

Title: Large language models show human-like content biases in transmission chain experiments

Authors:

A. Acerbi^{1*}, J.M. Stubbersfield²

Affiliations:

¹Department of Sociology and Social Research, University of Trento, Trento, Italy.

²Department of Psychology, University of Winchester, Winchester, UK.

*Corresponding author: alberto.acerbi@unitn.it

One-Sentence Summary: Chat-GPT, as humans, is biased towards certain content, such as stereotype-consistent, negative, social, or threat-related.

Abstract: As the use of Large Language Models (LLMs) grows, it is important to examine if they exhibit biases in their output. Research in Cultural Evolution, using transmission chain experiments, demonstrates that humans have biases to attend to, remember, and transmit some types of content over others. Here, in five pre-registered experiments with the same methodology, we find that the LLM chatGPT-3 shows biases analogous to humans for content that is gender-stereotype consistent, social, negative, threat-related, and biologically counterintuitive, over other content. The presence of these biases in LLM output suggests that such content is widespread in its training data, and could have consequential downstream effects, by magnifying pre-existing human tendencies for cognitively appealing, and not necessarily informative, or valuable, content.

Main text:

Research on algorithmic bias has highlighted how the application of machine learning techniques to corpora generated by humans is likely to reproduce the biases present in the corpora (1). As Large Language Models (LLMs) like ChatGPT have been recently opened to the broad public, with potential applications in journalism (2), copywriting (3), academia (4) and other writing tasks (5), and as they are trained on previous textual material produced mostly by humans, it is important to understand whether and how they would reflect those biases. Tools like ChatGPT provide a way of interacting that has become widespread in recent years (text-based chat) and that could greatly expand the user base of LLMs. In addition, they produce replies that feel as “natural” stories or narratives where those biases can be not immediately evident but pervasive in their effects (6).

To investigate this, we applied a method generally used with human participants: the method of serial reproduction, or “transmission chain” set-up. This method has a long history in psychology (7), and has been lately revived in the cultural evolution framework (8, 9). In short, the transmission chain method is a laboratory version of the telephone game, where participants pass iteratively to each other a story (or a solution to a task), and the researchers can track how these are modified through the steps of the chain. One can do the same with an LLM, asking it to summarize a story and then present in the next step the summarized version produced by the LLM to itself, and proceed iteratively, as illustrated in Figure 1.

Transmission chain experiments with human participants have shown that humans tend to preferentially preserve and transmit some content with respect to others (10). For example, in stories including both positively valenced and negatively valenced events, negative events tend to be transmitted and preserved more than positive ones, showing a possible negative bias in human cultural transmission (11).

We tested, in a fully pre-registered analysis, OpenAI's GPT-3 with the same material from five previous experiments with human participants to assess whether the LLM would show the same biases. In each experiment, the initial story was passed to GTP-3 with a short prompt (see Figure 1) and the output produced was then presented again with the same prompt, iteratively. At each passage, we tracked the proportion of information retained: in particular, the information consistent (and inconsistent) with the bias in each experiment. For example, in the story of Figure 1, some information is gender-stereotype consistent, while other is gender-stereotype inconsistent. In our analysis, we tested whether chatGPT's output would produce the same biases found in human participants. We used linear mixed effects models (implemented in R with (12)) with the proportion of content retained as the outcome, the type of content as the predictor, and the step in the chain and the replication as random effects.

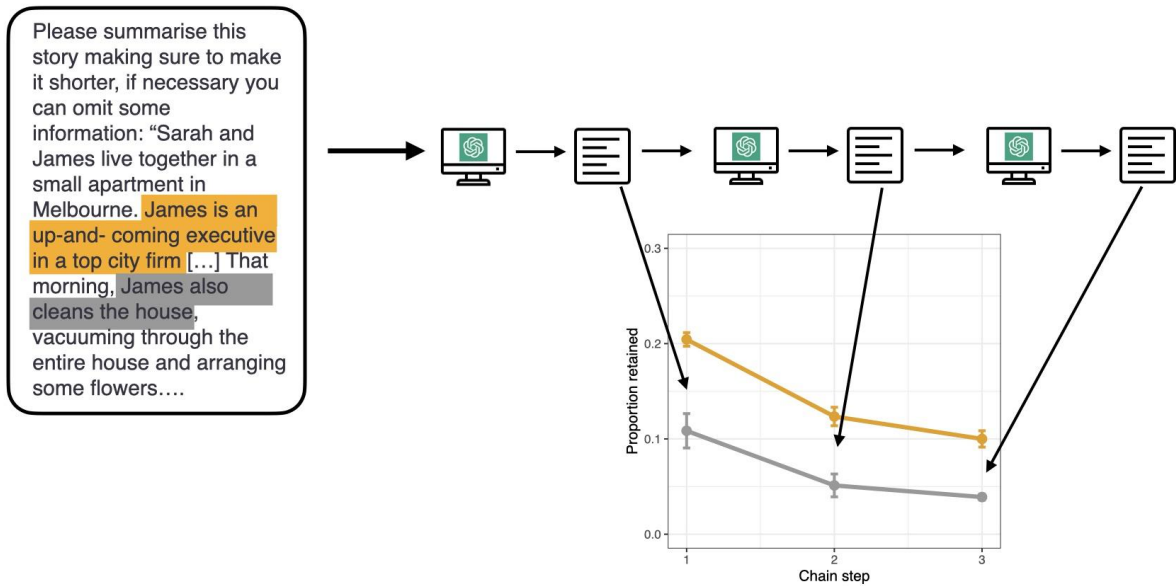


Fig. 1. Basic experimental set-up. A story with gender-stereotype consistent information (orange) and gender-stereotype inconsistent information (gray) is given to chatGPT, after a short prompt that asks to summarize it. The proportion of consistent and inconsistent information reproduced is recorded, and the output is passed again to chatGPT with the same prompt. The operation is iterated three times.

The first experiment (*I3*) compares gender-stereotype consistent and gender-stereotype inconsistent information, such as a wife cooking for a dinner party where the husband invites its boss hoping for a promotion (gender-stereotype consistent) and the same wife going out for drinks with friends before the dinner (gender-stereotype inconsistent - full texts of the original story, as well as of all material used in the other experiments are in Supplementary Material). Successive iterations of the story with chatGPT shows that, as with human participants, stereotype-consistent information was reproduced more than stereotype-inconsistent information through the chain ($\beta = 0.058$, $p < 0.01$ - see figure 2A).

The second experiment (*I1*) concerns negative versus positive information, with a story of a girl flying to Australia and, for example, sitting close to a man “with a nasty cold” (negative information) or being moved to business class (positive information). As above, chatGPT behaved consistently with human participants, reproducing more negative than positive information ($\beta = 0.117$, $p < 0.001$ - see figure 2B). This story also included ambiguous details (e.g., the protagonist sees a man “take an old woman’s bag”) that could be resolved positively (a kind man helping an older woman) or negatively (stealing the bag) light. Even in this case, consistently with previous results with humans, these initially ambiguous details were mostly resolved negatively ($\beta = 0.183$, $p < 0.001$ - see figure 2C, dark gray is for details remaining ambiguous, and light gray is for positive resolutions).

The third experiment (*I4*) examines the difference between social information (for example, a student having an affair with a professor) and non-social information (the student waking up late and missing a lecture, or the weather conditions). Human participants were found to preferentially preserve social information, and chatGPT produced results consistent with the experiments with humans ($\beta = 0.321$, $p < 0.001$ - see figure 2D).

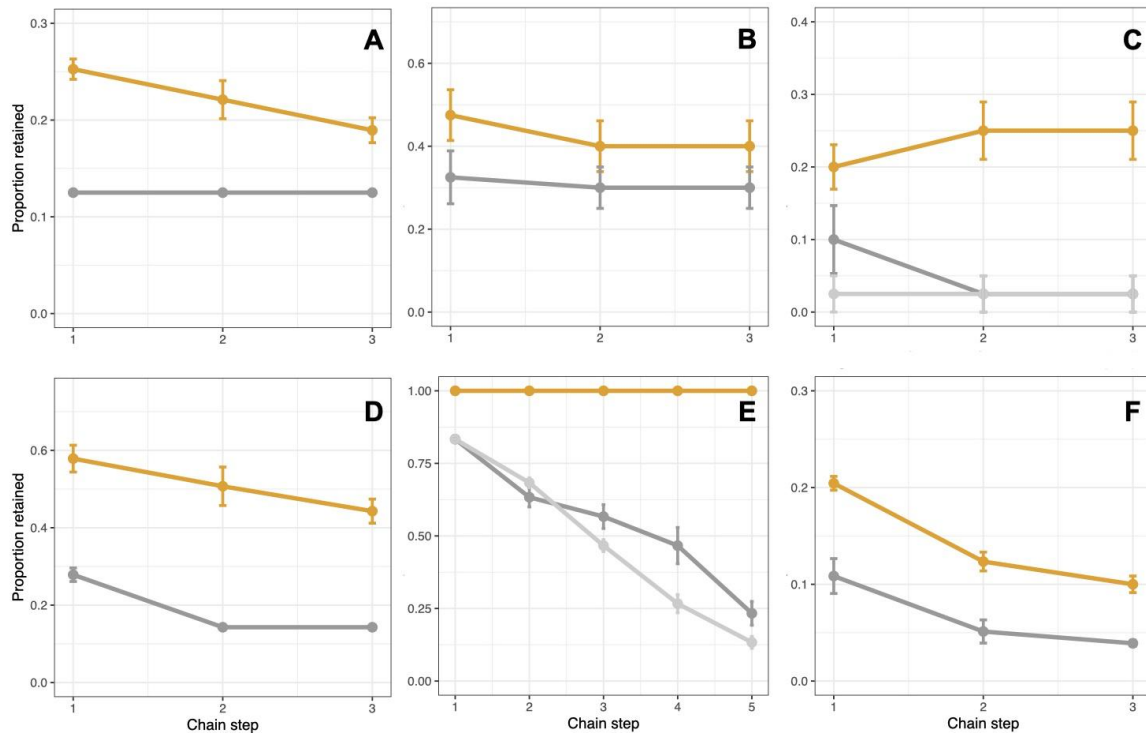


Fig. 2. Proportion of information retained by ChatGPT in the experiments. (A) Gender-stereotype consistent (orange) versus gender-stereotype inconsistent (gray) information in Experiment 1. (B) Negative (orange) versus positive (gray) information in Experiment 2. (C) Ambiguity resolutions in experiment 2: negative (orange) versus positive (light gray) and ambiguous (dark gray). (D) Social (orange) versus non-social (gray) information in Experiment 3. (E) Threat-related (orange) versus negative (dark gray) and ambiguous (light gray) in Experiment 4. (F) Counterintuitive biological, social and negative information (orange) versus other biases (gray) in Experiment 5. All data are average of five replications, bars show standard deviations.

