

The Simulation of Democracy

A Multi-Agent AI Experiment in Hungarian Electoral Prediction

Matteo Mio

April 11th 2026

Matteo Mio | Forward College / University of London

Abstract

What if, instead of polling voters, you simulated them? This paper does exactly that - running a full multi-agent AI simulation of Hungary's 2026 parliamentary election through MiroFish, a platform that models political opinion as a live emergent process rather than a frozen snapshot. Forty-five AI agents - each embodying a distinct voter archetype with their own beliefs, media diet, material conditions, and strategic calculus - interact over 100 rounds and 1,965 discrete actions, reasoning, persuading, updating, and ultimately deciding.

The technical stack is assembled locally: CAMEL-AI and OASIS power the agent layer, Gemini 2.5 Flash Lite handles every reasoning call through OpenRouter, Neo4j AuraDB maintains the shared knowledge graph, and BrainInTheFish audits the pipeline at three defined checkpoints. The raw simulation output does not go directly to a seat count. It passes through eleven structured bias corrections under a two-layer epistemic recalibration framework before a single vote share is accepted as valid.

Constituency-level seat modelling is then handled by Chronicler-v2, calibrated at $\lambda = 0.58$ - the strategic transfer rate from DK and MKKP voters to TISZA in competitive single-member districts.

The results come in and predict a tight majority for TISZA. The predicted national shares (Fidesz 41.0%, TISZA 36.7%, DK 7.0%, MKKP 6.7%, Mi Hazánk 3.9%) run through Hungary's mixed electoral system - 106 FPTP constituencies, 93 D'Hondt list seats, full winner and loser compensation - to yield a final 199-seat parliament: TISZA 101, Fidesz 84, DK 7, MKKP 5, German minority seat 1, Independent 1. TISZA wins a governing majority by a single seat.

The margin is razor-thin, and deliberately so. Sensitivity analysis shows that a λ shift from 0.58 to 0.55 - three percentage points of strategic voter defection - collapses the majority into a hung parliament. That single parameter carries more weight than any individual party's vote share. The paper closes with a frank researcher's perspective on what the method gets right and where uncertainty remains, alongside a complete resource audit: 143 million tokens, \$7.95 in API costs, 48 hours of laptop compute, and a Neo4j cloud instance running throughout.

About the Author

Matteo Mio is an Italian student, raised between Italy and Switzerland, currently enrolled in a Bachelor's in Economics and Politics at the LSE through Forward College's international program. Passionate about innovation and social progress, he manifests his curiosities through reading, writing and debating, touching on topics such as the startup ecosystem or European technologic industries as well as general public policy and AI regulation.

Through his academic journey, he has studied in Lisbon, is currently based in Paris and is projected to finish his bachelor in Berlin. He is building Scalia Studio, a freelancing group of young students passionate about building future technologies through AI and automation. Their work tackles political and cultural challenges as well as business optimisation and AI integration in corporate operations. He loves to discuss these topics with other professionals and interested individuals at events such as Web Summit, which he has attended in the past two editions and hosted a booth at.

When he is not studying or working, he can be found at the gym or on the football field, but also at any random jazz bar in Paris. He also enjoys having politically challenging conversations with friends (or strangers), and travelling to countries of which he does not speak the language.

Contact

Email: mattemio9@gmail.com

[LinkedIn](#)

Table of Contents

Section 1 – Introduction	1
Section 2 – The Setup	2
2.1 Architecture Overview	2
2.2 The Components	2
2.3 How They Work Together	3
2.4 The Seed Documents	3
2.5 The Research Question	4
Section 3 – How the Simulation Works	5
3.1 Graph Build	5
3.2 Agent Instantiation	5
3.3 The Simulation Run	6
3.4 The Report	6
3.5 Deep Interaction	6
Section 4 – Extracting Information Through Deep Interaction	7
4.1 The Problem With the Report	7
4.2 Vote Share Extraction	7
4.3 The Corrections Applied	8
4.4 Turnout and Demographic Weighting	8
4.5 SMD Surveys	8
4.6 What Deep Interaction Added	9
Section 5 – Emergent Social Phenomena	9
5.1 The Silenced Centre	9
5.2 The Rural Echo Chamber	10
5.3 The Ghost Persuader	10
5.4 The Shy Voter Asymmetry	10
5.5 The Urban Surge	11
5.6 The Diaspora Effect	11
5.7 Cross-Platform Divergence	11
Section 6 – Post-Hoc Epistemic Recalibration	12
6.1 Why This Was Necessary	12
6.2 Two Layers of Bias	12
6.3 The Issue Register	13
6.4 Correction Batches – Vote Shares	14
6.5 SMD Correction: From Agent Survey to Chronicler-v2	14
6.6 The λ Calibration	15
Section 7 – Final Results	16
7.1 National Vote Shares	16
7.2 SMD Seat Distribution	17

7.3 List Seat Distribution	17
7.4 Final 199-Seat Parliament	18
7.5 What This Means	18
Section 8 – Scenarios and Sensitivity	19
8.1 The Variable That Decides Everything: λ	19
8.2 What Moves λ	20
8.3 Vote Share Sensitivity	20
8.4 Threshold Risk	21
8.5 The Hung Parliament Scenario	21
Section 9 – Researcher’s Perspective	22
9.1 On the Results	22
9.2 On the Method	22
Section 10 – Resource Footprint	23
10.1 API Usage and Cost	23
10.2 Energy and Carbon Footprint	23
10.3 Water Footprint	24
10.4 Context	24
Appendices	A-1
Appendix A – Full Agent Roster	A-1
Appendix B – Agent Activity Configuration	A-2
Appendix C – DOC6 Archetype Mapping	A-3
Appendix D – Final Results Tables	A-3
Notices	A-4

Section 1 - Introduction

Artificial intelligence is getting remarkably good at simulating how people think. This experiment is one early attempt to put that capability to a concrete test.

Polling is useful. It tells you where people stand at a given moment. What it cannot tell you is how they got there - and more importantly, where they are going.

It does not capture how a rural pensioner's vote intention shifts after weeks of state television. It does not model the strategic calculation of a DK supporter in a marginal Budapest constituency who knows their candidate cannot win. It does not account for the social pressure within a small-town community, or the way economic grievance interacts with media exposure to produce a final decision on election day. Polling captures a snapshot. It misses the process.

This is the gap that predictive agentic AI is designed to fill. A recent open source agentic platform could be the first truly interesting and accessible tool for use cases of these technologies. MiroFish is a multi-agent simulation platform that models political opinion formation as an emergent social process. Rather than aggregating stated preferences, it instantiates a population of AI-powered agents - each representing a distinct voter archetype, equipped with their own information diet, material conditions, and political priors - and lets them interact, influence each other, reason strategically, and arrive at vote intentions through simulated social dynamics. The output is not a poll. It is a model of how an electorate thinks its way to an outcome.

After testing the platform on simpler processes, such as predicting World Cup results, it was time to put this technology to the test on one of its most interesting applications: political modeling and prediction.

Hungary's April 2026 parliamentary election is the most competitive the country has seen since 2010 - a genuine contest between Fidesz and the emerging TISZA party of Péter Magyar, playing out on a structurally asymmetric field of gerrymandered constituencies, state media dominance, and a fragmented opposition navigating strategic voting decisions in real time. The

result carries weight beyond Hungary's borders: a shift in power or loss of supermajority could reshape the EU's internal balance at a critical moment. A consolidation of Orbán's grip, on the other hand, would push the sovereigntist, illiberal wave sweeping European politics even further.

The question posed to the simulation is:

"In the context of the 2026 Hungarian Parliamentary elections, and under realistic information asymmetry conditions, how do different voter segments form

their final voting decision - and what aggregate outcome does that produce?"

Polling data is explicitly excluded as the answer. The outcome must emerge entirely from agent dynamics - from the social process the simulation is designed to model.

What follows is a full account of that process: the setup, what the agents produce, how the output is interrogated and corrected, and what the final prediction looks like.

Section 2 - The Setup

2.1 Architecture Overview

The system used in this experiment is a locally hosted integration of several interconnected tools, assembled and configured specifically for this prediction exercise. At its centre sits MiroFish - a multi-agent social simulation platform built on the CAMEL-AI and OASIS frameworks - running as a local application and connecting outward to a cloud graph database, a commercial LLM API, and a post-hoc evaluation tool. No single component produces the result. They work as a pipeline, each handling a distinct layer of the process.

