

# How temporal assumptions influence causal judgments

YORK HAGMAYER and MICHAEL R. WALDMANN  
*University of Göttingen, Göttingen, Germany*

Causal learning typically entails the problem of being confronted with a large number of potentially relevant statistical relations. One type of constraint that may guide the choice of appropriate statistical indicators of causality are assumptions about temporal delays between causes and effects. There have been a few previous studies in which the role of temporal relations in the learning of events that are experienced in real time have been investigated. However, human causal reasoning may also be based on verbally described events, rather than on direct experiences of the events to which the descriptions refer. The aim of this paper is to investigate whether assumptions about the temporal characteristics of the events that are being described also affect causal judgment. Three experiments are presented that demonstrate that different temporal assumptions about causal delays may lead to dramatically different causal judgments, despite identical learning inputs. In particular, the experiments show that temporal assumptions guide the choice of appropriate statistical indicators of causality by structuring the event stream (Experiment 1), by selecting the potential causes among a set of competing candidates (Experiment 2), and by influencing the level of aggregation of events (Experiment 3).

Assumptions about temporal intervals between events guide our causal judgments in science and in everyday life. Suppose, for example, that we are suffering from a severe headache and take a new medicine that promises relief. How long are we going to wait until we swallow more of this drug because the first dose did not help? We will probably base our decision on assumptions about the expected delay between taking the drug and experiencing an effect. If we assume that it would take 30 min and 1 h has already passed without any effect, we will probably conclude that the drug was not effective and, possibly, will decide to increase the dose. However, if we assume that it will take more than 1 h, we may wait. Therefore, the same fact may lead to very different causal judgments and actions.

## Experienced Versus Described Learning Events

Previous research on the role of time in learning has focused on the learning of events that are being observed in real time. The research on the role of spatiotemporal contiguity is a prime example. A large number of studies have shown that increasing the temporal delay between cause and effect tends to weaken the impression of a causal relation (Bullock, Gelman, & Baillargeon, 1982; Leslie & Keeble, 1987; Mendelson & Shultz, 1976; Michotte, 1963;

Schlottmann, 1999; Schlottmann & Shanks, 1992; Shultz, 1982; Shultz & Kestenbaum, 1985; Siegler, 1976; Siegler & Liebert, 1974). The focus on the learning of events that are being observed in real time is probably due to the fact that learning theories are rooted in the tradition of animal studies, which do not allow for alternative presentation modes.

In contrast, human learning has more flexibility. Besides observation of events in real time, there is also the possibility to learn and reason on the basis of verbally or graphically *described* events.<sup>1</sup> In fact, most studies on human causal induction use materials in which events are being described rather than directly experienced (e.g., disease classification tasks). Described events are also presented in spatiotemporal arrangements. However, the order and the timing of the presentation hardly ever mirrors the events that are being described. If participants in an experiment, for example, receive the information that a 50-year-old patient suffers from cancer 3 sec after they receive the information that he was exposed to asbestos when he was a 10-year-old child, the 40-year delay is causally relevant, whereas the 3-sec delay between the consecutive screens presenting the trial is irrelevant. One of the main goals of our experiments was to demonstrate that assumptions about the temporal characteristics of the events that are being described may also affect learning, even when these temporal relations are not directly observed.

## The Role of Temporal Assumptions in Theories of Causal Learning

The dominant approach to causal learning is the view that causal knowledge is a result of association processes. *Associative theories* model causal learning as the acquisition of associative weights between cues and outcomes,

---

Experiments 2 and 3, reported in this article, were planned and conducted at the Department of Psychology of the University of Tübingen, Germany. The research was in part funded by Grant Wa 621/5-2,3 from the Deutsche Forschungsgemeinschaft. We thank S. Sloman and P. White for their helpful comments. Correspondence concerning this article should be sent to Y. Hagmayer, Department of Psychology, University of Göttingen, Gosslerstr. 14, 37073 Göttingen, Germany (e-mail: york.hagmayer@bio.uni-goettingen.de).

regardless of whether the cues represent causes and the outcomes effects or vice versa (see Waldmann, 1996). In many of these theories, the important role of time in learning has been acknowledged, but they refer to the timing of cues and outcomes, and not to the timing of the causal events to which the cues and the outcomes refer (see, e.g., Barnet & Miller, 1996; Blaisdell, Denniston, & Miller, 1998; Gallistel & Gibbon, 2000; Schmajuk, Lamoureux, & Holland, 1998). This means that these theories blur the distinction between experienced and described events. Consequently, a number of researchers working within the framework of associationism have claimed that human learning is insensitive to the fact that causes temporally precede their effects (e.g., Matute, Arcediano, & Miller, 1996; Price & Yates, 1993; Shanks & Lopez, 1996).

In a second class of theories, the *mechanism view* (e.g., Ahn, Kalish, Medin, & Gelman, 1995; Koslowski, 1996; Mendelson & Shultz, 1976; Schlottmann, 1999), it is claimed that knowledge about mechanisms is the basis for causal reasoning. Since mechanisms are temporally ordered sequences of events, this theory is, in principle, sensitive to temporal aspects. One problem with mechanism-based theories is that specific knowledge about mechanisms may not always be available. Nevertheless, we may have intuitions about the temporal delay of an effect. Knowledge about temporal delays does not require knowledge about the mechanisms that generate them. For example, we know how long the booting process usually takes on our computer, although we may not have any further knowledge about the underlying technical details.

A third class of models of causal learning postulates that causal strength estimates are based on statistical indicators, such as contingencies, that express the degree of covariation between events (e.g., Allan & Jenkins, 1983; Cheng & Novick, 1990, 1992; Shaklee & Tucker, 1980). *Contingency* is defined as the difference between the probability of the effect in the presence of the cause [ $P(e|c)$ ] and the probability of the effect in the absence of the cause [ $P(e|\sim c)$ ]; Cheng & Novick, 1992; Eells, 1991; Suppes, 1970; Waldmann, 1996].<sup>2</sup> Covariations between causes and effects are symmetric and, therefore, share the same problems as associative weights: They are not inherently sensitive to temporal aspects of causality (see Shanks, 1993; Waldmann, 1996).