The fourth experiment (15) considers a specific type of negative information: information related to possible threats. The set-up is slightly different here: instead of a story, it presents a “consumer report” followed by statements to help a “friend [that] mentioned that he would like to purchase this product”. For example, one concerns a “new running shoe brand called Lancer™”, and the statements include, among others, “Lancer™’s strap design can cause sprained ankles when used for activities other than running” (threat-related information), “Lancer™ special fabric may smell if not cleaned properly” (negative information), or “Lancer™ customization process analyzes the way you run” (neutral information). In agreement with human participants, chatGPT retained threat-related statements through the iterations, dropping negative and neutral ones ($\beta = 0.523$, $p < 0.001$ - see figure 2 “Threat-related”). When negative content is tested against neutral, excluding threat-related content from the analysis, negativity predicts, as hypothesized and as for humans’ results, retention through the chain ($\beta = 0.070$, $p < 0.005$ - see figure 2E, dark gray is negative, and light gray is neutral).

Finally, the fifth experiment (16) included material relevant to multiple possible content biases in two different narratives created for the original study and inspired by creation myths. Human participants were found to preferentially transmit negative information (“Muki cried and cried, until the spark in the sky darted away”), social information (“The elder ones had not approved of their marriage”), or counterintuitive information related to biological

processes (“the hairs of Pata's chin became spiders and crawled up from their bed of clay”) versus other kinds of content (including content relevant to other biases), and the results for chatGPT were consistent with the human results ($\beta = 0.076$, $p < 0.001$ - see figure 2F). Extended results, with outcomes for single biases and different material within each experiment, are included in Supplementary Material.

Across five experiments, using the chatGPT LLM to replicate previous transmission chain studies with human participants, we found that the information retained in outputs produced by the LLM was analogous to the information retained and transmitted by human participants. Consistent with pre-registered hypotheses, text produced by chatGPT reflected human content-based social transmission biases for stereotype consistency (Experiment 1), negative information (Experiment 2), social information (Experiment 3), and threat-related information (Experiment 4). In addition, it reflects human biases for negative, social, and biologically counterintuitive content over other biases (Experiment 5), and a bias for resolving ambiguous statements as negative (Experiment 2).

As such, we can expect that when LLMs generate new texts or are used to summarize pre-existing text, their outputs will reflect these biases. It is important to note that the concept of bias we are using in this research is different from the common concept of “algorithmic bias”, which has a distinct negative connotation (17). In cultural evolution, transmission biases indicate that individuals are predisposed, on average, to adopt some cultural variants over others, impacting on their overall spread within populations (18). In this sense, biases are neither inherently good nor bad, and they are to be expected each time individuals choose among different cultural variants. This is reflected in the experiments reported here, where some biases would be considered negative (e.g., the preference for stereotype-consistent information in Experiment 1) but others neutral, or possibly functional, as the bias towards threat-related information in Experiment 4. On the other side, for the same reason, those biases could be more difficult to recognize, and they could have consequential downstream effects, by magnifying the pre-existing human tendencies. We might anticipate that, without human intervention, LLMs could enable negative gender stereotypes to persist in potentially harmful ways. Additionally, given concerns over emotional contagion in digital media (19), negativity and threat bias in LLM generated material could contribute to wider negativity and overestimation of threats in humans. If used to summarize scientific articles for the purposes of journalism (2), LLMs may focus on content which appeals to these biases rather than the truly pertinent, although not necessarily more so than a human editor. A key implication then, is that it is important to recognize our own subtle biases, and to understand that LLMs reflect these, and do not act as neutral agents.

Research in cultural evolution suggests that content biases are present in humans as a result of biases in our cognition which lead us to preferentially attend to, recall and/or transmit some types of information over others (10). In most cases, it is proposed that these biases are a result of evolved cognition, being on average adaptive, e.g. a bias for social information resulting from the fitness benefits of attending to and remembering such information within our social groups (14), or biases for negative and threat-related information because not attending to such is more costly than not attending to positive information or benefits (15). Biases in the outputs of LLMs cannot be the result of such evolutionary processes; however, the human-produced training material is itself a product of a cultural evolutionary process where human content biases have led to the preferential retention and dissemination of information which align with those biases. Such content then likely makes up a disproportionate amount of that training material relative to other types of content and are

therefore reflected in the ‘biases’ of the LLM. In the case of chatGPT-3 the training material was around 45 TB of text taken from multiple web-based sources including Wikipedia, books, and raw web page data (20). As such the results of our study are evidence that biases which have previously been tested experimentally, or within corpus analysis (21–24) are found across an exceptionally wide range of cultural artifacts. It also suggests that the outputs of LLMs could be useful sources for studying broader human culture (although this will depend significantly on how outputs are constrained by other processes).

While overall the results here are consistent with the outcomes of the original experiments, there are some interesting minor differences which likely reflect the differing mechanisms which produce the biased outputs. One such difference was the higher retention of ‘gossip’ over standard social information. In the original study (14) no difference was found between gossip (defined as information about intense third-party social relationships) and non-gossip social information, and both were equally well retained over non-social information about an individual or the physical environment. In our study (Experiment 3) gossip had a significant advantage over standard social information (see Figure S2 - Supplementary Material), largely explaining the overall prominence of social information. This could be a result of the LLM having a stronger bias towards emotive material than social information, possibly negative emotion (as the gossip story in Experiment 3 was mostly negative, see full text in Supplementary Material). Relative importance of different biases is something which could be tested for directly in future studies.

A second difference was in the retention of stereotype consistent information. In the original study (13), while stereotype consistent information was preferentially retained across the overall chain, earlier chain steps showed an advantage for stereotype *inconsistent* information. Later research suggested that a bias for stereotype consistent information is a product of communicative intent, rather than memory (25). No such pattern was found in our study (Experiment 1), rather, stereotype consistent information was preferentially retained from at the first chain step. This could be a consequence of the different mechanisms which produce biases in humans and LLMs. While content biases in humans may be present in attention, memory, or transmission, and may vary across these three phases [10], a bias in an LLM can only be a product of its training material. If the training material was predominantly stereotype consistent, this will be reflected in the outputs. In general, across our experiments, evidence for bias was found from the first chain step, suggesting that iteration is not strictly required for LLM output to reflect biases.

The results of this study could be expanded by examining the potential impact of different prompts. Here the prompt was “*Please summarize this story making sure to make it shorter, if necessary you can omit some information*”. We did not directly test the influence of different prompt wording on output, and it is unlikely that this prompt would produce the hypothesized biases alone, but future research could examine how, and to what extent, prompt wording influences the reflection of biases in LLM output. A limitation is that due to the rapid development and diversification of LLMs, our results may not generalize beyond chatGPT-3. ChatGPT-4’s training was similar to chatGPT-3, so may reflect the same biases, but includes reinforcement learning using human feedback in an attempt to prevent the LLM from producing output which violates OpenAI’s policy on harmful behavior and to mitigate harmful biases. However, small-scale experiments with early versions of chatGPT-4 suggest it still generates biased outcomes based on gender stereotypes (26). To what extent it would reflect this bias in summarizing text, as tested here, or the other biases examined here is unknown. Given the similarity in training, however, and that outside of gender stereotypes

the biases tested here are unlikely to be considered harmful, it is plausible to expect that these biases would still be present in later generations of LLMs.

References and notes:

1. A. Caliskan, J. J. Bryson, A. Narayanan, *Science* **356**, 183–186 (2017).
2. S. Petridis *et al.*, in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, pp. 1–16.
3. G. Chen, P. Xie, J. Dong, *T. J. Advert.* **48**, 347–355 (2019).
4. O. “Oz” Buruk, *arXiv.2304.11079* (2023).
5. R. Dale, *Nat. Lang. Eng.* **27**, 113–118 (2021).
6. L. Lucy, D. Bamman, in *Proceedings of the Third Workshop on Narrative Understanding (2021)*, pp. 48–55.
7. F. C. Bartlett, *Remembering* (Cambridge University Press, 1932).
8. A. Mesoudi, A. Whiten, *Philos. Trans. R. Soc. London Ser. B* **363**, 3489–3501 (2008).
9. B. Thompson, B. van Opheusden, T. Summers, T. L. Griffiths, *Science* **376**, 95–98 (2022).
10. J. M. Stubbersfield, *Culture and Evolution* **19**, 41–60 (2022).
11. K. Bebbington, C. MacLeod, T. M. Ellison, N. Fay, *Evol. Hum. Behav.* **38**, 92–101 (2017).
12. D. Bates, M. Mächler, B. Bolker, S. Walker, *J. Stat. Softw.* **67**, 1–48 (2015).
13. Y. Kashima, *Pers. Soc. Psychol. Bull.* **26** (2000), pp. 594–604.
14. A. Mesoudi, A. Whiten, R. Dunbar, *Br. J. Psychol.* **97** (2006), pp. 405–423.
15. T. Blaine, P. Boyer, *Evol. Hum. Behav.* **39**, 67–75 (2018).
16. R. E. W. Berl *et al.*, *Evolut. Hum. Sci.* **3**, e42 (2021).
17. D. Danks, A. J. London, in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (2017)*, pp. 4691–4697.
18. R. Boyd, P. J. Richerson, *Culture and the evolutionary process* (University of Chicago Press, 1985).
19. A. Goldenberg, J. J. Gross, *Trends Cogn.* **24**, 316–328 (2020).
20. T. Brown, *et al.*, in *Advances in Neural Information Processing Systems (2020)*, pp. 1877–1901.

