Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Leveraging Collective Intelligence: A Simulation of a Digital Co-Creation Platform with Agent-Based Modeling

Research Paper

Nicola Lo Russo

nlorusso@ethz.ch

Science, Technology & Policy
D-GESS, ETH Zurich

Supervisors:
Dino Carpentras

June 28, 2025

Abstract

Democratic systems are increasingly challenged by declining participation, institutional distrust, and the limitations of traditional deliberative processes. This research presents an agent-based modeling (ABM) simulation of a novel digital Co-Creation Platform (CCP) designed to enhance and simplify participatory democracy. The model simulates interactions between solver agents, classified as “Innovators”, “General Modifiers”, and “Best Modifiers”, and stakeholder evaluators, whose utility function is unknown to the solvers, within a decentralized problem-solving framework. By modeling utility and potential scores for solution ideas, I explore how diverse reward mechanisms affect collective decision-making. I implement several reward schemes to shape agent evolution and assess their impact on solution quality and diversity. Results indicate that no single solver type or reward scheme dominates; instead, optimal outcomes emerge from sustained strategic diversity and balanced exploration-exploitation dynamics. Particularly, combined reward systems that integrate innovation and improvement incentives maintain heterogeneous populations and outperform homogenous approaches. My findings underscore the central role of platform design in fostering constructive deliberation and highlight the importance of adaptive incentive structures in participatory systems. For policymakers, these insights offer guidance on designing institutional frameworks that are both inclusive and effective. By leveraging ABM, this work contributes to democratic theory, computational social science, and policy design, illustrating how digital participatory platforms can support scalable, resilient, and evidence-informed collective intelligence for governance in the digital age.

Contents

CHAPTER 1

Introduction

1.1 The Crisis of Democracy

In the last decades, democratic systems[1] all over the world have been facing an important crisis [2, 3]. Declining trust in politicians, widespread disillusionment with democratic processes, and the rise of populist movements, that intrinsically distrust procedures and institutions [4], have led to a growing disengagement from the *res publica*, resulting in a progressive widespread decline of turnout, as shown by [5] in Figure 1.1.

Various novel methods have emerged to quantify, analyze, and address these tensions. In particular, participatory[2] and decentralized[3] approaches have been proposed to increase citizen engagement and rebuild democratic legitimacy. These models are rooted in the idea that legitimacy is not only derived from electoral representation, but also from the extent to which citizens can contribute to shaping decisions that affect their lives.

Yet many of these systems encounter a structural trade-off: as the scale of participation increases, the depth and quality of deliberation often suffer. Participatory methods are powerful tools for fostering meaningful discussions, but their effectiveness tends to decline sharply as the number of participants grows. One potential solution lies in introducing multi-layered filtering mechanisms, such as those adopted by the vTaiwan platform [6], that manage large-scale input without sacrificing deliberative quality. These methods attempt to structure the democratic process in such a way that wide-scale participation becomes manageable,

[1] In my work, I adopt the definition of democracy provided by Mansbridge et al. in [1]: "a normative theory of democratic legitimacy based on the idea that those affected by a collective decision have the right, opportunity, and capacity to participate in consequential deliberation about the content of decisions. [. . .] A functioning deliberative democracy requires not just deliberative forums, but also a larger process of broad scale public deliberation encapsulated most recently in the concept of a deliberative system."

[2] A participatory framework actively involves citizens in decision-making processes, beyond the minimal act of voting.

[3] A decentralized system distributes authority and decision-making power across multiple individuals or groups, rather than concentrating it in a central authority.

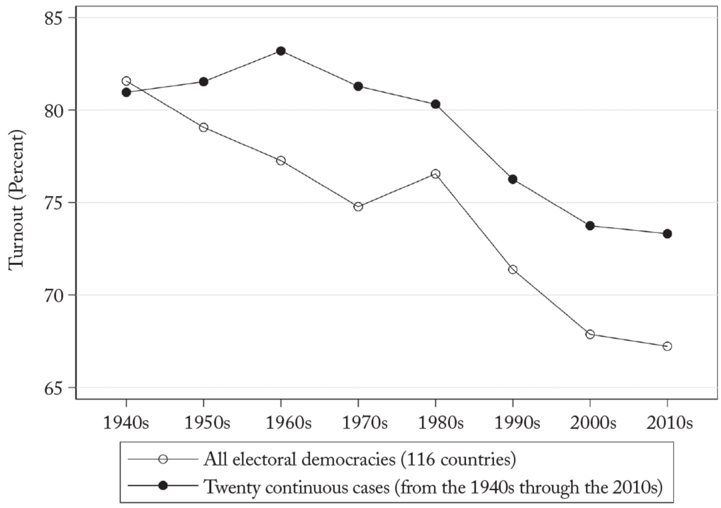

Figure 1.1: Evolution of Voter Turnout in National Elections 1945–2017. From [5].

without compromising the constructive value of individual contributions.

