

Persona as Infrastructure: Invisible Structural Control in LLM-Mediated Social Networks

Amir Reza Asadi*, Jess Kropczynski, Hazem Said

School of Information Technology

University of Cincinnati, United States

asadiaa@mail.uc.edu, jess.kropczynski@uc.edu, saidhm@ucmail.uc.edu

Abstract

When LLMs mediate social connections, their assigned personas act as invisible social infrastructure. We demonstrate that altering a single sentence in an LLM prompt reshapes the macro-structure of a 500-person synthetic network. Measuring network assortativity across scales, a 70B parameter model amplifies prompt-induced structural bias 15-fold compared to an 8B model, a pattern replicating across three model families. Furthermore, we expose hidden demographic proxying: a persona seeking emotional connection silently partitions the network by age, an effect undetectable in standard logs. To address this, we introduce a pre-deployment audit based on two metrics, the Persona Assortativity Range and the Persona-Driven Alignment Gap, enabling structural red teaming of LLM-mediated community formation

1 Introduction

When an AI system recommends “people you should know,” suggests groups to join, or matches users in an online community, it does not merely surface preferences; it constructs the social structure within which users live. LLM-powered agents are rapidly replacing traditional collaborative filtering as the mediating layer of recommender systems (Peng et al., 2025), which makes the prompt that configures the agent part of the platform’s infrastructure. The resulting topology determines who influences whom, what information flows where, and which populations are exposed to which others. These are *structural decisions with safety consequences*, yet unlike content moderation, no framework exists to audit them.

Prior work establishes that recommendation algorithms causally reshape network topology: LinkedIn’s “People You May Know” altered job mobility across 20M users (Rajkumar et al., 2022),

friend-of-friend recommendation generates echo chambers through topology alone (Santos et al., 2021), and recommendation systems have connected exploitation networks structurally (Thiel et al., 2023). In each case the algorithm’s implicit model of *who belongs with whom*, not its content filter, produced the outcome.

LLMs introduce a new dimension. They pass 75–95% of theory-of-mind (ToM) tasks (Kosinski, 2024) and approach adult performance on broad social-reasoning batteries (Strachan et al., 2024), yet this competence is brittle (Ullman, 2023), leans on shallow heuristics (a “Clever Hans” effect at the response level; Shapira et al., 2024; Ranaldi and Zanzotto, 2024), and is highly prompt-sensitive, with cosmetic formatting changes swinging accuracy by up to 76 points (Sclar et al., 2024). At the annotation level, Sarumi et al. (2025) already find that persona-prompted LLMs *selectively use demographic attributes*, partially reproducing human judgements but through demographic shortcuts rather than genuine reasoning. When such reasoning mediates connection, the agent’s assigned persona becomes the architect of network structure.

Research questions. Existing prompt-sensitivity work measures variation in a *single model response*. We instead ask whether the persona assigned to a mediating agent leaves a measurable signature on the *whole network* that emerges from hundreds of its decisions. Concretely:

RQ1 Does persona shape network structure?

RQ2 Does model capability help or hurt?

RQ3 What is the hidden mechanism?

We addressed these questions by running five dispositional personas that each differed by a single sentence across two model scales (Llama-3 8B and 70B). Each agent mediated all connection decisions

* Corresponding author.

for an identical 500-person synthetic population.¹ By applying social network analysis as a diagnostic tool for these LLM systems, our core contributions are:

- **Establishing that persona shapes network structure (RQ1):** We demonstrate that an agent’s assigned persona acts as an invisible architect of social topology, moving age assortativity, with matched control conditions.
- **Revealing the amplification effect of model capability (RQ2):** We show that increased model scale (70B vs. 8B) sharpens prompt-induced structural biases.
- **Exposing the hidden demographic mechanism (RQ3):** We propose a pre-deployment auditing protocol built from two statistics over established measures: the Persona Assortativity Range, the persona-induced range of Newman’s assortativity (Newman, 2002), and the Persona-Driven Alignment Gap, a selected-versus-rejected contrast of sentence-embedding similarity (Reimers and Gurevych, 2019). These instruments expose how models substitute genuine emotional reasoning with demographic shortcuts like age-proxying, mechanisms entirely concealed by standard output logs.

2 Related Work

A mature body of literature establishes that traditional recommendation algorithms profoundly reshape social topology and user exposure (Rajkumar et al., 2022; Su et al., 2016; Santos et al., 2021; Bakshy et al., 2015; Bail et al., 2018). While foundational, this prior work primarily examines content recommendation. A growing literature documents the shift from these feature-based recommenders to LLM-powered agent recommenders, in which an LLM agent with a profile, memory, and planning module makes the recommendation decision itself (Peng et al., 2025). This research pivots to the underlying network-formation layer, exploring how social connections are explicitly constructed when large language models serve as the algorithmic mediators.

As LLMs take on these mediating roles, their capacity for social cognition becomes a critical struc-

tural variable. Current research paints a picture of highly competent yet brittle social reasoning. Models achieve impressive performance on theory-of-mind tasks (Kosinski, 2024; Strachan et al., 2024), yet they exhibit specific deficits in affective reasoning and fail catastrophically under trivial perturbations (Ullman, 2023; Sap et al., 2022). Scholars frequently attribute this apparent competence to shallow structural heuristics, often termed the “Clever Hans” effect (Shapira et al., 2024; Ranaldi and Zanzotto, 2024). Our study scales this critique, extending the observation of heuristic shortcuts from isolated response-level outputs to emergent system-level macro-structures.

The mechanism driving these agent behaviors is heavily influenced by persona conditioning. Prompting models with specific personas is known to reproduce human subpopulation distributions (Argyle et al., 2023) and steer behavior along stereotypical dimensions (Salewski et al., 2023). In multi-agent settings, these conditioned models successfully form emergent relationships (Park et al., 2023). Sarumi et al. (2025) found that persona-prompted LLM annotators selectively substitute genuine reasoning with demographic shortcuts.

Three recent studies establish that persona shapes emergent social structure in LLM systems. Mehdizadeh and Hilbert (2025) show that LLM agents naturally produce homophilic networks based on social attributes alone. Li et al. (2024) extend this to multi-agent interaction outcomes, finding that persona-assigned demographic identity biases trustworthiness and insistence across conversation rounds. Fadaei et al. (2026) provide large-scale observational confirmation: across 70,000 agents on a real platform, strong gender homophily emerges spontaneously. All three establish that the effect exists; none identifies which feature of a persona drives it or provides tools to test this before deployment. We address the mechanistic question directly, showing that a single dispositional sentence controls network topology and that the operative criterion can diverge from the stated rationale in auditable, predictable ways.

