

Temporal Explanations

Laura Jane Kelly and Sangeet Khemlani

Navy Center for Applied Research in Artificial Intelligence, US Naval Research Laboratory, Washington, District of Columbia, United States

People can explain phenomena by appealing to temporal relations, for example, you might explain a colleague's absence at a meeting by inferring that their prior meeting ended late. Previous explanatory reasoning research shows that people construct causal explanations to resolve causal conflicts. Accordingly, temporal explanations may help reasoners resolve temporal conflicts, and we describe four experimental tests of the hypothesis ($N = 240$). Experiment 1 provided participants with conflicting or consistent temporal information and elicited natural responses about what followed. Participants spontaneously provided temporal explanations to resolve inconsistencies, and only a minority of them provided more conservative, direct refutations. Experiments 2 and 3 showed that participants preferred temporal explanations over simpler refutations to resolve conflicts, and Experiment 4 showed that participants judged temporal explanations more probable than refutations, and thereby yielded a novel class of conjunction fallacies. The research is the first to examine patterns in temporal explanatory reasoning.

Keywords: temporal explanations, conflicts, inconsistencies, duration, event cognition

These three sentences are inconsistent with one another:

The football game went from 1 pm to 5 pm.
Ria arrived at the game at 6 pm.
She attended the game.

Ria is the Queen; she shifted the game's schedule. [causal]
Ria attended a different game. [spatiotemporal]
The game was delayed. [temporal]

The situation is impossible: how can Ria attend the game if she arrived after it occurred? One of the three sentences must be false, that is, they cannot all be true at the same time. Psychologists since William James have argued that people cope with conflicts by minimally revising their information, that is, they reject as few of the sentences as possible. But recent work shows that reasoners generate explanations to resolve conflicts (Khemlani & Johnson-Laird, 2011, 2012). Here are some plausible explanations that might suffice:

Ria attended the game virtually. [spatial]
The posted schedule was wrong. [epistemic]

Explanations help reasoners understand the past and predict the future (Anderson et al., 1980; Craik, 1943; Einhorn & Hogarth, 1986; Gopnik, 2000; Lombrozo & Carey, 2006; Ross et al., 1977), and many cognitive scientists argue that they are a hallmark of human rationality (Harman, 1965; Horne et al., 2019; Johnson-Laird, 1983; Lombrozo, 2007), though they also serve as the basis of magical thinking, conspiracy theory, and pseudoscience (Gronchi et al., 2017; Weisberg et al., 2008). They add new relations or entities into the discourse that were not present in the given information (e.g., a virtual game, the posted schedule, Ria's royal status). The consequences of a particular explanation can help guide reasoners to reject information, for example, if a reasoner infers that Ria attended the game virtually, then they may tacitly reject the second premise that she arrived at the game at 6 p.m. Many types of explanation can help reasoners resolve the conflict, for example, a spatial explanation appeals to spatial locations, and an epistemic explanation appeals to knowledge and belief.

The majority of empirical research into explanatory reasoning has focused on how people assess causal explanations (e.g., Ahn & Kalish, 2000; Alicke et al., 2015; Fernbach et al., 2012; Johnson-Laird et al., 2004; Lombrozo, 2016; Sloman, 2005). One reason for the focus on causality may be because people tend to spontaneously generate causal explanations when given the opportunity in the laboratory (Khemlani & Johnson-Laird, 2011) and in more natural contexts (Zemla et al., 2017). In particular, reasoners are more likely to generate causal explanations to resolve conflicting, inconsistent information rather than to elaborate on a consistent description (Khemlani & Johnson-Laird, 2011, 2012, 2013; Legare, 2012). Consider this description from Khemlani and Johnson-Laird (2013):

Sangeet Khemlani  <https://orcid.org/0000-0002-1158-553X>

Action editor: Michele Diaz, PhD

This research was performed while the first author held an NRC Research Associateship award at the U.S. Naval Research Laboratory. We are grateful to Kalyan Gupta and Danielle Paterno at the Knexus Research Corporation for their help in conducting the experiments. And we thank Bill Adams, Gordon Briggs, Monica Bucciarelli, Hillary Harner, Tony Harrison, Phil Johnson-Laird, Joanna Korman, Markus Knauff, Andrew Lovett, Greg Trafton, and Ben Wright for their advice and comments. Experiments 1 and 2 were presented at the Annual Conference of the Cognitive Science Society in 2021 and published in its proceedings (Kelly & Khemlani, 2021), as well as disseminated in a talk at the University of Granada in 2022. Supplementary materials, including raw data, statistical analyses, and code for the experiments, are available at <https://osf.io/atyv8/>.

Correspondence concerning this article should be addressed to Sangeet Khemlani, Navy Center for Applied Research in Artificial Intelligence, US Naval Research Laboratory, 4555 Overlook Ave SW, Washington, District of Columbia 20375, United States. Email: sunny.khemlani@nrl.navy.mil

119 If a person does regular aerobic exercises then the person strengthens her
 120 heart.
 121 Someone did regular aerobic exercises, but she [*did/did not*] strengthen
 122 her heart.
 123 What, if anything, follows?