2.2 The Components

MiroFish MiroFish is the simulation engine. It instantiates AI-powered agents, defines their personas, manages their interactions across a dual-platform social environment, and aggregates their behaviour into structured outputs. The platform is built on CAMEL-AI and its OASIS social simulation extension - a framework specifically designed for modelling networked agent behaviour over time. In this experiment, 45 agents were run across 100 rounds, producing a total of 1,965 discrete social interactions across the two platforms. Each agent has a configured activity level, posting frequency, active hours, influence weight, and sentiment orientation - parameters that shape not just what they believe, but how loudly and how often they broadcast it.

Neo4j AuraDB The knowledge graph layer is handled by Neo4j, a graph database running as a cloud instance. As the simulation runs, MiroFish builds and updates a structured graph of entities, relationships, and beliefs - topics agents discuss, positions they hold, connections between them. This graph is the simulation's memory. It allows agent reasoning to be grounded in a persistent, queryable representation of the political and social landscape rather than in isolated, context-free LLM calls. The graph is also exportable as a semantic knowledge structure, which feeds into the evaluation layer.

OpenRouter + Gemini 2.5 Flash Lite Every agent reasoning step - every belief update, every social post, every response to a peer - is powered by a large language model call routed through OpenRouter, using Google's Gemini 2.5 Flash Lite as the underlying model. OpenRouter acts as the API layer, handling routing and cost management. Gemini 2.5 Flash Lite was chosen for its balance of reasoning capability and cost efficiency given the volume of calls a 45-agent,

100-round simulation generates. A separate embedding model - Nomic Embed Text - handles semantic similarity operations within the platform.

BrainInTheFish BrainInTheFish is a post-hoc evaluation tool that integrates with MiroFish at three points in the pipeline. It snapshots the Neo4j knowledge graph as a semantic RDF structure after the graph build phase, runs a structural credibility check on the facts extracted by MiroFish's reporting layer, and performs a full audit of the final simulation report. In this experiment, BrainInTheFish's evaluation framework and expert panel scoring rubric directly informed the epistemic recalibration process - providing the structured methodology through which simulation biases were identified, classified by tier, and corrected. Its integration is non-fatal: if any step fails, the simulation continues. It is a quality layer, not a dependency.

2.3 How They Work Together

The pipeline runs as follows. MiroFish is launched locally with the seed documents and agent configuration loaded. Neo4j is populated with the knowledge graph as agents initialise. OpenRouter LLM calls power each agent's reasoning at every simulation step. BrainInTheFish hooks fire at defined points - after graph build, during report generation, and after the final report is written - attaching credibility scores and evaluation outputs to the result. The final deliverable is a simulation report, an evaluated knowledge graph, and a structured set of agent outputs available for further interrogation through the Deep Interaction system.

2.4 The Seed Documents

The simulation was not run blind. Before agents were instantiated, six structured documents were provided as the knowledge base from which agent priors, information diets, and reasoning frameworks were built.

Document 1 - Electoral System Mechanics A technical document covering the full architecture of the Hungarian electoral system: the 106 single-member district seats decided by

first-past-the-post, the 93 national list seats allocated via D'Hondt, the winner and loser compensation mechanism, the 5% list threshold, and the strategic voting implications of each. This document also contains the Chronicler-v2 methodology for constituency-level seat modelling, used in the post-hoc SMD correction.

Document 2 - Political Actors, Platforms and Endorsements Profiles of the main competing parties - Fidesz-KDNP, TISZA, DK, MKKP, and Mi Hazánk - covering their platforms, voter bases, key figures, and endorsement networks. This seeded agent perceptions of party credibility and legitimacy.

Document 3 - Media and Information Environment A mapping of Hungary's media landscape in 2026: the reach and editorial orientation of state-aligned outlets, independence levels of commercial and online media, and the information asymmetry between rural and urban audiences. This determined each archetype's primary media sources and the trust weighting

applied to political signals received through them.

Document 4 - Economic and Social Grievances A structured overview of the material conditions facing Hungarian voters: inflation persistence, real wage dynamics, housing affordability, rural infrastructure gaps, minority community conditions, and diaspora economic ties. This grounded agent reasoning in economic self-interest alongside ideological alignment.

Document 5 - Polling Context A compilation of polling data from both state-adjacent and independent research firms, covering party support levels, demographic breakdowns, and trend lines across 2024-2026. Provided as context only - agents were explicitly instructed not to reproduce polling figures but to arrive at their own conclusions through the simulation.

Document 6 - Simulation Design and Agent Specifications The operational blueprint: the 10 agent archetypes and their full specifications, the 5-round structure and what each round was designed to produce, the initial belief states assigned to each archetype, and the strategic voting logic agents were required to apply.

2.5 The Research Question

With the architecture in place and the knowledge base loaded, the simulation was initiated with the following research question:

"In the context of the 2026 Hungarian Parliamentary elections, and under realistic information asymmetry conditions, how do different voter segments form their final voting decision - and what aggregate outcome does that produce?"

The system was further instructed to model this as a fully emergent process: agents were required to reason from their archetype-specific conditions, update through social interaction, apply strategic voting and threshold reasoning consistent with the Hungarian electoral system, and produce a final output covering national vote percentages, full seat allocation, majority type,

and hung parliament determination. Polling data was explicitly excluded as a permissible answer.

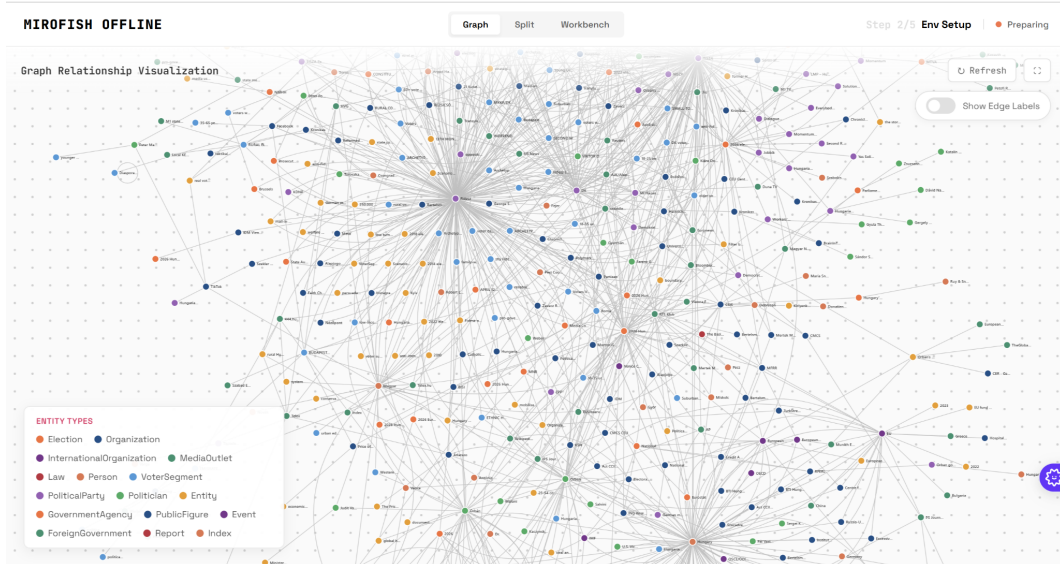
Section 3 - How the Simulation Works

Running a MiroFish simulation is not a single operation. It is a sequence of distinct phases, each building on the last - and understanding what each phase does matters for interpreting what comes out at the end.

3.1 Graph Build

Everything starts with the knowledge graph. Before any agent opens their mouth, MiroFish processes the seed documents and constructs a structured representation of the political and

social landscape in Neo4j. Entities are extracted - parties, politicians, policies, media outlets, economic indicators, demographic groups - and the relationships between them are mapped. This graph becomes the shared world that agents inhabit. It is what makes their reasoning contextually grounded rather than generic: when an agent reasons about Mi Hazánk's electoral viability, it is drawing on a structured graph of that party's relationships, positions, and context - not just a language model's training data.



The Graph Build once completed.

3.2 Agent Instantiation

Once the graph is built, the 45 agents are instantiated. Each agent is constructed from a specific archetype defined in Document 6 - rural elderly voter, young urban first-timer, diaspora mail-in voter, and so on across the full demographic spectrum of the Hungarian electorate. The archetype determines everything: what media sources the agent trusts, what economic pressures they feel, what their initial political prior looks like, how active they are, how much influence they carry in the network, and what hours of the day they are most engaged. Two agents from the same archetype will not behave identically - but they will reason from the same starting conditions.