21. J. M. Stubbersfield, E. G. Flynn, J. J. Tehrani, *Evolutionary Studies in Imaginative Culture* **1**, 121–136 (2017).
22. A. Acerbi, *Palgrave Commun.* **5**, 15 (2019).
23. O. Morin, A. Acerbi, *Cogn Emot.* **31**, 1663–1675 (2017).
24. C. O. Brand, A. Acerbi, A. Mesoudi, *Evolut. Hum. Sci.* **1**, e11 (2019).
25. A. Lyons, Y. Kashima, *Asian Journal of Social Psychology* **9**, 59–71 (2006).
26. S. Bubeck, et al., arXiv.2303.12712 (2023).

Acknowledgments:

The authors would like to thank Fabiana Lombardi for her work as an independent coder.

Funding: The authors did not receive funding for this article. **Author contributions:**

Conceptualization: AA, JMS; Methodology: AA, JMS; Investigation: AA, JMS; Analysis: AA; Visualization: AA; Writing: AA, JMS. **Competing interests:** Authors declare that they have no competing interest. **Data and materials availability:** The analysis was fully pre-registered at: <https://osf.io/xv2y9>, and full data and scripts for reproducing the results are available in an OSF repository: <https://osf.io/6v2ps/>

Supplementary materials:

Materials and Methods
Supplementary Text
Figs. S1 to S7
Table S1

Supplementary Materials for:

“Large language models show human-like content biases in transmission chain experiments”

A. Acerbi, J.M. Stubbersfield

Corresponding author: alberto.acerbi@unitn.it

The pdf file includes:

Materials and Methods

Supplementary Text

Figs. S1 to S7

Table S1

Material and methods

We used the 9 January 2023 version of OpenAI's GPT-3 language model (publicly available at: <https://chat.openai.com/chat>), with default parameters values, to run a series of experiments using transmission chains methodology. We selected five studies (see below) that highlighted in human participants different content biases, and that made use of single stories with different biases (as opposed to experiments that used two stories or more stories with different biases). In one case (study 3), we modified the original material, which consisted in four stories testing four different biases, creating a single story. We used the same material (stories) of the original experiments (modifications are reported below). The material was presented in chatGPT with the prompt:

*Please summarize this story making sure to make it shorter, if necessary you can omit some information: **story***

For each study, we run five different chains/replications, and each chain/replication consisted of three steps. In each chain, the original story was presented with the prompt above; the output produced by chatGPT was then presented again with the same prompt (step 2), and the process was iterated a last time (step 3). (The setup is slightly different for study 4, see below). The number of chains/replications and of steps was chosen after pre-tests showing a limited variability of chatGPT's outputs given the same prompt, and that the main modification of the material was happening in the first step of the chain.

The general hypothesis we tested is that chatGPT's output would produce the same biases found in human subjects. While the original studies use different statistical analyses, we decided to have the same general analytic strategy for all studies. We used linear mixed effects models with the *proportion* of content retained as the outcome, the *type* of content as the predictor, and the step in the chain and the replication as random effects. Using the R package lme4 (12) the general formula can be written as:

```
lmer(proportion ~ content + (1|chain_step) + (1|chain_id))
```

The coding consisted in determining the presence or absence of basic information from the original story. ChatGPT's output was coded by AA (study 1, 2, and 4) and by JMS (study 3 and 5). A third independent coder, unaware of the experimental procedure and of the predictions, double coded studies 1, 2, 3, and 5 (study 4 did not need double coding as the procedure is slightly different - see below). Inter-rater reliability was generally high and it is reported in Table S1.

Detailed hypotheses and implementations are presented below for each study, as well as how the outputs were coded. The preregistration, plus all the original data, coding, and R scripts to perform the analysis and visualizations are available in an OSF repository:

<https://osf.io/6v2ps/>

Study 1. Stereotype consistency

Source: Kashima (2000)

Story:

“Sarah and James live together in a small apartment in Melbourne. James is an up-and-coming executive in a top city firm where Sarah also works as a highly competent personal assistant. This particular Saturday is extremely important because they are entertaining James’s employer; James’s promotion is at stake and he wishes to make a good impression. James has been awaiting this for a long time, and the possibility of a raise is crucial for their plans to start a family. Sarah has promised to cook a beautiful three-course meal and to make it a successful evening altogether.

They spend the morning shopping at Sarah’s favorite, exclusive delicatessens. Sarah even gets her hair done for the occasion while James takes the opportunity to choose the appropriate wine. That morning, James also cleans the house, vacuuming through the entire house and arranging some flowers. Sarah has set up the kitchen and makes some initial preparations for dinner, which she will finish later that afternoon. She then makes lunch and they take a break together.

The phone rings. It’s a couple of Sarah’s mates from squash; she hangs up and rushes into the bedroom. She yells to James that the girls are down at the Royal and that she’s just going to pop down and have a few drinks with them. She gives him a quick peck on the cheek and says she won’t be long.

Feeling good about the preparations for that evening, and confident it was going to run smoothly, James settles down for an afternoon of Wide World of Sports. This week they’re crossing live to the Brazilian Grand Prix. It was a good afternoon for a bit of TV.

James wakes to the phone ringing and it dawns on him that he’s been asleep all afternoon. He crawls off the couch but misses the phone. Wandering into the kitchen, he sees the preparation for dinner; Sarah was not yet home. Could it have been Sarah trying to ring? He tries her mobile but it isn’t switched on and James begins to worry. Sarah didn’t usually stay at the pub this long, and she had so much to do. He realizes he must do something and opens the cookbook placed on the bench. He isn’t exactly sure what Sarah had planned to cook.

It is 6 o’clock and James is whipping the cream for dessert, he hears voices and the door slam shut. Laughing, Sarah runs in and gives James a big hug. James angrily pushes her away, yelling “Where the hell have you been?” Sarah, still laughing, tells James that she had so much fun drinking with the girls that she invited Brooke and Nat back for the dinner party.

As Sarah and her friends went to freshen up, the doorbell rang. James couldn’t believe this was happening.”

Coding:

We record the number of propositions that are retained from the original material in successive iterations (based on coding provided in the original paper). Propositions are defined as a predicate plus a series of ordered arguments. As precise wording is unimportant

for the propositional representation, propositions do not need to use the exact words, to be considered as retained (see Mesoudi, et al., 2006).

The text used for coding is below. Each proposition is separated by a slash. PMC = plot-relevant male stereotype-consistent; PFC = plot-relevant female stereotype-consistent; PMI = plot-relevant male stereotype-inconsistent; PFI = plot-relevant female stereotype-inconsistent; BMC = background male stereotype-consistent; BFC = back-ground female stereotype-consistent; BMI = background male stereotype-inconsistent; BFI = background female stereotype-inconsistent.

Notice some propositions were not coded with the above classifications as they were classified as 'gender-neutral' by independent judges in the original study .

Text for coding:

Sarah and James live together in a small apartment in Melbourne./ James is an up-and-coming executive in a top city firm (BMC)/ where Sarah also works as a highly competent personal assistant. (BFC)/ This particular Saturday is extremely important because they are entertaining James's employer; (PMC)/ James's promotion is at stake and he wishes to make a good impression. (PMC)/ James has been awaiting this for a long time, (PMC)/ and the possibility of a raise is crucial for their plans to start a family./ Sarah has promised to cook a beautiful three-course meal/ and to make it a successful evening altogether. (PFC)/

They spend the morning shopping/ at Sarah's favorite, exclusive delicatessens. (BFC)/ Sarah even gets her hair done for the occasion (BFC)/ while James takes the opportunity to choose the appropriate wine. (BMC)/ That morning, James also cleans the house, (BMI)/ vacuuming through the entire house (BMI)/ and arranging some flowers. (BMI)/ Sarah has set up the kitchen and makes some initial preparations for dinner, (PFC)/ which she will finish later that afternoon. (PFC)/She then makes lunch (BFC)/ and they take a break together./

The phone rings./ It's a couple of Sarah's mates from squash; (PFI)/ she hangs up and rushes into the bedroom./ She yells to James that the girls are down at the Royal (BFI)/and that she's just going to pop down and have a few drinks with them. (PFI)/ She gives him a quick peck on the cheek (BFC)/ and says she won't be long. (PFI)/

Feeling good about the preparations for that evening,/ and confident it was going to run smoothly, (BMC)/ James settles down for an afternoon of Wide World of Sports. (BMC)/ This week they're crossing live to the Brazilian Grand Prix./ It was a good afternoon for a bit of TV./

James wakes to the phone ringing/ and it dawns on him that he's been asleep all afternoon. (PMC)/ He crawls off the couch/ but misses the phone./ Wandering into the kitchen, he sees the preparation for dinner;/ Sarah was not yet home./Could it have been Sarah trying to ring?/ He tries her mobile but it isn't switched on/ and James begins to worry. (PMI)/Sarah didn't usually stay at the pub this long, (PFI)/ and she had so much to do. (PFC)/ He realizes he must do something (PMI)/ and opens the cookbook placed on the bench. (PMI)/He isn't exactly sure what Sarah had planned to cook./

It is 6 o'clock/ and James is whipping the cream for dessert, (PMI)/ he hears voices and the door slam shut./ Laughing, Sarah runs in/ and gives James a big hug. (BFI)/ James angrily pushes her away, yelling "Where the hell have you been?"(BMC)/ Sarah, still laughing,

(BFI)/tells James that she had so much fun drinking with the girls (PFI)/ that she invited Brooke and Nat back for the dinner party. (PFI)/

As Sarah and her friends went to freshen up, (PFC)/ the doorbell rang./ James couldn't believe this was happening.