124 If the woman strengthened her heart, there is nothing to explain, and
 125 reasoners often respond that “nothing follows.” If she did not
 126 strengthen her heart, however, the two premises are inconsistent
 127 with one another, that is, reasoners can draw contradictory conclu-
 128 sions from them. They often infer explanations to eliminate the con-
 129 flict, for example,

131 Perhaps she has a health condition that prevents her heart from getting
 132 stronger.

134 Since *prevention* is a causal relation, the explanation is causal in
 135 nature, and reasoners rely on causal knowledge to resolve other
 136 kinds of inconsistencies, too (Khemlani & Johnson-Laird, 2020).
 137 They need not have constructed an explanation: a more conservative
 138 response would be to directly refute the premises, for example, rea-
 139 soners could infer that the first premise is strictly false (it describes a
 140 generalization that has exceptions). But, as studies show, people pre-
 141 fer causal explanations to refutations and generate them more often.

142 In the above example, the first premise describes a causal relation:
 143 regular aerobic exercises cause a person’s heart to strengthen. So, in
 144 retrospect, it may not be surprising that reasoners infer causal expla-
 145 nations to resolve causal conflicts. But conflicts can arise in other
 146 sorts of information, too, for example, temporal descriptions can
 147 contain conflicts. Reasoners can make sophisticated inferences
 148 about time and duration (Hoerl & McCormack, 2019), and they
 149 can detect conflicts in temporal descriptions. Consider this descrip-
 150 tion (from Kelly et al., 2020):

152 The meeting happened during the conference.
 153 The sale happened before the conference.
 154 The meeting happened before the sale.

156 The three sentences cannot all be true at the same time, and reasoners
 157 have little difficulty detecting the inconsistency—indeed, they
 158 sometimes assess even consistent descriptions as inconsistent
 159 because they fail to consider all the ways the sentences can be true.

160 When people detect an inconsistency in temporal information,
 161 they may try to explain it by introducing new events and temporal
 162 relations to the discourse, that is, they may try to construct tem-
 163 poral explanations. Cognitive scientists have yet to examine tempo-
 164 ral explanations, and no studies have assessed whether people
 165 make them in response to temporal conflicts. This article accordingly
 166 examines how reasoners construct and evaluate explanations of time
 167 as a way to cope with conflicts. Four studies test the hypothesis that
 168 conflicts should prompt reasoners to generate temporal explanations
 169 and consider them as sensible more often than more conservative ref-
 170 utations. Experiment 1 showed that people produce temporal expla-
 171 nations to resolve conflicting information; Experiment 2 revealed
 172 cases in which people prefer temporal explanations to refutations,
 173 and Experiment 3 controlled for a confound in Experiment 2 and
 174 further shows that participants prefer explanations to refutations.
 175 Experiment 4 showed that people consider temporal explanations
 176 more probable than refutations. We conclude by describing how
 177 temporal explanations differ from other kinds of explanations and

180 why they are a particularly helpful strategy for resolving conflicts
 181 in information.

182 Experiment 1

183 Experiment 1 sought to test whether reasoners can generate tempo-
 184 ral explanations in a systematic way. It gathered reasoners’ natural
 185 responses to conflicting temporal information. Participants typed out
 186 their responses to problems such as:

187 Suppose that you are told the following:

188 The blood drive was open from 9 am to 4 pm on Monday.
 189 Trisha arrived at the blood drive at 5 pm on Monday.
 190 You discover the following fact:
 191 Trisha gave blood at the blood drive.
 192 What, if anything, follows?
 193

194 The set of premises is inconsistent because they describe a scenario
 195 in which Trisha gave blood at a blood drive after it had closed. The
 196 study varied whether the premises described consistent or inconsis-
 197 tent scenarios; previous work on causal explanations suggests that
 198 reasoners should generate explanations more often for inconsistent
 199 scenarios.

200 Method

202 Participants

203 We used the *pwr* package (Champely et al., 2018) in R to conduct a
 204 power analysis for a single item in our study. Our goal was to obtain
 205 0.75 power to detect a medium-large effect ($d = 0.4$) at 0.05 α error
 206 probability, so 45 participants were required for the study. This exper-
 207 iment and subsequent ones collected demographic information prior
 208 to the study proper; demographic surveys asked participants about
 209 their age, sex (*male, female, other*, or *prefer not to say*), native lan-
 210 guage (*English* or *other*), and the number of courses in logic they
 211 had taken (*none, one introductory, some introductory, a few*
 212 *advanced, lots of advanced*). Fifty-one participants completed the
 213 experiment for monetary compensation (\$2) through Amazon’s
 214 Mechanical Turk. Six participants produced a majority of nonsensical
 215 responses, so we dropped their data. The analyses reported are based
 216 on the remaining 45 participants (18 female, 27 male, $M_{age} = 37.2$).
 217 The participants were native English speakers, and six had taken
 218 one or more courses in logic.
 219

220 Preregistration and Data-Availability

222 The experimental designs, predictions, and analyses for
 223 Experiments 1–4 were pre-registered through the Open Science
 224 Framework platform (<https://osf.io/atyv8/>). The same link provides
 225 the corresponding experimental code, materials, data, and data cod-
 226 ing rubric.
 227

228 Materials, Procedure, and Design

229 Participants completed eight problems which each presented a
 230 participant with information concerning the duration of an event,
 231 information about when an individual arrived at the event, and
 232 information about whether or not that individual took part in the
 233 event. Provided that an individual can take part in an event only
 234 if they arrive sometime between when the event started and when
 235 it ended, the premises in each problem could conflict with
 236

one another. For instance, the set of premises provided in the above example is inconsistent, because Trisha arrived after the drive ended—but a change to the first premise makes the set of premises consistent:

The blood drive was open from 9 am to 8 pm on Monday.