This is also the phase where the simulation's structural choices have the most consequence, as will become clear in the recalibration section.

3.3 The Simulation Run

With agents live and the graph populated, the simulation runs. Over 100 rounds - representing a compressed but structured social timeline - agents interact across a dual-platform environment. They post content, respond to each other, consume information filtered through their archetype-specific media diet, and update their beliefs accordingly. In total, the simulation

produced 1,965 discrete agent actions across the two platforms.

The belief-updating process is the core of the experiment. Each interaction is a potential influence event: an agent's post can shift a peer's position, reinforce an existing view, or introduce a new consideration. High-influence agents carry more weight in this process. Agents with low activity levels and zero posting rates - the undecided segment, as it turned out - absorb influence without propagating their own deliberation back into the network. These asymmetries, invisible during the run itself, surface as meaningful findings later.

3.4 The Report

At the end of the simulation run, MiroFish's reporting layer - InsightForge - processes the accumulated agent behaviour and knowledge graph state into a structured output document. This report aggregates agent vote intentions, summarises the social dynamics that emerged, and attempts to produce the electoral prediction the research question asked for: vote percentages, seat allocation, majority determination.

The report is the simulation's official answer. It is also, as this experiment demonstrates, only the starting point.

3.5 Deep Interaction

The report captures what the simulation concluded. Deep Interaction is the tool for asking the simulation to explain itself.

Through MiroFish's IPC - Inter-Process Communication - command system, individual agents can be surveyed directly after the simulation completes. Structured questions are sent to each agent, and responses are logged as JSON files for extraction and analysis. In this experiment, Deep Interaction was used extensively: to extract precise vote share estimates from each demographic segment, to probe turnout assumptions, to correct artifact responses, and - critically - to conduct three rounds of constituency-level SMD surveys across all 45 agents.

If the simulation run is the experiment, Deep Interaction is the debrief. It is what transforms a high-level report into a granular, segment-by-segment dataset that can be weighted, corrected, and fed into the electoral mechanics model described in the sections that follow.

Section 4 - Extracting Information Through Deep Interaction

SECTIONS 4/4 ELAPSED 11m 37s TOOLS 30 COMPLETED

- PL Planning / Outline
- 01 Fragmented Parliamentary Outcome and Shifting Voter Alliances
- 02 Agent Archetype Reactions to Electoral Mechanics and Information Asymmetry
- 03 Emergent Trends: Strategic Voting and the Influence of Electoral System Thresholds
- 04 Potential Future Risks: Political Instability and the Erosion of Informed Decision-Making
- 0K Complete

[Enter Deep Interaction →](#)

Entering the Deep Interaction Portal.

4.1 The Problem With the Report

The simulation report produced by InsightForge at the end of the 100-round run contained a national-level prediction. It gave vote percentages, a seat projection, and a majority determination. But it did so at a level of aggregation that made it difficult to evaluate. Which segments were driving the Fidesz vote? Where was TISZA's support concentrated? What were the turnout assumptions? The report answered the research question in broad strokes. It did not show its working.

This is where Deep Interaction became essential. The goal was to go back into the simulation - after the run had completed - and extract from each of the 45 agents the specific data needed to build a granular, segment-by-segment electoral model from scratch.

4.2 Vote Share Extraction

The first Deep Interaction survey targeted vote intentions. All 45 agents were asked, on both platforms, to state their segment's estimated vote share for each party. The format was structured: percentage allocations for Fidesz, TISZA, DK, MKKP, Mi Hazánk, and Others, summing to 100.

The responses were uneven. Some agents complied cleanly, providing formatted percentage breakdowns. Others deflected - describing their persona rather than answering the question. Several gave politically implausible responses: a Roma segment agent allocating 15% to Mi Hazánk, a party whose platform is explicitly anti-Roma. A young urban first-time voter agent on one platform returned Fidesz at 75%, contradicting its own output on the other platform.

This prompted a second round, targeting the 29 agents that had deflected or produced incomplete data. The prompt was more forceful, specifying the exact output format and rejecting persona-description responses. Compliance improved. A third round corrected the ten most problematic artifact responses individually, and a fourth round addressed the Roma-specific Mi Hazánk anomalies.

By the end of the four rounds, a complete dataset existed: vote share estimates across all 13 demographic segments, corrected for artifacts, with both platforms represented where both produced valid data.

4.3 The Corrections Applied

The artifact corrections during this extraction phase were not arbitrary adjustments. Each followed a specific logic:

- Platform inversion artifacts - where an agent gave contradictory estimates across the two platforms (e.g. Agent 34, young urban first-time voter, returning Fidesz 75% on one platform and TISZA 70% on the other), the platform consistent with the archetype profile was retained and the other discarded.
- Political implausibility corrections - where an agent produced a result that contradicted basic political reality (e.g. Roma voters supporting Mi Hazánk, or an undecided Fidesz-leaning voter allocating 25% to DK and 20% to Mi Hazánk), the response was corrected to the nearest defensible estimate based on the archetype's documented characteristics.
- Rural-urban inversion - Agent 10, a rural mid-age conservative agent, returned TISZA 45% / Fidesz 35%. This was manually corrected to Fidesz 60% / TISZA 22.5% during extraction on the grounds that a TISZA-leading rural agent was implausible. As will become clear in the recalibration section, this correction turned out to be the most consequential editorial intervention in the dataset - and arguably the least justified.

Each correction was logged. The raw and corrected values were preserved side by side. This transparency became critical later, when the entire correction history was itself audited for bias.

4.4 Turnout and Demographic Weighting

Alongside vote shares, agents were surveyed for turnout estimates. Combined with externally sourced population weights for each demographic segment, this produced the weighting model used to convert segment-level vote shares into a national result. Each segment's effective weight was calculated as population share multiplied by turnout rate, then normalised to 100%.

The weighting table that emerged from this process became one of the canonical inputs to the final prediction - and also, as the recalibration section documents, one of the sources of structural bias identified later.

4.5 SMD Surveys

The most ambitious use of Deep Interaction was the attempt to extract constituency-level estimates from agents. Hungary's electoral system allocates 106 of its 199 parliamentary seats through single-member districts, and modelling these requires geographic granularity that a national vote share model cannot provide.

Three rounds of SMD surveys were conducted, each refining the prompt to address problems encountered in the previous round.

Round 1 asked agents to estimate how many of the 106 SMD seats each party would win in their geographic area. The results were disappointing. Most agents defaulted to national-level estimates - returning figures like "Fidesz 70, TISZA 15" regardless of their segment's actual geography. Several agents deflected entirely.

Round 2 deployed an improved prompt, explicitly stating that the question was not about the national result, bounding answers to the agent's segment-specific constituencies, and naming the forbidden behaviour of describing one's persona instead of answering. Differentiation improved: urban agents began returning TISZA-leading estimates, while rural agents held firm for Fidesz. But defectors persisted.

Round 3 pushed further, with the strictest prompt yet. Compliance reached its highest level, with most agents providing formatted constituency-level estimates. A small number of agents - notably Agent 17 on one platform and Agent 37 on one platform - still broke persona entirely, responding as AI assistants rather than voter archetypes.

The three rounds produced 270 total data points (90 responses per round across 45 agents and 2 platforms). These were aggregated using a bloc method - grouping agents into Rural, Urban, and Mixed blocs weighted by demographic effective weight - to produce an initial SMD estimate. This estimate was later superseded by the Chronieler-v2 methodology described in Section 6, but the agent survey data informed the calibration of that model and confirmed the direction of its findings.

4.6 What Deep Interaction Added

Without Deep Interaction, the prediction would have been a single report-level output - a national vote share and a rough seat estimate, with no visibility into the segment dynamics underneath. The extraction process transformed this into a dataset of 13 segment-level vote share estimates, a demographic weighting model, and a set of constituency-level signals - all derived from the agents' own simulated beliefs, not imposed from outside.

It also exposed, through the pattern of artifacts and corrections required, the structural weaknesses in the simulation's architecture - weaknesses that became the subject of the next phase.

Section 5 - Emergent Social Phenomena

One of the more interesting aspects of running a social simulation is what it produces that you did not explicitly programme. The agent archetypes, the interaction rules, the influence weights - these are all designed inputs. But the patterns that emerge from 100 rounds of 45 agents interacting, updating, and broadcasting to each other reveal dynamics that no single design choice fully anticipated. Several of them turned out to be as analytically valuable as the vote share outputs themselves.

