Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Reframing Memories of Academic Uncertainty With a Constructive-Developmental AI Reduces Expressed Uncertainty in STEM Learners

Maximilian Ramseyer

maximilian.ramseyer@tufts.edu

Tufts University

Ayanna Thomas

Tufts University

Article

Keywords: autobiographical memory, memory reconstruction, academic uncertainty, cognitive reappraisal, self-authorship, constructive-developmental theory, tentative language, LIWC, STEM education, large language models, generative AI agents

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11368354/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Reframing Memories of Academic Uncertainty With a Constructive-Developmental AI Reduces Expressed Uncertainty in STEM Learners

AUTHORS. Ramseyer, Maximilian[1]; Thomas, Ayanna[1]

AFFILIATIONS. [1] Department of Psychology, Tufts University, Medford, Massachusetts, USA

CORRESPONDENCE. Maximilian Ramseyer, maximilian.ramseyer@tufts.edu

Abstract

Conversational AI can alter what people later remember, yet its effects on how students remember their own academic experiences are untested. We examined whether a large language model (LLM)'s reframing of memories of academic uncertainty changes how STEM undergraduates later describe those memories. Participants recalled five such memories and received an LLM-generated response to each: a reframe instructed to move them toward a self-authoring perspective according to Constructive-Developmental theory, or a restatement that elaborated the uncertainty. One week later, Self-Authorship participants used less tentative language when redescribing their memories, whereas restatement produced no reliable change. LLM-generated vocabulary reappeared in later descriptions of the specific memories it addressed in both conditions. The reduction in tentative language remained after novel reframe vocabulary was removed from the Session 2 descriptions. Emotional valence did not differ by condition. A brief interaction with a theory-constrained AI agent can change how learners later describe academic uncertainty.

Keywords: autobiographical memory, memory reconstruction, academic uncertainty, cognitive reappraisal, self-authorship, constructive-developmental theory, tentative language, LIWC, STEM education, large language models, generative AI agents

Introduction

Generative AI agents have transformed the teacher-student dyad into a triad: large language model (LLM) agents now tutor students (Kestin et al., 2025) and are being built as companions that provide emotional and motivational support through structured cognitive reappraisal (Bhattacharjee et al., 2026; Held et al., 2025). Most college students use them in their coursework, with use rising under academic workload and time pressure (Abbas et al., 2024; von Garrel & Mayer, 2023). When a student describes an academic difficulty to a conversational agent, its response can offer an alternative interpretation of the experience, tailored to the student's own account (Sharma et al., 2023; Zhan et al., 2024). Whether such a response may alter how the student later remembers this experience must be considered in light of autobiographical memory's reconstructive nature (Bartlett, 1932; Conway & Pleydell-Pearce, 2000). Recall reassembles a personally experienced event in relation to the rememberer's current goals and self-concept (Conway, 2005), such that the meaning assigned to an event is partly reconstructed based on its current appraisal (Levine, 1997; Levine et al., 2001). This malleable aspect of memory is further evident in studies showing that content encountered after an experience can become incorporated into later recollections (Chan et al., 2009; Loftus, 2005; Marsh, 2007).

Generative AI is a new source of such content, as conversational agents and AI-edited media can introduce details that people later report as part of their own memories, with chatbot-induced distortions demonstrated to be present one week later (Chan et al., 2024; Pataranutaporn et al., 2025). Related findings also suggest that the content of a memory, and how it is later described, may be open to deliberate modification (Elsey et al., 2018; Speer et al., 2021). This possibility carries functional significance, because recollection of past events guides subsequent affect and behavior (Pillemer, 2003; Rubin et al., 2011). Academic settings offer a particularly consequential context in which to extend this finding, because they often confront learners with challenging experiences that are later remembered and revisited (Pillemer et al., 1996). These demands are especially pronounced in science, technology, engineering, and mathematics (STEM) coursework, where difficulty and student attrition are well documented (Chen, 2013; Seymour & Hewitt, 1997). The current investigation examines whether a personalized reframe generated by an LLM agent can update memories of one such difficulty often encountered in STEM educational settings: academic uncertainty.

Among the challenges of STEM learning, uncertainty is a recurrent and affectively salient one (Jordan & McDaniel, 2014; Moffett et al., 2021). Uncertainty refers to the perceived lack of predictability that accompanies unfamiliar or difficult material (Lamnina & Chase, 2019). It can also be aversive: intolerance of uncertainty is a transdiagnostic correlate of anxiety (Carleton, 2016), and a scoping review found uncertainty to be predominantly associated with anxiety, depression, and distress (Massazza et al., 2023). In our present study, participants are asked to retrieve various memories of academic uncertainty that may differ in narrative content but share the common feature of being an unresolved experience for the learner. This unresolved aspect is the target of the experiment's memory updating manipulation.

Reframing memories of uncertainty has broader significance because people draw on remembered experiences when anticipating their emotional responses to similar future situations (Morewedge et al., 2005; Wirtz et al., 2003). Furthermore, emotions and appraisals have been shown to influence academic engagement and achievement (Pekrun et al., 2017), making learners' responses to uncertainty relevant to their willingness to persist. Research on STEM attrition identifies students' course experiences and academic performance as factors associated with persistence in these fields (Chen, 2013; Seymour & Hewitt, 1997). Importantly, autobiographical recall of academic experiences has been shown to influence performance: in Selimbegović et al. (2011), recalling a memory of academic success or failure affects students' performance on a subsequent math test, with performance on difficult math problems depending on the type of content recalled immediately prior to the test. A follow-up study found that the memory exerts this effect by shaping the appraisal the student brings to the upcoming task, with fear of failure and challenge appraisal each mediating the link between recall and performance. The effect also carried over to tasks in subjects outside of the one recalled (Selimbegović et al., 2011; Selimbegović et al., 2016). These findings indicate that students' remembered version of earlier academic experiences in STEM can influence their expectation and performance in a new academic situation. Furthermore, they suggest that learners can improve their academic performance depending on which kind of academic memory content they retrieve.

In this study, we investigate whether engaging participants in an adaptive memory reframing exercise affects how they later recall their STEM-related academic memories of uncertainty, borrowing from the methodology introduced by Speer et al. (2021). In that study, participants described negative autobiographical memories and then either elaborated on their

positive meaning or recalled them naturally. When participants described the memories one week and two months later, positive elaboration produced greater increases in positive content. In the Positive condition, later memory descriptions also incorporated details introduced during elaboration (Speer et al., 2021).

The present study departs from this prior work in what the reframe is designed to modify. Rather than targeting the valence of the remembered experience, the intervention targets how the memory is organized in relation to the self. Constructive-Developmental theory distinguishes the content of memory, i.e., what a person thinks or feels about an experience, from its structure, or the way through which that memory is interpreted and organized (Kegan, 1994; Creamer et al., 2010). Self-authorship describes a form of organization in which individuals increasingly hold their interpretations, assumptions, and responses as objects of reflection and organize their responses from an internally generated framework of beliefs, values, and commitments (Baxter Magolda, 2008; Kegan, 1994). Longitudinal research has linked the development of self-authorship to a greater capacity to navigate uncertainty and circumstances outside one's control by distinguishing external events from one's interpretation and response to them (Baxter Magolda, 2008). Similarly, studies on decision-making have found that self-authored individuals evaluate competing perspectives in relation to internally defined beliefs rather than relying primarily on external formulas for what to think or do (Creamer & Laughlin, 2005; Laughlin & Creamer, 2007). Importantly, this developmental change in memory organization does not mean that memories must become more positive, as individuals may continue to experience fearful, unpleasant or negative affect while becoming more internally directed in how they make sense and respond to uncertainty (Kegan, 1994; Baxter Magolda, 2008). Applied to an academic memory centered on uncertainty, a self-authoring reframe does not prompt the participant to feel better about the uncertainty, but to reorganize it in relation to their own values, interpretations, and direction for action. We expect this shift to make uncertainty less organizing at later retrieval, expressed in less hedging or hesitant language, without obligatorily increasing positive or decreasing negative emotional valence.

To measure changes of uncertainty wording in memory narration, we use tentative language as the primary linguistic outcome. The tentative category in Linguistic Inquiry and Word Count (LIWC; tentat) was defined as a direct measure of expressed uncertainty, as it quantifies the

percentage of words classified as expressing hesitation (Boyd et al., 2022). An early validation study compared LIWC scores with independent judges' ratings of college students' writing about their adjustment to college. Within the emotional-disclosure group, LIWC tentative scores correlated with judges' ratings of tentative language across participants' essays ($r = .41$; Pennebaker & Francis, 1996). We use this category as a linguistic measure of expressed uncertainty in the descriptions participants provide at each session. According to Baxter Magolda's (2008) research, we expect Self-Authorship reframing of participants' academic memories to make uncertainty less prominent in subsequent narrative word choice, without necessarily altering remembered event details or significantly shifting affective valence.

As secondary linguistic outcomes, we examine whether Self-Authorship reframing may attenuate anxiety-related language, and whether anxiety might be sustained or increased when unresolved uncertainty is rehearsed in the Control condition (Kross & Ayduk, 2008, 2017). We also use LIWC's Analytic summary dimension, which distinguishes categorical language characterized by greater article and preposition use from more personal, dynamic language characterized by greater use of pronouns, auxiliary verbs, adverbs, conjunctions, and negations (Boyd et al., 2022; Pennebaker et al., 2014). Prior intervention research has reported more analytic language following cognitive reappraisal tasks compared with generic expressive writing (Colombo et al., 2026). We examine whether differences in anxiety-related and analytic language also emerge at later recall after participants engage in a self-authored reappraisal of their academic experience of uncertainty.

Whether interactions with LLM agents change how students later recall their memories of academic uncertainty has not been measured. Previous work found that an LLM chatbot that asked misleading questions during a simulated witness interview induced false memories of a video that participants had watched, and the false memories persisted one week later (Chan et al., 2024). In a separate study, participants who wrote essays with an LLM were less able than unassisted writers to quote their own essays afterward and reported lower ownership of them (Kosmyna et al., 2025). The first study shows that content an agent introduces can enter later recall of an event, whereas the second shows that text produced with an agent is recalled less accurately as one's own. In both studies, the agent's content concerned a video or an essay, but not an experience from the participant's own life. Neither tested an agent responding to a student's own account of an

academic experience, followed by their unassisted recall of it. The present study asks this question for students' memories of academic uncertainty.