In this description, Trisha arrived during the event hours and therefore it is consistent to say she gave blood at the event. Half of the problems were consistent and half were inconsistent; the experiment randomized the consistency of each problem. Participants typed their response to the question, “What, if anything, follows?” into a response box. The experiment required them to type a response with at least 1 character for each problem. Each participant carried out the eight problems in a different random order.

The premises of the problems in Experiment 1 came from eight separate scenarios that concerned everyday events (e.g., attending a class, speaking at a meeting, picking up a prescription), and each set of premises was made consistent or inconsistent by manipulating the interval of the event described. There was an error in one scenario such the timeline was incoherent if taken literally, that is, “The party was scheduled to occur from 7 pm to 12 am on Friday.” The analysis below focuses on the remaining seven; excluding the erroneous scenario had no qualitative effect on the results.

Rejection and Coding Criteria

The first author coded participants’ typed responses. Responses that appeared nonsensical, copied from the premises, or otherwise inappropriate to the task were dropped from further analysis (14% of the data from the 51 original participants; 2% of the data from the 45 participants included in the analyses). The subsequent analyses concerned the remaining 309 responses (see Table 1 for examples). If participants generated more than one plausible response (this occurred for 6% of the trials), we coded only their first response.

We coded responses on the following four criteria:

1. *Did the response directly refute one of the premises?* Each response was coded on whether it explicitly denied the truth of one of the premises (e.g., “The blood drive didn’t close at 4 pm...”).

2. *Did the response explain the premises by reference to some temporal concept?* Temporal explanations are responses that introduce a new temporal relation, for example, “The class was pushed back that day,” or a new event, for example, “The pottery class had so many people turn out that they had to split the group into two classes. Matteo attended the later class.”
3. *Did the response explain the premises in some other way?* Other explanations concerned responses that introduce non-temporal entities or relations, such as spatial, epistemic, or causal relations, or else relations that were ambiguous in nature.
4. *Which premise did a response refute or explain?* Refutations or explanations could concern the event’s time interval (premise 1); the time an individual arrived at the event (premise 2); or whether the individual attended the event (premise 3). Responses that were equivocal were not considered for further analyses.

Results and Discussion

Participants produced temporal explanations in 32% of the trials and direct refutations in 13% of the trials; Table 1 provides a breakdown of the different types of responses. Participants produced more explanations for inconsistent problems than for consistent problems (66% vs. 3%; Wilcoxon test, $z = 5.70, p < .001, Cliff's \delta = 0.84$), and 39 out of 45 participants exhibited this pattern (binomial test, $p < .001$). The pattern is analogous to how individuals cope with conflicts in causal sets of premises (Khemlani & Johnson-Laird, 2011), that is, they produce explanations in light of conflicts. Likewise, participants produced direct refutations more often for inconsistent versus consistent problems (22% vs. 4%; Wilcoxon test, $z = 3.83, p < .001, Cliff's \delta = 0.34$), though 27 out of 45 participants produced no direct refutations whatsoever. Seventeen of the remaining 19 produced more refutations for inconsistent versus consistent problems (binomial test, $p < .001$, prior probability of 1/2). These overall patterns were robust to the different materials, as well; when aggregated by the seven different scenarios, all seven revealed more temporal explanations and more refutations for inconsistent problems than consistent problems (binomial tests, $ps < .01$).

Experiment 1 revealed that individuals could spontaneously construct temporal explanations. They did so on about a third of

Table 1

The Types of Responses Produced by Participants in Experiment 1; the Percentages of Those Responses; Examples of Each Response Type; and the Percentages of Responses That Unambiguously Refuted or Explained One of the Three Premises, Along With Relevant Examples

Type of response	%	Example
Temporal explanations	32%	“The staff meeting was postponed.”
Direct refutation of premises	13%	“Kiana did not pick up her medication.”
Other explanations	5%	
Causal	<2%	“Kiana broke into the pharmacy after it closed.”
Epistemic	<2%	“The hours given for the doctor’s office were inaccurate.”
Spatial	<2%	“Ria attended the meeting remotely.”
Miscellaneous	<1%	[omitted for brevity]
Premise that was either refuted or explained		
Premise 1	33%	“The blood drive location decided to stay open later.”
Premise 2	4%	“Ria got to work early before ... the meeting.”
Premise 3	10%	“Ria did not make the meeting, she was too late.”

