



Transitive and pseudo-transitive inferences[☆]

Geoffrey P. Goodwin^{*}, P.N. Johnson-Laird

Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544, United States

Received 28 May 2007; revised 11 February 2008; accepted 17 February 2008

Abstract

Given that A is longer than B, and that B is longer than C, even 5-year-old children can infer that A is longer than C. Theories of reasoning based on formal rules of inference invoke simple axioms (“meaning postulates”) to capture such transitive inferences. An alternative theory proposes instead that reasoners construct mental models of the situation described by the premises in order to draw such inferences. An unexpected consequence of the model theory is that if adult reasoners construct simple models of typical situations, then they should infer transitive relations where, in certain cases, none exists. We report four studies corroborating the occurrence of these “pseudo-transitive” fallacies. Experiment 1 established that individuals’ diagrams of certain non-transitive relations yield transitive conclusions. Experiment 2 showed that these premises also give rise to fallacious transitive inferences. Experiment 3 established that when the context suggested alternatives to the simple models, the participants made fewer errors. Experiment 4 showed that tense is an important aspect of meaning which affects whether individuals draw transitive conclusions. We discuss the implications of these results for various theories of reasoning. © 2008 Elsevier B.V. All rights reserved.

Keywords: Transitive inference; Reasoning illusions; Mental models; Mental representation

[☆] This research was supported by a grant from the National Science Foundation (BCS 0076287) to investigate strategies in reasoning. For their helpful advice we thank: Ruth Byrne, Sam Glucksberg, Adele Goldberg, Uri Hasson, Cathy Haught, Georg Jahn, Sangeet Khemlani, Louis Lee, Stephen Newstead, Mike Oaksford, Sanna Reynolds, and two anonymous reviewers. Portions of this research were presented at the Fifth International Conference on Thinking, Leuven, Belgium, July 2004.

^{*} Corresponding author. Tel.: +1 609 258 9498; fax: +1 609 258 1113.

E-mail address: ggoodwin@princeton.edu (G.P. Goodwin).

1. Introduction

The ability to make transitive inferences is fundamental to human reasoning. The inference in the Abstract yielding the conclusion that A is longer than C is obviously valid, i.e., if the premises are true then the conclusion must be true too. Its validity relies on the transitivity of the relation, *longer than*. Transitivity is a logical property of some but not all relations. Some other relations are intransitive. For instance, from the fact that Bill is the father of Tom, and Tom is the father of Mike, it follows validly that Bill is *not* the father of Mike. And still other relations are neither transitive nor intransitive, but non-transitive, because they yield neither sort of valid inference. For instance, from the fact that Bob loves Anne, and Anne loves James, it follows neither that Bob loves James nor that Bob does not love James.

Transitivity is one of a set of logical properties that relations can possess (see e.g., Herrick, 1994). Relations may also be symmetric, so that if an entity, A, is related to another entity, B, then B stands in the same relation to A (e.g., *cousin of*), and some relations may also be reflexive, so that any entity stands in that relation to itself (e.g., *equal to*). They can also be differentiated in terms of their so-called ‘relational complexity’ – that is, in terms of the number of arguments that they take. For instance, the relation *longer than* takes two arguments and is therefore a binary relation, whereas the relation *between* takes three arguments and is therefore a ternary relation (see e.g., Goodwin & Johnson-Laird, 2005; Halford, Wilson, & Phillips, 1998). Relations also have specific meanings that in turn can yield further sorts of inference (see e.g., Chaffin & Herrmann, 1984).

Children as young as 5 years old are able to draw transitive inferences (Bryant & Trabasso, 1971). They acquire the knowledge that certain relations are transitive, and learn to reason accordingly. At around the age of 8, children sometimes over-generalize, and draw transitive conclusions for non-transitive relations such as *loves* and *kicks* (Kuczaj & Donaldson, 1982). However, this tendency largely disappears by the age of 10. Since Bryant and Trabasso’s (1971) demonstration, there has been controversy about the age at which children master transitive reasoning (see Andrews & Halford, 1998, 2002), but they undoubtedly do master certain transitive inferences sometime during childhood. Perhaps for this reason, the boundaries of adult competence in transitive reasoning have largely been unexplored (cf. De Soto & Kueth, 1958). Theories have been constructed to explain reasoning with transitive relations. They have focused on ‘three-term series’ problems, which have two premises interrelating three entities, and on the latencies of correct responses, because accuracy is so high (see e.g., Clark, 1969; De Soto, London, & Handel, 1965; Hummel & Holyoak, 2001; Hunter, 1957; Huttenlocher, 1968).

In contrast, our experiments examine both transitive and non-transitive relations. Their aim is to examine the consequences for relational inferences of a theory of human reasoning, and to answer two questions: How well do individuals discriminate between relations that validly yield transitive inferences and those that do not? How, if at all, do they mentally represent the property of transitivity? The answers to these questions should shed light on how reasoners are able to make transitive inferences.

A valid inference is easily defined in logic: “A valid inference is one whose conclusion is true in every case in which all its premises are true” (Jeffrey, 1981, p. 1). In other words, a valid inference admits no counterexample in which its premises are true but its conclusion is false (Beth, 1971, p. 10). One goal in formal logic is to capture valid inferences in a purely formal way. The goal is simple to achieve in logic, because logic deals with the implications of *sentences* in a formal language. But, it is difficult to achieve in everyday life, because validity concerns the *propositions* expressed by sentences in natural language as they are used in particular contexts – it is propositions that are true or false, not sentences. And a given sentence can express different propositions depending on its context (Jeffrey, 1981, p. 14). Unfortunately, there is no algorithm as yet for recovering the logical form of all the propositions expressible in natural language, and many modern logicians confronted with natural language have therefore found the notion of logical form unilluminating. As the logician, the late Jon Barwise (1989, p. 4) wrote: “Within the model-theoretic tradition [in logic], valid entailments are valid not in virtue of form, but in virtue of content”. Likewise, in his controversy with Fodor (1987), he proposed that everyday reasoning was not a formal process (Barwise, 1989, e.g., p. 159).

A long-standing debate within cognitive science similarly concerns whether human reasoning is based on form or content, and our study of transitive inferences addresses this debate. Several psychological theories of reasoning attempt to capture validity by way of formal rules of inference that are sensitive only to the logical form of the premises (e.g., Braine & O’Brien, 1991, 1998; Rips, 1994). Consider, for example, this transitive inference:

Ann is taller than Beth.
Beth is taller than Cath.
Therefore, Ann is taller than Cath.

Theories based on formal rules have to invoke an axiom (a so-called “meaning postulate”) in order to allow a formal proof that the conclusion follows from the premises. This axiom stipulates that the relation, *is taller than*, is transitive:

For any x, y, z , if x is taller than y , and y is taller than z , then x is taller than z . In the proof, the value of x is set to Ann, the value of y is set to Beth, and the value of z is set to Cath. The rest of the proof is easy. The premises match the antecedent conjunction in the axiom of transitivity, and so the conclusion follows in a single step. Errors may occur in the application of the formal rules, but they should be haphazard rather than systematic.

An alternative theory postulates instead that reasoning depends on content rather than form. According to this theory, reasoners use the meaning of the premises and any pertinent knowledge to envisage what is possible, and they represent each possibility in a mental model (e.g., Johnson-Laird & Byrne, 1991). They grasp the meaning of the statement that Ann is taller than Beth, and use it to construct a model in which Ann is represented as taller than Beth. They grasp that ‘Beth’ in the second premise is co-referential with an individual that is already represented in their model of the first premise, and so they can extend this model to represent that Beth is taller than Cath. The resulting model of the relative heights of Ann, Beth, and Cath yields, as an emergent

property, the transitive relation that Ann is taller than Cath. No alternative model of the premises is a counterexample to this conclusion, and so it is valid. Inferences about spatial and temporal relations can be drawn in this way, and the theory has been implemented in computer programs (see Johnson-Laird & Byrne, 1991, chap. 9).

According to the model theory, reasoners should keep the number of possibilities that they represent to a minimum and try to avoid constructing more than one model (Johnson-Laird, Legrenzi, Girotto, & Legrenzi, 2000; Ormerod & Richardson, 2003). In this way, they can prevent overloading the processing capacity of working memory (Baddeley, 1986). Individuals should also tend to build as simple models as possible, and they may overlook alternatives to them. The theory thus yields a novel prediction: individuals should be susceptible to systematic fallacies with transitive relations, despite their apparent simplicity. Fallacies of transitivity should arise from simple models of typical situations and from the failure to consider alternative models of the premises. The relations that elicit these fallacies should be those that have two crucial properties. First, they should be non-transitive in that they yield validly neither transitive nor intransitive conclusions. Second, they should elicit simple models of typical situations in which the transitive conclusion does hold. We refer to relations of this sort as *pseudo-transitive*. Consider this example:

Ann is a blood relative of Beth.
Beth is a blood relative of Chris.
What follows?