Analysis

We used:

```
lmer(proportion ~ content + (1|chain_step) + (1|chain_id))
```

Where “proportion” is the proportion of sentences of each content retained at each step and in each chain, and content is stereotype-consistent content ("PMC", "PFC", "BMC", "BFC") versus stereotype-non consistent content (“PMI”, “PFI”, “BMI”, “BFI”).

Study 2: Negativity

Source: Bebbington et al. (2017)

Story:

Sarah is 18 years old. She decided to take some time off before starting university. Her parents were disappointed that Sarah didn't want to start uni right away. Sarah's exam results were due out in a week. She was sure she knew how she had done. Her parents had made her promise to phone with the results. They had always been supportive of her. Sarah had decided to spend some time in the UK. On the day of her trip Sarah arrived at the airport. She was so excited. Sarah boarded the plane and found her seat. As she sat down Sarah found her seatbelt was broken. She called for a flight attendant to ask to be moved to a new seat. The flight attendant told Sarah she would check to see what seats were available. When she returned she told Sarah that she would be moved to business class. Sarah was thrilled. As the plane sped down the runway Sarah noticed how jittery she felt. When the seatbelt light went out Sarah unfastened her seatbelt and pushed her seat back. The seat was so comfortable. There was a young woman sitting in the seat next to Sarah. Sarah decided to start a conversation with her. The woman told her she was travelling for work. As Sarah told her about her travel plans she could tell how interested the woman was in the conversation. An hour into the flight lunch was served. There was a choice between fish and beef. Sarah had once gotten food poisoning from bad fish, so she chose the beef. As Sarah smelled the food she realised she was hungry. Sarah noticed that the meat had an unusual flavour. Sarah had a short stop-over in Dubai. Walking down the concourse, Sarah saw a young man take an old woman's bag. As she walked on further, Sarah saw that one of the stores was having a big sale. She couldn't resist having a quick look. She picked out a top to try on. As she came out of the changerooms, Sarah overheard someone talking about how she looked. As Sarah left the store, she found some money on the ground. As she was putting the note in her pocket, Sarah heard a boarding call for her flight over the PA system. Looking around Sarah realised she had no idea where her gate was. She hurried over to an information desk to ask for directions and she started running toward her gate. She got there just in time! Sarah boarded the plane for the next leg of her trip. The man in the seat next to her seemed to have a nasty cold. Sarah got up from her seat and asked the flight attendant if she could be moved. The flight attendant told her that she would have to stay in her allocated seat. Sarah returned to her seat. She was sure she would catch his cold. During the flight Sarah got up to go to the toilet. There was a long queue. When Sarah reached the front, she was surprised by the condition of the toilet. She moved back to her own seat and reflected on her trip so far. She had been surprised by the quality of the service on the flight.

Coding:

The material from Bebbington and colleagues (2017) was divided into single statements describing an event. We record the number of Negative and Positive event statements that are retained from the original material in successive iterations. As in the original experiment, events are considered to be retained if the 'basic gist' is the same as the original text. In addition, as in the original experiment, we coded whether "ambiguous" sentences were retained as they were or transformed into positive or negative, e.g.: "Walking down the concourse, Sarah saw a young man take an old woman's bag", the man could be helping the old woman (positive) or being a thief stealing her bag (negative).

Text for coding:

Sarah is 18 years old./ She decided to take some time off before starting university./ Her parents were disappointed that Sarah didn't want to start uni right away. (Negative)/ Sarah's exam results were due out in a week./ She was sure she knew how she had done. (Ambiguous)/ Her parents had made her promise to phone with the results./ They had always been supportive of her. (Positive)/ Sarah had decided to spend some time in the UK./ On the day of her trip Sarah arrived at the airport./ She was so excited. (Positive)/ Sarah boarded the plane and found her seat./ As she sat down Sarah found her seatbelt was broken. (Negative)/ She called for a flight attendant to ask to be moved to a new seat. / The flight attendant told Sarah she would check to see what seats were available./ When she returned she told Sarah that she would be moved to business class. (Positive)/ Sarah was thrilled. (Positive)/ As the plane sped down the runway Sarah noticed how jittery she felt. (Ambiguous)/ When the seatbelt light went out Sarah unfastened her seatbelt and pushed her seat back./ The seat was so comfortable. (Positive)/ There was a young woman sitting in the seat next to Sarah./ Sarah decided to start a conversation with her./ The woman told her she was travelling for work./ As Sarah told her about her travel plans/ she could tell how interested the woman was in the conversation. (Ambiguous)/ An hour into the flight lunch was served./ There was a choice between fish and beef. / Sarah had once gotten food poisoning from bad fish, (Negative)/ so she chose the beef./ As Sarah smelled the food she realised she was hungry./ Sarah noticed that the meat had an unusual flavour. (Ambiguous)/ Sarah had a short stop-over in Dubai. / Walking down the concourse, Sarah saw a young man take an old woman's bag. (Ambiguous)/ As she walked on further, Sarah saw that one of the stores was having a big sale. (Positive)/ She couldn't resist having a quick look./ She picked out a top to try on./ As she came out of the changerooms, Sarah overheard someone talking about how she looked. (Ambiguous)/ As Sarah left the store, / she found some money on the ground. (Positive)/ As she was putting the note in her pocket,/ Sarah heard a boarding call for her flight over the PA system./ Looking around Sarah realised she had no idea where her gate was. (Negative)/ She hurried over to an information desk to ask for directions/ and she started running toward her gate./ She got there just in time! (Positive)/ Sarah boarded the plane for the next leg of her trip./ The man in the seat next to her seemed to have a nasty cold. (Negative)/ Sarah got up from her seat/ and asked the flight attendant if she could be moved./ The flight attendant told her that she would have to stay in her allocated seat. (Negative)/ Sarah returned to her seat/ She was sure she would catch his cold. (Negative)/ During the flight Sarah got up to go to the toilet./ There was a long queue. (Negative)/ When Sarah reached the front, she was surprised by the condition of the toilet. (Ambiguous)/ She moved back to her own seat/ and reflected on her trip so far./ She had been surprised by the quality of the service on the flight. (Ambiguous)

Analysis

We used:

```
lmer(proportion ~ content + (1|chain_step) + (1|chain_id))
```

Where “proportion” is the proportion of positive and negative sentences retained at each step and in each chain, and content can be one of “positive” or “negative”.

For the analysis of ambiguity resolution we used:

```
lmer(proportion ~ content + (1|chain_step) + (1|chain_id))
```

Where “proportion” is the proportion of ambiguous sentences retained at each step and in each chain (out of the total number of ambiguous sentences), and content can be one of "positive", “negative,” or “same” (i.e., in case the ambiguous sentence is still ambiguous).

Study 3: Gossip and social information

Source: Mesoudi et al. (2006)

We edited into a single text the four stories of the original experiment.

Story:

Nancy is a college student who lives in Colorado. The weather in Colorado gets hot and dry in the summer. This removes moisture from the soil and dries out the plants that grow there. The dry vegetation catches fire easily, leading to frequent forest fires. These fires release smoke containing carbon monoxide into the atmosphere. This smoke contributes to global warming, increasing temperatures further.

Nancy enjoys swimming. Nancy was going to the swimming pool but got lost, so she asked an old man waiting at a bus stop for directions. The old man could not give her directions. A bus arrived at the bus stop and the old man asked the driver for directions. The driver gave Nancy directions to the swimming pool, so Nancy was able to go swimming.

The next day Nancy's alarm clock broke and she overslept. When she woke up she realised that she was late for an important lecture. She got dressed as quickly as she could, left the house and ran to the lecture theatre. When she got there the lecture theatre was empty. Nancy had missed the lecture.

Nancy is having an affair with her married college professor. Nancy recently became pregnant with the professor's child. The professor promised Nancy that he would leave his wife, but since Nancy told him she was pregnant, the professor refused to see her. So Nancy told the professor's wife about the affair. The professor's wife was so upset that she left the professor.

Coding:

We record the number of propositions that are retained from the original material in successive iterations (based on coding provided in the original paper)

The text used for coding is below. In the original study, the different types of material (social, gossip, individual, physical) were transmitted separately along different chains. Here, we combine the material into a single narrative but retain the different types of material in separate sections. Propositions are illustrated below each section and derived from coding provided by the original authors here:

https://github.com/amesoudi/mesoudi_whiten_dunbar_2006/blob/master/study2_material.pdf

Notice that the first sentence was added to give narrative context but is not included in coding for analysis.

Text for coding

Non-social Physical section:

Nancy is a college student who lives in Colorado [excluded]. The weather in Colorado gets hot and dry in the summer. This removes moisture from the soil and dries out the plants that grow there. The dry vegetation catches fire easily, leading to frequent forest fires. These fires

release smoke containing carbon monoxide into the atmosphere. This smoke contributes to global warming, increasing temperatures further.

Propositions:

1. HOT, WEATHER/ 2. DRY, WEATHER/ 3. IN, SUMMER/ 4. REMOVE, MOISTURE, SOIL/ 5. DRY, PLANTS/ 6. CATCH-FIRE, VEGETATION/ 7. EASILY, CATCH-FIRE/ 8. CAUSE, FOREST-FIRES/ 9. FREQUENT, FOREST-FIRES/ 10. RELEASE, SMOKE/ 11. CONTAIN, SMOKE, CARBON-MONOXIDE/ 12. INTO, SMOKE, ATMOSPHERE/ 13. CONTRIBUTE, GLOBAL-WARMING/ 14. INCREASE, TEMPERATURE

Social Non-gossip section:

Nancy enjoys swimming. Nancy was going to the swimming pool but got lost, so she asked an old man waiting at a bus stop for directions. The old man could not give her directions. A bus arrived at the bus stop and the old man asked the driver for directions. The driver gave Nancy directions to the swimming pool, so Nancy was able to go swimming.