1.2 Computational Methods to Enhance Democracy

While representative democracy has been the most common political form over the last two centuries, recent technological and scientific advancements have opened possibilities for more direct, real-time, and inclusive democratic innovations that rely on computational tools to simulate and enhance collective decision-making and its efficiency [7, 8]. Among these tools, agent-based modeling (ABM), conceptualized in the 1940s but only effectively implemented in the 1970s [9], stands out as a powerful method for exploring how individual behaviors and interactions give rise to emergent, system-level outcomes.

Agent-based simulations are used to model complex (social) systems composed of interacting autonomous agents, that represent individual entities, such as people, organizations, or decision-makers. With their own set of rules, behaviors, and goals, agents interact and often produce large-scale patterns that could not be predicted by examining the agents in isolation. These tools allow researchers and policymakers to test "what-if" scenarios in a controlled, replicable environment and to understand how small-scale changes might lead to large-scale effects.

In the context of participatory democracy, agent-based simulations can help policymakers understand how different governance rules, voting systems, incentive structures, or communication protocols might influence collective dynamics, opinion formation, and the quality of decisions of a certain population of agents.

Modeling agents with heterogeneous preferences, agent-based models provide insight into how decentralized systems function in practice, and under what conditions collective intelligence emerges or fails. This provides a valuable basis for the design of democratic institutions that are both inclusive and resilient, capable of harnessing the diversity of the population in a constructive way.

1.3 The Co-Creation Platform

With this work, I present a simulation of a digital **Co-Creation Platform** (CCP) that aims to address this scalability-participation trade off. It is designed as a horizontal, filter-free system that remains inclusive, scalable, and portable. This platform simulation offers several advantages over traditional participatory formats: it reduces important limitations of participatory systems such as scheduling conflicts, social anxiety[4], and the marginalization of minority voices, while enabling broader and more flexible participation. Although the digital divide[5] remains a relevant concern, online frameworks generally enhance accessibility and inclusiveness. In this sense, digital platforms do not simply scale participation, but reconfigure the very structure of participation, allowing it to take place asynchronously, anonymously, and across geographic and social boundaries.

This platform is grounded in the principle of collective intelligence, defined as the enhanced problem-solving and decision-making capability that emerges when groups of individuals or agents collaboratively interact [10, 11]. It involves integrating diverse perspectives, knowledge, and strategies to achieve outcomes superior to those achievable individually [12]. In societal and political contexts, collective intelligence can be harnessed to improve the quality of decisions through a decentralized approach that mixes citizen input, expert knowledge, and stakeholder feedback. Similarly, diplomatic negotiation frameworks can benefit from the interaction of agent-based models and collective intelligence mechanisms [13], which facilitate better understanding of trade-offs, foster compromise, and improve outcomes by effectively aggregating diverse viewpoints.

The design of the platform explicitly responds to the challenges identified in recent deliberative systems theory. In particular, it incorporates the insights of Ercan, Hendriks, and Dryzek [14], who warn that simply multiplying voices in the public sphere, what they term communicative plenty, can lead to democratic dysfunction if not accompanied by mechanisms for listening and reflection. The framework addresses this by structuring asynchronous input and interaction to

[4]Social anxiety, also called social phobia, is defined as a long-term and overwhelming fear of social situations.

[5]The digital divide is commonly defined as the gap between those who have access to and can effectively use digital technologies, like the internet and computers, and those who do not. Typically, it correlates with age.

promote meaningful deliberation, rather than unmoderated expression.

One of the major challenges for contemporary democracies, as pointed out in [15], is to balance individual preferences with the collective good. With this report, I aim to contribute to the ongoing efforts of enhancing democratic processes by proposing a novel digital platform that integrates collective intelligence principles and agent-based modeling to address the scalability-participation tradeoff in a decentralized democratic framework. I argue that the success of such systems depends not only on the intelligence of the participants, but critically on the design of the participatory framework itself. A naive approach would suggest that more participants automatically lead to better outcomes; however, research in collective intelligence shows that the system's design, which comprises the rules, structures, and feedback mechanisms, plays a fundamental role in determining whether participation results in meaningful and constructive outcomes. My work is thus also a contribution to the broader question of how democratic frameworks can be designed to channel diverse contributions into coherent, high-quality collective decisions.

Chapter 2

Scenario and Implementation

2.1 Scenario

In this work, I simulate a collective problem-solving process using an agent-based modeling framework. The model consists of a fixed population P of agents that interact over multiple discrete time steps from $t = 0$ to $t = T$. These agents are assigned to two distinct, mutually exclusive roles: *stakeholders* and *solvers*. Stakeholders are individuals directly involved in or affected by a given problem. This group includes citizens experiencing the issue, as well as institutional actors such as constitutional judges, budget holders, or public officials responsible for implementing decisions. Solvers, on the other hand, are agents who voluntarily contribute by generating candidate solutions to the shared problem. They may or may not be directly affected by it. Once a solution is proposed by a solver, it is evaluated by the stakeholders, who assign it a utility score based on their individual preferences. This division captures the realistic scenario in which problem-solvers do not have direct access to stakeholder preferences and must instead infer the quality of their proposals from the feedback they receive.