Applying the Constructive-Developmental definition of self-authorship consistently across participants and idiosyncratic memories is difficult to achieve using predetermined written instructions. We used a large language model (GPT 5.2-Thinking), instructed to move participants from Kegan's (1994) socialized mind toward a self-authoring mind, to generate the reframe. Large language models can be instructed within specified frameworks and have been used to support cognitive reappraisal and related psychological interventions, with structured interactions associated with reduced self-reported anxiety (Bhattacharjee et al., 2026; Held et al., 2025; Sharma et al., 2023). In this experiment, the customized model acted as an agent in the exchange rather than as a fixed script. For each of the participants' five memories, the agent received participants' descriptions of what made the uncertainty intense and what felt at stake in the memory, and then returned three reframes written for that memory. As a result, each participant received personalized reinterpretations of their own academic experience. The LLM agent was instructed to be a motivational companion that responded to the participant's academic uncertainty, and the experiment examines how the agent's response may affect participants' later memory of the uncertainty. In both conditions, STEM undergraduates recalled memories of academic uncertainty and received an LLM-generated response individualized to the content of their particular memory. In the Self-Authorship condition, the response reframed the uncertainty using developmental principles of self-authorship, whereas in the Control condition the model restated and elaborated the same uncertainty in first-person sentences without offering an alternative perspective. Participants in both conditions then retyped their selected response and completed four condition-specific elaboration prompts before re-describing the same memories one week later.

We tested two primary hypotheses. First, we hypothesized that a self-authoring reinterpretation of an autobiographical memory of academic uncertainty would reduce the uncertainty expressed when that same memory was described one week later. This was operationalized as a reduction in tentative, uncertain language relative to the Control condition. Second, we hypothesized that the content introduced through the LLM-generated reinterpretation would be incorporated into later recall of the specific memory episode. This was operationalized as novel reframe vocabulary appearing in the Session 2 memory description at rates exceeding

chance comparisons. Because both hypotheses concern the LLM agent's potential effect on later recall, we also tested whether any reduction in tentative language survived removing the LLM-generated vocabulary from the Session 2 narratives. Anxiety-related language, analytic language, and affective language were also examined as secondary outcomes.

Methods

Participants

Participants were young adults concentrating in STEM undergraduate studies, recruited by email from an introductory physics course and from the online recruitment platform SONA, which compensates participants with course credit or monetary payment. For SONA participants, STEM status was verified with an eligibility screener completed before sign-up. Respondents indicated whether they were a STEM major or prospective major, defined as any major in science, technology, engineering, or mathematics; psychology majors and courses were not counted as STEM, whereas neuroscience majors and courses were. Respondents who reported a STEM major, or no STEM major but at least two STEM courses, then listed all STEM courses they had taken, and the sign-up code for the study was released only to respondents who listed at least two courses. Respondents who reported no STEM major and fewer than two STEM courses were informed that they were not eligible. The screener also confirmed undergraduate status, and graduate students were screened out. Participants were randomly assigned to one of two conditions, a Self-Authorship condition or a Control condition (described below). We targeted 20 participants per condition, informed by prior experimental work on autobiographical memory updating (Speer et al., 2021). Fifty-three participants were randomized (Self-Authorship, $n = 28$; Control, $n = 25$). Exclusion criteria were failure to recall five memories of STEM-related uncertainty from the previous six months, failure to qualify as a STEM student, failure to return for the second session, and computer or technical problems that prevented completion. All exclusions occurred at data cleaning, when 11 participants (Self-Authorship, $n = 6$; Control, $n = 5$) were found to have no Session 2 memory narratives because they had not attended or completed Session 2, 10 with no Session 2 record and one who began Session 2 but recorded no memories; no participant was excluded under the other criteria. The analyzed sample comprised the 42 participants who

completed both sessions (Self-Authorship, $n = 22$; Control, $n = 20$). Demographic characteristics, overall and by condition, are presented in Table 1.

Ethics

This study was approved by the Tufts University Institutional Review Board under protocol STUDY00005471. All participants provided informed consent before taking part. The consent form disclosed that participants' written responses would be submitted to OpenAI's GPT-5.2 Thinking model to generate the reframes, and participants consented to this use of their responses.

Design and Procedure

The study used a two-session design with sessions separated by approximately one week (five to nine days), similar to Speer et al. (2021). Both sessions were administered online in Qualtrics, and no task in either session had a time limit. Session 1 lasted approximately 90 minutes and consisted of (1) a brief demographic survey, (2) structured recollection and examination of recent, personally experienced academic uncertainty, and (3) a guided, LLM-enabled reframing or Control exercise (Supplementary Note 1). The reframing exercise had two components in both conditions: selection of an LLM-generated reframe of each memory and a directed written elaboration on the selected reframe. The conditions were procedurally identical and differed only in the content of these components: in the Self-Authorship condition, the reframes recast the uncertainty as a manageable opportunity within a broader, self-authored perspective of choice, action, and effective response, and the elaboration directed participants toward motivation and next steps; in the Control condition, the reframes restated the uncertainty in synonymous terms without added perspective, and the elaboration kept participants within the unresolved uncertainty (see Reframing Intervention).

Memory Cueing and Elaboration

Participants first completed a cued autobiographical memory task, adapting an established procedure in which verbal prompts elicit specific events from participants' own lives (Crovitz & Schiffman, 1974). Previous studies have allowed participants to select cues from a larger list when generating autobiographical memories (Emmerdinger & Kuhbandner, 2018). In memory-reappraisal research, Speer et al. (2021) used life-event cues followed by brief written descriptions

and ratings of memory intensity and vividness. For the present study, cue content focused on uncertainty in STEM learning. Participants were shown a list of 30 life-event cues related to STEM uncertainty (e.g., "Missing a Core Concept"; Supplementary Note 2) and selected the five cues that triggered their most intense and vivid memories of uncertainty. The 30 event cues were presented in the same fixed order to all participants and participants then completed the corresponding memory elaboration in selection order. For each selected cue, participants verified through brief, one-click true/false follow-up questions whether the memory was intense, was still unresolved, reminded them of genuinely challenging uncertainty, and was distinct from previously selected memories. On each of the five trials, participants viewed one event cue and wrote three to five sentences describing what happened in that memory, followed by ratings of the uncertainty the memory evoked in the present moment (1 = not at all, 5 = extremely), the perceived probability that the uncertainty would be resolved within a week (yes / no / unsure), the present intensity of the uncertainty (1 = not at all, 5 = extremely intense), the vividness of the memory (1 = not vivid, 5 = extremely vivid), and the temporal direction of the uncertainty (experiencing something in the present / reliving something in the past / anticipating something in the future). Participants then examined the memory by writing two to three sentences identifying the source of its intensity ("Why does the uncertainty feel intense?") and two to three sentences identifying the core concern driving it ("What feels at stake to you in this situation?"), surfacing the deeper assumption underlying the uncertainty.

Reframing Intervention

At this point the conditions diverged. For each of the five selected memories, participants' responses describing the sources of their uncertainty were submitted to a custom LLM agent, which returned three candidate reframes. Both conditions used the same agent, implemented as a customized GPT (GPT-5.2 Thinking) using a zero-shot instruction prompt that engineered the output format to three numbered reframes. The LLM agent ran on OpenAI's ChatGPT platform from February 12, 2026 to August 17, 2026. Participants submitted their responses to the agent themselves by entering them in the custom GPT's chat box during the session, and the agent returned three numbered reframes in the same window. The instruction set for each condition is reproduced in Supplementary Note 1. The conditions differed in the instructions governing both the generated reframes and the subsequent participant elaboration. In the Self-Authorship

condition, the agent was instructed to rewrite each uncertainty as an agentic reframe that placed the participant in control of next steps. The instructions directed the agent to move the participant from Kegan's (1994) socialized mind toward a self-authoring mind without naming either. In the Control condition, the agent was instructed to restate the uncertainty from other angles in first-person sentences that emphasized the uncertainty the participant currently felt. The Control instructions prohibited solutions or therapeutic suggestions. Within their assigned condition, participants selected the response or responses that "most resonated" with them and retyped their preferred reframe in Qualtrics, with the option to combine responses or use their own words.

The second component of the manipulation was a directed elaboration on the selected reframe. In the Self-Authorship condition, participants described how the reframe made them feel more self-motivated, how it encouraged them emotionally, and one action they could take during the coming week to implement it. They also generated a short phrase or word of encouragement to remind themselves of the reframe when facing the uncertainty. In the Control condition, participants described how the reframe made them feel stuck, how it made them feel emotionally, and how they expected the uncertainty to affect them during the coming week. They then generated a short phrase or word identifying the situation's key unknowns. Both conditions asked participants to think about the coming week, but the Self-Authorship prompts emphasized how they could derive a sense of agency over their uncertainty, whereas the Control prompts maintained how their uncertainty might affect them. The elaboration prompts for each condition are reproduced in Supplementary Note 3.

In both conditions, participants then reported, for each memory, whether the reframe led them to examine and change their assumptions about the uncertainty event, whether they generated concrete and doable next steps, whether the agent provided a genuinely helpful way to see their situation differently, and how helpful they found the elaboration; these items recorded participants' evaluations of the task. Participants also reported how they selected their preferred reframe: typing one reframe word-for-word, pasting it, combining multiple reframes, using their own words inspired by the reframes, or another approach.

Session 2

Participants returned five to nine days after Session 1 to complete Session 2. They were asked to recall the same memories cued by the five event cues they had selected in Session 1 and again wrote narrative descriptions of those memories in three to five sentences.

Language Analysis

Participants' written memory narratives from Sessions 1 and 2 were analyzed using Linguistic Inquiry and Word Count (LIWC-22; Boyd et al., 2022). The primary linguistic outcome examined was tentative language, which measures the percentage of words signaling hesitation or uncertainty (e.g., maybe, perhaps, guess; Boyd et al., 2022; Tausczik & Pennebaker, 2010). We also examined analytic language, which is a standardized LIWC summary score derived from the Categorical-Dynamic Index (CDI; Pennebaker et al., 2014), with higher scores indicating more formal, categorical, and hierarchically organized language and lower scores indicating more personal, immediate, and narrative language. Anxiety language (emo_anx) was used to measure the frequency of words communicating fear, worry, and dread. We additionally examined positive and negative tone (tone_pos, tone_neg), which assesses sentiment, and positive and negative emotion (emo_pos, emo_neg), which assesses explicit emotion language. The LIWC models only score the memory descriptions before and after the intervention (“Please describe this specific memory of uncertainty in 3–5 sentences.”; "Please describe how you remember this event in 3–5 sentences."). Each dimension was scored separately for each of the five memories at both sessions. Within-memory change scores (Session 2 – Session 1) were used for descriptive figures, whereas inferential linear mixed-effects models used the session-level scores directly (see Statistics and Reproducibility).

