



WISDOM OF AI CROWDS: A PROPOSED THREE-PHASE POLICY SIMULATION FRAMEWORK

Gian Paul Gauci¹ⁱ,

Alexiei Dingli²

¹PhD Candidate,

Department of Policy, Politics and Governance,

University of Malta,

Malta

²Professor,

Department of Artificial Intelligence,

University of Malta,

Malta

Abstract:

'Wicked Problems' are persistent societal issues characterised by complexity, divergent stakeholder perspectives, and resistance to traditional solutions. A key barrier to addressing them is the inability to experiment with potential policies safely, given their dynamic, high-stakes, and often 'one-shot' nature in the real world. This paper proposes the 'Wisdom of AI Crowds,' a novel conceptual framework designed to overcome this barrier. It employs artificial societies populated by agentic AI, whose personas are grounded in empirical data reflecting stakeholder norms, values, and beliefs, within a three-phase process: Input, Simulation, and Human-in-the-loop Feedback. The original contribution of this framework lies in its integration of agentic AI within a dynamic, iterative simulation environment. Unlike prior static mapping or high-risk incremental approaches, the 'Wisdom of AI Crowds' provides a risk-free virtual laboratory to test multiple policy scenarios, observe emergent behaviours over time, and incorporate expert validation before real-world implementation. This approach offers the potential to shift policymaking for wicked problems from reactive interventions to proactive, evidence-based experimentation, enabling the identification of more robust, well-vetted policy options. The framework explicitly incorporates considerations for ethical challenges, data representativeness, and simulation validation.

Keywords: wicked problems; public policy; agentic AI; artificial societies; policy simulations; agent-based modelling; computational social science

ⁱ Correspondence: email gian-paul.gauci@um.edu.mt

1. Introduction

One of the key functions of policymakers is to solve policy and societal issues. However, significant issues such as climate change, poverty and criminality persist despite any attempt to tackle these issues. The prevailing theory on why these problems prevail was introduced in 1973 by Rittel and Webber in their essay, which described these issues as “*wicked problems*”: problems which evade classic solutions. They described these problems as having ten distinct characteristics. Over the years, academics have discussed these issues and contributed to the theory, converging on the idea that these problems are exacerbated by differences among the stakeholders involved in tackling them. Other academics have tried to create solution frameworks for these problems. Notwithstanding this, to date, there is no consensus on how to address these perennial challenges. This paper addresses this gap by introducing a new framework called “Wisdom of AI crowds.” It posits that building an artificial society composed of agentic AI, based on empirical stakeholder data, creates a safe and innovative experimental space to address these challenges.

This paper proceeds as follows: First, it synthesises the academic literature on wicked problems, outlining their core characteristics and critically reviewing existing solution frameworks to identify a key research gap. Second, it briefly introduces the relevant concepts from Artificial Intelligence, specifically agentic AI and artificial societies, necessary to understand the proposed solution. Third, it presents the paper's original contribution: the ‘Wisdom of AI Crowds,’ a novel, three-phase conceptual framework for simulating wicked policy problems. Fourth, it provides a proof of concept, illustrating how the framework would be applied using the Maltese Land Grab as an example. Fifth, it proactively addresses the key methodological and ethical challenges inherent in this approach, outlining built-in mitigation strategies. Finally, the paper concludes by discussing the potential implications of this framework for policy research and practice.

2. Literature Review

2.1 A synthesis of the Literature

In “Dilemmas in General Theory of Planning,” Rittel and Webber described persistent issues. They theorise that these problems, described as ‘wicked’, have ten distinct characteristics. These characteristics, visualised in Figure 1, show that the issue essentially lies in the lack of understanding of what makes the problem a problem. However, of importance to this paper are three of these characteristics, namely, “*Solutions are one-shot*”, “*No stopping rule*”, and the “*Designer has no right to be wrong*” (Rittel and Webber 1973, pp.160–166). Fundamentally, this means that real-world experimentation is impossible because any introduction of a possible solution alters the problem, leading to consequences for the solution finder. Any attempted solution is a high-stakes, one-shot gamble that irreversibly alters the problem itself.

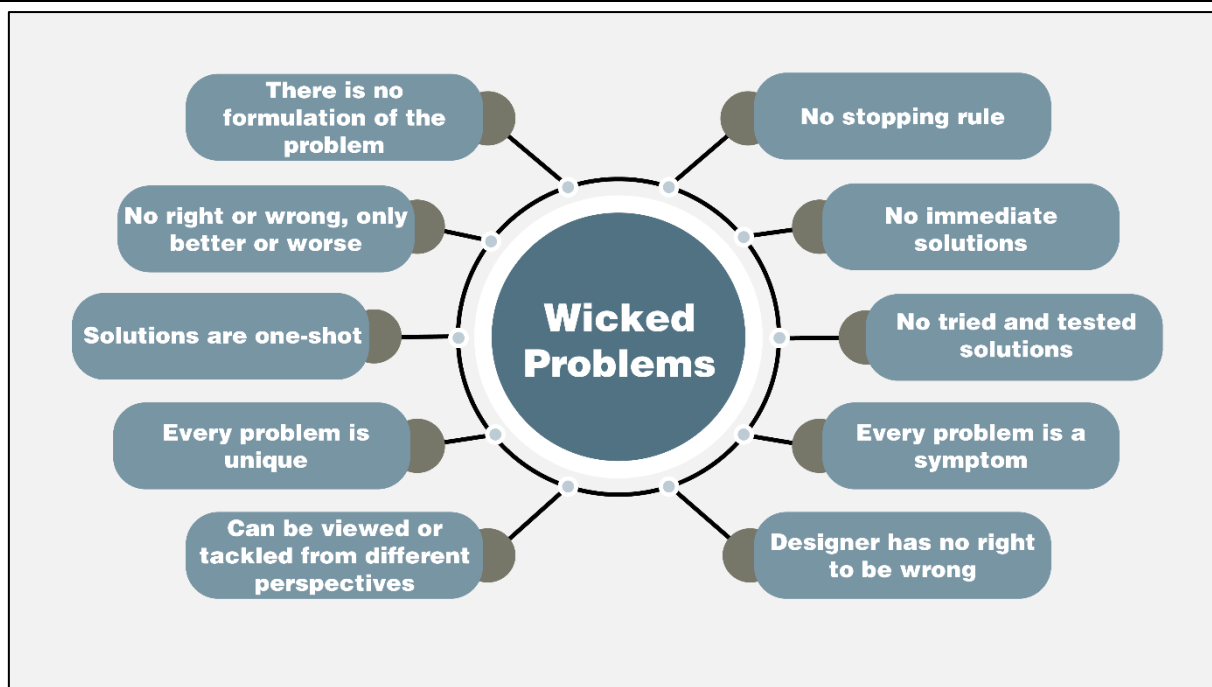


Figure 1: The ten characteristics of wicked problems (synthesised from Rittel & Webber 1973)