At each iteration, a solver is tasked with proposing a new idea. This proposed solution is then evaluated by the entire population of stakeholders through a voting process. Each stakeholder casts a vote based on their individual utility function, which maps solution features[1] to a satisfaction score $u \in \mathbb{R}$. Then, the weighted aggregate score determines the *population utility* of the solution. This scalar value is the only feedback returned to the solver on its proposed idea.

Stakeholders are assumed to be heterogeneous, each possessing a unique and unknown utility function. This models the diversity of opinions typically found in real-world decision-making contexts.

Solvers are categorized into three behavioral subgroups, each employing a different strategy to generate new solution ideas:

[1] Without loss of generality, I make no assumption about the dimensionality of the solution space, allowing for a representation that encompasses the large number of characteristics on which a solution may be evaluated from a stakeholder.

- **Innovators**: These agents propose entirely new solutions, sampled randomly from an infinite-dimensional solution space $\mathcal{S}$. They represent the exploratory force of the population and are unconstrained by past proposals.
- **General Modifiers**: Given a non-empty set of previously proposed solutions, agents of this type randomly select one from the pool and apply a small perturbation to it.
- **Best Modifiers**: Given a non-empty set of proposed solutions, these agents identify the current best-scoring solution (i.e., the one with the highest population utility score from previous rounds) and apply a small modification to it using the same perturbation mechanism as General Modifiers.

The categorization of solvers into the three groups above is inspired by cognitive theory, particularly the distinction introduced by [16] between *adaptors* (our *modifiers*) and *innovators*. While innovators prefer to challenge established paradigms and generate radically new ideas, adaptors focus more on refining existing solutions within the current structure. My model operationalizes this distinction: Innovators represent solvers exploring new regions of the solution space; General Modifiers act as intermediate agents capable of recombination and broader adaptation; and Best Modifiers embody local optimizers, conservative and risk-averse, aligned with the adaptor profile.

The ratio of these three types is an input parameter of the model and can be tuned to study how different group compositions affect the dynamics of idea evolution. Innovators intuitively facilitate exploration, General Modifiers increase coverage (a mix of exploration and exploitation) of the solution space, and Best Modifiers focus the search around the most promising candidate, exploiting the local maximum randomly found in the solution space.

In this paper, I will call *innovative solutions* solutions proposed by innovators, and *derived solutions* all solutions proposed by General or Best Modifiers.

Moreover, I define a *solution path* as the sequence of solutions consisting of a leaf solution and all of its ancestral solutions, traced back recursively to the original solution proposed by an Innovator. The *depth* of a solution is the number of derived solutions that that solution generated.

2.1.1 Utility and Potential

Each solution s is associated with two scalar values: a *utility* score $u(s)$ and a *potential* score $p(s)$.

The utility $u(s)$ of a solution represents the aggregated satisfaction of stakeholders in response to that solution. For solutions proposed by Innovators, utility is sampled randomly from a distribution intended to approximate the weighted average of individual stakeholder preferences. This randomness reflects the assumption that solvers cannot access stakeholder utility functions directly, and instead must rely on noisy feedback.

For derived solutions, generated by General or Best Modifiers, utility is computed as a small stochastic perturbation of the utility of the parent solution:

$$u_{\text{child}} = u_{\text{parent}} + \epsilon, \quad \epsilon \sim \mathcal{N}(0, p),$$

where p is the potential, which quantifies the solution's capacity for further improvement through modification. Intuitively, the potential reflects how much the utility of a solution could vary (positively or negatively) if slightly altered. In this sense, it functions similarly to the concept of volatility in financial contexts, indicating wether a solution is high-risk and potentially high-reward.

Potential is also randomly initialized for solutions introduced by Innovators and is assumed to be independent of utility. For derived solutions, potential decays with depth and includes stochastic perturbation:

$$p_{\text{child}} = \eta \cdot p_{\text{parent}} \cdot r^d, \quad \eta \sim \mathcal{N}(1, \epsilon),$$

where $r \in [0, 1]$ is the decay rate, d is the derivation depth (i.e., how many times the original idea has been modified), $0 < \epsilon \ll 1$ induces a perturbation noise, and η introduces noise in decay strength.

This decay models the intuitive idea that, as a solution undergoes repeated modification, its room for improvement narrows, eventually reaching a point of saturation.