LIWC Tentative Language Word-Count Sensitivity Analysis

Because LIWC scores are calculated from the proportion of words that pertain to each linguistic category, estimates from very short texts may be less stable. LIWC guidance warns that texts containing fewer than 25 words are often unreliable for linguistic analysis, as the LIWC2015 psychometric development analyses excluded texts containing fewer than 25 words (Pennebaker et al., 2015). To check the robustness of our results, we repeated the tentative language analysis as a sensitivity check after excluding the 10 narratives containing fewer than 25 words.

Tentative Language in the LLM-Generated Reframes

We tested whether the five LLM-generated reframes each participant selected in Session 1 differed in tentative language by condition. Reframe text was scored with LIWC-22 for tentative language, expressed as the percentage of words signaling hesitation or uncertainty. We modeled tentative language as a function of condition using a linear mixed-effects model with a random intercept for participant to account for the five reframes contributed by each participant. The estimated marginal means and the between-condition contrast for the Language Outcomes were derived from the fit model.

Statistics and Reproducibility

We evaluated the language outcomes via linear mixed-effect models, which allowed us to account for the hierarchical, repeated measures structure of the data, i.e., five memories per participant measured at two time points, and to retain all observations without averaging over memories. We modeled each outcome as a function of Condition (Self-Authorship vs. Control), Session (Session 1 vs. Session 2), and their interaction. We included random intercepts for participant and for memory nested within participant: Outcome ~ Condition × Session + (1|Participant) + (1|Participant:Memory). The Condition × Session interaction was the primary independent variable of theoretical interest, indicating whether the change in each outcome from Session 1 to Session 2 varied by condition. We fit the models using maximum restricted likelihood estimation in R 4.4.2 using the lme4 (1.1-37) and lmerTest (3.1-3) packages. Type III tests and Satterthwaite degrees of freedom were used to assess significance. Simple effects, i.e., the change in each outcome from Session 1 to Session 2 for both conditions, were assessed using the emmeans package (1.11.0) with Satterthwaite degrees of freedom.

We did not perform an a priori power analysis, since such an analysis would require a specific target effect size that could be based on previous research or determined by the minimum effect of interest (Lakens, 2022). As far as we are aware, there has been no published work that examines the impact of a self-authorship-based reframe on the tentative language used in subsequent recall of autobiographical memories using LIWC as a dependent measure. In the

absence of an empirically grounded target, we view this study as exploratory, designed to provide an initial estimate of the direction and general magnitude of the effect.

To quantify the magnitude of the Condition × Session interaction, we standardized the estimated coefficient from the model by dividing it by the total standard deviation of tentative language predicted by the model. We calculated the total standard deviation of tentative language as the square root of the sum of the estimated variances.

The post-intervention evaluation ratings were only collected during Session 1 and were not included in the longitudinal model. Each rating was averaged across the five memories recalled by each participant. We then compared the ratings of the Self-Authorship and Control groups for each rating using Welch's two-sample t-tests and Cohen's d as a standardized effect size estimate. Similarly, we used Welch's two-sample t-test and Cohen's d to compare the proportions of each participant's memories that contained a verbatim reproduction of a reframe, either typed verbatim or pasted directly, for the reframe selection question.

Reframe Lexical Intrusion Analysis

To evaluate whether the content of the LLM-generated reframes appeared in the subsequent recall of memories, we calculated a lexical intrusion score for each memory. To ensure that the lexical intrusion scores captured actual word repetition rather than superficial matches, we processed the text of each memory as follows: lowercased; removed punctuation, numbers, and stop words from the spaCy English list (e.g., "the," "and," "of"); split into word tokens of three or more letters; and stemmed to ensure that variations such as "learn," "learned," and "learning" were treated as identical lexical items (Manning et al., 2008; Porter, 1980). For each memory, novel reframe vocabulary was defined as stems appearing in the selected LLM-generated reframe but absent from both the participant’s Session 1 narrative and the texts originally submitted to the LLM. The intrusion rate was the proportion of these novel stems that appeared in the Session 2 narrative.

Because there was a possibility that overlap occurred due to random chance, the observed rate of intrusions were evaluated based upon two different comparisons: the participant's four remaining reframes, controlling for individual vocabulary and the task at hand, and 20 reframes randomly selected from other participants, controlling for vocabulary common to all participants.

For each participant, we calculated how often words from each memory's own reframe appeared at Session 2 than words from the comparison reframes, averaged this difference across memories, and tested whether the average was greater than zero.

Tentative Language After Removing Reframe Vocabulary

To test whether the tentative language effect depended on wording copied from the LLM-generated reframes, we removed the reframe words identified as lexical intrusions from each Session 2 narrative. The edited Session 2 narratives were rescored with LIWC-22, and the tentative language model described in Statistics and Reproducibility was refit with the original Session 1 scores.

Results

Sample Characteristics

Of the participants randomized across the two conditions, 42 completed both sessions and were included in the analyses (Self-Authorship, $n = 22$; Control, $n = 20$; Table 1). The sample was primarily in its first year of study (18 of 42; 42.9%) and was predominantly women (32 of 42; 76.2%), with a mean age of 19.6 years ($SD = 1.2$). The most frequently reported racial/ethnic identifications were White (19 of 42; 45.2%) and Asian (17 of 42; 40.5%), and 7 of 42 participants (16.7%) identified as Hispanic or Latino. The two conditions were broadly similar in age, gender, racial/ethnic composition, and grade level (Table 1).

Across the analyzed memory narratives, in the Control condition, Session 1 narratives averaged 80.9 words ($SD = 33.8$, range = 23–207) and Session 2 narratives averaged 68.8 words ($SD = 26.2$, range = 11–135). In the Self-Authorship condition, Session 1 narratives averaged 75.7 words ($SD = 27.0$, range = 25–182) and Session 2 narratives averaged 67.0 words ($SD = 27.7$, range = 21–165). Each participant contributed five narratives per session, yielding 100 narratives per session in the Control condition and 110 per session in the Self-Authorship condition (Table 2).

Descriptive statistics for the language variables of interest are presented in Table 3. Tentative language in the initial memory narrative was similar across conditions at Session 1 (Self-

Authorship: M = 3.54, SD = 2.91; Control: M = 3.08, SD = 2.48), but was lower in the Self-Authorship condition at Session 2 (M = 2.87, SD = 2.40) while increasing slightly in Control (M = 3.43, SD = 2.99). Analytic language decreased from M = 40.01 (SD = 24.46) to M = 32.78 (SD = 28.17) in Control, whereas it remained comparatively stable in Self-Authorship (Session 1: M = 36.95, SD = 25.75; Session 2: M = 36.41, SD = 26.82). Anxiety-related language increased in both conditions, from M = 0.30 (SD = 0.77) to M = 0.75 (SD = 1.28) in Control and from M = 0.42 (SD = 0.78) to M = 0.64 (SD = 1.13) in Self-Authorship. Means and standard deviations for positive and negative emotion and tone are also reported in Table 3.

Post-Intervention Ratings

Participants' evaluations of the task were collected after each Session 1 reframing exercise and averaged across the five memories (Table 4; Figure 1). We compared conditions with Welch's two-sample t tests, using Cohen's d as a standardized effect size. Self-Authorship participants rated the task more favorably than Control participants on all four items. Participants in the Self-Authorship condition rated the reframe as more helpful than did those in the Control condition (Self-Authorship M = 4.13, SD = 1.07; Control M = 2.95, SD = 1.12), a large and statistically reliable difference, $t(39.2) = 3.48$, $p = .001$, $d = 1.08$. They rated the elaboration exercise as substantially more helpful (Self-Authorship M = 4.48, SD = 1.00; Control M = 3.13, SD = 1.03), also a large and statistically reliable difference, $t(39.4) = 4.31$, $p < .001$, $d = 1.33$. They likewise reported examining and changing their assumptions about the uncertainty to a substantially greater degree (Self-Authorship M = 4.03, SD = 1.03; Control M = 2.72, SD = 1.02), $t(39.7) = 4.13$, $p < .001$, $d = 1.27$. They also reported generating concrete and doable next steps to a greater degree (Self-Authorship M = 4.49, SD = 1.03; Control M = 2.81, SD = 0.97), $t(39.9) = 5.42$, $p < .001$, $d = 1.67$.

On the reframe-selection item, participants typed a single reframe word-for-word or pasted it for 45 of 110 Self-Authorship memories (40.9%) and 42 of 100 Control memories (42.0%), wrote their own words inspired by the reframes for 36 (32.7%) and 22 (22.0%), combined multiple reframes for 25 (22.7%) and 32 (32.0%), and described another approach for 4 (3.6%) and 4 (4.0%) (Supplementary Table 1). The proportion of each participant's memories in which a reframe was reproduced verbatim did not differ between conditions (Self-Authorship M = .41, SD = .38; Control M = .42, SD = .32), $t(39.8) = -0.10$, $p = .920$, $d = -0.03$.

Change in Tentative Language

We modeled tentative language as a function of Condition, Session, and their interaction (Table 5). The Condition x Session interaction on tentative language was statistically significant, b = -1.02, SE = 0.45, $t(208)$ = -2.25, p = .026; Type III $F(1, 208)$ = 5.04, p = .026. The interaction coefficient was −0.37 when standardized by the model-implied total standard deviation of tentative language (2.71). Neither the main effect for Condition, F(1, 40) = 0.02, p = .894, nor the main effect for Session, F(1, 208) = 0.48, p = .487, were statistically significant. When decomposing the interaction (Figure 2), tentative language decreased from Session 1 to Session 2 in the Self-Authorship condition (−0.66 points, t(208) = −2.13, p = .034), whereas there was an increase in tentative language in the Control condition (+0.35 points), however this increase was not reliable, t(208) = 1.07, p = .286.

LIWC Word-Count Sensitivity Analysis

Because 10 of 420 narratives (2.4%) fell below the 25-word threshold used in prior LIWC psychometric analyses, we repeated the tentative analyses after excluding narratives containing fewer than 25 words. The Condition × Session interaction for tentative language remained statistically significant, $F(1, 205.6) = 5.40$, $p = .021$ (Supplementary Table 2), and was comparable to the result obtained in the full sample, $F(1, 208) = 5.04$, $p = .026$.

