The Temporality Effect in Thinking about What Might Have Been

Ruth M.J. Byrne & Ronan Culhane

Psychology Department,
University of Dublin, Trinity College,
Dublin 2, IRELAND
rmbyrne@vax1.tcd.ie

Alessandra Tasso

Psychology Department,
University of Padua, ITALY
psico06@ipdunivx.unipd.it

Abstract

When people think about what might have been, they construct a mental representation of the actual state of affairs, and they generate an imaginary alternative by carrying out minimal mutations to it. When they think about how an undesirable outcome might have been avoided, they mutate the events leading to the outcome in regular ways, for example, they undo the more recent event in a series of independent events. We describe a computer simulation of the cognitive processes that underlie these effects of temporality on counterfactual thinking that is based on the idea that reasoners construct contextualized models. We report the results of two experiments that show that the temporality effect arises because the first event provides the context against which subsequent events are interpreted. The experiments show that when the contextualizing role of the first event is decoupled from its temporal order the effect is eliminated, for both bad and good outcomes. The results rule out an alternative explanation based on the idea that the more recent event is 'fresh' in mind. The context effect in temporal mutability may shed light on the remaining primary phenomena of counterfactual thinking.

Thinking about Imaginary Situations

Thinking about what might have been -- *counterfactual* thinking -- depends on the comparison of an event with an imaginary alternative. People think about what might have been when they mull over the past, imagining how a situation could have turned out differently, e.g.,

1. If I had chosen a 4 instead of a 5, I would have won the lottery.

They also engage in similar hypothetical thinking when they speculate about the future, e.g.:

2. If I were to buy lots of tickets, my chances of winning the lottery would be better.

Counterfactual thinking allows us to think about situations that are different from the actual one and to go beyond reflections about the actual past or the predictable future. It is closely related to conditional reasoning (e.g., Byrne, 1989a; 1989b; Byrne and Johnson-Laird, 1992) and to suppositional reasoning (e.g., Byrne, Handley, and Johnson-Laird, 1995). But, the essence of counterfactual thought is the *comparison* between an actual situation and a temporarily supposed counterfactual one (e.g, Byrne and Tasso, 1994).

Thinking about imaginary alternatives is central to higher-level cognition -- problem solving, creativity, reasoning, and decision-making (e.g., Ginsberg, 1986; Macrae, 1992; Miller and McFarland, 1987; Roese, 1994). It underlies the construction of sub-goals in problem-solving; it underlies planning and intention, and thus the concepts of change and improvement at the heart of creativity; and it underlies the search for counterexamples in reasoning, that is, the search for alternative situations in which the premises of an argument could be true but the conclusion false.

The primary purpose of counterfactual thinking is to enable us to learn. We carry out comparisons between factual and imaginary situations to work out how a chosen situation could have turned out better (or worse) and so we have the capacity to learn from our mistakes (or our good fortune). As a result, we can develop our own theories of the causes of an outcome, and how to avoid similar outcomes in the future. Indeed, counterfactual thinking may underlie causal understanding (e.g., Chisholm, 1946; Goodman, 1973; Mackie, 1973; Skyrms, 1980), and causal attributions (e.g., McGill and Klein, 1993; Wells and Gavanski, 1989), and it is crucial to scientific, legal, and social progress (for a review see Kahneman and Miller, 1986). Moreover, counterfactual thinking gives rise to emotional experiences such as regret, blame, guilt, and worry, as well as hope, surprise, relief, and wonder (e.g., Landman, 1987; Johnson, 1986). These emotions emerge from the *comparison* of an actual situation to an imaginary alternative. The functions of counterfactual thinking are varied and wide-spread, and we suggest that its key purpose is to release the human mind from the limitations of facts. What the mind does when it engages in counterfactual thinking is to compare facts to their imaginary alternatives.

Two key questions have been addressed in trying to understand counterfactual thinking. One concerns how we make inferences from imaginary situations, and in particular, whether counterfactual reasoning is similar to factual reasoning (e.g., Byrne and Tasso, 1994; Lewis, 1973; Stalnaker, 1968). The other concerns how we generate imaginary scenarios (e.g., Kahneman and Tversky, 1982; Kahneman and Miller, 1986). It is this second question that we focus on in the current paper.

