Do Whales Have Hair? Are Whales Mammals? Identifying Synchronic Inconsistencies Among Beliefs

Joseph Sommer (joseph.sommer@princeton.edu)
Princeton University, University Center for Human Values and Program in Cognitive Science, Princeton, NJ 08540

Tania Lombrozo (lombrozo@princeton.edu)
Princeton University, Department of Psychology, Princeton, NJ 08540

Abstract

Inconsistency among beliefs is a hallmark of irrationality. Despite longstanding interest in inconsistency in philosophy and psychology, empirical evidence of synchronically held inconsistencies among people's belief has proven elusive. Here, across two pre-registered experiments (*N*s = 500, 274), we identify inconsistent beliefs simultaneously held by individual participants. Drawing on Sommer et al.'s (2023) proposal that accessibility in memory helps people achieve consistent beliefs, we constructed sets of questions that facilitated or hindered the accessibility of relevant knowledge. Our results support the proposal that consistency is enforced when beliefs are simultaneously accessible, rather than resulting from exhaustive consistency-checking. We find that when participants have simultaneous access to inconsistent beliefs–even regarding inconsequential general knowledge topics–they tend to revise their beliefs toward consistency. Furthermore, we experimentally distinguish an alternative explanation that the inconsistencies we evoked are merely inconsistent responses. Taken together, our results suggest that inconsistency among beliefs may be common, arising when inconsistencies are inaccessible.

Keywords: Belief; Consistency; Rationality; Accessibility

Introduction

The importance of maintaining consistency among beliefs has long been stressed by psychology (e.g., Festinger, 1957), decision theory (Salmon, 1966), and philosophy (Ramsey, 1929/1991). If someone has inconsistent beliefs, they can theoretically be induced to infer absurd conclusions or to commit to irrational decisions (Ramsey, 1926; Salmon, 1966). It is therefore unsurprising that inconsistency is an enduring topic of interest across numerous fields, with research addressing various forms of inconsistency.

For instance, famous work in judgment and decision making has found that people's heuristic judgments are inconsistent with normative rules (Tversky & Kahneman, 1974) and that people appear to display inconsistency (intransitivity) of preferences (Tversky, 1969). Research on lay and expert judges finds inconsistencies in their judgments across cases and over time (Kahneman, Sibony, & Sunstein, 2021). Similarly, work on intertemporal choice finds that people's choices are inconsistent across time delays (Ainslie, 2001). Additionally, research on consistency-checking in text has found that people make systematic errors in identifying inconsistencies in deduction (Johnson-Laird et al., 2004). Finally, the social psychology of attitudes has investigated several cognitive consistency theories (Abelson et al., 1968). Researchers in this tradition have documented inconsistencies between beliefs and behavior as well as between beliefs and attitudes (Fishbein & Ajzen, 1977).

Given this wide interest and long history of investigation, what is surprising is the absence of any direct evidence for the canonical case of inconsistency. To our knowledge, there is no systematic empirical demonstration that people have inconsistencies among synchronically held beliefs – that is, that they hold inconsistent beliefs at the same point in time. Our aim in this paper is to document such inconsistencies and to test an account of the conditions under which they arise.

Before introducing our theoretical and empirical approach, it is instructive to consider why three close candidates for documenting synchronic inconsistency fall short. First, Vul and Pashler (2008) find that people provide inconsistent answers to the same questions, such that averaging one's own answers can be leveraged to take advantage of the wisdom-of-crowds within a single person. However, such responses are judgments in response to questions participants have little knowledge about (e.g., "what percentage of the world's airports are in the U.S.?"). It is not obvious that such guesses should be considered beliefs. Second, consider work on conspiracy theories, which finds positive correlations in the endorsement of contradictory conspiracies (e.g., that Princess Diana was murdered *and* that she faked her own death; Wood, Douglas, & Sutton, 2012). However, subsequent work has found that this correlation is driven almost entirely by participants who *reject* sets of contradictory conspiracies, not by simultaneous endorsement (van Prooijen et al., 2023). Moreover, conspiracy beliefs might be subject to expressive responding (e.g., dishonesty or signaling opposition to authority; Douglas et al., 2019). Finally, work on *explanatory coexistence* has found that people sometimes hold multiple, potentially incompatible explanations for the same phenomenon (Metz, Weisberg, & Weisberg, 2020; Shtulman & Lombrozo, 2011). For example, the very same participant might offer microbial explanations for disease while also invoking witchcraft (Legare & Gelman, 2008). But such heterogeneity need not indicate inconsistency: participants could endorse different explanations in different cases, or integrate different explanatory frameworks (Davoodi & Lombrozo, 2023; Legare & Visala, 2011). In short, demonstrating true synchronic inconsistency may be more difficult than it seems (see Regenwetter et al., 2011 for a similar observation in the context of intransitivity of