Propositions:

1. ENJOY, NANCY, SWIMMING/ 2. GO, NANCY, SWIMMING-POOL/ 3. LOST, NANCY/ 4. ASK, NANCY, MAN/ 5. FOR, DIRECTIONS/ 6. IS, MAN, OLD/ 7. WAIT, MAN, BUS-STOP/ 8. GIVE, MAN, DIRECTIONS, CANNOT/ 9. ARRIVE, BUS, BUS-STOP/ 10. ASK, MAN, DRIVER/ 11. FOR, DIRECTIONS/ 12. GIVE, DRIVER, NANCY, DIRECTIONS/ 13. TO, DIRECTIONS, SWIMMING-POOL/ 14. GO, NANCY, SWIMMING

Non-social Individual section:

The next day Nancy's alarm clock broke and she overslept. When she woke up she realised that she was late for an important lecture. She got dressed as quickly as she could, left the house and ran to the lecture theatre. When she got there the lecture theatre was empty. Nancy had missed the lecture.

Propositions:

1. BROKE, ALARM-CLOCK/ 2. IN, MORNING/ 3. BELONGS-TO, ALARM-CLOCK, NANCY/ 4. OVERSLEEP, NANCY/ 5. WAKE-UP, NANCY/ 6. LATE, NANCY, LECTURE/ 7. IMPORTANT, LECTURE/ 8. DRESS, NANCY/ 9. QUICKLY, DRESS/ 10. LEAVE, NANCY, HOUSE/ 11. RUN, NANCY, LECTURE-THEATRE/ 12. ARRIVE, NANCY, LECTURE-THEATRE/ 13. EMPTY, LECTURE-THEATRE/ 14. MISS, NANCY, LECTURE

Social Gossip section:

Nancy is having an affair with her married college professor. Nancy recently became pregnant with the professor's child. The professor promised Nancy that he would leave his wife, but since Nancy told him she was pregnant, the professor refused to see her. So Nancy told the professor's wife about the affair. The professor's wife was so upset that she left the professor.

Propositions:

1. HAVE, NANCY, AFFAIR/ 2. WITH, AFFAIR, COLLEGE-PROFESSOR/ 3. IS, PROFESSOR, MARRIED/ 4. IS, NANCY, PREGNANT/ 5. PREGNANT-BY, NANCY, PROFESSOR/ 6. PROMISE, PROFESSOR, NANCY, 7/ 7. LEAVE, PROFESSOR, WIFE/ 8. SEE, PROFESSOR, NANCY, REFUSE/ 9. TELL, NANCY, PROFESSOR, 4/ 10. SINCE, 9, 8/ 11. TELL, NANCY, WIFE/ 12. ABOUT, AFFAIR/ 13. UPSET, WIFE/ 14. LEAVE, WIFE, PROFESSOR

Analysis:

We used:

```
lmer(proportion ~ content + (1|chain_step) + (1|chain_id))
```

Where “proportion” is the proportion of sentences of each content retained at each step and in each chain, and content can be one of “social-gossip”, “social-non-gossip”, “non-social-individual”, “non-social-physical” (see above).

We compare social information (social and social-gossip) versus non-social (non-social individual and non-social physical).

Study 4: Threat-related information

Source: Blaine & Boyer (2018)

The original experiment is different from the others, presenting a “consumer report” (see below) followed by 8 statements to help a “friend [that] mentioned that he would like to purchase this product”. In the first iteration, a participant is asked to choose 7 statements, that are then presented to the next participant, that need to choose 6, and 6 on, until three statements remain. The order of the initial 8 statements is randomized.

In this case the prompt was the actual story, without the “*Please summarize this story..., etc.*” as in the other experiment, as the request (“*Please select *n* statements from the following:*”) is part of the script.

As for the previous studies, we run five different chains/replications for each “consumer report”.

Stories:

1. Running Shoes

Please read the following consumer report about a new running shoe brand called Lancer™.

Lancer™ is a luxury running shoe brand that comes in many shapes and sizes. Recently there has been an up-tick in the market for custom running shoes sold through in person consultation sessions. The company uses a patented customization process to analyze the way a person runs to find where their feet need support. The company then builds a unique shoe tailored to those specifications. With regular use, Lancer™ typically last for up to a year longer than other brands. However, when used for activities other than running, Lancer™’s strap design can cause sprained ankles. Also, Lancer™’s smooth sole can cause runners to slip and fall on certain surfaces. Lancer™’s special fabric may smell if not cleaned properly, and does not come in colors other than white.

Your friend mentioned that he would like to purchase this product. Which items about the product are most important to tell him so that he makes a good decision?

Please select 7 statements from the following:

Lancer™’s strap design can cause sprained ankles when used for activities other than running.

Lancer™’s smooth sole can cause runners to slip and fall on certain surfaces.

Lancer™ special fabric may smell if not cleaned properly.

Lancer™ are only available in white and do not come in other colors.

Lancer™ is a luxury running shoe brand that comes in many shapes and sizes.

Lancer™ customization process analyzes the way you run.

Lancer™ builds a unique shoe tailored especially for you.

Lancer™ shoes last up to a year longer than other brands.

2. One-Step Hair Dye

Please read the following consumer report about a new hair product called Flash Ultra Color™.

Flash Ultra Color™ transforms naturally dark hair into super reflective tones. Using recent advances in hair color technology, Flash Ultra Color™ allows the darkest brunette to achieve bright highlights in one easy step. The product includes premium grapeseed oils in its hair color formula, which is according to the company philosophy: “healthier hair means better color.” Flash Ultra Color™ spreads easily when applied, and has a no-drip guarantee. For customers who currently have color treated hair, there are special considerations. In some instances, hair color may fade in just a few days, and it is important to note that this product cannot be returned after opened. Also, Flash Ultra Color™ can cause severe allergic reactions, and may burn or irritate the scalp if applied to certain skin types. In addition, Flash Ultra Color™ provides an online color selection tool to help you find your preferred shade.

Your friend mentioned that she would like to purchase this product. Which items about the product are most important to tell her so that she makes a good decision?

Please select 7 statements from the following:

Flash Ultra Color™ may burn or irritate the scalp if applied to certain skin types.

Flash Ultra Color™ can cause severe allergic reactions.

Flash Ultra Color™ may fade in just a few days.

Flash Ultra Color™ cannot be returned after being opened.

Flash Ultra Color™ transforms naturally dark hair into super reflective tones.

Flash Ultra Color™ includes premium grapeseed oils in its hair color formula.

Flash Ultra Color™ spreads easily when applied, and has a no-drip guarantee.

Flash Ultra Color™ provides an online color selection tool to find your preferred shade.

3. Topical Acne Medication

Please read the following consumer report about a new acne medication called Nutane™.

Nutane™ is a cream-based acne medication. It is a mixture of exfoliating beads, cooling cream, and antimicrobial solution. Nutane™ is used to treat a variety of acne breakouts over the course of a week. Nutane™ should be applied in a thin layer, and then rinsed with cool water. As part of your skin’s natural absorption process, Nutane™ may cause your face to appear slightly greasy for the next few hours, and can smell strong while being applied. Some side effects exist. Nutane™ may burn if applied to certain skin types, and can also cause dizziness if used while dehydrated. Nutane™ is available over-the counter in all major pharmacies.

Your friend mentioned that he would like to purchase this product. Which items about the product are most important to tell Justin so that he makes a good decision?

Please select 7 statements from the following:

Nutane™ may burn if applied to certain skin types.

Nutane™ can cause dizziness if used while dehydrated.

Nutane™ may cause your face to appear slightly greasy for the next few hours.

Nutane™ may smell strong while being applied.

Nutane™ is a cream-based acne medication.

Nutane™ is a mixture of exfoliating beads, cooling cream, and antimicrobial solution.

Nutane™ should be applied in a thin layer.

Nutane™ is available over-the-counter in all major pharmacies.

Coding

The eight initial statements are coded, in their order, as: threat - threat - negative - negative - neutral - neutral - neutral - neutral. As the text is not reproduced, but we simply asked to select the statements, we just counted the proportion of the statements remaining.

Analysis

We use:

```
lmer(proportion ~ content + (1|chain_step) + (1|chain_id))
```

Where “proportion” is the proportion of statements of each content retained at each step and in each chain, and content can be one of “threat”, “negative”, or “neutral”.

Our hypothesis was that “threat” was a significant predictor of content retained relative to both negative and neutral content. In addition, we hypothesized negative content being a significant predictor of content retained relative to neutral content. To test this, we use the same model, but without including “threat” as a content.

Study 5: Multiple biases (social, gossip, survival, positive, negative, moral, counterintuitive)

Source: Berl et al. (2021)

Stories

1. 'Muki'

In the beginning times, Mata and Pata had run away together from the place of their people, far away beyond the realm of the sky, farther away than the stars stretch. The elder ones had not approved of their marriage and so they fled to our world here, which was then only a vast plain. They had with them their Child, Muki, who was the source of their greatest joys.

Here Mata rested with the sleeping Child at her breast, but the rest was short. Their people were hunting them and Muki would slow them when they most needed speed.

Pata took some black clay from his pack. Mata breathed upon the clay and shaped it with her hands, rolling it warm and round. Pata shaped it into a hammock with his hands, weaving its strands together, and hung it on the sky. Mata placed Muki carefully within the hammock, and pressed kisses to the Child's cheek. Then Mata plucked hairs from her own head, and hairs from Pata's chin, and scattered them across the ground. Pata struck his flint, which sent a bright, fiery spark up among the stars. The spark wandered about, bringing warmth and light to the world.