When people undo an actual situation to create an imaginary one, their mutations are systematic and minimal (e.g., Pollock, 1986). When they think about how an outcome could have turned out differently, they tend to undo exceptional features, that is, actions that are unusual for the specified person rather than actions that are routine for them (e.g., Kahneman and Tversky, 1982), although the exceptionality effect is modulated by how exceptional the outcome itself is (e.g., Gavanski and Wells, 1989; Bouts, Spears, and Van der Pligt, 1992). People tend to undo the first cause in a causal sequence that led to the outcome rather than subsequent causes (e.g., Wells, Taylor, and Turtle, 1987). They tend to undo actions that a person carried out, e.g. switching stock to a (subsequently deteriorating) company, rather than inactions, e.g., leaving stock in a (subsequently deteriorating) company (e.g., Kahneman and Tversky, 1982; Landman, 1987). They undo events under an individual's voluntary control, e.g., stopping at a pub for a drink, rather than events outside their control, e.g., a fallen tree in the road (e.g., Girotto, Legrenzi, and Rizzo, 1991). Finally, they tend to undo the most recent action or event rather than an earlier one in an independent sequence of events (Miller and Gunasegaram, 1990).

Most researchers agree that the mutability of an element in a representation -- how easy it is to undo that aspect of the representation -- depends on how easy it is to think of alternatives to it. The question then becomes, what determines how easy it is to think of an alternative to a particular element in a representation? Our aim in the current paper is to outline a theory of counterfactual thinking that answers this question. We will sketch a computational simulation that captures the spirit of the theory, and report the results of two experiments that test its claims. We will describe the theory with reference to Miller and Gunasegaram's temporality effect, and so it is to a more detailed description of this effect that we now turn.

Temporal Mutability

Consider the following scenario (from Miller and Gunasegaram, 1990, p. 1111):

3. Imagine two individuals (Jones and Cooper) who are offered the following very attractive proposition. Each individual is asked to toss a coin. If the two coins come up the same (both heads or both tails), each individual wins £1,000. However, if the two coins do not come up the same, neither individual wins anything. Jones goes first and tosses a head; Cooper goes next and tosses a tail. Thus, the outcome is that neither individual wins anything.

When asked to undo the outcome, more than 80% of subjects agreed that the alternative of Cooper tossing a head came more readily to mind than Jones tossing a tail (Miller and Gunasegaram, 1990). They also judged that Cooper would experience more guilt, and would tend to be blamed more by Jones. Logically, of course, neither party should be considered more mutable, or more likely to experience any more guilt or desire to blame than the other, because the event is one of chance. The temporality effect has been postulated to play a role in many everyday judgements, such as the tendency for blackjack players to be averse to playing on the last box, the tendency for teams to sport their faster runner last in a relay race, and for people to wager more on their predictions than their postdictions (Miller and Gunasegaram, 1990).

We suggest that temporality affects mutability because of the mechanisms that underlie the construction of mental representations of actual and imaginary situations. What cognitive mechanisms could give rise to the temporality effect? We suggest that earlier events in a sequence provide the context against which subsequent events are interpreted. When reasoners construct a model of a situation, they try to represent as little as possible because of working memory constraints, and so some information is represented explicitly and some information is represented only implicitly (see Johnson-Laird and Byrne, 1991). We suggest that in scenarios such as the coin-toss one, the earlier event in a sequence *initializes* the model, that is, it identifies what the model is about and provides the cornerstone of its foundation. In this way, the earlier event in a sequence provides the context that ensures the coherency of the model. The contextualizing event in a model may configure and constrain the subsequent elements that are inserted into it. In everyday thinking, people's models may be continually changing to deal with new situations, and the cornerstone of a new situation may initialize a new model. Because the early event in the coin-tossing scenario provides the foundation stone of the model, it cannot be mutated without weakening the entire structure. For example, when reasoners attempt to undo the outcome that neither individual won any money in the coin-tossing scenario, they could flesh-out their models to be consistent with a number of counterfactual scenarios, e.g, Jones could have tossed tails instead of heads; or alternatively, Cooper could have tossed heads instead of tails. But, the temporality effect shows that people tend to flesh-out their models to be consistent with just one of the options: the one in which the contextualizing element of the model -- the first event -- is left untouched.