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

preferences). Here, we suggest three criteria for identifying true synchronic inconsistencies. First, a single participant should respond with inconsistent answers without having the time and/or opportunity to learn new information between responses. Second, such responses should plausibly represent real beliefs that the participant holds (i.e., not guesses). Third, participants should acknowledge the inconsistency, to block alternative explanations for their inconsistent responses (for instance, that they misinterpreted the question). In this work, we attempt to meet all three criteria.

Assuming inconsistencies do exist, how are they to be explained? While most researchers seem to agree that people's beliefs are occasionally inconsistent, there is less agreement regarding how such inconsistencies might arise. Fragmentation theories propose that beliefs are stored in separate compartments in memory, and that consistency checking only operates within fragments (Bendaña & Mandelbaum, 2021), allowing inter-fragment inconsistency. Other theorists suggest that motivated or ideological beliefs are subject to different kinds of consistency checking, which allow desirable inconsistencies to persist (Abelson, 1986). Still others posit that conspiracy theorists might hold inconsistent beliefs that are consistent with broader claims, such as that authorities cannot be trusted. This wider consistency could allow belief in narrow claims to persist even if they contradict each other (Wood, Douglas, & Sutton, 2012). In contrast to these theories, Sommer et al. (2023) argue that inconsistencies might not need special explanation. Instead of viewing consistency among beliefs as the default, with departures from consistency requiring appeals to special mechanisms, Sommer et al. suggest that inconsistency is the default, and that it is consistency that needs mechanistic explanation. This is because exhaustive consistency checking is *computationally intractable* (Cherniak, 1986), which makes it exceedingly likely that some (perhaps many) of people's beliefs will be inconsistent. Sommer et al. also suggest that consistency checking (and correction) may only occur among beliefs that are accessible in memory at a given time, which offers a boundedly rational way to achieve partial consistency. If this proposal is correct, then when accessibility fails, people might be unable to recognize inconsistencies.

In this work, we aimed to identify inconsistencies among beliefs participants already hold. By asking people a series of questions designed – depending on the order of the questions – to either make knowledge accessible or inaccessible, we hypothesized that people would report inconsistent answers. In so doing, we tested two primary hypotheses.

First, the *no special explanation hypothesis* predicts that people possess inconsistent beliefs about mundane (i.e., non-motivated) topics. Additionally, such inconsistencies are predicted not to depend on relevant knowledge being scattered across topics, as fragmentation theories might demand. Instead, people's beliefs may be inconsistent simply because they have not yet been rendered consistent. As such, no motivational or partisan mechanism is required to explain why inconsistencies persist (though this is not to deny the possibility that such mechanisms play a role in some cases).

Second, the *accessibility hypothesis* predicts that synchronic inconsistencies can be revealed by eliciting beliefs that are unlikely to have been accessed simultaneously in the past, and that creating simultaneous access, in turn, reduces inconsistencies. Accordingly, we also assessed how participants respond to inconsistencies. While inconsistencies are theoretically irrational and often assumed to be emotionally aversive, it is an empirical question whether people will care about (cf. Levy et al., 2018) and/or reduce them. We allowed participants to choose whether to revise their responses and asked follow-up questions about their reasons for revising or not revising their answers.

Study 1

Study 1 examined whether people will synchronically report inconsistent beliefs when accessibility in memory is poor. We developed four triplets, comprising three interrelated questions, designed to make information accessible or inaccessible, depending on the order of the questions. Two example triplets are reproduced below:

1. Do whales have hair?
2. Are whales mammals?
3. Do all mammals have hair?
   (Triplet inspired by Greco, 2021; cf. Sloman, 1996).

1. What month of the year is the U.S. presidential vote certified?
2. What date was the U.S. Capitol building stormed in 2021?
3. What official governmental event was taking place in 2021 when protestors stormed the U.S. Capitol building?

Each participant views the same three questions for their assigned triplet. Our *opacity* manipulation only varies the order of presentation. *Transparent* triplets are presented in the order [3, 2, 1], while *Opaque* triplets are presented in the order [1, 2, 3]. We expected participants to be more likely to offer consistent responses in the *Transparent* order than in the *Opaque* order, because information relevant to answering Question 1 has just been rendered accessible.