Analysis of Tentative Language in the LLM-Generated Reframes

In order to examine if the reduction in tentative language that was observed could be attributed to the type of language used by the LLMs to produce the reframes themselves, we conducted an analysis on the reframe text provided in Session 1. The Self-Authorship LLM-produced reframes contained significantly fewer instances of tentative language as compared to the Control LLM-produced reframes, $F(1, 40) = 76.04$, $p < .001$ (Supplementary Table 3). The estimated marginal means were 1.95 (SE = 0.30, 95% CI [1.35, 2.55]) for Self-Authorship and 5.67 (SE = 0.31, 95% CI [5.05, 6.30]) for Control, a difference of 3.72 percentage points with a standard error of 0.43, $t(40) = 8.72$, $p < .001$. Additionally, the direction of this difference was consistent with the Condition × Session interaction observed in the participants' own memory narratives, as those who had received the less tentative Self-Authorship reframes subsequently used less tentative language when recounting these memories.

Reframe Lexical Intrusion Analysis

The reframes added an average of 11.4 content words to each participant's vocabulary, which had not been used to describe their memory during Session 1 or in the text that they provided to the agent (Figure 3). At the time of Session 2 one week later, these novel reframe words appeared in their corresponding Session 2 memory description at rates above both chance baselines. Relative to the between-participant baseline, the average difference in percent between the reframe provided by the participant and the sample baseline was +1.54 percentage points, $t(41) = 3.23$, $p = .002$, $d = 0.50$. Relative to the within-participant baseline, the average difference in percent was +0.96 percentage points, $t(41) = 2.16$, $p = .036$, $d = 0.33$. The effect of the reframe existed in both conditions (Control: +2.08 points, $t(19) = 2.45$, $p = .024$, $d = 0.55$; Self-Authorship: +1.06 points, $t(21) = 2.16$, $p = .042$, $d = 0.46$) and did not significantly differ between them, Welch $t(30.7) = 1.04$, $p = .305$, 95% CI for the difference [−0.97, 3.01] percentage points. In comparison to the within-participant baseline, the above-chance gap was +0.99 points in Control and +0.93 points in Self-Authorship, and the conditions once again did not significantly differ, Welch $t(33.8) = 0.06$, $p = .951$, 95% CI [−1.80, 1.91] (Figure 4; Table 6). These results indicate that in both conditions participants included vocabulary from their selected reframe into their subsequent descriptions of their specific memory.

Change in Tentative Language After Removing LLM-Reframe Vocabulary

Because the Self-Authorship reframes contained less tentative language than the Control reframes and participants in both conditions reused reframe vocabulary at Session 2, the tentative language effect could reflect wording copied from the reframes. Novel reframe words were present in 64 of the 210 Session 2 narratives. After each memory's novel reframe words were removed from its Session 2 narrative, the Condition × Session interaction remained significant, $b = -1.00$, $SE = 0.45$, $t(208) = -2.20$, $p = .029$; Type III $F(1, 208) = 4.85$, $p = .029$ (Supplementary Table 4), comparable to the estimate from the full narratives, $b = -1.02$, $p = .026$. Neither the main effect of condition, $F(1, 40) = 0.01$, $p = .915$, nor the main effect of session, $F(1, 208) = 0.52$, $p = .474$, was significant. Tentative language declined from Session 1 to Session 2 in the Self-Authorship condition (-0.66 points, $t(208) = -2.12$, $p = .035$), whereas the change in the Control condition was positive and not reliable (+0.34 points, $t(208) = 1.03$, $p = .306$). The decline following the Self-Authorship reframe therefore did not depend on vocabulary carried over from the reframe.

Secondary Linguistic Outcomes

Across the secondary linguistic outcomes, there was no reliable evidence of condition-specific change (Supplementary Table 5). Analytic language showed no significant Condition × Session interaction, $F(1, 208) = 3.03$, $p = .083$, despite a decline in the Control condition and little change following Self-Authorship. Positive and negative emotion and tone likewise showed no significant interactions (all ps $\geq .468$), although negative emotion and negative tone increased across sessions. Anxiety-related language also increased across sessions, $F(1, 208) = 16.37$, $p < .001$, with no reliable difference between conditions, $p = .162$. Overall, Self-Authorship did not differentially alter analytic language, emotional valence, or anxiety-related language relative to Control.

Discussion

The present study asked whether reinterpreting an autobiographical memory of academic uncertainty changes how that same episode is described one week later. The lead finding is that Self-Authorship reframing reduced the tentative language participants used when they later described the same memory of academic uncertainty, whereas tentative language did not reliably change following the Control reframe. Two results are consistent with a change in how participants described their memories: first, we found no difference in the rate of reframe vocabulary that reappeared in both conditions at the second memory description, but we did find that tentative language declined only after Self-Authorship reframing; second, we also found that the interaction between condition and session remained statistically significant and nearly unchanged in effect size after each memory's novel reframe vocabulary was removed from its Session 2 narrative. Because novel vocabulary excluded stop-words and words the participant had already used, this removal does not rule out the reuse of tentative words that also appeared in the reframes. The reduction in tentative language at the later recall appears to be an effect of Self-Authorship reframing on how participants expressed uncertainty about these memories, although the present design does not establish the precise mechanism through which it occurred.

The lexical intrusion results further suggest that reframe content remained tied to the memory it addressed: LLM-generated vocabulary intrusion exceeded both the between-participant

comparison and the comparison drawn from each participant's other reframes. This specificity suggests that there was more than a general shift in participants' writing style, but it leaves open how Self-Authorship reframing reduced tentative language. Participants may have reorganized their understanding of the uncertainty, as the adult developmental literature account predicts (Kegan, 1994). Alternatively, because Self-Authorship reframes contained substantially less tentative language than Control reframes, participants may have incorporated less tentative phrasing by using their own words inspired by the LLM agent reframe.

These results are consistent with a broader literature that shows how information encountered after an event can influence how that event is later remembered. Post-event information can be seen in later reports of an experience (Loftus, 2005), and retrieval of an event prior to receiving new information can increase the probability that the new information will be incorporated into memory (Chan et al., 2009). Additionally, research on retelling also demonstrates that later memory relies upon not only on which details are rehearsed, but also on the perspective used to organize the retelling (Marsh, 2007). Specifically, Tversky and Marsh (2000) showed that using a biased perspective while retelling an event affected later recall even when the retelling included little specific information from the original event. Reappraisal may also have effects that persist over time. Denny et al. (2015) found that responses remained altered for one week later for stimuli participants had previously reappraised, but not for similar or novel stimuli that were exposed in the same manner. Collectively, this work illustrates that later remembering can include information introduced after an event, what is rehearsed during retrieval, and the perspective utilized to interpret the experience. The present study extends this work to autobiographical memories of academic uncertainty by evaluating whether changing how participants interpret uncertainty is reflected when they later describe that same experience.

The results also inform how generative AI agents function within educational contexts. One week after a single session, words generated by an agent for a specific memory were present when students later reported that memory in both conditions and at similar rates. This supports evidence that conversational agents can impact what individuals later report remembering (Chan et al., 2024; Pataranutaporn et al., 2025) and extends the findings beyond episodic details to students' descriptions of their own academic experiences. An agent's interpretation of a student's uncertainty persists outside the interaction that created it, as evidenced by its appearance in the

student's later description of the experience. Whether that is advantageous or not partly depends on what the agent was instructed to do. Both conditions used the same model and the prompting format and left similarly sized lexical traces, however, only the Self-Authorship procedure resulted in reduced tentative language. The difference between the two conditions was the instruction customization of each LLM agent rather than access to the agent, which is consistent with evidence indicating that the learning effects associated with LLM assistance depend on the constraints placed on it (Bastani et al., 2025; Kestin et al., 2025).

Concerns related to generative AI in education revolve around students offloading cognitive labor and retaining less of it (Darvishi et al., 2024; Fan et al., 2025; Grinschgl et al., 2021; Kosmyna et al., 2025). The outcome measured in the present study was unassisted memory description one week following the removal of the LLM agent. The reduction in tentative language survived withdrawal of LLM agent wording, which suggests that the effect was carried by the students' own language. This finding does not mean that students fully adopted the self-authorship perspective within their own thinking. However, it does indicate that the effect was not dependent on the LLM agent's language alone or on its continued presence in the experiment. Simultaneously, the results seen from the Control condition provides a cautionary note. An LLM agent that simply repeated students' uncertainty without adaptively reinterpreting it also led to lexical intrusions at a similar rate as did the Self-Authorship condition, but it did not result in a reduction in tentative language. When agents serve as socio-emotional companions, repeating a student's academic difficulty appears to preserve their sense of uncertainty as strongly as before, which highlights the need to consider how an agent is designed and engineered to interact with students.

Additionally noteworthy, the Self-Authorship condition did not produce a greater increase in positive emotion or positive tone, or a greater decrease in negative emotion or negative tone than the Control condition. This distinguishes the tentative language finding from the valence change reported by Speer et al. (2021), whose positive-meaning intervention increased positive emotion and positive content in later recollections. In our study, negative emotion and negative tone instead increased across sessions in both conditions. Anxiety-related language also increased across the sample. The increase was numerically smaller following the Self-Authorship reframe, but the interaction was not significant and does not provide reliable evidence that anxiety changed

differently between conditions, nor that the Self-Authorship reframe attenuated anxiety language. Analytic language also did not demonstrate a statistically significant condition-specific change. While these findings are not statistically reliable, they may be explained by the Constructive-Developmental definition of self-authorship, whose cultivation is expected to affect how individuals interpret their memories rather than necessarily influence the valence of their memories (Baxter Magolda, 2008).

There are other possible explanations for this reduction in tentative language. The most direct is that the conditions also differed in more than the generated reframe. Both groups received personalized LLM-generated responses, retyped a selected response, and completed the same number of elaboration steps. However, the content of those elaboration prompts differed. They varied in valence, temporal orientation, and the kinds of reflection and action they asked participants to consider. The more encouraging Self-Authorship prompts could have led participants to describe their memories more positively, with less tentative language as a by-product. The lack of a condition-specific change in emotional valence suggests that the Self-Authorship procedure did not produce this shift. The study cannot determine how much of the observed difference came specifically from the generated Self-Authorship reframe rather than different features embedded in the nature of the two reframes.