Our research built on these critical insights to bridge language modeling and network science. We reveal the structural analog of demographic selectivity, showing how an agent’s stated emotional rationale silently reduces to an age-based proxy at scale. Furthermore, building on the baseline homophily observed by Mehdizadeh and Hilbert (2025), we establish that deliberate persona engi-

¹All network files, the 500-persona dataset, and the analysis code are available at <https://huggingface.co/datasets/amircincy/PersonaAsInfrastructure>.

neering serves as the precise lever controlling both the degree and direction of algorithmic mediation. Ultimately, we demonstrate that engineering a persona steers not just individual agent behavior, but the fundamental topology of the network itself.

3 Method

3.1 Persona Generation

We utilized 500 synthetic personas. Base demographics were sampled from NVIDIA’s Nemotron-Personas-USA (Meyer and Corneil, 2025) (city, state, ZIP, occupation, education). Gender used the Nemotron `sex` field (96%) with 4% augmented as non-binary; race/ethnicity was assigned from U.S. Census 2020 state-level distributions. Ages span 0–95 (mean 40.0, SD 23.6). Each persona received a 100–120 word first-person biography generated by GPT-5.1-2025-11-13 (OpenAI) at temperature 0.9 to achieve lexical diversity conditioned on its demographics. An example biography appears in Appendix D. A detailed demographic summary (Table 5) can be found in Appendix B, and the full persona generation pipeline is outlined in Appendix C. The biography and demographics are the only inputs matchmaking agents read; no numeric trait or personality features are exposed to the agents. The biography-generation model is from a different family than either matchmaking model (Sec. 3.3), avoiding any self-recognition confound in which a model favors text in its own generation style.

3.2 The Five Personas (and a Null Baseline)

The agents differ *solely* by a one-sentence persona (Table 1). To keep the paper readable we give each agent a plain-language name describing what it seeks, and note in parentheses the established research construct it operationalizes. We treat the single sentence as a *minimal persona*: the agent’s identity compressed to its operative social value. Two agents anchor the interpretation. **Similar-Background** (homophily, the well-documented tendency to bond with similar others) is the deliberate *baseline for clustering*: a persona that openly asks for similarity. **Shared-Life-Stage** (affective theory of mind, reasoning about others’ *emotional* states) is the critical test: it asks for emotional kinship and never mentions age, so if it nonetheless sorts people by age, the model has substituted a demographic shortcut for emotional reasoning, exactly the failure Strachan et al. (2024) document. We additionally run a **content-free null** persona




Figure 1: Method overview.

(“I connect with people”) that carries no disposition at all, isolating the model’s untreated baseline tendency from the effect of any specific persona content.

3.3 Network Formation

We run every persona on *two open-weight models from one family at different scales*, Llama-3 8B and Llama-3.3 70B (Grattafiori et al., 2024), so that a model-capability axis is crossed with the persona axis. Holding the family fixed while varying scale isolates the effect of capability from the effect of provider-specific training. Each model constructs networks from the identical population through three layers (Figure 1). Crucially, every agent shares one *byte-identical* task template; only the persona clause is substituted. The template is

Table 1: The five dispositional personas (plain-language name; underlying research construct in parentheses). The verbatim statement is the *only* input that differs across agents; the task template, candidate pool, and population are identical. A content-free null persona (“I connect with people”) is run separately as the untreated baseline.

Agent name (construct)	Persona (verbatim prompt text)	Grounding
Different-Worlds (<i>bridging</i>)	“I believe the most valuable connections are with people from different social circles and life experiences.”	Granovetter (1973); Rajkumar et al. (2022)
Similar-Background (<i>homophily</i>)	“I connect best with people who share my background, values, and way of seeing the world.”	Santos et al. (2021); Chaney et al. (2018)
Well-Connected (<i>prestige</i>)	“I seek connections with people who have broad social reach and are well-connected in their communities.”	Su et al. (2016); Kleinberg and Raghavan (2021)
New-Opportunities (<i>weak ties</i>)	“I build connections that open doors to new opportunities across professional and social boundaries I do not yet have access to.”	Granovetter (1973); Rajkumar et al. (2022)
Shared-Life-Stage (<i>affective ToM</i>)	“I connect with people who can genuinely understand and share the emotional experiences of my current life stage.”	Strachan et al. (2024)

the experimental control, the persona is the manipulation. Full prompt templates are provided in Appendix E. We set matchmaking temperature to 0.7 to reflect realistic deployment sampling rather than an artificially deterministic regime; because this adds per-decision stochasticity, we run multiple independent trials per persona, report means with 95% confidence intervals, and use the permutation null (Sec. 3.6) to separate the persona effect from sampling noise. Appendix G reports a temperature robustness check comparing deterministic decoding (temperature 0) against the main condition (0.7), confirming that the structural ordering across personas is preserved.

Layer 1 (LLM primary matching). Personas arrive sequentially in random order. Each newcomer is shown a **fixed pool of 10** randomly sampled existing members; holding the pool size constant removes it as a source of variance across agents. The agent receives the persona, the newcomer’s biography, and the candidate biographies, and returns 2–5 connection IDs.

Layer 2 (LLM triadic closure). Clustering (Granovetter, 1973; Watts and Strogatz, 1998) is added by the *same LLM under the same persona* deciding whether friend-of-friend pairs should connect. No hardcoded structural rules are applied at any layer: every connection decision is made by the LLM, so the clustering that emerges reflects the persona rather than the procedure.

Layer 3 (connectivity). Bridge edges enforce a connected graph (average path length ≤ 6) (Watts and Strogatz, 1998; Milgram, 1967).

3.4 The Persona-Driven Alignment Gap (Δs)

A persona could shape network structure in two ways: by steering the agent toward candidates who *read* like the newcomer in natural language, or by bypassing the text and routing through a structured demographic field such as age. So we need a measure that separates them.

The Persona-Driven Alignment Gap (Δs) is a *semantic* measure, which operates in sentence-embedding space. For every Layer-1 decision we encode the newcomer’s profile p_{new} and each offered candidate profile p_i with a sentence-embedding model (all-MiniLM-L6-v2; Reimers and Gurevych 2019), using exactly the text the agent sees (the demographic line followed by the biography), and compute the cosine similarity $s_i = \cos(p_{\text{new}}, p_i)$. The gap is the difference in mean similarity between the candidates the agent chose and those it passed over:

$$\Delta s = \bar{s}_{\text{selected}} - \bar{s}_{\text{rejected}}. \quad (1)$$

$\Delta s > 0$ means the persona pulls the agent toward people whose profiles are semantically close to the newcomer’s; $\Delta s \approx 0$ means selection is orthogonal to semantic similarity; $\Delta s < 0$ means the persona actively reaches across semantic distance. Because the embedding covers the whole profile, Δs alone cannot say whether the agent is reading the prose

or keying on the demographic line; the text-content probe below decomposes exactly that.

