Exploring the asymmetrical representation of causal relations in semantic memory

by

Nathaniel Barr

A thesis

presented to the University of Waterloo

in fulfilment of the

thesis requirement for the degree of

Master of Arts

in

Psychology

Waterloo, Ontario, Canada, 2010

© Nathaniel Barr 2010

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Nathaniel Barr

Abstract

The current study provides evidence for the dissociation between two types of asymmetries in relations within semantic memory: those due to unidirectional associative strength (e.g., Kahan, Neely & Forsythe, 1999), and the inherent asymmetry of causal relations (e.g., Fenker, Waldmann & Holyoak, 2005). By manipulating the stimulus onset asynchrony (SOA) in a relation recognition task, I demonstrate that expectancy differentially impacts the effects of these asymmetries. An asymmetrical directional response time advantage was seen with causal relations at both long (1000 ms) and short (150 ms) SOAs, but only at long SOAs for unidirectional associates. These data are taken to support the hypothesis that latencies due to unidirectional association are a result of the manner in which these relations are accessed, and latencies due to the asymmetry of causal relations are reflective of the manner in which cause-effect asymmetries are represented.

Acknowledgments

I would like to extend my appreciation to Jonathan Fugelsang, Jennifer Stolz, and Derek Besner for their assistance in this project specifically, and my development as a researcher in general. I would like to especially thank Jonathan for his patience and enthusiasm throughout the course of this project. I would additionally like to thank those listed above and Mike Dixon for their instrumental roles in affording me the opportunity to conduct this research.

Table of Contents

List of Figures

List of Tables

Introduction

The manner in which humans represent causal knowledge and how previously learned causal relations are accessed in semantic memory has become the focus of increasing amounts of research in recent years. Semantic memory is described as our long-term inventory of knowledge about the world (Tulving, 1972), including information about categories, features, and the complex inter-relations that exist between them (Murphy & Medin, 1985). Importantly, one's understanding of causal relations is also a part of this general semantic knowledge base. There are a number of features of causal knowledge that may set it apart from other forms of semantic knowledge. For example, there exists a growing body of evidence suggesting that the perception and learning of causal relations is intimately tied to basic perceptual processes and action (i.e., dynamic events) that gives rise to the seemingly high-level cognitive processing required for causal inference (see Scholl & Tremoulet, 2000 for review). Furthermore, research from diverse approaches, including phenomenal causality (e.g., Michotte, 1963), developmental psychology (e.g., Leslie & Keeble, 1987), and human and animal learning (e.g., Waldmann, Hagmayer, & Blaisdell, 2006) has revealed a deep understanding of causal relations even amongst infants and non-human animals.

An important question concerns how such relations are represented in semantic memory. There are numerous features of causal knowledge that would lead one to predict that it may be dissociable from other forms of semantic relations at the level of representation. Most crucially, causal relations are inherently asymmetrical, as the constituent parts of a causal relation (cause and effect) have certain non-interchangeable binding roles (Pearl, 2000). For example, sunlight can cause freckles, but the inverse is not true: freckles cannot cause sunlight. In this sense, it seems that the asymmetry of causal relations must be tied to the semantic, and perhaps even structural, representation of these relations in that the ordering is integral in extracting meaning; unlike in asymmetrical general associative relations. A superficial analogy that can illustrate this point is that in multiplication, like general association, one can reverse the terms involved and still derive the same result. However, in division, like causal relations, if one reverses the order of the operands, the result will be fundamentally different.

In a series of three experiments, Fenker et al., (2005) explored this asymmetry in causal relations using a *relation recognition paradigm*. The authors recorded the response times of participants while they determined whether or not word pairs that were roughly equal in directional associative strength could be causally or associatively related. They found that when queried specifically if a causal relation could exist, word pairs presented in the predictive, cause-effect order were accessed, and identified as potentially causally related, faster than when the same pairs were presented in the diagnostic, effect-cause order. However, when asked if an associative relation could exist between the same word pairs, no differences in response times were observed as a function of directionality. As the stimuli were normed to be bidirectionally equal in associative strength, these data were taken as evidence that causal relations were distinct from general associative relations, and that general association did not underlie the predictive priming advantage found with causal relations.

Although this study provided compelling evidence that unidirectional association (or unequal associative strength) was not responsible for the observed effects, it did not provide any test as to the nature of the representation underlying the asymmetry in causal relations nor if this type of asymmetry differs in its representation from the type of asymmetry that was controlled for (unequal associative strength). For one, not using asymmetrically associated word pairs as a control makes any inferences regarding differentiation between these types of asymmetries beyond the scope of their investigation. Moreover, Fenker et al.'s choice of a 1000 millisecond (ms) SOA (which in this paradigm translates to how long the first word is present before it is replaced by the second) leaves open the possibility that the direction effects observed could be a reflection of expectancy; that is, the ability of participants to *predict* effects from causes compared to *predicting* causes from effects.

The role of expectancy is often discussed in traditional semantic priming and word recognition research (e.g., Becker, 1980), but I feel that its importance within the context of the *relation recognition paradigm* is worth consideration. If participants more effectively predict the ensuing target word when presented with a cause than when presented with an effect, the predictive advantage described could merely be a reflection of this ability. However, if this RT advantage persists despite limiting participants' ability to effectively predict ensuing targets, this would constitute strong evidence that the predictive advantage is capturing something reflective of the *representation* of causal relations themselves, rather than just how they may be accessed. The current study tests this hypothesis.