The Session 2 retrieval procedure provides another possible explanation. Participants received the same memory cue they had encountered during Session 1, which may have inadvertently reinstated the context in which they had previously encountered the reframe. In an eyewitness-memory study, mentally reinstating the context in which misleading information had been presented increased susceptibility to that information on a recognition test one week later (Roebers & McConkey, 2003). In our current experimental design, due to the repetition of the cue in both conditions, reinstatement can account for why reframe vocabulary was accessible in Session 2, however, it cannot solely explain why tentative language decreased only in the Self-Authorship condition. Re-use of reframe vocabulary does not necessarily imply that participants recalled the original event differently. Participants may have intentionally used a later interpretation to describe the event instead of recalling that interpretation as something they had thought or felt when the event occurred. This issue relates to source monitoring, which refers to determining whether information is derived from an individual's own experience or from

something encountered after the event (Johnson et al., 1993). Since we did not ask participants to judge the source of their memories, the occurrence of LLM-generated vocabulary intrusions establishes that the reframe content was incorporated into their subsequent memory descriptions, but this does not confirm that they experienced it as part of the original memory. We are unsure of the underlying mechanism of their memory description change given that we did not ask participants about the source of the language they used in Session 2.

The interaction between condition and session predicting tentative language remained significant after excluding narratives shorter than 25 words, strengthening confidence that the effect was not driven by unstable scores in very short texts. Together, the findings show that brief Self-Authorship reframing can change how students describe academic uncertainty one week later. The effect persisted after removing words identified as reused from the reframes. Given evidence that how people remember academic experiences can influence their subsequent task appraisals and performance (Selimbegović et al., 2011; Selimbegović et al., 2016), reduced expression of uncertainty during recall may have implications for how students appraise and perform on later academic tasks. Future research should test whether Self-Authorship reframing affects these outcomes and whether any effects are associated with reductions in tentative language at recall.

Several limitations restrict what we can conclude from this study. First, our sample was small and consisted solely of students from a single university. Our sample was also homogenous with regards to age and educational context. Given these constraints, the precision of our estimates is limited and it is unclear whether similar patterns will emerge among older students, students attending different educational institutions, or individuals experiencing varying types of uncertainty. Larger studies with more diverse samples are required to determine how widely these findings apply.

Second, the lexical incorporation measure captures only one type of similarity between a reframe and a later memory description. In our analysis, we stemmed the content words from both the reframe and the later memory description and searched for identical word forms across the two texts. While this method can identify direct reuse and closely related forms of the same word, it cannot detect a paraphrase using entirely new vocabulary. An individual could adopt the meaning of a reframe using entirely new vocabulary and receive no credit for doing so. Conversely, an individual could use a word form that appeared in a previous reframe but have it contain an entirely

different meaning in their later memory description. Future research could combine lexical matching with semantic similarity measures and human coding to distinguish between direct word reuse, paraphrased incorporation, and broader changes in how participants interpret their memories.

Third, since each LLM-generated reframe was created specifically for each participant and memory, the reframes varied across participants and memories. Personalization was an intentional aspect of the intervention, but it also implies that participants received neither identical language nor identical amounts of each self-authorship component. We did observe a clear difference between conditions in terms of the amount of tentative language present in the generated reframes, however, the present study cannot determine which characteristics of those reframes contributed to later changes. Future studies could standardize reframe length and structure while manipulating specific self-authorship components. Reframe content and elaboration instructions could also be manipulated independently to evaluate the relative contributions of each aspect of the reappraisal task. Some questions cannot be answered within this experimental design. Specifically, the study cannot establish whether a general positive reframe would produce an equivalent reduction in tentative language as the Self-Authorship reframe.

Finally, we studied how participants subsequently described their memories rather than whether the intervention influenced their behavior or generalized to new experiences. We do not know whether participants behaved differently when they later encountered academic uncertainty, whether they applied their reinterpretation to a new incident, or whether they maintained the effects for longer than one week. Investigating these questions would necessitate behavioral outcomes, new memories that were not directly reframed, and extended follow-up periods. Future studies could also inquire about participants' perceptions regarding the origin of certain ideas expressed in their later memory descriptions and manipulate retrieval cues employed during follow-up. These measures would assist us in determining whether reframe content became part of how participants remembered the incident, remained accessible as a distinct interpretation thereof, or was retrieved because of the Session 2 procedure prompting participants to recollect the prior exercise.

A reframe generated by an LLM agent after participants recalled an episode of academic uncertainty was still evident when they described that same memory one week later. Participants reused vocabulary that had been introduced specifically for that memory, and participants in the

Self-Authorship condition described their memories with less tentative, uncertain language than they had in Session 1. The memories did not become more positive in emotional valence relative to the Control condition. The findings suggest that a recently adopted interpretation of memories of academic uncertainty can remain present when a specific autobiographical experience is later described. This extends prior work on memory updating by showing that later recall may reflect not only new information about an event, but also a new way of expressing oneself in relation to the experience. The study also identifies a tractable method to test this process at the level of individual autobiographical memories by combining memory-specific reframing, delayed recall, linguistic change, and lexical incorporation. Future work can build on this design to test whether these changes persist over longer intervals, generalize to new experiences of uncertainty, and influence how students respond when uncertainty arises again. More broadly, the findings point to autobiographical memory as a promising target for interventions designed to change how learners make sense of difficult academic experiences. For generative AI agents in education, the findings show that a brief interaction with an instructed agent can change how learners later describe academic uncertainty. Delayed recall combined with lexical tracing offers a way to evaluate an LLM agent's influence on students' academic memories.

Data Availability

The de-identified data that support the findings of this study are available on Figshare at https://figshare.com/s/6172d39eac7ca378d7ca: participant-level ratings and reframe-selection responses, LIWC-22 scores for every memory narrative (original and with novel reframe words removed) and every selected reframe, lexical intrusion scores by memory, and the demographic summary in Table 1. Narratives, free-text responses, the LLM-generated reframes, row-level demographics, timestamps, and study identifiers are withheld to protect participants and are available from the corresponding author on reasonable request.

Code Availability

A Quarto document (reproduce_results.qmd) that reproduces each reported statistic, table, and data figure from the deposited data, including the mixed-effects models, the lexical intrusion tests, and the reframe-vocabulary removal analysis, is available at https://figshare.com/s/6172d39eac7ca378d7ca with a rendered copy. The scripts that computed the lexical intrusion scores and the reframe-word-removed narratives from the raw text are available from the corresponding author on request because they require the withheld narratives.

Acknowledgements

This material is based upon work supported by the U.S. National Science Foundation under Grant No. 2428640 (Growing Convergence Research). The funder played no role in study design, data collection, analysis and interpretation of data, or the writing of this manuscript.

Author Contributions

M.R.: Conceptualization, Methodology, Investigation, Formal Analysis, Writing – original draft.

A.T.: Conceptualization, Methodology, Writing.

Competing Interests

The authors declare no financial or non-financial competing interests.

References

Abbas, M., Jam, F. A., & Khan, T. I. (2024). Is it harmful or helpful? Examining the causes and consequences of generative AI usage among university students. *International Journal of Educational Technology in Higher Education, 21*, Article 10. https://doi.org/10.1186/s41239-024-00444-7

Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. Cambridge University Press.

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences, 122*(26), Article e2422633122. https://doi.org/10.1073/pnas.2422633122

Baxter Magolda, M. B. (2008). Three elements of self-authorship. *Journal of College Student Development, 49*(4), 269–284. https://doi.org/10.1353/csd.0.0016

Bhattacharjee, A., Suh, J., Chandra, M., & Hernandez, J. (2026). User perceptions of an LLM-based chatbot for cognitive reappraisal of stress: Feasibility study. arXiv. https://doi.org/10.48550/arXiv.2601.00570

Boyd, R. L., Ashokkumar, A., Seraj, S., & Pennebaker, J. W. (2022). *The development and psychometric properties of LIWC-22*. University of Texas at Austin. https://www.liwc.app

Carleton, R. N. (2016). Into the unknown: A review and synthesis of contemporary models involving uncertainty. *Journal of Anxiety Disorders, 39*, 30–43. https://doi.org/10.1016/j.janxdis.2016.02.007

Chan, J. C. K., Thomas, A. K., & Bulevich, J. B. (2009). Recalling a witnessed event increases eyewitness suggestibility: The reversed testing effect. *Psychological Science, 20*(1), 66–73. https://doi.org/10.1111/j.1467-9280.2008.02245.x

Chan, S., Pataranutaporn, P., Suri, A., Zulfikar, W., Maes, P., & Loftus, E. F. (2024). Conversational AI powered by large language models amplifies false memories in witness interviews. arXiv. https://doi.org/10.48550/arXiv.2408.04681

Chen, X. (2013). *STEM attrition: College students' paths into and out of STEM fields* (NCES 2014-001). National Center for Education Statistics, Institute of Education Sciences, U.S. Department of Education.