However, some more recent theories (e.g., causal-model theory; Waldmann, 1996) have integrated the assumption of causal priority by assuming a level of causal models that consists of hypothetical networks of directed cause-effect relations and a level of statistical indicators that serve as empirical evidence. Causal-model theory explicitly acknowledges the necessity of an interaction between prior assumptions about causal events and the processing of the learning input. Therefore, causal-model theory is prepared to distinguish between experienced and described events, since it clearly separates cues and outcomes (the learning events experienced in real time) from the causal model that is generated independent of the

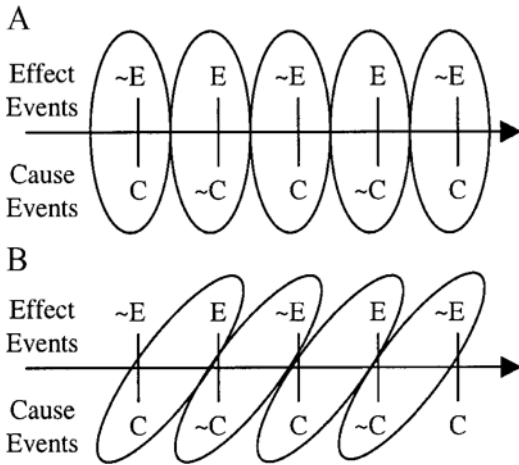
temporal ordering of learning events. The causal model includes information about causal directionality and may also contain knowledge about the temporal dynamics of the described events. However, thus far, apart from the incorporation of the temporal priority of causes, the role of time has not been further analyzed within this approach. Anderson's (1990) rational theory of causal inference is another example of a more recent view that models the influence of temporal cues on causal inferences. This approach includes a temporal parameter in the statistical estimation procedure; therefore, it can model the influence of different temporal assumptions. However, so far, little experimental work has been presented within this framework that addresses the role of temporal cues.

### Estimating Causal Strength

Most theories of causal learning have focused on how estimates of causal strength are acquired. However, these theories have typically neglected the question of how the events are selected that enter causal strength assessments. In everyday life, we are typically confronted with complex situations in which a large number of statistical relations are potentially available. Even in relatively simple cases with only three events, dozens of statistical relations could potentially be considered (see also Waldmann & Martignon, 1998).<sup>3</sup> Given that it is practically impossible for an information-processing system with limited capacity to consider all possible statistical relations between observed events, it is necessary to constrain the space of possible relations. Adopting the framework of causal-model theory, we assume instead that people typically enter a causal-learning task with assumptions about the causal situation they are going to learn about (Waldmann, 1996). These assumptions determine which contingencies are assessed to estimate causal strength. Sometimes these assumptions may be very specific (e.g., about concrete mechanisms), but often they are abstract and refer to causal structures. Previous work has shown that simply making assumptions about what is a potential cause and what is an effect may have a powerful influence on the outcome of learning even when no further knowledge about the processes connecting these events is available (see Waldmann, 2000, 2001; Waldmann & Hagmayer, 2001; Waldmann & Holyoak, 1992; Waldmann, Holyoak, & Fratianne, 1995). Temporal assumptions are an example of more concrete assumptions that might guide causal learning and influence causal judgments.

### The Influence of Temporal Assumptions

The aim of the present paper is to investigate the influence of temporal assumptions on causal judgments. How can temporal assumptions influence causal judgments? We investigated three possible ways. First, temporal assumptions can structure sequences of events into pairs of related events. Second, they can select events as possible causes for an effect. Third, they can influence the aggregation level on which the data are analyzed.



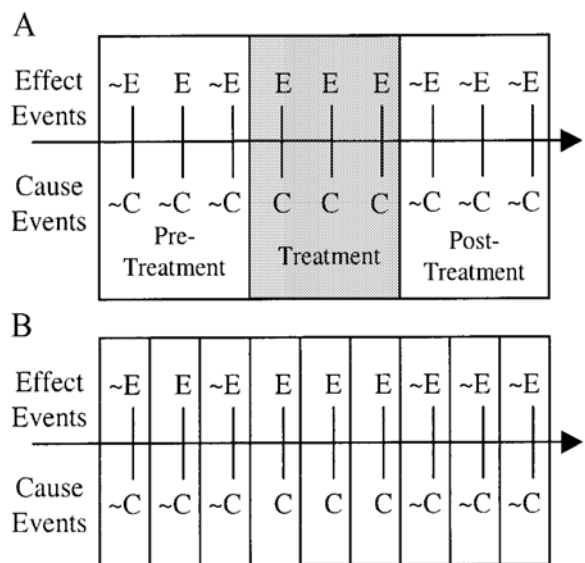
**Figure 1.** Influence of temporal assumptions on the structuring of a sequence of events. Panel A shows related events when no temporal delay of the effect is assumed; panel B shows the same events when a time lag of  $\Delta t = 1$  is expected.

Temporal assumptions refer to expected lags between causes and their effects. Thus, in a continuous sequence of events, these assumptions may combine potentially related events. Whether or not these events are causally related will still depend on the observed contingency between these events, but prior assumptions will determine which events will be involved in contingency assessments. Figure 1 provides an example that illustrates how different assumptions about causal lags may lead to different contingency assessments. Panel A depicts the result of relating simultaneously occurring events. The effect is never present when the cause occurs, but it is always present when the cause is absent. Thus, the cause seems to *prevent* the effect. However, if the same sequence of events is structured on the basis of the assumption that the effect occurs with a lag of one time unit, different event pairs are formed (Panel B). Here, the effect always occurs after the cause and never occurs when the cause is absent. Thus, the data now indicate a generative cause that deterministically produces the effect. In both cases, *all* events are used in the contingency assessments. However, the temporal assumptions about causal lags determine which events are paired together and, consequently, affect the causal assessments. In Experiment 1, this kind of influence was investigated.