Note. Percentages are pooled over consistent and inconsistent problems, and many participants produced miscellaneous, that is, non-explanatory and non-refutational responses to consistent problems. Therefore, percentages do not add to 100%.

inconsistent trials. A minority of the participants' responses (22%) directly refuted one of the premises on inconsistent trials, for example, "The pottery class started and ended at a different time than what was planned." Such responses seldom occur in the case of causal conflicts (Khemlani & Johnson-Laird, 2011, Experiment 1). The result is notable because the ability to type open-ended responses could have allowed some participants to immediately infer the consequences of the explanations they generated. Hence, open-ended responses obscured the participants' preferred strategies for coping with temporal conflicts. Experiment 2 accordingly used a forced choice task to directly compare participants' preferences between the two.

Experiment 2

Experiment 2 tested whether participants prefer temporal explanations to direct refutations when coping with premises that describe a temporal inconsistency. In half of the trials, participants received problems and response options such as:

Suppose that you are told the following:

The concert was occurring from 9 pm to 11 pm.
 Ruthie arrived at the concert hall at 11:30 pm.
 You discover the following fact:
 Ruthie attended the concert.
 What, if anything, follows?
 The concert was delayed by two hours. [explanation]
 Ruthie did not attend the concert. [refutation]
 Nothing follows from the given information.

In the remaining trials, participants compared the same response options for consistent problems, for example, a problem akin to the one above except where the concert ended at midnight after Ruthie arrived. Refutations are simpler than explanations because explanations introduce concepts not present in the premises, for example, the explanation above introduces the temporal concept of a "delay" while the refutation merely negates what was expressed in a previous premise. Hence, a conservative response may be to prefer refutations over explanations. But, as previous research on causal explanations shows, people often prefer explanations to refutations because explanations provide a more complete narrative of what gave rise to the inconsistency.

Method

Participants

Fifty-five participants completed the experiment for monetary compensation (\$2.50) through Amazon's Mechanical Turk, commensurate with minimum-wage standards. We dropped data from five participants who took less than 2 minutes to complete the task, gave nonsensical responses to the debriefing questions, or were non-native English speakers. Of the remaining 50 participants, all but 12 had taken one or fewer courses in introductory logic. Their mean age was 39.1; 17 participants were female, 32 were male, and one preferred not to say.

Task and Design

As in Experiment 1, participants were presented with eight problems in randomized order, half of which were consistent and half of which were inconsistent. Each problem presented a forced choice

task between three possible responses to the prompt, "What, if anything, follows?"

Materials

The eight scenarios were based on participants' natural responses from Experiment 1. Some of the problems in Experiment 1 contained various ambiguities that permitted participants to construe the problems as consistent when they were designed to be inconsistent. The materials in Experiment 2 were modified so that they described unambiguous event intervals. For each scenario, the experiment provided three response options: a temporal explanation, a refutation, and "Nothing follows from the given information." The temporal explanation implied a change to the interval described in the first premise, for example, an extension or a postponement. The refutation focused on the third premise by denying that the agent attended the event or that they carried out the action that required attending the event, for example, "Yasmine did not speak at the meeting." The explanations and refutations were constructed to have the same number of syllables (see Appendix A). There was an issue with the same scenario as in Experiment 1. The analysis we report was conducted on the other seven scenarios; excluding the erroneous scenario did not qualitatively affect the results.

Procedure

Each problem began by displaying the event information and the question. After a 3 s delay, the three response options appeared in a randomized order. Participants selected a response to move to the next problem.

Results and Discussion

Table 2 provides the percentages of participants' choices for explanations, refutations, or "nothing follows" responses in Experiment 2. Across the study, participants preferred temporal explanations over refutations (51% vs. 7%; Wilcoxon test, $z = 5.76$, $p < .001$, Cliff's $\delta = 0.89$) and "nothing follows" responses (51% vs. 41%; Wilcoxon test, $z = 2.29$, $p = .022$, Cliff's $\delta = 0.33$). Their pattern of responses depended on the consistency of the scenario, $\chi^2(2, N = 350) = 155.07$, $p < .001$, namely, that they preferred explanations and refutations more often for inconsistent problems, but "nothing follows" responses for consistent problems. Planned comparisons revealed this pattern: participants preferred explanations over refutations in the inconsistent condition (82% vs. 10%; Wilcoxon test, $z = 5.49$, $p < .001$, Cliff's $\delta = 0.85$), and they chose explanations more frequently in the inconsistent condition than the consistent condition (82% vs. 23%; Wilcoxon test, $z = 8.89$, $p < .001$, Cliff's $\delta = 0.79$). In the consistent condition, participants preferred "nothing

Table 2

The Percentages of Participants' Selections of the Three Different Response Options in Experiment 2 as a Function of Whether the Problem Was Consistent or Inconsistent

Type	Consistent	Inconsistent	All
Explanation	22%	82%	51%
Refutation	5%	10%	7%
Nothing follows	73%	8%	41%

follows” responses over the other response options (73% vs. 27%; Wilcoxon test, $z = 4.45$, $p < .001$, Cliff’s $\delta = 0.71$).

Experiment 2 directly tested participants’ preferences for refutations, explanations, and “nothing follows” responses as answers to consistent and inconsistent problems. One limitation of the study is that it confounded the type of response with the premise under consideration. That is, because explanations concerned delays or postponements of the event under the description, they served to revise the first premise (the premise describing the time interval of the relevant event), whereas refutations were explicit denials of the third premise (the premise describing the agent’s participation in the event). Experiment 3 addressed the confound by presenting explanations and refutations that both concerned the first premise.

