Published in final edited form as:

J Exp Psychol Learn Mem Cogn. 2009 March ; 35(2): 334–352. doi:10.1037/a0014929.

Spontaneous Assimilation of Continuous Values and Temporal Information in Causal Induction

Jessecae K. Marsh and

Department of Psychology, Texas Tech University

Woo-kyoung Ahn

Department of Psychology, Yale University

Abstract

Existing models of causal induction primarily rely on the contingency between the presence and the absence of a causal candidate and an effect. Yet, classification of observations into these four types of covariation data may not be straightforward because (a) most causal candidates, in real life, are continuous with ambiguous, intermediate values and because (b) effects may unfold after some temporal lag, providing ambiguous contingency information. Although past studies suggested various reasons why ambiguous information may not be used during causal induction, the authors examined whether learners spontaneously use ambiguous information through a process called causal assimilation. In particular, the authors examined whether learners willingly place ambiguous observations into one of the categories relevant to the causal hypothesis, in accordance with their current causal beliefs. In Experiment 1, people's frequency estimates of contingency data reflected that information ambiguous along a continuous quantity dimension was spontaneously categorized and assimilated in a causal induction task. This assimilation process was moderated by the strength of the upheld causal hypothesis (Experiment 2), could alter the overall perception of a causal relationship (Experiment 3), and could occur over temporal sequences (Experiment 4).

Keywords

causal reasoning; categorization; ambiguity interpretation; belief assimilation

People appear to learn many novel, causal relations without much difficulty: They do not need intensive training to learn that flipping a switch causes a light to be turned on or that eating spoiled food causes stomachaches. Such apparently straightforward causal learning is actually impressive, given that the environment does not always present a learner with neatly organized data. In the current set of studies, we examine one of the difficulties people have to face as a result; namely, that observations may not come readily classified as evidence for or against a causal hypothesis.

Let us illustrate this problem in the context of covariation-based models of causal induction. The focus of these models has been on how people use covariation data to learn causal relations (e.g., Cheng, 1997; Cheng & Novick, 1990; Collins & Shanks, 2002; Jenkins & Ward, 1965; Rescorla & Wagner, 1972; Shanks, Lopez, Darby, & Dickinson, 1996). In these models,

Correspondence concerning this article should be addressed to Jessecae K. Marsh, Department of Psychology, Box 2051, Psychology Building, Texas Tech University, Lubbock, TX 79409-2051. jessecae.marsh@ttu.edu.

This article was submitted as part of Jessecae K. Marsh's dissertation project. A discussion of Experiments 1 and 2 can be found in a section of the book chapter "Dynamic Interpretations of Covariation Data" in the edited book *Causal Learning: Psychology, Philosophy, and Computation* (2007).

covariation data are preclassified into four events: What happens to an effect when a causal candidate is present (ce and $\bar{c}\bar{e}$ in Table 1; the overbar indicates the absence of an event) and what happens to an effect when the causal candidate is absent ($\bar{c}e$ and $c\bar{e}$). In general, ce and $\bar{c}\bar{e}$ are interpreted as evidence for a generative causal relationship, and $\bar{c}e$ and $c\bar{e}$ are interpreted as evidence against such a hypothesis. An important question that has been neglected in previous studies is how observations are classified into these four types to begin with.

Consider the hypothesis that overwatering causes a potted plant to die. Something akin to submerging a potted plant in a tub of water would obviously constitute overwatering, whereas leaving a potted plant in the desert at noon would not. Yet, it is not always obvious how events that fall between these extremes come to be classified as instances of the presence of the relevant cause, versus instances of the absence of the relevant cause. Indeed, many, if not all, of our causal hypotheses involve relations among continuous variables (e.g., caffeine intake, amount of exercise, level of depression, etc.); therefore, the discrete categorical distinctions required by existing models of causal induction for these continuous variables may not be particularly feasible. For example, what amount of exercise is a cause of better sleep? Should the effect of drinking only a half cup of coffee be considered when testing whether caffeine causes alertness?

Causal variables may be difficult to parse into the type of 2×2 contingency table shown in Table 1, not only because the variables are continuous but also because events develop over time, blurring the boundaries of the presence or the absence of a cause and an effect. For instance, upon drinking a cup of coffee, many people become revitalized after some temporal delay. These events separated by a temporal delay could be interpreted as drinking coffee causing alertness after some delay (i.e., ce ; evidence for a generative causal relationship) or as the separate events of drinking coffee not causing alertness (i.e., $\bar{c}\bar{e}$), followed by not drinking coffee and alertness (i.e., $\bar{c}e$; two pieces of evidence against a generative causal relationship; see also Greville & Buehner, 2007; Shanks & Dickinson, 1988).

Methods for Dealing With Ambiguity in Causal Induction

How do people handle quantitatively or temporally ambiguous observations if causal induction is to be performed over categorical values?¹ There are two possibilities. First, a learner may ignore, or not be swayed by, any ambiguous observations and may rely only on unambiguous, easily classified information. This approach may save cognitive resources, if resolving ambiguities requires great cognitive effort. It also could be more prudent, in that causal learning would be based only on clean data.

