

Do Personal AI Conversations Reduce Loneliness?

Evidence from a Pre-Registered Experiment

Louis Fréget¹, Ariell Reshef², and Claudia Senik³

¹Université Paris-Dauphine (PSL); Cepremap

²CNRS, Paris 1 Pantheon-Sorbonne, Paris School of Economics

³Sorbonne University; Paris School of Economics; IUF

Version: July 30, 2026

Abstract. Generative AI is increasingly promoted as a remedy for loneliness. We test this claim in a large-scale pre-registered experiment on a large sample of French adults. Participants assigned to treatment were asked to hold at least one personal conversation a day with a generative AI tool of their choice for 28 days; the control group continued as usual. The intervention raised socio-emotional AI use substantially. This led to an increase in subjective loneliness, a decline in life satisfaction, and a rise in depressive affect and in objective isolation, with more dinners taken alone. The effects are consistently adverse across outcome categories.

JEL: C93, I31, I12, O33. **Keywords:** Personal AI conversations ; Generative AI; Subjective well-being; Loneliness; Social connectedness.

1 Introduction

Modern societies are experiencing a loneliness epidemic: in 2023, the U.S. Surgeon General declared social isolation a public health emergency, noting that half of American adults experience measurable loneliness (Murthy, 2023), and in 2023, 19% of young adults across the world reported having no one they could count on for social support, a 39% increase compared to 2006 (Helliwell et al., World Happiness Report 2025, chapter 2).¹ The World Health Organization has described loneliness as a pressing health threat (WHO, 2023).²

Alongside this crisis, generative AI has become pervasive. By mid-2025, more than half of U.S. adults aged 18–64 had used a generative AI tool (Bick et al., 2025). Its fastest-growing uses—emotional support, companionship, and relationship advice—place these systems squarely in roles once occupied by human interlocutors. While technology leaders have proposed AI as a direct remedy for loneliness, there is little evidence to support this claim.³ In this paper we test whether encouraging individuals to converse with AI about their personal lives reduces loneliness, or reinforces it.

The question is of first-order importance. First, generative AI has created a new form of socio-emotional interaction, and this use is diffusing rapidly. Its consequences for well-being and social connection are, as yet, largely unknown. Second, the policy stakes are concrete. If ordinary AI conversation improves the well-being of isolated adults at negligible cost, it could serve as a low-cost bridge to human support. If it displaces that support, leading to adverse outcomes, the welfare implications run the other way.

The effect of such AI conversations on subjective well-being is theoretically ambiguous. AI conversation may improve well-being: it offers companionship, allows users to articulate their feelings, and may prompt them toward human contact (complementary). It may also deteriorate subjective well-being: repeated attention to personal difficulties can encourage rumination, which has been shown to amplify anxiety and depression (Nolen-Hoeksema, 1991); an always-available interlocutor may substitute for the effort of reaching out to others (substitution); and an agreeable model may validate grievances with relatives rather than help resolve them. Which mechanism dominates is an empirical question.

In this paper, we report the results of a pre-registered randomized controlled trial conducted on a large sample of French adults. Participants assigned to treatment were asked to

¹In the United States, the share of adults reporting no close friends rose from 3% in 1990 to 12% in 2021 (Cox, 2021) and up to 17% in 2024 (Cox & Pressler, 2024).

²For example, social isolation is associated with mortality risks comparable to smoking fifteen cigarettes a day (Holt-Lunstad et al., 2010).

³For example, Mark Zuckerberg (CEO of Meta Platforms, a leading technology company) argues that most people want more social connections than they have, and that AI might partially fill this gap, albeit not as well as human relationships. He also suggests that AI can help people prepare for difficult interpersonal conversations, for instance with a partner or a manager. (Zuckerberg, 2025)

hold at least one personal conversation per day with a generative AI tool of their choice for 28 consecutive days on topics related to their personal lives: their day, their feelings, their goals, their concerns, or any other aspect of their life they would like to share. Control participants were asked to continue their usual behavior. We measure subjective well-being (life satisfaction, loneliness, happiness, and depression), as well as relationship satisfaction with partners, family, and friends, and objective social interactions (meals taken alone and in-person meetings) at baseline and four weeks later. We primarily estimate the intention-to-treat effect of encouraging personal AI conversations.