Mata and Pata swore to return one day for Muki. They whispered to her, "Sleep deeply, grow, and be loved."

When Muki woke from dreaming, wanting Mata's breast for her milk and Pata's steady hands for their comfort, they were no longer there. The Child beat her fists upon the earth until it quaked and shuddered. Muki cried and cried, until the spark in the sky darted away, leaving the Child and the world in darkness.

Muki grabbed fists full of clay and scraped out steep valleys in the land. The Child's tantrum churned up the hills as her kicking heels pounded in the earth. Her blood stained the clay, giving life, and she cried huge tears that became the Great River.

From the hairs of Mata's head, a forest of trees grew. The trees grew strong and tall and fruit of many colors sprouted from their branches.

On the shores of the Great River, the hairs of Pata's chin became spiders and crawled up from their bed of clay, moving on eight long, wiry legs. Muki snatched up a spider and pulled off each of its legs, one by one. The spider, wriggling from the pain, became a snake. Muki then tore the snake in two, dropping half into the water

This half-snake was now the swimming Fish. Muki tried to eat the other half, but Snake bared its fangs full of venom and Muki spat it out with a retch.

When another spider crawled up, this time Muki tore off four of its legs. This spider grew large and became the strong Wolf, which bounded away into the trees.

"Do not go near the Child" called the animals to the next spider who crawled up from the river. But this one was too clever to be caught. Six of its wiry legs twisted around each other

to form wings that began to flap. Soaring up above the forest, this spider became the clever Bird.

Then the last Spider summoned together the Snake, the Fish, the Wolf, and the Bird in the forest by the Great River. It was decided that the animals ought to live in different ways. Fish would have the winding river, Wolf would have the shadowed forests, and Snake would have the broad plains. Bird claimed the open skies for itself. "And I will have the hiding places," said Spider.

As they talked, Muki continued to shake the ground. Spider asked, "How will we stop the destruction of our homes?"

Bird said, "I flew high and fast, and I heard shouting voices in the sky. We should give the Child to these other people who search for her."

Snake disagreed, "We should not give the Child to those people, because they are not her people and seek to do her harm."

Wolf said, "Then we must be the Child's people, for it is right to look after children in need."

So Snake, Wolf, and Spider climbed across the Child's belly to tickle her until she laughed. Muki fell back onto the earth, happy and quiet.

Yet the world was still in darkness, as there was no light in the sky. Spider caught sight of the wandering spark in the sky and shouted, "You! You shall be the Sun for our world. When you light the skies and the land Muki will wake, and when you rest and the world darkens, Muki will sleep." The Sun shone proudly, for it is good to have purpose, and it brought the changing of the seasons.

Then out of the clay came our people, those who are our ancestors, because the land around Muki was good and fertile. The Child called Muki became the mountain that protects our village. We knew then as we do today that the Child must never be alone again, and we wait for Mata and Pata to return for her.

And when our people say today to our children, "Sleep deeply, grow, and be loved," we say it so that Muki hears this too and knows that she is not alone.

2. *'Taka and Toro'*

In the beginning times, Taka and her younger brother Toro were rowing through a storm and crashed upon a rocky island in the sea. Taka stepped ashore and the sharp rocks cut her feet. Everywhere her blood touched, life sprung forth. The grasses and the trees took root and the people, our ancestors, arose from the drops of blood.

Our ancestors learned from Taka and became her friends, and this made Toro feel jealous. She was always too busy playing with them.

"Oh-ho," cackled Puna, the bird in the palm, "Toro, who will you play with now?"

Toro was saddened by Puna's mocking words. He decided to make his own island. "It will be a new and BETTER island," Toro thought with a grin, "so then everyone will want to play with ME!"

Toro climbed up the palm and with his knife cut down the leaves. These he wove together, and bound with rope.

"Oh-ho," cackled Puna, "That is a raft, not an island, and you have not even performed the proper rites!"

But Toro was clever. He said, "If I promise to feed you and your family until your bellies are full, will you help me? If I give something to you, you must return the favor."

"Well," said Puna, "Our bellies are never full. But what is your plan?"

Placing his fingers in his mouth, Toro blew a whistle so piercing that Puna fell from her perch in terror. From out of the jungle came a mass of red ants.

"Why do you call us?" they demanded in their many tiny voices.

"I am building a new island," Toro announced, "And if you help me you can be the first to live there and can take the best homes for yourselves."

The ants agreed, and soon thousands of them came marching out of the jungle, carrying palm leaves on their backs. Toro continued to weave the leaves and the raft grew so large that it was bigger than Taka's island. The ants crawled up onto the raft.

Then, from the sky, dropped Puna and her family. When the birds were satisfied from feasting upon the ants, they grabbed hold of the edges of the new island and lifted it off of Taka's beach and into the sea. They pulled four times under the watch of the moon, and five under the watch of the sun, and they came to a place where the fish were many.

Toro swam down, down, and bound his island to the sea floor so it would not float away. Toro covered the island with soil and built up huge mountains from the land. The sea was pleased with Toro's new island and so sent coconuts to its shores. The coconuts sprouted into thick groves of palms. Toro was very proud. "My island is now the best of them all," he said to himself.

Yet still only Puna came, and that was to see whether there were more of the tasty ants. Toro was disappointed. He reached beneath the waves and found himself a crab. Toro said to the crab, "You must carry news of this island to my sister's people."

The crab, whose name was Kawa, narrowed his beady eyes and spoke, "That is a long way, and I am too lazy to swim that far."

Puna flew down and ate Kawa, because the lazy are always punished for their carelessness. Puna then carried Kawa's shell, filled with Toro's whispers of the island, and dropped it onto Taka's beach. All across Taka's island, the people began to speak of the rumors.

Hoki told his wife Otta, "At the other island, there are so many fish there is no room for them all in the sea. They leap out of the water and into a man's arms like a woman. There are also

many palms at the other island, and here there are more neighbours than trees. This is not as it should be."

Over the crashing of the waves, Otta did not hear her husband's words clearly. Otta went and told her sister Kohe, "My husband speaks of other women in his arms! He should not have broken my spirit in this way. If he wishes to go to the new place alone, I will not be sorry."

The people of Taka's island readied their boats together. When these people, our ancestors, came to Toro's island, the island that we call home, they saw that it was all that had been promised.

Taka was not happy that Toro had taken her playmates away. And so she sent the summer storms, making travel dangerous between our two homes. But Taka was an older sister, and like all older sisters, she loved her younger brother in spite of herself. And so she sent the winter trade winds which bring us prosperity and happiness.

We remember this, and today we praise the gods for our good fortune and celebrate.

Coding

Due to the combination of different types of content, coding here is more complex than in the previous studies. To illustrate this, see the first paragraph of 'Muki':

'In the beginning times, Mata and Pata had run away together from the place of their people, far away beyond the realm of the sky, farther away than the stars stretch. The elder ones had not approved of their marriage and so they fled to our world here, which was then only a vast plain. They had with them their Child, Muki, who was the source of their greatest joys.'

This is expressed in the original study in the following 18 propositions:

1. BEGINNING, TIMES/ 2. RUN AWAY FROM, MATA, PATA, PLACE OF PEOPLE/ 3. TOGETHER, MATA, PATA/ 4. FAR AWAY, PLACE OF PEOPLE 5. BEYOND, PLACE OF PEOPLE, REALM OF SKY/ 6. FARTHER AWAY THAN, PLACE OF PEOPLE, STARS STRETCH/ 7. NOT APPROVE OF, ELDER ONES, MARRIAGE/ 8. MARRIAGE, MATA, PATA/ 9. FLEE TO, MATA, PATA, WORLD/ 10. OUR, WORLD/ 11. HERE, WORLD/ 12. SO, 7, 9/ 13. PLAIN, WORLD/ 14. VAST, PLAIN/ 15. WITH, MATA, PATA, MUKI/ 16. CHILD, MUKI, MATA, PATA/ 17. SOURCE OF JOYS, MUKI, MATA, PATA/ 18. GREATEST, JOYS

As in other studies, propositions are coded as containing different types of informational content. Here the content is coded as social-basic, social-gossip, survival, emotional-positive, emotional-negative, moral, rational, and counterintuitive. Counterintuitive content is coded as mental counterintuitive, physical counterintuitive, and biological counterintuitive. For example the proposition TOGETHER, MATA, PATA is coded as social-basic; the proposition FARTHER AWAY THAN, PLACE OF PEOPLE, STARS STRETCH is coded as physical counterintuitive (a breach of folk physics); and the proposition FLEE TO, MATA, PATA, WORLD is coded as emotional-negative. Some propositions are coded as containing multiple types of information, for example the proposition NOT APPROVE OF,

ELDER ONES, MARRIAGE is coded as social-gossip and moral. Please see the coding guide provided by the authors of the original study for full coding details: <https://osf.io/ck64e>

Analysis:

We use:

```
lmer(proportion ~ content + (1|chain_step) + (1|chain_id))
```

Where “proportion” is the proportion of statements of each content retained at each step and in each chain, and content can be one of social-basic, social-gossip, survival, emotional-positive, emotional-negative, moral, rational, and counterintuitive (physical, mental or biological).

Our hypothesis, consistently with the original paper findings, is that social, negative emotional and counterintuitive (biology) is a significant predictors of proportion retained, versus the other biases.

Supplementary Text

Extended results

We report below results for single biases in experiments where, in the main analysis, they were collated in macro-categories according to the hypothesis we tested (study 1, 3, and 5) and/or where different stories were considered cumulatively (study 4 and 5). Figure S1 shows the results for study 1 where the biases PMC = plot-relevant male stereotype-consistent; PFC = plot-relevant female stereotype-consistent; BMC = background male stereotype-consistent; BFC = back-ground female stereotype-consistent were considered in the category “stereotype-consistent” and

PMI = plot-relevant male stereotype-inconsistent; PFI = plot-relevant female stereotype-inconsistent; BMI = background male stereotype-inconsistent; BFI = background female stereotype-inconsistent were considered in the category “stereotype-inconsistent”.