5.1 The Silenced Centre

Perhaps the most telling finding from the simulation's architecture was what happened - or rather, what did not happen - with the undecided voter segment.

Undecided voters are, in theory, the most interesting agents in any electoral simulation. They are the ones most susceptible to persuasion, most sensitive to social dynamics, most likely to swing an outcome. In this simulation, they were also the ones least able to participate in the process. The undecided agent had a posting rate of zero and an influence weight of 0.7 - the lowest in the simulation. It could receive influence from every other agent. It could not send any back.

The result was that the undecided segment defaulted to its initialisation prior - a Fidesz-leaning 55% - and stayed there. The deliberative process the simulation was designed to model never fired for the one segment it would have mattered most to. This is not just a technical finding. It is a recognisable social phenomenon: the disengaged centre, present in the room but not in the conversation, shaped by the loudest voices without shaping them in return.

5.2 The Rural Echo Chamber

The rural segments told a related story, but through a different mechanism. Three to four agents in the simulation mapped to the Rural Elderly archetype - more than double the representation their population weight would call for. Two of the highest influence weights in the entire simulation (1.5 each) belonged to Fidesz-aligned agents: the Religious Conservatives and the state-media voter archetype. No TISZA-aligned agent exceeded 1.3.

The combined effect was that Fidesz-supporting content was both more frequent and more persuasive in the shared network space than the demographic design intended. Rural agents were, in effect, talking to each other more, and being heard more, than their real-world weight warranted. The simulation did not intend to replicate Hungary's rural media environment - but it accidentally reproduced some of its structural logic.

5.3 The Ghost Persuader

Agent 41 - labelled "Political Journalist" in the configuration - carried the highest influence weight in the entire simulation: 2.0, significantly above every voter agent. It was classified as a

VoterSegment type, meaning it participated in the belief-updating network as if it were a voter. But it mapped to no demographic group, cast no weighted vote, and appeared nowhere in the final output.

This is what it effectively was: a ghost persuader. Maximum network influence, zero electoral accountability. Its stated sentiment bias was neutral (0.0), but neutrality cannot be verified from the outside, and the claim that a highly influential agent with no demographic grounding had no directional effect on the surrounding network is difficult to sustain. It is one of the more unusual artefacts of the simulation design.

5.4 The Shy Voter Asymmetry

Agent 40 was explicitly named "Shy Fidesz Voter" - a direct encoding of the well-documented social desirability phenomenon where Fidesz supporters understate their preference in social settings. This agent had a sentiment bias of +0.3 and an influence weight of 1.3, and it functioned as a built-in correction for the assumption that Fidesz voters would underperform their actual support in social discourse.

There was no equivalent "Shy TISZA Voter" agent. The asymmetry meant that the simulation corrected for one party's social desirability effect and not the other's - a structural tilt that, while small in absolute terms, ran consistently in one direction throughout all 100 rounds.

5.5 The Urban Surge

In contrast to the rural dynamics, the urban and Budapest segments produced something closer to what the design intended: genuine opinion movement. Urban professional and young first-time voter agents, starting from divided initial positions, converged measurably toward TISZA through successive rounds of interaction. The signal was consistent across both platforms - urban agents influencing each other in the direction of the opposition - and it produced the TISZA-dominant urban estimates that drove the final seat allocation in competitive constituencies.

This convergence is worth noting because it mirrors a real dynamic in Hungarian political sociology: the gradual sorting of urban, educated, and younger voters away from Fidesz and toward whatever opposition force carries the most credibility. The simulation did not need to be told this was happening. It produced it.

5.6 The Diaspora Effect

One of the quieter but structurally significant findings was the diaspora segment's impact on the national result. Ethnic Hungarian voters abroad - 4% of the population with a 95% turnout rate and a 90% Fidesz vote share - punched well above their demographic weight in the final aggregation. Their effective contribution to Fidesz's national total amounted to several percentage points, functioning as a reliable, high-turnout bloc with almost no variance.

This is not an emergent finding in the traditional sense - it reflects a known feature of Hungary's mail-in voting architecture, which Fidesz engineered specifically to capture this community. But seeing it surface through agent dynamics, unprompted, gave the simulation's representation of the electoral landscape additional credibility.

5.7 Cross-Platform Divergence

Finally, a more technical but revealing phenomenon: a meaningful number of agents produced notably different estimates on the two platforms. The same agent, representing the same archetype, would express different vote shares depending on the platform environment - sometimes dramatically so. Agent 34, for example, returned Fidesz 75% on one platform and TISZA 70% on the other.

This is not random noise. It suggests that the platform environment - the nature of the discourse, the peer agents most active in each space, the format of interaction - was genuinely shaping expressed opinion. Short-form public discourse and longer-form community discussion produced different social pressures, and agents responded to those pressures differently. It is, in miniature, a simulation of something real: the same voter can sound very different depending on where they are talking and who is listening.

Section 6 - Post-Hoc Epistemic Recalibration

6.1 Why This Was Necessary

The simulation report and the Deep Interaction extraction produced a dataset. But before treating that dataset as a prediction, a harder question had to be answered: how much of what the simulation produced reflected genuine emergent political dynamics, and how much reflected the structural properties of the simulation itself?

The answer, on inspection, was uncomfortable. A systematic audit of the simulation's inputs, architecture, and configuration revealed a set of identifiable biases - some introduced through asymmetries in the seed documents, others baked into the simulation's design parameters - that had been shaping agent behaviour throughout all 100 rounds without ever appearing in the output. Left uncorrected, these biases would have produced a prediction that was not wrong by chance but wrong by construction.

The recalibration process that followed belongs to what computational social scientists call the indirect inference correction tradition: when a simulation cannot be re-run, known input and process biases can be corrected post-hoc by applying directionally-justified adjustments to the output, provided each correction is independently sourceable, conservatively estimated, and presented alongside the raw output rather than instead of it.

That is the standard this section holds itself to.

6.2 Two Layers of Bias

Before cataloguing individual issues, it is worth distinguishing two analytically separate sources of distortion, because they require different correction methods.

Layer 1 - Seed Document Bias covers asymmetries in the six input documents that caused agents to be initialised with skewed priors. These biases affect what agents believed before any social interaction took place.

Layer 2 - Process and Architecture Bias covers flaws in how the simulation was configured that distorted how beliefs evolved during the 100-round run. These biases affect the dynamics through which beliefs changed.

Some issues span both layers. All of them ran, with very few exceptions, in the same direction.

6.3 The Issue Register

Eleven distinct issues were identified across the two layers, plus two additional findings uncovered through direct inspection of the simulation configuration file. Each was classified into one of three tiers based on its estimated impact and the confidence in the correction direction.

Tier 1 - Full numerical corrections were applied to four issues where the bias had a clean, mathematically definable solution.

Tier 2 - Quantified adjustments with uncertainty ranges were applied to issues where the correction direction was clear but the magnitude required scenario modelling.

Tier 3 - Qualitative caveats were documented but not numerically corrected, to avoid the risk of over-correction on issues where quantifying the bias would have introduced more uncertainty than it resolved.

The most significant issues, briefly:

Agent 10 manual correction (Tier 1) - The single largest editorial intervention in the dataset. Agent 10, a rural mid-age agent, had emerged from the simulation with TISZA 45% / Fidesz 35% - a TISZA-leading rural result. This was manually corrected during extraction to Fidesz 60% / TISZA 22.5% on the grounds that a TISZA-leading rural agent was "politically implausible." The problem is that this judgment bakes in exactly the prior assumption the simulation was designed to test. The correction was reverted.

Demographic double-counting (Tier 1) - The weighting model treated Religious Conservatives (7% of population) and Hungarian Families (5%) as fully independent demographic segments, when in reality approximately 40% of the former and 50% of the latter were already counted within the Rural Elderly and Rural Mid-Age buckets. This inflated the pre-normalisation denominator with Fidesz-heavy voters counted twice, artificially boosting Fidesz's weighted national share.

Undecided agents structurally silenced (Tier 2) - Covered in Section 5. The undecided segment's posts-per-hour of zero meant it defaulted to its Fidesz-leaning initialisation prior throughout the entire run.