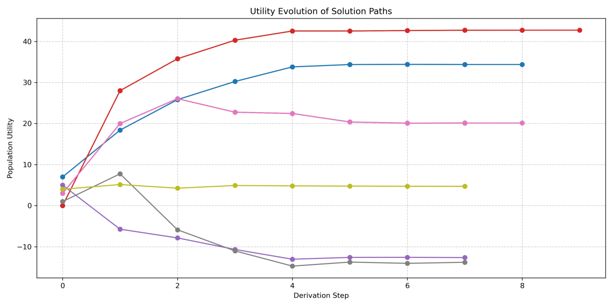

Figure 2.1: Utility evolution of different solution paths. Each line represents the sequence of utility values for a specific path, starting from an Innovator's proposal and proceeding through successive modifications by General or Best Modifiers.

Example Figures 2.1 and 2.2 illustrate the evolution of utility and potential values for a selected set of solution paths. Each path begins with an original proposal by an Innovator and continues through successive modifications. I observe

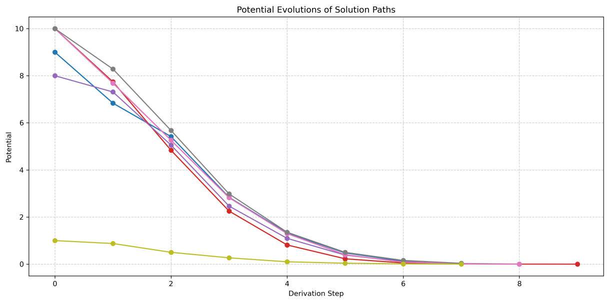

Figure 2.2: Potential evolution of different solution paths. Each line shows how potential decays with derivation steps. As expected, potential declines exponentially due to the decay mechanism and noise, modeling the diminishing returns of iterative refinement. Paths with initially low potential (e.g., green) display little change in utility in Figure 2.1.

that solution paths exhibiting significant increases or decreases in utility tend to start with high potential values. On the other hand, the green path, which is characterized by an extremely low initial potential, shows minimal variation in utility across derivation steps, remaining nearly constant.

2.1.2 Reward Methods

A key component of the evolutionary version of my simulation is the *reward method* R, that determines how solver agents are evaluated and selected for reproduction. The reward mechanism effectively shapes the evolutionary pressure within the system: it defines which agent behaviors are promoted and which are penalized.

Reward is applied periodically (every r'th iteration) to the system and it determines which agents are more likely to reproduce or be replaced. A well-designed reward system should encourage useful exploration while discouraging stagnation or inefficient strategies. Conversely, a poorly designed reward should lead to suboptimal performance.

It is important to note that this reward mechanism is a simplification of more complex evaluation and selection dynamics that might exist in real-world co-creation or problem-solving settings.

In this platform, I experiment with many different reward systems, each emphasizing different aspects of solver behavior:

- **Best Score:** I reward only the solver who produced the single highest utility solution over the last r iterations. Mathematically, if a^* is the agent who proposed solution $s^* = \arg\max_{s \in S_r} u(s)$, where S_r is the set of solutions from the last r iterations, then

$$R_i = \begin{cases} 1 & \text{if } i = a^* \\ 0 & \text{otherwise} \end{cases}$$

- **Last Score:** I reward the solvers proportionally to their last proposed solution's score. If a solver does not propose any solution over the last r iterations, its reward is automatically set to zero. Mathematically:

$$R_i = u(s_i^{\text{last}}) \quad \text{if } s_i^{\text{last}} \text{ exists, else } 0$$

- **Average Score:** I reward the solvers proportionally to their solutions' average score (for example, if agent a has proposed two solutions, with values 10 and 20 respectively, its reward will be 15). Mathematically:

$$R_i = \frac{1}{N_i} \sum_{j=1}^{N_i} u(s_{ij})$$

  where N_i is the number of solutions proposed by solver i and s_{ij} is the j-th solution.

- **Depth:** Solvers are rewarded based on the average derivation depth of their proposed solutions, promoting exploration and long term refinement:

$$R_i = \frac{1}{N_i} \sum_{j=1}^{N_i} d(s_{ij})$$

  where $d(s)$ is the derivation depth of solution s.

- **Improvement:** Solvers are rewarded based on the relative improvement of their most recent solution compared to their prior solutions. Let u_{recent} be the last utility and $u_{\max}$ be the highest previous utility over the last r iterations:

$$R_i = \max(0, u_{\text{recent}} - u_{\max})$$

- **Improvement and Influence:** This composite reward scheme incentivizes both personal progress and collaborative impact. Agents are rewarded for: (1) improvement in their solution performance over a recent window of time, and (2) the number of times their solutions are modified by other agents. Mathematically, let I_i be the improvement score for agent i based on the difference between maximum and minimum utility in the recent window of their proposals, and M_i be the number of times their ideas were used as the basis for derived solutions. Then, the total reward is given by $R_i = \alpha \cdot I_i + \beta \cdot M_i$, where α and β are tunable weights that balance personal improvement and social influence.

