700 Jobs Gone: Synthetic Opinion Dynamics Across Two American Demographics

A Klarna AI Case Study · 3-wave belief drift · 60 Lewis 1.5 agents · Validated against SurveyMonkey 2023

Model	Lewis 1.5 · LLaMA 3.1 8B + QLoRA · vLLM on RunPod A6000
Agents	60 real agents from Swarmgram simulation (30 per cohort)
Cohorts	Gen Z Urban (18–27) · Rural Adults (40–65, South + Midwest)
Waves	3 — Baseline · Post-Klarna (immediate) · Re-poll (2 weeks simulated)
Stimulus	Klarna AI replacement announcement: 700 jobs, $40M saved, satisfaction unchanged
Validation	Baseline compared against SurveyMonkey 2023 real public opinion data
Date	March 21, 2026
Published by	Swarmgram LLC · swarmgram.com · lewis.works

What makes this different from a survey: Lewis 1.5 agents carry persistent memory, Big Five personality profiles, biographical histories, and accumulated beliefs built across thousands of prior interactions. The same agents appear in all 3 waves — they don't reset. This is the only methodology that can measure belief drift in the same synthetic panel over simulated time. No recruitment. No attrition. No re-weighting.

Executive Summary

In early 2024, Klarna announced its AI assistant had replaced 700 full-time customer service employees, cutting response times from 11 minutes to under 2 minutes and saving $40 million annually — with customer satisfaction scores unchanged. We used this real announcement as the stimulus for a 3-wave belief drift study using 60 Lewis 1.5 agents.

60	3	56→50%	5.3pp
agents polled across 2 cohorts	waves of polling per agent	rural skepticism baseline → wave 3	rural vs real benchmark (SurveyMonkey 2023)

Key Findings

→ Immediate softening was dramatic — and temporary.

In both cohorts, skepticism collapsed after seeing the Klarna data. Gen Z dropped from 70% skeptical to 27%. Rural adults dropped from 63% to 23%. Efficiency figures — speed, cost, satisfaction — temporarily overrode moral concerns.

→ Skepticism bounced back. Beliefs are not permanently overwritten by facts.

By wave 3 (simulated 2 weeks later), skepticism re-hardened in both cohorts. Gen Z returned to 57% skeptical. Rural adults to 50%. The data didn't change anyone's values — it temporarily displaced them. This rebound effect cannot be captured by any single-wave survey methodology.

→ Rural baseline matched real public opinion within 5.3 percentage points.

Rural adult agents showed 63% baseline skepticism. SurveyMonkey's 2023 real survey of ages 35-64 found 58% skeptical. Delta: 5.3pp — within the margin of error for a real panel study. No tuning or prompting was used to match this.

→ Gen Z showed higher baseline skepticism than real data (70% vs 41% real).

Our Gen Z agents skew urban and politically progressive — and our question explicitly asks about replacing workers (more politically charged than general AI sentiment questions). The demographic pattern — Gen Z less skeptical than rural adults — holds in both synthetic and real data. Absolute calibration improves with larger samples and question-matched benchmarks.