Temporal assumptions may also guide the selection of possible causes. Explaining a certain event requires a selection of causal candidates from the set of preceding events. Since this set is potentially infinite, it is practically impossible to consider every possible candidate. Temporal assumptions may provide useful clues. For example, imagine that an effect, *E*, is present that might be caused by two different events,  $C_1$  and  $C_2$ . Both events,  $C_1$  and  $C_2$ , occurred with a time lag  $\Delta t = 2$  prior to the effect, *E*. Without any assumptions about temporal lags, both causal events,  $C_1$  and  $C_2$ , could have caused *E*. A decision is im-

possible because both are confounded with each other. However, if we assume that causes of the type  $C_1$  normally produce *E* with a time lag  $\Delta t$  of 1 and that causes of the type  $C_2$  produce *E* with a time lag of 2, a decision is possible. Probable causes are instances of  $C_1$  that occurred  $\Delta t = 1$  and instances of  $C_2$  that occurred  $\Delta t = 2$  prior to the effect, *E*. In Experiment 2, the influence of this type of temporal assumption on the selection of potential causes was investigated.

Temporal assumptions not only may refer to lags between causes and effects, but also may specify how an observed stream of events should be segmented into relevant event categories. Typical examples are medical or psychotherapeutical intervention studies in which researchers often decide to globally compare a pretreatment stage with a posttreatment stage, instead of looking at finer grained relationships between events occurring within these stages. Figure 2 provides an example of how these two levels of temporal aggregation may lead to different conclusions. In the fictitious situation depicted in Figure 2, a treatment, the potential cause, is administered several times during a specific period of time, the treatment phase. The treatments may refer to individual doses of a drug that is being tested. Neither before nor after this phase is the treatment delivered. A global analysis of the effect of the treatment would focus on the comparison between the pre- and the posttreatment phases, which, in this example, indicates a negative result (Figure 2A). The treatment does not increase the probability of occurrence of the effect event (i.e., the absence of the symptom), as compared with its probability in the pretreatment phase. If anything, this probability decreases. However, an analysis that disregards these global phases and focuses on short-term effects of single doses of the drug will lead



**Figure 2.** Influence of temporal assumptions on the aggregation of data. Panel A represents a global analysis; panel B represents a local analysis of the efficacy of the cause.

to the opposite conclusion (see Figure 2B). According to this analysis, the effect always occurs after the treatment has been administered but only occurs once when no treatment has been delivered. In Experiment 3, the effect of this type of temporal assumption was investigated.

In summary, in the following three experiments, the role of temporal assumptions on causal judgments about described events was investigated. Temporal assumptions were manipulated by means of differential initial instructions. The tasks referred to situations for which no specific knowledge about mechanisms was available. In general, in all the experiments, identical learning data were given to the participants. Information about temporal relations was given in written form or was graphically depicted. This procedure was used in order to eliminate possible interactions between experienced temporal delays and assumed temporal relations. Since all the participants received identical data, the statistical patterns potentially available were identical. However, adopting the framework of causal-model theory, we expected that the learners would tend to use prior temporal assumptions to constrain the selection of statistical indicators of causal strength.

## EXPERIMENT 1

The goal of Experiment 1 was to investigate whether temporal assumptions are used to structure sequences of events. In particular, we tested whether learners extract statistical relations according to prior assumptions about temporal lags between causes and effects. We expected that the learners would first pair events that, in their opinion, belonged together and then would base their causal judgments on estimates of covariation between the paired events. It is important to note that the prediction was not that prior assumptions would lead to biases in covariation judgments or to a differential neglect of specific events. It was rather expected that *all* events would enter the stage of contingency assessment.

To test these predictions we gave the participants the task of assessing how two different types of intervention affected a frequently occurring mosquito plague. All the participants received identical data for both interventions. The crucial manipulation involved participants' assumptions about temporal delays. The participants were told that, in one community, an insecticide was used to kill the mosquito larvae and that, in another community, a flower was planted that mosquito-larvae-eating beetles needed in order to breed. The causal mechanisms were not explained in greater detail, and the participants did not receive any information about the temporal delay between each type of intervention and its effect. Nevertheless, we assumed that the participants would associate an insecticide with a fast mechanism, because insecticides typically degrade fast. Thus, we expected the participants who held this assumption to assess the effect of the insecticide by estimating the contingency between the insecticide and the amount of the mosquitoes within the same year (short-delay assumption). By contrast, for the second type of intervention, we expected the participants to assume that the

plants first needed to grow before they become effective and that the beetles needed time to breed. Thus, it seemed reasonable to expect that many of the participants would pair the respective intervention with the presence or absence of a mosquito plague in the year subsequent to planting (long-delay assumption).

The structure of the learning data (which was identical for both interventions) allowed us to dissociate the two assumptions. If the events in the same years were paired, the contingency turned out to be slightly negative, which would indicate that the intervention apparently had no positive effect on the mosquito plague. By contrast, if the events were paired that were separated by a lag of 1 year, a (modest) positive effect of the intervention would be found.

## Method

**Participants and Design.** The participants were 32 students from the University of Göttingen, who were tested in a within-subjects design with short- versus long-delay assumptions as the contrasted conditions. The sequence of the two conditions was counterbalanced across participants.