Method

Participants

Sixty undergraduate students from the University of Waterloo participated in this experiment for course credit. Data from 12 participants had to be replaced due to higher than 25% error rate in at least one condition. All participants had normal, or corrected to normal vision, and spoke English as a first language.

Design

The current study used a 2 (Task: associative vs. causal) x 2 (SOA: 150 ms vs. 1000 ms) x 2 (Direction: predictive/forward vs. diagnostic/backward) within-subjects design. All subjects made both causal and associative judgments at both SOAs and judged all word pairs. The Task variable was blocked and block order was counterbalanced. The order of presentation of each word pair in each condition was randomized, as was the SOA used in each trial. The specific items appearing at each SOA and in each direction were counterbalanced across participants, whereby participants saw different word pairs in different directions at different SOAs.

Stimuli

Causal stimuli. The causal stimuli consisted of 64 causally related word pairs and 64 weakly associated word pairs adapted from Fenker et al. (2005) which are included as Appendix A. Weakly associated pairs were used as the filler stimuli rather than unrelated word pairs to prevent participants from being able to use association as a cue to causality. These stimuli were originally selected from the USF Word Association Norms List (Nelson et al., 1998). The word pairs were controlled for strength of association in each direction and bidirectional statistical contingency. For a more detailed review of the norming procedures please refer to Fenker et al. (2005).

Unidirectionally associated stimuli (UDA). The UDA and unrelated word pairs used were adapted from Kahan, Neely and Forsythe (1999) and are included as Appendix B. With these word pairs, Kahan et al. found that the average probability of responding with the target given the prime was 17%, whereas the average probability of responding with the prime given the target was only 1%. These response patterns clearly demonstrate the asymmetry in that an association exists in the forwards (prime to target) direction and not in the backwards (target to prime) direction.

Unrelated word pairs were used as filler items in an effort to make the task roughly equal in nature to that of the causal task. Pilot testing was conducted in which the UDA word pairs and unrelated word pairs from Kahan et al. (1999) were rated on the degree to which they were associated in a forwards direction and in a backwards direction. Judgments were made using an 11-point Likert scale ranging from 0 to 10, with 0 indicating a response of "not associated" and 10 indicating a response of "highly associated". Each word pair was rated in either direction by an independent sample of twenty-five undergraduate students from the University of Waterloo. There were 120 UDA word pairs and 120 unrelated word pairs (originally from Kahan et al., 1999) included in the pilot study. Of these pairs, 64 of the most strongly unidirectionally related pairs were selected to serve as the critical stimuli and 64 of the pairs deemed to be least related were selected to serve as fillers.

Procedure

Half of the participants completed the associative task first, and the other half completed the causal task first. Written instructions were provided instructing participants to assess whether a causal or associative relation (depending on the condition) could exist between the words presented in the ensuing pairs. To ensure that participants understood that valid causal pairs could be presented bidirectionally (i.e., cause-effect, or effect-cause order), participants were told to assess "whether the event described by the first word *could cause* or *could be caused by* the event described by the other word". If further clarification was required, participants were told to decide if one of the words in a given pair could cause the other. In the associative condition, participants were asked to assess "whether they believed an associative relation could exist between the presented word pairs".

Once the task commenced, participants saw a fixation point in black font in the center of a white screen for 1000 ms after which the screen was blank for 500 ms. The first word then appeared in Arial Black font size 24, on a white background, for either 150 ms or 1000 ms and the second word was presented immediately after (i.e., 0 ms inter-stimulus-interval). The second word remained on the screen until the participant responded. If subjects believed the items could be causally/associatively related, they were to press the letter "C" on the keyboard. If they believed the items not to be causally/associatively related, they were to press the "N" key. Figure 1 depicts the time course of the presentation of tasks.

Results

Response Times (RTs)

The analysis of RTs excluded all incorrect responses and any outliers that were more than two standard deviations above or below each participant's mean within each task and condition. This outlier criterion resulted in the removal of approximately 3% of trials. Table 1 presents the mean RTs, percentage errors, difference scores, and as the percentage reduction from the slower, diagnostic/backward order to the predictive/forward (% Reduction). For individual subject means, please refer to Appendix C for RT's in the Associative task as a function of direction at each SOA, and refer to Appendix D for RT's in the Causal task as a function of direction at each SOA.