We suggest that it is possible to de-couple the contextualizing role of the first event from its position in the sequence. If our explanation based on contextualized models is correct, then a separate context event prior to the two target events should result in both the first and second target event being perceived to be equally mutable. We constructed scenarios based on the following sort of content:

4. Imagine two individuals (Jones and Brady) who take part in a television game show, on which they are offered the following very attractive proposition. Each individual is given a shuffled deck of cards, and each one picks a card from their own deck. If the two cards they pick are of the same colour (i.e. both from black suits or both from red suits) each individual wins £1,000. However, if the two cards are not the same colour, neither individual wins anything.

Jones goes first and picks a black card from his deck. At this point, the game-show host has to stop the game because of a technical difficulty. After a few minutes, the technical problem is solved and the game can be restarted. Jones goes first again, and this time the card that he draws is a red card. Brady goes next and the card that he draws is a black card. Thus, the outcome is that neither individual wins anything.

The technical hitch device allows us to manipulate the stage-setting independently of the plays of the first and second player. According to our explanation, the first players pre-hitch choice provides the context against which the subsequent events are interpreted. Either of the post-hitch events is an initial candidate for mutability. Reasoners compute the similarity of the events to the context: each of the post-hitch events share some properties with the context and differ from it on others -- the first event shares the same player with the context, but differs in the colour of the card drawn; the second event shares the colour of the card drawn with the context, but differs in the identity of the player (and for simplicity, we assume that the player and the colour of the card drawn are equally salient). Because the events are both dissimilar from the context, either one is mutable. We predict that the temporality effect will be eliminated for this sort of scenario. In the next section we report computational and experimental evidence in support of this view.

The Production of Imaginary Models

We have written a computer program in LISP to simulate the cognitive processes that we suggest reasoners rely on when they think counterfactually, in the colour-card task described in 4 earlier. The program is called **IMP** (for Imaginary Models Production), and it generates imaginary scenarios. It identifies the most mutable event by comparing the events to the context, and it undoes the event by calibrating it to the context. The program takes as input a list of assertions of the following sort:

5. Jones picked black
   There was a technical hitch
   Jones picked black
   Brady picked red

It produces as output a set of models of the actual situation:

6. Jones black / Jones black Brady red

where the symbol "/" is a language-like tag corresponding to a technical hitch (see Johnson-Laird and Byrne, 1991, for comments on the use of propositional tags in models). It also produces a set of models corresponding to the counterfactual situation constructed by mutating the most mutable event:

7. Jones black / Jones black Brady black

The core of the program lies in a suite of mutability functions that identify the most mutable event by comparing it to the context event, and a suite of undoing functions that construct a counterfactual scenario by calibrating the most mutable event to the context. The main computation carried out by the mutability procedures is a comparison of each of the events with the context. Each event is compared with the context in terms of whether they share a common actor and whether they share a common card colour. The most mutable event is selected according to the following two principles: if one event is identical to the context, the other event is the most mutable; whereas if both events are dissimilar from the context they are both mutable (for further details, see Byrne, Culhane, and Tasso, 1995). For example, when the program is given the sentences in 5 earlier, it identifies the second event, in which Brady picked red, as the most mutable. Each of the two events is compared to the context: the first event is identical to the context and so it is immutable, the second event differs from the context and so it is the most mutable.

The main computation carried out by the undoing procedures is a comparison of the most mutable event to the context in order to alter it according to the following two principles: if they differ in the colour of the card picked (and so the actual scenario must be one in which the players lost), the most mutable event is made the same as the context in this respect (so that the counterfactual scenario is of a winning game); otherwise (the actual scenario is one in which the players won), the most mutable event is made different from the context (so that the counterfactual scenario is of a losing game). For example, the program produces a counterfactual scenario in which Brady picked black instead of red for the sentences in 5 earlier. It compares the colour of the card that Brady picked in the game with the colour of the card in the context event and because they differ, it changes Brady's choice to be the same as the context.

The program is designed to simulate the known effects of temporal order on counterfactual mutations. It produces the temporal order effect when it is given sentences in which there is no technical hitch: it identifies the first event as being the context, as well as being the first target event. It implements our ideas on the mental representations and cognitive processes that lead to the effects of temporality on mutability. The theory embodied in the computer program makes predictions about a variety of situations in which the temporality effect will occur and situations in which it will be eliminated. The program simulates these effects and we have corroborated the predictions empirically as we will now outline.