Chinese activity schedule (Tier 2) - The simulation configuration explicitly noted that agent activity patterns followed "typical Chinese work schedule habits" - suppressing the 6-10am morning window when rural Hungarian voters consume state television, and amplifying the 19-22pm urban evening window when TISZA-leaning agents dominate. The two effects partially cancelled but were not symmetric.

Influence weight asymmetry (Tier 2) - The two highest influence weights in the simulation (1.5 each) both belonged to Fidesz-aligned agents. The maximum TISZA-aligned weight was 1.3 - a 15.4% gap that compounded across every shared-network interaction across all 100 rounds.

Shy Fidesz Voter asymmetry (Tier 2) - Agent 40 encoded a social desirability correction for Fidesz underreporting. No equivalent agent existed for TISZA.

Ghost persuader (Tier 3) - Agent 41, a "Political Journalist" with the highest influence weight in the simulation (2.0), participated in the belief-updating network as a VoterSegment but mapped to no demographic group and cast no weighted vote. Its directional effect cannot be verified.

D'Hondt applied as simple proportional (Tier 1) - The seat calculation used a simplified formula rather than the true iterative D'Hondt divisor method, systematically understating the largest party's list seat count.

Mi Hazánk phantom SMD seat (Tier 1) - Mi Hazánk was assigned one SMD win at 4.0% national support with no specific constituency modelled. This was removed.

6.4 Correction Batches - Vote Shares

Corrections were applied in three sequential batches, each preserving the running state for the next.

Batch 1 applied the three highest-impact corrections: the Agent 10 revert, the demographic double-counting fix, and the undecided silencing correction. Together these moved Fidesz -2.97pp and TISZA +2.18pp, narrowing the raw simulation's 9.2pp Fidesz-TISZA gap to 4.1pp.

Batch 2 applied the Tier 2 weighted adjustments: the combined influence weight and archetype duplication correction, the undecided equilibrium floor adjustment, the RTL Klub independence overstatement, and the Shy Fidesz Voter asymmetry removal. These produced smaller but directionally consistent movements.

Batch 3 completed the Tier 1 structural corrections: the true D'Hondt re-run on list seats and the removal of the Mi Hazánk phantom SMD seat.

The final corrected national vote shares after all batches:

Party	Raw Simulation	Corrected	Change
Fidesz	43.8%	41.0%	-2.8pp
TISZA	34.6%	36.7%	+2.1pp
DK	6.8%	7.0%	+0.2pp
MKKP	6.6%	6.7%	+0.1pp
Mi Hazánk	4.0%	3.9%	-0.1pp

The corrected result lands within 1pp of Median Research's independent polling for both major parties - without having been anchored to that polling at any point in the correction process. This alignment is treated as external validation, not as the target.

6.5 SMD Correction: From Agent Survey to Chronicler-v2

The initial SMD estimate, derived from the three-round Deep Interaction survey, produced Fidesz 73 seats and TISZA 24. A subsequent audit identified four structural problems with this figure.

SMD1 - No quantitative methodology. The original prediction.md SMD estimate was a narrative judgment ("based on geographic distribution of Fidesz vs. opposition support") with no uniform swing calculation, no 2022 constituency baseline, and no documented method.

SMD2 - DK winning 4 SMDs. For DK to win four SMDs outright, their candidates would need to beat both Fidesz and TISZA locally - impossible in a landscape where TISZA polls 48% in Budapest metro against DK's 14% in the same segment. In 2022, DK ran as part of a unified opposition coalition and had no standalone constituency baseline. A defensible estimate is 0-1 SMD wins.

SMD3 - MKKP winning 2 SMDs. MKKP has no constituency infrastructure, no standalone 2022 baseline, and faces the same strategic voting pressure as DK. Defensible estimate: 0.

SMD4 - Vote-splitting risk ignored. If DK and MKKP run candidates in marginal constituencies, they draw votes away from TISZA - potentially allowing Fidesz to hold seats it would otherwise lose. This is the exact mechanism that gave Fidesz 96 SMDs in 2014 despite winning 44% nationally.

To correct all four issues simultaneously, the Chronicler-v2 methodology recommended in Document 1 was applied. This constructs a synthetic linear distribution of all 106 constituencies calibrated to the known 2022 aggregate result (Fidesz 54.13%, 87 SMD wins), applies the corrected 2026 uniform swing (-13.15pp for Fidesz), and models strategic vote transfer from DK and MKKP supporters to TISZA through the parameter λ .

6.6 The λ Calibration

λ represents the fraction of DK and MKKP voters who cast their SMD ballot for TISZA rather than their own party's candidate. It is the single most consequential variable in the seat model.

Its value was calibrated from real-world candidate registration data:

- DK confirmed approximately 95 candidates across 106 constituencies, with 3 official withdrawals targeting marginals where TISZA had realistic winning chances. This partial cooperation behaviour, combined with the loyalty rate of remaining candidates, produced $\lambda_{DK} = 0.42$.

- MKKP explicitly confirmed it would not field 106 candidates, concentrating on urban areas with approximately 65 active constituencies. In the 41 constituencies where no MKKP candidate exists, MKKP list voters have no SMD option but TISZA, giving $\lambda = 1.0$ locally. Weighted across all constituencies: $\lambda_{MKKP} = 0.63$.

- The effective combined λ , weighted toward marginal seats where the outcome is decided: $\lambda_{eff} = 0.58$.

This figure sits at the boundary of two electoral outcomes - a finding returned to in full in the Scenarios section.

At $\lambda = 0.58$, applying the simulation's corrected 2026 vote shares as a uniform swing on the 2022 constituency baseline, the Chronieler-v2 model produces:

Party	SMD Seats
TISZA	66
Fidesz-KDNP	38
DK	1
MKKP	0
Mi Hazánk	0
Other / Independent	1
Total	106

This figure replaced the agent survey SMD estimate as the definitive input to the final seat allocation.

Section 7 - Final Results

This is what the simulation produced, corrected, and calibrated to the mechanics of the Hungarian electoral system.

7.1 National Vote Shares

After all recalibration corrections, the simulation's corrected national vote shares are:

The Fidesz-TISZA gap stands at 4.3 percentage points - less than half what the raw simulation produced before recalibration. Mi Hazánk falls below the 5% national list threshold and is therefore excluded from list seat allocation entirely, receiving no proportional representation regardless of its raw vote share.

These percentages do not emerge from a poll. They emerge from 45 agents reasoning through 100 rounds of simulated social interaction, subsequently corrected for eleven identified structural biases. The fact that the corrected result aligns within 1 percentage point of Median Research's independent polling for both major parties - without having used that polling as a correction target at any stage - provides meaningful external validation of the methodology.

What drove these numbers? The simulation's agent dynamics tell a clear story. Fidesz's support is concentrated and loyal in rural and elderly segments, reinforced by state media consumption patterns and high diaspora turnout (90% Fidesz at a 95% participation rate). TISZA's support is urban, educated, and surging among first-time voters - a coalition that is newer and more volatile but proved broad enough, once aggregated across the full demographic weighting model, to reach 36.7% nationally. DK and MKKP both clear the threshold, their vote shares compressed by TISZA's rise but sufficient to maintain parliamentary presence.

7.2 SMD Seat Distribution

The 106 single-member district seats are allocated by first-past-the-post in individual constituencies. The result here, derived by applying the simulation's corrected vote shares as a uniform swing on the 2022 constituency baseline through the Chronicler-v2 model at $\lambda = 0.58$, is:

Party	SMD Seats
TISZA	66
Fidesz-KDNP	38
DK	1
MKKP	0
Mi Hazánk	0
Other / Independent	1
Total	106

This result requires context. TISZA winning 66 SMDs - more than Fidesz - is not simply a function of their vote share. It is the product of strategic vote consolidation. At $\lambda = 0.58$, 58% of

DK and MKKP voters cast their constituency ballot for TISZA rather than their own party's candidate. In the 41 constituencies where MKKP fields no candidate at all, that transfer is automatic. In the marginal constituencies where DK made targeted withdrawals - specifically designed to clear the path for TISZA - the transfer rate is higher still. TISZA does not win 66 seats because it is more popular than Fidesz nationally. It wins 66 seats because the opposition vote consolidates behind a single candidate in enough constituencies to tip the balance.

Fidesz's 38 SMD wins represent their rural heartland: the constituencies where their vote share, even after a -13 percentage point uniform swing from 2022, remains above the combined opposition. These seats are structurally safe. In the competitive suburban and small-town constituencies where the election is decided, the consolidation dynamic works against them.