Alternatively, people may use ambiguous evidence during causal induction, which is the position purported in this article. Specifically, we propose that a learner spontaneously categorizes ambiguous observations into one of the four types of evidence listed in Table 1, on the basis of the hypothesis that the learner holds about a causal relationship, and as a result, ambiguous observations become used during causal learning. We refer to this process as *causal assimilation*. For example, imagine that a potted plant owner currently believes that overwatering (e.g., 500 ml of water in her definition) kills small potted plants, whereas using a small amount of water (e.g., 100 ml of water) does not. The plant owner might be initially uncertain whether 300 ml of water constitutes overwatering. Under the causal assimilation proposal, if 300 ml of water had killed her plant (i.e., e), she would interpret 300 ml as overwatering (i.e., c), but if the plant had survived (i.e., $\bar{e}$), she would interpret 300 ml as not

¹It is possible that participants may not categorize the ambiguous information and may instead learn a continuous relationship between the cause and the effect. The current results, however, argue against this possibility by demonstrating the spontaneous categorization of ambiguous information. We thank an anonymous reviewer for suggesting this additional implication of our results.

over-watering (i.e., $\bar{c}$). Conversely, if she holds an opposite causal belief (i.e., overwatering does not kill potted plants) then the former event would be interpreted as $\bar{c}e$, and the latter would be interpreted as $c\bar{e}$. In this way, the identical event is interpreted differently, depending on both the currently held hypothesis and its pairing with the effect. Similarly, in parsing temporally developing events such as c , followed by a temporal gap, followed by e , those who believe that c causes e may classify the sequence as ce , and those who believe that c does not cause e may classify the same sequence as $c\bar{e}$ followed by $\bar{c}e$. Once ambiguous information receives these types of categorical interpretations, it can be used in causal induction.

We have outlined two possible alternatives to how people might handle ambiguous causal information: People might ignore or use ambiguous causal events. In the next sections, we review previous studies supporting each of these two possibilities.

Interpretation of Ambiguous Information Due to Personal Beliefs

Work in attitude formation has demonstrated that personal beliefs can lead to the biased assimilation of information, akin to the top-down causal assimilation we are proposing. In a classic study, Lord, Ross, and Lepper (1979) found that people openly accepted data that aligned with their personal beliefs about capital punishment, while thoroughly questioning and critiquing conflicting evidence. Similar biased assimilation of information has been found in interpreting the information's validity (e.g., Munro & Ditto, 1997; Munro et al., 2002) and strength (e.g., Anderson, 1995) and the behavioral source of data (e.g., Sedikides & Anderson, 1992).

Yet, these demonstrations may be limited to instances in which people have a vested interest in interpreting ambiguities. Indeed, Lord et al. (1979) predicted assimilation in their experiment when "subjects with strong initial attitudes are confronted with empirical data concerning a controversial social issue" (p. 2100). The underlying catalyst that makes biased assimilation occur in attitude formation studies may be this preservation of a personal belief of importance. Beliefs that do not resonate with the same level of import and personal involvement as, for example, views on capital punishment may not promote biased assimilation. Thus, the relatively impersonal beliefs formed in a typical causal induction study may not hold the power to drive the assimilation of ambiguous information, as could be suggested by the previous social literature.

Interpretation of Ambiguous Information Due to Categorization

To the contrary, in colder, more impersonal categorization and similarity judgments tasks in which stimuli not related to strong personal values were used, evidence has been provided for top-down influences on the interpretation of ambiguous features of objects. For instance, Medin, Goldstone, and Gentner (1993) found that when an object that was ambiguous as to whether it had three or four appendages was paired with a three-appendage object, the ambiguous object was described as having three appendages, whereas when the same ambiguous object was paired with a four-appendage object, the ambiguous object was described as having four appendages. Similarly, Archambault, O'Donnell, and Schyns (1999) found that learning to categorize objects at either a specific level or a general level altered the ability to detect changed objects in a classic change detection task, presumably because different features of the objects were being selected and encoded during the categorization task. The effect of the categorization of an object on feature perception has been demonstrated in a variety of other studies (e.g., Quinn, Schyns, & Goldstone, 2006; Schyns, Goldstone, & Thibaut, 1998; Schyns & Oliva, 1999; Schyns & Rodet, 1997; Wisniewski & Medin, 1994; see also Corneille, Huart, Becquart, & Brédart, 2004; Dunning & Sherman, 1997; Eberhardt, Dasgupta, & Banaszynski, 2003; Gawronski, Geschke, & Banse, 2003;

Kunda & Sherman-Williams, 1993, for similar demonstrations of the effects of categorization in the social domain).

Within the causal learning field, categorization has been shown to affect the interpretation of causal efficacy. In Waldmann and Hagmayer's (2006) experiments, participants learned the category membership of exemplars, followed by information about whether each exemplar causes a certain effect. After this learning, participants completed a test phase in which they had to estimate the likelihood that particular exemplars caused the effect. There were critical exemplars, for which all participants learned identical causal information in the second phase but had learned different category membership in the first phase. Participants who initially learned to categorize these exemplars into the category that mostly produced the effect provided higher causal ratings than did those participants who learned to categorize the same exemplars into the category that mostly did not produce the effect, although all participants later learned identical causal information about these exemplars. That is, prior categorization affected subsequent causal learning (see also Lien & Cheng, 2000).

Note that what these previous studies demonstrated is how prior categorization provided by the experimenter affects the interpretation of information. However, the problem that we started out with was whether and how learners would spontaneously categorize events into the absence or presence of a causal candidate. In the absence of precategorization, assimilating ambiguous values while learning novel causal relations might result in overwhelming cognitive demands.

Furthermore, people may be unwilling to assimilate ambiguous causal candidates based on information about the presence of the effect because people know that the same effect can have multiple, distinctive causes. For example, sneezing can be caused by cats, dust, feathers, viruses, or cold weather. Lumping all of these causal candidates together just because they cause the same effect would certainly be nonsensical. To the extent that people understand that an effect can be caused in multiple ways, they may be reluctant to assimilate ambiguous information based on the status of an effect alone. Thus, the current study empirically tests whether causal assimilation spontaneously takes place despite these possibly preventative factors.

Problems in Parsing Events Over Time

Events that unfold over time present another case of ambiguity in causal induction. The sequence of a cause event followed by a temporal gap and then the effect event can be interpreted as ce happening over time, or as $c\bar{e}$ followed by $\bar{c}e$. Can a held causal belief differentially invoke these two interpretations, as illustrated in continuous quantities?

