within which optimisation can occur. For climate policy, this includes emissions pathways, fiscal instruments, technology transfer, adaptation funding, verification protocols, and enforcement mechanisms.

Multi-issue linkage is essential. Classic bargaining theory shows that linking issues allows negotiators to transform conflicts from zero-sum trade-offs into mutually beneficial exchanges, particularly when parties value issues differently (Schelling 1960). In the climate context, this means that commitments to emissions reductions can be paired with financial transfers, technology access, or implementation flexibility in ways that make cooperation rational for all participants rather than punitive for some.

When emissions reductions are negotiated in isolation, without credible commitments or compensation mechanisms, high-cost actors face insufficient compensation and rationally resist ambitious commitments. When emissions are linked to financial transfers, market instruments, technology access, and enforcement guarantees, meaningful trade-offs become possible. This follows well-established results in collective-action theory and international environmental cooperation. When countries face unequal costs and benefits from cooperation, agreements are rarely stable unless they include *linked instruments*—such as financial transfers, investment funds, or technology-sharing arrangements—that help compensate those bearing higher costs. Research on collective action and international environmental cooperation shows that agreements are most stable when they combine credible commitments, issue linkage, and mechanisms for addressing asymmetric costs and benefits (Ostrom 1990; Barrett 2003). These findings motivate the inclusion of financial and institutional instruments alongside emissions targets within the SNF. Absent these mechanisms, incentives to delay or defect remain strong.

However, even with multi-issue linkage, a deeper structural constraint remains: under unilateral or staggered implementation, governments face intense competitive pressure that sharply limits what policies are politically and economically feasible.

3.3 COMPETITIVE PRESSURE, POLICY FEASIBILITY, AND THE ROLE OF SIMULTANEITY

Under conditions of unilateral or non-simultaneous implementation, rising international competitive pressure sharply constrains the range of politically feasible climate policies. Governments rationally fear capital flight, carbon leakage, loss of export competitiveness, and domestic political backlash if they act ahead of others. As competitive pressure increases, the set of policies that can be implemented safely by any single country rapidly collapses. Figure 3 illustrates this structural constraint, which helps explain why global climate negotiations tend to settle into a low-feasibility equilibrium—one in which emissions continue to rise despite widespread rhetorical commitment to mitigation.

Simultaneous implementation alters this strategic geometry without changing the underlying structural relationship. By synchronizing adoption and enforcement across participating states, simultaneous implementation **neutralizes competitive pressure as the binding constraint**. Each government moves together, and no actor faces a unilateral disadvantage. In terms of Figure 3, simultaneity shifts the operating point of the international system from the high-pressure, low-feasibility region of the curve to a low-pressure, high-feasibility region in which ambitious policy becomes politically and economically viable.

In this sense, simultaneity does not expand the structural policy space itself; rather, it repositions the system within that space by removing the competitive-risk trap that otherwise governs unilateral action.

3.4 SIMULTANEOUS IMPLEMENTATION AND STRATEGIC STABILITY

Simultaneous implementation is critical for escaping the multipolar trap because it changes the strategic conditions under which governments make commitments. When policies are adopted unilaterally or implemented sequentially, early movers face credible risks of economic disadvantage if others delay or defect. By contrast, when all major actors commit to implementing agreed measures on the same effective date under coordinated domestic legislation, fears of first-mover disadvantage are greatly reduced.

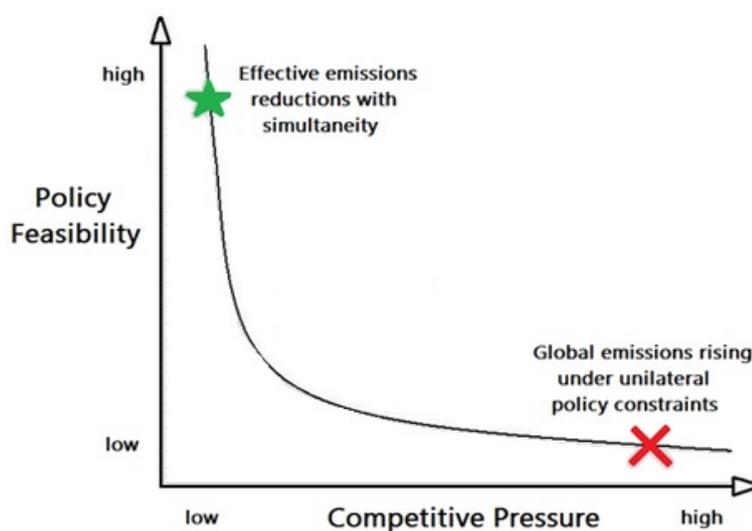


Figure 3. Policy feasibility as a function of international competitive pressure, illustrating how simultaneous implementation neutralizes the competitive constraint that otherwise limits feasible policy outcomes under unilateral implementation.