This study is, to our knowledge, the first to test the causal effect of open-ended, naturalistic AI conversation—as opposed to scripted tasks or dedicated mental-health applications—on loneliness and subjective well-being. The existing causal evidence comes from guarded settings: dedicated mental-health applications with clinical safeguards (Angelucci et al., 2026), companion-specific products used by self-selected users (De Freitas et al., 2026), laboratory-style experiments with assigned interaction modes and scripted prompts and tasks (Fang et al., 2025), and structured chatbots explicitly designed around psychological well-being strategies (Schöne et al., 2025). Participants use ordinary generative AI, chosen by themselves, for unscripted personal conversation, and we never observe the content of these conversations. The paper relates to the literature in health economics on how digital technologies shape well-being (Allcott et al., 2020; Braghieri et al., 2022) and to the emerging literature on the effects of generative AI (Noy & Zhang, 2023), to which it adds evidence on naturalistic socio-emotional use.

Our findings can be summarized as follows. Treated participants complied with the request, increasing the number of days on which they held personal conversations with AI by about 5-6 days. These conversations did not, however, improve their mental health. In fact, subjective loneliness increased, while life satisfaction and affect declined. These effects consistently point in the same adverse direction; the single-item loneliness measure remains highly significant after multiple-testing correction, and the remaining outcomes remain significant at least at the 10% level. Objective social interactions and relationship quality were also adversely affected, although these effects are less precisely estimated. The sole significant subgroup contrast indicates that the harm is more pronounced among partnered, as opposed to single respondents. Taken together, the results are more consistent with substitution and rumination than with support and complementarity.

2 Data and Design

Setting and sample. We recruited adults residing in France from an online panel (recruited by the survey company Bilendi). In Wave 1 (January 5–19, 2026), participants completed a baseline questionnaire covering demographics and all outcomes and were randomized to treatment or control in equal proportions. The treatment period ran from January 22 to February 18, 2026. In Wave 2 (February 19–28, 2026), participants completed the endline questionnaire covering the same outcome measures.⁴ Randomization was at the individual level. Of those who passed the baseline attention check, 12,356 respondents (6,253 control and 6,103 treatment) completed both waves and form the analysis sample. Endline completion was similar across both arms: 81.7% of control participants and 81.3% of treatment participants completed Wave 2 ($p = 0.56$). Covariate balance is good, as attested by Table 1, which reports Imbens–Rubin normalized differences for the primary analysis sample, all below 0.04 in absolute value and well under the conventional threshold of 0.25. Table 2 compares the sample to the French population: it is older, more educated, and more likely to be partnered, and reports somewhat lower well-being at baseline.⁵

Intervention. Participants assigned to treatment were asked to hold at least one brief personal conversation each day for 28 consecutive days with a generative AI tool of their choice (for example, ChatGPT, Le Chat, or Mistral) on personal topics such as their day, feelings, goals, or concerns. Control participants were asked only to maintain their usual behavior. We did not access the content of any conversation. Compliance was self-reported at endline: participants indicated the total number of days on which they had held a personal AI conversation over the 28-day period. To help participants maintain an accurate count, both groups were provided with a daily online log in which they recorded whether they had spoken with an AI about their personal life each day; the log automatically tallied their running total. Weekly SMS reminders prompted both groups to update the log.

First stage. The intervention shifted behavior as intended. Assignment to treatment increased the self-reported number of days with personal AI conversation by about 5.8 days over the 28-day window (compared to a control median of 3 days).⁶ (compared to a control

⁴During the experimental period (January 22 – February 18, 2026), several capable generative AI tools were freely available, including ChatGPT (GPT-4o, OpenAI), Claude 3.5 Sonnet (Anthropic), Gemini 1.5 Flash (Google), and Le Chat (Mistral). Premium subscriptions—ChatGPT Plus, Claude Pro, and Gemini Advanced—offered higher usage limits for around \$20 per month.

⁵The over-representation of older adults is consistent with feedback from the survey company, which noted that younger respondents were less likely to commit to a month-long daily task.

⁶Respondents who could not recall their exact number of personal AI conversation days were allowed to skip this question, resulting in 892 missing values (7.2% of the analysis sample). Under mean imputation

median of 3 days). As pre-specified, we also report the effect on a broad index of AI use frequency across all domains, which increased by 0.53 standard deviations ($F = 798$). As an exploratory measure, we additionally report the effect on a socio-emotional sub-index summing three items — emotional support, companionship, and relationship advice, each on a 1–6 frequency scale — which increased by 0.51 standard deviations ($F = 654$).