There are reasons to be skeptical about such top-down assimilation for temporal events. First, temporal contiguity is crucial in causal perception, and even a short temporal lag between events can disrupt this perception (e.g., Lagnado & Sloman, 2006; Michotte, 1946; Schlottmann & Shanks, 1992; Shanks & Dickinson, 1988; Shanks, Pearson, & Dickinson, 1989). As such, even when a learner believes that c causes e , the learner may be more likely to interpret c followed by a temporal gap and then e as a sequence of $c\bar{e}$ and $\bar{c}e$, rather than as ce .

Second, assimilation over a temporal dimension could mean that a learner should treat varying temporal delays equivalently. For instance, once a positive relationship between c and e is believed to exist, c followed by e sequences may be perceived as ce events, regardless of the length of the temporal gap between c and e . However, recent research has shown that people discriminate varying temporal gaps when provided with explicit background knowledge about a temporal delay (e.g., Allan, Tangen, Wood, & Shah, 2003; Buehner & May, 2002, 2003; Buehner & McGregor, 2006; Hagmayer & Waldmann, 2002). In Buehner and May (2004), for

example, after viewing a system in which a light bulb illuminated 4 s after an attached switch was flipped, participants provided much higher causal strength ratings if the cover story they had been given matched the delay (e.g., this is a light bulb that takes a while to warm up before lighting) than if the cover story did not match the data (e.g., this is a light bulb that should light immediately upon flipping the switch). Thus, when expecting short delays between *c* and *e*, participants appeared to have given less weight to data that displayed longer causal lags, instead of assimilating and using the longer delays in their causal theory. Similarly, Hagmayer and Waldmann (2002, Experiment 2) demonstrated that participants utilized only data about events that were temporally consistent with the mechanism they were provided with, rather than using all data.

These studies, however, may not serve as evidence against the causal assimilation of temporal lags, as the tasks and manipulations may have inadvertently presented situations in which it was impossible or highly improbable for causal assimilation to occur. For instance, in these experiments, the mechanisms underlying delays were spelled out to the participants (e.g., Buehner & McGregor, 2006); the presence of a delay and the absence of a delay were compared, rather than different degrees of delays (e.g., Buehner & May, 2004); and verbal summaries of temporal information were used instead of real-time, presented temporal data (e.g., Hagmayer & Waldmann, 2002). Although these studies make a valid point that people can selectively use events varying along temporal lags, these findings may not be conclusive evidence that people would be so sensitive to temporal differences that they would be unwilling to use a causal hypothesis to interpret ambiguous temporal gaps. In particular, temporal information acquired at the early stage of causal learning may not *necessarily* overwhelm the ability to interpret subsequent gaps in different ways, depending on the top-down hypothesis as predicted by causal assimilation.

Summary

To summarize, existing models of causal induction primarily rely on the contingency between the presence and the absence of a causal candidate and an effect. Yet, classification of observations into these four types of events may not be straightforward because most causal candidates, in real life, are continuous in nature with ambiguous intermediate values and because effects may unfold with some temporal lag, providing ambiguous contingency information. As such, ambiguous information may not be used in causal induction. In the current study, we attempt to provide an empirical demonstration of the causal assimilation process: namely, whether learners will use their current causal hypothesis to spontaneously categorize ambiguous information into the four types of events.

Experiment 1 serves as an initial test to examine whether causal assimilation spontaneously takes place; if so, people would categorize ambiguous data differently depending on their causal hypotheses. In Experiment 2, we test how this process changes, depending on the strength of observed causal relations. In Experiment 3, we investigate how causal assimilation may affect overall causal strength estimates. In Experiment 4, we examine whether people would assimilate data across real-time temporal dimensions in light of their governing hypothesis. Taken together, our goal in the current study is to establish the existence of the causal assimilation phenomenon by demonstrating that causal assimilation occurs due to a learner's current hypothesis, that causal assimilation matters in that it affects causal strength judgments, and that causal assimilation is not limited to only a certain type of dimension.

Overview of Experiments 1, 2, and 3

The general procedural logic used in Experiments 1, 2, and 3 is as follows. Participants were presented with stimuli that were designed to unambiguously represent the different types of

evidence captured in the traditional 2×2 contingency table of Table 1. Using the example of the Set 2 stimuli of Figure 1, participants would see tall bacteria (i.e., C) and short bacteria (i.e., $\bar{C}$) co-occur with the presence or the absence of an effect event (e.g., a protein's presence). We use an upper case C or $\bar{C}$ to represent polar values on a given dimension that were actually observed in the trial sequences (similarly, upper case E is used to represent the presence of effects actually observed in the trial sequence). Along with these events, participants were presented with instances that were our ambiguous stimuli: namely, items that did not clearly instantiate the presence or the absence of the candidate cause. In the Set 2 stimuli of Figure 1, for instance, such ambiguous candidate causes (A) would be instantiated as middle-height bacteria. After observing a trial-by-trial sequence of intermixed, well-defined, and ambiguous contingency information (e.g., CE, $\bar{C}\bar{E}$, AE), participants were asked to assess how many pieces of information they actually had seen corresponding to each of the four types of events outlined in Table 1 (e.g., for the ce judgment pertaining to the second material set in Table 1, the question read, "In how many cases were the bacteria TALL and the protein PRESENT?"). Critical to the current hypotheses, participants were never provided with definitions of what counts as each type of information (e.g., tall). This undefined cause that participants were asked to assess is referred to as a lower case c or $\bar{c}$. The main question of interest is whether participants' frequency estimates reflected the incorporation of ambiguous evidence into participants' representations of c or $\bar{c}$, as predicted by the causal assimilation account.