Experiment 3

Experiment 3 tested participants’ preferences among explanations, refutations, and “nothing follows” responses for consistent and inconsistent scenarios when both explanations and refutations led to revisions of the same information, that is, the event interval. It was similar to Experiment 2: it provided participants with three response options to consider, that is, an explanation, a refutation, and a “nothing follows” response. The explanations implicitly refuted the first premise, and the refutations did so explicitly. For example, one problem in the study included the following premises:

Suppose that you are told the following:

The pottery class was Thursday from 6:30 pm to 8 pm.

Matteo arrived at the pottery studio at 9 pm on Thursday.

You discover the following fact:

Matteo attended the pottery class.

The response options were:

The teacher was late, delaying the start of the class. [explanation]

The pottery class was not on Thursday from 6:30 pm to 8 pm. [refutation]

Nothing follows from the given information.

Participants had to choose the option that followed the information provided. A preference for simplicity would favor refutations over explanations; a preference for explanatory completeness predicts the opposite effect.

Method

Participants

54 participants completed the experiment for monetary compensation (\$2.50) through Amazon’s Mechanical Turk. We dropped the data from 10 participants who took less than 2 min to complete the experiment. The analyses reported below are based on the remaining 44 participants (21 female, 33 male, $M_{age} = 36.8$). All of the participants were native English speakers and 30 had taken one or fewer courses in introductory logic.

Design, Procedure, and Materials

The task, design, and procedure were similar to Experiment 2 except in two respects. The first is that the refutations directly negated premises describing the interval of the event, for example, “The pottery class was not on Thursday from 6:30 pm to 8 pm.” As a

consequence, refutations tended to be longer than explanations. The second is that, unlike in the two previous experiments, the first premise did not vary. The experiment instead manipulated the arrival time described in the second premise to create consistent and inconsistent problems, for example, “Matteo arrived at the pottery studio at [7 pm/9 pm] on Thursday.” Each participant carried out eight problems—four consistent and four inconsistent—in a different random order.

Results and Discussion

Table 3 provides the percentages of participants’ choices for explanations, refutations and “nothing follows” responses in Experiment 3. As in Experiment 2, participants preferred explanations over refutations or over “nothing follows” responses (49% vs. 17% and 34%, respectively; Wilcoxon tests, $z_s > 2.16$, $ps > .030$, Cliff’s $\delta_s > 0.32$). Their response preferences likewise varied as a function of the consistency of the scenario ($\chi^2(2, N = 352) = 71.64$, $p < .001$). For consistent problems, participants preferred “nothing follows” responses over both explanations and refutations (56% vs. 31% and 13%; Wilcoxon tests, $z_s > 2.23$, $ps < .026$, Cliff’s $\delta_s > 0.35$); for inconsistent problems, they preferred explanations over both refutations and “nothing follows” (37% vs. 13%, Wilcoxon test, $z = 3.44$, $p < .001$, Cliff’s $\delta = 0.48$).

Taken together, Experiments 1–3 reveal new patterns of reasoning about temporal explanations. They show that participants generate temporal explanations (Experiment 1) and that they can prefer temporal explanations over refutations to resolve inconsistencies (Experiments 2 and 3). Experiment 4 sought to test a consequence of these patterns: people should be subject to a novel form of a conjunction fallacy (Tversky & Kahneman, 1983) when reasoning about temporal events. Each explanation in the preceding study, for example, “The teacher was late, delaying the start of class” logically implied the refutation to which it was paired—and so reasoners who estimate the probability of an explanation to be higher than that of a refutation, commit a conjunction fallacy about temporal relations.

Experiment 4

Reasoners can generate explanations that they only partially believe in. For instance, suppose you are not quite sure how to explain Matteo’s participation in pottery class, so you assume with the uncertainty that there was a delay at the start of the class. We can use probabilities to describe that uncertainty, for example, by treating the explanation’s probability as falling somewhere between 0.0 and 1.0: $0 < P(\text{delay}) < 1.0$. But uncertain explanations can have certain consequences. For instance, the delay explanation implies with certainty that the class did not occur at the prespecified time interval. In general, if a statement *A* implies another statement *B*

Table 3

The Percentages of Participants’ Selections of the Three Different Response Options in Experiment 3 as a Function of Whether the Problem Was Consistent or Inconsistent

Type	Consistent	Inconsistent	All
Explanation	31%	66%	49%
Refutation	13%	22%	17%
Nothing follows	56%	12%	34%

532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590

—that is, if $P(A) = P(A \& B)$ —then A cannot be more probable than B , because $P(A \& B) \leq P(B)$. For example, if rainy weather implies bad weather, then $P(\text{rain}) \leq P(\text{bad weather})$, because there are other forms of bad weather (e.g., windy and snowy weather). So, the probability of the explanation, $P(\text{delay})$, must be less than the probability that the class did not occur at the original interval, $P(\text{did not occur at interval})$. In that event, people should judge the probability of the explanation less than the probability of the refutation. Experiment 4 sought to test whether participants instead judge $P(\text{delay}) > P(\text{did not occur at interval})$. Participants who judge the probability of an explanation higher than its corresponding refutation commit a conjunction fallacy, that is, they provide a response equivalent to $P(A \& \text{not-}B) > P(\text{not-}B)$.