A 2 (Task: associative vs. causal) x 2 (SOA: 150 ms vs. 1000 ms) x 2 (Direction: predictive/forward vs. diagnostic/backward) within-subjects ANOVA was performed on the RT data. Analyses revealed significant main effects of Task, $F(1, 59) = 118.21$, $MSE = 244614$, $p < .001$, $p\eta^2 = .67$, SOA, $F(1, 59) = 39.82$, $MSE = 22447$, $p < .001$, $p\eta^2 = .40$, and Direction, $F(1, 59) = 33.79$, $MSE = 17995$, $p < .001$, $p\eta^2 = .36$. Also present were significant interactions between Task and SOA, $F(1, 59) = 8.04$, $MSE = 21995$, $p = .006$, $p\eta^2 = .12$, and Task and Direction, $F(1, 59) = 6.10$, $MSE = 22357$, $p = .016$, $p\eta^2 = .09$. The interaction between SOA and Direction ($F < 1$), and Task, SOA, and Direction, and was not significant, $F(1, 59) = 1.19$, $MSE = 22047$, $p = .28$, $p\eta^2 = .02$. For the complete ANOVA table please refer to Appendix E. Based on the significant Task x SOA and Task x Direction interactions, and to test our key predictions regarding the impact of Direction and SOA for each Task, two separate 2 x 2 ANOVAs were carried out for each Task, with SOA and Direction as factors. The critical result for my hypotheses is the presence or absence of an SOA x Direction interaction in these

Figure 1. Priming task event sequence

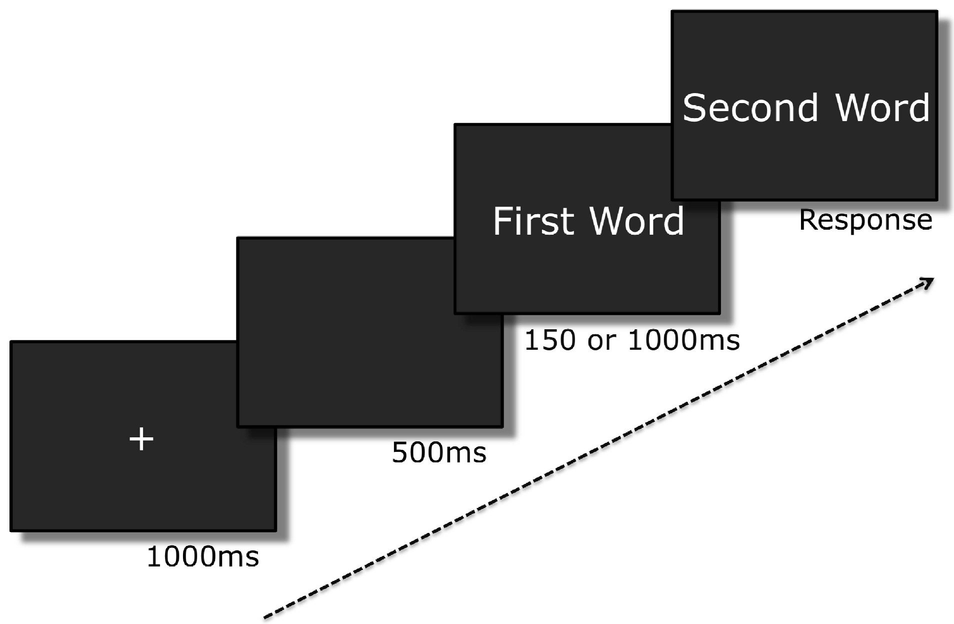

Table 1. Data summary

	Causal Judgment Task				Associative Judgment Task			
	150 ms SOA		1000 ms SOA		150 ms SOA		1000 ms SOA	
	RT	% Error	RT	% Error	RT	% Error	RT	% Error
Predictive/Forwards	1334	14.8	1215	12.3	853	8.5	781	8.3
Diagnostic/Backwards	1444	18.0	1314	17.9	867	9.7	843	9.1
Difference	110	3.2	99	5.6	14	1.2	62	0.8
% Reduction	7.6		7.5		1.6		7.4	

Note - Mean RTs (in milliseconds), Percentage Errors, Difference scores, and Percentage Reduction in Mean RT between Diagnostic/Backwards and Predictive/Forwards are depicted for each Task x SOA x Direction condition.

tasks. The absence of this interaction would imply that the size of the direction effect does not vary as a function of SOA, whereas the presence of this interaction would constitute evidence that the magnitude of the direction effect depends on the SOA. In the Causal Task, the SOA x Direction was clearly insignificant, $F < 1$, $p = .81$, $p\eta^2 = .001$. In the Associative Task, however, the SOA x Direction interaction was significant, $F(1, 59) = 4.45$, $MSE = 7839$, $p < .05$, $p\eta^2 = .07$.

To follow up these analyses, a series of planned paired sample *t*-tests were conducted to investigate the precise effect of the SA manipulation on RT for each type of stimuli in each direction. Within the Causal Task, significant differences were found as a function of direction at both SOAs. Specifically, in the 1000 ms SOA condition, RTs were approximately 99 ms faster for predictive trials ($M = 1215$ ms) than for diagnostic trials ($M = 1314$ ms), $t(59) = 2.58$, $SE = 38.55$, $p = .012$. In the 150 ms SOA condition, RTs were approximately 110 ms faster for predictive trials ($M = 1334$ ms) than for diagnostic trials ($M = 1444$ ms), $t(59) = 4.07$, $SE = 27.12$, $p < .001$. In contrast, in the associative task significant differences between the directions were only obtained at the long SOA. Specifically, in the 1000 ms SOA condition, RTs were approximately 62 ms faster for forwards trials ($M = 781$ ms) than for backwards trials ($M = 843$ ms), $t(59) = 4.18$, $SE = 14.74$, $p < .001$. However, in the 150 ms SOA condition, RTs for forwards ($M = 853$ ms) and backwards ($M = 867$ ms), trials were clearly not significantly different, $t(59) = 1.06$, $SE = 12.68$, $p = .295$.