Pretest of Materials

The first step before the main experiments was to construct stimuli that are perceived as ambiguous. The participants for this pretest were 17 undergraduate Vanderbilt University students who did not participate in any of the main experiments. We initially constructed 12 sets of materials, each consisting of two anchors separated from each other along a continuous dimension (e.g., a 0.8 cm long anchor vs. a 8.8 cm long anchor) and 15 different test stimuli varying systematically between those anchors on that dimension (e.g., test items that were 1.3 cm, 1.8 cm, 2.3 cm, etc., in height). For each material set, we constructed a rating scale with one anchor (e.g., the 0.8 cm long anchor) placed on one end of the scale, labeled as a 1, and the other anchor (e.g., the 8.8 centimeters long anchor) placed on the opposite end of the scale, labeled as a 9. The participants' task was to place each of the 15 test items on the scale to indicate the item's similarity to the anchor items. A rating of 1 corresponded to something *very similar to Target A* in _____ and a rating of 9 corresponded to something *very similar to Target B* in _____. Targets A and B referred to the two anchors, and the blank was replaced with the dimension that varied between the two targets (e.g., height). With this scale, a test item equally similar to both anchors should be given a rating of 5, the midpoint of the scale. In each set, the 15 test items were presented and rated in a random order. On the basis of the pretest results, four material sets were selected for the main experiments. Each of these sets had a test item rated as not significantly different from the similarity scale's midpoint of 5 (one-tailed *t* tests, $p > .15$; see Figure 1 for the *M*s and *SD*s), and this item was also the actual objective midpoint between the two anchors (e.g., for the previous example, this would be something that was 4.8 cm in height). In the experiments reported below, these stimuli sets were described as bacteria varying either in spacing of membrane structure (Set 1), in height (Set 2), in membrane layer spacing (Set 3), or in outer shape (Set 4). To summarize, the pretest verified that the A values used in the main experiments were perceived to be equally similar to the two extreme values.

Experiment 1: Does Causal Assimilation Spontaneously Happen?

The goal in Experiment 1 was to test whether causal assimilation will happen; that is, whether causal candidates ambiguously defined over a continuous dimension would be spontaneously categorized into one of the four types of events that traditionally serve as the basis for causal induction. We used two pairings of the ambiguous causal candidates and the effect. In the AE

condition, instances of A were paired with E. If participants use the hypothesis that C causes E to assimilate the ambiguous observations then when A was paired with E, A would be interpreted as an instance of c. This interpretation would result in AE trials being categorized as ce trials, increasing the number of ce trials that participants would report that they had seen. In the A $\bar{E}$ condition, however, A was paired with the absence of the effect, and thus, A would be interpreted as $\bar{c}$, increasing reported $\bar{c}\bar{e}$ trials. Thus, the frequency estimates participants make about the well-defined trial types can illuminate whether ambiguous information is spontaneously incorporated, as predicted by our proposed causal assimilation account.

Method

Participants—Twenty undergraduate students from Yale University participated, either for pay or for partial fulfillment of an introductory psychology course requirement.

Procedures—Each participant viewed and made judgments for four trial blocks in a counterbalanced order. Each block instantiated one of the four conditions described below in the *Design and stimuli* section and used in a counterbalanced manner one of the four sets of stimuli selected through the pretest described earlier and shown in Figure 1 (e.g., bacteria varying in height). Each block began with a cover story that described the events depicted in the upcoming trials (i.e., participants were informed that they would see information about strains of bacteria that co-occurred with the presence of a compound, such as a protein or a mercury compound). The instructions described the temporal order of events (i.e., bacteria introduction preceded compound presence) in order to prevent confusion about the possible direction of causality. Participants were told that their task was to observe pieces of information depicting the co-occurrence of the bacteria and the presence of a compound, to later assess the causal relationship between one form of the bacteria and the presence of the compound (e.g., “determine the extent to which tall bacteria cause the presence of the protein”).

Then, participants observed 60 trials consisting of well-defined and ambiguous stimuli as described below. On each trial, participants saw a version of a bacteria type (e.g., a short bacterium in Figure 1) and its potential effect (e.g., the presence of a protein) presented simultaneously. Participants did not know how many trials they would observe for a given block.

After observing all the trials of a block, participants were asked to report how many trials they saw for each of the four cells of Table 1 (e.g., “In how many cases was the bacteria tall and the protein was present?”). The four frequency estimate questions were presented one at a time on the screen, with the presentation order of questions being randomized for each participant and each block (cf. Marsh & Ahn, 2003). Participants were not alerted ahead of time as to what frequency estimates they would be asked to make. Therefore, at least during the first block, participants did not know that they would not be asked to make estimates for the A trials (e.g., middle height bacteria). If participants knew ahead that they would only make four frequency estimates, they could have felt forced to incorporate the A trials into one of those four evidence types. By not informing participants ahead of time what estimates they were to make, we could minimize possible demand characteristics to assimilate ambiguous trials. It should be also noted that at no time throughout the experiment were participants given example forms or definitions of the descriptors (e.g., tall) used in the relationship of interest. In addition, there was no mention of the range of values for that dimension (e.g., “You will see bacteria that vary in height,” with no mention of any height ranges).

After making the frequency estimates, participants judged the overall causal strength between c and e on a scale of 0 (not a cause) to 100 (strongly causes), followed by a confidence rating for this judgment on a scale of 1 (not at all confident) to 7 (very confident). The causal strength estimate task was included to be in accordance with the participants’ instructed goal for

observing the trial frequencies, but the causal strength estimates were not of primary interest in this study, and this issue will be taken up in Experiment 3.² Lastly, participants were asked to rate the similarity of the ambiguous causal candidate to each of the anchor values used for C and $\bar{C}$, by the same method as in the pretest. This perceptual similarity posttest was added to see how the causal assimilation process might affect the perceived similarity of the ambiguous stimuli to the well-defined forms of the stimuli.

Upon completing the frequency, causal strength, confidence, and perceptual posttest judgments at the end of each block, participants moved to the next block. Participants moved at their own pace through all stages of the experiment. The experiment was administered on eMac computers via the RSVP experiment program (Williams & Tarr, 2001).