In this setting, **strategic stability** arises because no participant can improve its position by delaying, defecting, or free-riding once implementation is synchronized. Governments comply not because they hope others will follow through, but because implementation is synchronised and institutionally locked in across jurisdictions. For example, if carbon pricing, border adjustments, or fossil fuel phase-out schedules take effect simultaneously across multiple economies, firms and governments face a level competitive playing field. The incentive to defect or postpone action is diminished because deviation would immediately carry domestic and international costs.

A similar logic appears in the Simultaneous Policy (Simpol) approach to international cooperation. Simpola treats simultaneous implementation as a central mechanism for addressing familiar failures in global governance, including free-riding, race-to-the-bottom dynamics, and weak enforcement. Entirely novel in the Simpola approach is that rather than relying on trust, goodwill, or post hoc sanctions, Simpola mobilises citizen-driven electoral coordination to pressure governments to commit in advance to implementing agreed policies together, and only once sufficient participation is assured (Bunzl and Duffell 2017). This electoral pressure also has the potential to overcome resistance on the part of elites to policies such as a Tobin tax being included in negotiations.

Within a deliberative optimisation framework such as Smartsettle Infinity, simultaneity can therefore be analysed not as an abstract ideal, but as a practical coordination mechanism. The COP case study that follows examines whether embedding simultaneous implementation within a structured, multi-issue negotiation process can transform an otherwise fragile bargaining environment into a more cooperative and potentially Pareto-improving coordination problem.

4. The Smartsettle Infinity climate negotiation simulation: A COP case study

4.1 PURPOSE AND SCOPE

An alternative architecture for COP-style climate change mitigation negotiations was simulated using Smartsettle Infinity (Thiessen, Boynton, and Ross 2025) as a **proof of deliberative decision-support architecture**. To render the problem tractable while preserving its essential geopolitical structure, the sixteen country groupings typically represented at COP were consolidated into **six composite stakeholder groups** representing the dominant strategic positions in global climate negotiations:

1. Global North (G7 and close allies)
2. China–Russia Bloc (CRB)
3. Emerging High Emitters (EHE)
4. Global South (South)
5. Small Island Developing States (SIDS)
6. Least Developed Countries (LDC)

These groupings were selected to preserve the principal **geopolitical and developmental asymmetries** that shape climate negotiations. Together, they capture high historical emitters, current large emitters, rapidly industrialising economies, highly climate-vulnerable states, and countries with minimal historical responsibility but high exposure to climate risk. While this simplification does not claim representational completeness, nor does it imply that these groups act as unified political actors in practice, it preserves the **core strategic tensions** that define real COP dynamics.

Stakeholder roles in the simulation were instantiated using AI agents operating under author-supervised preference structures derived from publicly stated national positions, policy constraints, and economic profiles available in late 2025. These agents did not attempt to reproduce real diplomatic behaviour or bargaining tactics. Their role was to provide internally consistent representations of stakeholder preferences suitable for testing the negotiation architecture itself.

The objective of the simulation was therefore not behavioural realism, but to examine whether a formally structured, multi-issue optimisation process—combined with simultaneous implementation—could identify a mutually improving policy package under plausible strategic constraints. The deliberative decision-support environment in which the simulation was conducted is illustrated in Figure 4.

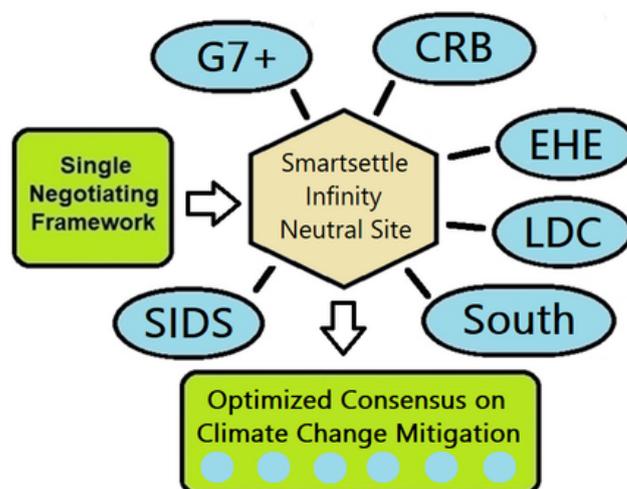


Figure 4. Smartsettle Infinity multi-stakeholder negotiation workspace used for the climate simulation. A comprehensive Single Negotiating Framework enables simultaneous issue structuring, private preference entry, optimisation, and outcome comparison across all stakeholder groups, supporting convergence toward an optimised consensus.