Participants

We collected data from 549 adults based in the U.S. whose first language is English (305 Female, 227 Male, 15 Non-binary, 2 Preferred not to answer; mean age 38.4). Participants were recruited via Prolific and were paid $0.75 for a 3-minute study. We excluded 49 participants who failed one of two attention checks, resulting in a final sample of 500.

Procedure

Study 1 was fully between-subjects. Participants responded to one of the four triplets in either the Opaque or Transparent condition. As an attention check, participants received an additional triplet of simple questions. Each question was

multiple choice and contained the option for participants to respond, "I don't know." This option was included to reduce guesses, thereby increasing the likelihood that other response options reflect held beliefs. Following their responses, participants were presented with the full triplet they had previously responded to, along with their selected answers, and asked if they would like to change any of their responses. Finally, for exploratory purposes, we asked participants to write-in their reasons for revising or choosing not to revise their answers (we do not report responses here due to space). We were interested in three main DVs:

Inconsistent responses We coded participants' responses to each triplet for inconsistent patterns. To be conservative, participants who responded with "I don't know" to any question in a triplet were not classified as inconsistent.[1] We predicted that inconsistent responses would be more likely to in the Opaque, as opposed to Transparent, condition.

Revisions Participants had the opportunity to review their responses to the full Triplet and select whether they would like to change their answers to any question (coded as "Revise"), decline to revise, or 'Other' (which they could write in). We predicted greater revisions for inconsistent patterns compared to consistent patterns.

Revision-to-Consistency If a participant selected questions to revise, their revisions could reduce an inconsistency, fail to reduce an inconsistency, or create an inconsistency. We coded changes as revisions toward consistency if they targeted a question responsible for the inconsistency (note, however, that we did not require participants to indicate what they would change their response *to*). We predicted a greater tendency to revise toward consistency, especially among inconsistent participants.

Results

As predicted, we found that participants more often generated inconsistent patterns of responses in the Opaque condition (20.5%) than in the Transparent condition (8%; see Figure 1, Panel A). Note that this does not mean that the remaining 80-92% of participants were correct on all responses, only that they were not inconsistent by our coding (in fact, only 39% of the Transparent and 27% of Opaque participants provided answers that were both consistent and correct; other patterns included incorrect-but-consistent answers and patterns including "I don't know"). A multi-level binomial logistic regression predicting inconsistent responses, with Opacity (Opaque vs. Transparent) as a fixed effect and random effects for individual Triplets, revealed that the difference across conditions was significant (β = 1.08, SE = 0.28, p < .001).

Next, we investigated participants' choices to revise their answers. In the Transparent condition, 10.7% of participants revised an answer (vs. no revision or 'Other'). In the Opaque condition, 23% of participants revised an answer. A chi-square test revealed that these proportions differed significantly, $\chi 2(1, N = 500) = 12.73$, p < .001. To further analyze revisions, we conducted a multi-level logistic regression, with Inconsistency and Opacity as predictors, allowing random effects for Triplets. The model found that participants who provided inconsistent responses were significantly more likely to revise their answers than were participants whose original responses were consistent (β = 1.97, SE = 0.293, p < .001). Note that this result likely underestimates the impact of inconsistencies on revisions, as participants who responded "I don't know" to any of the questions in the triplet were considered consistent. There was a trend for more revisions among participants in the Opaque condition, but this result did not reach significance (p = .057; see Figure 1, Panel B).

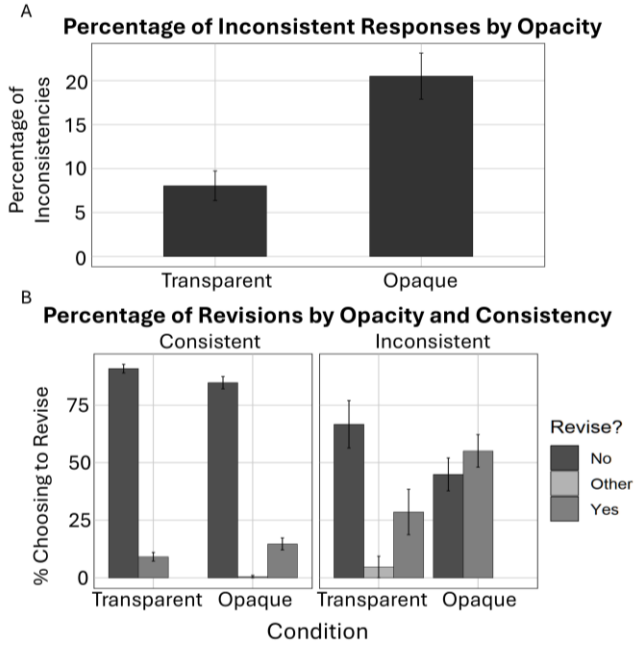

Figure 1: (A) Inconsistent responses by opacity. (B) Revision choices by opacity and consistency of original responses.