Design and stimuli—Participants observed 60 trials in a given block. The order of these 60 trials was randomized, and a black screen lasting for 500 ms was inserted to create a clear delineation between trials. Forty of the 60 trials used the two anchor values on the variant dimension, or C and $\bar{C}$ (e.g., the tallest and the shortest bacteria in Figure 1). C and $\bar{C}$ were paired with the presence or absence of the effect event in such a way as to suggest a strong, positive relationship between the cause and the effect (i.e., a $\Delta P = .8$;³ see Table 2 for the actual frequencies). The remaining 20 trials of the 60 total trials presented A (i.e., the form of the stimuli that was pretested to fall directly between the two anchor values; shown as the medium items in Figure 1). To be more ecologically valid and to help combat any boredom produced from repeatedly viewing the exact same form of stimuli, we developed two variations for each of C, $\bar{C}$, and A, one with a slightly greater value and the other with a slightly lesser value than C, $\bar{C}$, and A, along the critical dimension. For example, in material Set 2 of Figure 1, the original value of C used in the pretest was used along with one slightly taller and one slightly shorter variation. There was no overlap between the values used for C, $\bar{C}$, and A within each material set. The last column of Figure 1 shows rough ideas for the distributions of variations. The original values of C, $\bar{C}$, and A were used for 40% of the trials, and the two variations were each used for 30% of the trials.

There were two experimental conditions differing in how the 20 trials involving A were presented. In the AE condition, A was always paired with E, and in the $A\bar{E}$ condition, A was always paired with $\bar{E}$ (see Table 2.) If assimilation occurs, the A trials in the AE condition would be incorporated into the ce evidence type, whereas the A trials in the $A\bar{E}$ condition would be incorporated into the $\bar{c}\bar{e}$ evidence type.

Two additional control conditions were used for the following reasons. Differences in frequency estimates could result not only from causal assimilation, but also from sources such as faulty memory for similar events. For instance, finding increased ce estimates in the AE condition, compared with the $A\bar{E}$ condition, may be simply because the AE condition presents more trials overall involving the presence of E than the $A\bar{E}$ condition does (see Table 2). Such inequalities could bias participants to inflate estimates of certain cells in each condition. The control conditions included to test this possibility were identical to the two ambiguous conditions in all manners, except for what was presented in the place of A. The picture of the ambiguous causal candidate was replaced with a question mark and the word *unknown*, thereby providing no information about the type of bacteria present. Thus, there were two unknown

²Assimilation would result in an increase in the amount of positive evidence in both the AE and $A\bar{E}$ conditions (see Table 2). However, the predicted resulting change in overall causal strength is quite small and similar across conditions. For example, assimilating all twenty ambiguous trials into either the ce or $\bar{c}\bar{e}$ trial types results in the ΔP between C and E increasing from .8 to .85 and the causal power (see Cheng, 1997) rising from .889 to .894 in the $A\bar{E}$ condition and to .944 in the AE condition. Correspondingly, causal strength estimates in Experiment 1 did not differ over conditions, as reported in the *Results* section.

³ $\Delta P = p(E|C) - p(E|\bar{C})$, where E is the effect and C is the cause. That is, ΔP equals the probability of the effect occurring when the cause is present minus the probability of the effect occurring when the cause is absent.

conditions: One in which the unknown causal candidate (U) was paired with E (UE condition) and one in which (U)/U was paired with $\bar{E}$ ($U\bar{E}$ condition).

Assimilation is not predicted for these unknown trials because the unknown status does not provide any individuating information that could be used in a top-down manner for causal assimilation to occur.⁴ Akin to how a plant owner may not attribute a potted plant's death to overwatering if she is unaware of how much water was provided to the plant because of the existence of possible alternative causes of death (e.g., dehydration), participants in our study may not assimilate the unknown trials because of their lack of individuating information. In short, the unknown conditions provide comparison conditions that are matched for the overall number of trials involving the presence or the absence of effects and delay between trials, while still possessing some element of ambiguity in a section of the trial sequence (i.e., the ambiguity of what U may be). If assimilation occurs in the ambiguous conditions, then the ce frequency estimate is expected to be inflated in the AE condition over the ce estimate in the UE condition. Likewise, assimilation should result in the $\overline{ce}$ frequency estimate for the $A\bar{E}$ condition being higher than the same estimate in the $U\bar{E}$ condition. If assimilation does not occur, then no differences would be found between the ambiguous and the unknown conditions.

Results

Frequency estimates—Participants' average frequency estimates for the four types of events can be found in Table 2. There was no effect of order or materials in this experiment or in any of the following experiments. All analyses are collapsed across these measures.

The results presented in Table 2 indicate that participants did incorporate ambiguous information into their frequency estimates, as predicted by the causal assimilation account. Namely, participants' estimates for the ce evidence type in the AE condition were inflated over the same estimate in the UE condition, as was the $\overline{ce}$ estimate for the $A\bar{E}$ condition, as compared with the $U\bar{E}$ condition. No other evidence type in either of the ambiguous conditions showed similar inflations. Also as predicted, no evidence of assimilation was found in either of the unknown conditions. The following analyses provide statistical support for these claims.

First, we compared the frequency estimates in the AE condition with the estimates in the UE condition, using a 2 (condition: AE and UE) $\times$ 4 (evidence type: ce , $\overline{ce}$, $\bar{ce}$, and $c\bar{e}$) repeated measures analysis of variance (ANOVA) over the frequency estimates. There was a significant main effect of evidence type, $F(3, 57) = 137.2, p < .001, \eta_p^2 = .87$, and a significant main effect of condition, $F(1, 19) = 14.0, p < .001, \eta_p^2 = .42$. It is important to note that a significant interaction between condition and evidence type was found, $F(3, 57) = 4.15, p = .01, \eta_p^2 = .17$.⁵ This interaction arose because the ce estimate was significantly higher in the AE condition ($M = 27.5$) than in the UE condition ($M = 20.0$), $t(19) = 3.57, p = .002, d = 0.80$. There were no other significant differences between the AE and the UE conditions, ($ps > .20$).⁶

⁴Research in social judgeability theory showed that people avoid using personal beliefs (specifically, stereotypic beliefs) to interpret a person's behavior when little information is provided about the person (Yzerbyt, Schadron, Leyens, & Rocher, 1994). However, if some type of individuating information is known, regardless of how scant, participants will make judgments based on their beliefs (e.g., Corneille, Leyens, Yzerbyt, & Walther, 1999; Yzerbyt, Leyens, & Schadron, 1997). Thus, prior beliefs are applied only when there is some individuating information (regardless of its informativeness) that gives license to do so.