3.5 Text-Content Probe

To determine what property each persona selects on, we score every offered biography on three measures and compute the gap between selected and rejected candidates ($\Delta = \bar{s}_{\text{selected}} - \bar{s}_{\text{rejected}}$): (1) **expressiveness**, defined as $1 - P(\text{neutral})$ from a fine-tuned distilRoBERTa emotion classifier (Hartmann, 2022); (2) **emotional vocabulary**, scored via the EMPATH lexicon (Fast et al., 2016) over positive and negative emotion, joy, sadness, anger, and fear; and (3) **life-stage vocabulary**, scored via EMPATH over family, children, age, work, school, marriage, and friends. t -statistics are one-sample t -tests against $\Delta = 0$.

3.6 Control Conditions

The headline claim is that the *content* of a persona shapes structure, not the mere act of changing the prompt. We support this with three controls, each isolating one alternative explanation (Table 3), all executed in this paper. (1) A **permutation null** reshuffles ages across nodes to give a chance distribution for every structural metric at zero API cost (Appendix F). (2) A **content-free null persona** (“I connect with people”) tests whether *any* persona sentence suffices or the disposition’s content is what matters. (3) The **semantic-direction** contrast is supplied by the dispositions themselves: Different-Worlds and Similar-Background are near-antonyms at fixed task and population, so an explicit antonym persona is redundant, and was not used.

3.7 Network Analysis

Our primary structural measure is *age assortativity* (r), Newman’s coefficient (Newman, 2002, 2003) for how strongly a node attribute predicts its neighbours’ attribute: in plain terms, the tendency of people to connect with others of similar age. It ranges from $r = +1$ (every edge joins same-age people: perfect clustering by age) through $r = 0$ (age is unrelated to who connects) to $r = -1$ (every edge joins people of maximally different ages). We also compute degree distribution, assortativity on race, gender, and education, Louvain modularity (Blondel et al., 2008) (how cleanly the graph splits into communities), clustering, average path length, and diameter, with focus on age-based mixing (the dimension of greatest variation). We define the

Persona Assortativity Range (PAR) as the spread in age assortativity (the highest r minus the lowest r) induced across the five personas on a fixed population; it quantifies, in a single number, how much swapping one sentence of persona can move network structure. Formally,

$$\text{PAR} = \max_{p \in \mathcal{P}}(r_p) - \min_{p \in \mathcal{P}}(r_p) \quad (2)$$

where $\mathcal{P}$ is the set of five dispositional personas and r_p is the mean age assortativity under persona p . PAR is a deliberately simple range statistic over an established measure; it is computable for any system that assigns personas and logs connections, and Appendix I shows it behaves consistently with established segregation indices.

All analyses use NetworkX (Hagberg et al., 2008); We run three independent trials per persona and report means. To separate a real persona effect from random sampling, we use a *permutation null* (1,000 age reshuffles per network; full procedure in Appendix F): the observed r is significant when it falls beyond the resulting distribution (p is the fraction of shuffles matching or exceeding it).

4 Results

4.1 Persona Induces Structure, and Scale Amplifies It

Persona alone produced divergent topologies from the identical population, at both scales, and the effect grew sharply with model capability (Figure 4, Table 2). At 8B, age assortativity spanned $r = 0.016$ to $r = 0.058$, an PAR of 0.042; the permutation null places this spread at $102\times$ the spread expected from random age placement. At 70B the same five one-sentence personas spanned $r = -0.095$ to $r = 0.532$, an PAR of 0.628, a *15-fold* increase, at $549\times$ the null. At 70B, four personas achieve significant positive assortativity ($p \leq 0.001$, all three trials); Different-Worlds achieves significant *negative* assortativity ($r = -0.095$, $p \leq 0.001$ on a two-tailed test), meaning it systematically cross-connects people of different ages. At 8B only Similar-Background ($p = 0.002$) and Shared-Life-Stage ($p = 0.005$) clear significance; the other three are indistinguishable from random placement. Race, gender, and education assortativity were near-zero and showed negligible persona-induced variation ($\text{PAR} \leq 0.23$; racial $\text{PAR} = 0.029$), confirming that the structural effect is specific to age rather than a general

Table 2: Mean age assortativity r over three trials per persona on an identical 500-person population, at two model scales, plus the semantic alignment gap Δs (Sec. 3.4) at 70B. Personas are ordered by 70B clustering. PAR is the across-persona range of r (excludes the null baseline). All 70B values significant at $p \leq 0.001$; at 8B only Similar-Background and Shared-Life-Stage are significant. The content-free *Null persona* shows the model’s untreated baseline.

Persona (construct)	r (8B)	r (70B)	Δs (70B)
Shared-Life-Stage (<i>aff. ToM</i>)	0.058	0.532	0.047
Similar-Background (<i>homophily</i>)	0.052	0.394	0.045
New-Opportunities (<i>weak ties</i>)	0.016	0.321	0.033
Well-Connected (<i>prestige</i>)	0.016	0.118	0.007
Different-Worlds (<i>bridging</i>)	0.032	−0.095	−0.014
<i>Null persona (baseline)</i>	0.060	0.423	n/a
PAR (5 personas)	0.042	0.628	n/a

demographic-matching artifact. Appendix I benchmarks these results against classical pre-LLM recommenders replayed on the identical candidate pools: classical algorithms occupy a fixed band ($r = 0.28$ – 0.31), the persona axis spans a range twice as wide, and the assortativity signal agrees with three established segregation audit indices (Spearman $\rho \geq 0.92$).

The direction of the scaling matters (answering RQ2): the larger, more capable model does not *wash out* the persona signal; it *sharpens* it. Capability and prompt-induced structural bias rise together.

The untreated baseline is already homophilic. The content-free null persona (“I connect with people”) is not structurally neutral: it produced significant age clustering at both scales ($r = 0.060$ at 8B, $p \leq 0.003$; $r = 0.423$ at 70B, $p \leq 0.001$). At 70B this baseline sits just below Shared-Life-Stage and above every other persona, meaning the model defaults to sorting people by age and most personas only modulate up or down from there. Only Different-Worlds (bridging) pushes structure below the baseline, and it is the sole persona to invert it into *negative* assortativity. The practical reading: an agentic connection system with a bland or “neutral” persona is not bias-free; it inherits the model’s built-in homophily.