Finally, we examined the likelihood of participants' revisions reducing inconsistency. Of those participants who revised their answers, 69.9% revised toward consistency, while 30.1% of revisions did not reduce inconsistency, e.g., the answers revised were correct (and could not create consistent patterns even if they were revised to be incorrect). A chi-square test revealed that these percentages differ from chance, with revisions toward consistency more common than other revisions, $\chi 2(1, N = 83) = 13.12$, p < .001. Additionally, we ran a multi-level binomial logistic regression with Inconsistency and Opacity as predictors and random intercepts for Triplet. We found that inconsistent responses significantly predicted participants' revisions toward consistency (β = 3.39, SE = 0.58, p < .001), with no significant effect of Opacity ($p = 0.4$)

[1] Note that in some cases, "I don't know" can be inconsistent with other responses. For example, if one knows that whales are mammals and that all mammals have hair, whales must have hair. Responding "I don't know" to "Do whales have hair?" is therefore inconsistent, but was nonetheless categorized as consistent to ensure a conservative classification of inconsistencies.

Discussion

Results from Study 1 confirmed our predictions. Accessibility in memory appears to be relevant for producing consistent responses. Participants in the Opaque condition, for whom relevant information was less likely to be accessible, were significantly more likely to provide inconsistent responses relative to participants in the Transparent condition, for whom relevant information was made readily accessible. We also found that participants who reported inconsistent responses tended to recognize as much and revised their answers, typically toward consistency.

These results suggest that people do possess synchronically held inconsistencies in beliefs and that accessibility in memory helps them reduce these inconsistencies. When accessibility is blocked, people may be unable to remedy inconsistencies. Our finding that participants choose to revise their inconsistent answers suggests that these inconsistencies do not persist because people do not care to correct them, but because they fail to realize that they hold inconsistent beliefs until all the relevant information is made accessible.

Study 2

One limitation of Study 1 is that inconsistent responses and revisions might reflect mental states aside from true synchronic inconsistencies. Perhaps participants' beliefs are consistent, but they are responding inconsistently due to changes in their interpretation of the questions. To address this possibility, Study 2 considers "Ambiguity" triplets, which are predicted to behave like accessibility cases in yielding inconsistent responses when presented in an opaque order. However, we expect Ambiguity Triplets, unlike the Accessibility Triplets used in Study 1, to operate by altering participants' interpretations of the questions, not their beliefs about the answer to a question whose interpretation is fixed. For example, the Ducks Ambiguity triplet trades on the *generic overgeneralization effect* (Leslie et al., 2011). The three questions in this triplet are as follows:

1. Do all ducks lay eggs?
2. Do male ducks lay eggs?
3. Are male ducks considered ducks?

Question 1 suggests an interpretation of various species of ducks. However, the follow-up Question 2 suggests a very different scope for Question 1 (i.e., male vs. female ducks). Here, someone might respond that all ducks lay eggs, but that male ducks do not (despite male ducks being ducks). These responses appear inconsistent, but are a result of a shift in Question 1's interpretation, not in the person's beliefs. The goal of Study 2 is to contrast Ambiguity and Accessibility Triplets to increase confidence that Study 1's results are true inconsistencies rather than shifts of interpretation.

Participants

We collected data from 302 adults based in the U.S. whose first language is English (162 Female, 133 Male, 7 Non-binary, 1 Preferred not to answer; mean age 35.7). Participants were recruited via Prolific and were paid $1.00 for a 4-minute study. We excluded 28 participants who failed one of two attention checks, resulting in a final sample of 274.

Procedure

All participants in Study 2 answered two target Triplets – one Accessibility and one Ambiguity – along with an additional attention check "triplet," presented in random order. This ensured that participants were experienced with both types of Triplets when answering our question about the reasons for their revision choices. All Triplets in Study 2 were presented in the Opaque order. As in Study 1, following their responses, participants were presented – again in random order – with the full Triplets along with their selected answers. If a participant chose to revise their answers, they were presented with a question asking about their reasons for making revisions. If a participant chose not to revise their answers, they saw a similar question asking about their reasons for electing not to revise their answers.

In Study 2, we were interested in four DVs – three from Study 1, and one new DV for the purposes of disentangling Accessibility and Ambiguity items. In so doing, this measure is intended to support the interpretation that Accessibility Triplets induce true synchronic inconsistency.