Naïve reasoners, i.e., those with no training in logic, should tend to construct a simple model of siblings, or linear descendants, and to draw the transitive conclusion that Ann is a blood relative of Chris. But, the inference is invalid, as a counterexample shows: Anne is Beth's mother, and Chris her father, but her mother and father are not blood relatives. Pseudo-transitive relations contrast with transitive relations, such as *is taller than*, which have only models that yield transitive conclusions, with relations that are intransitive, such as *father of*, which have no models that yield transitive conclusions, and with regular non-transitive relations, such as *loves*, which have typical models that do not yield transitive conclusions.

The model theory makes one further prediction. Individuals should be more likely to refrain from pseudo-transitive inferences if they are given a context that elicits alternatives to simplistic models, e.g., they are told to bear in mind the consequences of marriage on kinship. Previous studies of transitive reasoning have focused on valid inferences (e.g., Bryant & Trabasso, 1971; Clark, 1969; Huttenlocher, 1968; Knauff & Johnson-Laird, 2002). They have not systematically examined pseudo-transitive inferences, and so we carried out four experiments to investigate them.

2. Experiment 1

Experiment 1 established the existence of pseudo-transitive relations, that is, relations that are not transitive, but that yield simple models of typical situations in

which a transitive conclusion holds. To establish these relations, we presented participants with two separate tasks with a putative set of such relations. First, they were presented with pairs of premises in the form of a three-term series problem. Unlike typical three-term series problems, we selected relations for them that were non-transitive, yet that were likely to yield typical models supporting transitive conclusions. Participants were asked to draw a single diagram that was consistent with the premises, and we predicted that they would tend to draw diagrams in which a transitive conclusion holds. Second, in order to demonstrate that these relations are not always judged as transitive, participants were presented with a set of three alternative diagrams of the premises. Two of these diagrams were consistent with the premises, and one yielded a transitive conclusion, whereas the other did not. The third diagram was inconsistent with the premises. Participants had to evaluate whether or not each diagram was consistent with the premises. With pseudo-transitive relations, they should judge both of the consistent diagrams as consistent, irrespective of whether they yielded a transitive conclusion. Together, the two tasks aimed to show that individuals' spontaneous or preferred mental models of the pseudo-transitive relations yielded transitive conclusions, even though they judged diagrams in which the transitive conclusions did not hold as consistent with the premises.

2.1. Methods

2.1.1. Participants

Twenty-five undergraduate participants (20 female, 5 male) from Princeton University participated in the experiment as part of a course requirement.

Table 1

The percentages of transitive diagrams that the participants drew to represent the premises in Experiment 1, and the percentages of their judgments of consistency for the three sorts of diagrams (transitive, non-transitive, inconsistent)

Premises	Percentages of transitive diagrams drawn by the participants	Percentages of judgments of 'consistent' for the three sorts of diagram		
		Transitive	Non-transitive	Inconsistent
Fred is a blood relative of Bob's. Bob is a blood relative of Anne's.	68	100	96	4
Tim is on Mark's left. Mark is on Fred's left.	96	100; 98	100; 98	0; 9
Ken is behind John. John is behind the car.	100	100	76	0
In a particular series of events, Planet Z went round planet X. Planet X went round planet Y.	48	88	84	8
In a sports race, Alvarez overtook Underwood. Underwood overtook Henderson.	86	100	92	12

Note. The second set of consistency judgments for the relation, *on X's left*, show the judgments of a separate group of participants, who judged a corrected set of diagrams.

2.1.2. Design and materials

Five putative non-transitive relations were presented to participants, and these were: *is a blood relative of*, *is behind*, *is on X's left*, *went round*, *overtook*. The relations as they were presented to participants are shown in Table 1. We constructed a single set of terms as arguments for these relations, and they were presented to all the participants. The terms in the relations were constant across the two different tasks: drawing diagrams and judging diagrams for consistency. The figure of the premises was also constant, and was always: A–B; B–C.

The order of the two tasks was counter-balanced: half the participants received the drawing task first, and half the participants received the evaluation task first. The order of the problems was randomized for each participant, and this order was constant across the two tasks. In the judgment task, the participants saw three different diagrams for each problem: one in which the transitive conclusion held, one in which the transitive conclusion did not hold, and one that was inconsistent with the premises. Appendix A displays the three sorts of diagram for each of the five relations.¹ For each relation, given that there were three diagrams, there were six possible orders in which they could have been presented, and one of these six orders was selected at random for each participant.

2.1.3. Procedure

The participants were tested individually, and they used pencil and paper for all the problems. The instructions for the two tasks were presented on an initial sheet of paper. The key instructions for the drawing task were as follows:

You will be presented with five problems that contain two premises each. The premises will concern the relation between three different people or objects in terms of some property. Your task is to draw a diagram that represents a situation or state of affairs that is consistent with the two premises, i.e., just draw one diagram representing a single possibility consistent with the premises. Try to make each diagram as clear and informative as possible, so that someone who did not know what the premises were would be able to interpret it.

For the evaluation task, the key instructions were as follows:

You will be presented with five problems that contain two premises each. The premises will concern the relation between three different people or objects in terms of some property. Your task is to judge whether a series of diagrams that you are presented with are consistent with the premises. For each diagram, you will be asked to indicate whether it is consistent with the premises (i.e., it

¹ There are four diagrams for the relation, *on X's left*, because the atypical model that was presented in the experiment in fact (mistakenly) yielded a transitive conclusion (although it still required an intrinsic frame of reference to be judged correctly). A separate group of 43 undergraduate participants from the same population as the main study judged just the set of models for the relation, *on X's left*, including this time an atypical model in which the transitive conclusion did not hold (see Table 1).

represents a state of affairs that could be the case, given that the premises are true), or inconsistent (i.e., it represents a state of affairs that could not be the case given that the premises are true). For each of the five problems you will be asked to judge three separate diagrams. Remember that certain diagrams may be consistent with the premises, even if they were not the ones that first came to mind for you.

For both tasks, the instructions were adjusted in minor ways when the task was presented second.

2.2. Results

Table 1 presents the data from both tasks with the five relations. Participants tended to draw diagrams that yielded transitive conclusions, as shown in the second column of Table 1. Some of the diagrams were not interpretable as either yielding a transitive conclusion or not, and this ambiguity occurred for 20% of the diagrams for the *blood relative* problem, and for 12% of the diagrams for the *overtook* problem. Hence, the data in Table 1 are the percentages from the set of interpretable diagrams only. The tendency to draw a transitive diagram rather than a non-transitive one was reliable on Binomial tests (one-tailed) for most relations: *blood relative* (68% transitive, $p < .09$), *on X's left* (96% transitive, $p < .00001$), *behind* (100% transitive, $p < .00001$), *went round* (48% transitive, $p = 1.0$), *overtook* (86% transitive, $p < .001$).

The participants also tended to judge correctly that both sorts of consistent diagram, transitive and non-transitive, were indeed consistent with the premises. For all relations, the transitive diagrams tended to be judged as consistent slightly more often than the non-transitive diagrams, but this difference was not reliable except for the *behind* relation (Binomial test, two-tailed, $p < .04$). In accordance with the main prediction, the non-transitive consistent diagrams were judged as consistent with the premises reliably more often than chance for each of the five relations (Binomial tests, one-tailed, all $ps < .01$). As the table shows, the participants very rarely judged that the inconsistent diagrams were consistent with the premises.

2.3. Discussion

The results established the existence of a set of pseudo-transitive relations, that is, non-transitive relations that tend to elicit diagrams that support transitive conclusions. Yet, the participants judged them to be consistent with other diagrams that did not support transitive conclusions. Hence, the experiment demonstrates that there is a class of relations that tend to elicit transitive models, but that are in fact non-transitive, as the participants, in effect, conceded in judging non-transitive diagrams as consistent with them. In short, pseudo-transitive relations are real.