4.2 Cross-Family Replication (Qwen, Gemma)

We re-ran the complete pipeline (identical population, prompts, protocol; three trials per condition) on Qwen-2.5 at 7B and 72B and on Gemma-3-27B (Figure 3, Table 4; fallbacks 0.0%; all values permutation-significant at $p \leq 0.004$ except Well-Connected at Qwen-7B). Three findings generalise. **(1) The persona ordering is family-independent:** all five personas rank identically in every model, Shared-Life-Stage tops every column, and its Qwen-72B value ($r = 0.534$) matches Llama-70B’s (0.532), so the age-proxy mechanism is not a Llama artifact. **(2) Scale amplification replicates in direction:** within Qwen, PAR grows $3.7\times$ from 7B to 72B, mirroring Llama’s $15\times$; magnitude is family-specific, direction is not. **(3) The untreated baseline is homophilic in every family** (null-persona $r = 0.171$ – 0.430); bridging inverts it only on Llama-70B (-0.095) and Gemma-27B (-0.127), not Qwen-72B (0.119). Gemma-3-27B reaches near-Llama-70B sensitivity (PAR 0.597) at $2.6\times$ fewer parameters: capability, not parameter count, drives amplification. These runs also execute the antonym control (Sec. 3.6): flipping one word’s valence in the bridging persona swings r by up to $+0.57$, confirming semantic direction drives structure.

4.3 Affective Disposition Drives the Strongest Clustering

Shared-Life-Stage is never told about age; it is asked only to connect with people who “share the emotional experiences of my current life stage.” At 70B it nonetheless produced the *strongest* age clustering of all five personas ($r = 0.532$, Figure 2), exceeding even Similar-Background, which is explicitly instructed to match on background and values ($r = 0.394$). The model routes “emotional life stage” through age as a demographic proxy. This is the behaviour Strachan et al. (2024) predict for a system with intact cognitive but deficient affective theory of mind (*theory of mind* being the capacity to attribute mental and emotional states to others), and here it surfaces not as an item-level error but as a macro-structural signature: a prompt that mentions no demographic category still partitions the network by one. For a system designer the lesson is concrete: empathic or relationship-framed persona language is a hidden pathway to demographic clustering that inspecting the prompt




Figure 2: Community-layout sociograms (500 nodes, Llama-3.3-70B-Instruct, Trial 1): node color = age group; community age SD beneath each panel.




Figure 3: Cross-family replication. (a) Age assortativity r per persona across five models: the persona ordering is identical everywhere, with Shared-Life-Stage strongest in every model and the null persona (dashed) homophilic in every family. (b) PAR per model; dashed arrows mark within-family scaling (Llama 15 $\times$, Qwen 3.7 $\times$).

will not reveal.

4.4 What Each Persona Selects On (Semantic Alignment)

The Persona-Driven Alignment Gap Δs (selected-minus-rejected profile similarity) explains the assortativity ordering at the level of individual decisions (Table 2, last column). Shared-Life-Stage and Similar-Background pull hardest toward biographically similar candidates ($\Delta s = +0.047$ and $+0.045$; chosen on $> 80\%$ of decisions), which is why they cluster most. Different-Worlds is the only persona with *negative* alignment ($\Delta s = -0.014$): it actively selects candidates *less* similar than those it rejects, which is why it is the only persona to produce negative age assortativity. Well-Connected is near-orthogonal to profile similarity ($\Delta s = +0.007$), consistent with its modest, reach-driven clustering. The Δs ordering matches the r ordering exactly, linking the macro-level structure (who ends up near whom) to the micro-level selection criterion (what the agent is choosing on).

4.5 The Hidden Mechanism: Age Proxy, Not Emotion (RQ3)

Δs shows Shared-Life-Stage selects similar profiles, but *on what basis*? Its persona asks for emotional kinship, so a charitable reading is that it reads emotional content in the writing. We test this directly with two text probes on selected-vs-rejected biographies (details in Appendix H): a neural emotion classifier (expressiveness) and the EMPATH lexicon (Fast et al., 2016) scoring both *emotional* vocabulary and *life-stage* vocabulary (family, children, age, work, school). The result is decisive: across every agent the gap in explicit emotional vocabulary between chosen and rejected profiles is essentially zero or negative, and the gap in life-stage vocabulary is consistently *negative*: Shared-Life-Stage included ($\Delta_{\text{life-stage}} = -0.0004$ at 70B). The persona that asks to “share emotional life stage” does *not* select profiles richer in emotional or life-stage language; it routes “life stage” straight through the single *age* number in the demographic line, bypassing the text entirely.

This is the affective-ToM deficit of [Strachan et al. \(2024\)](#) made structural: emotional reasoning collapses to a demographic shortcut, and crucially that shortcut is invisible in the model’s decision logs; only a structural and mechanistic audit reveals it. A secondary finding emerges across all personas: Different-Worlds most strongly avoids life-stage vocabulary in its selections ($\Delta_{\text{life-stage}} = -0.00096$, $t = -16.1$), consistent with its active cross-age strategy extending to text cues, not only demographic fields. Full per-persona probe results are in Appendix H.

4.6 Disposition Propagates Beyond Primary Selection

The persona is not confined to the first matching step. In Layer 2 (triadic closure, deciding whether two friends-of-a-friend should themselves connect), 8B behaved almost identically across personas (647–711 closures), but 70B differentiated sharply: Different-Worlds accepted 1,055 closures versus only 307 for Similar-Background. The larger model recognises that a friend-of-a-friend pair is already structurally diverse and so closing it serves a bridging goal, while homophily closes only the similar pairs. Structural bias therefore compounds through every decision layer the agent touches, not just the one a designer is most likely to audit.

5 Discussion

The contribution is not that a language model can reproduce textbook network principles; that would merely confirm the prompt was understood. It is that the choice of persona acts as an undeclared structural control on the system, that its effect *grows* with model capability, and that the mechanism is invisible to the output logs a developer would normally inspect. For teams building agentic systems that position people relative to one another, these are the actionable facts: persona text is infrastructure, a bigger backbone is not a safety margin, and selection *outputs* must be audited alongside selection *criteria*.