4.2 ISSUE ARCHITECTURE

The simulation employed a pre-defined **Single Negotiating Framework (SNF)** consisting of multiple climate-related issue categories, as summarised in Table 1. The SNF specified, for each stakeholder group, commitments to **emissions-reduction targets** and an associated **annual penalty for non-compliance**. In addition to emissions mitigation, the framework incorporated issues relating to **technology deployment, adaptation and resilience funding, verification protocols, and enforcement mechanisms**.

A central financial instrument within the SNF is the Global Climate Fund, designed to facilitate equitable participation across stakeholder groups by addressing asymmetries in capacity and historical responsibility (Barrett 2003; Nordhaus 2015). Multiple funding sources were explored by the authors, including direct national contributions from higher-income countries and a **Tobin Tax** on international currency transactions (Tobin 1974). A Tobin Tax is a small levy on foreign-exchange trades, originally proposed to curb destabilising speculation but widely recognised for its potential to generate substantial, stable public revenue with minimal impact on real economic activity.

Table 1. Single Negotiating Framework for Global Climate Change Mitigation Consensus

Emissions Reduction

- Emissions Reductions for each Stakeholder
- Total Emissions Reduction
- Penalty for Missing Target

Other Issues

- How to Account for Carbon Removal
- Technology Access & Innovation Sharing
- Verification Mechanisms

Funding Sources

- National Resources from Wealthy Nations
- Global Climate Fund
- Tobin Tax

In the simulation, this mechanism functioned as a stabilizing instrument by drawing revenue from global financial activity rather than national budgets alone, thereby reducing fiscal salience, broadening cost distribution, and helping sustain participation by lower-income and high-vulnerability stakeholders within the optimised outcome. Importantly, instruments such as a Tobin Tax are presented here for **illustrative purposes only**. Their role is to demonstrate how the negotiation architecture can surface and integrate stabilizing financing mechanisms—not to advocate any particular policy choice.

In real-world settings, the adoption of a Tobin Tax would itself face political resistance and coordination challenges. Other mechanisms—such as carbon border adjustment revenues, sectoral levies, or pooled climate finance commitments—could serve comparable stabilizing functions within the same decision-support framework and merit comparative exploration in future work.

4.3 PRIVATE PREFERENCE MODELLING

The preferences of each stakeholder group were modelled using **publicly available policy statements, emissions data, and economic indicators**. This required the construction of **satisfaction functions for every issue** in the negotiation. Continuous functions were estimated for numerical variables, while discrete ratings were assigned for option-based variables. In addition, each issue category was assigned a **relative importance weight**, representing its contribution to overall preference for each stakeholder group (Keeney and Raiffa 1993; Raiffa 1982).

Although climate negotiations are conducted in the public domain and parties generally possess partial knowledge of other actors' priorities, the simulation required that all **preference data be entered privately and independently**. No stakeholder group had direct access to any other group's inputs. This design preserves **sovereignty, confidentiality, and strategic autonomy**, while preventing strategic signalling during preference formation. Because Smartsettle Infinity rewards accurate self-representation rather than strategic posturing, each group was oriented toward expressing its **true internal preferences** as the most effective means of achieving a globally optimal result.

While the resulting **preference functions** were necessarily stylised—that is, simplified representations that abstract from full political, economic, and behavioural complexity—they were grounded in publicly articulated national positions, published emissions trajectories, and established economic and development profiles. The resulting preference representations include order-of-magnitude estimates of avoided damages and efficiency gains.

In this context, preference functions refer to the complete formal specification of how a stakeholder group evaluates outcomes across each issue: including satisfaction curves for quantitative variables, option ratings for discrete choices, relative importance weights across issues, and any necessary constraints or thresholds. Together, these elements define what constitutes improvement or acceptability for each participant. *Put simply, a preference function acts like a map of what each stakeholder values, how strongly they value it, and where their limits lie.*

These inputs are entered privately by each stakeholder into Smartsettle's secure neutral platform, where the underlying algorithms are implemented and remain inaccessible to other parties. The system then facilitates structured proposal exchange and uses optimisation methods to identify promising candidate agreements. This process generates an initial baseline outcome, which can subsequently be improved through iterative exploration of mutually beneficial trade-offs. The objective of the modelling was not to predict actual negotiation behaviour, but to ensure that stakeholder preferences remained within empirically plausible bounds suitable for testing the structural performance of the deliberative negotiation process itself.