We did not examine the diagrams individuals draw for straightforwardly non-transitive relations such as *loves* or *bites*. However, a moment's reflection should convince readers that, in contrast to the pseudo-transitive relations, there is little likelihood of constructing transitive models for such relations. The difference in

the models that these relations elicit should lead to different inferences being drawn from them. Hence, our subsequent experiments investigated the sorts of conclusion that individuals draw from premises based on three sorts of relation: pseudo-transitive relations, transitive relations, and relations that are not transitive (some that are non-transitive and some that are intransitive). The theory predicts that individuals tend to rely on the first model of the premises that comes to mind. For pseudo-transitive relations, this model should typically yield a transitive conclusion, and so individuals should tend to draw invalid transitive conclusions. They may sometimes be able to refrain from such conclusions if they actively engage in a search for counterexamples. The initial models for transitive relations and for relations that are not transitive should yield transitive conclusions and conclusions that are not transitive, respectively.

3. Experiment 2

This experiment tested whether adult reasoners spontaneously draw their own transitive conclusions from premises based on pseudo-transitive relations. It examined genuine transitive relations, such as *is taller than*, relations that are not transitive, such as *loves*, and pseudo-transitive relations, such as *is a blood relative of*.

3.1. Methods

3.1.1. Participants

Twenty-four Princeton University undergraduates (13 female, 9 male, 2 unrecorded) participated either to fulfill part of a course requirement, or for payment of \$8.

3.1.2. Design and materials

Each participant received a booklet of 13 two-premise problems. Their task was to state what, if anything, followed validly from each set of premises. Each premise asserted a binary relation, and the two premises together referred to three separate individuals or entities. There were five problems based on the pseudo-transitive relations: *went faster than*, *behind*, *in front of*, *overtaken*, *blood relative of*; five problems based on the transitive relations: *heavier than*, *more expensive than*, *saltier than*, *sharper than*, *taller than*; and three problems based on relations that are not transitive: *envies*, *loves*, *father of*. The pseudo-transitive and transitive problems were each of the general form:

B relation A.
C relation B.
What follows?

But, in order to vary the figure of the premises, the order of the two premises was reversed for the problems that were not transitive. The transitive and pseudo-transi-

tive problems were presented in two blocks in counter-balanced orders, with the problems that were not transitive presented in the second (*loves*), seventh (*envies*), and ninth (*father of*) positions in the 13 trials. Otherwise, the order of the problems within each block was random for each participant.

3.1.3. Procedure

The participants were tested individually, and carried out the experiment using paper and pencil. The instructions were presented on an initial sheet of paper, and the key instructions were as follows:

Given that both premises are true, your task is to write down what follows validly (if anything) from the premises – that is, to write down what must be true given that the premises are true. If you can, try and draw a conclusion that relates people or items that are not mentioned in the same premise. If you think that nothing follows from the premises, write down ‘nothing’.

3.2. Results

Fig. 1 presents the percentages of transitive conclusions to each of the pairs of premises. The participants drew transitive conclusions from 98% of the transitive relations, from 72% of the pseudo-transitive relations, and from 8% of the relations that are not transitive (Page's *L* test, $z = 6.35$, $p < .0001$). Seventeen out of the 24 participants yielded exactly this predicted trend (Binomial test, $p < 1$ in 100 million) and of the remainder, six drew the same number of transitive conclusions from the transitive and pseudo-transitive relations, and one participant drew more transitive conclusions from the pseudo-transitive relations. Every participant drew at least one pseudo-transitive conclusion, though some of these relations yielded more transitive conclusions than others (Friedman nonparametric analysis of variance, $\chi^2 = 16.5$, $p < .005$). All of the pseudo-transitive problems yielded transitive conclusions more than half of the time, which is well above chance, given that the task was to generate a conclusion spontaneously rather than to select one from a given set. Further evidence that participants were not merely drawing conclusions by chance is their highly accurate performance with the transitive and non-transitive premises.

We suspected that reasoning from transitive problems in the first block might increase the tendency to draw pseudo-transitive inferences in the second block. But, there was no such effect, and the means for the pseudo-transitive problems showed a trend, which was not reliable, in the opposite direction.

3.3. Discussion

Individuals drew a high percentage of transitive conclusions from the pseudo-transitive relations – relations which are not transitive, but which tend nevertheless to elicit transitive models. The clearest case was with the pseudo-transitive relation *blood relative of*, which, as we have already shown, is not transitive. The other




Fig. 1. The percentages of transitive conclusions in Experiment 2. Of the 13 sets of premises, five sets were transitive: (1) *heavier than*, (2) *more expensive than*, (3) *saltier than*, (4) *sharper than*, (5) *taller than*, five sets were pseudo-transitive: (1) *went faster than*, (2) *behind*, (3) *in front of*, (4) *overtook*, (5) *blood relative of*, and three sets were not transitive: (1) *envies*, (2) *loves*, (3) *father of*. (For the relations that were not transitive, there are no columns visible in slots 3, 4, and 5 of the figure because there were no transitive conclusions drawn for *father of* (slot 3), and because there were no relations in slots 4 and 5).

pseudo-transitives are also not transitive. Consider, for instance, the following premises:

Armstrong overtook Hamilton.
Ullrich overtook Armstrong.

Our participants inferred that Ullrich overtook Hamilton. But, on the contrary, Armstrong may have led from the start, then been overtaken by both Ullrich and Hamilton riding in parallel, before finally overtaking them both. Both premises are true, but Ullrich never overtook Hamilton. It is hard to resist the transitive conclusion, because to do so calls for the construction of a complex and much less obvious model of the premises. Consider, similarly, the premises:

John is behind the car.
Steve is behind John.

It does not necessarily follow that Steve is behind the car. For instance, John could be behind the car but facing with his back to the car's trunk. Steve could be in front of the car, facing towards the windscreen, in which case he is behind John.

The relation *went faster than* is also not transitive, in part because of the use of the past tense. This usage illustrates an important point about transitive relations. If such a relation is stated in the past tense in at least one of the premises, and without specific reference to an exact time in the past, it becomes non-transitive. Past tense statements in English are often interpreted as referring to events that occurred in the order in which they are described, e.g., John entered the room, and he sat down in the chair. However, even this interpretative convention was not enough to prevent our participants from representing the premises in a simple transitive way. For instance, given the premises:

- B went faster than A.
- C went faster than B.

the participants interpreted B's speed as the same in both premises, i.e., they probably assumed that the two premises referred to the same event. This assumption makes for simpler models, and is reasonable in many cases, but it leads to error when it is overruled by tense.

Overall, the results corroborate the model theory, but they do not rule out theories based on formal rules of inference. Individuals might have erred by over-generalizing axioms of transitivity, and applying them to relations for which they were not appropriate. Experiment 3 aimed to test a further prediction of the model theory, which should prove more problematic for formal theories. It investigated whether certain contexts affect the prevalence of pseudo-transitive conclusions. According to the model theory, contexts that induce participants to build an alternative to those models that first come to mind should reduce the tendency to draw transitive conclusions from pseudo-transitive relations.

If context does matter, formal rule theories would require some additional mechanisms to take account of it. Some logicians, however, have argued that in a formal language, whether or not an inference is valid is determined from its representation alone without recourse to the context or circumstances of the inference (see e.g., Barwise, 1989, p. 158). The contextual effects of the sort predicted here would seem at the very least to require different meaning postulates to be triggered depending on the context in which relational assertions occur rather than on their logical form alone. As far as we can tell, no such mechanism has been proposed in the literature.

4. Experiment 3

This experiment also used pseudo-transitive relations, transitive relations, and relations that are not transitive. It tested the effect of contexts designed to elicit alternative models to those of typical situations. If reasoners rely on content to make

inferences and tend to draw pseudo-transitive conclusions because they think of typical situations, then a context that brings to mind atypical cases or counterexamples should inhibit pseudo-transitive inferences.

There were three sorts of pseudo-transitive relation. For a first sort, the *temporal* problems, individuals tend to assume wrongly that events happen at the same time, e.g., *went faster than*. The context designed to evoke alternative models accordingly interpolated the phrase, “and then”, between the premises. For a second sort of pseudo-transitive, the *orientation* problems, individuals tend to construct ego-centric spatial models from relations, such as *in front of*. The context accordingly stated that objects can have intrinsic parts, such as intrinsic fronts. For a third sort of pseudo-transitive, the *typicality* problems, individuals tend to construct models of typical cases, e.g., *blood relative of*. The context accordingly alerted them to less typical possibilities. The experiment examined participants’ inferences both with and without these contexts. It also examined the effects of negation, in an effort to demonstrate the generality of pseudo-transitive inferences. The negation of a transitive relation, as in *not taller than*, is transitive, but it causes difficulties in reasoning from three-term series premises (Clark, 1969). Hence, negation might decrease pseudo-transitive inferences, because a negative relation, such as is *not taller than*, has a disjunctive meaning (i.e., *equal to or shorter than*) that might elicit alternative models. However, the theory predicts that negation should not eliminate pseudo-transitive inferences, because it is still effortful to construct alternative models which invalidate the transitive conclusion.