Outcomes. Our primary outcomes measure subjective well-being: life satisfaction, happiness, and depression, each on a 0-10 scale, and loneliness, measured both with a single 0-10 item and with the three-item UCLA scale (Hughes et al., 2004). These integer scales are standard in the economics of well-being (Kahneman & Deaton, 2010) and predict subsequent behavior in an approximately linear fashion (Kaiser & Oswald, 2022). Secondary outcomes measure social connectedness: satisfaction with partners, family, and friends, and objective loneliness, specifically the number of meals taken alone and the number of in-person meetings with friends or relatives in the preceding seven days.⁷ We also record willingness-to-accept (WTA) to forgo AI for one week and how pleasant participants found the conversations (0-10).

Empirical strategy. The primary estimand is the intention-to-treat (ITT) effect of assignment on endline outcomes. Our main specification is an analysis of covariance (ANCOVA)⁸:

$$Y_{i2} = \alpha + \beta T_i + \gamma Y_{i1} + \mathbf{X}_{i1}'\boldsymbol{\theta} + \varepsilon_{i2}, \quad (1)$$

where Y_{i2} is the outcome at endline, T_i is treatment assignment, Y_{i1} is the baseline value of the same outcome, and $\mathbf{X}_{i1}$ is a vector of pre-specified covariates (age, gender, education, socio-economic status, telework frequency, relationship status, and household composition). We report robust standard errors, and each outcome is estimated on its own complete cases, so the estimation sample varies across outcomes. We pre-specified three outcome families—subjective well-being, relationship quality, and objective interaction—and report Romano-Wolf stepdown p -values within each.⁹ As pre-specified robustness checks, we report

with a missing value indicator, as pre-specified in the pre-analysis plan, the first-stage coefficient remains large and highly significant (5.41, SE = 0.18, $p < 0.001$, $N = 12,356$), though mechanically attenuated toward zero relative to the complete-cases estimate (5.83).

⁷Perceived social support was pre-specified in the relationship-quality family, but it was not collected in the endline questionnaire due to miscommunications with the survey company. Note that this omission cannot affect the qualitative conclusions, as including perceived social support in the Romano-Wolf family would make the correction more conservative, and Family 2 outcomes are already non-significant after multiplicity adjustment.

⁸ANCOVA dominates difference-in-differences in efficiency when the baseline-endline autocorrelation is below one. (McKenzie, 2012)

⁹We report the pre-specified experimental analyzes in Section 4.2 of the pre-analysis plan.

a difference-in-means specification without baseline adjustment (Appendix Table A3), and Lee (2009) bounds to assess sensitivity to differential attrition (Appendix Table A4). Our results are largely robust to these checks.¹⁰

3 Results

Baseline associations. We first document the cross-sectional negative relationship between socio-emotional AI use and well-being at baseline (Appendix Table A1). These associations are descriptive and subject to selection, as the decision to use AI for companionship is not random. In particular, reverse causality might be at play. Our experiment is intended to eliminate this selection through randomization.

Subjective well-being. The experimental estimates are reported in Table 4. Assignment to daily personal AI conversation increases subjective loneliness, on both the single-item measure (0.17 points on a 0-10 scale, or 0.06 SD; $p < 0.001$) and the UCLA-3 scale (0.03 SD; $p = 0.02$). It reduces life satisfaction (by 0.06 points, or 0.03 SD; $p = 0.03$) and raises reported depression on a 0-10 scale (0.08 points, or 0.03 SD; $p = 0.04$); the effect on happiness is negative but not statistically significant. After Romano–Wolf correction within each family of outcomes, the single-item loneliness effect remains highly significant ($p = 0.001$), while the life-satisfaction, UCLA-3, and depression effects are statistically significant at the 10% level ($p \approx 0.08$). The point estimates are modest in size, but they are uniformly signed: greater exposure to personal AI conversation decreases subjective well-being.