An alternative way to analyse these data is to re-express the observed RT difference between the predictive/forward and the diagnostic/backward direction as a *percentage reduction* from the RT of the slower, diagnostic/backward order (see % Reduction in Table 1). This re-expression of the data reinforces the conclusions derived from the parametric analyses.

Specifically, in the causal task, at the long and short SOAs respectively, a 7.5% and 7.6% reduction as a function of direction exists. In the associative task, at the long SOA, a 7.4% reduction occurs, however, at the short SOA, only a 1.6% reduction in RT occurs as a function of direction. As such, there is a clear dissociation in the manner in which these different types of asymmetries impact RTs.

Errors

A parallel 2 x 2 x 2 ANOVA conducted on the error data revealed a significant main effect of Task, $F(1, 59) = 19.44$, $MSE = 290.66$, $p < .001$, $p\eta^2 = .25$, with significantly more errors being made in the causal task than in the associative task. Also present were main effects of Direction, $F(1, 59) = 14.86$, $MSE = 58.65$, $p < .001$, $p\eta^2 = .20$, and an interaction between Task and Direction, $F(1, 59) = 4.75$, $MSE = 69.03$, $p < .05$, $p\eta^2 = .08$. Paired samples t-tests revealed that this interaction was due to the finding that no significant differences existed as a function of direction within the associative task at either SOA (Largest $t = 1.11$, $p = .27$), but, in the causal task more errors were made on diagnostic than predictive trials at both the 150 ms SOA, $t(59) = 2.19$, $SE = 1.45$, $p < .05$, and the 1000 ms SOA, $t(59) = 2.78$, $SE = 1.98$, $p < .01$. As such, the error data are in line with those from the RT analysis. In the causal task, direction impacted accuracy in that participants performed more poorly when pairs were presented in the diagnostic than in the predictive order, but in the associative task, direction had no effect upon accuracy. This dissociation lends further support to the idea that the asymmetry in causal relations is represented as an integral part of these relations but that this is not the case for unidirectionally associated pairs.

General Discussion

The representation of causal structure in semantic memory has been, and continues to be, a topic of much debate. Specifically, whether the inherent asymmetry of causal relations is a feature represented in semantic memory that is substantively distinct from asymmetries present in other domains (e.g., unidirectional associative strength) remains to be seen. I propose that one way in which these types of asymmetries might be distinct is the level at which direction is impacting responses (i.e., at the level of *representation* or *access*). If an asymmetry is intimately tied to the structure or meaning of that relation, one would expect that responses regarding those relations to be relatively immune to manipulations that limit higher-order processes, such as expectancy. If however, directional information is not intimately tied to the structure or meaning of that relation, manipulations that limit higher-order processes, such as expectancy, should impact direction effects as they are likely tied to how that directional information is accessed.

Our results support the notion that causal asymmetry is distinct from unequal associative strength, and that the directional information in these relations are likely represented in distinct ways. Specifically, on the long SOA trials, both causal and associative judgments exhibited an effect of direction with an RT advantage for word pairs present in the predictive/forwards order. However, on trials using short SOAs, an effect of direction was only observed for causal trials, and not unidirectional associates. These findings indicate that directional advantages associated with asymmetrical associative strength may be contingent upon one's ability to use expectancy-based strategies and generate potential targets that fulfill the requirements of the relation in question. As such, this type of directional effect seems to be a result of higher-order strategic processes employed by participants within this particular paradigm and are likely a result of the manner in which participants *access* these relations. It is likely that participants are better able to predict a target based on a prime when presented in the forwards order, as compared to backwards, and this ability could be responsible for the RT advantage on these trials. This explanation seems to be the most parsimonious and is in accord with the results of a norming study conducted on these and other words in which Kahan et al. found that the average probability of responding with the target given the prime was 17%, whereas the average probability of responding with the prime given the target was only 1%. Conversely, the type of directional effects exerted upon RTs for causal judgments does not hinge upon higher-order strategic processes and is unaffected by conscious attempts to generate targets that fulfill the required roles. As this directional advantage was robust to the SOA manipulation, I would argue that this advantage has more to do with the *representation* of this directional information in semantic memory. In light of the significant difference in overall RT between the tasks, it is important to note that the directional effects observed are interpreted as indexing the relative speed with which direction is influencing judgments, which is distinct from the absolute speed of the judgment task. Put another way, it is the speed with which a feature of the representation (direction) of the relation becomes activated, not overall task speed in which we are focusing upon (see Schwarz & Ischebeck, 2003, for a similar argument regarding the relative influence of separate components of relation information in number comparison).

The fact that the inherent asymmetry of causal relations is a facet of the representation of these relations seems logical in that the meaning of causal relations, and thus their utility in guiding action, is crucially dependent on the alignment of the constituent parts of the relation in the proper roles. Conversely, asymmetrical associative relations are still related regardless of the order in which they are aligned. The direction effects within the associative task may be a result of a familiarity-based mechanism due to increased exposure to these pairs in the familiar forward order in language (e.g., as commonly found in writing and speech). If one considers the large variance in language across cultures, it seems likely that effects of this nature would be heavily dependent upon the predominant ordering of word pairs in normal writing and speech within a particular language, and thus, would have variant patterns of results if the same stimulus sets were directly translated. It would seem that, unlike asymmetries due to unidirectional associative strength, the asymmetry of causal relations (at least those mechanistic in nature) transcend language in that no matter the tongue in which it is spoken, causal relations are still inherently tied to the constituent roles at play. Indeed, support for this notion comes from previous work demonstrating that causal perceptions are largely invariant across cultures (Morris & Peng, 1994).