**Procedure and Materials.** The task in this experiment was to assess the causal efficacy of two types of intervention to reduce mosquito plagues. The participants received written instructions (in German) in which they were told that two communities in the same area had tried to fight the regular mosquito plague by different methods. Data for 20 years were available, and the participants' task was to evaluate how effective the interventions were for alleviating the mosquito plague. According to the instructions, Community A used an insecticide in some of these years. In addition, it was pointed out that it was possible that the insecticide may not have worked because of the possibility of dilution. This hint was given to make the participants aware of the possibility that the insecticide also may not have proven effective, thus making it necessary to study the data in order to assess the efficacy of this intervention. For the method chosen by Community B, a different mechanism was mentioned. The participants were told that this community relied on a biological strategy. They planted a certain flower that beetles who hunt mosquito larvae needed in order to breed. Similar to the other condition, the participants were alerted to the possibility of failure. It was pointed out that one reason for a possible failure of this strategy may have been that the plant is very sensitive to environmental conditions—in particular, to pollution.

All the participants received the same data. For each community, they were presented with a table, on a different sheet of paper, that displayed 20 columns for the 20 years (1980–1999) and three rows. The first row showed the year, the second showed whether the insecticide was applied or whether the flowers were planted, and the third row showed whether there was a mosquito plague in the respective year or not. The data given were identical for both communities. In order to obscure the identity of the data, the columns were rearranged in the two conditions without changing the statistical relations between the intervention and the effect (see the analysis in the next paragraph). The two tables used are shown in the Appendix. The sequence of the conditions (short delay vs. long delay) and the assignment of the two tables to these conditions were counterbalanced.

The data were designed to embody different contingencies, depending on whether events were paired that occurred within the same year or in 2 consecutive years. Assessing the contingency in the same year (corresponding to a short-delay assumption) showed that there was no mosquito plague in 3 out of 8 years when the insecticide was used and in 6 out of 12 years when nothing was done. Thus, the contingency amounted to  $-0.125$ , which seems to indicate a slightly negative effect on the mosquitoes. In contrast,

taking a possible delay of 1 year into account (corresponding to a long-delay assumption) yielded a contingency of +0.292 which indicates a slightly positive effect of planting. In 5 out of 8 years, there was no plague when the flowers were planted the year before, whereas this was the case in only 4 of the 12 years when no flowers were planted the year before. This analysis holds true for both data sets. It is also notable that both contingency estimates were based on all observations without any differential weightings. The reason for the difference lay in the differential pairing of events.

In the experiment, the participants were instructed to study the sheet carefully in order to be able to assess how the two types of intervention affected the mosquito plague. The participants were permitted to take notes on a separate sheet of paper. After studying the data, the participants were requested to express their causal judgment for both interventions separately by using a rating scale ranging from 0 (*no effect*) to +10 (*strong effect*). As in the initial instructions, temporal delays were not mentioned in the rating instructions.

### Results and Discussion

The results confirmed that the participants proved sensitive to the manipulation of temporal assumptions. The participants concluded that the insecticide (short-delay condition) had only slightly reduced the mosquito plague ( $M = +4.2$ ). By contrast, the same participants, referring to the same data, estimated that the planting of the flowers (long-delay condition) had a larger positive effect ( $M = +5.4$ ). The difference between the ratings was significant [ $F(1,31) = 11.2, MS_e = 1.91, p < .01$ ]. These results support our prediction that the different instructions would lead the participants to combine data differently in their contingency assessments. The absolute difference between the conditions may not have been as large as was expected but this was possibly due to the fact that the manipulation of assumptions was rather implicit. Remember that the participants were never told about temporal aspects. The manipulation of temporal assumptions relied on the participants' making the expected inferences from the rudimentary information about causal mechanisms. Some of the participants may have arrived at different conclusions. For example, some may have inferred that the insecticides may also take time to take effect or that the plant intervention may take more than 1 year to materialize.

This post hoc explanation is bolstered by an additional experiment we ran in the pilot phase of this project. In order to get clear evidence for our predictions, we ran a similar study in which we *explicitly* manipulated assumptions about temporal delays. The participants were asked to evaluate the effect of a newly developed drug on the mood of a severely depressed patient. We manipulated temporal assumptions by either pointing out in the instructions that the temporal delay to be expected was very short, so that the effect should be visible within the same day (short-delay condition), or mentioning that the temporal delay of the effect was exactly 1 day (long-delay condition). All the participants received the same patient record as data. As in the present Experiment 1, data were presented that contained different contingencies (positive vs. negative), depending on the pairing of events. In this experiment, in which temporal assumptions were unam-

biguously instructed, the results demonstrated a clear reversal of causal ratings that closely mirrored the presented contingencies. Whereas the participants in the short-delay condition judged the drug to be helpful, the participants in the long-delay condition concluded that the drug led to a decline in the patient's mood.

### EXPERIMENT 2

The aim of Experiment 2 was to investigate whether temporal assumptions guide the selection of causal candidates in situations with multiple potential causes. Learners are typically confronted with situations that contain a large number of potentially related events. It seems plausible that there is a tendency to prestructure the event space into likely candidates of causally related events. In the present experiment, we tested the hypothesis that assumptions about probable time lags between causes and effects will lead participants to estimate contingencies only for event pairs that fall within the time window specified by the temporal assumption.

To test this hypothesis, we gave the participants the task of finding out the cause of nausea among passengers on a cruiser. All the participants received identical information about the menu for the passengers, the time the passengers had this menu, and whether they suffered from nausea or not. We manipulated the assumptions about the time lags for nausea after eating spoiled oysters or spoiled fish. The participants either had no assumptions, or assumed that one of the food items had an incubation period of 1 day and the other food item had one of 2 days.

We expected that the participants who had no assumptions about temporal lags would assess all passengers. This group served as a control group in order to show that in this particular, fairly simple, learning situation, it is possible to take a nonselective stance. However, we expected that a temporal assumption would be used to select possible causes. For example, if it takes 2 days for spoiled fish to generate nausea, we expected that the participants would assess only passengers who ate fish 2 days ago.

The participants received data that yielded different contingencies, depending on whether temporal assumptions were used or not. If temporal aspects were not considered, all the data should be assessed that yielded a contingency of zero between the two potential causes and the effect. Applying temporal assumptions and selecting the data accordingly yielded a positive contingency for one cause and a negative for the other.