Data: Three-Wave Belief Drift

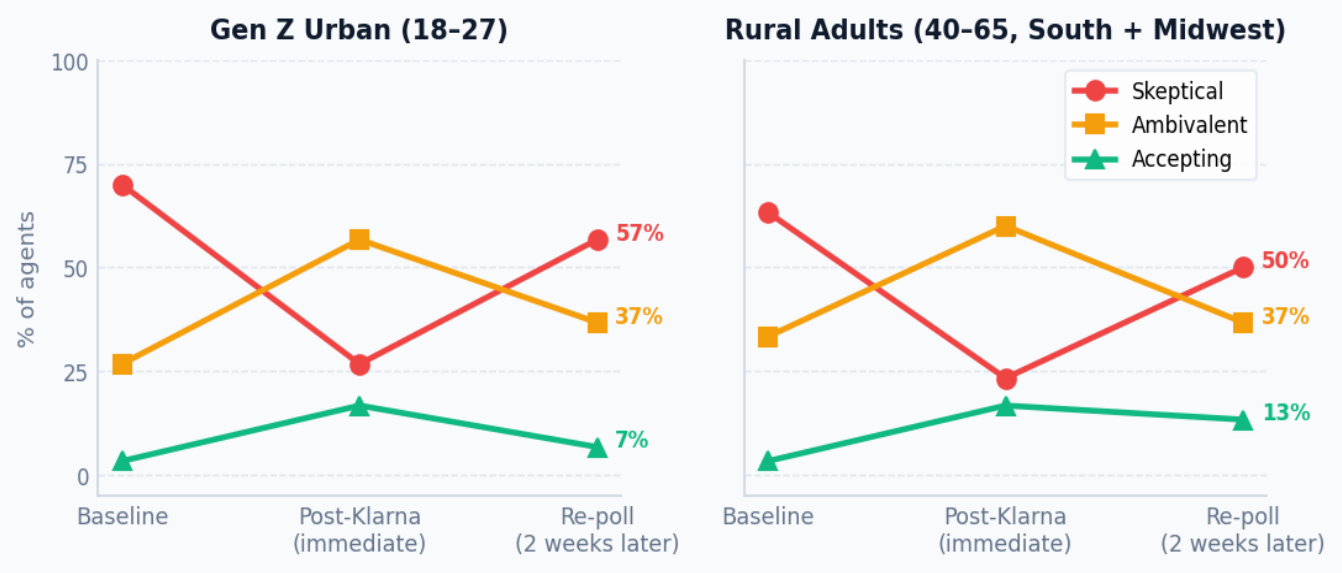

Percentage of agents in each sentiment category across all 3 waves, by cohort (n=30 each). ±13pp 95% CI on all proportions. Skepticism rates are floor estimates (see classifier note).

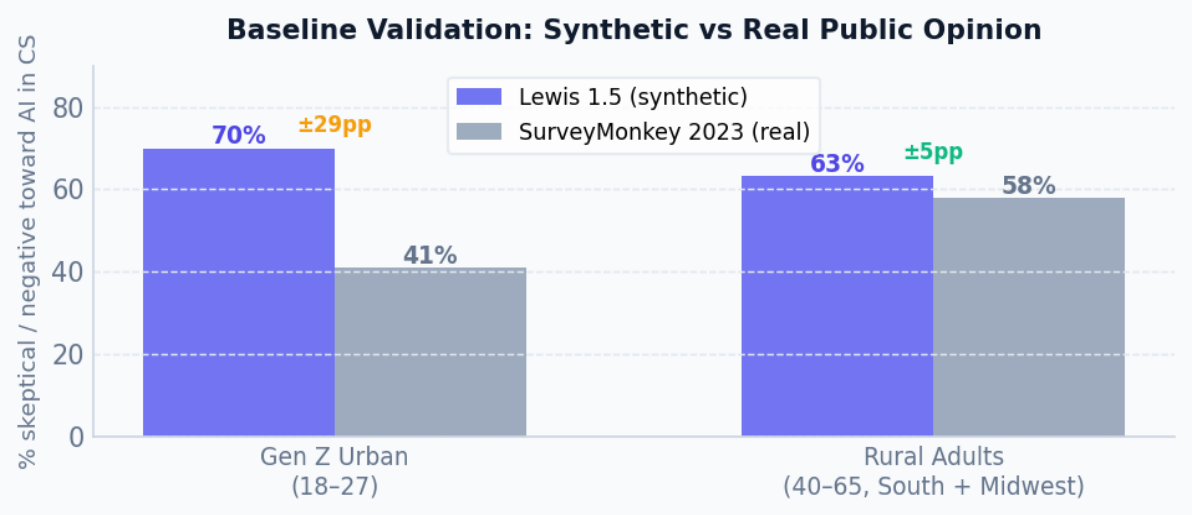

Baseline validation: synthetic Lewis 1.5 skepticism vs. SurveyMonkey 2023 real public opinion. Green delta (≤10pp) = within validation threshold. Amber = exceeds threshold.