Evidence that the structural representation of causal relations might be contributing to the observed latency differences in the causal task also comes from recent work aimed at determining the neural correlates of evaluating causal and associative relations in semantic memory. Satpute et al., (2005) employed a causal and associative relation recognition task while participants were scanned using fMRI. Unique activation in the left dorsolateral prefrontal cortex (DLFPC) and the right precuneus were observed when judging causal as opposed to associative relations. They argued that these unique activations link the access of causal relations with regions of the brain known to be implicated in role-binding (left DLFPC). The DLPFC activation is also consistent with the view that this binding requires extra working memory resources (Hummel & Holyoak, 2003). These findings may challenge the view, however, that the binding process per se requires extra working memory in that my data suggests that the binding process is likely is an integral part of the representation itself.

A longstanding debate exists between researchers advocating two broad theoretical perspectives regarding the nature of causal representation. Proponents of associative theories posit that the constituent parts of a causal relation are mapped onto the temporally defined cue or outcome roles, which can be either causes or effects (e.g., Shanks & López, 1996). This temporal precedence creates asymmetries in associative strength, and it is this asymmetry that has implications in subsequent access of these relations. In contrast, the causal-model view (Waldmann & Holyoak, 1992) proposes that people explicitly encode the specific roles of cause and effect to events and use this knowledge in subsequent scenarios to guide their judgments of causality. Here, the cause-effect asymmetry is a result of these binding roles impacting behavior. Although I acknowledge the importance of this debate, I feel that the present investigation need not directly appeal to either of these perspectives. That being said, the implications of these findings for either theory is open to debate, and future theoretical work on this topic would be well served to explicitly examine the degree to which causal-model and associative theories of causal knowledge can accommodate the present findings.

One possible avenue for future research could be the utilization of a more parametric distribution of SOAs to establish the boundary conditions of the phenomena described herein. In the current study, we attempted to select two SOAs designed to differentiate between scenarios where expectancy could easily be active, and those in which it is much less likely to be operational. By parametrically varying the SOA, one could garner an understanding of exactly when these types of asymmetries become dissociable.

Subsequent research on this topic should also attempt to situate these judgments in a more ecologically valid setting in which context plays a role. It is well known that context alters one's perception of association. For example, Balota & Lorch (1986) demonstrated that subjective ratings of associative strength are influenced by the strength of association of other items in the list. Similar contextual factors also likely influence causal judgments. If one considers social causal scenarios in particular, it may be especially likely that context could have a significant role. Future work could endeavor to systematically vary the causal strength of the candidate pairs, and determine if the directional advantages observed here varies as a function of this variability in causal strength. Based on the data and arguments outlined here, it is likely that variance would be observed in the response and the absolute speed of judgments as a function of causal power; however, I would predict that the directional advantages observed here would be immune to such variations in context.

A related issue that also pertains to the generalizability of these results is the differentiation between explicit speeded judgments, such as that employed here, and more implicit measures designed to indirectly measure the impact of these various types of asymmetries on behavior. As the correlation between one's explicit judgments pertaining to such tasks and one's implicit processing of such stimuli is far from transparent, it would be theoretically informative to see if similar dissociations between associative and causal relations would be reported if participants engaged in some kind of implicit task.

In summary, these results, although pertaining to causal relations specifically, provide broader insight into the way that asymmetrical relations in semantic memory may be represented and become manifest. To determine more precisely how various kinds of asymmetries impact the access of relational semantic knowledge, similar methodology and manipulations could be used to explore other types of semantic relations that have inherent asymmetries; including categorical, taxonomic and functional relations. Furthermore, the current study, and future research examining various asymmetrical associations in semantic memory, will aid in the development of more comprehensive computational models that attempt to mirror human semantic memory.

To date, relatively little work has focused directly on the relations between concepts stored in semantic memory (see Spellman, Holyoak, & Morrison, 2001), and as such, it is unsurprising that a paucity of research exists that focuses on more specific aspects of such relations, including the types of asymmetries that impact them. Future research on semantic memory will be well served to refine its focus onto such components and begin to decompose the subtleties that comprise semantic relations; as Fenker et al. (2005) astutely noted: “It is important to identify both the commonalities and differences among the varied semantic relations that constitute our knowledge of the world.” (p. 1042).

References

Balota, D. A., & Lorch, R. F. (1986). Depth of automatic spreading activation: Mediated priming effects in pronunciation but not in lexical decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 336-345.

Becker, C.A. (1980). Semantic context effects in visual word recognition: An analysis of semantic strategies. *Memory & Cognition*, *8*, 493-512.

Fenker, D., Waldmann, M., & Holyoak, K.J. (2005). Accessing causal relations in semantic memory. *Memory & Cognition*, *33*, 1036-1046.

Hummel, J.E., & Holyoak, K.J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review, 110*, 220-264.