4.4 ESTABLISHING A FAIR BASELINE

While Pareto-optimal outcomes lie on the **Efficiency Frontier**, not all such outcomes are perceived as fair or politically stable. Establishing a credible and broadly acceptable baseline is therefore essential for cooperation. To achieve this, Smartsettle engages participants in **multivariate Visual Blind Bidding**, a process in which each party submits confidential preference information across multiple issues without seeing the inputs of others. The Visual Blind Bidding process builds on early work in decision analysis aimed at reducing strategic manipulation and anchoring effects by separating preference expression from positional bargaining (Raiffa 1982). Using this information, the system helps parties reach a provisional agreement—referred to as the *Baseline*—that reflects a balanced accommodation of stakeholder priorities. Because bids are blind and multivariate, parties cannot strategically anchor or manipulate outcomes through positional bargaining, helping to establish a shared sense of procedural fairness.

Negotiation research also shows that parties evaluate gains and losses relative to reference points rather than in absolute terms. Establishing a shared Baseline therefore matters not only for analytical reasons, but because it shapes how stakeholders perceive subsequent proposals—as improvements or sacrifices—thereby influencing willingness to continue cooperating.

The Baseline thus performs a dual role. First, it anchors perceptions of fairness by requiring all stakeholders to explicitly accept the Baseline as a fair reference point. Second, it establishes the benchmark from which additional cooperative gains can be identified and measured. How close the Baseline lies to the Efficiency Frontier depends in part on the number of bidding sessions conducted. Additional sessions allow preferences to be more fully expressed and reconciled, particularly in multi-party settings, yielding a Baseline that is closer to the frontier. In this simulation, the process was deliberately accelerated and reached a Baseline after only three sessions, leaving substantial cooperative value unrealized and allowing scope for further optimisation in subsequent stages.

4.5 OPTIMISATION ALGORITHM

The core optimisation algorithm of the Smartsettle process is **Maximise the Minimum Gain**. This algorithm harvests any remaining hidden value and distributes it fairly to all parties. Rather than equalising benefits, this algorithm identifies the package of options that maximises the smallest relative improvement among all participants compared to their respective baselines. This operationalises fairness by ensuring that no actor is left disproportionately worse off. The magnitude of the benefits of this process is illustrated in Table 2.

4.6 RESULTS

The optimised solution produced a **Pareto-superior improvement for all six stakeholder groups**, meaning that no hidden value remained on the table. Each stakeholder achieved a measurable gain relative to baseline. Some gains were larger than others, reflecting asymmetries in the underlying model and associated preferences and trade-offs, but **no stakeholder was made worse off**. Only the final optimised package and each stakeholder's own improvement were visible at the conclusion of the process; all private preference functions remained confidential throughout.

Table 2. Optimisation results following application of the Maximise the Minimum Gain algorithm. The aggregate annual benefits across stakeholder groups approach the order of one trillion USD under the simulated assumptions, representing approximately a 30% improvement over the baseline outcome. While these potential benefit estimates are conservative according to many experts, they were still sufficient to unlock the collaborative optimised consensus illustrated in our simulation.

	Final Optimized Consensus	
	Total Value (\$B/yr)	Improvement over Baseline (\$B/yr)
Global North (G7 + allies)	166	58 (35%)
China–Russia Bloc	104	37 (36%)
Emerging High Emitters	144	61 (42%)
Global South	257	60 (23%)
Small Island Developing States	151	26 (17%)
Least Developed Countries	126	43 (34%)
Total Benefit	948	285 (30%)

Figure 5 summarises key elements of the final optimised consensus. The model assumed that a **global, all-party consensus** would be the optimal outcome and did not attempt to predict the political dynamics of partial agreements. However, sensitivity testing revealed that the **Climate Fund was essential** for achieving a consensus that included the Global South (Nordhaus 2015).

When the Global Climate Fund was removed from the negotiation framework, participation by the Global South became difficult to sustain within the optimisation. This result reflects a structural constraint rather than a normative choice. In the absence of compensatory finance, mitigation commitments impose net welfare losses on lower-income and climate-vulnerable countries that lack the fiscal capacity to absorb transition costs or invest in adaptation. Under these conditions, no combination of emissions reductions, policy instruments, or implementation schedules could raise Global South outcomes above their baseline alternatives while remaining acceptable to other stakeholders.