Our estimates are of the opposite sign to those of Angelucci et al. (2026), who find gains of 0.3 SD from a purpose-built AI mental health application administered to women in psychological distress—a very different population and product from the general-purpose AI tools studied here. A more relevant comparison is Allcott et al. (2020), who find that completely deactivating Facebook for four weeks improved subjective well-being by 0.09 SD. Our nudge—which shifted personal AI conversation by only about 6 days over 28—produces an adverse effect on loneliness of roughly 0.06 SD, or about two-thirds of that magnitude.

Relationships and objective contact. The same pattern appears in measures of real-world connection (Table 5). Treated participants report, on average, 0.06 more lunches and

¹⁰The difference in-means specification yields directionally identical results across all outcomes (Appendix Table A3), though point estimates and standard errors differ somewhat without baseline adjustment. Lee (2009) bounds confirm the robustness of the primary results to differential attrition. With a trimming proportion of only 0.45%, reflecting a difference in response rates of 0.4 percentage points ($p = 0.56$), the bounds are tight. Both bounds are positive and significant for the single-item loneliness measure, and positive for UCLA-3 and depressive affect.

0.06 more dinners taken alone per week, roughly 0.06 fewer in-person meetings with friends or relatives, and satisfaction with family relationships declines by 0.06 points. The effects on partner and friendship satisfaction are negative but not statistically significant. Within their families, the p-values of these estimates lie below 0.13 after the Romano-Wolf correction, while the effect on the number of dinners alone is statistically significant at the 10% level. Overall, however, all estimates move in the same adverse direction as the subjective outcomes.

Participants’ assessment of AI. The adverse effect does not appear to reflect a negative experience with the tool (Table 6). Treated participants rate the conversations as slightly more pleasant than controls (0.10 SD; $p < 0.001$), and they are no more likely to report that the AI took their side in conflicts. The decline in subjective well-being, therefore, does not seem to stem from an adverse experience or from validation dynamics that inflate grievances; if anything, participants found the conversations mildly pleasant even as their subjective well-being declined. Although these outcomes are only asked of participants who used AI during the experiment, as pre-specified in the registered questionnaire, baseline covariates remain well balanced within each estimation sample, with all standardized differences below 0.10 (See Appendix Table A2).¹¹

Heterogeneity. Among the pre-specified moderators, only one interaction is statistically significant (Table 7). The adverse effect on subjective loneliness is concentrated among participants who were partnered at baseline; single respondents are essentially unaffected, and the couple-versus-single interaction is strongly significant ($p = 0.001$). The remaining pre-specified splits—baseline loneliness, age, gender, and telework—show no heterogeneity once the full sample is used. The concentration of harm among partnered subjects is consistent with AI conversation displacing existing human interaction among respondents who had a live-in relationship to begin with, rather than filling a pre-existing void.

4 Discussion

The experiment provides an exogenous manipulation of socio-emotional AI use and a consistent set of results. The intervention raised personal AI conversation, with an adverse effect on subjective well-being, including loneliness, depressive affect, and life satisfaction. Objective isolation increased, with more dinners being taken alone. This configuration is

¹¹These outcomes are estimated without baseline adjustment, as the baseline is only defined for participants who had already used AI for the relevant type of conversation prior to the experiment—a highly selected subsample. Results are qualitatively unchanged when restricting to this subsample and controlling for the baseline value.

more consistent with substitution and rumination than with complementarity and support. The heterogeneity analysis supports this interpretation, as it shows that the adverse effect of personal AI conversations is concentrated among partnered respondents—those in a live-in relationship—rather than among singles.

One key limitation of our study is the duration of the experiment. It is only one month, and the effects of longer exposure may differ. Because we encourage rather than mandate use, our estimates are ITT effects of a nudge toward AI conversation, not the effects of intensive use. More importantly, one may ask whether the adverse effects that we uncover reflect AI conversation specifically or simply the act of daily personal reflection—which the treatment also induces. An alternative self-reflection control arm, such as diary-writing or discussion groups, would help disentangle the two, and we encourage future work to include one. Still, the volume of personal reflection itself may be part of the mechanism: an always-available, responsive interlocutor may induce people to dwell on their problems more than they otherwise would. The fact that outcomes worsen rather than improve suggests that an agreeable interlocutor is not the same as a helpful one.