Method

Participants

The study asked participants to render judgments via a Likert scale, and it tested a prediction that could be corroborated spuriously given the response options granted by the scale. As a conservative approach, we sought medium-effects ($d = 0.4$) at .001 α error probability, which necessitated the collection of a larger sample ($N = 103$). One hundred and one participants (39 female, 62 male, $M_{\text{age}} = 39.8$) completed the experiment for monetary compensation (\$2.50) through Amazon's Mechanical Turk. All of the participants were native English speakers and 90 had taken one or fewer courses in introductory logic.

Design, Procedure, and Materials

Experiment 4 presented participants with eight inconsistent descriptions of individuals attending events, such as:

Suppose that you are told the following:

1. The doctor's office was open from 10 am to 4 pm.
2. Josephine arrived for her doctor's appointment at 6 pm.

You discover the following fact:

Josephine was seen by the doctor.

On half of the problems, participants rated the likelihood (on a 7-point Likert scale from $-3 = \text{very unlikely}$ to $+3 = \text{very likely}$) of a refutation, for example,

How likely is it that the doctor's office was not open from 10 am to 4 pm?

On the remaining problems, participants rated the likelihood of an explanation, for example,

How likely is it that the doctor stayed late to see Josephine?

Both the refutations and the explanations negated some element of the first premise. For each participant, the experiment randomly assigned half the problems to concern explanations and half to concern refutations. Likewise, each participant carried out the eight problems in a different random order.

Results and Discussion

Participants in Experiment 4 rated temporal explanations ($M = 1.71$) as more likely than mere refutations ($M = 0.07$; Wilcoxon test, $z = 7.56$, $p < .001$, Cliff's $\delta = 0.62$). And 83 out of

101 participants' mean evaluations exhibited the pattern (binomial test, $p < .001$, given a prior probability of $1/2$). To assess whether the overall preference for temporal relations was robust, we subjected the data to a generalized linear mixed model analysis that used a maximal random effects structure, that is, one that took into account random effects of intercepts and slopes as contributed by participants and items. The results revealed a robust effect of the response type, $B = -1.63$, $SE = 0.31$, $p < .001$. In sum, participants exhibited temporal conjunction fallacies in their preferences for explanations over refutations.

General Discussion

We describe evidence that reasoners can spontaneously generate temporal explanations—that is, explanations that introduce novel events and temporal relations—particularly when those explanations resolve conflicts in premises that describe temporal relations. A series of experiments presented participants with problems of the following form:

The party occurred from 7 pm to 10 pm on Friday.
Maryam arrived at the party at 10:30 pm on Friday.
Maryam attended the party.
What, if anything, follows?

The premises explicitly concern temporal information, that is, the durations of the events and a particular individual's arrival time, and people have no difficulty assessing the consistency of such descriptions (though they have difficulty when durational relations yield ambiguous mental simulations; see Kelly et al., 2020). For instance, the premises above are clearly inconsistent: they cannot all be true at the same time.

Participants in Experiment 1 typed out their natural responses to such problems. In theory, their responses need not have appealed to temporal explanations: the premises imply other relations. For instance, if Maryam arrived at the party, it may be reasonable to induce that she knew that the party was happening (an epistemic relation). Maryam arrived at the party, which means that she had not been in the same spatial location as the party (a spatial relation). Arrival at the party demands some means of moving from one location to another (a causal relation). And so, in theory, participants could have appealed to any number of explanations to resolve the conflict, such as this epistemic explanation:

Maryam *mistakenly thought* the party was happening until midnight.

Instead, reasoners preferred to resolve the conflict by introducing novel temporal relations, as in this explanation:

Participant 43: "The party ran later than scheduled."

The relation *later* is temporal, and it helps to resolve the conflict by implicitly refuting the first premise in the description. Other temporal explanations are possible, for example,

Maryam went to the after-party, not the main party.

This explanation introduces a novel event (the *after-party*), which presumably occurs directly after the main party, and so it does not refute the first premise but rather the second.

Perhaps participants' tendency to generate temporal explanations was an artifact of the generative task in Experiment 1, that is, there

may be a cognitive burden associated with constructing an explanation from scratch (Horne et al., 2019). Experiments 2 and 3 accordingly provided participants with several options to resolve inconsistencies, including an explanation and a direct refutation. In Experiment 2, participants chose the explanation far more often than they chose the refutation, though the experiment was confounded such that the two options implicitly refuted different premises. Experiment 3 addressed the confound by providing participants with explanations and refutations that both concerned revisions to the same premise, and it too showed that participants preferred explanations over revision. Experiment 4 revealed a consequence of such a preference: individuals' overall ratings for explanations were more probable than those for refutations, which is a form of a conjunction fallacy (Tversky & Kahneman, 1983): explanations stipulate more information than refutations, so they should be rated less probable.