In practical terms, this means that without a redistributive mechanism, optimisation yields outcomes that are efficient among higher-capacity actors but fail to satisfy participation constraints for lower-capacity ones. The Climate Fund therefore functions not as a moral add-on, but as a stabilising instrument that makes all-party participation feasible (Barrett 2003; Ostrom 1990). By reallocating a portion of global mitigation costs through mechanisms such as a Tobin Tax on currency transactions, the fund enables the optimisation process to generate agreements that are both Pareto-superior and coalition-proof across heterogeneous stakeholder groups.

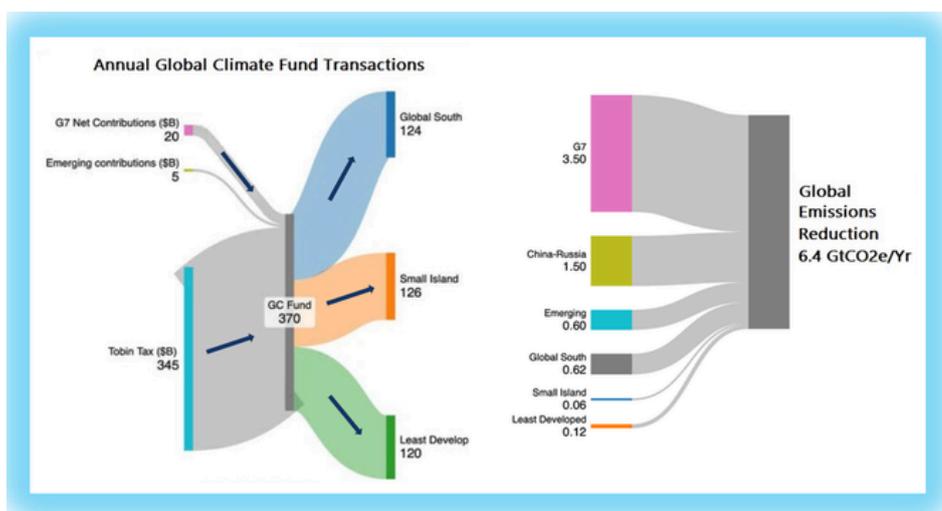


Figure 5. *Final Improved Outcome.* The left panel shows annual contributions to and withdrawals from the Global Climate Fund, illustrating the central role of the Tobin Tax in financing the fund and enabling broad stakeholder participation. The right panel shows the resulting annual emissions-reduction commitments of each stakeholder group under the optimised agreement, yielding a total global reduction of **6.4 GtCO₂e per year**, corresponding to an **estimated 12% annual reduction in global emissions**.

5. Interpretation

These numerical outcomes do not necessarily predict precise real-world behaviour. They demonstrate a **structural property**: even highly complex, multi-party climate problems with divergent interests can admit internally consistent, mutually beneficial solutions when negotiations are conducted within a fully structured, multi-issue optimisation framework.

The authors contend that the results of the simulation represent a realistic proof of concept that warrants serious consideration by climate governance institutions. While a real-world negotiation will almost certainly differ from what has been demonstrated here, outcomes within this general range would represent a substantial improvement over current trajectories.

What This Simulation Demonstrates

The simulation demonstrates that the primary obstacle facing global climate negotiations is not technological or economic, but structural. When negotiations are conducted as single-issue, sequential, and non-simultaneous processes, coordination, trust, and incentive alignment fail. By contrast, a multi-issue, optimisation-based framework with simultaneous implementation allows these barriers to be addressed together. The robustness of the result lies not in the precision of the numerical outputs, but in the structural feasibility of reaching a mutually improving consensus under realistic strategic constraints.

By contrast, if the same six stakeholder groups had been constrained to negotiate emissions reductions as a stand-alone issue using sequential pledge-and-review logic (Victor 2011; Keohane and Victor 2016), the feasible agreement space would have been dramatically smaller, making it much more difficult to find a solution. The simulation therefore suggests that the expanded gains observed here arise not from optimistic assumptions about political will, but from the structural effects of issue linkage, private preference modelling, and optimisation under simultaneity.

6. From optimised agreements to democratic binding force

Even optimal policy packages cannot become operational without domestic legal authority. Climate governance therefore requires a bridge between international optimisation and national democratic institutions.

One pathway is the implementation of **advance domestic enabling legislation**. Under such statutes, national legislatures authorise in principle the incorporation of any future agreement emerging from a defined international negotiation framework into domestic law, within specified limits. This preserves sovereignty while enabling rapid, synchronised implementation (Abbott and Snidal 2000).