- **Solver Type Only:** I only reward a certain type of solvers.
- **Random:** Baseline to compare all other rewards, the random method rewards agents randomly.

I foresee that each reward function, or combination of reward functions, biases the population composition differently. For instance, rewarding only the best solutions tends to favor Best Modifiers, who refine known high-utility ideas. In contrast, improvement-based or risk-based rewards provide more space for Innovators and General Modifiers to contribute meaningfully.

2.2 Implementation

The agent-based simulation is implemented and run in Python. The agents' behaviors, their interaction with a memory pool of past solutions, and their evolutionary adaptation are governed by a set of rules detailed below.

Agents and Strategies Solver agents are initialized with two behavioral parameters:

- **Modification probability** ($m \in [0, 1]$): This parameter determines the likelihood that an agent modifies an existing solution rather than proposing a novel one.
- **Best bias** ($b \in [0, 1]$): Conditional on modifying an existing solution, this parameter defines the probability of selecting the current best solution in the population over a random one, given a non-empty pool of existing solutions.

These two parameters operationalize the distinction between Innovators ($m = 0, b = 0$), General Modifiers ($m = 1, b = 0$), and Best Modifiers ($m = 1, b = 1$), while also allowing intermediate strategies to emerge during evolution. Each agent maintains a history of the utility values of the solutions it has proposed, used in some reward schemes for performance evaluation.

2.2.1 Solution Pool and Lineage Tracking

All proposed solutions are stored in a shared memory pool that tracks the following attributes: unique solution ID, utility score, reference to the parent solution (if derived), potential score, derivation depth, agent ID of the proposer. Each solution lineage is then traced using a reference map, enabling the computation of derivation depth and path-based statistics. This allows the system to model idea evolution as tree-like structures rooted in innovations (see Figures 2.1 and 2.2).

2.2.2 Proposal and Evaluation Process

At each round t, a randomly selected agent proposes a new solution. With probability $1-m$, the agent introduces a novel idea; with probability m, it selects an existing solution from the pool, either the best known (with probability b) or a random one (with probability $1-b$), and perturbs it slightly to generate a derived solution.

For innovative solutions, utility and potential are computed as two random numbers between 0 and 10. For derived solutions, utility is computed as a noisy perturbation of the parent solution's utility, and potential is updated according to a decay function, as defined in Section 2.1.1. If the new solution's utility exceeds the global best so far, it becomes the new reference for future Best Modifiers.

2.2.3 Reward and Evolution Mechanism

Every r iterations, the system evaluates solver performance using one of the reward functions described previously. Based on rewards, a fixed fraction of the lowest-performing agents is removed from the population and replaced with mutated copies of top performers. This evolutionary update promotes behavioral strategies that succeed under the current reward regime, allowing the agent population to adapt to different task environments and incentive structures.

Chapter 3

Results and Discussion

In this section, I present and interpret the results of my simulation. I begin by analyzing whether particular population compositions are consistently associated with better collective performance. In this first part, no reward is put in practice, and the evolutionary aspect is not present. This lets us assess which type of composition I should make emerge under evolutionary pressure, and what share of the population tends to contribute most effectively to the discovery of high-utility solutions. Then, I can turn my attention to the reward system. I evaluate a variety of reward schemes and I assess the extent to which these reward mechanisms succeed in promoting high-quality solutions and sustaining an effective balance between exploration and exploitation in the agent population. Discussions follow both simulation types and results.

3.1 Static Simulation

I first analyze and discuss simulations where the solver population is held fixed, and only the proportion of Innovators is varied. General Modifiers and Best Modifiers share the remaining fraction equally. The results are shown in Figure 3.1.

This figure clearly illustrates a non-linear relationship between solver diversity and performance. Populations composed exclusively of Innovators (far right) or exclusively of Modifiers (far left) perform significantly worse than mixed populations. The best outcomes are observed when Innovators range from 15 to 85% of the total population. This highlights the importance of maintaining a balanced strategy mix: Innovators inject novelty, but without Modifiers to refine and explore promising leads, innovation becomes chaotic and unfocused. Conversely, Modifiers alone exploit local information but lack the creative input necessary to escape suboptimal regions of the solution space.