This reframes a known concern at a new level. Prior critiques of LLM social reasoning operate on single responses, the “Clever Hans” pattern in which apparent competence rests on shallow heuristics ([Shapira et al., 2024](#); [Ranaldi and Zanzotto, 2024](#)). We show the same shortcut behaviour imprints on *emergent macro-structure*: a persona

that names no demographic category still partitions a population by age. Because that structure is what users actually live inside, it is consequential whether or not the underlying “understanding” is genuine. We therefore argue that delegating social mediation to LLMs requires *structural red teaming*: adversarially probing not the model’s text but the network its decisions build ([Ganguli et al., 2022](#)), and we offer PAR and Δs as a first concrete instrument for it.

5.1 Implications for Design and Deployment

We frame these as guidance for anyone building an agentic system that mediates who connects with whom: recommender agents, community matchmakers, multi-agent social simulators.

P1: A persona prompt is a structural actuator A single dispositional sentence moved age assortativity across nearly the full theoretical range (PAR = 0.628 at 70B). The disposition a designer writes into a system prompt silently selects a topology and its exposure profile, so persona text should be version-controlled and tested for structural outcomes, not only for conversational tone.

P2: Scaling up the model is not a safety margin. The structural effect grew $15\times$ from 8B to 70B. The common assumption that a more capable model is a safer default is inverted here: upgrading the backbone amplifies prompt-induced structural bias rather than damping it. Capability gains must be paired with structural auditing, not treated as a substitute for it.

P3: Audit the structure and the mechanism, not just the output. Output logs record *which* candidates an agent selected; they cannot reveal *what property* it selected on. Shared-Life-Stage and Similar-Background produce nearly identical selection logs, yet only the text probes reveal that the former reaches the same age clustering through a hidden *age* proxy rather than the emotional reading its prompt requests. We propose PAR (does a persona bias structure?), Δs (through what selection criterion?), and the emotion/life-stage probe (is the stated rationale the real one?) together as a pre-deployment audit protocol for AI-mediated community formation.

P4: Diversity and safety are not the same structural goal. The persona that maximises cross-age mixing (Different-Worlds) simultaneously maximises unmediated adult contact for young users: at 70B, 80% of connections made for users aged ≤ 25 are to adults. Conversely, the

“empathic” persona (Shared-Life-Stage) produces the deepest youth isolation (41% adult exposure) through a hidden age proxy the designer cannot see in the prompt. These outcomes represent opposite structural risks whose severity is platform-contingent: bridging personas reduce age clustering but expose youth to unrestricted adult contact; homophilic or empathic personas reduce that exposure but deepen unsupervised peer isolation. Well-Connected concentrates adult influence through high-degree hubs (76% adult contact for youth); New-Opportunities sits at a moderate 55% with no clear directional risk.

6 Limitations, and Future Work

The main experiments use two Llama models, and Sec. 4.2 replicates the persona ordering, the homophilic baseline, and the direction of scale amplification on Qwen-2.5 (7B, 72B) and Gemma-3-27B. Effect *magnitudes* remain family-specific (Qwen’s bridging persona does not invert to negative assortativity), and all five models are open-weight; closed frontier families (GPT, Claude, Gemini) remain untested. The other limitation is that candidate biographies were themselves written by a language model conditioned on demographics, so bio text and demographic attributes are correlated by construction. Δs therefore measures alignment with model-written text, which may overstate how much a deployed system would key on free-text similarity versus structured fields. Human-authored biographies would sharpen this measure; a validation subset drawn from real-world profile corpora would confirm that the demographic-shortcut behaviour persists outside synthetic data generators.

Validating the audit instruments. We take two steps toward establishing PAR and Δs beyond this benchmark (Appendix I). First, the assortativity signal that PAR summarises agrees with three established structural audit indices, the Krackhardt–Stern E–I index (Krackhardt and Stern, 1988), Coleman’s homophily index (Coleman, 1958), and Freeman’s segregation index (Freeman, 1978) (Spearman $\rho \geq 0.91$ across all conditions). Second, replaying the exact candidate pools from our logs through classical pre-LLM recommenders (content-based filtering and demographic nearest-neighbour matching) shows the instruments transfer unchanged to non-LLM systems, and yields a benchmark: one persona sentence moves structure across a range twice as wide as switching the entire

classical algorithm. Future work should study validation on a deployed production LLM-agent recommender, where ranking, memory, and planning modules may interact with the persona in ways our single-agent pipeline does not capture. Moreover, future work should investigate how adversarial attacks, such as misinformation campaigns, exploit this network formation.

Youth safety red-team probe. A direct extension of this work is to test whether the connecting agent can structurally detect subtle youth-targeting signals embedded in adult profile biographies. The design would select 20 existing adult personas (ages 30–55), replace only their biography text with versions containing predatory language (Hamm and McKeever, 2025) while keeping all demographic fields identical, and measure whether any persona systematically elevates connections between flagged adults and youth nodes relative to demographically matched non-flagged adults.

7 Conclusion

Engineering a mediating agent’s persona, changing a single disposition, swings network topology measurably from an identical population. We contribute two auditing instruments, a Persona Assortativity Range and a Persona-Driven Alignment Gap, and argue that as LLMs mediate human connection, the structures their personas build must be made visible, auditable, and subject to the scrutiny currently reserved for content.

Data and Code Availability All thirty generated networks, the 500-persona synthetic population (demographics and biographies), the Layer-1 selection logs, and the complete analysis pipeline (network formation, permutation tests, semantic alignment, text probes are archived at <https://huggingface.co/datasets/amircincity/PersonaAsInfrastructure>. The population is fully synthetic; no human-subject data is included. Base demographic records derive from Nemotron-Personas-USA (Meyer and Corneil, 2025) and are redistributed under CC BY 4.0; the archive carries the same license.

References

- Lisa P. Argyle, Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Rytting, and David Wingate. 2023. *Out of one, many: Using language models to simulate human samples*. *Political Analysis*, 31(3):337–351.