Contextualized Models and Temporal Mutability

We have carried out a series of experiments to test the role of contextualized representations in temporal mutability and we report the results of two experiments here. The first experiment was designed to test two alternative explanations of the role of temporality in the mutability of events. Our suggested explanation emphasises the nature of the mental representations that reasoners construct: the earlier event provides the context against which the subsequent events are interpreted, and so the earlier event is not available for mutation. The contextualizing role of the first event can be separated from its temporal position in the sequence by the 'technical hitch' scenario described earlier. We predict that the temporality effect will occur when the first target (post-hitch) event is the same as the context (pre-hitch) event (i.e., Jones draws black on each occasion), and we predict that the effect will be eliminated when the first target event is different from the context event (i.e., Jones draws black prior to the hitch, and draws red the next time).

We compared our explanation to an alternative one. An alternative possibility is that the more recent event is "fresh" in mind: it is encountered first in a backward search through the entries to the representation. Such a last-in, first-out principle may operate because of working memory constraints. A different set of predictions follows from this view. The second event is the more recent one in both of the technical hitch scenarios, regardless of the similarity or dissimilarity of the first target event to the context event, and so the temporality effect should be observed in each version.

We tested these alternative explanations in an experiment in which 75 undergraduate students from various departments in Trinity College, University of Dublin participated voluntarily. We gave the *different-context* version of the scenario described in 4 earlier to a group of 36 subjects: In this version, Jones draws black prior to the hitch, and he draws red after the hitch is resolved; Brady draws black. We gave the *same-context* version to another group of 39 subjects: In this version Jones draws black prior to the hitch, and he draws black after the hitch is resolved; Brady draws red. The subjects completed the sentence "Jones and Brady could each have won £1000 if only one of them had picked a different card, for instance if..." They also answered two questions about who they would predict would experience more guilt, and who will blame the other more. We expected that their beliefs about the players emotions would correspond to their beliefs about the mutability of the individuals.

The results corroborated our proposal that the temporality effects arise because the first event provides the context against which the subsequent events are interpreted (see Byrne, et al, 1995, for details). As Table 1 shows, in the same-context condition more subjects undid the second event overall (59%)[1] rather than the first (23%); whereas in the different-context condition the effect was eliminated (44% versus 42%) and this interaction is reliable [Meddis (1984) quicktest z' (n = 63) = 1.64, $p < 0.05$]. We found a similar pattern of results for the answers to the questions about emotions. More subjects attributed guilt to the second player than to the first in the same-context version (77% versus 10%), but not in the different-context version [44% versus 31%; Meddis quicktest z' (n = 61) = 2.59, $p < 0.01$]. More subjects considered that the first player would blame the other more than the second in the same-context version (51% versus 13%); unexpectedly, more subjects also considered the first player to blame the other more in the different-context version [50% versus 25%; Meddis quicktest z (n = 52) = 1.07, p = non-significant]. [However, in a subsequent experiment based on related materials, the elimination of the temporality effect was observed for questions about both guilt and blame, as well as for the sentence completion task (see Byrne et al, 1995)].

The experiment shows that the temporality effect arises because of the contextualizing role of the first event in the construction of a mental representation, and it is eliminated when the context is de-coupled from the serial position of the event. The results go against the alternative explanation that the temporal order effect arises because of the recency of the second event's entry into the representation. The results lend support to our suggestion that reasoners construct minimal models that are based on a solid foundation, the cornerstone laid by the earlier events in a sequence.

Table 1: The percentages of mutations of the first event and the second in Experiment 1.

Context	*Same* Black..hitch; Black..Red	*Different* Black..hitch; Red..Black
Undoing		
Second	59	44
First	23	42
Guilt		
Second	77	44
First	10	31
Blame		
Second	13	25
First	51	50

Temporal Mutability and Good Outcomes

Previous research has shown that various phenomena, such as the tendency to mutate exceptional events rather than normal ones, occur as readily for undoing positive outcomes as for negative ones (see also Landman, 1987; Johnson, 1986). Counterfactual thinking is not only for learning from mistakes but also for learning from our good fortunes. Our aim in the next experiment was to provide the first test of whether temporality affects situations of good fortune as well as situations of bad fortune. We relied on a scenario similar to that in the first experiment, but in this version, the individuals won.