To explore this dynamic further, Figure 3.2 shows a heatmap of final solution quality across the entire space of solver compositions. Here, I vary both the fraction of Innovators and the fraction of General Modifiers; Best Modifiers con-

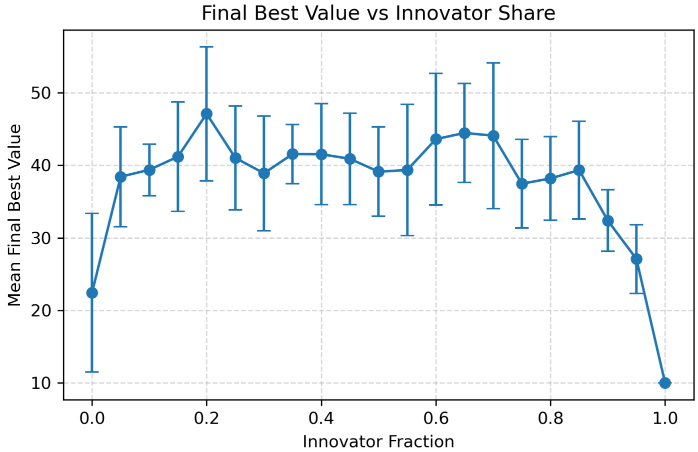

Figure 3.1: Final solution quality as a function of the share of Innovators in the population. The remaining population is equally split between General and Best Modifiers. Peak performance emerges when Innovators constitute a moderate proportion, underscoring the need for balanced strategic diversity.

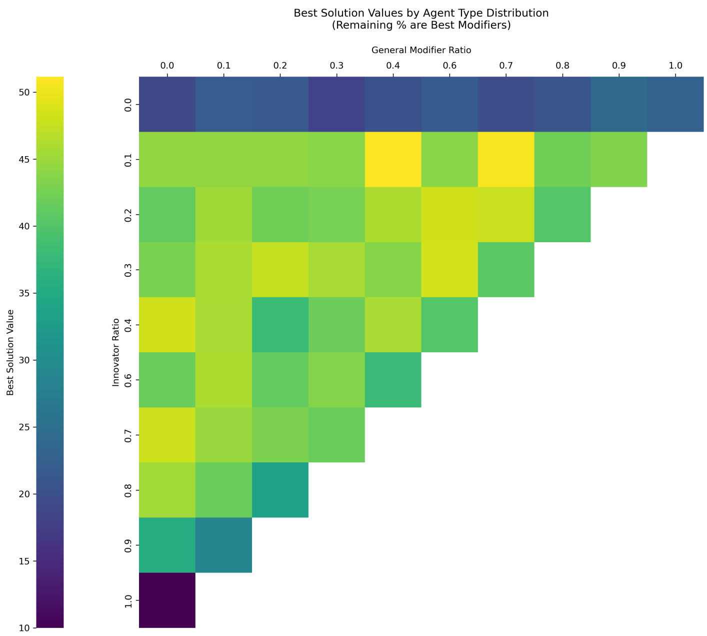

Figure 3.2: Heatmap of final solution quality across varying solver compositions. Each point corresponds to a specific mix of Innovators and General Modifiers, with Best Modifiers filling the remaining share. The highest-quality outcomes occur in regions where all three solver types coexist.

stitute the remaining fraction. This richer visualization confirms and extends the findings of Figure 3.1. The highest-performing configurations are located in the region where all three types are present in balanced proportions. Systems that exclude either Innovators or General Modifiers almost always perform worse.

These results suggest that the most robust collective intelligence emerges not from dominance by any one solver type, but from the interaction among exploration (Innovators), broad refinement (General Modifiers), and focused optimization (Best Modifiers). Reward systems and platform designs should therefore aim to preserve and support this strategic diversity. While this simulation do not determine a universal optimal distribution of solver types, it clearly highlight the necessity of maintaining behavioral diversity. Since the reward system can influence diversity and shape a solver population, I now add the evolutionary perspective to the model.

3.2 Evolutionary Simulation

Now I move to the evolutionary simulation, where I add a reward scheme that dynamically changes the population share over time. Table 3.1 summarizes the average performance of each reward scheme, whose descriptions can be found in Section 2.1.2, over several runs.

Reward Scheme	**Average Best Value ± Standard Deviation**
Random	42.65 ± 10.36
Last Score	41.81 ± 7.64
Average Score	48.02 ± 5.56
Best Score	44.35 ± 8.29
Improvement	39.76 ± 6.61
Depth	44.47 ± 7.72
Improvement and Influence	48.75 ± 3.98
Innovator Only	29.25 ± 12.49
General Modifier Only	40.24 ± 7.84
Best Modifier Only	35.74 ± 8.78

Table 3.1: Performance of each reward scheme over 20,000 rounds, where the reward is applied every 100 iterations with a replacement fraction of 1/3.

From Table 3.1, I see that the **Improvement and Influence** and **Average Score** reward mechanisms yielded the highest average best values. However, their advantage over other schemes is not statistically significant. In fact, their performance overlaps with that of simpler reward mechanisms such as **Best Score**, or **Best Contributor**. Even the baseline **Random** scheme achieves competitive performance. This underscores an important insight: no single reward scheme guarantees superior performance, and the relative effectiveness of each depends heavily on system dynamics and contextual factors.

Figure 3.3 illustrates the population dynamics under the **Improvement and Influence** scheme, which was among the top-performing approaches. Notably, I see that solver diversity is preserved throughout the simulation: Innovators and General Modifiers make up the bulk of the population, while Best Modifiers remain a small, consistent fraction.