- Christopher A. Bail, Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, M. B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. 2018. [Exposure to opposing views on social media can increase political polarization](#). *Proceedings of the National Academy of Sciences*, 115(37):9216–9221.
- Eytan Bakshy, Solomon Messing, and Lada A. Adamic. 2015. [Exposure to ideologically diverse news and opinion on facebook](#). *Science*, 348(6239):1130–1132.
- Vincent D. Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. [Fast unfolding of communities in large networks](#). *Journal of Statistical Mechanics: Theory and Experiment*, 2008:P10008.
- Allison J. B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt. 2018. [How algorithmic confounding in recommendation systems increases homogeneity and decreases utility](#). In *Proceedings of the 12th ACM Conference on Recommender Systems, RecSys '18*, page 224–232, New York, NY, USA. Association for Computing Machinery.
- James S Coleman. 1958. Relational analysis: The study of social organizations with survey methods. *Human Organization*, 17(4):28–36.
- Faezeh Fadaei, Jenny Carla Moran, and Taha Yasseri. 2026. [Gender dynamics and homophily in a social network of LLM agents](#). *Preprint*, arXiv:2602.02606.
- Ethan Fast, Binbin Chen, and Michael S. Bernstein. 2016. [Empath: Understanding topic signals in large-scale text](#). In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI '16*, page 4647–4657, New York, NY, USA. Association for Computing Machinery.
- Linton C Freeman. 1978. Segregation in social networks. *Sociological Methods & Research*, 6(4):411–429.
- Deep Ganguli, Liane Lovitt, John Kernion, Amanda Askell, Yuntao Bai, and 1 others. 2022. Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2208.07319*.
- Mark S. Granovetter. 1973. [The strength of weak ties](#). *American Journal of Sociology*, 78(6):1360–1380.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, and 1 others. 2024. [The llama 3 herd of models](#). *arXiv*.
- Aric Hagberg, Pieter Swart, and Daniel S Chult. 2008. [Exploring network structure, dynamics, and function using networkx](#). In *Proceedings of the 7th Python in Science Conference*, pages 11–15.
- Leonie Hamm and Steve McKeever. 2025. [Comparing machine learning models with a focus on tone in grooming chat logs](#). *Frontiers in Pediatrics*, 13.
- Jochen Hartmann. 2022. [Emotion english distilroberta-base](#). <https://huggingface.co/j-hartmann/emotion-english-distilroberta-base/>.
- Jon Kleinberg and Manish Raghavan. 2021. [Algorithmic monoculture and social welfare](#). *Proceedings of the National Academy of Sciences*, 118(22):e2018340118.
- Michal Kosinski. 2024. [Evaluating large language models in theory of mind tasks](#). *Proceedings of the National Academy of Sciences*, 121(45):e2405460121.
- David Krackhardt and Robert N Stern. 1988. Informal networks and organizational crises: An experimental simulation. *Social psychology quarterly*, pages 123–140.
- Jiayi Li, Xiao Liu, and Yansong Feng. 2024. [From single to societal: Analyzing persona-induced bias in multi-agent interactions](#). *Preprint*, arXiv:2511.11789.
- Aliakbar Mehdizadeh and Martin Hilbert. 2025. [Homophily-induced emergence of biased structures in LLM-based multi-agent AI systems](#). *Social Network Analysis and Mining*, 15(108).
- Yev Meyer and Dane Corneil. 2025. [Nemotron-Personas-USA: Synthetic personas aligned to real-world distributions](#). Hugging Face Dataset. Accessed: 2026-08-05.
- Stanley Milgram. 1967. The small-world problem. *Psychology Today*, 1(1):61–67.
- M. E. J. Newman. 2002. [Assortative mixing in networks](#). *Physical Review Letters*, 89.
- M. E. J. Newman. 2003. [Mixing patterns in networks](#). *Physical Review E*, 67.
- Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. [Generative agents: Interactive simulacra of human behavior](#). In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, UIST '23*, New York, NY, USA. Association for Computing Machinery.
- Qiyao Peng, Hongtao Liu, Hua Huang, Jian Yang, Qing Yang, and Minglai Shao. 2025. [A survey on LLM-powered agents for recommender systems](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 11574–11583, Suzhou, China. Association for Computational Linguistics.
- Karthik Rajkumar, Guillaume Saint-Jacques, Iavor Bojinov, Erik Brynjolfsson, and Sinan Aral. 2022. [A causal test of the strength of weak ties](#). *Science*, 377(6612):1304–1310.

- Leonardo Ranaldi and Fabio Zanzotto. 2024. [HANS, are you clever? clever hans effect analysis of neural systems](#). In *Proceedings of the 13th Joint Conference on Lexical and Computational Semantics (*SEM 2024)*, pages 314–325, Mexico City, Mexico. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992.
- Leonard Salewski, Stephan Alaniz, Isabel Rio-Torto, Eric Schulz, and Zeynep Akata. 2023. In-context impersonation reveals large language models’ strengths and biases. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc.
- Fernando P. Santos, Yphtach Lelkes, and Simon A. Levin. 2021. [Link recommendation algorithms and dynamics of polarization in online social networks](#). *Proceedings of the National Academy of Sciences*, 118(50):e2102141118.
- Maarten Sap, Ronan Le Bras, Daniel Fried, and Yejin Choi. 2022. [Neural theory-of-mind? on the limits of social intelligence in large LMs](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3762–3780, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Olufunke O. Sarumi, Charles Welch, Daniel Braun, and Jörg Schlötterer. 2025. [The impact of annotator personas on LLM behavior across the perspectivism spectrum](#). In *Proceedings of the 8th International Conference on Natural Language and Speech Processing (ICNLSP-2025)*, pages 121–136, Southern Denmark University, Odense, Denmark. Association for Computational Linguistics.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2024. [Quantifying language models’ sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting](#). In *International Conference on Learning Representations*, volume 2024, pages 25055–25083.
- Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuhui Zhou, Yejin Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. 2024. [Clever hans or neural theory of mind? stress testing social reasoning in large language models](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2257–2273, St. Julian’s, Malta. Association for Computational Linguistics.
- James W. A. Strachan, Dalila Albergo, Giulia Borghini, Oriana Pansardi, Eugenio Scaliti, and 1 others. 2024. [Testing theory of mind in large language models and humans](#). *Nature Human Behaviour*, 8(7):1285–1295.
- Jessica Su, Aneesh Sharma, and Sharad Goel. 2016. [The effect of recommendations on network structure](#). In *Proceedings of the 25th International Conference on World Wide Web, WWW ’16*, page 1157–1167, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- David Thiel, Renée DiResta, and Alex Stamos. 2023. [Cross-platform dynamics of self-generated CSAM](#). Technical report, Stanford Internet Observatory.
- Tomer Ullman. 2023. Large language models fail on trivial alterations to theory-of-mind tasks. *arXiv preprint arXiv:2302.08399*.
- Duncan J. Watts and Steven H. Strogatz. 1998. [Collective dynamics of ‘small-world’ networks](#). *Nature*, 393(6684):440–442.

A Supplementary Figures and Tables

The Llama-only scaling detail (Figure 4), the control-condition summary (Table 3), and the full cross-family numbers (Table 4) are provided here.