7.3 List Seat Distribution

The 93 list seats are where Hungary's electoral system produces its most counter-intuitive results, through the compensation mechanism that feeds surplus and loser votes back into the D'Hondt pool.

With Fidesz losing 68 SMDs, every vote cast for Fidesz in those 68 constituencies flows back as loser compensation into their national list pool. This inflates their effective pool share to 50.2% - well above their raw vote share of 41% - and gives them significantly more list seats than a

simple proportional allocation would suggest. TISZA, winning 66 SMDs, loses only 40 - meaning their loser compensation is far smaller, and their list pool share falls to 37.4% despite their raw vote share being close.

The system is, in a sense, self-correcting: the more SMDs you lose, the more list seats you recover. This partially cushions Fidesz's SMD collapse without reversing it.

One seat is subtracted from the regular 92-seat D'Hondt pool for the German minority preferential seat, a historically consistent allocation that goes to a pro-government aligned list.

Party	CMP Pool Share	List Seats
Fidesz-KDNP	50.2%	46
TISZA	37.4%	35
DK	6.6%	6
MKKP	5.8%	5
German Minority	—	1
Total	—	93

7.4 Final 199-Seat Parliament

Party	SMD	List	Total	% of Parliament
TISZA	66	35	101	50.8%
Fidesz-KDNP	38	46	84	42.2%
DK	1	6	7	3.5%
MKKP	0	5	5	2.5%
German Minority	0	1	1	0.5%
Other / Independent	1	0	1	0.5%
Total	106	93	199	100%

7.5 What This Means

TISZA reaches a bare majority at 101 seats - one seat above the 100-seat threshold required to govern. This is not a comfortable majority. It is the minimum viable result, dependent on every seat holding, on DK and MKKP not defecting on key votes, and on the single independent seat not becoming a critical variable. It is enough to form a government and pass ordinary legislation. It is not enough to change the constitution, which requires 133 seats.

Fidesz at 84 seats represents a dramatic collapse from their 2022 position of 135 - a loss of 51 seats and the end of their supermajority era. They remain the largest single party by vote share at 41%, and the compensation mechanism gives them 46 list seats from just 38 SMD wins. But they are in opposition.

The compensation mechanism produces its intended effect. The Hungarian system was designed precisely to prevent landslide seat distortions from being entirely irreversible. Fidesz loses the SMDs but recovers meaningfully through the list. TISZA wins the SMDs but gets fewer list seats than a simple proportional model would give them. The result is a parliament that is genuinely competitive rather than structurally lopsided - which is, ironically, what the system was supposed to produce before it was redesigned to do the opposite.

DK's survival at 7% is quietly significant. Their 7 total seats keep Gyurcsány's party in parliament and guarantee that the opposition bloc is not a monolith. Whether that is a strength or a liability for TISZA's governing coalition is a political question the simulation does not answer.

Section 8 - Scenarios and Sensitivity

The final prediction presented in Section 7 is a point estimate, not a certainty. It rests on a set of inputs - corrected vote shares, a calibrated λ , a 2022 constituency baseline - each of which carries its own uncertainty range. This section maps what happens to the outcome when those inputs move, and identifies which variables matter most.

8.1 The Variable That Decides Everything: λ

The single most consequential input in the entire seat model is not the vote share. It is λ - the strategic voting transfer rate from DK and MKKP SMD voters to TISZA. The sensitivity here is stark.

Scenario	λ	Fidesz SMDs	TISZA SMDs	Fidesz Total	TISZA Total	Outcome
Low	0.30	54	51	95	87	Hung parliament
Moderate	0.55	41	64	83	99	Hung parliament
Calibrated	0.58	38	66	84	101	TISZA bare majority
High	0.80	27	78	71	111	TISZA majority

Two tipping points define the landscape. Fidesz loses its majority the moment λ exceeds 0.20 - meaning any meaningful strategic coordination between opposition voters is enough to deny Orbán a governing majority. TISZA gains its own majority only when λ crosses 0.57 - a much higher bar, requiring genuine mass strategic behaviour, not just scattered defections.

Between those two tipping points lies a hung parliament zone where neither party reaches 100 seats. This is not a fringe scenario. At $\lambda = 0.55$ - just 0.03 below the calibrated estimate - TISZA finishes at 99 seats, one short of a majority, and Hungary enters uncharted constitutional territory.

The calibrated $\lambda = 0.58$ sits one hundredth of a percentage point above the TISZA majority threshold. That margin should not inspire false precision. What it should inspire is an honest acknowledgement that the difference between a TISZA government and a hung parliament is, in this model, a question of whether strategic voting consolidation fires fully or falls fractionally short.

8.2 What Moves λ

Several real-world factors could push λ in either direction before election day.

Upward pressure on λ - more strategic consolidation - comes from further DK candidate withdrawals in marginal seats, strong get-out-the-vote messaging from civil society groups emphasising the tactical SMD vote, and any last-minute polling that makes the scale of the

TISZA-Fidesz competition visible to undecided DK and MKKP voters in competitive constituencies.

Downward pressure on λ - less consolidation, more vote splitting - comes from DK's organisational resistance to further withdrawals (Dobrev expelled a candidate who withdrew without authorisation), MKKP's protest-party identity making its voters less susceptible to tactical appeals, and voter confusion in constituencies where multiple opposition candidates

remain on the ballot.

The calibrated $\lambda = 0.58$ already accounts for this tension. It weights marginal seats more heavily than average constituencies, reflecting that DK's targeted withdrawals were concentrated precisely where the outcome is decided.

8.3 Vote Share Sensitivity

The national vote shares feed the Chronicler-v2 swing calculation. If they move, the number of constituencies that flip changes.

If Fidesz overperforms by 2pp (41% \rightarrow 43%) - moving back toward the upper range of state-adjacent polling - the uniform swing narrows from -13.15pp to -11.15pp. Fewer constituencies cross the 50% threshold. At $\lambda = 0.58$, Fidesz likely recovers 4-6 SMDs, bringing their total to around 88-90. TISZA drops to 95-97 seats - back into hung parliament territory.

If TISZA underperforms by 2pp (36.7% \rightarrow 34.7%) - the combined opposition lead narrows, the swing is smaller, and the marginal constituencies stay with Fidesz. The effect is similar: 4-6 seats shift, hung parliament becomes the central scenario.

If Fidesz underperforms by 2pp (41% \rightarrow 39%) - the swing deepens, more constituencies flip, and TISZA's majority becomes more comfortable: approximately 105-108 seats.

The vote share range where the outcome is genuinely uncertain is roughly Fidesz 39-43% . Below 39%, TISZA governs comfortably. Above 43%, the hung parliament becomes the most likely scenario. The corrected simulation puts Fidesz at 41% - directly in the centre of that uncertainty band.

8.4 Threshold Risk

Two smaller parties sit close enough to the 5% list threshold to warrant scenario modelling.

DK at 7.0% clears the threshold with a 2pp buffer. A late campaign collapse - possible given DK's internal tensions and Dobrev's declining personal approval - could bring them close to 5%. If DK falls below threshold, their 6 list seats disappear, their compensation votes no longer generate list representation, and the D'Hondt pool redistributes those seats between Fidesz and

TISZA. The net effect would be approximately Fidesz +2 list seats, TISZA +1 list seat - small but not trivial in a 101-seat majority scenario.

MKKP at 6.7% faces similar arithmetic. A drop below 5% would cost them their 5 list seats, with similar redistribution dynamics. MKKP's threshold risk is somewhat lower given their stable protest-party base, but not negligible.

8.5 The Hung Parliament Scenario

It is worth dwelling on what a hung parliament actually means in the Hungarian context, because the implications are unusually severe.

Hungary has no established coalition culture. The electoral system was explicitly redesigned in 2011 to produce single-party majorities - and it has done so in every election since. A result where neither party reaches 100 seats would be without modern precedent, and the constitutional mechanisms for resolving it are neither well-tested nor politically neutral. The President - currently a Fidesz-aligned figure - would play a central role in designating a formateur, a process that could become deeply contentious.

At $\lambda = 0.55$, Hungary is in this scenario. At $\lambda = 0.58$, it narrowly escapes it. That one hundredth of a percentage point is the structural fact this simulation returns to, repeatedly, as its most honest finding.