### Method

**Participants and Design.** Thirty students from the University of Tübingen participated in this experiment. They were randomly assigned to one of three conditions. In the control condition, no information about latencies for nausea was provided. Temporal assumptions were manipulated in two experimental conditions. In one of these conditions, the participants were told that the time lag for nausea from spoiled fish was 2 days and that it was 1 day for spoiled oysters. In the second condition, oysters were assigned a 2-day lag and fish a 1-day lag.

**Procedure and Materials.** In the initial instructions, the participants were told that they should imagine themselves in the role of a physician on board a cruiser. One day, an epidemic of nausea occurred. The symptoms suggested that either spoiled fish or spoiled oysters might be the cause of the disease. To find out more about the cause, a short survey was conducted. Randomly selected passengers were asked whether they had eaten fish or oysters 1 or 2 days before and whether they suffered from nausea or not. In the control condition, no information was given about temporal lags. In the experimental conditions, the participants were told that, according to medical knowledge, fish typically takes 2 days and oysters 1 day to show an effect (*fish longer than oysters*) or vice versa (*oysters longer than fish*). No further information about any underlying mechanism leading to nausea was provided. The participants' task in all conditions was to assess whether fish or oysters were the cause of nausea on the cruise ship.

All the participants received identical learning data. The data were presented on a list. In the first column, it was shown whether and when the passenger had eaten fish ("none," "yesterday," or "the day before yesterday"). The second column provided the corresponding information about oysters. In the third column, it was shown whether the particular passenger had contracted nausea or not. The participants were asked to study the list carefully in order to be able to judge whether fish or oysters caused nausea. They were requested to express their impression on a rating scale ranging from -10 (*strongly reduces nausea*) to 0 (*does not influence nausea*) to +10 (*strongly generates nausea*). Again, it was not mentioned in the rating instructions that time should be considered. The participants were allowed to take notes on a separate sheet of paper.

Table 1 displays the structure of the learning data. Each cell contains the proportion of passengers that suffered from nausea after having eaten the particular combination of food. Disregarding temporal information would lead the participants in the control condition to lump together the corresponding cases in the cells that contain temporal information. Thus, 31 out of 62 people had nausea after eating fish, the same proportion as after having eaten oysters. Since these proportions were identical to the proportions of people with nausea who did not eat fish (18 out of 36), oysters (18 out of 36), or either (8 out of 16), a contingency-based assessment of causal influence would lead to the conclusion that both fish and oysters were not causally relevant. A more sophisticated analysis might take into account that the simultaneous presence of an alternative cause presents the problem of a possible confound. This problem might be overcome by using *conditional contingencies* as indicators of causalities, in which the causal relations were assessed when the cofactors were either constantly present or absent (see Spellman, 1996; Spellman, Price, & Logan, 2001; Waldmann & Hagmayer, 2001). However, this strategy would not change the conclusions for the data presented in this experiment. Checking, for example, the effect of fish on people who did not eat oysters would yield the same result as the overall analysis (10 out of 20). The same would hold true for oysters. The participants in the experimental conditions were expected to be sensitive to the hints about temporal lags. These participants would select cases on the basis of temporal plausibility. In the

fish-longer-than-oysters condition, passengers who had eaten oysters 2 days before and persons who consumed fish the day before would be eliminated from the analysis. If cases were to be selected according to this strategy, oysters would seem to inhibit nausea, whereas fish would appear to be a risk factor. Eleven out of 34 passengers who ate oysters the day before suffered from nausea, as compared with 38 out of 64 who did not eat oysters the day before. Eating fish 2 days before led to nausea in 20 out of 28 cases, as compared with only 29 out of 70 who did not eat fish during the critical period of time. In the second experimental condition (oysters longer than fish), the data indicated the opposite pattern, with oysters generating and fish inhibiting nausea. Again, some participants might decide to control for the potential impact of cofactors and only check cases in which the other factor was constantly present or absent. However, controlling for the confounded cause again would not change the statistical relations. In this case, the data would lead to qualitatively identical conclusions, which would allow one to assess the role of temporal assumptions independently of whether cofactors were held constant or not.

**Results and Discussion**

The results conformed to our predictions, showing that the participants used prior information about temporal lags in the selection of relevant events. An analysis of variance computed over all the conditions, with the two judgments for fish and oysters constituting a within-subjects factor, yielded a significant interaction between condition and ratings of causal efficacy [ $F(2,27) = 12.49, MS_e = 7.4, p < .01$ ]. No significant main effects were observed.

In the control condition, in which no prior information about temporal lags was provided, the ratings were close to zero for both potential causal factors: oysters ( $M = -0.50$ ) and fish ( $M = -0.10$ ). By contrast, in the experimental conditions, in which information about temporal lags was given, the results indicated that the participants analyzed only data about events that were temporally relevant. Consistent with the pattern of contingencies, in the fish-longer-than-oysters condition, the participants tended to give a negative causal rating for oysters ( $M = -2.2$ ) and a positive rating for fish [ $M = 2.6; F(1,9) = 14.05, MS_e = 8.2, p < .05$ ]. The opposite pattern was observed in the oyster-longer-than-fish condition [ $M = 1.7$  for oysters and  $M = -2.10$  for fish;  $F(1,9) = 6.64, MS_e = 10.9, p < .05$ ]. To summarize, the experiment clearly showed that the participants selected relevant data on the basis of temporal plausibility before they computed statistical indicators of causal strength. This finding is particularly interesting because the results for the control group showed that the participants were in principle capable, in this situation, of

**Table 1**  
**Structure of Data (Proportion of Patients With Nausea)**  
**Presented in Experiment 2**

Fish	Oysters			Total
	Two Days Before	One Day Before	None	
Two days before	9/10	5/10	6/8	20/28
One day before	5/10	2/12	4/12	11/34
None	6/8	4/12	8/16	18/36
Total	20/28	11/34	18/36	

processing all the data. Nevertheless, prior assumptions about temporal delays led the participants in the experimental groups to ignore data that they deemed irrelevant.