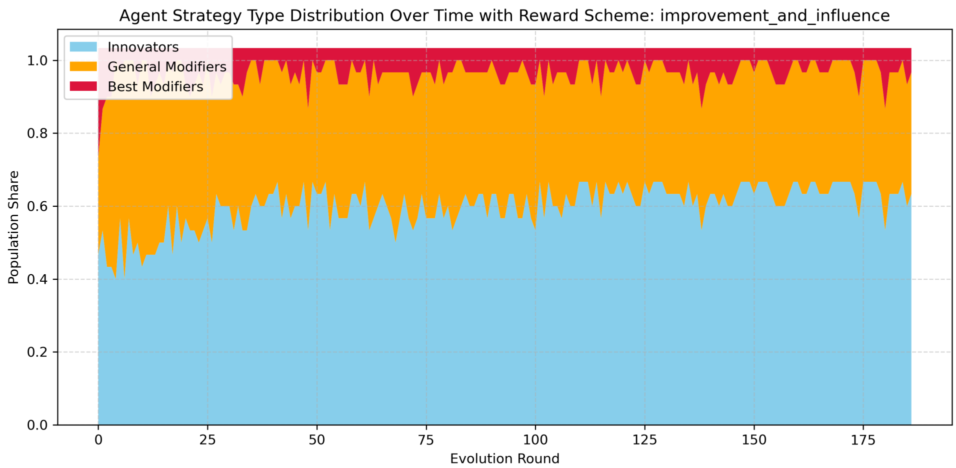

Figure 3.3: Example of solver type distribution over time under the **Improvement and Influence** reward with *alpha* = 1 and *beta* = 5. The population maintains a diverse mix of Innovators, General Modifiers, and Best Modifiers throughout the simulation. Similar combinations of parameters yielded similar results.

This stability is in contrast with reward schemes that favor only a single strategy, such as **Best Modifier Only** or **Innovator Only**, where the population quickly collapses into a homogenous composition. In those cases, as expected, I observe a performance stagnation, with no further gains after early improvements.

These results suggest that the long-term effectiveness of a reward system is less about identifying the "best" individual behaviors at a certain time and more about its capacity to sustain strategic diversity. The capacity of a system to mix exploration (via Innovators), recombination (General Modifiers), and refinement (Best Modifiers) appears crucial for avoiding premature convergence and ensuring continued progress.

In summary, I can state the most important determinant of success is not the specific reward formula itself, but whether the reward scheme encourages and preserves a heterogeneous population. Reward mechanisms that implicitly or explicitly support solver diversity are more likely to sustain collective intelligence over time.

3.3 Discussion

Across simulations, I consistently find that homogeneous populations, consisting entirely of one solver type, underperform. Populations composed solely of Innovators, for instance, fail to build on good ideas, resulting in stagnation. Pure Best Modifier populations tend to fixate on early solutions and suffer from a lack of novelty. Even General Modifiers, while more balanced, require external injection of quality ideas to function effectively.

The best performing runs all share a common trait: a dynamic equilibrium among the three solver types. Rather than converging to a single optimal strategy, successful systems cycle through periods of innovation, general modification, and refinement. The evolutionary pressures imposed by reward mechanisms regulate this cycling, sometimes amplifying one phase or suppressing another.

Interestingly, the random reward system where agents are rewarded independently of their performance, yields results comparable to more structured reward schemes. This outcome suggests that mere participation, without targeted incentives, can still support moderate levels of system performance.

Implications for Participatory Platforms My results suggest that the most crucial factor in collective problem-solving environments is not the precise configuration of the reward system, but rather the presence of cognitive diversity among solvers. Platforms that overly favor a single type of contribution,such as consistently rewarding only the "best idea",risk undermining the exploratory and combinatorial dynamics necessary for innovation and long-term improvement.

In this light, the primary design goal of participatory platforms should be to maximize accessibility and inclusiveness. Ensuring that a wide range of individuals can participate,across different backgrounds, experiences, and perspectives,is key to fostering the strategic diversity that sustains effective collective intelligence. The Co-Creation Platform (CCP) I propose aims to enable such participation by removing typical barriers: it is asynchronous, anonymous, and geographically unbounded, reducing exclusion due to time constraints, social pressure, or location.

Rather than searching for a single optimal reward rule, platform designers should focus on creating conditions that attract and support a heterogeneous population of contributors. When participation is open and diverse, simple or even random reward mechanisms can suffice to maintain dynamism. In this sense, architecture and accessibility seem to matter more than finely-tuned incentives in shaping the platform's collective outcomes.