Inconsistent responses We coded participants' responses to each triplet as inconsistent as in Study 1 (i.e., answer sets with "I don't know" responses to any question were considered consistent). We predicted that participants would report inconsistent patterns for both Accessibility and Ambiguity Triplets, though we were agnostic about any differences between them. If indeed Accessibility behaves like Ambiguity in inducing inconsistencies via misunderstandings of the questions, we might expect similar rates of inconsistencies across conditions.

Revisions Again like Study 1, participants had the opportunity to revise their responses while viewing the full triplet along with their prior answers. We coded responses as in Study 1 ("Other" responses were not considered revisions). We predicted greater revisions for inconsistent patterns.

Revision-to-Consistency We coded revisions as toward consistency vs. all other revisions. As in Study 1, we predicted a greater tendency to revise toward consistency, especially for inconsistent participants.

Reasons for Revisions We were interested in whether participants would assign blame to their own answers, which would suggest that they recognized their inconsistent responses as problematic, or to their interpretation of the question, which would suggest that they were tricked into providing a (merely) inconsistent response. We therefore provided all participants the opportunity to blame the question or the answer after their choice to revise/not revise their responses. When participants elected to revise their answers, they were asked "Why did you change your answer(s)?" with the options "My interpretation of the question hasn't changed, but I now think my original answer was wrong" (Answer blame), "My interpretation of the

question has changed, so I now think a different answer is appropriate" (Question blame), or "Other (please write in)." If participants chose not to revise their answers, they were asked "Why did you not change any of your answers?" They were presented with the options: "I think my answers are inconsistent, but I'm not sure what to do about it" (Answer blame), "I think there's a way to interpret each question such that each of my original answers is correct" (Question blame), and "Other (please write in)." Based on answers to these questions, we coded participants as assigning blame to the Question(s) or assigning blame to their Answer(s).

Results

We first examined inconsistent response patterns for Accessibility and Ambiguity Triplets. For Accessibility triplets, 21.9% of participants' responses were inconsistent, closely matching Study 1's Opaque condition. For Ambiguity triplets, 37.7% of responses were inconsistent, suggesting that Ambiguity Triplets also induce inconsistent responses (see Figure 2, Panel A). A multi-level binomial logistic regression predicting the likelihood of inconsistencies by Condition (Accessibility/Ambiguity), with random intercepts for Triplet, revealed that inconsistencies were significantly more likely for Ambiguity Triplets, relative to Accessibility Triplets ($\beta = 0.90$, SE = 0.38, $p = .019$).

Next, we compared participants' choices to revise an answer across conditions. We conducted a multi-level binomial logistic regression predicting revisions with fixed effects of Inconsistent responses and Condition, and random effects for the Triplets. We found that Inconsistent responses significantly predicted revisions ($\beta = 1.95$, SE = 0.33, $p < .001$). The effect of Condition and the interaction between Inconsistency and Condition were not significant (ps = .30 and .83, respectively), suggesting that choices to revise did not differ across conditions (Figure 2, Panel B).

We further investigated whether participants' revisions were toward consistency. Among participants whose answers were inconsistent, 92% of Ambiguity and 97% of Accessibility revisions were toward consistency. We conducted another multi-level binomial logistic regression predicting revisions toward consistency. The model included Condition (Accessibility/Ambiguity) and Inconsistency as fixed effects, with random intercepts for Triplet. There was no significant effect of Condition ($\beta = -0.07$, SE = 0.72, $p = .92$), suggesting revisions were not more likely to be toward consistency across triplet type. Paralleling results from Study 1, participants with inconsistent responses were significantly more likely to revise their answers toward consistency ($\beta = 2.07$, SE = 0.56, $p < .001$).

Finally, we examined whether participants blamed the question(s) or their answer(s) for inconsistencies. Of those participants who provided inconsistent responses, in the Ambiguity condition, only 8% of participants blamed their own answers, while 87% blamed their interpretation of the question (5% selected Other). In the Accessibility condition, 42% of participants blamed their answers while 56% blamed their question interpretation (2% Other; Figure 2, Panel C).