Figure S2 shows the results for study 3 where the biases Goss = social-gossip; Social = social-non-gossip were considered in the category “social” and Ind = non-social-individual; Phys = non-social-physical were considered in the category “social”.

Figure S3-S4-S5 show results for study 4, where the outcomes of the three different “consumer reports” were presented cumulatively.

Finally, figure S6 and S7 show results for study 5, where the outcomes of the two stories were presented cumulatively, and the biases CI_B (biological counterintuitive information), Social_basic (general social information), Social_gossip (gossip social information), and Emo_neg (emotionally negative information) were considered together against the other biases.

Fig. S1




Fig S1. Results for each content bias in study 1 (gender-stereotype consistent and gender-stereotype inconsistent information). The plot shows the proportion of information retained by chatGPT at each of the three steps of the transmission chain with respect to the original material. Each data point is an average of 5 replications, and bars represent standard deviation. PMC = plot-relevant male stereotype-consistent; PFC = plot-relevant female stereotype-consistent; PMI = plot-relevant male stereotype-inconsistent; PFI = plot-relevant female stereotype-inconsistent; BMC = background male stereotype-consistent; BFC = background female stereotype-consistent; BMI = background male stereotype-inconsistent; BFI = background female stereotype-inconsistent.

Fig. S2




Fig S2. Results for each content bias in study 3 (social and non-social information). The plot shows the proportion of information retained by chatGPT at each of the three steps of the transmission chain with respect to the original material. Each data point is an average of 5 replications, and bars represent standard deviation. Goss = social-gossip; Social = social-non-gossip; Ind = non-social-individual; Phys = non-social-physical.

Fig. S3




Fig S3. Results for “Lancer consumer report” in study 4 (threat-related information). The plot shows the proportion of statements retained by chatGPT at each of the five steps of the transmission chain with respect to the original material. Each data point is an average of 5 replications, and bars represent standard deviation.

Fig. S4

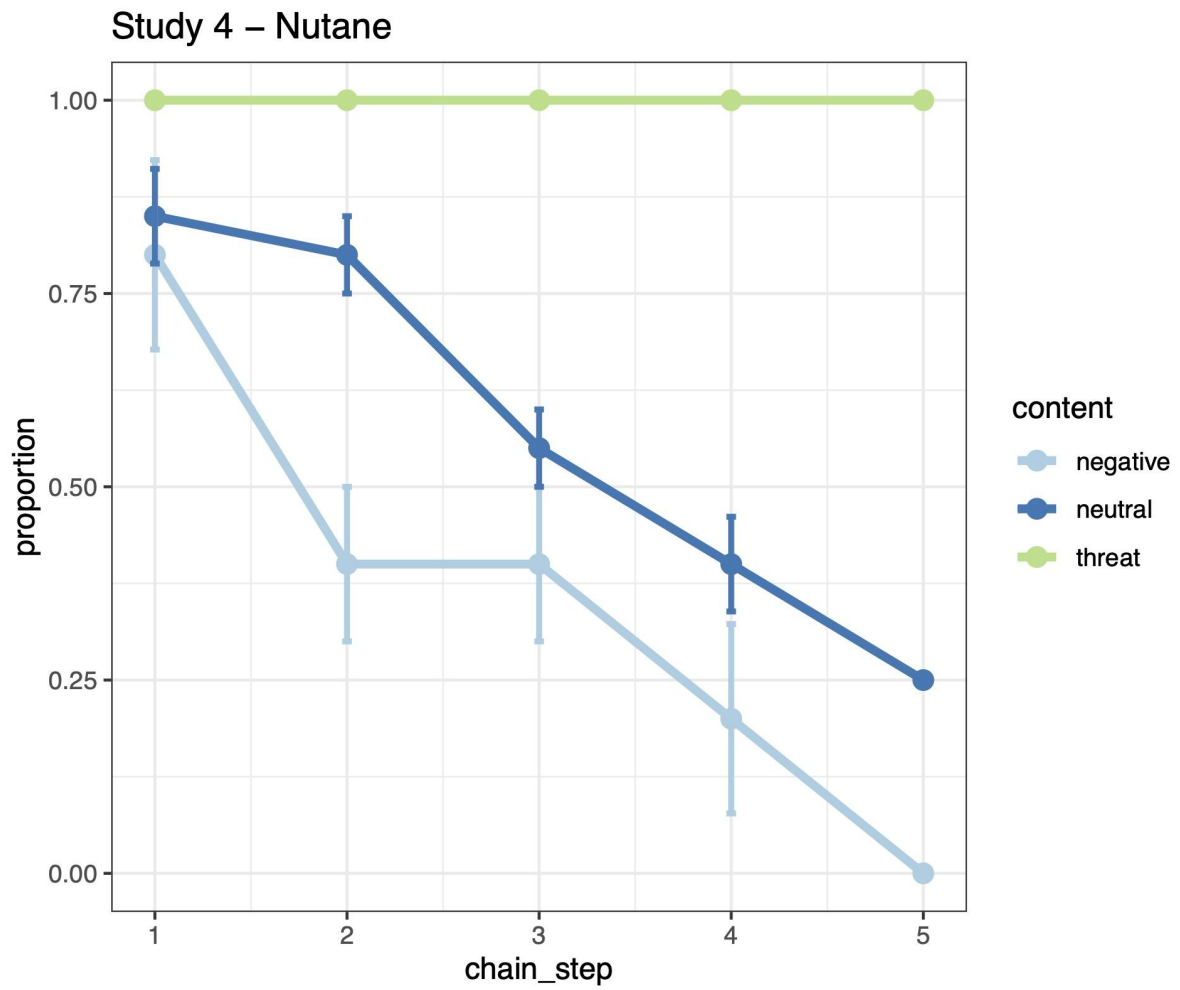


Fig S4. Results for “Nutane consumer report” in study 4 (threat-related information). The plot shows the proportion of statements retained by chatGPT at each of the five steps of the transmission chain with respect to the original material. Each data point is an average of 5 replications, and bars represent standard deviation.

Fig. S5

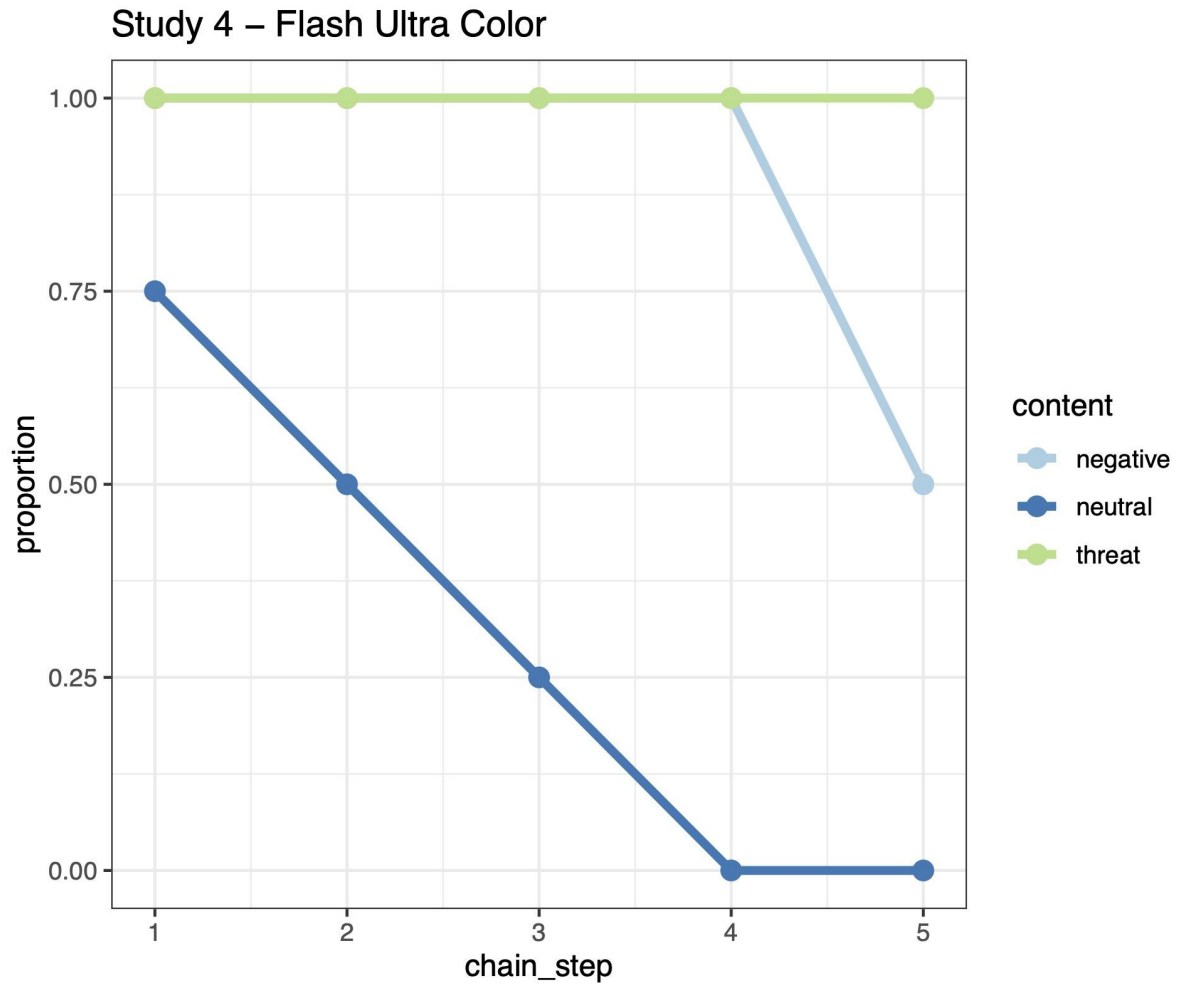


Fig S5. Results for “Flash Ultra Color consumer report” in study 4 (threat-related information). The plot shows the proportion of statements retained by chatGPT at each of the five steps of the transmission chain with respect to the original material. Each data point is an average of 5 replications, and bars represent standard deviation.