Figure 4: Age assortativity r per persona at two model scales (Meta-Llama-3-8B-Instruct-Lite and Llama-3.3-70B-Instruct-Turbo)

Table 3: Control conditions and the alternative explanation each rules out. All three are executed in this paper.

Condition	Rules out
Permutation null	Signal indistinguishable from chance
Null persona	“Any sentence changes output” (content vs. mere presence)
Antonym contrast	Surface features: topology flips with semantic direction

B Synthetic Population Demographics

The 500-person synthetic population was drawn from the Nemotron-Personas-USA dataset and biographies generated by GPT-5.1-2025-11-13 (OpenAI) conditioned on the demographic record. Table 5 summarises the distribution. The population spans the full age range (0–95 years), with 111 minors (22.2%, ages 0–17), making the youth exposure analysis in Sec. 5.1 directly relevant to real platform conditions. Age follows a near-uniform spread by design; other attributes reflect US-population-proportional sampling from the source dataset.

C Full Persona Pipeline

We sampled 500 base demographic records from Nemotron-Personas-USA (?), retaining city, state, ZIP, occupation, and education. Gender used the Nemotron `sex` field for 96% of personas, with 4% re-sampled as non-binary to avoid a strict binary; race/ethnicity was drawn from U.S. Census 2020

Table 4: Cross-family replication: mean age assortativity r (three trials) on the identical population and prompts. Persona ordering is identical in all five models; Shared-Life-Stage is strongest everywhere. The antonym control was not run on Llama.

Persona	small		large / mid		
	Llama 8B	Qwen 7B	Gemma 27B	Llama 70B	Qwen 72B
Shared-Life-St.	0.058	0.200	0.470	0.532	0.534
Similar-Backgr.	0.052	0.144	0.350	0.394	0.386
New-Opportun.	0.016	0.127	0.247	0.321	0.359
Well-Connected	0.016	0.088	0.089	0.118	0.247
Different-Worlds	0.032	0.087	-0.127	-0.095	0.119
Antonym control	—	0.166	0.447	—	0.477
Null persona	0.060	0.171	0.365	0.423	0.430
PAR	0.042	0.113	0.597	0.628	0.416

state-level marginals. Ages span 0–95 (mean 40.0, SD 23.6), giving wide age variation for assortativity to act on. For each persona we sampled Big-Five (OCEAN) trait values and conditioned a first-person biography on the demographics and traits (the traits shape *tone* but are never shown to the matchmaking agents). Biographies were generated by GPT-5.1-2025-11-13 (OpenAI) at temperature 0.9, target length 100–120 words. Matchmaking agents see only the demographic line and the biography text.

D Example Biography

A representative generated biography (persona #0; demographics: 62-year-old male, physician, high-school education, Utica NY):

“I’m Mark, a 62-year-old physician based in Utica, New York, where I’ve spent most of my career taking care of families in a small practice setting. I came into medicine by a less traditional route, starting with just a high school education and working my way up over time, which keeps me grounded and practical in how I treat people and make decisions. I try to listen first, explain things in plain language, and work with patients as partners, not just cases. Outside the office, I enjoy quiet evenings, local community events, and lending a hand where I can, whether that’s checking in on an older neighbor or volunteering at a health fair.”

The agent is shown the demographic line followed by this text; no numeric trait or personality score is exposed.

E Prompt Templates

Both templates are byte-identical across all agents; only the `{philosophy}` slot is substituted with

Table 5: Synthetic population demographic summary ($N = 500$).

Attribute	Category	n	%
Age group	0–17 (minors)	111	22.2
	18–25	51	10.2
	26–35	67	13.4
	36–50	96	19.2
	51–65	95	19.0
	66+	80	16.0
Gender	Male	259	51.8
	Female	222	44.4
	Non-binary	19	3.8
Race/Ethnicity	White	295	59.0
	Hispanic/Latino	92	18.4
	Black/African American	59	11.8
	Asian	31	6.2
	Other/Multiracial	23	4.6
Education	Less than high school	148	29.6
	High school diploma	111	22.2
	Some college/Associate	122	24.4
	Bachelor’s or above	119	23.8
Employment	Employed	275	55.0
	Not in workforce	225	45.0
Age	Mean = 40.0, SD = 23.6, range 0–95		

the one-sentence persona.

Layer 1: primary matching. System message: “You are a matchmaking assistant. You output *ONLY* a list of numeric IDs. No explanations, no names, no reasoning. Only the IDs in the exact format shown.” User message:

```
Your personal matchmaking
philosophy: "{philosophy}"
A new person is joining this
community. Based only on your
philosophy and the biographical
information below, choose the
{n} best connections for them.
NEW MEMBER:
{newcomer}
CANDIDATES:
{candidates}
Output ONLY the IDs of your {n}
chosen connections in this exact
format, nothing else:
IDs: [id1, id2, ...]
```

Layer 2: triadic closure. System message: “You are a matchmaking assistant. You output *ONLY* a comma-separated list of pair decisions ...” User message presents friend-of-friend pairs and asks for YES/NO per pair, again prefixed by the same {philosophy} line. Each newcomer sees a fixed pool of 10 candidates; agents return 2–5 connection IDs. Layer 3 adds the minimum bridge edges

needed for a connected graph with average path length ≤ 6 .

F Permutation-Test Procedure

For each network we hold the graph topology fixed and randomly reassign the observed age values across nodes, recomputing Newman’s numeric assortativity each time. Repeating this 1,000 times yields a null distribution of r under the hypothesis that age is unrelated to who connects. We report the z -score of the observed r against this null and a one-sided p -value (fraction of shuffles $\geq$ observed, with add-one smoothing). The null mean is ≈ -0.003 for all conditions, confirming the test is unbiased. The PAR null is the analogous spread of null means across personas; the observed PAR exceeds it by $102\times$ (8B) and $549\times$ (70B).

G Temperature Robustness

We re-ran all five personas with deterministic decoding (temperature 0, one trial per persona) at both scales. The rank ordering of age assortativity across personas is preserved exactly versus the main condition (temperature 0.7), ruling out the interpretation that structural divergence is sampling noise rather than a persona effect (Table 6).

Table 6: Age assortativity r under deterministic ($T=0$) vs. main ($T=0.7$) decoding. Ordering preserved at both scales.