We predict that the temporality effect will be observed with situations of good fortune: according to our view reasoners compare the two target events to the context in terms of the similarity of both the actor and the colour of the card drawn. For example, in one of the scenarios in the experiment, Jones draws red prior to the hitch, and he draws

[1] The percentages are based on a combined score, e.g., 59% is based on the undoings of the second event *only* (i.e., "if Brady drew black") combined with the undoings of the second event followed by the first event (i.e., "if Brady drew black or Jones drew red") -- see Byrne, et al (1995). The pattern of results remains essentially the same when the analysis is based on the former alone.

red after the hitch is resolved; Brady draws red. The first event is identical to the context in both actor and card colour; the second event differs in terms of actor, and so it is more mutable.

An alternative account makes a different prediction. A possible explanation for the temporality effect and its elimination in the previous experiment is that reasoners mutate the 'odd-one-out' in a sequence. Reasoners undo whichever of the two target events is the odd-one-out of the three events in terms of the colour of the card drawn, in other words card-colour is the salient comparison-point in these scenarios. According to this view, of course, the temporality effect should only occur in situations where there *is* an odd-one-out, and so there should be no temporality effect in situations where the individuals win (because the two target events contain the same colour card). Our first aim in the experiment was to test these two alternative explanations. On the assumption that temporality would affect mutability in situations of good fortune, our second aim was to show that the temporality effect can be eliminated in the manner established in the previous experiment, when both target events differ from the context.

We tested these alternative explanations in an experiment in which 94 subjects from Trinity College, University of Dublin participated voluntarily. We presented one group of 46 subjects with a scenario in which Jones draws red prior to the hitch, and he draws red after the hitch is resolved; Brady draws red. Thus, the outcome is that each individual wins £1000. We gave a second group of 48 subjects a version of the story in which Jones draws black prior to the hitch, and he draws red after the hitch is resolved; Brady draws red. In this case both events are dissimilar from the context: the first event differs in the colour of the card, and the second event differs in both the colour of the card and the actor. We expect that both events are mutable and so there should be no temporality effect. The subjects completed the following fragment: "After the draw both Jones and Brady reflected on how lucky they had been. After all, if one of them had picked a different card they might neither have won the £1000; for instance, if..." They also answered the question "Who would you predict would experience more relief at having won -- Jones or Brady?"

Table 2: The percentages of mutations of the first event and the second in Experiment 2.

Context	*Same* Red..hitch; Red..Red	*Different* Black..hitch; Red..Red
Undoing		
Second	59	33
First	26	38
Relief		
Second	59	42
First	35	44

The results corroborated both our predictions, as Table 2 shows. In the same-context condition more subjects undid the second event (59%) rather than the first (26%); whereas in the different-context condition the effect was eliminated [33% versus 38%; Meddis quicktest z' $(n = 73) = 1.9$, $p < 0.05$]. More subjects attributed relief to the second player than to the first in the same-context version (59% versus 35%), but not in the different-context version (42% versus 44%), although the interaction misses reliability [Meddis quicktest z $(n= 84) = 1.29$, p = non-significant].

The experiment shows that temporality affects mutability in situations of good fortune as well as situations of bad fortune. It corroborates our suggestion that reasoners compare the target events to the context event in terms of both the actor and the colour of the card, and not just in terms of the colour of the card, as the odd-one-out explanation suggests. It also shows that the effect can be eliminated readily by de-coupling the contextualizing role of the first event from its position in the sequence.

Conclusions

We have outlined a theory of counterfactual thinking about what might have been that is based on the idea that reasoners construct a model of the actual situation and produce an imaginary alternative by making minimal mutations to it. Our theory suggests that the temporality effect on mutability -- the tendency for people to undo the more recent event in a series of independent events -- arises because of the nature of the mental representations that people construct. We suggest that the first event in a sequence plays a contextualizing role, providing the cornerstone for the model, and so it is immutable. We have implemented this theory in a computer program and we have carried out several experiments that show that when the contextualizing role of the first event is separated from its position in the sequence of target events, the effect is eliminated. We have corroborated our view that the mutability of events depends on cognitive mechanisms such as the contextualization of models for one sort of problem in this series of experiments. We suggest that the influence of context in the construction of models of counterfactual scenarios may have implications for our understanding of the other primary phenomena of counterfactual thinking.

Acknowledgements

This research was supported by a grant from the Arts and Social Sciences Benefactions Fund of Trinity College, University of Dublin awarded to Ruth Byrne. We are grateful to Mark Keane, Phil Johnson-Laird, and Susana Segura-Vera for their help.

References

Bouts, P., Spears, R., and van der Plight, J. (1992). Counterfactual processing and the correspondence between events and outcomes. *European Journal of Social Psychology,* 22, 387-396.