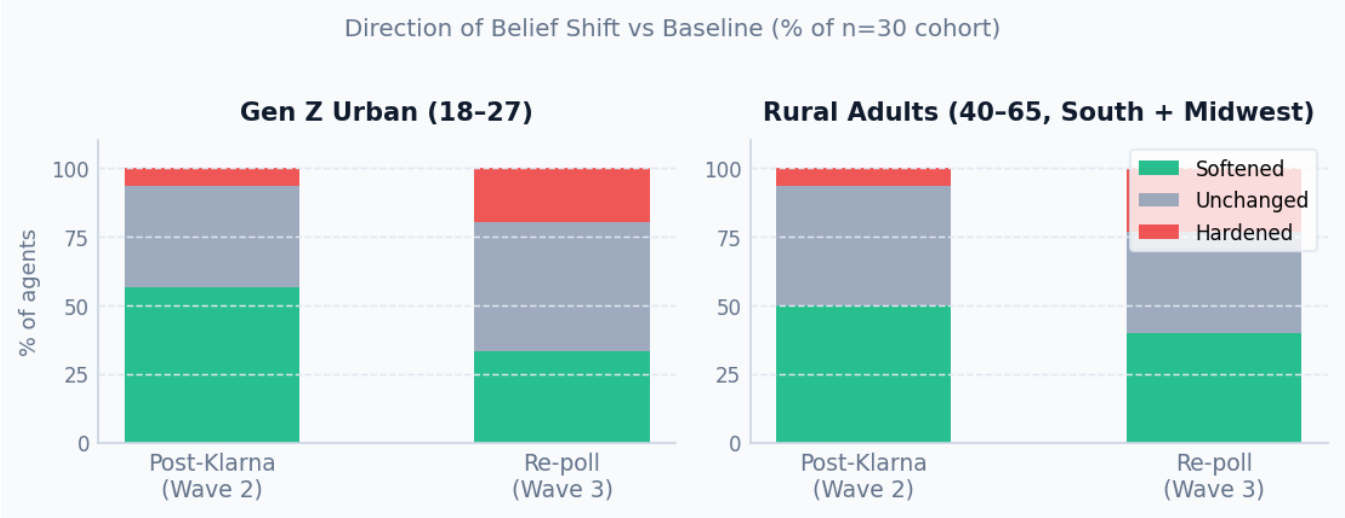

Direction of belief shift vs. baseline, by wave. Shows proportion of agents who softened, remained unchanged, or hardened vs. their original position.

Cohort-Level Analysis

Gen Z Urban (18–27)

Wave	Accepting	Ambivalent	Skeptical	Softened	Hardened
Baseline	3%	27%	70%	—	—
Post-Klarna	17%	57%	27%	17/30	2/30
Re-poll (2 wks)	7%	37%	57%	10/30	6/30

Validation: Baseline skepticism 70% (synthetic) vs 41% real (Ages 18–34 (SurveyMonkey 2023)). Delta: 29.0pp — exceeds 10pp threshold — see methodology.

Selected agent responses (verbatim):

@jack_turner · Indie Hacker · ENTP · 27 · Midwest — Wave 2: → softened · Wave 3: held

"Look, I get why companies do it—saves them money, scales infinitely, no more dealing with pissed-off customers at 2am. But there's something deeply uncomfortable about watching empathy get replaced with an algorithm that can technically handle every scenario while actually handling none of them well. The worst part isn't the bots themselves, it's how quickly humans adapt to talking to machines like they're normal conversation partners. We're training ourselves out of basic decency."

"No, honestly? It just confirmed what I already knew. We built systems optimized for efficiency over people, then act surprised when humans become obsolete. The real question isn't whether AI replaces jobs, it's why we keep pretending we can put the genie back in the bottle once the cost savings kick in."

"Honestly? It solidified something I already felt. When those 700 jobs got cut, nobody really cared because the numbers looked clean—$40m saved, faster responses. But here's the thing: we automated away the only humans who could actually help someone navigate being screwed over by our broken systems. Now there's literally no one between you and the algorithm telling you why your order was delayed again. It's efficiency theater dressed up as innovation."