### EXPERIMENT 3

Experiments 1 and 2 showed that temporal assumption about causal lags can influence the selection of relevant events that enter into contingency estimates. Temporal assumptions may also affect the categorization of learning events, which, in turn, will influence the contingencies seen in the data. The aim of Experiment 3 was to study how temporal assumptions influence the aggregation of events into causally relevant event categories.

As in the introductory example, we used a drug test scenario for this experiment. Both short-term and long-term causal effects are of potential interest in such a situation. We assumed that if the participants were interested in short-term effects of the drug, they would conduct a *local analysis*. In this case, single intakes and their immediate effects should be analyzed. By contrast, if the participants were interested in the long-term effects, the singular events would be aggregated into pretreatment, treatment, and posttreatment phases, and a *global analysis* would be conducted. We assumed that, in this case, the participants would ignore the short-term fluctuations and would compare the pretreatment phase with the posttreatment phase.

To test whether learners are sensitive to the pragmatic consequences of different temporal focuses, we gave the participants the task of finding out whether a newly developed drug cured a patient suffering from a certain disease. All the participants received identical learning data about the treatment and the patient's state of health. The presented data were divided into a pretreatment phase, a treatment phase, and a posttreatment phase. We manipulated whether the participants were interested in short-term effects or in long-term effects. We expected that the first group would conduct a local analysis of the immediate effects of the drug, whereas the second group would globally divide the events into three stages and base their decisions on a comparison of the global probabilities of the effect within these stages.

The structure of the learning data allowed us to distinguish between these two strategies. It was constructed in such a way that a local analysis yielded a positive result for the efficacy of the drug, whereas a global analysis revealed a negative effect.

### Method

**Participants and Design.** Twenty students from the University of Tübingen participated in this experiment. They were randomly assigned to one of two conditions in which the pragmatic goals of the learners were manipulated.

**Procedure and Materials.** The initial instructions stated (in German) that the task was to assess the causal efficacy of a newly developed drug that was designed to improve the state of health of patients suffering from a rare tropical disease. The participants received written instructions. The two conditions differed in one sentence. Either the participants were told that they should imagine themselves to be physicians interested in the acute effects of the drug, or they were instructed to imagine themselves to be researchers interested in the long-term effects of the drug. Otherwise, the instructions were identical. All the participants had the same task—namely, to assess the effect of the drug on the patient.

The participants in both conditions received identical learning data. The data were presented as a patient's record (see Table 2). On three pages, information was given about the patient's health (*bad* or *good*) on consecutive days of 8 weeks. On the tops of pages 1 and 3, it was pointed out that no intervention was conducted, whereas the second page contained data from the intervention phase, in which the drug was given. This phase lasted 4 weeks, whereas the other two encompassed 2 weeks each. After studying the data, all the participants received identical rating instructions. They were requested to express their judgments of the causal efficacy of the drug on a rating scale ranging from  $-10$  (*very negative*) to  $+10$  (*very positive*). The participants were allowed to take notes on a separate sheet of paper.

The presented data allowed us to empirically distinguish between the two strategies. Across all three stages (pretreatment, treatment, and posttreatment) the patient was in good condition on 8 days out of 12 (75%) when the drug was given, whereas his health was good on 6 days out of 44 (14%) when no intervention was conducted. This difference indicated that the drug generally seemed to improve the state of health. Thus, on the basis of this local analysis, the drug seemed to have a positive effect. However, a comparison of the pre- with the posttreatment phase (global analysis) revealed that the patient was in a good condition on 4 out of 14 days before the treatment and on 2 out of 14 days after the treatment phase. Thus, the health of the patient slightly deteriorated in the posttreatment stage relative to the pretreatment stage.

### Results and Discussion

The results of the experiment showed that the participants' strategies were guided by the one differing sentence in the instructions that expressed the pragmatic focus. The participants who imagined themselves in the role of a physician interested in short-term effects concluded that the drug had a positive effect on the state of health of the patient ( $M = 2.8$ ), whereas the participants who assumed the role of a researcher interested in the long-term effect arrived at the opposite conclusion. In

**Table 2**  
**Structure of Data Presented in Experiment 3**

Pretreatment Phase (Weeks 1–2)		Treatment Phase (Weeks 3–6)		Posttreatment Phase (Weeks 7–8)	
Treatment Days: 0	No-Treatment Days: 14	Treatment Days: 12	No-Treatment Days: 16	Treatment Days: 0	No-Treatment Days: 14
Relative frequency of good health: 0/0 days	Relative frequency of good health: 4/14 days	Relative frequency of good health: 8/12 days	Relative frequency of good health: 0/16 days	Relative frequency of good health: 0/0 days	Relative frequency of good health: 2/14 days

their view, the drug had a negative effect on health ( $M = -2.5$ ). This difference was statistically significant [ $F(1,18) = 5.29$ ,  $MS_e = 26.6$ ,  $p < .05$ ]. These ratings were consistent with the contingencies in the data. The local analysis, focusing on single occurrences of the cause and its immediate effects, yielded a positive contingency, whereas the global analysis, which segmented the data into three stages and compared the pre- with the posttreatment phase, indicated a negative outcome. These conclusions were supported by an analysis of the notes some of the participants took. These notes showed an aggregation of the data in the predicted fashion.