4.1. Methods

4.1.1. Participants

Thirty Princeton University undergraduates (19 female, 11 male) participated in the experiment as part of a course requirement.

4.1.2. Design and materials

Each participant stated what, if anything, followed validly from each of 20 pairs of premises presented in the following form:

A relation B.
B relation C.
What follows?

Two blocks of 10 problems were presented. In one block, the problems were presented with a short, linguistic context as part of the problem; and in the other block, there was no such context. The contexts were designed to elicit alternative models of pseudo-transitive premises, and therefore to reduce the likelihood of transitive inferences. For the other sorts of relation, the contexts should not have any reliable effect, and we included them to control for the possibility that any preamble to the premises might make reasoners less likely to draw conclusions. Appendix B presents the 20 problems used in the experiment with the contexts designed to reduce transitive inferences from pseudo-transitive premises.

The order of the two blocks, context vs. no context, was counter-balanced. Each block contained six pseudo-transitive problems, two transitive problems, and two problems that were not transitive. One of the blocks contained the first relation in each of the following pairs, and the other block contained the second relation:

Pseudo-transitive relations:

Temporal: faster than, slower than; caught up with, overtook.

Orientation: in front of, behind; on X's left, on X's right.

Typicality: blood relative of, connected electrically to; converged on, went round.

Transitive relations: taller than, sharper than; more expensive than, saltier than.

Relations that are not transitive: loves, envies; father of, mother of.

The assignment of the relations in these pairs to block (context vs. no context), was counter-balanced – in the first assignment, all of the relations that occurred first in the pairs above were presented in the block with a context, with the relations that occurred second presented in the block without context; and in the second assignment, this was reversed. We also manipulated the presence of negation by creating two overall assignments of negation. An initial assignment was made so that half of the problems in each context block were negated at random, and the two problems in the complementary pairs were of opposite polarity. A second assignment was then made by switching the assignment of negation for all problems. Eight separate between-subjects versions of the experiment were therefore created by crossing block order (context first or second), the two assignments of relations to context, and the two assignments of negation to the relations. The order of the problems within each block was randomized for each participant.

4.1.3. Procedure

The participants were tested individually, and the problems were presented on an IBM computer, running the Eprime program. The premises were presented one at a time, and participants pressed a key to receive each premise in turn, and then the question, “what follows?”. The key instructions were the same as in the first experiment. Participants were told to “read and understand both premises fully before you proceed to the question”. They typed their response to each problem.

4.2. Results

The participants drew 91% transitive conclusions from the transitive relations, 52% from the pseudo-transitive relations, and 3% from the relations that were not transitive (Page's *L* test, $z = 7.68$, $p < .0001$). Twenty-nine out of the 30 participants yielded exactly this predicted trend (Binomial test, $p < .167^{29}$). The remaining participant drew no transitive conclusions for the problems that were not transitive, and an equal number of transitive conclusions for the genuinely transitive and pseudo-transitive problems. Every participant drew two or more pseudo-transitive conclusions. Fig. 2 presents the percentages of transitive conclusions for each of the three




Fig. 2. The percentages of transitive conclusions in Experiment 3 for transitive relations, pseudo-transitive relations, and relations that are not transitive, depending on whether or not a contextual phrase occurred. The contexts for the pseudo-transitive relations were designed to elicit atypical models that blocked transitive inferences from them; and the contexts for the other sorts of relations were fillers.

sorts of problem presented with and without the contexts designed to elicit alternative models for the pseudo-transitive problems. The contexts reliably reduced the tendency to draw pseudo-transitive conclusions (only 41% of them in comparison with 62% without the contexts; Wilcoxon test, $z = 2.9$, $p < .005$, one-tailed), but produced no reliable effects on the other sorts of relation. Their effects on pseudo-transitives depended on the type of relation. They produced less transitive inferences for the orientation problems (58% vs. 38%, Wilcoxon test, $z = 2.05$, $p < .025$, one-tailed), and for the typicality problems (49% vs. 11%, Wilcoxon test, $z = 3.33$, $p < .001$, one-tailed), but not for the temporal problems (for these problems the trend was in the reverse direction, 63% vs. 81%, Wilcoxon test, $z = 1.66$, ns).

Negative relations were less likely to yield transitive conclusions both for transitive (97% vs. 85%, Wilcoxon test, $z = 2.65$, $p < .01$, two-tailed) and pseudo-transitive relations (59% vs. 42%, Wilcoxon test, $z = 2.02$, $p < .05$), but not for the relations that were not transitive. But, as the means indicate, erroneous transitive conclusions were still drawn frequently for pseudo-transitive relations when they were negative. And, indeed, the decrease in transitive conclusions for the pseudo-transitive relations was in fact no greater than the decrease for transitive relations ($p > .6$).

We analyzed the latencies of responses from the presentation of the premises to the participants' first keystroke in typing their response. The problems that were

not transitive (27.95 s) took less time than the transitive problems (34.36 s, Wilcoxon test, $z = 2.13$, $p < .04$) and marginally less time than pseudo-transitive problems (31.19 s, Wilcoxon test, $z = 1.86$, $p < .07$, two-tailed). There was no difference between transitive and pseudo-transitive problems. In addition, there were no differences among the different sorts of relation when only the correct responses were analyzed (28.28, 32.52, and 30.91 s, respectively).

4.3. Discussion

The results replicated those of Experiment 2: the participants drew transitive conclusions from the pseudo-transitive relations. The overall percentage (52%) was somewhat lower than in the first experiment (72%), since here the problems included those with negative relations and those with contexts, which both reduced pseudo-transitive inferences.

The effects of context suggest that individuals normally construct models of the typical situations from a description, but a pertinent cue can help them to consider alternative possibilities. The context for the temporal relations, interpolated the phrase, “and then”, between the two premises, but it did not reliably reduce the proportion of pseudo-transitive inferences. It was the shortest context, and evidently insufficient to prompt consideration of the quite complicated sequence of events that refutes the pseudo-transitive conclusions (see Section 3.3). There was no reliable effect of context for the transitive relations or for the relations that were not transitive. This result is as predicted, because the contexts in these cases should not affect the interpretation of these relations. But, the absence of a reliable effect matters, because it shows that the effect of context on the pseudo-transitive premises was not merely a result of an implied instruction for the participants to “think harder” or to “be more cautious”. What mattered was the specific content: it prompted them to think of alternative models that refuted the transitive conclusions.

Negation reduced the tendency to draw transitive conclusions for both pseudo-transitive and transitive relations. For the pseudo-transitive relations, it may have prompted reasoners to think of alternative models that invalidated the transitive conclusion. However, this effect was small, and so negation is no antidote to pseudo-transitive inferences. Although no alternative models invalidate the transitive conclusions from the transitive relations, the negated premises are consistent with more models than the affirmative premises. Hence, negation may have made it more difficult for individuals to draw a conclusion that held across all models.

5. Experiment 4

Formal theories of reasoning depend on meaning postulates that capture the transitivity of relevant relations, and so a relation such as *taller than* can have its own meaning postulate or be “tagged” with a general and abstract meaning postulate that applies to all transitive relations (e.g., Bar-Hillel, 1967). Our previous experiments point to several difficulties with this view as a psychological theory, and establish

some new constraints that more nuanced formal theories need to take into account. A further constraint is shown by the effect of tense: few relations are invariably transitive. If two premises based on *taller than* are both in the present tense, as in the earlier inference about Ann and Cath, then the relation is transitive. But, if they are both in the past tense then reasoners are likely to think of counterexamples, because individuals can grow at different rates. Given that at one time Bob was taller than Mark, and that at some other time Mark was taller than Fred, no valid conclusion can be drawn about whether Bob is, or indeed was ever, taller than Fred. And if one premise is in the present tense and the other in the past tense, then the possibility of such counterexamples should be still more salient. These observations affect the viability of the “tagging” of various relations with meaning postulates – since the same relation may be transitive or not depending on the tense of the proposition in which it occurs. Formal theories must therefore posit some machinery that allows different tags depending on the tense of the premises.