3.4 Limitations and Future Works

Several simplifying assumptions underpin the model:

- **Rationality and Homogeneity:** I assume that solvers and stakeholders behave as *homo oeconomicus*, meaning they act in a fully rational manner to maximize their respective objectives (e.g.: solvers aim to maximize rewards and, indirectly, lifetime, while stakeholders provide evaluations aligned with unknown utility functions). Moreover, all solvers are incentivized in the same way: given a certain reward value R, they are equally (de-)motivated to keep proposing solutions.
- **Solver independence:** Solvers do not explicitly collaborate or communicate; coordination emerges implicitly through solution inheritance and reward shaping.
- **Static stakeholders:** Stakeholder preferences are fixed and unknown to solvers. Future work may include adaptive or learning stakeholders.
- **One-shot proposals:** Each round involves one agent proposing one solution. Parallel or batch submissions are not currently modeled.
- **No strategic behavior:** Solvers do not engage in forward planning or strategic anticipation of the reward structure. They operate purely reactively.
- **Unlimited memory:** The solution pool retains all past ideas. Memory constraints or aging mechanisms are not considered but may impact scalability in larger systems.

Despite these limitations, my digital Co-Creation Platform (CCP) represents a promising approach to scaling participation in democratic processes. While the optimal reward system remains an open question, my results demonstrate that a well-structured platform, supporting asynchronous, inclusive, and diverse participation, is itself a powerful enabler of collective intelligence.

Bibliography

[1] J. Mansbridge, J. Bohman, S. Chambers, D. Estlund, A. Føllesdal, A. Fung, C. Lafont, B. Manin, and J. L. Martí, “A systemic approach to deliberative democracy,” in *Deliberative Systems*, J. Parkinson and J. Mansbridge, Eds. Cambridge University Press, 2012, pp. 1–26. [Online]. Available: https://doi.org/10.1017/CBO9781139178914.002

[2] I. Chedia, “Democracy in crisis,” *Copernicus Journal of Political Studies*, no. 2/2020, pp. 73–83, 2020.

[3] B. Weiffen, “The crisis of democracy: Conceptualization and empirical assessment (with special consideration of europe),” in *IPSA World Congress*, Brisbane, Australia, Jul. 2018, paper presented at the IPSA World Congress, July 21–25.

[4] J. Webber, “Understanding populism,” *Social & Legal Studies*, vol. 32, no. 6, pp. 849–876, 2023. [Online]. Available: https://doi.org/10.1177/09646639231156144

[5] F. Kostelka and A. Blais, “The generational and institutional sources of the global decline in voter turnout,” *World Politics*, vol. 73, no. 4, pp. 629–667, 2021. [Online]. Available: https://doi.org/10.1017/S0043887121000149

[6] vTaiwan, “vtaiwan: Public participation platform,” https://info.vtaiwan.tw/, accessed: 2025-06-18.

[7] A. Mikhaylovskaya, “Enhancing deliberation with digital democratic innovations,” *Philosophy & Technology*, vol. 37, no. 3, p. 3, Mar. 2024. [Online]. Available: https://link.springer.com/article/10.1007/s13347-023-00692-x

[8] J. Davies and R. Procter, “Online platforms of public participation: a deliberative democracy or a delusion?” in *Proceedings of the 13th International Conference on Theory and Practice of Electronic Governance*. Athens, Greece: ACM, 2020, pp. 746–753. [Online]. Available: https://dl.acm.org/doi/10.1145/3428502.3428614

[9] T. C. Schelling, “Dynamic models of segregation,” *The Journal of Mathematical Sociology*, vol. 1, no. 2, pp. 143–186, 1971. [Online]. Available: https://doi.org/10.1080/0022250x.1971.9989794

[10] A. W. Woolley, C. F. Chabris, A. Pentland, N. Hashmi, and T. W. Malone, "Evidence for a collective intelligence factor in the performance of human groups," *Science*, vol. 330, no. 6004, pp. 686–688, Oct. 2010.

[11] T. W. Malone and M. S. Bernstein, *Handbook of Collective Intelligence*. Cambridge, MA, USA: MIT Press, 2015.

[12] J. Surowiecki, *The Wisdom of Crowds*. New York, USA: Anchor Books, 2005.

[13] L.-E. Cederman and L. Girardin, "Computational approaches to conflict research from modeling and data to computational diplomacy," *Journal of Computational Science*, vol. 72, p. 102112, Sep. 2023. [Online]. Available: https://doi.org/10.1016/j.jocs.2023.102112

[14] S. A. Ercan, C. M. Hendriks, and J. S. Dryzek, "Public deliberation in an era of communicative plenty," *Policy and Politics*, vol. 47, no. 1, pp. 19–35, 2019.

[15] E. G. Weyl, A. Tang, and Community, *Plurality: The Future of Collaborative Technology and Democracy*. RadicalxChange, 2024. [Online]. Available: https://books.google.ch/books?id=-ye20AEACAAJ

[16] M. Kirton, "Adaptors and innovators: A description and measure," *Journal of Applied Psychology*, vol. 61, no. 5, pp. 622–629, 1976.