Overall, the results temper an optimistic view of personal AI conversations. The proposition that ordinary AI conversation can serve as a low-cost remedy for loneliness does not survive a naturalistic test; encouraging daily AI conversation left participants somewhat lonelier and more isolated. Our results suggest that encouraging ordinary, unstructured AI conversations is more likely to worsen subjective well-being than to improve it.

Acknowledgments

This study was pre-registered on the AEA RCT Registry (AEARCTR-0016749). We thank Victoire Chatain for her help with the initial data cleaning and for cross-validating the analysis, as well as the seminar participants at LEGOS for their comments. The authors declare no conflict of interest. The study was funded by the French National Research Agency (ANR-17-EURE-0001 and ANR-21-CE26-0019-03). This study was approved by the IRB-PSE (protocol 2025-044).

Table 1 Baseline balance between treatment and control

	Control	Treatment	Norm. diff
Age (years)	52.13	51.73	-0.028
Female (%)	54.66	54.22	-0.009
Tertiary education (%)	25.30	24.41	-0.020
In a relationship (%)	70.43	71.11	0.015
Household size	2.43	2.42	-0.007
Children under 15	1.74	1.73	-0.031
Economically active (%)	62.39	64.26	0.039
Socio-emotional AI use (z)	0.00	0.02	0.016
Life satisfaction (0–10)	6.44	6.43	-0.001
Loneliness, UCLA-3 (3–9)	4.79	4.82	0.018
Loneliness, single item (0–10)	3.50	3.55	0.020
Happiness (0–10)	6.41	6.44	0.012
Depression (0–10)	3.01	3.06	0.019
Observations	6253	6103	12356

Notes. Means by arm at baseline on the primary analysis sample ($N = 12,356$). The last column reports the Imbens–Rubin (2015) normalized difference: $\Delta = (\bar{X}_T - \bar{X}_C) / \sqrt{(s_T^2 + s_C^2)/2}$. Values below 0.10 in absolute value indicate negligible imbalance; the conventional threshold for concern is 0.25 (Imbens & Rubin, 2015). All normalized differences are below 0.04 in absolute value. Endline completion is unrelated to treatment ($p = 0.56$). The socio-emotional AI-use index is defined only for baseline AI users ($n = 6,756$); remaining observations have $z = 0$ (non-user centred).

Table 2 Sample characteristics vs. French population

	Sample	France	Difference
<i>Sociodemographics</i>			
Women (%)	54.4	51.6	+2.8
Age 18–34 (%)	14.2	28.1	−13.9
Age 55+ (%)	47.2	36.5	+10.7
Higher education, long (%)	24.9	16.0	+8.9
Partnered (%)	70.8	57.5	+13.3
Child <15 at home (%)	26.5	26.5	0.0
Household size	2.43	2.10	+0.3
<i>Well-being at baseline</i>			
Life satisfaction (0–10)	6.44	6.60	−0.16
Happiness yesterday (0–10)	6.43	6.90	−0.47
Depression yesterday (0–10)	3.04	2.10	+0.94
Loneliness, single item (0–10)	3.52	2.70	+0.82
Observations	12,356		

Notes. Sample statistics are computed on the primary analysis sample ($N = 12,356$). French population benchmarks for sociodemographics are from INSEE (2023 census); benchmarks for subjective well-being are from the CAMME survey (<https://www.insee.fr/fr/metadonnees/source/serie/s1208>). Higher education (long) corresponds to Master or PhD degrees.

Table 3 First stage: effect of assignment on personal AI conversations

	Days of personal conversation (0–28)	Broad AI use index (z)	Socio-emotional AI use (z)
Treatment	5.832*** (0.195)	0.533*** (0.019)	0.505*** (0.020)
Control median	3	—	—
First-stage F	894.2	797.8	653.7
Observations	11,464	12,356	12,356

Notes. ITT effect of assignment on compliance, measured at endline. Days of personal conversations with AI are self-reported. The socio-emotional index sums three endline items (emotional support, companionship, relationship advice), each on a 1-6 frequency scale. The index is standardized using the control-group mean and standard deviation. Both regressions control for pre-specified covariates. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4 Effect of daily personal conversations with AI on subjective well-being

