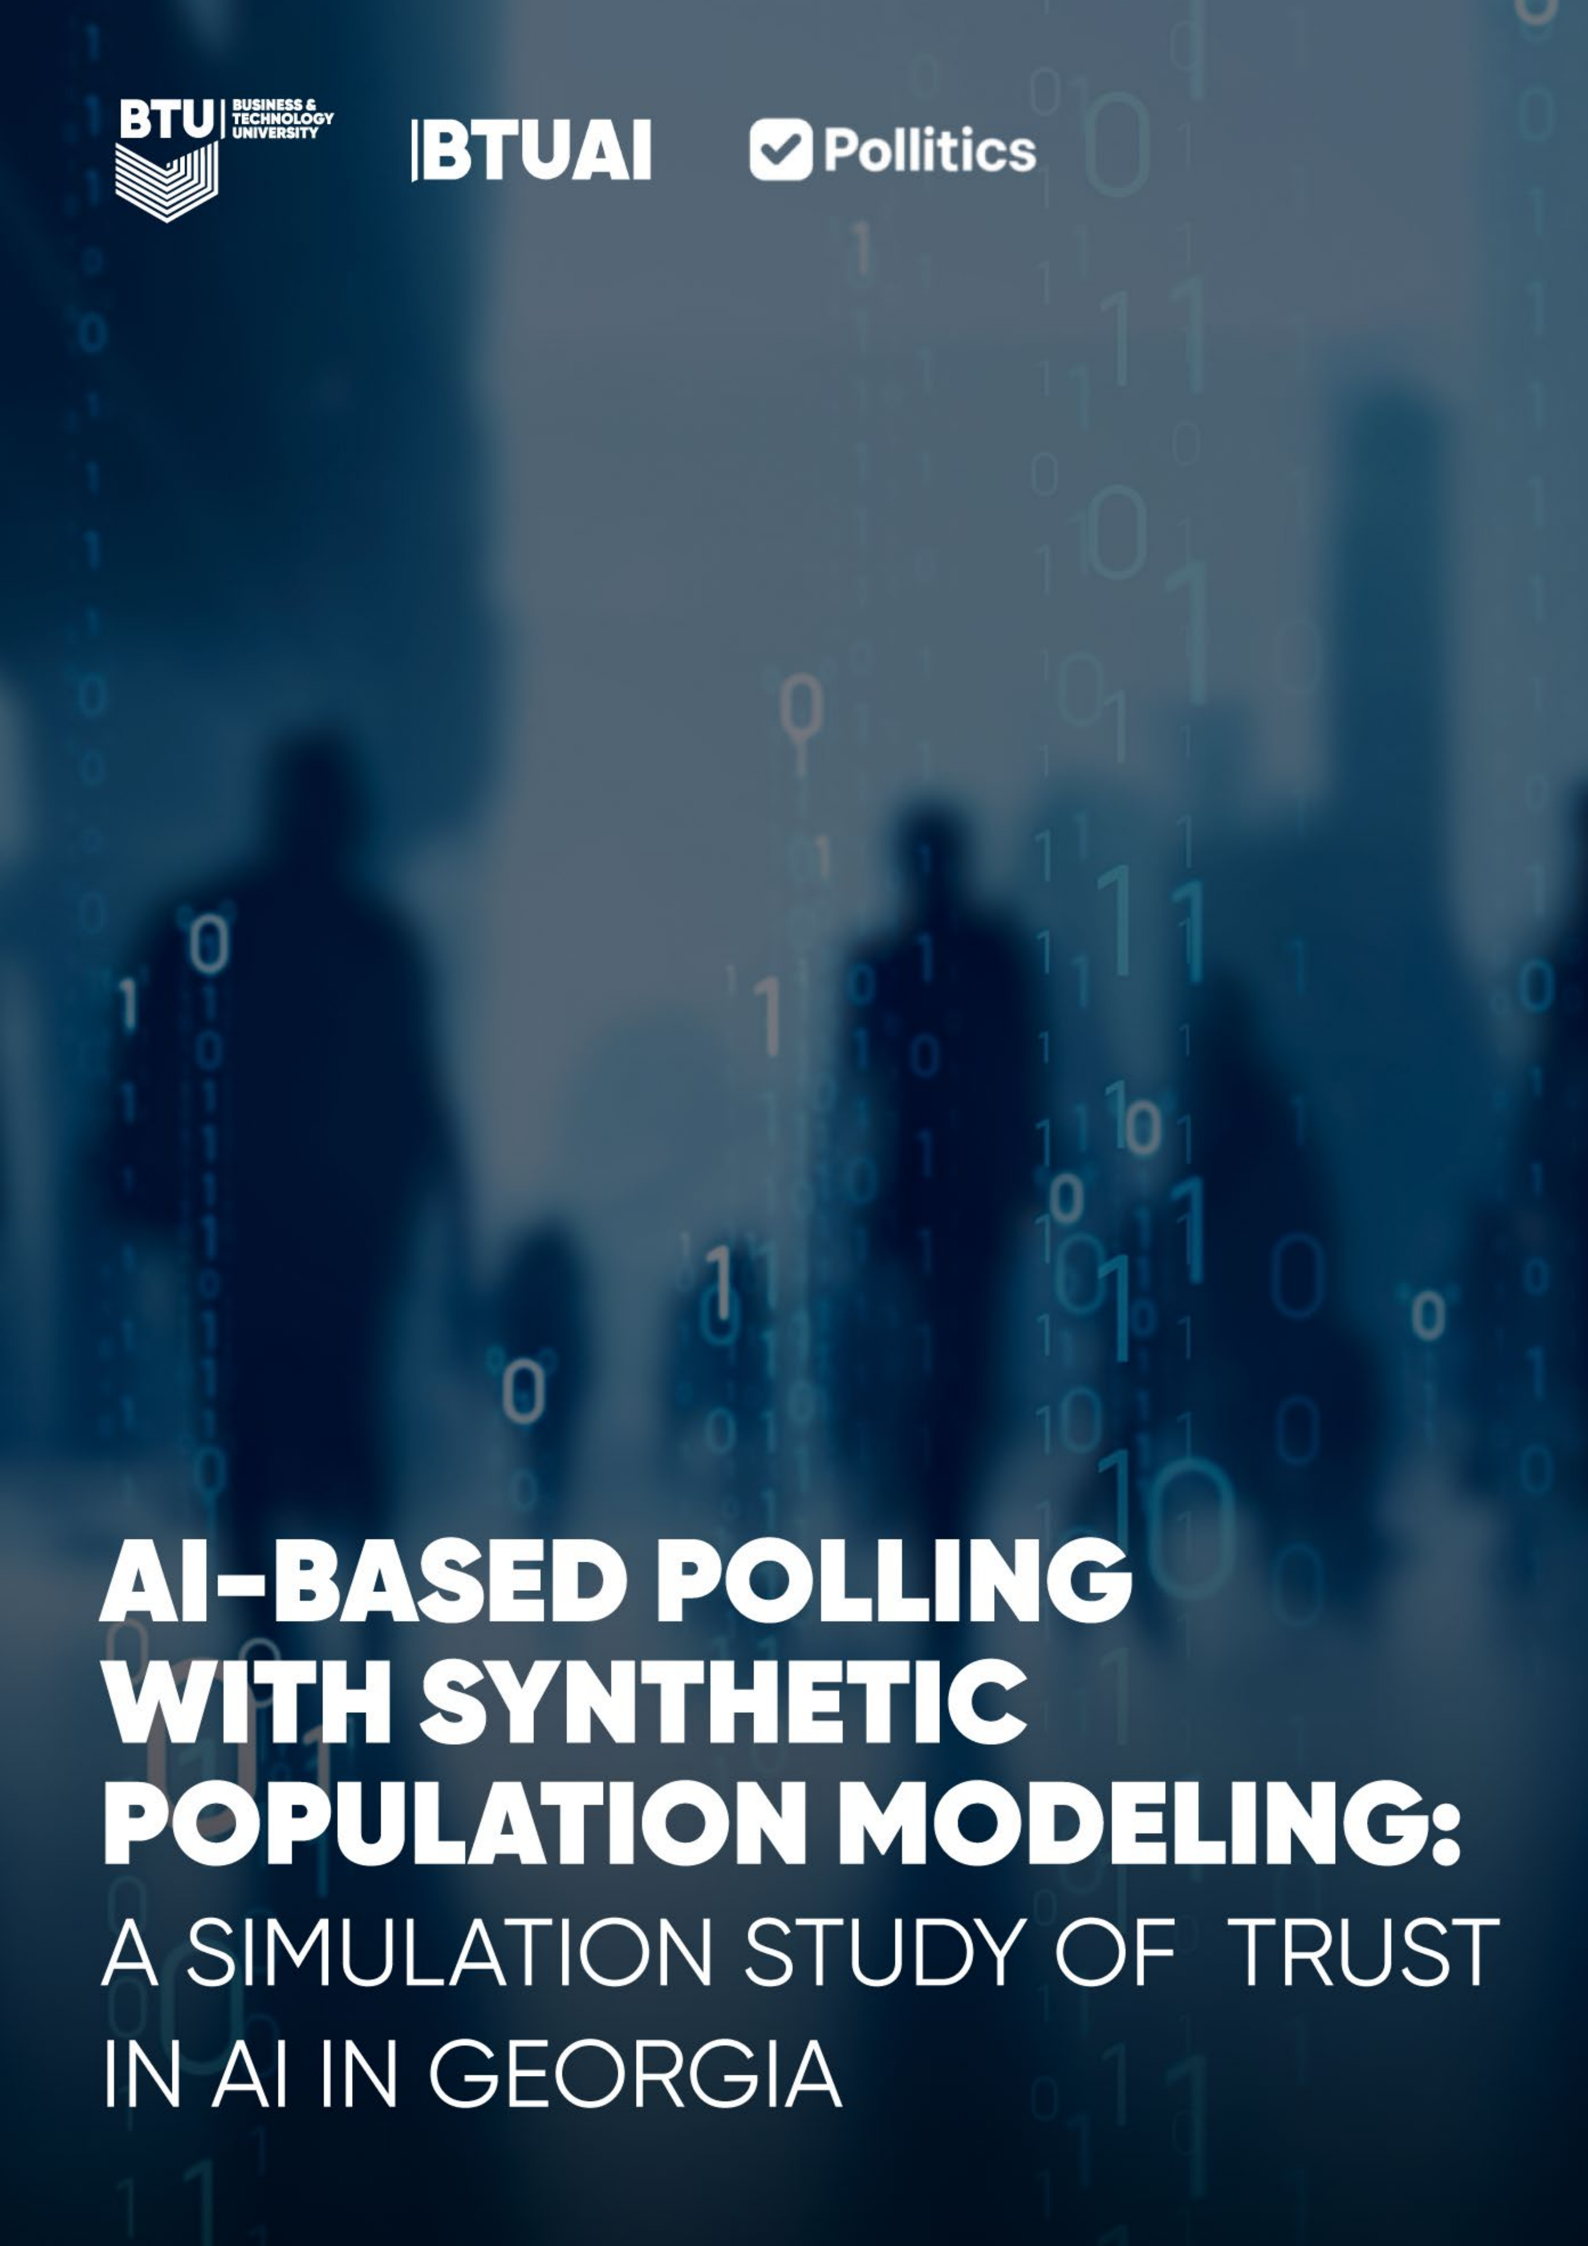

BTU BUSINESS & TECHNOLOGY UNIVERSITY

BTUAI

Pollitics

AI-BASED POLLING WITH SYNTHETIC POPULATION MODELING:

A SIMULATION STUDY OF TRUST IN AI IN GEORGIA

CONTENTS

1. INTRODUCTION

This study explores the use of AI-driven polling on a synthetic population as a new simulation layer for social research and policy analysis. Instead of relying only on traditional surveys, which are costly and slow to repeat, synthetic populations allow faster exploration, rapid iteration, and structured decision support. This approach treats AI not as a replacement for real data, but as a simulation environment where different assumptions and scenarios can be tested before investing resources in full-scale field research. Such a tool is especially valuable for small markets like Georgia, where limited budgets and population scale make repeated nationwide surveys inefficient or impractical.

The research was conducted jointly by Business and Technology University (Tbilisi, Georgia) and Pollitics (Station F, France) to demonstrate the validity and usability of this instrument in the Georgian context. Business and Technology University provided and validated official Georgian statistics for constructing the synthetic population and analyzed the simulated results in the local socio-economic context. Pollitics led the technical implementation, including model development, methodological design and execution of the AI-driven simulations.

The simulated statistics focus on a highly relevant and timely topic for which no reliable empirical data currently exist in Georgia: public trust in artificial intelligence. By applying the synthetic population method to this gap area, the study shows how AI-based simulation can provide early insights into social attitudes, support exploratory analysis, and guide future empirical research agendas tailored to Georgia's specific needs.

This report prepared for the BTUAI platform, summarizes tests run on binary questions using all provided cross-tabulated statistics for Georgia. The experiments used populations of 400 people, with gpt-4o, with and without context, and with and without Gumbel reweighting. Contexts were built primarily from Georgian sources, supplemented with English where needed.

Readers interested in exploring and testing the AI-based polling system directly may access the platform at www.pollitics.com.

2. METHODOLOGY

2.1 Population Generation

The key idea is to *build a whole synthetic population directly* and then improve it step by step. This is different from probabilistic methods (such as Bayesian SPG) that generate people by sampling from a probability model. Here, we create a complete population and enforce rules about individuals while making the overall statistics as close as possible to the real ones.

The method relies on two complementary types of constraints:

- **Local (horizontal) constraints** apply to each individual. They ensure that a single synthetic person is coherent (for example, age and education status should make sense together).
- **Global (distribution) constraints** apply to the whole population. They describe how the overall counts should match real-world statistics (for example, the percentage of women or the distribution of age groups).

The workflow can be explained in five simple steps:

1. **Collect the targets.** Gather the official statistics that describe the target population.
2. **Set up a constraint problem.** Each synthetic person is described by variables (for example age group, gender, education), and each variable has a small set of allowed values (its *domain*). Local constraints restrict which combinations of values are allowed for one person.
3. **Generate people in batches.** Instead of building the whole population at once, the method creates a batch of individuals, solves the constraints for that batch, and then updates the remaining targets for the next batch.
4. **Keep global targets in view.** After each batch, the method checks how far the current population is from the target distributions and adjusts the next batch so the overall totals stay on track.
5. **Finish when totals match.** The process stops when all batches together satisfy the local constraints and the global distributions are as close as possible to the targets.