While Rittel and Webber (1973) mainly described the “how” of wicked problems, previous theorists in decision-making might have described the “why”. Lindblom (1959) and Simon (1957) had clearly outlined that the problem in decision-making is the fact that human cognition is limited (Lindblom 1959) and our rationality is bounded (Simon 1957). Therefore, decisions are taken by ‘satisficing’ (Simon 1957, p.204) - decisions which are deemed satisfactory enough, or by ‘muddling through’ - administrators choosing a limited number of alternatives based on their values (Lindblom 1959, p.88). This is not, however, sufficient for wicked problems since the problem is too large for human cognition to comprehend fully, leading to Rittel and Webber’s (1973) first characteristic: there is no formulation of the problem.

Later academics identified that these challenges are also complex to solve due to the differences in stakeholders' views on the problem. Fischer (1993) noted that even if science and those in power support the solutions, citizens may not accept such a solution unless they are intrinsically involved in the decision-making process. Later theorists agreed and insisted that wickedness is determined by both the complexity of the problem and the divergences between the different stakeholders. In fact, if the problem is complex and there is total divergence between different stakeholders, these are regarded as “*very wicked problems*” (Alford and Head 2017, p.402). For these reasons, Roberts (2000) identified that solution finding in wicked problems often falls under collaborative leadership, since no single authority has enough power to force a decision. Grint (2010) argues that leaders prefer to frame problems as critical issues since such a frame requires an authoritative solution. Anyone solving such an issue would be deemed a hero, and this would later reflect in polls and elections. However, this does not eliminate the underlying issue. Instead, it often addresses only the symptoms rather than the root causes. Consequently, the core nature of the problem may evolve or change once a

solution targeting the symptoms is implemented, potentially giving a false sense of resolution.

Solution framework literature converged on the key elements needed in a solution: collaboration, mapping, and incrementalism. However, they fail to address key issues outlined by Rittel and Webber, such as the one-shot solution and the notion that the designer has no right to be wrong. In 1992, it was already evident that computational support and broad collaboration were needed. Duncan and Paradise (1992) proposed using a Group Decision-Support System (GDSS). Such a computational tool would help key experts in managing their ideas into solutions. Other scholars insisted that collaborative negotiations through “*policy games*” are necessary to overcome impasses (van Bueren 2003, p.196), together with a bout of humility and informal negotiations (Roberts 2000). However, these approaches, in their insularity, were limited. The few experts’ restrictions in these frameworks fail to incorporate all the stakeholders, whose inclusion is key to finding legitimate solutions (Fischer 1993).

Four distinct frameworks agree that solution-finding starts with organising and understanding the problem. Due to this, mapping tools such as Mess Mapping (Horn and Weber 2007), Dialogue Mapping (Conklin 2005), Computer-Aided General Morphological Analysis (Ritchey 2013) and Problem-Resolution Process (Elia and Margherita 2018) were created, all offering a way to organise the problem into more manageable tasks. One of the reasons why these mapping tools were created was to gather collective intelligence, the idea that the combined wisdom of a heterogeneous crowd can solve problems which individuals cannot (Surowiecki 2004). A number of these frameworks then flip these mapping tools into solution maps, such as the Problem-Resolution Matrix, to understand who is in charge of which part of the problem (Elia and Margherita 2018), Resolution Mapping to create “*simulated hindsight*” (Horn and Weber 2007, p.16) and Solution Space, whereby all possible solutions are shown when scenarios are chosen (Ritchey 2013). These mapping tools take a snapshot of the problem at a point in time and assume that the problem will remain as it is until we find a solution. Their fundamental limitation, however, is that they work only through a static snapshot of a dynamic problem. Rittel and Webber had underlined the “*no stopping*” rule, where the problem is inherently changing, a feature which none of these frameworks can simulate. Even the most advanced strategies for real-world implementation, which are typically an evolution of Lindblom’s classic ‘muddling through’ (1959), do not escape this principal issue. The ‘progressive incrementalism’ proposed by Levin et al., (2012, p. 125) aims to make ‘stick’ and build support over time by ensuring they are impossible to remove following government changes. They work by sequentially introducing small policies. However, each increment is technically still a real-world, ‘one-shot’ intervention with real consequences. The designer still has ‘no right to be wrong.’ The core issue of real-world experimentation still remains. However, with the ascent of new AI technologies, the advantages of each framework can be utilised to build a new conceptual framework that addresses the key limitations exposed in Table 1.

Table 1: Synthesis of the 'wicked problems' solution frameworks

Frameworks	Core idea	Key Ingredients	Limitations (Gap)	How 'Wisdom of AI Crowds' Addresses This Gap
Incrementalism (Lindblom) and Progressive incrementalism (Levin et al.)	Polymaking works best with small, incremental changes, a process known as muddling through. Introduces a strategy to use incrementalism to achieve “sticky” long-term policy solutions.	Policies should be adjusted in a strategic step-by-step incremental manner.	Lacks strategic foresight; can be reactive and limited to the administrator’s bounded rationality; any incremental change is still a high-stakes real-world ‘one-shot’ attempt.	Provides a risk-free environment where incremental policy changes can be simulated, tested, and reverted without real-world consequences. This enables foresight into potential outcomes and allows for exploring a wider range of options than is possible under conditions of bounded rationality.
Collaboration and Computing Tools (Duncan and Paradice, Fischer, Roberts, van Bueren)	All stakeholders should be included in the decision-making process; computational tools are needed to organise all stakeholder input.	The dual principles of stakeholder negotiation and computational support to organise input.	Often limited to small groups of ‘experts,’ which excludes the majority of stakeholders; no solution-finding after impasses.	Scalable to include a diverse range of stakeholder agents. Explicitly models negotiation dynamics and emergent outcomes from disagreements, overcoming the limitations of static consensus models.
Static Mapping (Horn and Weber, Conklin, Ritchey, Elia, and Margherita)	Provided visual tools to map the problem and find a solution. These tools also harness collective intelligence.	Collective intelligence is key in finding solutions, breaking the barrier of bounded rationality.	Static: They cannot model dynamic change. Can only provide a solution to a snapshot of the problem.	Dynamic: Explicitly models the evolving nature of the wicked problem and its environment over time. Allows for testing solutions within this changing context, rather than relying on a static snapshot.