Colombo, B., Cohen, J., Currier, R., & Caravita, S. C. (2026). Using an app to promote emotional regulation and well-being in young adults. A pilot study. *Journal of Behavioral and Cognitive Therapy, 36*(2), Article 100573. https://doi.org/10.1016/j.jbct.2026.100573

Conway, M. A. (2005). Memory and the self. *Journal of Memory and Language, 53*(4), 594–628. https://doi.org/10.1016/j.jml.2005.08.005

Conway, M. A., & Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the self-memory system. *Psychological Review, 107*(2), 261–288. https://doi.org/10.1037/0033-295X.107.2.261

Creamer, E. G., Baxter Magolda, M. B., & Yue, J. (2010). Preliminary evidence of the reliability and validity of a quantitative measure of self-authorship. *Journal of College Student Development, 51*(5), 550–562. https://doi.org/10.1353/csd.2010.0010

Creamer, E. G., & Laughlin, A. (2005). Self-authorship and women's career decision making. *Journal of College Student Development, 46*(1), 13–27. https://doi.org/10.1353/csd.2005.0002

Crovitz, H. F., & Schiffman, H. (1974). Frequency of episodic memories as a function of their age. *Bulletin of the Psychonomic Society, 4*(5B), 517–518. https://doi.org/10.3758/BF03334277

Darvishi, A., Khosravi, H., Sadiq, S., Gašević, D., & Siemens, G. (2024). Impact of AI assistance on student agency. *Computers & Education, 210*, Article 104967. https://doi.org/10.1016/j.compedu.2023.104967

Denny, B. T., Inhoff, M. C., Zerubavel, N., Davachi, L., & Ochsner, K. N. (2015). Getting over it: Long-lasting effects of emotion regulation on amygdala response. *Psychological Science, 26*(9), 1377–1388. https://doi.org/10.1177/0956797615578863

Elsey, J. W. B., Van Ast, V. A., & Kindt, M. (2018). Human memory reconsolidation: A guiding framework and critical review of the evidence. *Psychological Bulletin, 144*(8), 797–848. https://doi.org/10.1037/bul0000152

Emmerdinger, K. J., & Kuhbandner, C. (2018). Testing memories of personally experienced events: The testing effect seems not to persist in autobiographical memory. *Frontiers in Psychology, 9*, Article 810. https://doi.org/10.3389/fpsyg.2018.00810

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology, 56*(2), 489–530. https://doi.org/10.1111/bjet.13544

Grinschgl, S., Papenmeier, F., & Meyerhoff, H. S. (2021). Consequences of cognitive offloading: Boosting performance but diminishing memory. *Quarterly Journal of Experimental Psychology, 74*(9), 1477–1496. https://doi.org/10.1177/17470218211008060

Held, P., Pridgen, S. A., Szoke, D. R., Chen, Y., Akhtar, Z., & Amin, D. (2025). AI-facilitated cognitive reappraisal via Socrates 2.0: Mixed methods feasibility study. *JMIR Mental Health, 12*, Article e80461. https://doi.org/10.2196/80461

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114*(1), 3–28. https://doi.org/10.1037/0033-2909.114.1.3

Jordan, M. E., & McDaniel, R. R. (2014). Managing uncertainty during collaborative problem solving in elementary school teams: The role of peer influence in robotics engineering activity. *Journal of the Learning Sciences, 23*(4), 490–536. https://doi.org/10.1080/10508406.2014.896254

Kegan, R. (1994). *In over our heads: The mental demands of modern life*. Harvard University Press.

Kestin, G., Miller, K., Klales, A., Milbourne, T., & Ponti, G. (2025). AI tutoring outperforms in-class active learning: An RCT introducing a novel research-based design in an authentic educational setting. *Scientific Reports, 15*, Article 17458. https://doi.org/10.1038/s41598-025-97652-6

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. arXiv. https://doi.org/10.48550/arXiv.2506.08872

Kross, E., & Ayduk, O. (2008). Facilitating adaptive emotional analysis: Distinguishing distanced-analysis of depressive experiences from immersed-analysis and distraction.

Personality and Social Psychology Bulletin, 34(7), 924–938. https://doi.org/10.1177/0146167208315938

Kross, E., & Ayduk, O. (2017). Self-distancing: Theory, research, and current directions. In J. M. Olson (Ed.), *Advances in experimental social psychology* (Vol. 55, pp. 81–136). Academic Press. https://doi.org/10.1016/bs.aesp.2016.10.002

Lakens, D. (2022). Sample size justification. *Collabra: Psychology, 8*(1), Article 33267. https://doi.org/10.1525/collabra.33267

Lamnina, M., & Chase, C. C. (2019). Developing a thirst for knowledge: How uncertainty in the classroom influences curiosity, affect, learning, and transfer. *Contemporary Educational Psychology, 59*, Article 101785. https://doi.org/10.1016/j.cedpsych.2019.101785

Laughlin, A., & Creamer, E. G. (2007). Engaging differences: Self-authorship and the decision-making process. *New Directions for Teaching and Learning, 2007*(109), 43–51. https://doi.org/10.1002/tl.264

Levine, L. J. (1997). Reconstructing memory for emotions. *Journal of Experimental Psychology: General, 126*(2), 165–177. https://doi.org/10.1037/0096-3445.126.2.165

Levine, L. J., Prohaska, V., Burgess, S. L., Rice, J. A., & Laulhere, T. M. (2001). Remembering past emotions: The role of current appraisals. *Cognition and Emotion, 15*(4), 393–417. https://doi.org/10.1080/02699930125955

Loftus, E. F. (2005). Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning & Memory, 12*(4), 361–366. https://doi.org/10.1101/lm.94705

Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to information retrieval*. Cambridge University Press. https://doi.org/10.1017/CBO9780511809071

Marsh, E. J. (2007). Retelling is not the same as recalling: Implications for memory. *Current Directions in Psychological Science, 16*(1), 16–20. https://doi.org/10.1111/j.1467-8721.2007.00467.x

Massazza, A., Kienzler, H., Al-Mitwalli, S., Tamimi, N., & Giacaman, R. (2023). The association between uncertainty and mental health: A scoping review of the quantitative literature. *Journal of Mental Health, 32*(2), 480–491. https://doi.org/10.1080/09638237.2021.2022620

Moffett, J., Hammond, J., Murphy, P., & Pawlikowska, T. (2021). The ubiquity of uncertainty: A scoping review on how undergraduate health professions' students engage with uncertainty. *Advances in Health Sciences Education, 26*(3), 913–958. https://doi.org/10.1007/s10459-021-10028-z

Morewedge, C. K., Gilbert, D. T., & Wilson, T. D. (2005). The least likely of times: How remembering the past biases forecasts of the future. *Psychological Science, 16*(8), 626–630. https://doi.org/10.1111/j.1467-9280.2005.01585.x

Pataranutaporn, P., Archiwaranguprok, C., Chan, S. W. T., Loftus, E., & Maes, P. (2025). Synthetic human memories: AI-edited images and videos can implant false memories and distort recollection. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1–20). Association for Computing Machinery. https://doi.org/10.1145/3706598.3713697

Pekrun, R., Lichtenfeld, S., Marsh, H. W., Murayama, K., & Goetz, T. (2017). Achievement emotions and academic performance: Longitudinal models of reciprocal effects. *Child Development, 88*(5), 1653–1670

Pennebaker, J. W., Boyd, R. L., Jordan, K., & Blackburn, K. (2015). *The development and psychometric properties of LIWC2015*. University of Texas at Austin. https://doi.org/10.15781/T29G6Z

Pennebaker, J. W., Chung, C. K., Frazee, J., Lavergne, G. M., & Beaver, D. I. (2014). When small words foretell academic success: The case of college admissions essays. *PLOS ONE, 9*(12), Article e115844. https://doi.org/10.1371/journal.pone.0115844

Pennebaker, J. W., & Francis, M. E. (1996). Cognitive, emotional, and language processes in disclosure. *Cognition and Emotion, 10*(6), 601–626. https://doi.org/10.1080/026999396380079

Pillemer, D. B. (2003). Directive functions of autobiographical memory: The guiding power of the specific episode. *Memory, 11*(2), 193–202. https://doi.org/10.1080/741938208

Pillemer, D. B., Picariello, M. L., Law, A. B., & Reichman, J. S. (1996). Memories of college: The importance of specific educational episodes. In D. C. Rubin (Ed.), *Remembering our past: Studies in autobiographical memory* (pp. 318–338). Cambridge University Press. https://doi.org/10.1017/CBO9780511527913.013

Porter, M. F. (1980). An algorithm for suffix stripping. *Program, 14*(3), 130–137. https://doi.org/10.1108/eb046814

Roebers, C. M., & McConkey, K. M. (2003). Mental reinstatement of the misinformation context and the misinformation effect in children and adults. *Applied Cognitive Psychology, 17*(4), 477–493.

Rubin, D. C., Dennis, M. F., & Beckham, J. C. (2011). Autobiographical memory for stressful events: The role of autobiographical memory in posttraumatic stress disorder. *Consciousness and Cognition, 20*(3), 840–856. https://doi.org/10.1016/j.concog.2011.03.015

Selimbegović, L., Régner, I., Huguet, P., & Chatard, A. (2016). On the power of autobiographical memories: From threat and challenge appraisals to actual behaviour. *Memory, 24*(10), 1382–1389. https://doi.org/10.1080/09658211.2015.1111908

Selimbegović, L., Régner, I., Sanitioso, R. B., & Huguet, P. (2011). Influence of general and specific autobiographical recall on subsequent behavior: The case of cognitive performance. *Journal of Experimental Social Psychology, 47*(1), 72–78. https://doi.org/10.1016/j.jesp.2010.08.011

Seymour, E., & Hewitt, N. M. (1997). *Talking about leaving: Why undergraduates leave the sciences*. Westview Press.

Sharma, A., Lin, I. W., Miner, A. S., Atkins, D. C., & Althoff, T. (2023). Human–AI collaboration enables more empathic conversations in text-based peer-to-peer mental health support. *Nature Machine Intelligence, 5*(1), 46–57. https://doi.org/10.1038/s42256-022-00593-2

Speer, M. E., Ibrahim, S., Schiller, D., & Delgado, M. R. (2021). Finding positive meaning in memories of negative events adaptively updates memory. *Nature Communications, 12*, Article 6601. https://doi.org/10.1038/s41467-021-26906-4

Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology, 29*(1), 24–54. https://doi.org/10.1177/0261927X09351676

Tversky, B., & Marsh, E. J. (2000). Biased retellings of events yield biased memories. *Cognitive Psychology, 40*(1), 1–38. https://doi.org/10.1006/cogp.1999.0720

889 von Garrel, J., & Mayer, J. (2023). Artificial intelligence in studies: Use of ChatGPT and AI-
890 based tools among students in Germany. *Humanities and Social Sciences*
891 *Communications, 10*, Article 799. https://doi.org/10.1057/s41599-023-02304-7
892 Wirtz, D., Kruger, J., Napa Scollon, C., & Diener, E. (2003). What to do on spring break? The
893 role of predicted, on-line, and remembered experience in future choice. *Psychological*
894 *Science, 14*(5), 520–524. https://doi.org/10.1111/1467-9280.03455
895 Zhan, H., Zheng, A., Lee, Y. K., Suh, J., Li, J. J., & Ong, D. C. (2024). Large language models
896 are capable of offering cognitive reappraisal, if guided. arXiv.
897 https://doi.org/10.48550/arXiv.2404.01288
898

Tables and Figures

Table 1. Participant Demographics by Condition.