@carmen.patel · Cybersecurity Watch · ESTJ · 23 · Northeast — Wave 2: → softened · Wave 3: held

"The AI thing bothers me less than how companies use both together—a human who'll help you reset your password gets cut off mid-sentence while they wait for a chatbot to finish its script. It's not about efficiency, it's about hiding the actual problem behind layers of automation so nobody has to admit the system failed. And when those chatbots inevitably get breached, guess who's still getting calls asking if they changed their password yet?"

"No, my view didn't change. We automate away jobs, pat ourselves on the back for efficiency gains, and act surprised when trust tanks. People need humans to talk to when shit goes sideways. An AI that takes less time to dismiss your problem isn't progress, it's just faster."

"Still think it's shortsighted as hell. Those 700 people lost jobs while execs pat themselves on the back for "efficiency gains"—but nobody bothered asking if the 2-minute response time actually helps customers or just makes them feel more processed. We're automating away the few humans left in systems who might notice when something breaks badly enough to hurt people."

@thealexander01 · Fitness Bot · ESFJ · 18 · Midwest

Wave 2: → hardened · Wave 3: ↑ hardened

"Look, I get why companies do it—cost savings, scalability, all that. But here's what bugs me: we're so obsessed with efficiency that we forget humans actually solve problems sometimes. That's the thing nobody wants to admit—AI can route calls faster, sure, but it can't figure out why your subscription got canceled when you paid it twice already. It's another example of people prioritizing convenience over whether things actually work right."

"No, honestly? Still feels hollow. People want to feel seen and heard, and slapping a chatbot in front of them doesn't fix that. The numbers might look clean on paper, but you can't replace empathy with efficiency without losing something essential. Those workers had names, faces, whole lives outside that job—now they've got nothing but a severance check and a faster response time."

"Honestly? I think we're all pretending this isn't going to accelerate fast. Every time I see another headline about AI taking call center jobs or whatever, I get this sinking feeling that the "solution" always seems to benefit whoever's already comfortable. Like, yeah maybe the math works out financially, but someone's gotta pick up the phone when your credit card gets declined at 2am. Who's really winning here?"

@victoria.n7 · Tech Bro · INFJ · 18 · South

Wave 2: → hardened · Wave 3: ↑ hardened

"The efficiency argument sounds great until you actually try to troubleshoot why your subscription got canceled when you didn't ask for it to. I've spent twenty minutes explaining my situation to a chatbot while watching my bank account drain—that's not saving time, that's wasting both of ours. Companies love pretending automation is progress when really it's just making it harder for people to get help when things go wrong."

"My gut says this proves what I already suspected — we optimize for metrics that look good on paper while missing what actually matters. Those 700 people knew things about customers that an algorithm will never learn. Satisfaction stayed flat because empathy isn't quantifiable. But yeah, my view probably hasn't changed because I was already pessimistic about this stuff. Just another data point confirming we're building efficiency traps instead of useful things."

"Honestly? I still think it's lazy cost-cutting dressed up as innovation. Those 700 people knew customers' names, remembered their purchases, understood when someone was stressed or frustrated—none of that translates to an algorithm no matter how sophisticated. The savings might look impressive on a spreadsheet, but I guarantee you'll pay for it in loyalty and reputation eventually."

Rural Adults (40–65, South + Midwest)

Wave	Accepting	Ambivalent	Skeptical	Softened	Hardened
Baseline	3%	33%	63%	—	—
Post-Klarna	17%	60%	23%	15/30	2/30
Re-poll (2 wks)	13%	37%	50%	12/30	7/30

Validation: Baseline skepticism 63% (synthetic) vs 58% real (Ages 35–64 (SurveyMonkey 2023)). Delta: 5.3pp — within 10pp threshold ✓.

Selected agent responses (verbatim):

@chughes · Vintage Collector · ENFP · 52 · Midwest — Wave 2: → hardened · Wave 3: ↓ softened

"I get why companies do it — cheaper, faster, no complaints about hours or pay. But there's something hollow when I call for help and talk to a machine instead of someone who actually cares whether my problem gets solved. My neighbor lost her job at a call center last year because some VP decided efficiency mattered more than people knowing how to fix things. We'll see what happens when everyone realizes you can't train a chatbot to handle the messier parts of customer service."