Several of the pseudo-transitive relations that we investigated in the previous experiments relied on the use of the past tense for their lack of transitivity (e.g., *went round*, *went faster than*, *overtook*). In this experiment, we explicitly manipulated tense to determine whether it affects transitive inferences. We selected a set of relations that are transitive in the present tense, such as *taller than*. Our account yields three main predictions. First, there should be a trend in the proportion of transitive inferences depending on the tenses of the two premises. When they are both in the present tense, as shown in the previous experiments, they should yield a high proportion of transitive inferences. When they are both in the past tense, they suggest the possibility of counterexamples, and so the proportion of transitive inferences should decline. And when one premise is in the past tense and one premise is in the present tense, the saliency of counterexamples should be high, with a still greater decline in the proportion of transitive inferences.

A second, more subtle, prediction of the theory is that the sorts of entities in the transitive relations should affect inferences. We manipulated whether the entities were liable to change quickly or slowly with respect to the relation in question. For instance, mountains are unlikely to change quickly with respect to the relation *steeper than*, whereas the inclines of treadmills are likely to change quickly. Hence, counterexamples to a transitive conclusion from premises in the past tense based on *steeper than* should be easier to envisage when the entities in the relation are treadmills rather than mountains. An analogous prediction applies to the other transitive relations in the experiment.

The third prediction concerns the nature of the mental models that individuals are likely to construct. To envisage a counterexample for a relation that is transitive in the present tense, but not in the past tense, calls for a dynamic mental model in which the terms in the premises change over time with respect to the relation in question. Consider, for instance, the premises:

Bob was taller than Mark.
Mark was taller than Fred.

The conclusion that Bob is taller than Fred is open to a counterexample in which Fred has grown in height since the time at which the second premise held, and so Bob isn't taller than him. It takes time to construct a dynamic model of this sort. In contrast, a counterexample to a relation that is never transitive, regardless of tense, such as *knows*, does not depend on the construction of a dynamic mental model. Hence, the third prediction was that correct, "no valid conclusion," responses should take longer to make to premises containing transitive relations in the past tense than to premises containing relations that are never transitive.

5.1. Methods

5.1.1. Participants

Twenty-five undergraduate participants (19 female, 6 male) from Princeton University took part in the experiment for course credit.

5.1.2. Design and materials

Three-term series problems in the figure A–B, B–C, were created out of a set of nine relations. Of these nine relations, six were transitive in the present tense: *taller than*, *steeper than*, *bigger than*, *sharper than*, *brighter than*, *warmer than*, and three were not: *respects*, *is X's advisor*, *knows*. The relations were presented in three separate blocks. One block presented both premises in the present tense, another block presented both in the past tense, and yet another block (the *mixed* block) presented one premise in the present tense, and one in the past tense. Each block contained 15 problems: each of the six transitive relations occurred in two sorts of problem, and each of the three relations that were not transitive occurred once. The two sorts of transitive problem were based on the sorts of entities that occurred in the relation: one sort could change quickly with respect to the relation in question (e.g., treadmill inclines for *steeper than*), and the other sort could change only slowly with respect to the relation in question (e.g., mountains for *steeper than*). The premises containing the two sorts of entities were matched exactly in terms of their overall number of syllables. The three relations that were not transitive in any tense had a mean syllable length of 14.3 for each premise, which was greater than the mean syllable length of 11.7 for each premise with the other relations. This difference, which was deliberate, is an important control for a comparison that we make in Section 5.2. Table 2 presents the full set of relations and entities.

The order of the three blocks (present tense, past tense, mixed tense) was fully counter-balanced in the six possible orders. For three of these orders, the mixed tense problems were presented with the present tense premise first and the past tense premise second, whereas for the other three, they were presented in the reverse order. The order of problems within a block was random for each participant.

We constructed three overall sets of names. Each set consisted of 15 unique subsets of three names for the three entities in each problem. The three sets of names were used once for each participant – one set for each of the three tense blocks. One set of names was used consistently in the first block of problems participants

Table 2

The 15 different relations used in Experiment 4, presented in mixed tense, and the percentages of transitive conclusions drawn for each of the three blocks

Relational Premises	Past	Mixed	Present
Christina was taller than Jennifer. Jennifer is taller than Theresa.	56	28	96
John's snowman was taller than Bob's snowman. Bob's snowman is taller than Mark's snowman.	64	20	96
Collington Mountain was steeper than Jenkinsville Mountain. Jenkinsville Mountain is steeper than Richardson Mountain.	60	36	100
Bill's treadmill incline was steeper than Paul's treadmill incline. Paul's treadmill incline is steeper than Nick's treadmill incline.	63	28	92
Jack's family was bigger than Tom's family. Tom's family is bigger than Fred's family.	64	28	96
Lisa's balloon was bigger than Julie's balloon. Julie's balloon is bigger than Tanya's balloon.	64	40	92
The broad-sword was sharper than the meat knife. The meat-knife is sharper than the scalpel.	64	36	100
Anne's pencil was sharper than Jill's pencil. Jill's pencil is sharper than Dawn's pencil.	64	20	96
Planet Hera was brighter than planet Kuma. Planet Kuma is brighter than planet Lincoln.	56	36	96
The sewing room was brighter than the living room. The living room is brighter than the dining room.	72	28	96
The Green Ocean was warmer than the Dark Ocean. The Dark Ocean is warmer than the West Ocean.	60	32	92
Tracy's coffee was warmer than Ellen's green tea. Ellen's green tea is warmer than Wendy's cider.	68	40	96
<i>Professor Christensen respected Professor Zimmerman. Professor Zimmerman respects Professor Buchanan.</i>	0	4	0
<i>Senator McFarland was Senator Benjamin's advisor. Senator Benjamin is Senator Donovan's advisor.</i>	0	0	0
<i>Doctor Alvarado knew Doctor McNamara. Doctor McNamara knows Doctor Villalobos.</i>	8	4	0

Note. The relations in italics are those that are non-transitive in the present tense.

received, which allowed for a controlled between-subjects comparison of the first block of problems by tense (past, present, and mixed). The remaining two sets of names were varied across the second and third blocks participants received, so that they were presented evenly for the second and third blocks, and so that they were presented evenly across each tense block.

5.1.3. Procedure

The participants were tested individually, and the experiment was carried out on an IBM compatible computer running the Eprime program. The instructions were presented on the computer screen, and the key instructions were as follows:

You will be presented with 45 problems, each containing a pair of premises. These will concern some kind of relation between three different people or objects. Given that both premises are true, your task is to draw a conclusion about what follows validly (if anything) from the premises – that is, to type a conclusion about what **MUST** be true given that the premises are true. If you think that nothing follows from the premises, type in ‘nothing’.

Both premises were presented simultaneously, along with the question “what follows?” Participants had to press a key that allowed them to type in their conclusion, and the instructions emphasized that they should do so only when they knew what conclusion they were going to type. Participants’ responses were timed from the moment they were presented with the premises, until the time they hit the key to enter their conclusion, but they were not aware that their responses were being timed.

5.2. Results

Fig. 3 presents the percentages of transitive conclusions for the three sorts of premises: present tense (96%), past tense (63%), and mixed tense in which one premise was in the present tense and one premise was in the past tense (31%). This trend was highly reliable (Page’s $L = 330.5$, $z = 4.31$, $p < .0001$).

The results were different for transitive inferences drawn only in the first block of trials: present tense (92%), past tense (97%), and mixed tense (55%), and so they were drawn less often only for the problems with mixed tense. The effect of tense was reliable (Kruskal–Wallis test, $\chi^2(2) = 7.43$, $p < .03$). But the difference arose because problems in the past or present tense combined yielded more transitive conclusions than did those in mixed tense (Mann–Whitney $U = 31.00$, $z = 2.73$, $p < .01$).