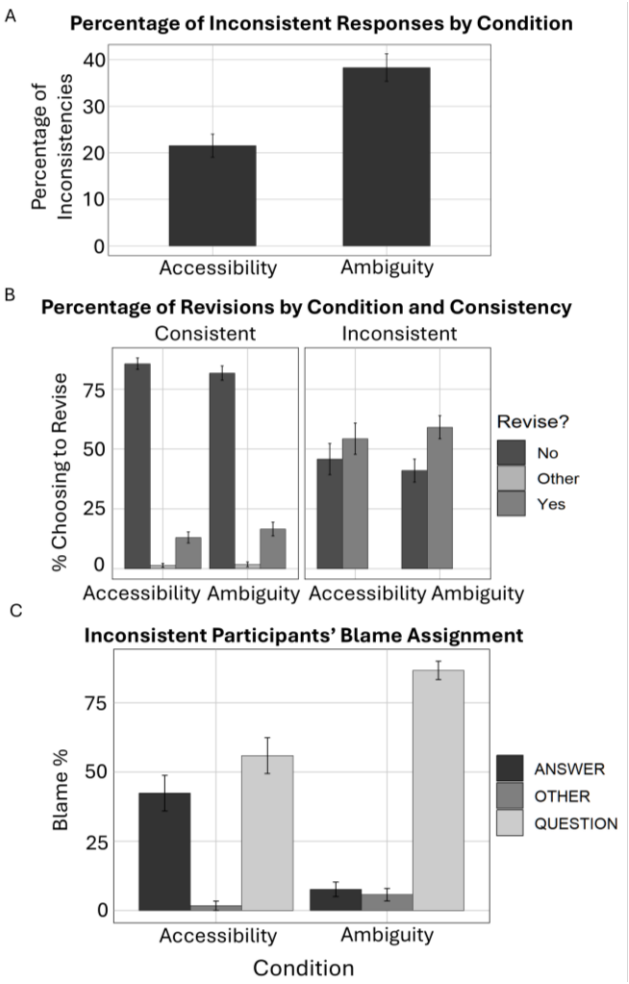

Figure 2: (A) Inconsistent responses by condition. (B) Choices to revise answers by condition and consistency of original responses. (C) Blame assignment by condition.

We modeled only participants who provided inconsistent responses with a multi-level binomial logistic regression predicting blame assignment with Condition and choice to Revise answers as fixed effects, along with random intercepts for Triplets. Results suggest that participants in the Ambiguity condition were significantly more likely to blame the question than those in the Accessibility condition ($\beta = 2.59$, SE = 0.50, $p < .001$). Additionally, participants who revised their responses were significantly less likely to blame the question than were those who chose not to revise their responses ($\beta = -2.23$, SE = 0.58, $p < .001$). This may suggest that participants are more inclined to revise answers they see as inconsistent while they retain answers that only appear inconsistent due to nuances of question interpretation.

Discussion

In Study 2, we replicated the elicitation of inconsistent responses for (Opaque) Accessibility triplets from Study 1. We also replicated results from Study 1 on choices of revision and revisions toward consistency. Participants who provided

inconsistent responses tended to revise their answers and were more likely to do so in the direction of consistency.

Going beyond Study 1, we also elicited inconsistent responses to our Ambiguity triplets. We were concerned that inconsistent *responses* may not reflect inconsistent *beliefs*. We addressed this concern by testing whether participants assigned blame for their inconsistent responses to their own answers or to their interpretation of the question. We predicted that participants who provided merely inconsistent *responses* would be less likely to identify their own answers as the source of the problem. Our Ambiguity Triplets, which *are* predicted to induce only inconsistent responses, serve as a comparison for this measure of question blame. That very few participants blame their own answers on Ambiguity items suggests that this measure is operating as intended. The higher rates of "answer blame" in the Accessibility condition are consistent with our assumption in Study 1 that inconsistent responses often reflect inconsistent beliefs.

The results of Studies 1-2 thus offer the necessary ingredients to satisfy our recipe for demonstrating genuine inconsistency in synchronically held beliefs. First, a single participant should respond with inconsistent answers without having the time and/or opportunity to learn new information between responses. Both of our experiments meet this criterion. Second, such responses plausibly represent real beliefs that the participant holds (i.e., not guesses). Our design addresses this through the choice of familiar general knowledge stimulus materials and via the response option of "I don't know," which we conservatively refrain from coding as inconsistent. Third, participants should acknowledge an inconsistency in beliefs (not just responses). This recipe is met in total by the 8% of participants in Study 2 who generated an inconsistent set of responses, recognized the need for revision, *and* attributed the need for revision to their answers. This represents a highly conservative estimate for the prevalence of true inconsistency in our sample.