Outcome	ITT (orig. units)	ITT (SD)	Control mean	RW p	N
Life satisfaction (0–10)	-0.058** (0.027)	-0.030** (0.014)	6.98	0.0750	12356
Loneliness, UCLA-3 (3–9)	0.049** (0.022)	0.027** (0.012)	4.67	0.0750	12356
Loneliness, single item (0–10)	0.168*** (0.040)	0.057*** (0.013)	3.44	0.0002	12356
Happiness (0–10)	-0.018 (0.031)	-0.009 (0.015)	6.99	0.5550	12356
Depression (0–10)	0.083** (0.041)	0.030** (0.015)	2.96	0.0750	12356
Baseline outcome & covariates	Yes	Yes			

Notes. ANCOVA ITT estimates (Eq. 1): each row is a separate regression of the endline outcome on treatment assignment, controlling for the baseline value of the outcome and pre-specified covariates. Following the pre-analysis plan, 0–10 and 0–7 outcomes are reported in original units (column 1); the UCLA-3 summed index (range 3–9) is standardized. Column 2 reports all effects in control-group baseline standard deviations for comparability. Each outcome is estimated on its own complete cases, so N varies (last column). Robust standard errors in parentheses. Column 4 reports Romano–Wolf stepdown p -values within the subjective well-being family (4,999 bootstrap replications). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5 Effect of daily personal AI conversations on social connectedness

Outcome	ITT	Control mean	RW p	N
<i>Relationship quality (0–10)</i>				
Partner satisfaction	-0.059* (0.033)	7.50	0.1244	8594
Family satisfaction	-0.058** (0.029)	7.22	0.1126	12246
Friendship satisfaction	-0.004 (0.029)	6.98	0.8856	12356
<i>Objective contacts (days out of 7)</i>				
Lunches alone	0.061* (0.034)	2.83	0.1256	12356
Dinners alone	0.064** (0.029)	2.15	0.0844	12356
Meetings with friends/relatives	-0.056* (0.030)	2.11	0.1256	12356
Baseline outcome & covariates	Yes			

Notes. ANCOVA ITT estimates in original units, each on its own complete cases (N in the last column). Partner satisfaction is estimated on the subsample partnered at baseline. Family satisfaction excludes respondents who indicated at baseline that the question was not applicable (no close family; $n = 110$). Romano–Wolf stepdown p -values are computed separately within the relationship-quality family and within the objective-interaction family. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6 Effect on how participants value and experience AI

Outcome	ITT (SD)	Control mean/median	N
Conversations pleasant (0–10)	0.097*** (0.022)	6.63	7,826
AI “took my side” in conflicts	0.024 (0.026)	6.37	5,744
WTA to forgo AI for one week	-0.021 (0.025)	50	5,587
Covariates	Yes		

Notes. ITT estimates on standardized endline outcomes, controlling for pre-specified covariates without baseline adjustment. The estimation samples are smaller than the primary analysis sample ($N = 12,356$) because these questions were only asked of participants who had used generative AI during the experiment, and respondents could select “not applicable” for the pleasantness and conflict items if they had not had the relevant type of conversation. For WTA, the estimation sample is smaller because the question was only asked of AI users, and a substantial share of users did not respond. For WTA, the control-group median is reported rather than the mean, as the distribution has a very long right tail (mean = 41,078 euros). Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7 Heterogeneity in the effect on subjective loneliness (UCLA-3, SD)

Moderator	Effect (A)	Effect (B)	Interaction (p)
<i>Baseline loneliness ($A = High \geq 5, B = Low < 5$)</i>			
High loneliness (≥ 5) vs Low (< 5)	0.018 (0.020)	0.038 (0.014)	0.397
<i>Age group (effect vs. 35–44 reference, $p =$ interaction with ref.)</i>			
18–24	-0.078 (0.092)	—	0.196
25–34	0.064 (0.039)	—	0.725
35–44 (ref.)	0.047 (0.029)	—	—
45–54	0.020 (0.028)	—	0.505
55+	0.019 (0.016)	—	0.398
<i>Gender ($A = Men, B = Women$)</i>			
Men vs Women	0.018 (0.017)	0.035 (0.017)	0.492
<i>Relationship status ($A = Couple, B = Single$)</i>			
Couple vs Single	0.053 (0.014)	-0.036 (0.023)	0.001***
<i>Telework ($A = Below\ median, B = Above\ median$)</i>			
Telework: below vs above median	0.033 (0.014)	0.013 (0.023)	0.463
Observations	12,356		