⁵The sphericity assumption was violated for the main effect of evidence type as well as the interaction of evidence type with condition.

Both of these findings remained significant when Greenhouse–Geisser corrections were used: $F(1.9, 36.5) = 137.2, p < .001, \eta_p^2 = .87$, and $F(2.1, 40.4) = 4.15, p = .021, \eta_p^2 = .42$, respectively. In all future analyses, if a violation of the sphericity assumption was found, it did not result in a change in the overall findings unless explicitly noted.

⁶Bonferroni corrections specific to these comparisons and all other similar t tests showed the same results throughout the article, unless specifically noted.

Similar comparisons were made for the A $\bar{E}$ and U $\bar{E}$ conditions, with a 2 (condition) $\times$ 4 (evidence type) repeated measures ANOVA for the frequency estimates. There were significant main effects for evidence type, $F(3, 57) = 174.1, p < .001, \eta_p^2 = .90$, and condition, $F(1, 19) = 22.6, p < .001, \eta_p^2 = .54$, as well as a significant interaction, $F(3, 57) = 7.69, p < .001, \eta_p^2 = .28$. The $\bar{c}\bar{e}$ estimate was significantly higher in the A $\bar{E}$ condition ($M = 28.8$) than in the U $\bar{E}$ condition ($M = 17.8$), $t(19) = 4.50, p < .001, d = 1.01$. There was no other significant difference between the A $\bar{E}$ and the U $\bar{E}$ conditions ($ps > .20$).

Separate analyses were completed to determine whether participants not only were differentiating their frequency estimates between the ambiguous and the unknown conditions but were also differentiating their estimates from the actual number of observed trials. Difference scores representing the amount of assimilation were calculated for each participant by subtracting the actual trial frequency for an estimate type from the estimated frequency for the corresponding evidence type. For each of the four conditions, these difference scores were submitted to one-way repeated measures ANOVA, with evidence type as a factor. For the AE condition, a significant main effect was found, $F(3, 57) = 6.35, p = .001, \eta_p^2 = .25$. One-sample t tests showed that the $\bar{c}\bar{e}$ difference score for the AE condition was significantly higher than zero, $t(19) = 4.46, p < .001, d = 0.99$. No other significant comparisons were found for the AE condition ($ps > .10$). The same ANOVA was completed for the A $\bar{E}$ condition, and again, a significant effect was found, $F(3, 57) = 8.43, p < .001, \eta_p^2 = .31$. Here, the $\bar{c}\bar{e}$ difference score was significantly higher than zero, $t(19) = 5.08, p < .001, d = 1.14$. No other comparison was significant for the A $\bar{E}$ condition ($ps > .10$). In line with our proposal, the frequency estimates from the unknown conditions did not significantly differ from the actual frequencies of C and $\bar{C}$. These same difference score ANOVAs were nonsignificant for the unknown conditions (UE condition: $p = .56, \eta_p^2 = .035$; U $\bar{E}$ condition: $p = .63, \eta_p^2 = .029$).⁷

Causal strength estimates did not differ across conditions (AE: $M = 79.2, SD = 17.7$; A $\bar{E}$: $M = 78.5, SD = 16.9$; UE: $M = 69.1, SD = 27.0$; U $\bar{E}$: $M = 75.5, SD = 19.7$). A 2 (cause type: ambiguous vs. unknown) $\times$ 2 (effect status: E vs. $\bar{E}$) repeated measures ANOVA showed neither significant main effects (cause type: $p = .17, \eta_p^2 = .097$; effect status: $p = .42, \eta_p^2 = .035$) nor a significant interaction ($p = .42, \eta_p^2 = .034$). Similar analyses for the confidence estimates for the causal strength also showed no significant main effects (cause type: $p = .35, \eta_p^2 = .048$; effect status: $p = .98, \eta_p^2 = 0$) nor a significant interaction ($p = .67, \eta_p^2 = .010$; AE: $M = 5.11$; A $\bar{E}$: $M = 5.03$; UE: $M = 4.74$; U $\bar{E}$: $M = 4.83$).

Perceptual similarity posttest judgments—The perceptual post-test ratings were examined for whether the act of classifying ambiguous stimuli with a well-defined causal

⁷It is possible that the wording of the frequency questions prevented assimilation in the unknown conditions. Participants may have felt that instances of the unknown cause appearing with the effect were *probably* examples of $\bar{c}$ but did not include such cases in their estimates because the frequency question asked, “In how many cases *were* [emphasis added] the bacteria tall?” An alternative version of Experiment 1 ($N = 25$) was run with the only change being made to the wording of the frequency questions: Participants were asked, “In how many cases *do you think* [emphasis added] the bacteria were tall?” This change should have allowed participants to include the unknown candidate causes in their frequency estimates if they only resisted previously because of the question wording. Contrary to this possibility, participants’ estimates showed the same pattern as in Experiment 1, in that the $\bar{c}\bar{e}$ estimates were significantly higher in the AE condition than the UE condition and $\bar{c}\bar{e}$ estimates were significantly higher in the A $\bar{E}$ condition than in the U $\bar{E}$ condition ($ps < .025, ds > 0.40$). As also seen in Experiment 1, the $\bar{c}\bar{e}$ estimate was significantly higher than 18 (the observed trial frequency) in the AE condition ($p = .01$) but not the UE condition ($p > .42$), and the $\bar{c}\bar{e}$ estimate was significantly higher than 18 in the A $\bar{E}$ condition ($p < .005$) but not the U $\bar{E}$ condition ($p > .90$). That is, participants’ estimates for the unknown conditions did not differ from the observed amount of trials as had been observed in Experiment 1. Participants’ perceptual posttest estimates showed the same pattern as in Experiment 1, with estimates significantly differing from the midpoint of 5 in the AE condition, $t(31) = 3.30, p = .002, d = 0.58$, and the A $\bar{E}$ condition, $t(33) = 2.24, p = .032, d = 0.39$, but not the unknown conditions ($ps > .5$).