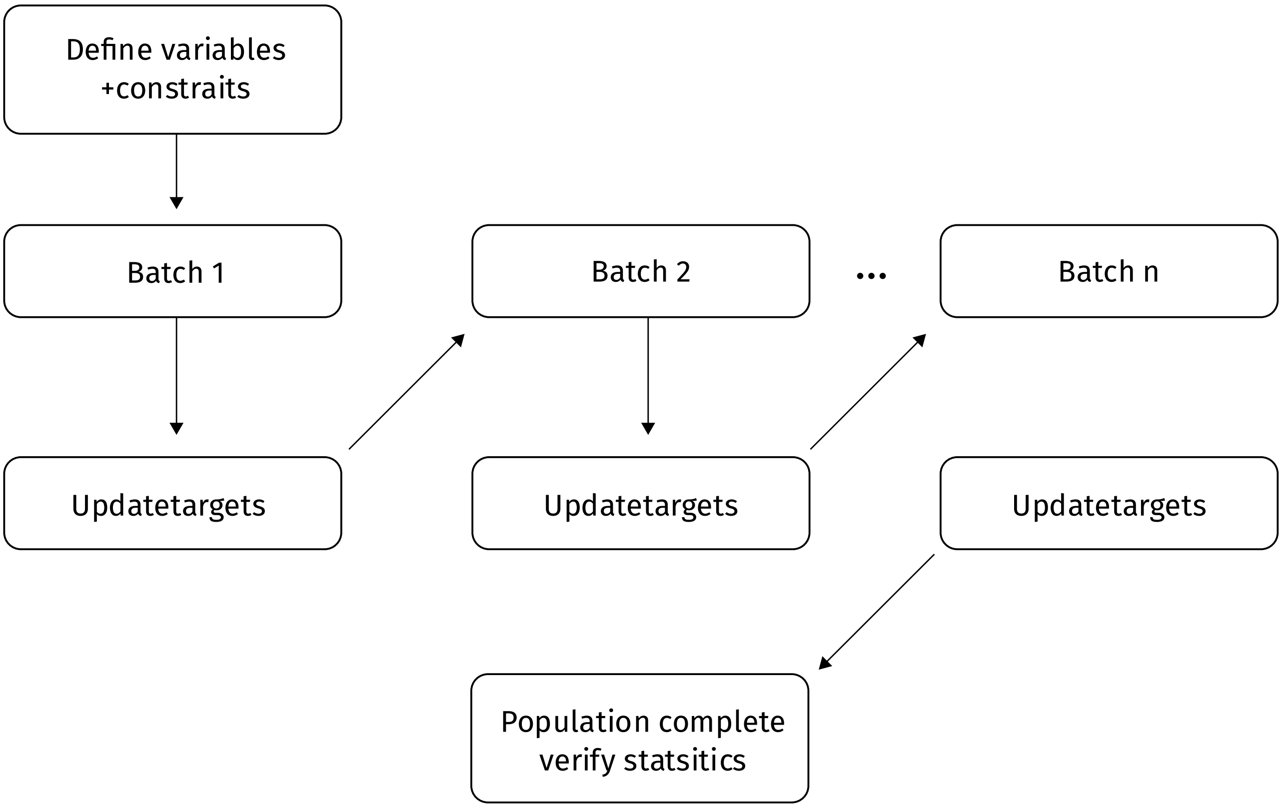

Figure 1: Constraint-based population generation with batch cycles until completion.

Cross-distributions. To enforce cross-distributions, we first normalize the percentages of distribution B with respect to the categories of distribution A, so that each joint distribution is consistent with the marginal totals. Then, for each individual, we add constraints that enforce uniqueness across the corresponding cross-distribution variables. For example, if we model *AGE*, *AGE+GENDER*, and *AGE+INCOME*, the individual must share the same *AGE* value across all three characteristics. This consistency is imposed by *table constraints* that forbid assignments where the AGE values differ across these cross-distribution attributes.

2.2 Binary Questions, Context Generation and Votes

For the binary questions, each generated persona is provided to a large language model (LLM). The persona is optionally enriched with a short *context*, and the model returns a probability distribution over the three answers: “Yes”, “No”, and “Don’t know”. These probabilities are then used to simulate individual votes and aggregate results.

The context generation step builds a concise, readable background for each persona and question. It starts with an online request that retrieves information available up to the present time from the training data of the voting LLM and, when relevant to the question, arranges that information in chronological order. The goal is to ground the response with up-to-date, question-relevant context while keeping the narrative short and focused.

The context is organized around two key characteristics selected based on the question (for example, *AGE* and *GENDER*). For every possible combination of these two characteristics, the system gathers two complementary kinds of information:

- *FACTS*: neutral, verifiable information (e.g., measured outcomes, official statistics, or documented events).
- *VIEWPOINTS*: commonly expressed perspectives or opinions associated with that group, without asserting that they are factual.