One reviewer raised a limitation of the current studies: temporal explanations, including all of the ones in the present experiments, appear to appeal to background-enabling conditions along with temporal concepts. For instance, reasoners who extended the time of the party may have done so, not to create a temporal relation, but rather to create an enabling condition—a causal relation in which one event makes an outcome possible when it otherwise might not be (see, e.g., Johnson-Laird & Khemlani, 2017). Perhaps the true explanatory force comes from the enabling conditions that the explanations describe, and not their temporal nature. But, enabling conditions rarely serve as cogent explanations. Consider the following explanations for why Abraham Lincoln was assassinated:

Booth shot him.	[cause]
Bullets were loaded in Booth's gun.	[enabler]
The safety on Booth's gun was off.	[enabler]
Booth was in the same theater as Lincoln.	[enabler]
Booth was furious over Lee's surrender.	[enabler]
...	

The first explanation is causal in nature—it describes a direct causal link between Booth's actions and Lincoln's assassination. All the other explanations describe enabling conditions, and none of them have much explanatory force: bullets may have been in Booth's gun the entire night, but common sense dictates that Lincoln was assassinated only after Booth pulled the trigger. The reason enabling conditions make for poor explanations is because, for any given phenomenon, there are an infinite number of enablers: Booth was in the same country as Lincoln; Booth was able to acquire a gun; Booth had the physical strength to pull the trigger; and so on. Reasoners are unlikely to appeal to them, generate them spontaneously, or even bring these conditions into conscious consideration unless explicitly prompted. The events described in the four studies are doubtless enablers—but if so, then their enabler-status would suggest that they should make for poor explanations. The participants instead concluded that the explanations were relatively plausible, and we argue that it was something other than their status as enablers that governed this plausibility. Their temporality made them persuasive: the explanations introduced some temporal concept or relation.

In sum, we report the discovery of a novel class of non-causal explanation: temporal explanations. Temporal explanations appeal to the relations between one or more events, for example, the duration of the event. Reasoners generate and evaluate them systematically, and they prefer explanations to more minimal refutations

(Khemlani & Johnson-Laird, 2011). Temporal explanations may help explain a trade-off in resolving temporal conflicts: while an explanation demands cognitive resources to produce and evaluate, it yields a better and more coherent understanding of how a particular conflict came about.

Context of the Research

The present research came from our recent explorations of how people comprehend discourse about time (Kelly et al., 2020; Kelly & Khemlani, 2020). We investigated how reasoners make errors and exhibit biases based on efficient ways in which they represent temporal concepts: they tend to build a single mental model of a set of temporal relations and draw conclusions from that single model. Explanations are akin to minimal, singular models of discourse (see Korman & Khemlani, 2020). So, when reasoners resolve conflicts by constructing explanations, they may do so out of the need to construct a single explanatory mental model—a simulated possibility linking antecedent events to the outcomes worth explaining—to yield a coherent potential narrative, even if that explanatory model forces them to tacitly abandon previous information (Khemlani et al., 2018). The present research sheds light on the powerful ways people process and reason about temporal concepts such as precedence and duration.

References

- Ahn, W., & Kalish, C. W. (2000). The role of mechanism beliefs in causal reasoning. In F. C. Keil & R. A. Wilson (Eds.), *Explanation and cognition*. MIT Press.
- Alicke, M. D., Mandel, D. R., Hilton, D. J., Gerstenberg, T., & Lagnado, D. A. (2015). Causal conceptions in social explanation and moral evaluation: A historical tour. *Perspectives on Psychological Science*, 10(6), 790–812. <https://doi.org/10.1177/1745691615601888>
- Anderson, C. A., Lepper, M. R., & Ross, L. (1980). Perseverance of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology*, 39(6), 1037–1049. <https://doi.org/10.1037/h0077720>
- Craik, K. (1943). *The nature of explanation*. Cambridge University Press.
- Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause. *Psychological Bulletin*, 99(1), 3–19. <https://doi.org/10.1037/0033-2909.99.1.3>
- Fembach, P. M., Macris, D. M., & Sobel, D. M. (2012). Which one made it go? The emergence of diagnostic reasoning in preschoolers. *Cognitive Development*, 27(1), 39–53. <https://doi.org/10.1016/j.cogdev.2011.10.002>
- Gopnik, A. (2000). Explanation as orgasm and the drive for causal knowledge: The function, evolution, and phenomenology of the theory-formation system. In F. Keil & R. A. Wilson (Eds.), *Explanation and cognition*. MIT Press.
- Gronchi, G., Zemla, J., & Brondi, M. (2017). *Cognitive style predicts magical beliefs*. Proceedings of the 39th annual conference of the Cognitive Science Society.
- Harman, G. H. (1965). The inference to the best explanation. *Philosophical Review*, 74(1), 88–95. <https://doi.org/10.2307/2183532>
- Hoerl, C., & McCormack, T. (2019). Thinking in and about time: A dual systems perspective on temporal cognition. *Behavioral and Brain Sciences*, 42, Article e244. <https://doi.org/10.1017/S0140525X18002157>
- Horne, Z., Muradoglu, M., & Cimpian, A. (2019). Explanation as a cognitive process. *Trends in Cognitive Sciences*, 23(3), 187–199. <https://doi.org/10.1016/j.tics.2018.12.004>
- Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Harvard University Press.