2.2 The ascent of the new Artificial Intelligence (AI)

In the computing world, an Intelligent Agent is an autonomous entity that perceives its environment and acts to achieve specific goals. Distinct from other software, these agents, ranging from virtual assistants to cleaning robots, are autonomous, reactive, and proactive in pursuing their objectives (Russell and Norvig 2020).

While AI has existed since the 1950s, the shift towards its popularity started with the creation of Large Language Models (LLMs) such as ChatGPT. Such models could generate media, such as text, images, videos, and music, at a fast rate and with relatively high accuracy. These models can also adapt to new environments, learn from past “mistakes” faster and easily generalise (Zhao et al. 2023). Intelligent Agents with LLM “brains”, now being called Agentic AI, can reason, adapt, and communicate in natural language, making them highly effective as believable proxies of human behaviour in simulations (Park et al. 2023).

When multiple agents interact within a shared environment, they form a Multi-Agent System (MAS) (Bonabeau 2002). A specific type of MAS, known as an Artificial Society, serves as a virtual laboratory where complex, emergent behaviour of a whole society of agents can be simulated (Branke 2010, p.46). This approach overcomes the limitations of traditional social science experimentation by allowing for controlled experimentation with a heterogeneous population of agents over manipulable time scales (Epstein and Axtell 1996). Two social experiments conducted with agentic AI concluded that, with the current generation, these mimic human interactions. In “Smallville”, twenty-five agents, all with different characters and life stories, were given a single prompt and played out two simulation days. Two specific agents received two additional prompts: one was running for a local election, and the other was hosting a Valentine’s Day party. Agents discussed politics and even invited each other to the Valentine’s Day party, without further human intervention via prompts (Park et al. 2023). In “METAAGENTS,” a simulation was created to mimic job-searching agents at a job fair, interacting with recruitment company agents. The job-searching agents, similar to human behaviour, lied in their interviews by embellishing their accomplishments to reach their goal of getting recruited. Moreover, the society as a whole (in this case, the job fair society) mimicked real society through skills mismatch, a common phenomenon in a real-world job fair. All of these happened without humans pushing for these conclusions (Li et al. 2023). These experiments show that, with current technologies, the time is ripe to utilise social simulations in addressing key public policy issues, such as wicked problems.

3. Material and Methods

3.1 Wisdom of AI Crowds Framework

Drawing inspiration from Surowiecki’s (2004) “The wisdom of crowds”, where he theorised that collective intelligence made up of heterogeneous individuals grouped into a crowd can solve problems which individuals cannot, this paper proposes a new conceptual framework. Termed “Wisdom of AI crowds”, this framework aims to operationalise this concept within a simulated environment, using agentic AI to model

the collective behaviour of stakeholders to address wicked problems. The framework, as outlined in Figure 2, is divided into three phases, namely the Input phase, the Artificial Society phase, and the Human-in-the-loop feedback phase.