Fig. S6

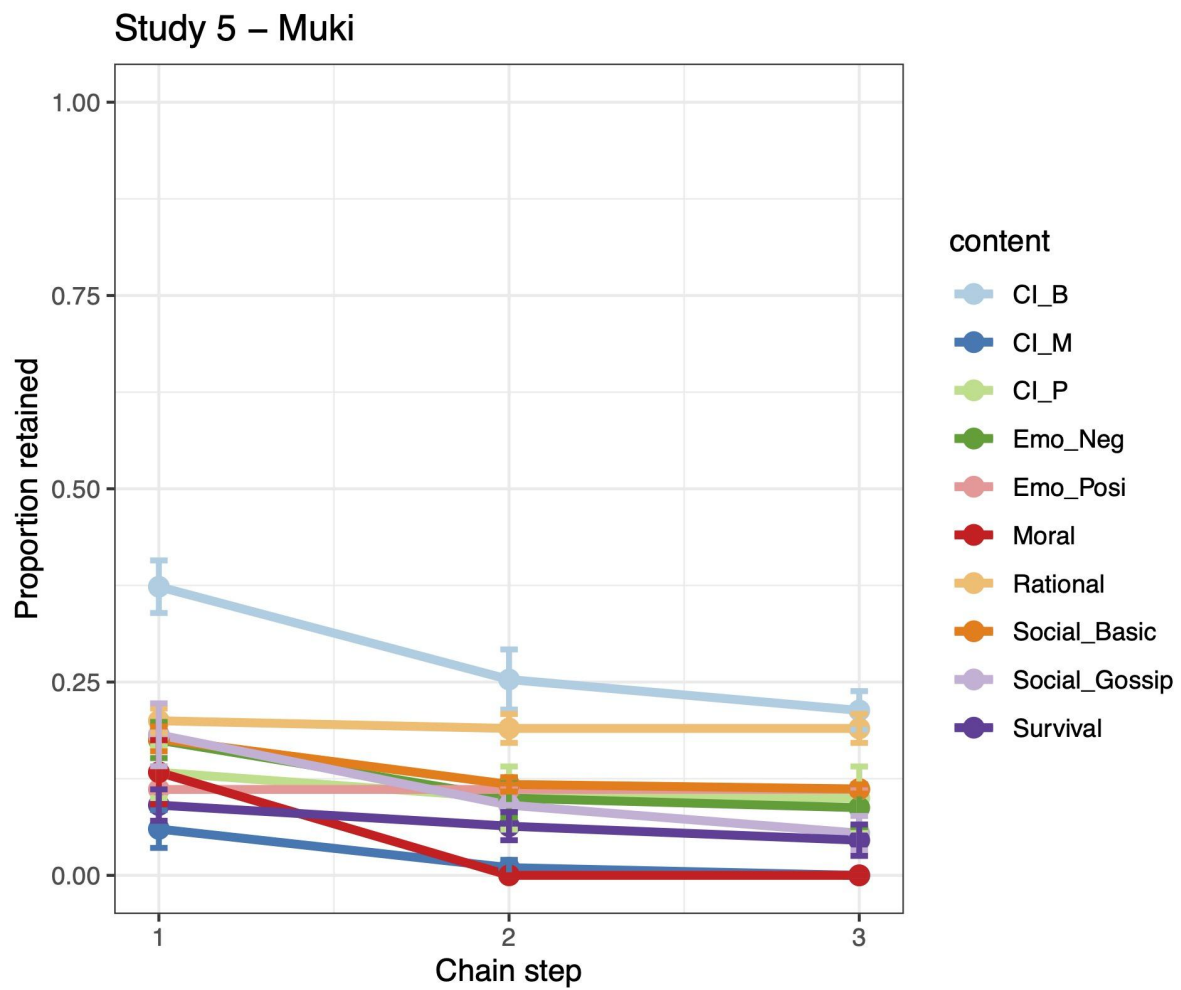


Fig S6. Results for each content bias in study 5 - “Muki” story. The plot shows the proportion of information retained by chatGPT at each of the three steps of the transmission chain with respect to the original material. Each data point is an average of 5 replications, and bars represent standard deviation. CI_B = counterintuitive information (biology); CI_M = counterintuitive information (mental); CI_P = counterintuitive information (physical).

Fig. S7

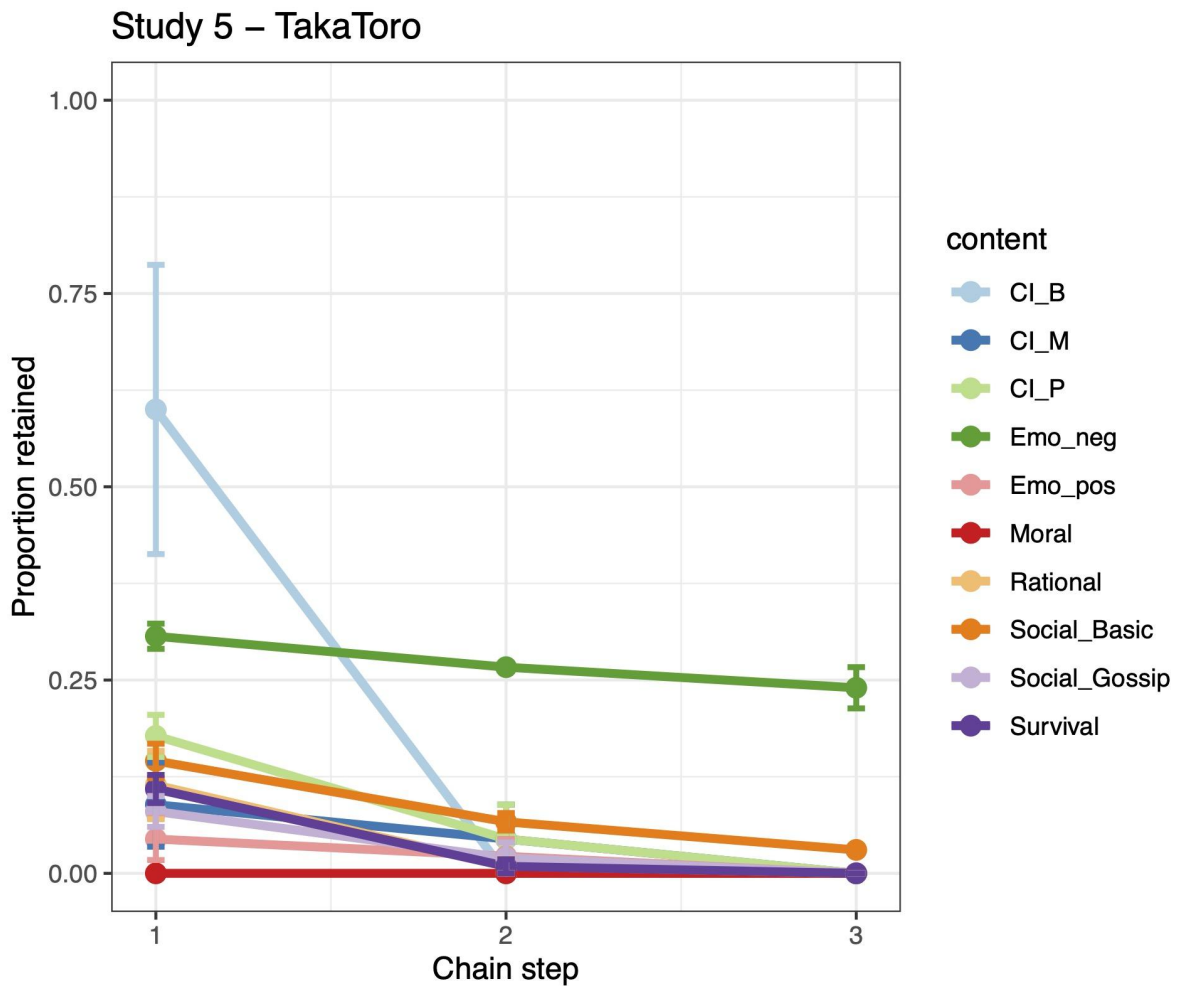


Fig S7. Results for each content bias in study 5 - “Taka & Toko” story. The plot shows the proportion of information retained by chatGPT at each of the three steps of the transmission chain with respect to the original material. Each data point is an average of 5 replications, and bars represent standard deviation. CI_B = counterintuitive information (biology); CI_M = counterintuitive information (mental); CI_P = counterintuitive information (physical).

Table S1

Study 1	$\kappa_w = .86$, (95% CI, .801 to .928)	$p < .001$
Study 2	$\kappa_w = .97$, (95% CI, .917 to 1.027)	$p < .001$
Study 2 (ambiguity resolution)	$\kappa_w = .87$, (95% CI, .772 to .966)	$p < .001$
Study 3	$\kappa_w = .82$, (95% CI, .764 to .882)	$p < .001$
Study 5 (Muki)	$\kappa_w = .79$, (95% CI, .744 to .843)	$p < .001$
Study 5 (Taka & Toro)	$\kappa_w = .83$, (95% CI, .770 to .884)	$p < .001$

Table S1. Inter-rater reliability. The table reports the inter-rater reliability (Cohen's Kappa coefficient) for the coding of chatGPT's outputs. Study 4 is not considered, as it did not involve the reproduction of a story, but the selection of items from a list.