candidate altered the subsequent perception of the ambiguous stimuli. If a similarity shift occurs due to causal assimilation then integrating the ambiguous trials with the ce estimate in the AE condition would have resulted in A being perceived as more similar to C than to $\bar{C}$. Similarly, in the $A\bar{E}$ condition, A would be perceived as more similar to $\bar{C}$ than to C. The similarity ratings for the different materials were recoded so that ratings close to 9 represented high similarity to C, whereas ratings close to 1 represented high similarity to $\bar{C}$ for all material sets. As shown in Figure 2, A was rated as more similar to C in the AE condition ($M = 6.33$, $SD = 1.78$) and as more similar to $\bar{C}$ in the $A\bar{E}$ condition ($M = 3.70$, $SD = 2.05$). No such shift was found in the unknown conditions (UE: $M = 5.30$, $SD = 2.08$; $U\bar{E}$: $M = 4.98$, $SD = 2.11$; see Figure 2). A 2 (cause type: ambiguous vs. unknown) $\times 2$ (effect status: E vs. $\bar{E}$) repeated measures ANOVA over the posttest judgments showed no main effect of cause type ($p = .65$, $\eta_p^2 = .01$) and showed a significant main effect of effect status, $F(1, 19) = 7.99$, $p = .011$, $\eta_p^2 = .29$. This effect should be interpreted in terms of a significant interaction, $F(1, 19) = 6.67$, $p = .018$, $\eta_p^2 = .26$. As evident in Figure 2, the posttest similarity judgments shifted only in the ambiguous conditions and not in the unknown conditions. The shift in similarity ratings was also found when comparing judgments with the absolute neutral midpoint of the similarity scale. One-sample t tests in which judgments in each condition were compared to the neutral midpoint of the similarity scale (i.e., 5) showed significant differences in the AE condition, $t(19) = 3.33$, $p = .004$, $d = 0.74$, and in the $A\bar{E}$ condition, $t(19) = 2.83$, $p = .011$, $d = 0.63$. No significant differences were found in the unknown conditions (UE: $p = .52$, $d = 0.14$; $U\bar{E}$: $p = .95$, $d = 0.01$), replicating the pretest.

Discussion

In Experiment 1, we tested whether people would spontaneously incorporate ambiguous information into estimates of the traditional breakdown of causal contingency information. Unlike previous demonstrations of the influence of top-down theories on ambiguous stimuli in purely cognitive tasks (e.g., Medin et al., 1993; Waldmann & Hagmayer, 2006), ambiguous information was not preclassified for participants, and therefore, the demand to assimilate ambiguous information in this study was much weaker. Participants in this experiment could have ignored or discounted ambiguous values and not incorporated such information into their estimates. Ignoring ambiguous information would have resulted in frequency estimates for the four evidence types not differing from the observed trial frequencies of the polar causal candidates or not differing from the unknown conditions. Alternatively, participants could have used a governing hypothesis about the causal relationship to interpret ambiguous candidates and spontaneously incorporate them into frequency estimates. The results of Experiment 1 support the latter account: Frequency estimates were inflated for the evidence types predicted by the causal assimilation account.

The causal assimilation of ambiguous information seems to be a preferential strategy and not an artifact of the demands of the experimental design. Participants had no knowledge of which frequency estimates they would make before the estimation sequence began and, therefore, did not know they would not be asked to estimate the ambiguous pairings as a separate class of events. It can therefore be concluded that participants did not include ambiguous information in the well-defined frequency estimates simply because they felt they must, given that there was no separate opportunity to estimate the ambiguous causal candidate trials.

In addition, participants demonstrated that causal assimilation was not a necessary consequence of the experimental design or tasks used in Experiment 1 because no assimilation was found in the unknown conditions. For instance, increased ce estimates in the AE condition or increased $\bar{ce}$ estimate in the $A\bar{E}$ condition could have been there merely because participants saw more trials involving E or $\bar{E}$, respectively. However, no such increases were observed in

the unknown conditions in which participants also observed proportionally more trials of E or $\bar{E}$ in the UE or the $U\bar{E}$ condition, respectively.

Lack of assimilation in the unknown conditions also demonstrates limits in the well-known confirmation bias. Participants could have used their governing hypothesis to interpret an unknown trial in a fashion similar to how ambiguous trials may have been interpreted: If the effect is present then the unknown event is most likely to have been an instance of the cause. The lack of assimilation in the unknown conditions and the presence of assimilation in the ambiguous conditions seem to reflect a natural preference to interpret ambiguously defined information in light of a governing causal hypothesis rather than a process of incorporating any ambiguity into estimates solely because of an overarching confirmation bias.

The perceptual posttest judgments provide yet another interesting way of assessing the influence of assimilation on causal reasoning. Within the biased assimilation literature, the process of interpreting evidence through a held belief has been found to have lasting results on such things as memory for a situation (Sedikides & Anderson, 1992) and attitude polarization for the given belief (Lord et al., 1979; Munro & Ditto, 1997). Previous cognitive studies showed that categorizing an object could result in a change in the object's internal representation by making it appear more similar to other items in the category (Goldstone, Lippa, & Shiffrin, 2001; Huttenlocher, Hedges, & Vevea, 2000; Livingston, Andrews, & Harnad, 1998). Similarly, in Experiment 1, the act of spontaneously categorizing the ambiguous AE trials with the CE trials resulted in A being rated as more similar to C in the perceptual posttest, and grouping $A\bar{E}$ trials with $\bar{C}\bar{E}$ trials resulted in A being rated as more similar to C. The unknown conditions show that this same participant pool rated the ambiguous values as not differing from the neutral midpoint of the scale when no such spontaneous categorization had occurred.