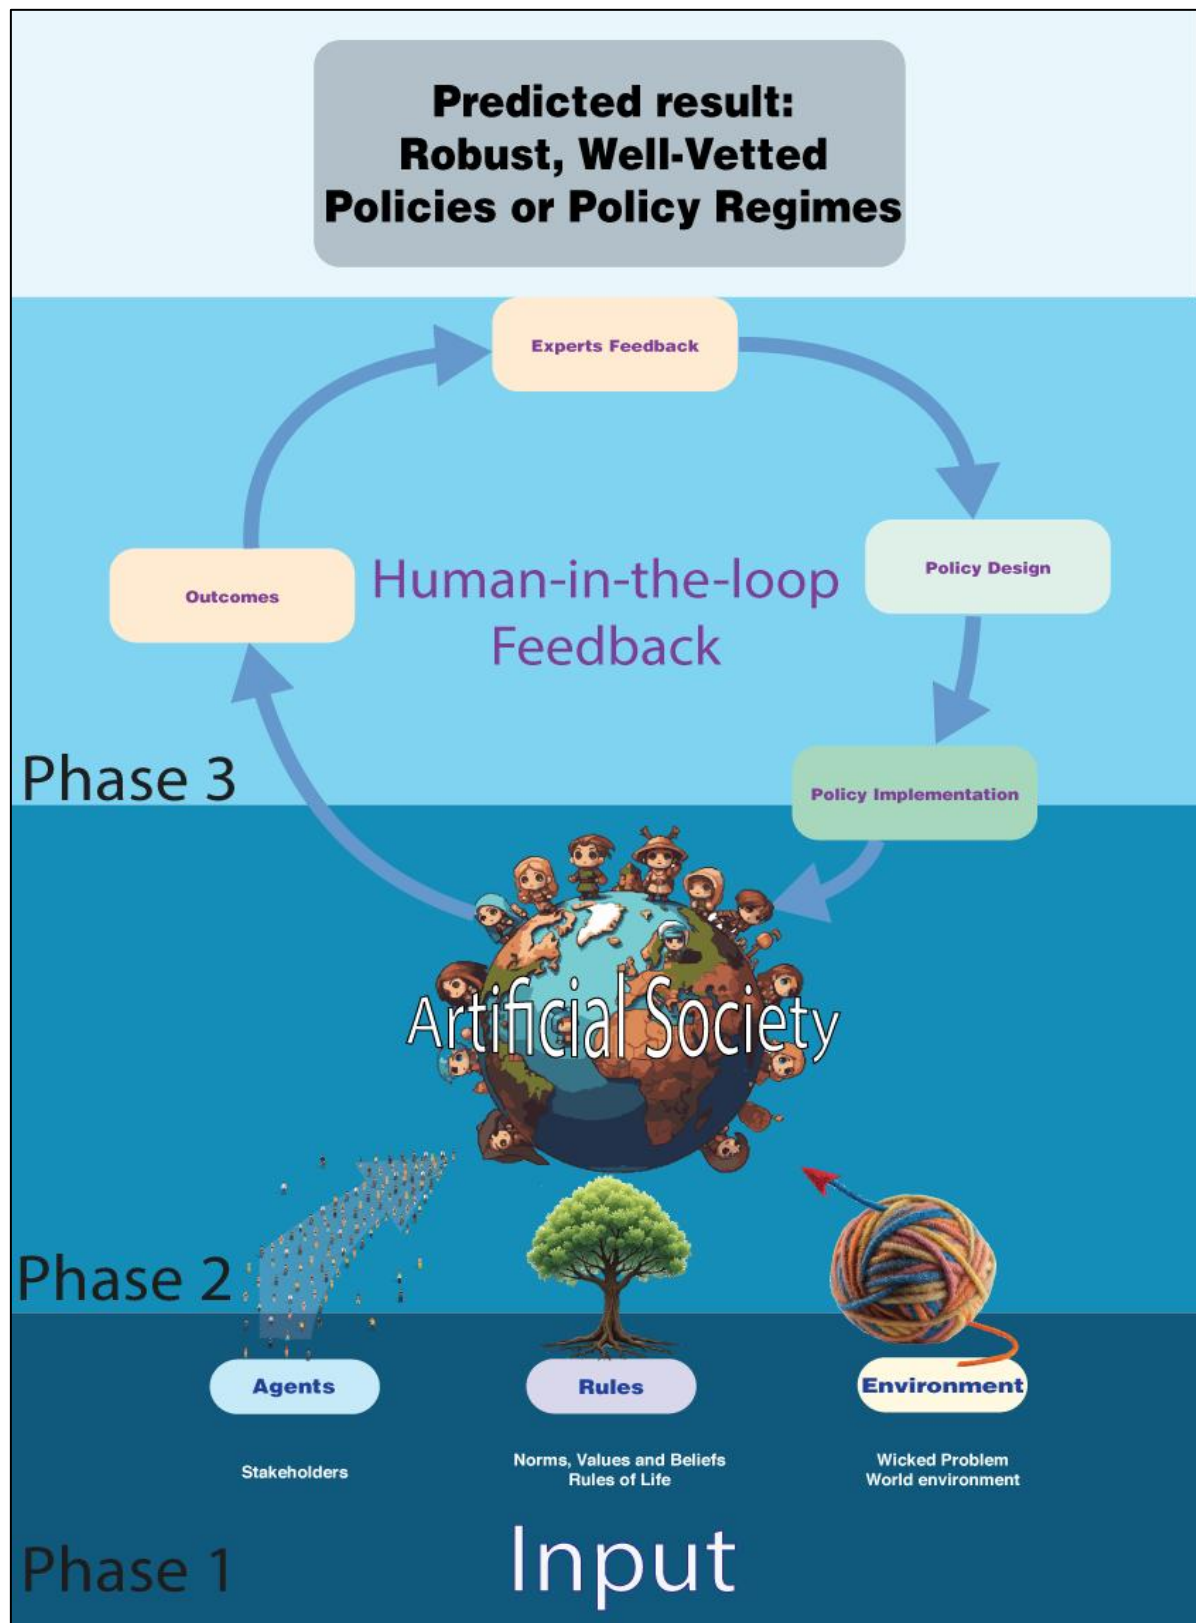


Figure 2: The Conceptual Framework

For better understanding, this framework will be split into a checklist which will then be described according to the phases.

Phase 1: Input and Setup

1. **Define Scope and Identify Stakeholders:** Clearly define the wicked problem's boundaries (e.g. Land Grab in Malta) and map out primary and secondary stakeholders.
2. **Gather Primary Stakeholder Data:** Conduct interviews and collect documentary evidence (laws, policies) related to primary stakeholders.
3. **Gather Secondary Stakeholder Data:** Design and deploy surveys to collect demographic, psychographic (norms, values, beliefs), and behavioural data for the broader public.
4. **Gather Environment and Rules Data:** Collect contextual documents (legislation, media reports, academic articles) defining the problem's environment and formal rules.
5. **Synthesise Inputs and Create Prompts:** Utilise an AI research assistant (e.g. NotebookLM) to process all gathered data and generate detailed prompts defining the agents (personas), the environment, and the rules for the artificial society.

Phase 2: Simulation and Experimentation

6. **Initialise Artificial Society:** Load the agent, environment, and rules prompts into the simulation platform.
7. **Conduct Policy Runs:** Introduce specific policy interventions (expert-derived or AI-generated) by modifying simulation parameters (rules/environment) and run the simulation for a defined period (simulating months/years), logging agent behaviour and outcomes.
8. **Iterate Scenarios:** Repeat step 7 for multiple policy variations and scenarios, including those suggested by the LLM itself.

Phase 3: Validation and Refinement

9. **Expert Review:** Analyse simulation logs and outcomes, then present findings (correlation between policies and promising outcomes) to human experts for validation regarding plausibility, realism, and policy feasibility.
10. **Iterate or Conclude:** Based on expert feedback, either refine the inputs (prompts, policies) and return to Step 6 for another simulation cycle, or conclude the process, documenting the most robust, well-vetted policy options identified.

3.2 Phase 1: The Input Phase

In the Input phase, information is gathered to fill the artificial society. Epstein and Axtell (1996) had outlined that an artificial society needs three elements to work: agents, environment, and rules.

The agents represent all stakeholders who are directly or indirectly involved with a specific wicked problem being tested. These can be broadly divided into two types,

based on their role in the policy process: the primary stakeholders, such as policymakers, political parties and pressure groups who directly shape policy; and the secondary stakeholders representing society at large, whose collective reactions and sentiments ultimately determine a policy's legitimacy and success. To build these primary stakeholder agents, information needs to be gathered from interviews and documentary evidence. In the case of the secondary stakeholder agents, the society's norms, values, and beliefs need to be gathered. To do this, demographic, psychographic, and behavioural data need to be collected through surveys. To mitigate bias, sampling for primary stakeholders will use purposive sampling to identify key institutional and civil society actors. For secondary stakeholders, a stratified sampling approach based on demographics will be used to ensure the survey data is representative.