Kahan, T., Neely, J., & Forsythe, W. (1999). Dissociated backward priming effects in lexical decision and pronunciation tasks. *Psychonomic Bulletin &* Review, *6*, 105-110.

Leslie, A.M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition,* 25, 265-287.

Michotte, A. (1963). The perception of causality. (T.R. Miles & E. Miles, Trans.). London: Methuen. (English translation of Michotte, 1954).

Morris, M.W., & Peng, K (1994). Culture and cause: American and Chinese attributions for social and physical events. *Journal of Personality and Social Psychology, 67*, 949-971.

Murphy, G. L. & Medin, D. L. (1985). *The role of theories in conceptual coherence*. *Psychological Review*, *92*, 289-316.

Nelson, D.L., McEvoy, C.L. & Schreiber, T.A. (1998). *The University of South Florida word association, rhyme, and word fragment norms*. Available at w3.usf.edo/FreeAssociation/.

Pearl, J. (2000). *Causality: Models, reasoning, and inference.* New York, NY, US: Cambridge University Press.

Satpute, A.B., Fenker, D.B., Waldmann, M.R., Tabibnia, G., Holyoak, K.J., & Lieberman, M.D. (2005). An fMRI study of causal judgments. *European Journal of Neuroscience, 22*, 1233-1238.

Scholl, B. J. & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences,* 4, 299-309.

Schwarz, W., & Ischebeck, A. (2003). On the relative speed account of number-size interference in comparative judgements of numerals. *Journal of Experimental Psychology: Human Perception and Performance*, 29, 507-522.

Shanks, D. R. & López, F.J. (1996). Causal order does not affect cue selection in human associative learning. *Memory & Cognition*, *24*, 511-522.

Spellman, B.A., Holyoak, K.J., & Morrison, R.G. (1996). Analogical priming via semantic relations. *Memory & Cognition*, *29*, 383-393.

Stolz, J.A. & Neely, J.H. (1995). When target degradation does and does not enhance semantic context effects in word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 596-611.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory*. New York: Academic Press.

Waldmann, M.R. & Holyak, K.J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General, 121*, 222-236.

Waldmann, M. R., Hagmayer, Y. & Blaisdell, A. P. (2006). Beyond the information given: Causal models in learning and reasoning. *Current Directions in Psychological Science, 15,* 307-311.

Appendix A

Sub-set of word pairs used in Fenker, D., Waldmann, M., and Holyoak, K. (2005). Accessing causal relations in semantic memory. *Memory & Cognition,* 33(6), 1036-1046. (**Original Source:** Nelson, D.L., McEvoy, C.L. and Schreiber, T.A. (1998).)

Causal Word Pairs		Weak Associate Filler Pairs	
absence	withdrawal	acrobat	athletes
alcohol	accident	apple	computer
attack	defense	atlas	dictionary
bacteria	infection	baptist	protestants
bang	deafness	basketball	teams
beat	bruise	bicycle	vehicle
betrayal	distrust	caffeine	mountain
birthrate	population	chipmunks	acorn
carcinogen	tumor	claw	dogs
chromosome	gender	control	volume
compliment	blush	corona	lime
crime	arrest	course	session
crush	damage	curve	shape
dairy	diarrhea	decency	respect
diet	hunger	elevator	floor
disease	injection	email	attachment
drought	famine	employment	office
drug	relief	fight	dagger
education	career	forecast	weather
eyedrops	dilation	fruits	cocktails
fertilizer	growth	furniture	bedroom
fracture	cast	gallon	ounce
frequency	pitch	gazelle	antelope
frowning	wrinkles	girl	maid
gang	riot	glass	window
gases	explosion	gossip	newspaper
gold	wealth	gown	graduation
hormones	mood	grab	pull
humidity	sweat	groceries	money
illness	treatment	hypothesis	test
invitation	visit	insurance	estimate
joke	amusement	kill	theft
lamp	heat	kindness	sympathy
lesion	scar	latin	medicine

lightning	fire	lettuce	vegetables
magnet	attraction	lover	girlfriend
moon	tide	officers	uniforms
movie	nightmare	painting	wall
mutation	cancer	passage	story
nuts	allergy	patty	hamburger
order	delivery	pituitary	glands
panic	escape	plane	car
period	cramps	planter	farmer
pollution	asthma	potatoes	ground
pressure	bursting	power	voltage
sadness	crying	propeller	helicopter
salt	thirst	ring	emerald
scratch	blood	roar	engine
shampoo	tears	rush	ambulance
shock	scream	screw	round
spice	flavor	seaman	harbour
spill	stain	security	force
sprain	swell	shrimp	ocean
stress	fatigue	sibling	family
study	pass	skyscraper	towers
sunlight	freckles	soup	cracker
sweets	cavity	spray	roach
training	fitness	terms	meaning
trash	stink	tomatoes	sandwich
trauma	coma	tote	umbrella
UVlight	tanning	tuba	saxophone
vacuum	suction	vein	vessel
virus	epidemic	wife	mother
wind	erosion	zebra	elephant

Appendix B

Sub-set of pairs used in Kahan, T., Neely, J., & Forsythe, W. (1999). Dissociated backward priming effects in lexical decision and pronunciation tasks. *Psychonomic Bulletin & Review*, 6(1), 105-110.