The theory predicted that correct “no valid conclusion” responses to the transitive relations in the past tense or in mixed tenses should be faster when the entities in the premises can change quickly rather than slowly, e.g., treadmill inclines vs. mountains, with respect to steepness. The predicted difference in latencies was marginal for the premises in the past tense (7.7 vs. 10.8 s, $z = 1.52$, $p < .07$, one-tailed), and not reliable for premises in mixed tenses (15.1 vs. 12.9, $z = 0.00$, $p = 1.0$). But, there was little power for premises in the past tense: only seven subjects contributed data, and six of them showed the predicted pattern. With greater power, the predicted difference might emerge. However, it did not occur in the *percentages* of transitive conclusions drawn either. For premises in the past tense, a reliable difference occurred in the direction opposite to the prediction – the participants drew fewer transitive conclusions (and thus gave more non-transitive responses) from premises with




Fig. 3. The percentages of transitive conclusions drawn from pairs of premises in the present tense, in the past tense, and in mixed tenses, in Experiment 4. (There were zero transitive conclusions drawn for relations that were not transitive and in the present tense, hence the absence of a column in the figure).

slow-changing terms than from premises with fast changing terms (60% vs. 66%, Wilcoxon test, $z = 2.06$, $p < .04$). There was no reliable difference for the mixed tense problems. In sum, there was only weak evidence for the predicted difference in the ease of envisaging counterexamples. However, as predicted, individuals were quicker to envisage static rather than dynamic counterexamples. They took more time to respond “no valid conclusion” for transitive relations in the past tense or in mixed tenses, which require dynamic counterexamples, than to make the same response for non-transitive relations, which do not require dynamic counterexamples (12.0 vs. 10.6 s, Wilcoxon test, $z = 2.25$, $p < .03$). This difference occurred despite the greater length of the non-transitive premises than the transitive premises (a mean difference of 2.6 syllables). The need to envisage entities changing their properties added extra processing time.

5.3. Discussion

In accordance with the model theory’s main prediction, Experiment 4 showed that individuals were sensitive to the impact of tense in drawing conclusions from pairs of relational premises. For relations such as *taller than*, which are transitive in the present tense, individuals drew transitive conclusions almost always. But, when these relations occurred in premises in the past tense, the percentage of transitive

conclusions dropped strikingly. It dropped still further when the two premises had mixed tenses – one in the present tense and one in the past tense.

Although transitive conclusions were drawn less often for problems in the past tense than for those in the present tense, the percentages were roughly equivalent in the first block of problems, in which transitive conclusions were drawn for both sorts of problem. This tendency dropped for the problems in the past tense in the two subsequent blocks. This result suggests that an encounter with mixed tenses served as a cue to the possibility of counterexamples for problems in the past tense. Without such an explicit reminder that tenses can differ, participants tended to think of transitive models of premises in the past tense.

There was little evidence to support the model theory's second prediction: premises with entities that can change the relevant property quickly should yield counterexamples faster than premises with entities that can change the relevant property slowly. However, the theory's third prediction was corroborated: the complexity of the counterexamples to transitive conclusions affected the time it took participants to reject the conclusions. Dynamic counterexamples are those in which a series of events happen over time, e.g., individuals change their relative heights over time, whereas static counterexamples are those in which no such series has to occur. To withhold conclusions took participants longer for premises with dynamic counterexamples than for premises with static counterexamples.

6. General discussion

Many studies have shown that even children can draw valid transitive conclusions in three-term series problems, such as:

Ann is taller than Beth.
Beth is taller than Cath.
Therefore, Ann is taller than Cath.

These inferences are so easy that alternative theories jostle to explain them, and the data hardly determine which theory gives the best account of them. Our aim was to try a different tack and to examine a class of three-term series problems that even adults might get wrong. We showed first that certain relations are pseudo-transitive, that is, they are not transitive, but individuals tend to draw diagrams of them that support transitive conclusions, while at the same time accepting that such relations do yield counterexamples to transitive conclusions (Experiment 1). If individuals draw such diagrams, they are likely to have envisaged corresponding mental models representing the same sort of situations. We then showed that individuals do make invalid transitive inferences from pseudo-transitive premises (Experiment 2). For example, given the problem:

Fred is a blood relative of Bob's.
Bob is a blood relative of Anne's.

What follows?

they tended to draw the conclusion:

Fred is a blood relative of Anne's.

The conclusion is invalid in the case of father, son, and mother, but individuals are more likely to think of simple, typical relations, such as lineal descendants or siblings. The experiment also included genuine transitive relations – from which the participants drew transitive conclusions; and relations that are not transitive – from which the participants did not draw transitive conclusions. Hence, the results cannot be explained in terms of some general tendency to draw conclusions from premises laid out in the same way as the preceding example. The fallacies seem to occur because individuals overlook possibilities – the single most prevalent error in all sorts of thinking (Johnson-Laird, 2006) – and, in the present case, they fail to think of alternative possibilities that are counterexamples to their conclusions.

These results strengthen and generalize previous unsystematic and post hoc observations. De Soto and Kuethe (1958), for instance, observed that participants often incorrectly ascribed symmetry to relations that are not symmetric (e.g., *makes at least as much money as*) and transitivity to relations that are not transitive (e.g., *can usually beat at ping-pong*). However, the experimental instructions in their study encouraged participants to draw probable rather than deductively valid conclusions, and De Soto and Kuethe (1959) examined explicitly probabilistic conclusions. In contrast, our instructions emphasized that conclusions *must* be true given that the premises are true.

If individuals err because they think of simple and typical situations and overlook alternative possibilities, then one remedy should be to cue them to think of alternatives that refute invalid transitive conclusions. For example, a cue that reminds them that relations can be created by marriage should inhibit the fallacy with *blood relative* that we illustrated above. We showed that such cues did reliably reduce the fallacies without affecting other sorts of inference (Experiment 3).

Logicians often write as though a relational word in everyday life is transitive (e.g., Quine, 1974, p. 159). In their ideal language of the eternal present – the first-order predicate calculus – the claim is true. But, even such a simple example as “taller than”, shows that the time at which a relation holds is crucial. The following premises illustrate the issue:

Christina was taller than Jennifer.

Jennifer is taller than Theresa.

The choice of the past tense in the first premise seems to raise the possibility that matters may have changed since the time at which the relation held. And this possibility blocks the transitive inference that Christina is taller than Theresa. Since the time that the relation in the first premise held, Jennifer may have grown in height, so that the Christina is currently the shortest of the three. We showed that when seemingly transitive relations occur in premises of mixed tenses, such as in the pre-

ceding example, individuals no longer draw transitive conclusions from them (Experiment 4). And after a block of trials with mixed tenses, they are even inhibited about making the inferences from premises that are both in the past tense. Nevertheless, they still draw these inferences when both premises are in the present tense. It takes time to envisage such possibilities as that Jennifer has grown taller than Christina, and so the participants in the experiment took longer to decide that nothing followed from the premises than to draw the same conclusion based on premises that are obviously not transitive, such as:

Doctor Alvarado knows Doctor McNamara.
Doctor McNamara knows Doctor Villalobos.

The theory of mental models predicted the main phenomena that we observed, but is there perhaps an alternative theory that can explain them? We consider two possible ways of accounting for the phenomena: ambiguity and default assumptions, and then consider how formal rule theories might accommodate them.

The first way to try to account for the phenomena is to argue that pseudo-transitive relations are ambiguous: they have a sense that is transitive and a sense that is not transitive. Only the transitive sense invites a transitive inference (Politzer, 2004). This hypothesis is applicable to, say, a relation such as, *behind* that has both a sense concerning an individual's point of view, which is transitive, and a sense concerning the intrinsic parts of an object, which is not transitive. Hence, an assertion such as:

Steve is behind John

can be true from the speaker's point of view, but false in the intrinsic sense if the two of them are standing face to face. A sensible test for ambiguity is accordingly that a sentence can be true in one sense, but false in another. This test shows that *blood relative* is not ambiguous. There are different sorts of blood relation, some of which are transitive (*sister of*), and some of which are not (*father of*). But, a statement asserting a blood relation between two individuals is true if *any* of the different types of blood relation holds. Hence, the term *blood relative* refers to any sort of blood relation, and so it has a single meaning.

A variant on an account based on ambiguity is to argue instead that individuals in carrying out our tasks *instantiate the relations*, i.e., they replace the relations with more specific ones. Given, say, that they instantiate both premises based on *blood relative* with the relation *sister of*, they are entitled to infer the transitive conclusion. Moreover, as long as individuals instantiate the premises with the same specific relation, they are entitled to infer the transitive conclusion about the more general *blood relative* relation. However, in Experiment 1, the participants had to draw the first diagram that came to mind for the pseudo-transitive premises. Their diagrams showed that they had no reliable bias to treat the two premises as instantiating the same more specific relation. Their diagrams were consistent with such an assumption on only 36% of trials. Hence, there is no evidence for a universal instantiation of blood relative with the same more specific relation. Instead, participants often drew typical family trees of blood relatives

without any explicit process of instantiation, or they instantiated the premises with different specific relations. Moreover, no such process of instantiation seems possible for the other pseudo-transitive relations.