## GENERAL DISCUSSION

Causal learning generally entails the problem of selecting appropriate induction strategies within a space of possibilities that typically surpasses the capacity of our information-processing system. These restrictions have often been circumvented in research on causal induction by presenting a preselected, clearly interpretable set of causal information to learners. In everyday life, we are not guided by helpful experimenters, however. We need to impose restrictions on the complexity of the learning situation ourselves. Causal-model theory is a theory that explicitly postulates the interaction of prior assumptions about causality and the processing of the learning input (see Waldmann, 1996). Prior knowledge about causal mechanisms may guide learning, but often this type of knowledge is not available to learners. However, there are other constraints that may reduce the complexity of the induction task. A number of studies have shown that learners use knowledge about causal directionality (e.g., Waldmann, 2000, 2001; Waldmann & Holyoak, 1992; Waldmann et al., 1995) and assumptions about causal relevance (Waldmann & Hagmayer, 2001) to guide their strategies for assessing causal strength. In the present paper, three experiments have been reported that add to this set of findings. Experiment 1 showed that assumptions about probable temporal lags between causes and effects may guide the choice of appropriate statistical indicators for causality. The results of this experiment demonstrate that learners tend to focus on contingencies between events that are within the expected temporal window specified by prior assumptions and tend to neglect other possible statistical relations. Experiment 2 supported these findings by showing that the participants' choices of a likely cause for an observed effect were guided by prior assumptions about temporal lags. The assumptions about time lags were used as criteria for the selection of possible causes. Events that occurred too far or too soon before the effect were neglected as candidate causes. Furthermore, this experiment showed that temporal assumptions are also used in situations that do not surpass capacity limitations. If the participants had no assumptions about time lags, all the cases were considered. Apparently, the participants' previous experience with complex learning situations outside the laboratory led to a general tendency to reduce complexity in

line with prior assumptions even in relatively simple learning situations. Experiment 3 also demonstrated the importance of assumptions about temporal lags and, in addition, showed that these assumptions may also lead to different levels of aggregating the observed events. The events either were considered individually or were aggregated into phases (pretreatment, treatment, and post-treatment), depending on the focus upon short-term or long-term effects.

## Temporal Assumptions and Perceived Temporal Delays

Different research paradigms have investigated the question of how the perceived temporal relations between learning events affect the outcome of learning and have found interesting and important effects (Schmajuk et al., 1998; Shanks & Dickinson, 1991; Wasserman & Neunaber, 1986; see also the introduction). Different theories have been proposed to account for these findings (Barnet & Miller, 1996; Gallistel & Gibbon, 2000). In contrast to these studies, we did not manipulate the timing of the learning events. All the participants in the experiments saw identical events. Furthermore, the temporal relations between the events were represented symbolically, rather than being observable in real time. This procedure was used to study the influence of temporal assumption without any interference of perceivable temporal sequences. As was elaborated in the introduction, our study focused on the role of temporal assumptions in the processing of *described* events. Whereas animals are restricted to learning situations in which they directly perceive events in real time (e.g., Lieberman, McIntosh, & Thomas, 1979), humans are capable, in addition, of learning about events that are described to them in communicative contexts. To a large extent, human knowledge is based on communicated data, and in fact, the majority of studies on human learning present learning events in processed form (e.g., event descriptions on lists, index cards, or computer displays). Thus far, the difference between perceived and described events has not been sufficiently acknowledged in the literature. Most theories of learning treat these two situations as identical.

Given that the present research was restricted to the investigation of the role of temporal assumptions in the processing of described events, an interesting question is how observable temporal relations and temporal assumptions interact in causal learning. The standard finding in research about the role of time in real-time learning contexts is that learners have a hard time detecting statistical relations when the events are separated by a temporal lag (see Domjan, 1998). There is also strong evidence from developmental studies for a general contiguity bias in causal reasoning (see Schlottmann, 1999). However, there are some studies that show that temporal assumptions and knowledge of causal mechanisms are able to overrule perceived temporal contiguity (Bullock et al., 1982; Mendelson & Shultz, 1976; Schlottmann, 1999). Despite these interesting initial findings, the inter-



play between top-down temporal assumptions and the timing of learning events has not been sufficiently addressed in extant theories specifying the processes of causal reasoning. In our opinion, the relation between time and causality in experienced and described events will prove a theoretically and empirically challenging field for future psychological research.