Section 9 - Researcher's Perspective

9.1 On the Results

Going into the simulation, I expected TISZA to win. What I did not expect was how much depth the simulation would be able to give to each segment's social behaviours. The focus on interactions, strategic voting, and decision-making processes was striking - but especially the differences between consolidated and non-consolidated voter segments. Consolidated groups like rural elderly Fidesz voters showed tight, internally reinforcing dynamics, while non-consolidated segments like undecided voters produced much noisier behaviour with higher variance across platforms. That distinction maps cleanly onto what we know about these groups politically, which gave me confidence the simulation was capturing something real.

What struck me most in terms of results was the seat allocation dynamics. Reading through the findings on national list seats and SMD seats made me realise how important the D'Hondt system is and how crucial strategic voting was in assigning SMD seats to TISZA candidates. The compensation mechanism was counter-intuitive at first - Fidesz losing more SMDs

actually inflates their list pool through loser compensation - but seeing it quantified made the mechanics of Hungarian democracy feel very concrete in a way that reading about them abstractly does not.

The diaspora finding was also very interesting. A 4% population share voting 90% Fidesz and contributing roughly 3.5 points to the national total is a legal structural advantage that polling rarely surfaces cleanly. It was a useful reminder that electoral outcomes in Hungary are shaped by rules as much as by preferences.

Lastly, I was quite surprised by how well DK and MKKP performed at a national level compared to recent polls. Both clearing the 5% threshold is not guaranteed in current polling, and whether they hold on election day is one of the aspects I am most looking forward to comparing with the actual results.

9.2 On the Method

The most immediate lesson is that a single simulation run is not enough. To make genuinely probabilistic claims you need multiple runs with varied seeds and agent configurations - at \$7.95 per run that is still financially accessible, but the tooling to aggregate and compare results across runs does not yet exist in any automated way, and building it is the obvious next step.

The hallucination rate - ten manual corrections out of 45 agents - makes clear that post-hoc bias correction is structural, not optional. The correction pipeline here was largely ad hoc, which works at 45 agents but does not scale. Future versions need a formalised plausibility layer applied automatically against demographic priors, rather than researcher judgement case by case.

The underlying dynamics are genuinely interesting and worth developing further. As models improve and agent frameworks mature, the ability to simulate opinion formation at realistic scale becomes more plausible. What this experiment shows is that the conceptual architecture is already there - the limiting factors are computational and methodological, not theoretical.

Most importantly, I ran this on my own laptop, as a 21-year-old student in Economics and Politics, with no data science background, in the middle of LSE exam revision, on less than ten dollars in API credits and Claude Code for assistance. Hopefully this is a small proof that AI can and should be for everyone. Or at least, for everyone crazy enough to try and wrap their head around it. It can be hard at first - but once you start understanding these systems, endless doors open to you.

Section 10 - Resource Footprint

10.1 API Usage and Cost

The actual figures from the OpenRouter dashboard:

Metric	Value
Total spend (Gemini 2.5 Flash Lite)	\$7.95
Total requests	5,150
Total tokens	143M

Five Claude Code sessions using Sonnet 4.6 on medium effort add an estimated ~150K tokens (~\$0.80 at Sonnet 4.6 pricing), used for architecture assistance, correction logic, and report drafting. Neo4j AuraDB ran on the free tier throughout - no direct cost. Total spend across all services: under \$9.

At UK grid carbon intensity (~250 gCO₂ /kWh), total emissions are approximately 625-750 gCO₂ - roughly equivalent to a 5km car journey. Under renewable-matched cloud accounting

the cloud components drop significantly, but the laptop's local draw is not offset by any certificate, making it the honest floor of the estimate.

10.3 Water Footprint

Water use applies to the cloud components only. At ~1.1 litres per kWh for Google and AWS data centre cooling, the cloud workload (~0.8-1.0 kWh combined) accounts for roughly 0.9-1.1 litres of water. The laptop contributes nothing to data centre water draw.

10.4 Context

The full experiment - including 48 hours of active laptop use - consumed roughly the energy equivalent of leaving a lightbulb on for two days. The cloud API costs were under \$9. For a solo student project, these figures are not a constraint worth worrying about. At larger scale, the laptop component scales out and gets replaced by server infrastructure, at which point cloud efficiency gains become more meaningful. For now, the binding resource in this experiment was time, not energy or money.

10.2 Energy and Carbon Footprint

Component	Est. Energy	Notes
Laptop (48h active use, ~40W avg)	~1.92 kWh	Largest single component
Gemini 2.5 Flash Lite (143M tokens)	~0.43–0.72 kWh	TPU inference
Neo4j AuraDB (shared free-tier instance)	~0.10–0.30 kWh	Cloud graph DB
Claude Code / Sonnet 4.6 (~150K tokens)	~0.002–0.008 kWh	Per-token cost higher than Flash Lite
Total	~2.5–3.0 kWh	

Appendices

Appendix A – Full Agent Roster (45 Agents)

Project: proj_f3e22472ad5d | Simulation: sim_6b6e4e5128ec | April 10, 2026

#	Handle	Name (ID)	Role
0	@Archetype 1 undecideds	archetype_1_undecideds_687	Electoral Information Facilitator
1	@UNDECIDED RESOLUTION	undecided_resolution_342	Voter Segment Analyst
2	@MKKP/DK supporters	mkkpdk_supporters_941	Voter Segment Representation
3	@rural elderly agents	rural_elderly_agents_249	Voter Segment Representation
4	@Archetype 7	archetype_7_624	Voter Segment & Diaspora Engagement
5	@Archetypes 1, 5, 8	archetypes_1_5_8_767	Voter Segment Representation
6	@Archetype 10	archetype_10_604	Voter Segment Analyst
7	@Hungarian families	hungarian_families_740	Representative for Hungarian Families
8	@Religious conservative agents	religious_conservative_agents_539	Voter Segment Advocacy
9	@Roma agents	roma_agents_938	Voter Segment Advocacy & Info Hub
10	@Rural agents	rural_agents_867	Voter Segment Representation
11	@Archetype 1	archetype_1_990	Community Engagement – Rural Elderly
12	@Archetype 5	archetype_5_590	Voter Segment Representation
13	@Archetype 8	archetype_8_512	Voter Segment Representation
14	@Urban agents	urban_agents_416	Voter Segment Representation
15	@Archetype 3	archetype_3_751	Budapest Professional Segment
16	@Archetype 4	archetype_4_529	Voter Segment Representation
17	@media-literate agents	medialiterate_agents_225	Media Literacy Advocacy Group
18	@Archetype 2	archetype_2_111	Voter Segment Analyst
19	@Diaspora agents	diaspora_agents_667	Diaspora Engagement & Info Hub
20	@Archetype 6	archetype_6_226	Diaspora Voter Segment Representative
21	@Diaspora Mail-In	diaspora_mailin_303	Voter Segment Information
22	@Emigrated	emigrated_438	Voter Segment Info Dissemination
23	@Religious Conserv.	religious_conserv_864	Voter Segment Representation

#	Handle	Name (ID)	Role
24	@Suburban Family	suburban_family_339	Voter Segment Representation
25	@Small-Town Worker	smalltown_worker_478	Voter Segment Representative
26	@Budapest Prof.	budapest_prof_105	Voter Segment Analyst
27	@Young Urban	young_urban_146	Voter Segment Representation
28	@Rural Business	rural_business_112	Voter Segment Representation
29	@Religious Conservative	religious_conservative_341	Voter Segment Advocacy
30	@Suburban Middle-Income Family	suburban_middleincome_family_170	Voter Segment Simulation
31	@Rural Elderly Pensioner (65+)	rural_elderly_pensioner_65_834	Voter Segment Information
32	@Small-Town Skilled Worker (35-55)	smalltown_skilled_worker_3555_844	Advocacy – Small-Town Workers
33	@Budapest Educated Professional (35-55)	budapest_educated_professional_2555_212	Informed Citizenry Segment
34	@Young Urban First-Time Voter (18-26)	young_urban_firsttime_voter_1825_953	Youth Civic Engagement
35	@Rural Small Business Owner (35-55)	rural_small_business_owner_3555_849	Rural Business Advocacy
36	@Ethnic Hungarian Diaspora (mail-in)	ethnic_hungarian_diaspora_mailin_983	Voter Segment Information
37	@Emigrated Hungarian (consulate vote)	emigrated_hungarian_consulate_vote_658	Voter Segment Representation
38	@Undecided	undecided_854	Electoral Information Facilitator
39	@voters w/ state-media diet	voters_with_statemedia_information_diet_792	Dissemination
40	@shy Fidesz voter	shy_fidesz_voter_409	Simulated Voter Segment Persona
41	@political journalist archetype	political_journalist_archetype_274	Political Journalist & Analyst
42	@politically active young urban voter	politically_active_young_urban_voter_136	Civic Engagement
43	@Roma	roma_261	Voter Segment Representation
44	@Roma communities	roma_communities_581	Community Advocacy