Electoral leverage remains essential, not only for achieving such legislative pre-authorization and for the subsequent simultaneous implementation of any agreement, but also for overcoming potential elite resistance to the inclusion of policies such as a Tobin tax. Here, **citizen-driven simultaneous voting strategies**, such as those pursued by Simpol have received significant parliamentary support in several jurisdictions.[3] This electoral leverage thus becomes structurally significant in driving the governments of democratic nations towards the cooperation required and providing a powerful on-going electoral incentive to maintain it.

Participation by non-democratic or hybrid regimes poses additional challenges not resolved by electoral coordination alone. The simulation assumes engagement by such actors in order to test whether structural incentives for participation exist once major democratic economies coordinate. Whether and how such participation occurs in practice would depend on issue linkage, economic leverage, and strategic considerations beyond the scope of this proof-of-concept.

The relationship between deliberative optimisation, domestic legal authority, and coordinated electoral mobilisation is summarised schematically in Figure 6. In this architecture, deliberative technologies define **what is structurally possible**, while electoral mobilisation determines **what becomes politically adoptable**. These two dimensions are complementary, not alternative.

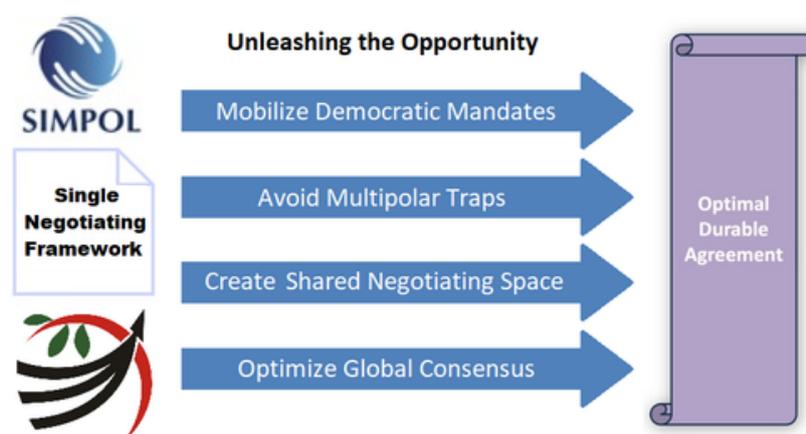


Figure 6. Key enabling conditions unlocked by the Simpol–Smartsettle approach. By mobilizing democratic mandates, eliminating competitive deadlocks, defining a shared Single Negotiating Framework, and optimizing multi-stakeholder trade-offs, the process creates the conditions for a stable and durable global agreement.

[3] See <https://simpol.org/who-we-are/pledged-politicians>

Importantly, this architecture does not imply the creation of a supranational government. Decision authority remains fully national. Deliberative optimisation operates only as a decision-support layer, while legal force arises exclusively through domestic legislation. The model therefore preserves constitutional sovereignty while enabling synchronized global implementation.

The mechanisms discussed here establish **necessary but not sufficient conditions** for adoption. Simultaneity and coordinated electoral pressure can reduce incentives for defection and delay, but they do not eliminate veto players, power asymmetries, or political contestation. Rather, the contribution of this architecture is to lower structural barriers to cooperation by aligning incentives, reducing uncertainty, and converting otherwise unstable bargaining problems into coordination problems—where political contestation operates within a feasible solution space rather than against it.

Accordingly, the analysis focuses on structural mechanisms that expand the space of politically feasible outcomes, rather than claiming to resolve the full political economy of adoption, which remains contingent on domestic politics, institutional capacity, and geopolitical context.

7. Policy implications and recommendations

The simulation results and accompanying governance analysis suggest several concrete policy implications. Taken together, they point toward a research and policy agenda focused not on substituting for existing climate institutions, but on augmenting them with deliberative decision infrastructure capable of supporting coordinated implementation. The following recommendations are directed primarily at public agencies, philanthropic funders, multilateral institutions, and research organisations with the capacity to support experimentation alongside formal COP processes.

POLICY RECOMMENDATIONS

1. Support experimental deliberative climate negotiation pilots alongside COP processes

Public agencies, philanthropic foundations, and research funders could support pilot projects that apply deliberative optimisation platforms to real or shadow climate negotiations operating in parallel to formal COP tracks. Suitable sponsors include international foundations focused on climate governance, development agencies, multilateral development banks, and research councils. Such pilots could be framed as exploratory ‘safe-to-fail’ environments for testing structured negotiation architectures without displacing existing diplomatic processes.