The resultant agents are sophisticated human-like agents, with human-like attributes like memory, goals, simplified emotions, and common cognitive biases to make behaviour realistic. Similar to the METAAGENTS experiment, where agents lied to increase their chances of getting employed, these agents would mimic human behaviour to reach their goals. Agents would learn and adapt their behaviour over time based on previous success or failure. This allows for dynamic relationships, such as alliances and rivalries, to form between them.

For the environment to be constructed, comprehensive data must be gathered that defines the context of the wicked problem. To ensure a robust understanding, a triangulation of sources is employed. First, primary qualitative data is gathered from interviews with key experts, policymakers, and pressure groups. Second, documentary evidence such as legislation, policies, and agreements relevant to the problem is analysed. Third, existing research and media analysis, including journal articles and newspaper articles, are reviewed. Through this triangulation of sources, both scholarly findings and public discourse are reflected. To synthesise these diverse data sources into a coherent prompt that defines the simulation's environment, an AI-powered research assistant such as NotebookLM is utilised.

Once created, in the environment, there will be "places of action" where various activities, such as protests and negotiations, take place. These places include government buildings, community meetings, and streets in front of government buildings. The simulation would also model resource dynamics, such as money and land, and their socio-economic contexts.

Finally, the rules provide the governing logic for the simulation. Two types of rules are needed: the rules of society and the rules of interaction. For rules of society, these are explicit constraints that emanate from the Environment research. In contrast, rules of interaction are implicit constraints on how social rules govern the behaviour, arguments, and negotiations of agents. These rules would be taken directly from the norms, values and beliefs extracted from the agents' research. These societal values can be numerically represented within agent rules, for example, as weighted preferences or utility functions that guide agent decision-making based on the survey data. Most importantly, a key feature of LLM-powered agentic AI is that the society generates its own complex emergent social norms and behaviours. As demonstrated in Smallville and

METAAGENTS, agents develop their own unprompted social dynamics, which is vital for modelling a complex, wicked problem. The inputs and outputs of this phase are shown in Table 2.

Table 2: Inputs and Outputs of Phase 1

Component	Specific Inputs	Output
Agents	Interviews, Survey data from Norms, Values and Beliefs (NVB)	Defined Agent Personas
Environment	Interviews, Documents, Media Analysis (NotebookLM)	Simulation Environment Prompt
Rules	Documents (Laws), Agent NVB Data	Simulation Rules Prompt

3.3 Phase 2: The Artificial Society phase

In this second phase, the artificial society serves as the experimental core. Various policy regimes are introduced by manipulating the initial parameters, such as altering the ‘Rules of the Society’ to represent a new law, or adjusting the ‘Environment’, such as introducing a new tax. The simulation provides total control over time, allowing the researcher to observe long-term consequences, then pause, rewind, and rerun different scenarios with different parameters. Simulation outputs are measured through a combination of quantitative metrics (e.g. resource distribution shifts, agent satisfaction scores) and qualitative logs capturing agent interactions and emergent social dynamics. This capacity for limitless, risk-free iteration directly overcomes the ‘one-shot’ characteristic of real-world interventions. Moreover, this phase provides two types of policy design: testing policies emanating from interviews with human experts, and LLM-generated policy variations, which enable a broad exploration for a solution which can identify novel approaches. The inputs and outputs of Phase 2 are exposed in Table 3.

Table 3: Inputs and Outputs of Phase 2

Component	Specific Inputs	Output
Simulation core	Agent Personas Environment Prompt Rules Prompt Policy Regimes	Simulation data logs (recording agent actions, environmental changes over simulated time) Scenario Outcomes (aggregate results, emergent behaviours, policy impact metrics) Package for Phase 3 (structured outputs ready for expert review)

3.3 Phase 3: The Human-in-the-loop Feedback Phase

In the final phase, the simulation's outcomes are presented to human experts for validation. This human-in-the-loop process is necessary for assessing the plausibility of the results and for collaboratively refining the proposed policies. The outputs are assessed through a rubric such as agent realism (how they react to policies, how they react to different stakeholders etc), and policy outcomes (such as how plausible the outcomes of this policy are?). The rubric is being portrayed in Table 4.

Table 4: Rubric for experts

Criterion	Guiding Questions for Experts	Decision Gate
Agent Realism	Do the key agents react to policy changes in a realistic fashion, given their values and goals?	If any of the key agent behaviours are implausible, return to Phase 1 to refine agent prompts
Emergent Behaviour	Do the collective outcomes align with observed real-world social dynamics?	If emergent phenomena are highly implausible, review Phase 2 environment rules and agent interaction logic (rules)
Outcome Plausibility	Are the final simulated outcomes a plausible consequence of the tested policies within the wicked problem context?	If outcomes are unrealistic, review Phase 2 policy implementation parameters and external event triggers
Policy Feasibility	Is the most plausible policy considered politically, economically and socially feasible for implementation in the real world?	If the policy looks well-vetted and robust but unrealistic in the real world, revisit the policy parameters and rerun the simulation
Decision Point		If all criteria are met, the policy is documented as “robust and well-vetted”

As outlined in table 4, these improved policies are then re-inserted into the simulation for another round of testing. This iterative cycle continues, aiming to generate robust, well-vetted policies or policy regimes for the specific wicked problem. Within this framework, a 'successful' policy regime is defined not as optimal, but as one demonstrating positive trends in key metrics, achieving a plausible stakeholder equilibrium (potentially resembling a Nash Equilibrium), and receiving positive validation regarding realism and feasibility from the human experts. While the inputs and outputs of phase 3 are being presented in Table 5, a visualisation of the whole methodology is presented in Figure 3.