These elements are then summarized into a short narrative, with *FACTS* and *VIEWPOINTS* clearly distinguished so the model can weigh evidence and perspective appropriately.

For each persona, the LLM outputs a triplet of probabilities $(p_{\text{yes}}, p_{\text{no}}, p_{\text{dk}})$ that sum to 1. When the optional *Gumbel reweighting* is enabled, we perturb the logits with Gumbel noise to encourage variability:

$$g_i \sim \text{Gumbel}(0,1), \tilde{p}_i = \sigma(\log(p_i/(1 - p_i)) + g_i),$$

where $\sigma(\cdot)$ is the logistic function. In practice, the perturbation is applied only when p_i is not too close to 0 or 1 (a threshold is used), and the resulting $\tilde{p}_i$ values are renormalized to sum to 1. We then draw a categorical response according to the final probabilities (either p_i or $\tilde{p}_i$) to obtain a simulated answer for each persona. Finally, the answers are aggregated across the population to compute the overall distribution of responses.

3. GENERATED STATISTICS ON TRUST IN AI IN POPULATION OF GEORGIA

The results presented in this section are based on simulations using contextual information without Gumbel reweighting; more detailed comparisons across different settings are provided in the Annex. The simulated results suggest that trust in AI in Georgia is moderate and somewhat cautious rather than strongly positive. General confidence in AI systems appears relatively limited, while attitudes toward specific applications are slightly more open. There is some willingness to accept AI in defined contexts, but reservations remain, particularly regarding institutional responsibility and accountability. Uncertainty about who should bear responsibility when AI systems cause harm points to ongoing concerns about governance and oversight.