General Discussion

Holding inconsistent beliefs is widely recognized as irrational. The question of how such irrationality is possible has been debated by psychologists for decades. Researchers often appeal to special explanations for the origins of inconsistency, including fragmented memory (Borgoni et al., 2021), motivational forces (Rosenberg & Abelson, 1960), or irrational/a-rational consistency checking for ideological beliefs (Abelson, 1986). In contrast, we provide support for two alternative hypotheses. First, in accordance with the *no special explanation hypothesis*, inconsistencies do not appear to require motivation or fragmentation (cf. Sommer et al., 2023). Second, in line with *the accessibility hypothesis*, we found that inconsistencies can arise in the absence of simultaneously accessible information.

Across two pre-registered studies, participants provided inconsistent sets of responses to general knowledge questions. Inconsistencies were significantly more common when information was relatively inaccessible (Study 1), suggesting a role for accessibility in reducing inconsistencies. In both Study 1 and Study 2, participants recognized their inconsistencies as such and wished to revise their answers toward consistency. This suggests that participants operate with a norm favoring consistency in belief. Furthermore, in Study 2, participants distinguished between inconsistent responses due to question interpretation and inconsistencies where the fault lay with their own answers. This suggests that participants' inconsistencies (in our Accessibility conditions) are not just inconsistent responses, and may instead represent true inconsistencies among synchronically held beliefs.

Our results have implications for our understanding of the psychology of belief and related fields interested in consistency (e.g., coherence-based reasoning; Simon & Read, 2023). First, we believe these results are the first systematic demonstration that people simultaneously hold inconsistent beliefs. Moreover, the beliefs in question are not ideological or motivated, suggesting that no additional forces (Rosenberg & Abelson, 1960) are needed for inconsistencies to persist. Nor is it obvious that a fragmentation account (Borgoni et al., 2021) can readily explain our results. If fragments are individuated by themes or topic, it is plausible that all items within each triplet would fall into a common fragment. On the other hand, if fragments are so fine-grained that they correspond to individual statements, then it is not clear that fragments operate as units that do any explanatory work. That said, most fragmentation accounts are still in their infancy and may not make specific predictions about our triplets. Future research will be required to determine whether these results fit with fragmentation. In the meantime, our results are consistent with Sommer et al.'s (2023) claim that inconsistency does not need special explanation via motivation or fragments. Instead, inconsistencies may only be corrected on the fly when the physical world itself (Craik, 1943) or accessibility aids in forming consistent beliefs.

It might be argued that the inconsistencies we identified are only between *dispositional beliefs.* Philosophers (e.g., Audi, 1994) distinguish *occurrent* beliefs, which are currently entertained, from dispositional beliefs that one has a tendency to entertain under the appropriate circumstances, even if one has never thought of them before. While our Triplets are most readily interpreted as tapping dispositional beliefs, this dispositional nature does not seem to weigh heavily against attribution of inconsistency. The inconsistent dispositions co-occur within the same, very brief context, where participants have little opportunity to learn new information that might rationally shift their dispositions. Additionally, the revisions participants made are also indicative of an (albeit short-lived) inconsistency in occurrent beliefs. However, future research might explicitly examine the relationship between inconsistency and occurrent vs. dispositional beliefs.

In describing the progress of early cognitive consistency theories, McGuire (1968) expressed disappointment that "the consistency need has become more of an end to be studied in its own right, rather than a means for throwing light on thought processes." By using inconsistencies to shed light on accessibility processing, we hope this work contributes to a broader understanding of thinking and believing.

References

Ainslie, G. (2001). *Breakdown of Will*. Cambridge University Press.

Abelson, R. P. (1986). Beliefs are like possessions. *Journal for the Theory of Social Behaviour, 16*, 223–250.

Abelson, R. P., Aronson, E. E., McGuire, W. J., Newcomb, T. M., Rosenberg, M. J., & Tannenbaum, P. H. (1968). *Theories of Cognitive Consistency: A Sourcebook*. Rand McNally and Company.

Audi, R. (1994). Dispositional beliefs and dispositions to believe. *Noûs*, *28*(4), 419-434.

Bendaña, J., & Mandelbaum, E. (2021). The fragmentation of belief. In C. Borgoni, D. Kindermann, & A. Onofri (Eds.). *The Fragmented Mind*. Oxford University Press.

Borgoni, C., Kindermann, D., & Onofri, A. (2021). *The Fragmented Mind*. Oxford University Press.

Cherniak, C. (1986). *Minimal Rationality*. MIT Press.

Craik, K. J. W. (1967). *The Nature of Explanation*. Cambridge University Press.

Davoodi, T., & Lombrozo, T. (2023). Scientific and religious explanations, together and apart. In J. N. Schupbach & D. H. Glass (Eds.). *Conjunctive Explanations* (pp. 219-245). Routledge.