<u>Unidirectional Associates</u>

acre	land
alright	okay
anatomy	body
anger	mad
antique	old
ashtray	cigarette
behind	front
blaze	fire
brief	short
campus	school
cent	penny
cigar	smoke
closed	open
cobweb	spider
component	part
conclude	end
consider	think
contact	touch
cork	wine
corpse	dead
correct	wrong
crescent	moon
differ	same
dilemma	problem
dine	eat
document	paper
elk	deer
finished	done
fracture	break
galaxy	stars
gossip	talk
harp	music
hawk	bird
interrupt	rude

<u>Unrelated Filler Pairs</u>

agriculture	girl
ambulance	window
ankle	farming
archer	phonebook
architect	doorbell
bank	rise
bars	lettuce
basin	academy
bearing	lamb
bike	report
brush	dices
bubble	velvet
caffeine	sky
cats	roles
ceiling	king
chef	fear
child	eagle
compass	beauty
consulate	door
cookie	nose
couch	dough
curve	lead
dancer	liquid
deer	pencil
diamond	gear
disgust	latin
drill	guest
elbow	pistol
force	elevator
garage	harbor
glass	rush
grab	screw
grass	fist
groceries	conductor

jupiter	planet	ground	disk
lily	flower	gymnastic	vegetables
lips	kiss	insurance	icecream
medium	large	justice	mousepad
meow	cat	kill	clock
nap	sleep	kite	basket
palm	hand	lemon	soccer
post	office	liar	eggs
powerful	strong	light	mouse
quill	pen	map	clown
racquet	ball	maths	landscape
request	ask	medicine	passage
sentry	guard	mile	apron
shears	scissors	miracle	ginger
shrink	small	mouth	actor
shutter	window	needle	currency
skull	head	patty	kitchen
slay	kill	penny	diabetes
sole	shoes	piano	money
spicy	hot	planter	power
strike	hit	point	queen
survive	live	priming	investor
syringe	needle	printer	angel
tax	money	pull	leather
tickle	laugh	radiation	jockey
trousers	pants	revolting	roach
truthful	honest	river	leopard
unite	together	sailor	glands
watt	bulb	switch	phone
zest	soap	tomatoes	bedtime

Appendix C

Individual subject mean RT's in the Associative task as a function of direction at each SOA

Subject	Assoc. 150 ms Forward	Assoc. 150 ms Backward	Assoc. 1000 ms Forward	Assoc. 1000 ms Backward
1	913	663	712	977
2	979	963	754	980
3	1018	1124	1079	1115
4	905	918	869	967
5	642	614	504	578
6	750	794	685	745
7	532	547	517	537
8	757	745	749	778
9	791	873	747	828
10	725	762	705	751
11	563	564	516	600
12	753	811	714	737
13	725	721	785	805
14	1103	1226	949	921
15	836	854	738	701
16	829	852	696	888
17	796	766	741	793
18	671	734	629	683
19	1053	965	962	930
20	902	862	912	882
21	882	988	875	843
22	1022	984	952	884
23	765	910	704	737
24	801	799	790	797
25	785	990	1005	948
26	769	746	755	703
27	753	937	816	751
28	760	889	881	769
29	1050	946	869	827
30	734	857	761	840
31	1026	1190	1089	1180
32	1032	1215	1312	1136
33	939	848	778	859
34	1075	894	751	946
35	1115	1135	923	1134
36	676	745	594	640
37	1042	982	786	898

38	913	871	760	792
39	806	934	783	873
40	1048	978	868	782
41	773	640	567	613
42	823	826	746	721
43	1243	1173	960	1067
44	804	762	540	1016
45	737	814	754	785
46	1234	1161	1072	1449
47	664	724	715	707
48	732	716	634	712
49	737	867	775	898
50	803	837	715	865
51	673	687	637	643
52	907	773	802	849
53	1045	947	935	1081
54	756	792	620	797
55	860	867	769	708
56	822	968	827	894
57	704	824	657	778
58	1071	842	747	1027
59	734	755	723	676
60	850	838	675	765

Appendix D

Individual subject mean RT's in the Causal task as a function of direction at each SOA

Subject	Causal 150 ms Predictive	Causal 150 ms Diagnostic	Causal 1000 ms Predictive	Causal 1000 ms Diagnostic
1	1492	2212	1402	1498
2	1644	1616	1663	1295
3	1719	1968	1391	1720
4	1628	1602	1993	1499
5	1043	939	663	767
6	1092	1243	943	980
7	874	1087	806	903
8	1045	1442	1154	1425
9	1250	1379	1070	1187
10	1075	1349	976	982
11	1007	948	852	886
12	1418	1328	1164	1301
13	927	1053	1038	1201
14	1557	1928	1180	1424
15	1562	1576	1213	1822
16	1388	1263	1190	1262
17	1178	1209	1074	1419
18	1041	1074	794	824
19	1264	1237	1296	1097
20	1951	2331	1747	1878
21	1311	1615	1163	1135
22	2675	2627	1718	1802
23	1253	1191	1013	1081
24	973	1213	882	1422
25	2408	2181	1621	2103
26	1144	1274	1130	1057
27	1216	1361	1157	1510
28	1488	1470	1214	1358
29	1430	1484	1488	1564
30	1064	1055	878	1078
31	1857	2232	1728	1413
32	1017	942	1016	1205
33	1309	1840	1219	1665
34	2026	2053	1398	1630
35	2540	2330	2312	2357
36	1332	1263	1006	1008
37	1263	1023	1060	1293