In general, an account based on ambiguity or instantiation cannot explain pseudo-transitive inferences. It also fails to explain why individuals in Experiment 1 judged diagrams of pseudo-transitive relations that did not yield transitive conclusions as nevertheless consistent with the relations (on 91% of trials). Yet, in Experiment 3, the participants drew transitive conclusions from the same relations (on 55% of trials). Binomial tests on each of the five relations, comparing the complement of the percentage of accepted non-transitive diagrams in Experiment 1 with the percentage of transitive conclusions in Experiment 3, yielded results ranging from $p < .02$ to $p < .001$. Although this comparison is made across different experiments, the participants in both experiments were from the same population. This result poses a substantial difficulty for an explanation of pseudo-transitivity based on invited inferences from one sense of ambiguous relations.

A second way to try to account for the pseudo-transitive fallacies is based on default assumptions. Tsal (1977), for instance, observed that individuals often assumed by default that an unknown relation was transitive and symmetric: they knew nothing about the relation because it was denoted by a meaningless symbol. Similarly, individuals drew transitive conclusions from quantified premises, such as: *Most X are Y, Most Y are Z* (Newstead, Pollard, & Griggs, 1986). Hence, individuals might draw a pseudo-transitive conclusion from relations, such as *blood relative*, because they assume transitivity by default. The default assumption, however, is overruled by the linguistic cue to marriage. One problem for this account is that individuals do not make a default assumption that all relations are transitive. As our results show, they make no such assumption for relations that are intransitive. So, until the account provides a mechanism to explain which relations are assumed to be transitive, and how context can block the assumption, this approach appears to be no more than a re-description of the results rather than a prediction of them.

In Section 1, we described theories of deductive reasoning based on formal rules of inference akin to those of logic (Braine & O'Brien, 1991, 1998; Rips, 1994). To account for mental proofs of transitive inferences, these theories postulate the existence in the mind of axioms, so-called “meaning postulates”, for the transitivity of relations, e.g.:

For any x, y, z , if x is taller than y , and y is taller than z , then x is taller than z .

Like logic, a simple version of formal theories posits that these postulates are indeed axiomatic: they apply in all circumstances. Hence, if you have the axiom for a relation, such as *blood relative*, then you should make transitive inferences from the relation, barring momentary lapses in performance. And if you do not have the axiom, then you can never make transitive inferences from the relation, barring occasional guesses or other illogical responses. Axioms should not switch on and off like a light, but instead they should have universal application, because they capture the fundamental logical properties of words (Bar-Hillel, 1967).

Our results yield a dilemma for formal rule theories as they are currently formulated. On the one hand, if individuals draw pseudo-transitive conclusions from a

relation such as *blood relative of*, then they have acquired an axiom that the relation is transitive. It follows that they should draw a transitive conclusion for the relation whenever they encounter premises that fit its antecedent condition. But, in this case, how are they able, as our results show, to refrain from the inference in a context that reminds them that individuals can be related by marriage? On the other hand, if individuals do not draw a pseudo-transitive conclusion in certain contexts, then they do not have the transitive axiom for the relation. But, in this case, how are they able to draw the transitive conclusion in the absence of the context?

The moral is that formal theories need to devise more nuanced ways to handle relational inferences. And this moral is reinforced by the results of Experiment 4 showing that the tense of premises affects whether or not a relation is transitive. Meaning postulates must take tense into account. They could be tagged to indicate that transitivity occurs only in the present tense. But, the variation from one participant to another, and from one block to another in Experiment 4, implies a corresponding variation in the availability of such tags, which poses an explanatory challenge for formal theories. It will not do to add a principle that, for instance, past tense undermines transitivity. If the exact moment in the past is specified for a relation such as *taller than*, the relation once again becomes transitive.

To handle both the effects of tense and those of context, more sophisticated machinery is required, which takes account of both the *meanings* of assertions, and their *contexts*. Such a move would add an extra component to the fundamental core of formal rule theories, which rely, not on meanings or contexts, but on logical forms. Formal rule theories have made some steps towards incorporating this kind of pragmatic machinery for conditional inference (see e.g., Braine & O'Brien, 1991). But, so far, this move has been confined to pragmatic principles that apply only to individual words. The present results demand a more thorough-going approach, in which such principles apply to the meanings of whole assertions, including their tenses.

The model theory predicts the phenomena in a consistent way. It invokes no axioms, no invited inferences, and no default assumptions of transitivity. Transitivity is an emergent property from the construction of models based on the meanings of premises and knowledge of context. When individuals build a model of an assertion such as, *Ann is taller than Beth*, the model represents the difference in heights iconically. A further assertion that *Beth is taller than Cath* yields a model in which Ann is taller than Cath. The transitive conclusion emerges from the model. No alternative model that is faithful to the meaning of the premises refutes the conclusion, and so it is valid. When individuals construct models from the relation *blood relative of*, they envisage a typical case, such as lineal descendants. The model yields a transitive conclusion, and they draw this invalid conclusion. When they are forewarned that some people are related by marriage, they can envisage a case that does not yield a transitive conclusion. And so they refrain from the invalid inference. Likewise, premises in the past tense alert them to the possibility of counterexamples, because entities can change their properties over time. But, these models can be difficult to envisage, and so participants do sometimes draw transitive conclusions from premises in the past tense. This difficulty also explains why the response of “no valid conclusion” took longer for transitive relations in the past tense or in mixed tenses than for

non-transitive relations. What is central throughout the process of inference is the meanings of the premises, and the models that individuals construct from them.

Individuals tend to construct typical models that satisfy premises, but what determines which models are typical? A major factor is likely to be their frequency of occurrence, and it may be possible to integrate probabilistic conceptions of reasoning within the theory (cf. Oaksford & Chater, 2007). Indeed, some recent formulations of the model theory in other domains of inference have parameterized its basic representational mechanisms (e.g., Oberauer, 2006). For relational inferences too, the process of model construction might be probabilistic, and dependent on background knowledge about the base-rates of certain real-world relations. For example, the probability that an individual builds a model of *blood relatives* that licenses transitivity may reflect the frequency with which this relation is actually transitive in experience. On this view, relations differ on a probabilistic continuum, in which the likelihood of drawing a transitive conclusion depends on knowledge about the environment, or at least the current representation of the environment. The hypothesis could be tested by biasing participants' representations of base-rate frequencies of certain relations, and examining their resulting inferences. Other factors in determining what individuals take to be typical may include a tendency to minimize the number of referents and relations in a model, and the complexity of these relations (Goodwin & Johnson-Laird, 2005; Halford et al., 1998). Ultimately, however, chance and unforeseeable factors will also affect what models become typical. Even granted the inclusion of a probabilistic component, our account does not dispense with the notion of logical validity, in contrast to more radical proposals in this vein (see e.g., Oaksford & Chater, 2007). It is still invalid to infer a transitive conclusion from the relation *blood relative* regardless of the reasonableness, probability, or adaptive value of such a conclusion.

The model theory's account of inference is parsimonious because it accounts for the vagaries of inference, even those that hinge on context. Imagine, for example, four individuals seated down one side of a table as in Leonardo's depiction of the last supper. A relation such as *on X's right-hand side* is transitive in this case. In contrast, if the four individuals are seated round a small circular table, the relation is intransitive. If the circular table were larger, however, then transitivity would extend over a small number of individuals, but gradually break down with a larger number as they got further round the table. And, if the individuals were seated on a bench, but some were facing in opposite directions from others, then transitivity would break down at once: two people could both be on each other's right-hand side. To capture these vagaries using formal rules calls for an indefinite number of axioms for the same relation. At one end of the spectrum, the relation is intransitive (the small round table), next it is transitive over three individuals but not four (a slightly larger round table), and so on, up to a relation that is unboundedly transitive (a long rectangular table). In fact, what matters in these examples is the context of the utterance – the actual seating arrangement to which it is known to refer, and whether or not one individual can be truly described as on another's right-hand side. A single meaning of *right-hand side* captures all these vagaries given a model of the seating arrangement. Transitivity is an emergent property from models, but, as pseudo-transitivity shows, it sometimes emerges in an illusory way.


Appendix A. The three sorts of diagram for the five relations in Experiment 1

Premises:

Fred is a blood relative of Bob's.

Bob is a blood relative of Anne's.

Transitive, consistent diagram Non-transitive, consistent diagram Inconsistent diagram




Premises:


Tim is on Mark's left.

Mark is on Fred's left.