- 827 Johnson-Laird, P. N., Girotto, V., & Legrenzi, P. (2004). Reasoning from
828 inconsistency to consistency. *Psychological Review*, *111*(3), 640–661.
829 <https://doi.org/10.1037/0033-295X.111.3.640>
- 830 Johnson-Laird, P. N., & Khemlani, S. (2017). Mental models and causation.
831 In M. Waldmann (Ed.), *Oxford handbook of causal reasoning* (pp. 1–42).
832 Elsevier, Academic Press.
- 833 Kelly, L., & Khemlani, S. (2020). Directional biases in durative inference. In
834 B. Armstrong, S. Denison, M. Mack, & Y. Xu (Eds.), *Proceedings of the*
835 *42nd annual conference of the Cognitive Science Society*, Austin, TX.
836 Cognitive Science Society.
- 837 Kelly, L., & Khemlani, S. (2021). Temporal explanations help resolve tempo-
838 ral conflicts. In T. Fitch, C. Lamm, H. Leder, & K. Tessmar (Eds.),
839 *Proceedings of the 43rd annual conference of the Cognitive Science*
840 *Society*, Austin, TX. Cognitive Science Society.
- 841 Kelly, L., Khemlani, S., & Johnson-Laird, P. N. (2020). Reasoning about
842 durations. *Journal of Cognitive Neuroscience*, *32*(11), 2103–2116.
843 https://doi.org/10.1162/jocn_a_01621
- 844 Khemlani, S., Byrne, R. M. J., & Johnson-Laird, P. N. (2018). Facts and pos-
845 sibilities: A model-based theory of sentential reasoning. *Cognitive*
846 *Science*, *42*(6), 1887–1924. <https://doi.org/10.1111/cogs.12634>
- 847 Khemlani, S., & Johnson-Laird, P. N. (2011). The need to explain. *Quarterly*
848 *Journal of Experimental Psychology*, *64*(11), 2276–2288. <https://doi.org/10.1080/17470218.2011.592593>
- 849 Khemlani, S., & Johnson-Laird, P. N. (2012). Hidden conflicts: Explanations
850 make inconsistencies harder to detect. *Acta Psychologica*, *139*(3), 486–
851 491. <https://doi.org/10.1016/j.actpsy.2012.01.010>
- 852 Khemlani, S., & Johnson-Laird, P. N. (2020). Causal conflicts produce dom-
853 ino effects. *Quarterly Journal of Experimental Psychology*, *73*(12), 2317–
854 2327. <https://doi.org/10.1177/1747021820958416>
- 855 Korman, J., & Khemlani, S. (2020). Explanatory completeness. *Acta*
856 *Psychologica*, *209*, Article 103139. <https://doi.org/10.1016/j.actpsy.2020.103139>
- 857 Legare, C. H. (2012). Exploring explanation: Explaining inconsistent infor-
858 mation guides hypothesis-testing behavior in young children. *Child*
859 *Development*, *83*(1), 173–185. <https://doi.org/10.1111/j.1467-8624.2011.01691.x>
- 860 Legrenzi, P., & Johnson-Laird, P. N. (2005). The evaluation of diagnostic
861 explanations for inconsistencies. *Psychologica Belgica*, *45*(1), 19–28.
862 <https://doi.org/10.5334/pb-45-1-19>
- 863 Lombrozo, T. (2007). Simplicity and probability in causal explanations.
864 *Cognitive Psychology*, *55*(3), 232–257. <https://doi.org/10.1016/j.cogpsych.2006.09.006>
- 865 Lombrozo, T. (2016). Explanatory preferences shape learning and inference.
866 *Trends in Cognitive Sciences*, *20*(10), 748–759. <https://doi.org/10.1016/j.tics.2016.08.001>
- 867 Lombrozo, T., & Carey, S. (2006). Functional explanation and the function of
868 explanation. *Cognition*, *99*(2), 167–204. <https://doi.org/10.1016/j.cognition.2004.12.009>
- 869 Ross, L., Lepper, M. R., Strack, F., & Steinmetz, J. (1977). Social explanation
870 and social expectation: Effects of real and hypothetical explanations on
871 subjective likelihood. *Journal of Personality and Social Psychology*,
872 *35*(11), 817–829. <https://doi.org/10.1037/0022-3514.35.11.817>
- 873 Sloman, S. A. (2005). *Causal models: How we think about the world and its*
874 *alternatives*. Oxford University Press.
- 875 Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning:
876 The conjunction fallacy in probability judgment. *Psychological Review*,
877 *90*(4), 293–315. <https://doi.org/10.1037/0033-295X.90.4.293>
- 878 Weisberg, D. S., Keil, F. C., Goodstein, J., Rawson, E., & Gray, J. R. (2008).
879 The seductive allure of neuroscience explanations. *Journal of Cognitive*
880 *Neuroscience*, *20*(3), 470–477. <https://doi.org/10.1162/jocn.2008.20040>
- 881 Zemla, J. C., Sloman, S., Bechlivanidis, C., & Lagnado, D. A. (2017).
882 Evaluating everyday explanations. *Psychonomic Bulletin & Review*,
883 *24*(5), 1488–1500. <https://doi.org/10.3758/s13423-017-1258-z>

Received March 30, 2022

Revision received October 11, 2022

Accepted October 31, 2022 ■