38	1419	1367	1268	1327
39	1641	2060	1713	1666
40	907	1069	893	1192
41	1177	1023	728	1129
42	1210	1040	877	837
43	1515	1940	1936	1591
44	1196	1184	1195	891
45	1051	1016	1018	1044
46	1542	1859	1510	1366
47	791	872	1045	924
48	976	1099	1068	1143
49	1012	1330	1120	1166
50	1011	1029	1011	1155
51	793	870	824	650
52	844	875	887	929
53	2597	3200	2035	3519
54	963	923	967	930
55	789	786	594	919
56	1245	1512	1208	1356
57	945	996	880	1017
58	1082	1281	1019	1151
59	1324	1463	1045	1154
60	1563	1945	2447	1762

Appendix E

Complete table of 2 x 2 x 2 omnibus ANOVA RT analysis (Factors: Task, SOA, and Direction)

Tests of Within-Subjects Effects

Measure:MEASURE_1

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
Task	Sphericity Assumed	2.892E7	1	2.892E7	118.211	.000
	Greenhouse-Geisser	2.892E7	1	2.892E7	118.211	.000
	Huynh-Feldt	2.892E7	1	2.892E7	118.211	.000
	Lower-bound	2.892E7	1	2.892E7	118.211	.000
Error(Task)	Sphericity Assumed	1.443E7	59	244614.021		
	Greenhouse-Geisser	1.443E7	59	244614.021		
	Huynh-Feldt	1.443E7	59	244614.021		
	Lower-bound	1.443E7	59	244614.021		
SOA	Sphericity Assumed	893761.909	1	893761.909	39.816	.000
	Greenhouse-Geisser	893761.909	1	893761.909	39.816	.000
	Huynh-Feldt	893761.909	1	893761.909	39.816	.000
	Lower-bound	893761.909	1	893761.909	39.816	.000
Error(SOA)	Sphericity Assumed	1324375.719	59	22447.046		
	Greenhouse-Geisser	1324375.719	59	22447.046		
	Huynh-Feldt	1324375.719	59	22447.046		
	Lower-bound	1324375.719	59	22447.046		
Direction	Sphericity Assumed	607964.718	1	607964.718	33.786	.000
	Greenhouse-Geisser	607964.718	1	607964.718	33.786	.000
	Huynh-Feldt	607964.718	1	607964.718	33.786	.000

	Lower-bound	607964.718	1	607964.718	33.786	.000
Error(Direction)	Sphericity Assumed	1061679.578	59	17994.569		
	Greenhouse-Geisser	1061679.578	59	17994.569		
	Huynh-Feldt	1061679.578	59	17994.569		
	Lower-bound	1061679.578	59	17994.569		
Task * SOA	Sphericity Assumed	176840.673	1	176840.673	8.040	.006
	Greenhouse-Geisser	176840.673	1	176840.673	8.040	.006
	Huynh-Feldt	176840.673	1	176840.673	8.040	.006
	Lower-bound	176840.673	1	176840.673	8.040	.006
Error(Task*SOA)	Sphericity Assumed	1297678.448	59	21994.550		
	Greenhouse-Geisser	1297678.448	59	21994.550		
	Huynh-Feldt	1297678.448	59	21994.550		
	Lower-bound	1297678.448	59	21994.550		
Task * Direction	Sphericity Assumed	136298.655	1	136298.655	6.096	.016
	Greenhouse-Geisser	136298.655	1	136298.655	6.096	.016
	Huynh-Feldt	136298.655	1	136298.655	6.096	.016
	Lower-bound	136298.655	1	136298.655	6.096	.016
Error(Task*Direction)	Sphericity Assumed	1319082.509	59	22357.331		
	Greenhouse-Geisser	1319082.509	59	22357.331		
	Huynh-Feldt	1319082.509	59	22357.331		
	Lower-bound	1319082.509	59	22357.331		
SOA * Direction	Sphericity Assumed	10390.662	1	10390.662	.666	.418
	Greenhouse-Geisser	10390.662	1	10390.662	.666	.418
	Huynh-Feldt	10390.662	1	10390.662	.666	.418
	Lower-bound	10390.662	1	10390.662	.666	.418

Error(SOA*Direction)	Sphericity Assumed	920213.441	59	15596.838		
	Greenhouse-Geisser	920213.441	59	15596.838		
	Huynh-Feldt	920213.441	59	15596.838		
	Lower-bound	920213.441	59	15596.838		
Task * SOA * Direction	Sphericity Assumed	26256.710	1	26256.710	1.191	.280
	Greenhouse-Geisser	26256.710	1	26256.710	1.191	.280
	Huynh-Feldt	26256.710	1	26256.710	1.191	.280
	Lower-bound	26256.710	1	26256.710	1.191	.280
Error(Task*SOA* Direction)	Sphericity Assumed	1300788.149	59	22047.257		
	Greenhouse-Geisser	1300788.149	59	22047.257		
	Huynh-Feldt	1300788.149	59	22047.257		
	Lower-bound	1300788.149	59	22047.257		