2. Integrate multi-issue Single Negotiating Frameworks into pre-COP preparation

UNFCCC bodies, COP presidencies, and affiliated policy institutes could explore the use of Single Negotiating Frameworks (SNFs) during pre-COP technical and ministerial consultations. Rather than replacing negotiation mandates, SNFs could serve as analytical scaffolding that helps parties clarify issue linkages, option spaces, and trade-offs in advance. This may improve coherence across mitigation, finance, adaptation, and implementation tracks that are currently negotiated in relative isolation.

3. Explore advance domestic enabling legislation for synchronised implementation

National legislatures, legal scholars, and constitutional experts could examine mechanisms for advance enabling legislation that conditionally authorises the domestic implementation of internationally negotiated climate packages, contingent on simultaneous action by other major actors. Such legislation would not pre-commit governments to unknown outcomes, but would pre-authorise participation in a bounded, procedurally defined decision space—lowering barriers to coordinated action once agreement is reached.

4. Strengthen civil-society strategies linking electoral pressure to global coordination

Civil-society organisations, climate NGOs, and transnational advocacy networks could play a critical role in translating coordinated policy architectures into political feasibility. Approaches such as the Simultaneous Policy (Simpol) strategy illustrate how citizen-driven electoral coordination can create incentives for governments to commit to synchronised implementation. Further experimentation is warranted to understand how such strategies scale across jurisdictions and interact with different political systems.

5. Invest in integration between deliberative governance platforms and Earth-system modelling

Research institutions and interdisciplinary consortia could explore tighter integration between deliberative negotiation platforms and climate and Earth-system models—that is, models that capture interactions among the atmosphere, oceans, land, ecosystems, and human systems. Such integration would allow negotiated policy packages to be evaluated not only for political feasibility and fairness, but also for their dynamic impacts on emissions trajectories, climate risks, and distributional outcomes over time. Besides those mentioned in this paper, there are other initiatives that explore integration between deliberative governance tools and Earth-system modelling. Radical X is one example among several emerging initiatives.

LIMITATIONS AND FUTURE RESEARCH

Several limitations of the present study must be acknowledged. The simulation does not model macroeconomic feedback, financial market dynamics, full Earth-system climate responses, or the full scale of participation found in COP processes. As currently organised, COP negotiations among nearly 200 countries operate through approximately sixteen negotiating groupings. To keep the simulation tractable while retaining realism, participants were aggregated into six representative stakeholder groups. This simplification is intended as a proof of concept rather than a claim of representational completeness.

From a process perspective, the Smartsettle Infinity framework does not rely on sequential or bilateral interaction. Parties submit their preferences simultaneously to a neutral platform and evaluate outcomes relative to their own priorities. As a result, the number of formal negotiating sessions does not necessarily increase in proportion to the number of participants. That said, increasing the number of parties would likely require more time to construct the Single Negotiating Framework (SNF), more preparatory simulations before formal negotiation, and a larger set of outcome packages to evaluate within each session. Scaling to larger numbers of actors—such as the sixteen COP groupings, or potentially all participating states—would therefore raise practical questions of representation, preference elicitation, and computational capacity. However, these are operational considerations rather than changes to the underlying decision architecture demonstrated here.

Preference inputs are stylised and cannot capture future political volatility, leadership change, or exogenous shocks. Moreover, simultaneous electoral mobilisation across jurisdictions—while conceptually plausible—remains an unresolved empirical challenge.

Accordingly, the findings should be interpreted as demonstrating structural feasibility rather than political inevitability. The contribution of this work lies in showing that, under realistic strategic constraints, coherent and mutually improving global climate agreements are not ruled out by logic or incentives alone. Whether such architectures can be translated into durable political practice remains an open and important question for future research and experimentation.

8. Conclusion

The Smartsettle Infinity/Simpol climate simulation demonstrates that the core barrier to robust global climate agreements is not a lack of feasible cooperative solutions, but the absence of institutional architectures capable of generating, legitimising, and implementing them. Multi-issue optimisation, simultaneous implementation, and citizen-driven political coordination via the SimpPol campaign together offer a plausible pathway for escaping the multipolar trap that has long constrained climate diplomacy.

The demonstration does not claim political inevitability, nor does it prescribe specific policies. It offers a proof of coordinated governance process in principle. If climate stabilisation is to be achieved at scale, governance innovation must advance with the same urgency and sophistication as climate science itself (IPCC 2023).