Notes. Each panel interacts treatment with one pre-specified moderator in the main ANCOVA on UCLA-3 loneliness (control-group SD units), controlling for the baseline outcome and all covariates. For binary moderators, Effect (A) and Effect (B) report the treatment effect in each subgroup, and the last column reports the p -value of the interaction term. For the age-group moderator (5 bins), the reference category is 35–44; Effect (A) reports the treatment effect in each bin relative to the reference, and the last column reports the p -value of the interaction with the reference group (“—” for the reference itself). All age-group interactions are non-significant. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

References

- Allcott, H., Braghieri, L., Eichmeyer, S., & Gentzkow, M. (2020). The welfare effects of social media. *American Economic Review*, 110(3), 629–676.
- Angelucci, M., Fabregas, R., & Vazquez, A. (2026). *The well-being effects of digital mental health care*. IZA Discussion Paper No. 18538.
- Braghieri, L., Levy, R., & Makarin, A. (2022). Social media and mental health. *American Economic Review*, 112(11), 3660–3693.
- Cox, D. A. (2021). *The state of American friendship: Change, challenges, and loss*. Survey Center on American Life, American Enterprise Institute. <https://www.americansurveycenter.org/research/the-state-of-american-friendship-change-challenges-and-loss/>
- Cox, D. A., & Pressler, S. (2024). *Disconnected: The growing class divide in American civic life*. Survey Center on American Life, American Enterprise Institute. <https://www.americansurveycenter.org/research/disconnected-places-and-spaces/>
- De Freitas, J., Oğuz-Uğuralp, Z., Uğuralp, A. K., & Puntoni, S. (2026). AI companions reduce loneliness. *Journal of Consumer Research*, 52(6), 1126–1148.
- Fang, C. M., Liu, A. R., Danry, V., Lee, E., Chan, S. W. T., Pataranutaporn, P., Maes, P., Phang, J., Lampe, M., Ahmad, L., & Agarwal, S. (2025). *How AI and human behaviors shape psychosocial effects of extended chatbot use: A longitudinal randomized controlled study* [Working paper]. MIT Media Lab / OpenAI.
- Bick, A., Blandin, A., & Deming, D. (2025, November 13). The state of generative AI adoption in 2025. *Federal Reserve Bank of St. Louis, On the Economy Blog*. <https://www.stlouisfed.org/on-the-economy/2025/nov/state-generative-ai-adoption-2025>.
- Holt-Lunstad, J., Smith, T. B., & Layton, J. B. (2010). Social relationships and mortality risk: A meta-analytic review. *PLOS Medicine*, 7(7), e1000316.
- Hughes, M. E., Waite, L. J., Hawkey, L. C., & Cacioppo, J. T. (2004). A short scale for measuring loneliness in large surveys. *Research on Aging*, 26(6), 655–672.
- Imbens, G. W., & Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press.

- Kahneman, D., & Deaton, A. (2010). High income improves evaluation of life but not emotional well-being. *Proceedings of the National Academy of Sciences*, 107(38), 16489–16493.
- Kaiser, C., & Oswald, A. J. (2022). The scientific value of numerical measures of human feelings. *Proceedings of the National Academy of Sciences*, 119(42), e2210412119.
- Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Review of Economic Studies*, 76(3), 1071–1102.
- McKenzie, D. (2012). Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics*, 99(2), 210–221.
- Murthy, V. H. (2023). *Our epidemic of loneliness and isolation: The U.S. Surgeon General’s advisory on the healing effects of social connection and community*. U.S. Department of Health and Human Services.
- Nolen-Hoeksema, S. (1991). Responses to depression and their effects on the duration of depressive episodes. *Journal of Abnormal Psychology*, 100(4), 569–582.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192.
- Schöne, J. P., Salecha, A., Lyubomirsky, S., Eichstaedt, J., & Willer, R. (2025). *Structured AI dialogues can increase happiness and meaning in life* [Working paper]. <https://psyarxiv.com/2bf7t>
- Zuckerberg, M. (2025, April 29). *Interview with Dwarkesh Patel* [Video podcast]. YouTube. <https://www.youtube.com/watch?v=rYXeQbTuVl0>
- World Health Organization. (2023, November 15). *WHO launches commission to foster social connection*. <https://www.who.int/news/item/15-11-2023-who-launches-commission-to-foster-social-connection>