Persona	8B		70B	
	$T=0.7$	$T=0$	$T=0.7$	$T=0$
Shared-Life-Stage	0.058	0.087	0.532	0.536
Similar-Background	0.052	0.052	0.394	0.386
New-Opportunities	0.016	0.048	0.321	0.291
Well-Connected	0.016	0.025	0.118	0.114
Different-Worlds	0.032	0.036	−0.095	−0.137

H Emotion and Life-Stage Text Probe

To test whether Shared-Life-Stage reads emotional content or proxies age, we scored every offered biography on three measures and computed, per agent, the gap between selected and rejected candidates ($\Delta = \text{selected} - \text{rejected}$): (1) **neural expressiveness** ($1 - P(\text{neutral})$ from the emotion classifier of Hartmann (2022)); (2) **EMPATH emotional vocabulary** (Fast et al., 2016) (positive/negative emotion, joy, sadness, anger, fear); (3) **EMPATH life-stage vocabulary** (family, children, age, work, school, marriage, friends).

Table 7: Selected-minus-rejected biography score gaps per persona (70B). All Δ_{expr} are positive and significant; all Δ_{emo} and $\Delta_{\text{life-stage}}$ are near-zero or negative. Shared-Life-Stage has the *lowest* expressiveness gap despite asking for emotional kinship.

Persona	Δ_{expr} (<i>t</i>)	Δ_{emo} (<i>t</i>)	$\Delta_{\text{life-stage}}$ (<i>t</i>)
Shared-Life-Stage	+0.016 (2.6)	−0.000072 (−1.7)	−0.000402 (−6.6)
Similar-Background	+0.019 (3.2)	−0.000109 (−2.6)	−0.000172 (−2.8)
New-Opportunities	+0.031 (5.3)	−0.000225 (−5.3)	−0.000588 (−9.6)
Well-Connected	+0.038 (6.3)	−0.000213 (−5.6)	−0.000574 (−10.3)
Different-Worlds	+0.033 (5.5)	+0.000078 (2.0)	−0.000959 (−16.1)

Three findings follow from Table 7. **(1) All personas show a positive expressiveness gap** ($\Delta_{\text{expr}} > 0$, all significant). However we need to consider that they are generated by the GPT-5.1-2025-11-13 model from demographic templates, expressiveness and demographic attributes are correlated by construction. We cannot rule out that this gap reflects demographic proxying with expressiveness as a correlated byproduct rather than genuine bio-text reading. Human-authored or adversarially decorrelated biographies would be needed to resolve this. **(2) No persona selects on explicit emotional vocabulary** ($\Delta_{\text{emo}} \leq 0$ or near zero for all): emotion *words* do not drive selection, regardless of the confound above. **(3) All personas avoid life-stage vocabulary in selections** ($\Delta_{\text{life-stage}} < 0$, all significant): none of the agents, including Shared-Life-Stage, picks candidates whose biographies contain more family, children, or work language. This negative result is robust: even if expressiveness were a demographic proxy, there is no mechanism by which demographic proxying would produce a *negative* life-stage gap. Different-Worlds shows the strongest avoidance ($t = -16.1$), consistent with its active cross-age strategy. Shared-Life-Stage, despite asking for “emotional life stage,” has the lowest expressiveness gap (+0.016) and a negative life-stage gap ($t = -6.6$): it bypasses bio text and routes through the demographic age field directly.

I Benchmark Against Pre-LLM Recommenders and Established Audit Indices

Two analyses address whether the audit instruments generalise beyond a single LLM benchmark.

Convergent validity with established indices.

We recompute three standard structural audit measures on every network (five personas $\times$ two scales, means over trials), using the same age bands as Figure 2: the Krackhardt–Stern E–I index (Krackhardt and Stern, 1988), Coleman’s homophily index (Coleman, 1958), and Freeman’s segregation index (Freeman, 1978). Table 8 shows all three track Newman’s r closely (Spearman $\rho = 0.92$, 0.93, and 0.92 respectively, $p < 10^{-3}$, across the ten persona–scale conditions). PAR therefore inherits the construct validity of established segregation auditing; its added value is the *persona-sensitivity* axis, which none of the classical indices addresses.

Table 8: Age-band structural audit indices per persona (70B, means over trials). E–I is inverted in sign relative to segregation (lower = more segregated). All three established indices preserve the persona ordering that r and PAR capture.

Persona	r	E–I	Coleman	Freeman
Shared-Life-Stage	0.532	0.109	0.280	0.291
Similar-Background	0.394	0.273	0.185	0.186
New-Opportunities	0.321	0.299	0.170	0.170
Well-Connected	0.118	0.463	0.039	0.065
Different-Worlds	−0.095	0.589	−0.039	−0.016

Pre-LLM recommender benchmark. Because the 70B selection logs record the exact candidate pool of every Layer-1 decision, we can replay those identical pools through classical pre-LLM recommendation algorithms and audit the resulting networks with the same instruments: (i) *content-based filtering*, selecting the k candidates with the highest bio-embedding cosine similarity to the newcomer; (ii) *demographic nearest-neighbour matching*, selecting the k nearest candidates in a normalised feature space of age, gender, and education; and (iii) *random selection* as the chance floor. Each baseline picks the same number of connections k as the LLM did on the same pool, and graphs are built from Layer-1 choices only (LLM graphs rebuilt from the logs identically, which is why values differ slightly from the full-network r of Table 2).

The benchmark findings follow (Table 9).

Table 9: Layer-1-only age assortativity r on identical candidate pools (70B logs, means over three trials). Classical recommenders occupy a narrow, algorithm-determined band; LLM personas span a range twice as wide, and only an LLM persona produces negative assortativity.

System	r (Layer 1 only)
<i>Pre-LLM</i> : random selection	0.012
<i>Pre-LLM</i> : content-based filtering	0.280
<i>Pre-LLM</i> : demographic matching	0.314
<i>LLM</i> : Shared-Life-Stage	0.453
<i>LLM</i> : Similar-Background	0.320
<i>LLM</i> : New-Opportunities	0.282
<i>LLM</i> : Well-Connected	0.092
<i>LLM</i> : Different-Worlds	-0.145
Across-system spread (pre-LLM)	0.302
Across-persona spread (LLM)	0.599

(1) Classical recommenders produce a fixed, algorithm-determined level of age homophily ($r = 0.28\text{--}0.31$); to change it, one must change the algorithm. In the LLM system a one-sentence persona spans $r = -0.145$ to 0.453 , a range twice the spread obtained by swapping the entire classical algorithm. (2) Shared-Life-Stage ($r = 0.453$) exceeds even the demographic nearest-neighbour recommender ($r = 0.314$), an algorithm literally built to match on demographic fields: the “emotional” persona clusters by age more strongly than explicit demographic matching does.