Transitive, consistent diagram:




Non-transitive, consistent diagram(s):



The first group of participants was presented with the first combination of names: Fred Tim Mark which actually forms a transitive diagram if an egocentric interpretation of the relation is made. The second group was presented with the corrected diagram: Tim Fred Mark, which does not form a transitive diagram on either an egocentric or a deictic interpretation of the relation.

Inconsistent diagram:




Premises:


Ken is behind John.

John is behind the car.


Transitive, consistent diagram:



Non-transitive, consistent diagram:



Inconsistent diagram:




Premises:

In a particular series of events,



Planet Z went round planet X.

Planet X went round planet Y.


Transitive, consistent diagram Non-transitive, consistent diagram Inconsistent diagram*Premises:*

Alvarez overtook Underwood.

Underwood overtook Henderson.

Transitive, consistent diagram:*Non-transitive, consistent diagram:*

Inconsistent diagram:



Appendix B

The 20 problems in Experiment 3 with the contexts (in italics) designed to reduce the proportion of transitive inferences from the pseudo-transitive problems. In the experiment, half the problems in each block were negated (e.g., Mick is not taller than Ben).

Problem type	Block 1	Block 2
Pseudo- transitives (temporal)	Sue went faster than Anne, <i>and then</i> Anne went faster than Jess. Smith caught up with Clark, <i>and then</i> Clark did not catch up with Jones.	Ted went slower than Phil, <i>and then</i> Phil went slower than John. Tom overtook Chris, <i>and then</i> Chris overtook Dale.
(orientation)	<i>People and certain objects have a front side and a back side. As a result,</i> Anne is in front of Kim Kim is in front of the car. <i>At a small, round table,</i> Mary is on Claire's left Claire is on Beth's left.	<i>People and certain objects have a front side and a back side. As a result,</i> Frank is behind John John is behind the car. <i>At a small, round table,</i> Bob is on Mark's right Mark is on Ted's right.

(continued on next page)

Appendix B (*continued*)

Problem type	Block 1	Block 2
(simple structure)	<i>People can be related either by blood or by marriage. As a result,</i> Mel is a blood relative of Kate's. Kate is a blood relative of Jill's. <i>Sub-atomic particles are highly volatile, and may be acted on by many different forces. As a result,</i> The electron converged on the pion. The pion converged on the muon.	<i>An apparatus may be composed of multiple electric circuits. As a result,</i> The battery is connected electrically to the apparatus. The apparatus is connected electrically to the light. <i>Asteroids sometimes cross through the solar system. In one such instance,</i> The asteroid went round the planet. The planet went round the star.
Transitives	<i>Height is determined primarily by genes. As a result,</i> Mick is taller than Ben. Ben is taller than Fred. <i>The price of clothes is not always based on their aesthetic appeal. As a result,</i> The trousers are more expensive than the shoes. The shoes are more expensive than the braces.	<i>The sharpness of a blade depends partly on the density of the materials it is made from. As a result,</i> The knife is sharper than the axe. The axe is sharper than the sword. <i>The saltiness of food depends partly on how it is prepared. As a result,</i> The fish is saltier than the lamb. The lamb is saltier than the beef.
Not transitives	<i>Emotional relationships between people may not be symmetrical. As a result,</i> Max loves Jane. Jane loves Bill. <i>A man can be the father of many, but he himself can only have one father. As a result,</i> Bob is the father of James. James is the father of Ned.	<i>Emotional relationships between people may not be symmetrical. As a result,</i> Keith envies Liz. Liz envies Steve. <i>A woman can be the mother of many, but she herself can only have one mother. As a result,</i> Kay is the mother of Sal. Sal is the mother of Jen.

References

- Andrews, G., & Halford, G. S. (1998). Children's ability to make transitive inferences: The importance of premise integration and structural complexity. *Cognitive Development*, 13, 479–513.
- Andrews, G., & Halford, G. S. (2002). A cognitive complexity metric applied to cognitive development. *Cognitive Psychology*, 45, 153–219.
- Baddeley, A. D. (1986). *Working memory*. Oxford: Clarendon Press.
- Bar-Hillel, Y. (1967). Dictionaries and meaning rules. *Foundations of Language*, 3, 409–414.
- Barwise, J. (1989). *The situation in logic*. Stanford, CA: Center for the Study of Language and Information Publications.
- Beth, E. W. (1971). *Aspects of modern logic*. New York: Humanities Press.
- Braine, M. D. S., & O'Brien, D. P. (1991). A theory of if: A lexical entry, reasoning program, and pragmatic principles. *Psychological Review*, 98, 182–203.
- Braine, M. D. S., & O'Brien, D. P. (Eds.). (1998). *Mental logic*. Mahwah, NJ: Erlbaum.
- Bryant, P. E., & Trabasso, T. (1971). Transitive inferences and memory in young children. *Nature*, 232, 456–458.
- Chaffin, R., & Herrmann, D. J. (1984). The similarity and diversity of semantic relations. *Memory & Cognition*, 12, 134–141.
- Clark, H. H. (1969). Linguistic processes in deductive reasoning. *Psychological Review*, 76, 387–404.
- De Soto, C. B., & Kuethe, J. L. (1958). Perception of mathematical properties of interpersonal relations. *Perceptual and Motor Skills*, 8, 279–286.
- De Soto, C. B., & Kuethe, J. L. (1959). Subjective probabilities of interpersonal relationships. *Journal of Abnormal and Social Psychology*, 59, 290–294.
- De Soto, C. B., London, M., & Handel, S. (1965). Social reasoning and spatial paralogic. *Journal of Personality and Social Psychology*, 2, 513–521.
- Fodor, J. (1987). A situated grandmother?. *Mind and Language* 2, 64–81.
- Goodwin, G. P., & Johnson-Laird, P. N. (2005). Reasoning about relations. *Psychological Review*, 112, 468–493.
- Halford, G. S., Wilson, W. H., & Phillips, S. (1998). Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behavioural and Brain Sciences*, 21, 803–831.
- Herrick, P. (1994). *The many worlds of logic*. Orlando, FL: Harcourt Brace.
- Hummel, J. E., & Holyoak, K. J. (2001). A process model of human transitive inference. In M. Gattis (Ed.), *Spatial schemas and abstract thought* (pp. 279–305). Cambridge, MA: MIT Press.
- Hunter, I. M. L. (1957). The solving of three term series problems. *British Journal of Psychology*, 48, 286–298.
- Huttenlocher, J. (1968). Constructing spatial images: A strategy in reasoning. *Psychological Review*, 75, 286–298.
- Jeffrey, R. (1981). *Formal logic: Its scope and limits*. New York: McGraw-Hill.
- Johnson-Laird, P. N. (2006). *How we reason*. New York: Oxford University Press.
- Johnson-Laird, P. N., & Byrne, R. M. J. (1991). *Deduction*. Hillsdale, NJ: Erlbaum.
- Johnson-Laird, P. N., Legrenzi, P., Girotto, V., & Legrenzi, M. S. (2000). Fallacies in reasoning about consistency. *Science*, 288, 531–532.
- Knauff, M., & Johnson-Laird, P. N. (2002). Visual imagery can impede reasoning. *Memory & Cognition*, 30, 363–371.
- Kuczaj, S. A., II, & Donaldson, S. A. III, (1982). If the boy loves the girl and the girl loves the dog, does the boy love the dog? The overgeneralization of verbal transitive inference skills. *Journal of Psycholinguistic Research*, 11, 197–206.
- Newstead, S. E., Pollard, P., & Griggs, R. A. (1986). Response bias in relational reasoning. *Bulletin of the Psychonomic Society*, 24, 95–98.
- Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford, UK: Oxford University Press.

- Oberauer, K. (2006). Reasoning with conditionals: A test of formal models of four theories. *Cognitive Psychology*, 53, 238–283.
- Ormerod, T. C., & Richardson, J. (2003). On the generation and evaluation of inferences from single premises. *Memory & Cognition*, 31, 467–478.
- Politzer, G. (2004). Reasoning, judgement and pragmatics. In Ira A. Noveck & D. Sperber (Eds.), *Experimental pragmatics* (pp. 94–115). Hampshire, UK: Macmillan.
- Quine, W. V. O. (1974). *Methods of logic* (3rd ed.). London: Routledge & Kegan Paul.
- Rips, L. J. (1994). *The psychology of proof*. Cambridge, MA: MIT Press.
- Tsal, Y. (1977). Symmetry and transitivity assumptions about a nonspecified logical relation. *Quarterly Journal of Experimental Psychology*, 29, 677–684.