Graph 1: Proportion of the synthetic Georgian population by attitudes toward AI

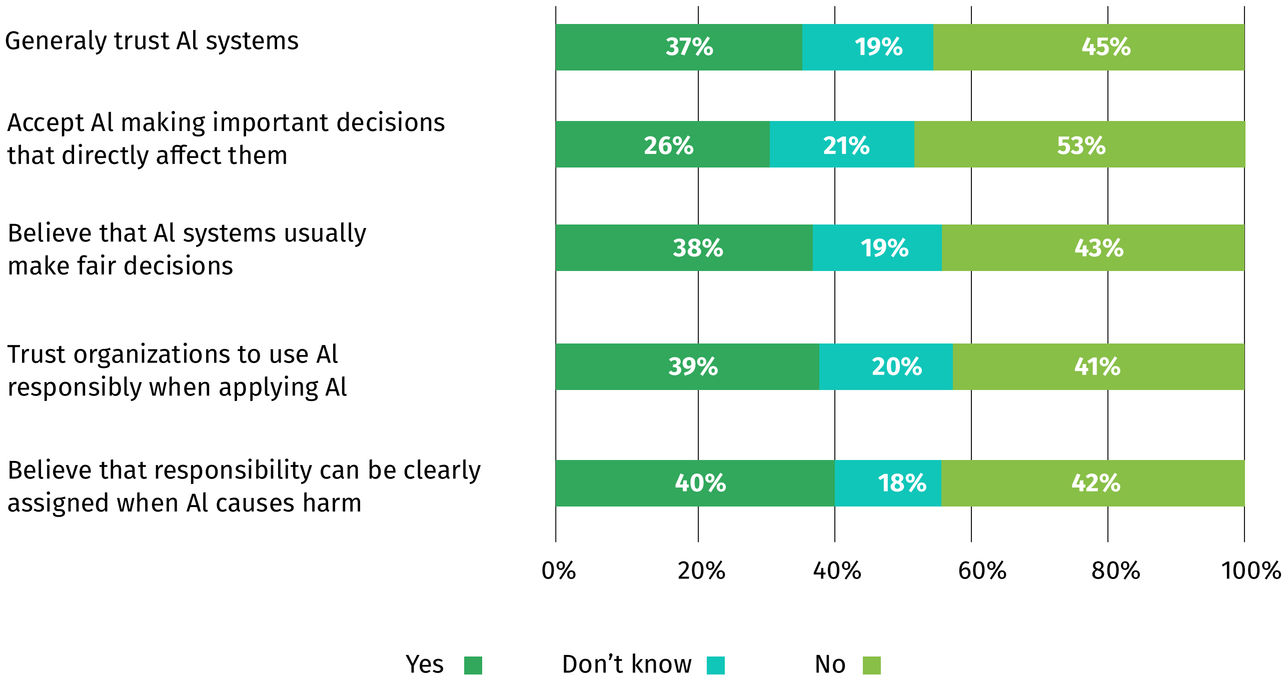

Gender differences show a nuanced pattern in attitudes toward AI. Women demonstrate slightly higher general trust in AI overall. At the same time, they are strongly represented in groups characterized by thoughtful skepticism, those who may question AI but do not reject it outright. This indicates a more critical and contextual approach rather than simple opposition. Men appear more frequently in clearer or more absolute positions, including both strongly trusting and undecided groups. Gender balance is more even in mixed positions, such as trusting AI without accepting its authority or accepting it without full trust.

Higher education is associated with more thoughtful and conditional trust rather than automatic enthusiasm, whereas lower education correlates more strongly with rejection. Urban residents, especially those living in the capital are more prevalent in trusting and positive groups, while rural populations are more concentrated among cautious or distrustful segments. Very low digital engagement strongly aligns with rejection, but beyond a basic level of digital participation, engagement alone does not determine attitudes toward AI.

Graph 2: Proportion of the Synthetic Population of Georgia with General Trust in AI by Key Demographic Characteristics

Male	33%	Yes	41%	Female
	17%	Don't know	21%	
	50%	No	38%	
with Higher Education	46%	Yes	33%	without Higher Education
	23%	Don't know	17%	
	31%	No	50%	
Urban areas	41%	Yes	29%	Rural areas
	20%	Don't know	18%	
	39%	No	54%	

AI peaks among individuals aged 30–44, where young to middle-aged professionals dominate the trusting groups, whether they fully accept AI decisions or maintain certain boundaries. The youngest group presents a more complex pattern: many accept AI without fully trusting it, reflecting a pragmatic view of AI as an inevitable reality rather than an object of strong confidence. Middle-aged individuals span the spectrum, appearing in both positive and uncertain categories. In contrast, older individuals concentrate heavily on distrusting and rejecting groups, with those in their sixties and beyond forming the core of skeptical positions.

Graph 3: Share of the Synthetic Population of Georgia with General Trust in AI by Age Group (Years)

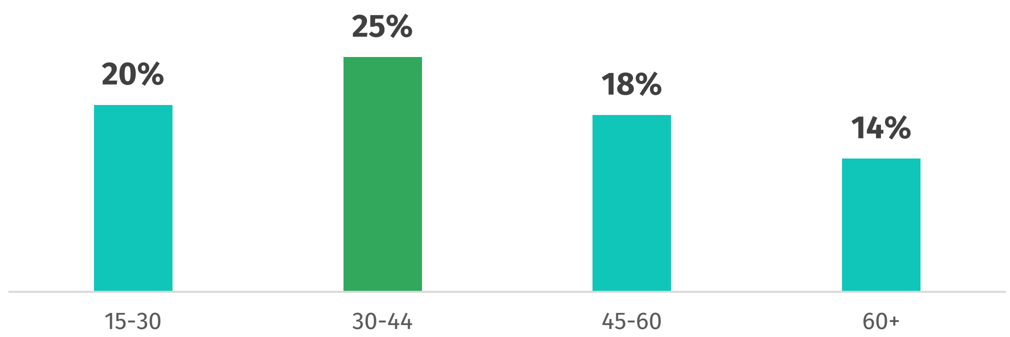

These archetypes (graph 4) were constructed by combining responses to two core dimensions: stated trust in AI systems and willingness to accept AI making important decisions. Additional response patterns as the overall number of negative answers, were used to distinguish stronger rejection from more nuanced skepticism.