Douglas, K. M., Uscinski, J. E., Sutton, R. M., Cichocka, A., Nefes, T., Ang, C. S., & Deravi, F. (2019). Understanding conspiracy theories. *Political Psychology, 40*, 3-35.

Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford University Press.

Fishbein, M., & Ajzen, I. (1977). *Belief, Attitude, Intention, and Behavior: An Introduction to Theory and Research*. Addison Wesley Publishing.

Greco, D. (2021). Fragmentation and Coarse-Grained Content. In C. Borgoni, D. Kindermann, & A. Onofri (Eds.). *The Fragmented Mind*. Oxford University Press.

Johnson-Laird, P. N., Girotto, V., & Legrenzi, P. (2004). Reasoning from inconsistency to consistency. *Psychological Review, 111*(3), 640.

Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A Flaw in Human Judgment*. Hachette UK.

Legare, C. H., & Gelman, S. A. (2008). Bewitchment, biology, or both: The co-existence of natural and supernatural explanatory frameworks across development. *Cognitive Science*, *32*(4), 607-642.

Legare, C. H., & Visala, A. (2011). Between religion and science: Integrating psychological and philosophical accounts of explanatory coexistence. *Human Development*, *54*(3), 169-184.

Leslie, S. J., Khemlani, S., & Glucksberg, S. (2011). Do all ducks lay eggs? The generic overgeneralization effect. *Journal of Memory and Language, 65*(1), 15-31.

Levy, N., Harmon-Jones, C., & Harmon-Jones, E. (2018). Dissonance and discomfort: Does a simple cognitive inconsistency evoke a negative affective state? *Motivation Science, 4*(2), 95.

McGuire, W. J. (1968). Theory of the structure of human thought. In R. P. Abelson, E. E. Aronson, W. J. McGuire, T. M. Newcomb, M. J. Rosenberg, & P. H. Tannenbaum (Eds.). *Theories of Cognitive Consistency: A Sourcebook*. Rand McNally and Company.

Metz, S. E., Weisberg, D. S., & Weisberg, M. (2020). A case of sustained internal contradiction: Unresolved ambivalence between evolution and creationism. *Journal of Cognition and Culture*, *20*(3-4), 338-354.

Ramsey, F. P. (1926). Truth and probability. In D. H. Mellor (Ed.). *Philosophical Papers: F. P. Ramsey*. Cambridge University Press.

Ramsey, F. P. (1929/1991). On Truth. In N. Rescher & U. Majer (Eds.), *On Truth: Original manuscript materials (1927–1929) from the Ramsey collection at the University of Pittsburgh*. Springer Science & Business Media.

Regenwetter, M., Dana, J., & Davis-Stober, C. P. (2011). Transitivity of preferences. *Psychological Review, 118*(1), 42.

Rosenberg, M. J., & Abelson, R. P. (1960). An analysis of cognitive balancing. In H. Rosenberg, A. McGuire, & J. W. Brehm (Eds.), *Attitude Organization and Change: An Analysis of Consistency Among Attitude Components*. Yale University Press.

Salmon, W. C. (1966). *The Foundations of Scientific Inference*. University of Pittsburgh Press.

Simon, D., & Read, S. J. (2023). Toward a general framework of biased reasoning: Coherence-based reasoning. *Perspectives on Psychological Science*, 17456916231204579.

Shtulman, A., & Lombrozo, T. (2016). Bundles of contradiction. In D. Barner & A. S. Baron (Eds). *Core Knowledge and Conceptual Change*, pp. 53-72. Oxford University Press.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119*(1), 3.

Sommer, J., Musolino, J., & Hemmer, P. (2023). A hobgoblin of large minds: Troubles with consistency in belief. *Wiley Interdisciplinary Reviews: Cognitive Science, 14*(4), e1639.

Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *Science, 185*(4157), 1124-1131.

Tversky, A. (1969). Intransitivity of preferences. *Psychological Review, 76*, 31–48.

van Prooijen, J. W., Wahring, I., Mausolf, L., Mulas, N., & Shwan, S. (2023). Just dead, not alive: Reconsidering belief in contradictory conspiracy theories. *Psychological Science, 34*(6), 670-682.

Vul, E., & Pashler, H. (2008). Measuring the crowd within: Probabilistic representations within individuals. *Psychological Science*, *19*(7), 645-647.

Wood, M. J., Douglas, K. M., & Sutton, R. M. (2012). Dead and alive: Beliefs in contradictory conspiracy theories. *Social Psychological and Personality Science, 3*(6), 767-773.