Characteristic	Overall (N = 42)	Self-Authorship (n = 22)	Control (n = 20)
Age			
M (*SD*)	19.6 (1.2)	19.7 (1.3)	19.6 (1.1)
Range	18–22	18–22	18–21
Gender			
Man	10 (23.8%)	5 (22.7%)	5 (25.0%)
Woman	32 (76.2%)	17 (77.3%)	15 (75.0%)
Race/Ethnicity			
Asian	17 (40.5%)	9 (40.9%)	8 (40.0%)
Asian, White	1 (2.4%)	0 (0.0%)	1 (5.0%)
Black or African American	3 (7.1%)	2 (9.1%)	1 (5.0%)
Prefer not to answer	2 (4.8%)	1 (4.5%)	1 (5.0%)
White	19 (45.2%)	10 (45.5%)	9 (45.0%)
Hispanic/Latino			
No	34 (81.0%)	17 (77.3%)	17 (85.0%)
Prefer not to answer	1 (2.4%)	1 (4.5%)	0 (0.0%)
Yes	7 (16.7%)	4 (18.2%)	3 (15.0%)
Grade/Year			
First Year	18 (42.9%)	11 (50.0%)	7 (35.0%)
Second Year	5 (11.9%)	2 (9.1%)	3 (15.0%)
Third Year	12 (28.6%)	5 (22.7%)	7 (35.0%)
Fourth Year	7 (16.7%)	4 (18.2%)	3 (15.0%)

Note. Includes only participants who completed both sessions. Values are n (%) unless otherwise indicated.

Table 2. Narrative Word Counts by Condition and Session.

Condition	Session	Participants	Narratives	*M*	*SD*	Min	Max
Self-Authorship	Session 1	22	110	75.7	27.0	25	182
Self-Authorship	Session 2	22	110	67.0	27.7	21	165
Control	Session 1	20	100	80.9	33.8	23	207
Control	Session 2	20	100	68.8	26.2	11	135

Note. Word counts are LIWC-22 word counts (WC) for each memory narrative. Each participant contributed five narratives per session. Ten narratives (2.4%) contained fewer than 25 words; see Supplementary Table 2 for the analysis excluding them.

Table 3. Descriptive Statistics for LIWC-22 Language Variables by Condition and Session.

Language variable	Self-Authorship Session 1	Self-Authorship Session 2	Control Session 1	Control Session 2
Tentative language (tentat)	3.54 (2.91)	2.87 (2.40)	3.08 (2.48)	3.43 (2.99)
Analytic language (Analytic)	36.95 (25.75)	36.41 (26.82)	40.01 (24.46)	32.78 (28.17)
Anxiety language (emo_anx)	0.42 (0.78)	0.64 (1.13)	0.30 (0.77)	0.75 (1.28)
Positive emotion (emo_pos)	0.45 (0.87)	0.35 (0.69)	0.54 (0.98)	0.53 (1.14)
Negative emotion (emo_neg)	0.83 (1.19)	1.34 (1.63)	0.66 (0.97)	1.24 (1.63)
Positive tone (tone_pos)	1.79 (1.76)	1.72 (2.02)	2.00 (1.84)	1.98 (1.77)
Negative tone (tone_neg)	1.31 (1.54)	2.02 (1.84)	1.23 (1.22)	1.74 (1.93)

Note. Values are means with standard deviations in parentheses, computed across memory narratives (Self-Authorship: n = 22 participants, 110 narratives per session; Control: n = 20 participants, 100 narratives per session). Tentative, anxiety-related, emotion, and tone scores are LIWC-22 percentages of words; Analytic is the LIWC-22 summary score (0–100).

Table 4. Post-Intervention Ratings by Condition.

Measure	Self-Authorship *M* (*SD*)	Control *M* (*SD*)	Mean diff.	*t*(*df*)	*p*	*d*
Reframe Helpfulness	4.13 (1.07)	2.95 (1.12)	1.18	3.48 (39.2)	.001	1.08
Elaboration Helpfulness	4.48 (1.00)	3.13 (1.03)	1.35	4.31 (39.4)	< .001	1.33
Examined & Changed Assumption	4.03 (1.03)	2.72 (1.02)	1.31	4.13 (39.7)	< .001	1.27
Concrete Next Steps	4.49 (1.03)	2.81 (0.97)	1.68	5.42 (39.9)	< .001	1.67

Note. Self-Authorship n = 22; Control n = 20. Values are means with standard deviations in parentheses. Mean diff. = Self-Authorship − Control. p values are from Welch t tests; d = Cohen's d.

925 **Table 5.** Linear Mixed-Effects Model of Tentative Language by Condition and Session.

926

Effect	*b* (*SE*)	*t* (*df*)	*p*	Type III *F* (*df*)	*p*
Intercept	3.08 (0.30)	10.14 (77.0)	< .001	—	—
Condition (Self-Authorship vs. Control)	0.46 (0.42)	1.10 (77.0)	.276	0.02 (1, 40)	.894
Session (Session 2 vs. Session 1)	0.35 (0.33)	1.07 (208)	.286	0.48 (1, 208)	.487
Condition × Session	−1.02 (0.45)	−2.25 (208)	.026	5.04 (1, 208)	.026

927 *Note.* LIWC-22 tentative language (percentage of words) modeled as tentat ~ Condition × Session + (1 |
928 participant) + (1 | participant:memory), fit by restricted maximum likelihood (N = 420 narratives; 210
929 memories; 42 participants). *b* = unstandardized coefficient with Control and Session 1 as reference levels;
930 *t* and Type III *F* tests use Satterthwaite degrees of freedom. Random-effect variances (*SD*): memory
931 within participant 1.52 (1.23); participant 0.47 (0.68); residual 5.35 (2.31); REML criterion 2008.3.
932 Standardized by the model-implied total standard deviation (2.71), the Condition × Session coefficient
933 equals −0.37.
934

935 **Table 6.** Reframe Lexical Intrusion Analysis.

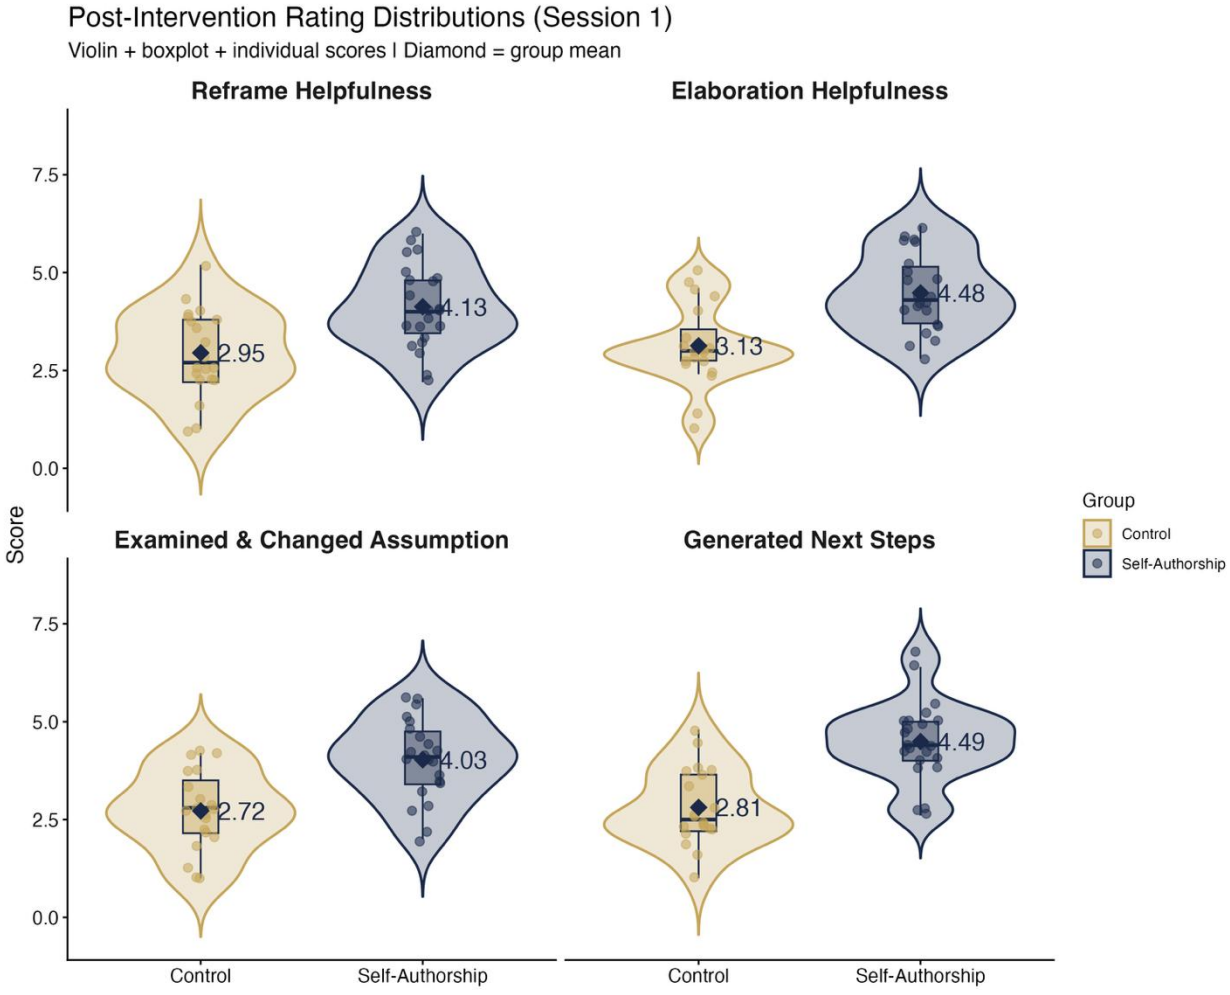
936

Analysis	Group / Comparison	Above-Chance Gap	*t*(df)	*p*	*d*	BF_{10}
Other-person chance baseline	All participants	+1.54 percentage points	3.23 (41)	.002	0.50	13.5
Own-other-memory chance baseline	All participants	+0.96 percentage points	2.16 (41)	.036	0.33	1.4
Other-person chance baseline	Control	+2.08 percentage points	2.45 (19)	.024	0.55	2.5
Other-person chance baseline	Self-Authorship	+1.06 percentage points	2.16 (21)	.042	0.46	1.5

Analysis	Group / Comparison	Above-Chance Gap	*t*(df)	*p*	*d*	BF_{10}
Between-condition comparison, other-person baseline	Control vs. Self-Authorship	+1.02 percentage points	1.04 (30.74)	.305	—	—
Between-condition comparison, own-other-memory baseline	Control vs. Self-Authorship	+0.06 percentage points	0.06 (33.83)	.951	—	—