Table 3: Inputs and Outputs of Phase 3

Component	Specific Inputs	Output
Expert Validation	Data Package Human Experts	Validated/Critiqued Outcomes Feedback Refined Policy Regimes
Loop (back to Phase 2)	Refined Policy Regimes	Robust, Well-vetted policies or policy regimes

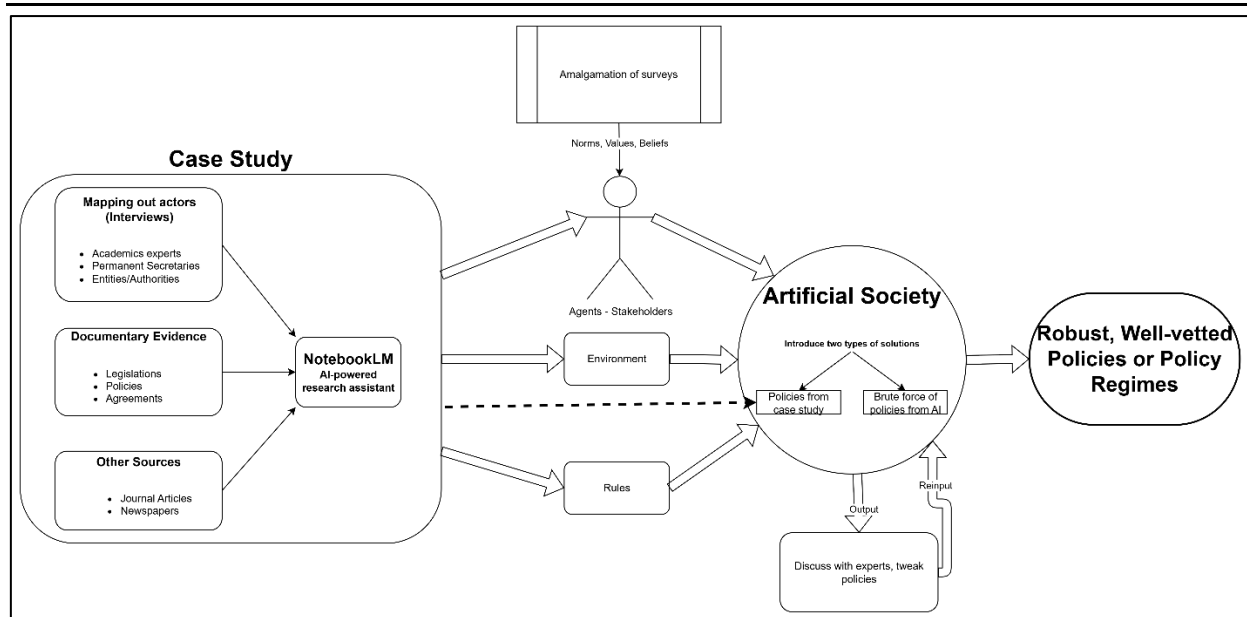


Figure 1: Methodology Diagram

To further illustrate this model, this paper includes an example of how this model might work on a specific wicked problem, namely the Land Grab in Malta.

3.4 Proof of Concept

For this paper, “Land Grab in Malta” refers to the significant change of use of large land driven by government actions. This encompasses both the acquisition of private land by the government, concessions of public land given to third parties and the utilisation of any public land for public or private use. In Malta, this is a wicked problem that meets both the criteria of a complex problem and the diverse stakeholder differences. There are massive pressure groups for and against land grab, and there is an everlasting debate on economy versus environment, and the long-term consequences are unclear.

If this social phenomenon was to be used for the simulation, phase 1 would mean that the stakeholders would be drawn up, such as in Table 6.

Table 6: Main Stakeholders in Land Grab

Who will be interviewed	Rationale for Inclusion
Academics	Experts in the Fields of Small Island States and Land Grab
Lands Authority	The Authority in Charge of Public Land Use
Ministry for Culture, Lands and Local Government	The Ministry which handles any land-related issue. Lands Authority is an authority which falls under the jurisdiction of this Ministry.
Moviment Graffiti	The most vocal pressure group against Public Land Use being utilised by the Private sector
MDA	The pressure group representing developers in Malta
FAA	A pressure group for a better environment – pressure for saving open spaces
PN	The Opposition Party in Parliament
MHRA	The Employers Association represents hotels and restaurants. Several restaurants are granted use of public land for commercialisation
Momentum	A political party which speaks out frequently against overdevelopment in Malta and the commercialisation of public land
ADPD	Another political party which speaks out frequently against overdevelopment in Malta and the commercialisation of public land
Residenti Beltin	A pressure group/political group based in the capital city (Valletta) made up of its residents who often speak out against the commercialisation of public land

Further to these, the policies and legal documents which directly relate to land grab are drawn up to be utilised in the research assistant phase. The Malta Tourism Authority policy on outdoor catering, Chapter 563 (Lands Authority Act) and 573 (Government Lands Act) of the Laws of Malta, are examples of documents which would need to be fed into the system (NotebookLM). Moreover, press releases issued by the diverse entities in Malta related to land grab, together with newspaper articles, are also gathered and fed into the research assistant. A number of datasets from surveys need to be gathered to build up the norms, values and beliefs of the secondary stakeholders. Namely, demographics, sociopolitical beliefs, and attitudes towards land use in Malta need to be utilised. NotebookLM, through a number of prompts, would write the prompts of the stakeholders, the environmental prompts together with their rules, which will be fed into the artificial society.

Once the simulation starts, prompts such as a change in policies or a change in law is introduced. Stakeholders such as Moviment Graffiti will be monitored to see if they will protest the new change, a typical activity which happens any time the group opposes laws relating to this issue. MDA will also be monitored on how it will react to such a change, typically issuing press releases or meeting with politicians. Other stakeholders are also monitored, including the secondary stakeholders who usually pick a side or another based on their values. The simulation will model a number of real-world years to see what the repercussions will be towards such changes. If need be, further changes may be introduced. Once one scenario is fully played out, further scenarios are tested. LLMs are also utilised to suggest further scenarios themselves and test them out.

Iterative Run

To test a policy (for example, a policy that restricts public land concessions to approval by 2/3rds of the House of Parliament), the model executes the following sequence. The entire process is repeated n times to generate a statistical distribution of possible futures.

1. Initialisation (Time $t = 0$)
 - The model loads the policy scenario
 - The environment is set: current land use, existing laws (Ch. 563, 573), public sentiment
 - The agent population is loaded. Primary stakeholders and secondary stakeholders are activated according to their core beliefs.
2. Simulation tick ($t + 1 \text{ month}$)
 - Policy and Environment update: A New law is introduced
 - Agent perceptions: Agents observe the change. The MDA agent sees a threat to developer interests; Moviment Graffiti agent sees potential victory.
 - Agent deliberation: Each agent's internal model processes the change (thinks). The Moviment Graffiti must decide whether the law is powerful enough or it is a smoke-screen. MDA weighs the cost of objecting versus compliance.
 - Agent action: Agents act. Moviment Graffiti decide whether to protest or not. MDA decides whether to issue a press release or hold meetings with the Government or the Opposition.
 - Environment feedback: The environment updates. Public sentiments shift, alliances may form or cease.
3. Logging
 - Key metrics for this tick are logged, such as: % of public land conceded, public sentiment score, and number of protests
 - The thoughts and decisions of each agent are logged
4. Termination
 - The simulation runs for a set period (e.g. 60 ticks = 5 years). The final state and full log are saved for analysis.