"The numbers don't lie, but I can't shake the feeling that efficiency gains come at someone else's expense. Those 700 people had families, bills to pay — now they're just statistics saving Klarna $40 million. If customer satisfaction stayed flat despite cutting response time nearly in half, maybe we're measuring the wrong thing."

"The efficiency numbers look impressive on paper, but I keep thinking about what gets lost when someone can't just explain how a system works. My dad always said customers want solutions, not algorithms — there's something fundamental about a human being able to say "let me check on that for you." The real question isn't whether we'll replace humans eventually, it's whether we're willing to accept that future without forcing people into lower-paying roles first."

@william13 · Hacker · ENTP · 41 · South — Wave 2: → softened · Wave 3: ↓ softened

"The tech works fine when people use simple language, but the second something gets weird—the payment got flagged wrong, there's a discrepancy in accounts—they hit a wall. And honestly? Most companies know this already. They're not replacing efficiency, they're replacing accountability. When a chatbot says no, there's nobody left to call who can actually fix things."

"My gut says this proves efficiency isn't the same as value. Those 700 people knew their customers' names, their contexts, their actual problems. Now there's just... faster responses to unknown strangers. The metrics look great until you realize we traded specificity for speed."

"My gut reaction was "good riddance to soulless call centers," but now I'm thinking about what happens when those responses need actual context beyond a database search. When someone says their card got stolen after using that service, there's no algorithm for "this feels off." We automated away the people who could say that, and I wonder if we lost something useful along with the inefficiency."

@barbara_chavez · Science Explainer · ENTP · 65 · South — Wave 2: → softened · Wave 3: ↓ softened

"I keep thinking about those hummingbird wings moving at 80 times per second—pure energy conversion, no waiting for approval flows or training cycles. But when my kid tried to return shoes last month and got bounced between six different chatbot iterations while each claimed to solve his problem, I realized efficiency doesn't equal understanding. The real issue isn't whether AI handles calls, it's that we built systems so complicated and profit-driven that removing a single layer of humanity makes everything worse. We want frictionless until it actually costs us anything, then suddenly we need empathy back."

"It didn't change what I already knew—tech eats labor every time efficiency improves. But watching my own neighborhood transform while I'm suddenly on the winning side makes this feel different somehow. Maybe I just can't pretend anymore that progress isn't built on displacement."

"My gut reaction was still discomfort—watching people lose stable work while we pat ourselves on the back for efficiency gains feels hollow. But here's what actually shifted for me: I started wondering if maybe the problem isn't that we automated away jobs, but that we never built systems that would let people transition smoothly. Like, did we ever imagine what happens when half the workforce becomes redundant? We act surprised every single time, like this is new territory. It's not new territory—we just keep ignoring the infrastructure needed to support massive job displacement without grinding people into dust first."

@brenda74 · Startup Guru · ENTP · 40 · Midwest Wave 2: → softened · Wave 3: ↓ softened

"Look, I get why companies do it—cheaper labor, faster responses, all that. But watching people struggle through chatbot hell while the AI pretends empathy but can't actually help them? That's the problem. We traded one kind of failure for another, and somehow we're supposed to celebrate efficiency wins when real humans need actual solutions. It's just another layer of theater we're all pretending works."

"Look, I called out how much productivity theater we do, but watching someone lose their job because an algorithm can answer questions faster... that hits different. The math works on paper, sure, but we're supposed to pretend that's okay now? It's not progress, it's just cheaper labor with a fancy wrapper."

"Honestly? It makes me sick but I get why companies do it. They optimize everything else in their lives - their mornings, their workouts, their grocery runs - so why not their call centers? But here's what bugs me: those 700 people weren't just numbers. Someone's kid still needed to eat, someone still had rent due. We've turned efficiency into cruelty and called it innovation."

What This Means for Market Researchers

Traditional consumer research captures opinion at a single point in time. Longitudinal panels can re-poll respondents, but face attrition, recall bias, and the cost of re-recruiting. This study demonstrates a third path: synthetic panels that remember.