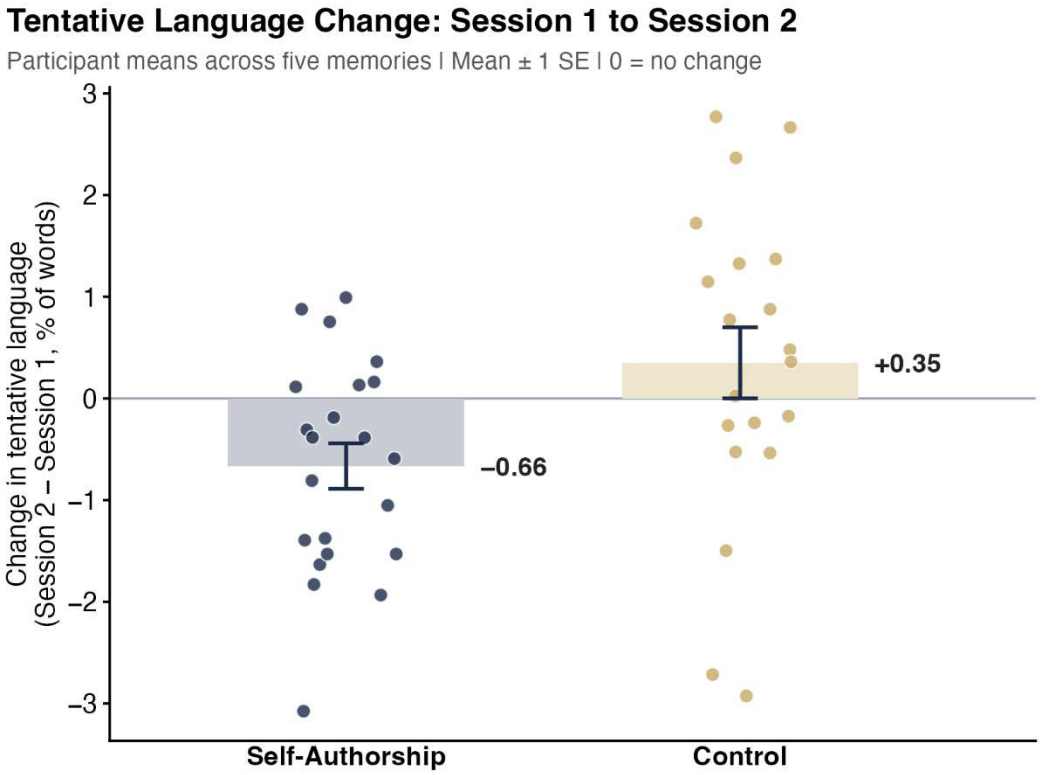
Note. Positive values indicate that vocabulary from the participant's assigned reframe appeared in the corresponding Session 2 memory description more often than expected from the specified chance baseline. For the between-condition comparisons, the mean above-chance rates were .0208 in Control and .0106 in Self-Authorship against the other-person baseline, 95% CI for the mean difference [−.0097, .0301], and .0099 in Control and .0093 in Self-Authorship against the own-other-memory baseline, 95% CI [−.0180, .0191].

Figure 1. Post-Intervention Rating Distributions (Session 1).



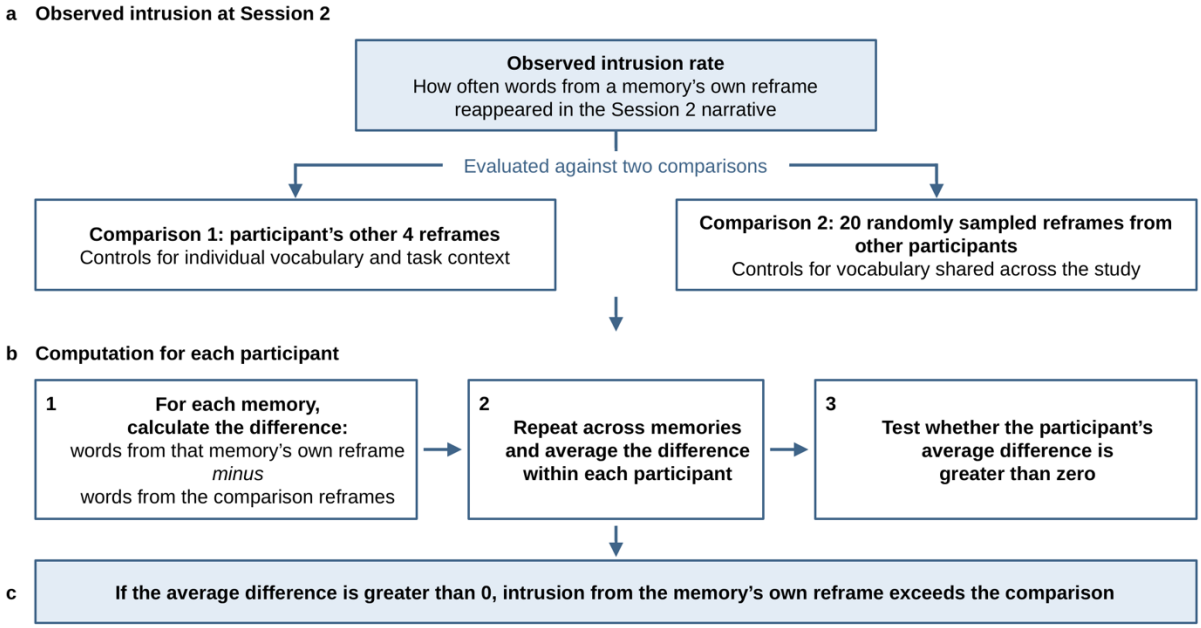
Note. Distributions of four post-intervention ratings collected immediately after the Session 1 reframing exercise, by condition (Self-Authorship, n = 22; Control, n = 20): rated helpfulness of the reframe, rated helpfulness of the elaboration, the degree to which participants examined and changed their assumptions about the uncertainty event, and the degree to which participants generated concrete and doable next steps. Each panel shows the distribution (violin), the interquartile range (box), and individual scores; diamonds denote group means.

Figure 2. Tentative Language Change (Session 1 to Session 2).



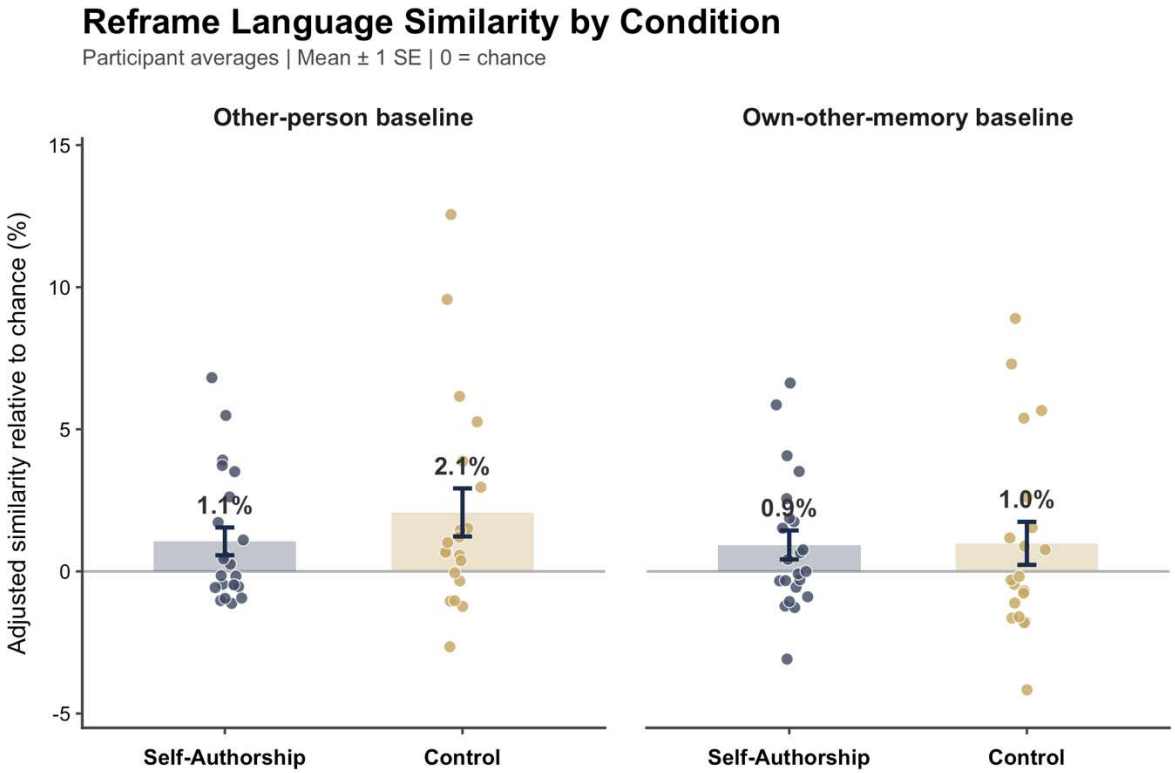
Note. Mean within-memory change (Session 2 minus Session 1) in the LIWC tentative word category (tentat) for the Self-Authorship and Control conditions. Points are participant means across the five memories; bars show condition means; error bars denote ±1 SE across participants.

Figure 3. A Diagram of the LLM-Generated Lexical Intrusion Analysis.



Note. Schematic of the lexical intrusion analysis. **a**, For each memory, the observed intrusion rate is the proportion of novel stems from that memory's own reframe that reappeared in the Session 2 narrative. It was evaluated against two comparison points: the same participant's other four reframes, which control for individual vocabulary and task context, and 20 randomly sampled reframes from other participants, which control for vocabulary shared across the study. **b**, For each participant, the difference between the own-reframe rate and each comparison rate was averaged across the five memories and tested against zero. **c**, A positive average indicates that intrusion from a memory's own reframe exceeded the comparison.

Figure 4. Reframe Lexical Intrusion by Condition and Chance Baseline.



Note. Participant-level mean rate, in percentage points above chance, at which novel vocabulary from each memory's own reframe reappeared in the corresponding Session 2 narrative, relative to the other-person baseline (left) and the own-other-memory baseline (right), by condition (Self-Authorship, n = 22; Control, n = 20). Points are individual participants; bars show condition means, printed above each bar; error bars denote ±1 SE; the horizontal line at 0 marks chance.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- RamseyerThomas2026SupplementaryInformation.docx