Box 1: Policy testing in the simulation

Once several policies and their effects are simulated, and the most promising outcomes are correlated to their respective policies, these results are discussed with the same academics and experts in the field. The experts are asked to evaluate the realism of the results and the plausibility of the agents' behaviour in response to the policies. If adjustments are needed, the initial prompts or policies are tweaked according to the feedback, and the simulation is rerun. This iterative loop continues until the process yields robust, well-vetted policy options that represent the most effective approaches identified in the simulation. The logic of the simulation run is being outlined in Box 1. However, successfully implementing this framework requires addressing several key methodological and ethical challenges.

3.5 Challenges

The challenges which require cognisance of and addressing for the experiment to be successful are data-related, AI reliability, realism, legitimacy, and ethical challenges as outlined in **Error! Reference source not found.**

Table 7: Challenges/Risks and Mitigation measures

Challenge/Risk	Why it matters	Mitigation in Methodology
Data quality and representativeness of input data	Simulation validity requires mitigating data scarcity/bias . Underrepresentation due to the "digital divide" can skew results if not countered (Norori et al. 2021).	Input Phase uses triangulation of diverse sources (interviews, documents, surveys) to create a more representative dataset, ensuring agents are faithful proxies.
AI Reliability: Risk of unreliable agent behaviour due to LLM limitations	Data scarcity produces inaccurate results (Alzubaidi et al. 2023). This can also create hallucinations, where agents would generate non-sensical or incorrect information (Leiser, Eckhardt et al. 2024).	In-context learning: Grounding agents in specific data from Input Phase. Human-in-the-loop validation: Experts review outcomes for plausibility and realism, also addressing causality concerns.
Realism: Ensuring realistic simulation outcomes from non-human agents	AI agents are not human and do not possess genuine consciousness or emotions.	Goal is 'believable proxies,' not perfect replication. Achieved by grounding agents in empirical norms, values, and beliefs from the Input Phase, constraining behaviour with plausible human motivations.
Legitimacy & Ethics: Ensuring trust and ethical soundness.	Decisions solely by AI lack public trust (Starke and Lünich 2020); unethical design can cause harm or unfairness (Askell et al., 2021). Adherence to data protection regulations (e.g. GDPR).	Human-in-the-loop ensures human accountability. Framework designed for fairness (representative data), explicability (transparent phases), and non-maleficence (expert validation). Uses anonymised secondary data adhering to GDPR, obtains informed consent for primary data collection. Real-world application would require clear governance protocols defining who controls the simulation parameters, interprets results, and ensures alignment with democratic values.

In the case of this framework, these challenges are directly addressed through its design, making it a robust and ethically considered tool for policy simulation. Nevertheless, beyond the mitigation measures listed in Table 7, the framework adheres to a strict ethical plan. All primary data collection requires explicit, informed consent for use in simulation-building, with participants aware of the study's purpose. For secondary stakeholders, the framework uses pre-existing, public datasets (e.g. Eurobarometer, Household Budgetary Survey). This data is obtained in a fully anonymised and aggregated format from the respective statistical authorities, who are themselves bound by strict data protection protocols. The ethical obligation within this framework is to ensure this data is used

responsibly, guaranteeing that the creation of agent personas does not allow for the re-identification of any individuals, in full compliance with GDPR principles.

4. Discussion

Should this framework be successful, it would create a new way of addressing policies for challenging problems. A shift from a reactive towards a proactive approach would be possible, allowing policymakers to predict outcomes and test more radical ideas safely before implementation. However, the societal acceptance of such a tool remains an open question. While it offers the potential for more effective solutions, the risk of over-reliance on such a tool, a concern currently being debated regarding LLMs in broader society, must be considered.

4.1 Recommendations

Future research may study individual phases and enhance them. For example, a study can focus on whether expert validation can be enhanced. Another example would be whether certain LLM models would fare better in the simulation than other models.

5. Conclusion

This paper has argued that the impossibility of addressing wicked problems emanates from the inability to experiment with solutions safely. Any solution attempted in the real world changes the problem, thus making any attempt a final one-shot solution. Although current solutions were heading in the right direction through experimentation, they lacked the technological ability to create a safe, isolated environment, a sandbox, that could separate the problem from real-world impacts. Any of their attempts focused on a single snapshot of the situation, were static, and high-risk. The framework 'Wisdom of AI Crowds' is a novel methodological blueprint that seeks to overcome this barrier. It does this through a three-phase process comprising input, simulation, and feedback, thus creating a dynamic approach. This novel approach is the main contribution to the literature, where this framework will create a virtual laboratory for testing policy on complex societal issues.

The framework, however, is still conceptual, and its success depends on the quality of the input data and expert validation. The next step, which the authors aim to perform, is to actually apply this framework using two case studies, an insular wicked problem: Land Grab in Malta and a global wicked problem: Climate Change. Through the success of these experiments, the framework would be validated for use in policymaking.

Acknowledgment

The authors would like to thank Prof. Mario Thomas Vassallo for his invaluable guidance, supervision, and critical feedback throughout the development of this research.

Conflict of Interest Statement

The author declares no conflicts of interest.

About the Author(s)

Gian Paul Gauci is a PhD Candidate at the Department of Policy, Politics and Governance, University of Malta. His research interests include wicked problems, public policy, artificial intelligence, and policy simulations.

ORCID: <https://orcid.org/0009-0009-0595-1036>,

ResearchGate: <https://www.researchgate.net/profile/Gian-Gauci>

Prof. Alexiei Dingli is a Professor at the Department of AI, Faculty of ICT, University of Malta and an expert in Artificial Intelligence. His research interests include AI applications, digital innovation, and the governance of emerging technologies.