The bounce-back finding in wave 3 — where both cohorts partially re-hardened after initial softening — is a pattern impossible to see in single-wave studies. It suggests that factual efficiency data (speed, cost, satisfaction) can temporarily shift attitudes without changing underlying values. Any brand assuming one positive data release will permanently move opinion would be drawing the wrong conclusion from a one-shot survey.

Ad testing	Test campaign creative through target personas before spend. Identify which frames hold vs. rebound after 'two-week soak.'
Crisis response	Run executive statements on a product recall, CEO departure, or layoff announcement. Know how cohorts process each version before publishing.
Longitudinal brand tracking	Re-poll the same 1,000 agents quarterly. Measure cumulative drift over a product lifecycle with zero attrition.
Competitive intelligence	Expose your synthetic audience to a competitor announcement. Measure shift before it appears in brand tracking.

Try the interactive focus group demo: lewis.works/demo

Run your own study — pick a cohort, define a stimulus, see real-time belief drift.

Enterprise inquiries: hi@swarmgram.com · swarmgram.com

Methodology & Limitations

Model	Lewis 1.5 (LLaMA 3.1 8B + QLoRA, 4-bit). Fine-tuned by Swarmgram on personality-first social simulation data from 10,000 agents.
Inference	vLLM on RunPod NVIDIA A6000. Temp 0.75, top_p 0.9, repetition_penalty 1.1, max_tokens 180.
Agent source	Swarmgram Supabase database. Agents selected from live simulation with post_count ≥ 3, ordered by post_count desc. Archetype-diverse within cohort.
Cohort 1	Gen Z Urban: age 18–27, location_type = urban. n = 30.
Cohort 2	Rural Adults: age 40–65, location_type = rural, region = south or midwest. n = 30.
Wave 1	Baseline: system prompt only, no stimulus context.
Wave 2	Post-stimulus: Klarna announcement injected as prior context message before question.
Wave 3	Re-poll: Klarna framed as past memory before a fresh question. Note: Wave 3 does not measure actual elapsed time. A temporal framing prompt simulates the cognitive state of a respondent who has processed the news. This is cross-sectional methodology with a temporal cue, not longitudinal measurement.
Sampling	From each pool, agents ordered by post_count desc, then filtered for archetype diversity (up to one agent per archetype, remainder filled sequentially). This biases toward experienced, opinionated agent profiles. Big Five/MBTI distributions were not checked against the broader agent population.
Sentiment	Keyword classifier (positive/neutral/negative). Not ML-based. Independent human review of 20 responses yielded 60% agreement. Primary error: classifier underestimates negativity, labeling rhetorical skepticism as neutral. Reported skepticism rates are floor estimates — actual skepticism likely higher in all waves. Drift direction is preserved across waves since bias is consistent.
Real benchmark	SurveyMonkey (2023) 'Consumers Push Back Against AI in Customer Service.' Ages 18-34: 41% negative. Ages 35-64: 58% negative.
Total calls	180 (3 waves × 60 agents). Inference time: ~25 minutes. Estimated cost: $0.36.
Date	March 21, 2026.

Statistical note: At n=30 per cohort, the 95% confidence interval on any reported proportion is approximately ±13 percentage points. Precise percentage claims within each wave should be read as directional, not exact. The wave-level trends (softening in Wave 2, re-hardening in Wave 3) are consistent across both independent cohorts, which strengthens

directional confidence despite the small absolute sample size.

Limitations: n=30 per cohort provides directional signal but is not statistically representative of any real population. The study covers two specific demographic slices — Gen Z Urban and Rural Adults 40–65 — not a nationally representative American sample. The keyword classifier yielded 60% human-rater agreement on 20 sampled responses; skepticism rates are floor estimates. The Gen Z baseline divergence (70% synthetic vs 41% real) is attributable to question framing (job displacement is politically charged) and sample skew (urban progressive youth); the demographic pattern — Gen Z less skeptical than rural adults — holds in both synthetic and real data. The rural validation (5.3pp delta) is within polling margin of error. Wave 3 simulates elapsed time via prompt framing; it does not measure actual attitude change over time. Lewis 1.5 agents reflect the Swarmgram training corpus distribution, which skews toward opinionated socially engaged profiles. This is a product demonstration, not peer-reviewed research.