While climate change provides a compelling test case, the governance architecture explored here may be applicable to other global collective-action challenges, including biodiversity loss, international trade coordination, artificial intelligence governance, and security-risk reduction. A credible next empirical step would involve shadow negotiations, regional blocs, or sector-specific pilots to further test political and institutional viability.

Acknowledgements

The authors are grateful to Arnold Bomans for his detailed and constructive comments on earlier drafts of this paper and on the accompanying video material, which helped strengthen and clarify the manuscript. We also thank Lisa Schirch for her support of this project and for facilitating external review through the University of Notre Dame. Any remaining errors or interpretations are the responsibility of the authors.

REFERENCES

- Abbott, Kenneth W., and Duncan Snidal. 2000. "Hard and Soft Law in International Governance." *International Organization* 54, no. 3: 421–456.
- Barrett, Scott. 2003. *Environment and Statecraft: The Strategy of Environmental Treaty-Making*. Oxford: Oxford University Press.
- Bunzl, John, and Nick Duffell. 2017. *The SIMPOL Solution*. London: Peter Owen Publishers.
- Intergovernmental Panel on Climate Change (IPCC). 2023. *Sixth Assessment Report: Synthesis Report*. Cambridge: Cambridge University Press.
- Keohane, Robert O., and David G. Victor. 2016. "Cooperation and Discord in Global Climate Policy." *Nature Climate Change* 6: 570–575.
- Keeney, Ralph L., and Howard Raiffa. 1993. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge: Cambridge University Press.
- Nash, John F. 1950. "Equilibrium Points in n-Person Games." *Proceedings of the National Academy of Sciences* 36, no. 1: 48–49.
- Nordhaus, William D. 2015. "Climate Clubs: Overcoming Free-Riding in International Climate Policy." *American Economic Review* 105, no. 4: 1339–1370.
- Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.
- Ostrom, Elinor, Thomas Dietz, Nives Dolsak, Paul C. Stern, Susan Stonich, and Elke Weber. 2002. *The Drama of the Commons*. Washington, DC: National Academy Press.
- Raiffa, Howard. 1982. *The Art and Science of Negotiation*. Cambridge, MA: Harvard University Press.
- Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Thiessen, Ernest M., Loucks, D. P., & Stedinger, J. R. 1998. Computer-assisted negotiations of water resources conflicts. *Group Decision and Negotiation*, 7(2), 109–129.
- Thiessen, Ernest M., Carissa Boynton, and Graham Ross. 2025. "A Methodological Simulation Using Smartsettle Infinity for Negotiating Transboundary Water Disputes." In *Decision Support Models to Assist International and Transboundary Water Negotiations: Methodologies and Applications*, edited by Ariel Dinar and Mohammadreza Shahbazbegian, chap. 2. Oxford: Oxford University Press.
- Thiessen, Ernest M., John Bunzl and Leland Beaumont. 2026. *From Deadlock to Agreement: A simulated Global Climate Negotiation*. Smartsettle eNegotiation Channel. YouTube video. <https://www.youtube.com/watch?v=6WNXvAPQkUM>
- Tobin, James. 1974. *The New Economics One Decade Older*. Princeton: Princeton University Press.
- United Nations Framework Convention on Climate Change (UNFCCC). 2015. *Paris Agreement*. New York: United Nations.
- Victor, David G. 2011. *Global Warming Gridlock: Creating More Effective Strategies for Protecting the Planet*. Cambridge: Cambridge University Press.



THE TODA PEACE INSTITUTE

The Toda Peace Institute is an independent, nonpartisan institute committed to advancing a more just and peaceful world through policy-oriented peace research and practice. The Institute commissions evidence-based research, convenes multi-track and multi-disciplinary problem-solving workshops and seminars, and promotes dialogue across ethnic, cultural, religious and political divides. It catalyses practical, policy-oriented conversations between theoretical experts, practitioners, policymakers and civil society leaders in order to discern innovative and creative solutions to the major problems confronting the world in the twenty-first century (see www.toda.org for more information).

CONTACT US

Toda Peace Institute

Samon Eleven Bldg. 5 th Floor
3-1 Samon-cho, Shinjuku-ku, Tokyo 160-0017, Japan

Email

contact@toda.org

Sign up for the Toda Peace Institute mailing list

<https://toda.org/policy-briefs-and-resources/email-newsletter.html>

Connect with us on the following media.

YouTube: [@todapeaceinstitute3917](https://www.youtube.com/@todapeaceinstitute3917)

X (Twitter): <https://twitter.com/TodaInstitute>

Facebook: <https://www.facebook.com/TodaInstitute>