Online Appendix

Table A1 Socio-emotional AI use and well-being at baseline (correlational)

	Life sat. (0–10)	Loneliness UCLA-3	Loneliness item (0–10)	Happiness (0–10)	Depression (0–10)
Socio-emotional AI use (z)	−0.050** (0.022)	0.193*** (0.017)	0.450*** (0.027)	−0.023 (0.022)	0.386*** (0.028)
Controls	Yes	Yes	Yes	Yes	Yes
R^2	0.061	0.086	0.097	0.040	0.068
Observations	12,356	12,356	12,356	12,356	12,356

Notes. OLS at baseline on the primary analysis sample. Each column is a separate regression of the outcome on the standardized socio-emotional AI-use index, controlling for age, gender, education, relationship status, number of friends at age fifteen, and household size. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A2 Baseline covariate balance within AI report estimation samples

	Pleasantness ($N = 7,826$)	Sided with me ($N = 5,744$)	WTA ($N = 5,587$)
Age	−0.033	−0.031	−0.066
Female (%)	−0.008	−0.028	−0.011
Tertiary educ. (%)	−0.015	−0.030	+0.007
In a relationship (%)	−0.013	−0.050	−0.016
Household size	+0.029	+0.008	+0.000
Economically active (%)	−0.025	−0.018	−0.045
Life satisfaction (0–10)	−0.054	−0.070	−0.025
Loneliness, UCLA-3	+0.043	+0.063	−0.000

Notes. Imbens–Rubin normalized differences within each estimation sample in Table 6.

Table A3 Robustness: difference-in-means vs. ANCOVA

Outcome	ANCOVA (primary spec.)			Difference-in-means		
	ITT	(SE)	<i>N</i>	ITT	(SE)	<i>N</i>
<i>Subjective well-being (Family 1)</i>						
Life satisfaction (0–10)	-0.058**	(0.027)	12356	-0.057	(0.035)	12356
Loneliness, UCLA-3 (std.)	0.049**	(0.022)	12356	0.074**	(0.032)	12356
Loneliness, single item (0–10)	0.168***	(0.040)	12356	0.204***	(0.053)	12356
Happiness (0–10)	-0.018	(0.031)	12356	-0.006	(0.037)	12356
Depression (0–10)	0.083**	(0.041)	12356	0.121**	(0.050)	12356
<i>Relationship quality (Family 2)</i>						
Partner satisfaction (0–10)	-0.059*	(0.033)	8594	-0.099**	(0.046)	8597
Family satisfaction (0–10)	-0.058**	(0.029)	12246	-0.070*	(0.039)	12356
Friendship satisfaction (0–10)	-0.004	(0.029)	12356	-0.025	(0.038)	12356
<i>Objective contact (Family 3)</i>						
Lunches alone (0–7)	0.061*	(0.034)	12356	0.107**	(0.047)	12356
Dinners alone (0–7)	0.064**	(0.029)	12356	0.073	(0.049)	12356
Meetings with friends/relatives (0–7)	-0.056*	(0.030)	12356	-0.030	(0.033)	12356

Notes. The ANCOVA specification (columns 1–3) controls for the baseline value of the outcome and pre-specified covariates. The difference-in-means specification (columns 4–6) regresses the endline outcome on treatment assignment only, without baseline adjustment (pre-specified robustness check). Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A4 Lee (2009) bounds on treatment effects (subjective well-being)

Outcome	Lower bound		Upper bound	
	Coef.	95% CI	Coef.	95% CI
<i>Subjective well-being (Family 1)</i>				
Life satisfaction (0–10)	-0.088	[-0.198, 0.022]	-0.043	[-0.123, 0.037]
Loneliness, UCLA-3 (std.)	0.066*	[-0.008, 0.140]	0.093*	[-0.001, 0.187]
Loneliness, single item (0–10)	0.189***	[0.080, 0.298]	0.234***	[0.118, 0.350]
Happiness (0–10)	-0.037	[-0.137, 0.063]	0.008	[-0.080, 0.096]
Depression (0–10)	0.108*	[-0.003, 0.218]	0.153**	[0.033, 0.273]

Notes. The trimming proportion is 0.45%, reflecting a difference in response rates of -0.4 percentage points between treatment and control ($p = 0.56$). Bootstrap standard errors, 2,000 replications. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.