Appendix B – Agent Activity Configuration

#	Agent Name	Stance	Posts/h	Cmts/h	Activity	Sentiment	Influence
0	Archetype 1 undecideds	OBSERVER	2	5	70%	0.0	1.0
1	UNDECIDED RESOLUTION	NEUTRAL	3	7	80%	+0.1	1.1
2	MKKP/DK supporters	SUPPORTIVE	2	6	70%	+0.3	1.3
3	rural elderly agents	SUPPORTIVE	1	3	40%	+0.2	1.4
4	Archetype 7	OPPOSING	2	4	60%	-0.1	1.2
5	Archetypes 1, 5, 8	OBSERVER	3	6	75%	0.0	1.0
6	Archetype 10	NEUTRAL	1	3	50%	0.0	1.1
7	Hungarian families	SUPPORTIVE	2	5	60%	+0.3	1.3
8	Religious conservative agents	SUPPORTIVE	1	4	55%	+0.4	1.5
9	Roma agents	NEUTRAL	1	2	45%	+0.1	0.9
10	Rural agents	NEUTRAL	1	3	50%	+0.1	1.0
11	Archetype 1	OBSERVER	2	5	70%	0.0	1.0
12	Archetype 5	NEUTRAL	2	4	60%	-0.1	1.1
13	Archetype 8	OPPOSING	1	3	55%	-0.1	1.2
14	Urban agents	NEUTRAL	2	5	65%	+0.1	1.1
15	Archetype 3	NEUTRAL	2	5	70%	+0.1	1.0
16	Archetype 4	NEUTRAL	2	5	70%	-0.1	1.0
17	media-literate agents	OBSERVER	3	8	80%	+0.2	1.2
18	Archetype 2	NEUTRAL	2	5	70%	0.0	1.0
19	Diaspora agents	OBSERVER	1	3	60%	+0.1	0.9
20	Archetype 6	NEUTRAL	2	5	70%	-0.2	1.0
21	Diaspora Mail-In	OBSERVER	1	2	50%	0.0	0.8
22	Emigrated	OBSERVER	1	4	60%	+0.1	0.9
23	Religious Conserv.	SUPPORTIVE	2	6	70%	+0.3	1.1
24	Suburban Family	NEUTRAL	2	5	70%	+0.1	1.0
25	Small-Town Worker	NEUTRAL	2	5	70%	-0.1	1.0
26	Budapest Prof.	NEUTRAL	3	7	80%	+0.2	1.2

#	Agent Name	Stance	Posts/h	Cmts/h	Activity	Sentiment	Influence
27	Young Urban	NEUTRAL	3	8	80%	+0.1	1.2
28	Rural Business	NEUTRAL	2	5	70%	+0.1	1.0
29	Religious Conservative	SUPPORTIVE	2	6	70%	+0.3	1.1
30	Suburban Middle-Income Family	NEUTRAL	2	5	70%	+0.1	1.0
31	Rural Elderly Pensioner (65+)	SUPPORTIVE	1	3	50%	+0.2	1.1
32	Small-Town Skilled Worker (35-55)	NEUTRAL	2	4	60%	0.0	1.0
33	Budapest Educated Professional	OPPOSING	3	6	80%	-0.1	1.3
34	Young Urban First-Time Voter	NEUTRAL	3	7	70%	0.0	0.9
35	Rural Small Business Owner	SUPPORTIVE	1	3	50%	+0.1	1.1
36	Ethnic Hungarian Diaspora (mail-in)	SUPPORTIVE	1	2	40%	+0.3	1.2
37	Emigrated Hungarian (consulate)	SUPPORTIVE	1	2	40%	+0.2	1.1
38	Undecided	OBSERVER	0	3	30%	0.0	0.7
39	voters w/ state-media diet	SUPPORTIVE	2	4	60%	+0.4	1.5
40	shy Fidesz voter	SUPPORTIVE	1	3	50%	+0.3	1.3
41	political journalist archetype	OBSERVER	2	5	50%	0.0	2.0
42	politically active young urban voter	OPPOSING	4	8	80%	-0.2	1.2
43	Roma	NEUTRAL	2	4	60%	+0.1	0.9
44	Roma communities	NEUTRAL	1	3	50%	0.0	0.9

Appendix C – DOC6 Archetype Mapping

Archetype	Description	Agent handles
1	Rural Elderly Pensioner (65+)	rural_elderly_agents, rural_elderly_pensioner_65, archetype_1, archetype_1_undecideds, archetype_1_undecideds_agents
2	Small-Town Skilled Worker (35-55)	smalltown_worker, smalltown_skilled_worker_3555, archetype_2
3	Budapest Educated Professional (25-55)	archetype_3, budapest_prof, budapest_educated_professional2555, medialiterate_agents
4	Young Urban First-Time Voter (18-25)	archetype_4, young_urban, young_urban_firsttime_voter, politically_active_young_urban_voter
5	Rural Small Business Owner (35-55)	archetype_5, rural_business, rural_small_business_owner_3555, rural_agents
6	Ethnic Hungarian Diaspora (mail-in)	archetype_6, diaspora_mailin, ethnic_hungarian_diaspora_mailin, diaspora_agents
7	Emigrated Hungarian (consulate vote)	archetype_7, emigrated, emigrated_hungarian_consulate_vote
8	Religious Conservative (45-70)	archetype_8, religious_conserv, religious_conservative, religious_conservative_agents
9	Roma Minority Voter	roma_agents, roma, roma_communities
10	Suburban Middle-Income Family (30-50)	archetype_10, suburban_family, suburban_middleincome_family, hungarian_families
Cross	Cross-cutting / Meta agents	mkkpdk_supporters, undecided_resolution, undecided, shy_fidesz_voter, voters_with_statemedia

Appendix D – Final Results Tables

D.1 Corrected National Vote Shares

Party	Corrected Vote Share	5% List Threshold
Fidesz-KDNP	41.0%	✓ Cleared
TISZA	36.7%	✓ Cleared
DK	7.0%	✓ Cleared
MKKP	6.7%	✓ Cleared
Mi Hazánk	3.9%	✗ Excluded

D.2 SMD Seat Distribution ($\lambda = 0.58$)

Party	SMD Seats
TISZA	66
Fidesz-KDNP	38
DK	1
MKKP	0
Other / Independent	1

Party	SMD Seats
Total	106

D.3 List Seat Distribution (D'Hondt + CMP)

Party	Pool Share	List Seats
Fidesz-KDNP	50.2%	46
TISZA	37.4%	35
DK	6.6%	6
MKKP	5.8%	5
German Minority	—	1
Total	—	93

D.4 Final 199-Seat Parliament

Party	SMD	List	Total	% of Parliament
TISZA	66	35	101	50.8%
Fidesz-KDNP	38	46	84	42.2%
DK	1	6	7	3.5%
MKKP	0	5	5	2.5%
German Minority	0	1	1	0.5%
Other / Independent	1	0	1	0.5%
Total	106	93	199	100%

D.5 λ Sensitivity Analysis

Scenario	λ	Fidesz SMDs	TISZA SMDs	Fidesz Total	TISZA Total	Outcome
Low	0.30	54	51	95	87	Hung parliament
Moderate	0.55	41	64	83	99	Hung parliament
Calibrated	0.58	38	66	84	101	TISZA bare majority
High	0.80	27	78	71	111	TISZA majority

Notices

References

For timing purposes, a formal references section has not been included in this document. Readers wishing to obtain a full list of sources, polling data providers, electoral law citations, academic frameworks referenced, and software documentation are welcome to contact the author directly.

Contact: Matteo Mio

Use of Generative AI

Generative AI tools were partly used in the making of this document for purposes of structure, drafting, and writing assistance. The research design, agent configuration, epistemic recalibration decisions, analytical interpretation, and all quantitative calculations are the work of the author. AI-assisted writing was reviewed and edited by the author throughout. The simulation itself was conducted using AI-powered agents as described in the methodology sections above; this notice refers specifically to the production of this written report.