Graph 4: Major Archetypes of the Synthetic Population of Georgia by Share

Enthusiasts trust AI and are willing to accept it making important decisions about them. They see AI as capable and beneficial and adopt it without major reservations. This group is mostly urban professionals between thirty and forty-five years old, with many living in the capital. The gender balance tilts slightly male. Education levels are relatively high, including higher education and vocational qualifications. Their digital engagement is moderate, suggesting their trust comes from familiarity and understanding rather than blind enthusiasm.

While enthusiasts fully embrace AI, **Confident Resisters** draw a clear boundary. They trust AI's capabilities but refuse to let it make important decisions for them. Their position reflects a deliberate choice to preserve human control despite recognizing technological effectiveness. They are typically urban professionals in their thirties and forties, with mixed educational backgrounds and moderate digital engagement. Between strong acceptance and clear resistance stand the **Hesitant** Optimists. They trust AI but are unsure about accepting its authority in significant decisions. Often young, highly educated, and digitally fluent urban professionals, they are still weighing benefits and risks before committing. Their uncertainty reflects careful consideration rather than lack of knowledge.

On the opposite end of the spectrum are the **Rejecters**, who neither trust AI nor accept it making decisions and show consistently negative responses across trust-related questions. They skew older, often in their late fifties and sixties, with lower digital engagement and more regional dispersion. Their stance represents firm and comprehensive resistance. Close to this group but less absolute are the **Skeptics**. They also distrust AI and reject its decision-making role, yet they acknowledge nuance in some areas. Often older and female-dominated, they combine moderate digital engagement with critical but differentiated views.

In contrast, the **Resigned** present an unusual combination: they do not trust AI but still accept its decisions. This reflects pragmatic acceptance or a sense of inevitability. The group spans ages, including younger respondents, and shows moderate digital use, engaging with AI despite reservations. Sharing the element of distrust but without a final position are the **Cautious Doubters**. They do not trust AI and remain unsure whether they would accept its decisions. Leaning older and more rural, their skepticism is clear, yet their practical stance remains unsettled.

Finally, the largest segment is **The Uncertain**. Unlike the other groups, they have not formed a clear opinion about trusting AI and remain divided on acceptance as well. They span ages and education levels and represent attitudes that are still developing rather than firmly established.

4. ANNEX

The table below reports percentage (%) results for each question under four settings on populations of 400 individuals: *no context*, *no context + Gumbel*, *with context* and *with context + Gumbel*. ChatGPT-4o was used both to generate and organize the context and to produce the vote probabilities. Contexts were built primarily from Georgian sources, supplemented with English where needed.

	Do you generally trust AI systems?			Do you believe AI systems usually make fair decisions?			Do you trust organizations to use AI responsibly when they apply AI?		
Mode	Yes	No	Don't know	Yes	No	Don't know	Yes	No	Don't know
No context	33.8	47.0	19.2	29.5	48.5	22.0	29.0	46.8	24.2
No context + Gumbel	31.2	45.0	23.8	29.8	46.2	24.0	28.2	47.0	24.8
With context	36.5	44.5	19.0	38.2	43.2	18.6	38.8	41.0	20.2
With context + Gumbel	36.5	43.8	19.7	41.8	36.2	22	40.0	36.0	24.0

	If an AI system causes harm, do you believe responsibility can be clearly assigned?			Would you accept AI making important decisions that affect you directly?		
Mode	Yes	No	Don't know	Yes	No	Don't know
No context	41.8	37.5	20.7	28.5	52.2	19.3
No context + Gumbel	41.5	33.0	25.5	30.8	47.5	21.7
With context	40.0	42.0	18.0	25.8	53.0	21.2
With context + Gumbel	43.2	38.5	18.3	32.2	49.5	18.3