Byrne, R.M.J. (1989a). Suppressing valid inferences with conditionals. *Cognition*, 31, 61-83.

Byrne, R.M.J. (1989b). Everyday reasoning with conditional sequences. *Quarterly Journal of Experimental Psychology,* 41A, 141-166.

Byrne, R.M.J. and Johnson-Laird, P.N. (1992). The spontaneous use of propositional connectives. *Quarterly Journal of Experimental Psychology,* 44A, 89-110.

Byrne, R.M.J., Handley, S.J., and Johnson-Laird, P.N. (1995). Reasoning with suppositions. *Quarterly Journal of Experimental Psychology.* In press.

Byrne, R.M.J. and Tasso, A. (1994). Counterfactual reasoning: inferences from hypothetical conditionals. In A. Ram and K. Eiselt (Eds). *Proceedings of the sixteenth annual conference of the Cognitive Science Society.* Hillsdale: Erlbaum.

Byrne, R.M.J., Culhane, R. and Tasso, A. (1995). *Cognitive Processes in Imaginary Thinking: Mutability and Temporality in Counterfactual Scenarios.* Mimeo, Trinity College, University of Dublin, Ireland.

Chisholm, R.M. (1946). The contrary-to-fact conditional. *Mind,* LV, 289-307.

Gavanski, I. and Wells, G.L. (1989). Counterfactual processing of normal and exceptional events. *Journal of Experimental Social Psychology,* 25, 314-325.

Ginsberg, M. L. (1986). Counterfactuals. *Artificial Intelligence,* 30, 35 - 79.

Girotto, V., Legrenzi, P., and Rizzo, A. (1991). Event controllability in counterfactual thinking. *Acta Psychologica,* 78, 111-133.

Goodman, N. (1973). *Fact, Fiction and Forecast.* Third Edition. New York: Bobbs-Merrill.

Johnson, J. (1986). The knowledge of what might have been: affective and attributional consequences of near outcomes. *Personality and Social Psychology Bulletin,* 12, 51-62.

Johnson-Laird, P.N. and Byrne, R.M.J. (1991). *Deduction.* Hove, UK, and Hillsdale, N.J.: Erlbaum.

Kahneman, D. and Miller, D. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review,* 93, 136-153.

Kahneman, D., and Tversky, A. (1982). The simulation heuristic. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), *Judgement under uncertainty: heuristics and biases.* (pp. 201-208). New York: Cambridge University Press.

Landman, J. (1987). Regret and elation following action and inaction: affective responses to positive versus negative outcomes. *Personality and Social Psychology Bulletin,* 13, 524-536.

Lewis, D. (1973). *Counterfactuals.* Oxford, Blackwell.

Mackie, J.L (1973). *Truth, probability, and paradox.* Oxford: Clarendon.

Macrae, C. N. (1992). A tale of two curries: counterfactual thinking and accident-related judgements. *Personality and Social Psychology Bulletin,* 18, 84-87.

McGill, A.L. and Klein, J.G. (1993). Contrastive and counterfactual reasoning in causal judgement. *Journal of Personality and Social Psychology,* 64, 897-905.

Meddis, R. (1984). *Statistics using ranks.* Oxford: Basil Blackwell.

Miller, D.T. and Gunasegaram, S. (1990). Temporal order and the perceived mutability of events: implications for blame assignment. *Journal of Personality and Social Psychology,* 59, 1111- 1118.

Miller, D.T. and McFarland, C. (1987). Counterfactual thinking and victim compensation: a test of norm theory. *Personality and Social Psychology Bulletin,* 12, 513-519.

Pollock, J.L. (1986). *Subjunctive Reasoning.* Dordrecht: Reidel.

Roese, N.J. (1994). The functional basis of counterfactual thinking. *Journal of Personality and Social Psychology,* 66, 805-818.

Skyrms, B. (1980). *Causal Necessity.* New haven: Yale University.

Stalnaker, R. C. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in logical theory.* Oxford: Basil Blackwell.

Wells, G.L., and Gavanski, I. (1989). Mental simulation of causality. *Journal of Personality and Social Psychology,* 56, 161-169.

Wells, G.L., Taylor, B.R., and Turtle, J.W. (1987). The undoing of scenarios. *Journal of Personality and Social Psychology,* 53, 421-430.