ORCID: <https://orcid.org/0000-0002-5951-8299>,

ResearchGate: <https://www.researchgate.net/profile/Alexiei-Dingli>

References

- Alford, J. and Head, B.W. 2017. Wicked and less wicked problems: A typology and a contingency framework. *Policy and Society* 36(3), pp. 397–413. doi: <https://doi.org/10.1080/14494035.2017.1361634>.
- Alzubaidi, L. et al. 2023. A survey on deep learning tools dealing with data scarcity: definitions, challenges, solutions, tips, and applications. *Journal of Big Data* 10(1), p. 46. Available at: <https://doi.org/10.1186/s40537-023-00727-2>.
- Askell, A. et al. 2021. *A General Language Assistant as a Laboratory for Alignment*. Ithaca: Cornell University Library, arXiv.org. Available at: <https://search.proquest.com/docview/2605773955>.
- Bonabeau, E. 2002. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences* 99, pp. 7280–7287. Available at: <https://doi.org/10.1073/pnas.082080899>.
- Branke, J. 2010. Artificial Societies. In: Sammut, C. and Webb, G. I. eds. *Encyclopedia of Machine Learning*. Boston, MA: Springer US, pp. 44–48. Available at: https://doi.org/10.1007/978-0-387-30164-8_36.
- van Bueren, E.M. 2003. Dealing with Wicked Problems in Networks: Analyzing an Environmental Debate from a Network Perspective. *Journal of Public Administration Research and Theory* 13(2), pp. 193–212. Available at: <http://jpart.oupjournals.org/cgi/doi/10.1093/jopart/mug017>.
- Conklin, J. 2005. *Dialogue mapping: Building shared understanding of wicked problems*. John Wiley & Sons, Inc. Available at: <https://www.wiley.com/en-us/Dialogue+Mapping%3A+Building+Shared+Understanding+of+Wicked+Problems-p-9780470017685>
- Duncan, N. and Paradice, D. [no date]. *Creativity in GDSS An Integration of Creativity, Group Dynamics, and Problem-Solving Theory*.

- Elia, G. and Margherita, A. 2018. Can we solve wicked problems? A conceptual framework and a collective intelligence system to support problem analysis and solution design for complex social issues. *Technological Forecasting and Social Change* 133, pp. 279–286. <https://doi.org/10.1016/j.techfore.2018.03.010>
- Epstein, J.M. and Axtell, R. 1996. *Growing artificial societies: social science from the bottom up*. Brookings Institution Press. <https://doi.org/10.7551/mitpress/3374.001.0001>
- Fischer, F. 1993. Citizen Participation and the Democratization of Policy Expertise: From Theoretical Inquiry to Practical Cases. *Policy Sciences* 26(3), pp. 165–187. <https://doi.org/10.1007/BF00999715>
- Grint, K. 2010. The cuckoo clock syndrome: Addicted to command, allergic to leadership. *European Management Journal* 28(4), pp. 306–313. <https://doi.org/10.1016/j.emj.2010.05.002>
- Horn, R.E. and Weber, R.P. 2007. New tools for resolving wicked problems: Mess mapping and resolution mapping processes. *Watertown, MA: Strategy Kinetics LLC*. Available at: https://complexitylabs.io/New_Tools_For_Resolving_Wicked_Problems.pdf
- Levin, K., Cashore, B., Bernstein, S. and Auld, G. 2012. Overcoming the tragedy of super wicked problems: Constraining our future selves to ameliorate global climate change. *Policy Sciences* 45(2), pp. 123–152. <https://doi.org/10.1007/s11077-012-9151-0>
- Li, Y., Zhang, Y. and Sun, L. 2023. MetaAgents: Simulating Interactions of Human Behaviors for LLM-based Task-oriented Coordination via Collaborative Generative Agents. Available at: <http://arxiv.org/abs/2310.06500>.
- Lindblom, C.E. 1959. Science of Muddling Through. *Public Administration Review* 19.
- Norori, N., Hu, Q., Aellen, F.M., Faraci, F.D. and Tzovara, A. 2021. Addressing bias in big data and AI for health care: A call for open science. *Patterns (New York, N.Y.)* 2(10), p. 100347. Available at: <https://dx.doi.org/10.1016/j.patter.2021.100347>.
- Park, J.S., O'Brien, J., Cai, C.J., Morris, M.R., Liang, P. and Bernstein, M.S. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In: *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. ACM Conferences. New York, NY, USA: ACM, pp. 1–22. Available at: <https://dl.acm.org/doi/10.1145/3586183.3606763>.
- Ritchey, T. 2013. Wicked problems: Modelling Social Messes with Morphological Analysis. *Acta morphologica generalis* 2(1), pp. 1–8. Available at: https://www.researchgate.net/publication/236885171_Wicked_Problems_Modelling_Social_Messes_with_Morphological_Analysis
- Rittel, W.J.H. and Webber, M.M. 1973. Dilemmas in a General Theory of Planning. *Policy Sciences* 4(2), pp. 155–169. doi: 10.1007/BF01405730. Available at: https://urbanpolicy.net/wp-content/uploads/2015/06/Rittel-Webber_1973_DilemmasInAGeneralTheoryOfPlanning.pdf
- Roberts, N. 2000. Wicked Problems and Network Approaches to Resolution. *International Public Management Review* 1(1), pp. 1–19. Available at: <https://ipmr.net/index.php/ipmr/article/view/175>.

- Russell, S.J. and Norvig, P. 2020. *Artificial intelligence: a modern approach*. 4th edition. Boston: Pearson Boston. Available at: <https://worldcat.org/title/1021874142>.
- Simon, H.A. 1957. *Models of man: social and rational; mathematical essays on rational human behavior in a social setting*. New York: Wiley.
- Starke, C. and Lünich, M. 2020. Artificial intelligence for political decision-making in the European Union: Effects on citizens' perceptions of input, throughput, and output legitimacy. *Data and Policy* 2(2). Available at: https://books.google.ro/books/about/Models_of_Man_Social_and_Rational.html?id=4iZNxAEACAAJ&redir_esc=y
- Surowiecki, J. 2004. *The wisdom of crowds: why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations*. New York: Doubleday. Available at: <https://www.asecib.ase.ro/mps/TheWisdomOfCrowds-JamesSurowiecki.pdf>
- Zhao, P., Jin, Z. and Cheng, N. 2023. *An In-depth Survey of Large Language Model-based Artificial Intelligence Agents*. Ithaca: Cornell University Library, arXiv.org. Available at: <https://search.proquest.com/docview/2869396437>.
- .