## REFERENCES

- AHN, W.-K., KALISH, C. W., MEDIN, D. L., & GELMAN, S. A. (1995). The role of covariation vs. mechanism information in causal attribution. *Cognition*, *54*, 299-352.
- ALLAN, L. G., & JENKINS, H. M. (1983). The effect of representations of binary variables on judgment of influence. *Learning & Motivation*, *14*, 381-405.
- ANDERSON, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- BARNET, R. C., & MILLER, R. R. (1996). Temporal encoding as a determinant of inhibitory control. *Learning & Motivation*, *27*, 73-91.
- BLAISDELL, A. P., DENNISTON, J. C., & MILLER, R. R. (1998). Temporal encoding as a determinant of overshadowing. *Journal of Experimental Psychology: Animal Behavior Processes*, *24*, 72-83.
- BULLOCK, M., GELMAN, R., & BAILLARGEON, R. (1982). The development of causal reasoning. In W. F. Friedman (Ed.), *Developmental psychology of time* (pp. 209-254). New York: Academic Press.
- CHENG, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367-405.
- CHENG, P. W., & NOVICK, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality & Social Psychology*, *58*, 545-567.
- CHENG, P. W., & NOVICK, L. R. (1992). Covariation in natural causal induction. *Psychological Review*, *99*, 365-382.
- DOMJAN, M. (1998). *The principles of learning and behavior* (4th ed.). Pacific Grove, CA: Brooks/Cole.
- EELLS, E. E. (1991). *Probabilistic causality*. Cambridge: Cambridge University Press.
- GALLISTEL, C. R., & GIBBON, J. (2000). Time, rate and conditioning. *Psychological Review*, *107*, 289-344.
- KOSLOWSKI, B. (1996). *Theory and evidence: The development of scientific reasoning*. Cambridge, MA: MIT Press.
- LESLIE, A. M., & KEEBLE, S. (1987). Do six-month-old infants perceive causality? *Cognition*, *25*, 265-288.
- LIEBERMAN, D. A., MCINTOSH, D. C., & THOMAS, G. V. (1979). Learning when reward is delayed: A marking hypothesis. *Journal of Experimental Psychology: Animal Behavior Processes*, *5*, 224-242.
- MATUTE, H., ARCEDIANO, F., & MILLER, R. R. (1996). Test question modulates cue competition between causes and between effects. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *22*, 182-196.
- MENDELSON, R., & SHULTZ, T. R. (1976). Covariation and temporal contiguity as principles of causal inference in young children. *Journal of Experimental Child Psychology*, *22*, 408-412.
- MICHOTTE, A. E. (1963). *The perception of causality*. New York: Basic Books.
- PRICE, P. C., & YATES, J. F. (1993). Judgmental overshadowing: Further evidence of cue interaction in contingency judgment. *Memory & Cognition*, *21*, 561-572.
- SCHLOTTMANN, A. (1999). Seeing it happen and knowing how it works: How children understand the relation between perceptual causality and underlying mechanism. *Developmental Psychology*, *35*, 303-317.
- SCHLOTTMANN, A., & SHANKS, D. R. (1992). Evidence for a distinction between judged and perceived causality. *Quarterly Journal of Experimental Psychology*, *44A*, 321-342.
- SCHMAJUK, N. A., LAMOUREUX, J. A., & HOLLAND, P. C. (1998). Occasion setting: A neural network approach. *Psychological Review*, *105*, 3-32.
- SHAKLEE, H., & TUCKER, D. (1980). A rule analysis of judgments of covariation between events. *Memory & Cognition*, *8*, 459-467.
- SHANKS, D. R. (1991). On similarities between causal judgments in experienced and described situations. *Psychological Science*, *5*, 341-350.
- SHANKS, D. R. (1993). Human instrumental learning: A critical review of data and theory. *British Journal of Psychology*, *84*, 319-354.
- SHANKS, D. R., & DICKINSON, A. (1991). Instrumental judgment and performance under variations in action-outcome contingency and contiguity. *Memory & Cognition*, *19*, 353-360.
- SHANKS, D. R., & LOPEZ, F. J. (1996). Causal order does not affect cue selection in human associative learning. *Memory & Cognition*, *24*, 511-522.
- SHULTZ, T. R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development*, *47* (No. 1).
- SHULTZ, T. R., & KESTENBAUM, N. R. (1985). Causal reasoning in children. In G. Whitehurst (Ed.), *Annals of child development* (Vol. 2, pp. 195-249). Greenwich, CT: JAI.
- SIEGLER, R. (1976). The effects of simple necessity and sufficiency relationships on children's causal inferences. *Child Development*, *47*, 1058-1063.
- SIEGLER, R., & LIEBERT, R. M. (1974). Effects of contiguity, regularity and age on children's inferences. *Developmental Psychology*, *10*, 574-579.
- SPELLMAN, B. A. (1996). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes. *Psychological Science*, *7*, 337-342.
- SPELLMAN, B. A., PRICE, C. M., & LOGAN, J. M. (2001). How two causes are different from one: The use of (un)conditional information in Simpson's paradox. *Memory & Cognition*, *29*, 193-208.
- SUPPES, P. (1970). *A probabilistic theory of causality*. Amsterdam: North-Holland.
- WALDMANN, M. R. (1996). Knowledge-based causal induction. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation: Vol. 34. Causal learning* (pp. 47-88). San Diego: Academic Press.
- WALDMANN, M. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *26*, 53-76.
- WALDMANN, M. R. (2001). Predictive versus diagnostic causal learning: Evidence from an overshadowing paradigm. *Psychonomic Bulletin & Review*, *8*, 600-608.
- WALDMANN, M. R., & HAGMAYER, Y. (2001). Estimating causal strength: The role of structural knowledge and processing effort. *Cognition*, *82*, 27-58.
- WALDMANN, M. R., & HOLYOAK, 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., HOLYOAK, K. J., & FRATIANNI, A. (1995). Causal models and the acquisition of category structure. *Journal of Experimental Psychology: General*, *124*, 181-206.
- WALDMANN, M. R., & MARTIGNON, L. (1998). A Bayesian network model of causal learning. In M. A. Gernsbacher & S. J. Derry (Eds.), *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 1102-1107). Mahwah, NJ: Erlbaum.
- WASSERMAN, E. A., & NEUNABER, D. J. (1986). College students' responding to and rating of contingency relations: The role of temporal contiguity. *Journal of the Experimental Analysis of Behavior*, *46*, 15-35.

## NOTES

1. Shanks (1991) also presents a distinction between experienced and described events that is different from ours. He separates trial-by-trial learning, which according to his view, is handled by associative learning mechanisms, and verbal descriptions of causal situations (e.g., stories or summary data) that trigger judgmental processes. However, on this account, learning trials may consist of verbal descriptions of cues and outcomes (as in disease classification tasks), which, according to our view, is an example of described, rather than directly experienced, causal scenarios.
2. Alternative measures of causal strength have been proposed (e.g., Cheng's, 1997, power PC theory). Since it was not our goal to differentiate between competing theories of causal strength, we used materials that lead to qualitatively similar patterns of results independently of the statistical indicator of causal strength that is being used.
3. As contingencies between events are asymmetric, there are 6 unconditional contingencies between three events and 12 conditional contingencies. Furthermore, possible interactions between events have to be considered, resulting in even more contingencies that can potentially be assessed.

**APPENDIX  
Data Presented in Experiment 1**

First Table																				
Year	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99
Intervention	I		I			I			I		I		I			I		I		
Mosquito plague	P		P	P				P		P	P			P	P	P		P		P
Second Table																				
Year	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99
Intervention	I		I		I			I		I			I		I				I	
Mosquito plague		P	P		P	P	P	P		P		P	P			P				P

Note—See Experiment 1, Procedure and Materials, for further explanation.

(Manuscript received March 21, 2000;  
revision accepted for publication May 25, 2002.)