The above findings suggest that causal assimilation may be an online process of interpreting and reclassifying ambiguous stimuli that alters the perception of ambiguous items. At the same time, assimilation could take place because people were changing their definitions of what constitutes a cause, so that ambiguous information would fit the definition (e.g., see Murphy, 1988, for context-dependent changes in meaning of adjectives). Using our Set 2 materials as an example, seeing the ambiguous height bacteria with e could cause participants to open up their boundaries of what is defined as tall (i.e., c) to include anything that is very tall to medium in height. It would be intriguing in a future study to further examine whether changes in definitions of causal candidates and changes in perception of ambiguous stimuli work in combination to produce assimilation or whether one process is more dominant than the other under certain conditions.

Experiment 2: Assimilation and the Strength of Causal Hypotheses

A crucial aspect of the described causal assimilation is that a held hypothesis serves as the basis for assimilation. If this tenet is true, then altering the strength between C and E on which a hypothesis is built could alter the amount of assimilation observed. In Experiment 2, the covariation between C and E (i.e., well-defined trials) was varied to be perfect ($\Delta P = 1.0$), moderate ($\Delta P = .6$), weak ($\Delta P = .3$), or none ($\Delta P = 0$), to test this possibility. There are several reasonable outcomes for how ambiguous information may be assimilated across these changes in covariation, illustrated in Table 3 in their most extreme forms.

One possibility (Possibility 1 in Table 3) is that causal assimilation is an all or none process such that as long as people believe in a positive causal relationship, whether weak or strong, all ambiguous information should be incorporated to the same extent. Thus, a similar pattern and amount of assimilation as seen in Experiment 1 would be expected across the perfect,

moderate, and weak conditions. Because there is no causal relationship depicted in the zero contingency condition, no assimilation would be predicted.

The second possibility (Possibility 2 in Table 3) is that the strength of the relationship between C and E will determine into which evidence types the ambiguous information is classified. To explain, the AE trials could potentially belong to either ce or $\bar{c}\bar{e}$ because the effect is present in both of these cases. When assimilation takes place, people may divide AE trials between ce and $\bar{c}\bar{e}$ in a way that mirrors the observed proportion of CE or $\bar{C}\bar{E}$ trials. Thus, the same total amount of ambiguous evidence would be incorporated into frequency estimates across the positive contingency conditions, but which estimates the ambiguous information is incorporated into (i.e., ce and $\bar{c}\bar{e}$) would change depending on the given contingency.

The third possibility is that once a learner believes in a positive causal relationship, AE trials would always be interpreted as ce and not as $\bar{c}\bar{e}$ (i.e., differing from Possibility 2 above), but the amount of ambiguous information incorporated into ce frequency estimates would decrease as the strength of the relationship between C and E decreases (differing from Possibility 1 above; see Possibility 3 in Table 3 for an illustration). Although not the focus of the current study, there are several plausible reasons why this might be the case. At lower contingency strengths, people may be less confident in the overall hypothesis on which causal assimilation is based. Also, the negative information that is frequently encountered in weaker contingency conditions may delay the formation of a hypothesis of how C and E are related, which in turn, would delay causal assimilation. Both of these possibilities would result in fewer ambiguous trials being incorporated into frequency judgments as the contingency between events decreases. Experiment 2 contrasts the possible outcomes illustrated in Table 3 to determine how the strength of a causal relationship affects assimilation of ambiguous evidence.

Method

Participants—Thirty-two Yale University undergraduates participated, either for partial fulfillment of an introductory psychology course requirement or for monetary compensation.

Materials and procedures—The same set of materials, instructions, and procedures were used as in Experiment 1 except for the following modifications. (a) The $A\bar{E}$ condition was not included because it was a mirror image of the AE condition in Experiment 1 and, therefore, redundant to demonstrating assimilation. (b) Participants were presented with four different conditions in which the covariation between C and E was manipulated (zero, weak, moderate, perfect) as specified in Table 3. Participants completed the four conditions in a counterbalanced order. (c) The unknown control conditions were not included in Experiment 2, and participants' estimates were compared directly with the observed trial frequencies because we found in Experiment 1 that participants were fairly accurate in frequency estimates; the $c\bar{e}$ and $\bar{c}e$ estimates, the ce estimate for the $A\bar{E}$ condition, and the $\bar{c}\bar{e}$ estimate for the AE condition did not differ from observed frequencies.

Results and Discussion

Frequency estimates—Table 3 shows the mean frequency estimates for this experiment. To allow us to compare the amount of assimilation across the contingency conditions, difference scores were calculated by subtracting the observed frequencies of CE, $\bar{C}\bar{E}$, $\bar{C}E$, and $C\bar{E}$ from participant's reported frequency estimates for ce , $\bar{c}\bar{e}$, $\bar{c}e$, and $c\bar{e}$, respectively. Higher difference scores should reflect inflation in an evidence estimate over the number of observed trials. The mean difference scores for each evidence type across the conditions, as shown in Figure 3, were most consistent with Possibility 3 outlined in Table 3. The ce estimates were more inflated over observed CE trials at higher contingencies than at lower contingencies, suggesting greater assimilation of AE trials into ce estimates for higher contingencies. Other

difference scores remained flat, staying close to 0, demonstrating that participants did not indiscriminately inflate frequencies as contingency increased and did not assimilate AE trials into $\bar{c}e$ estimates, as proposed in Possibility 2.

A 4 (contingency: zero, weak, moderate, perfect) $\times$ 4 (evidence type: ce , $\bar{c}\bar{e}$, $\bar{c}e$, $c\bar{e}$) repeated measures ANOVA conducted over the difference scores found significant main effects of evidence type, $F(3, 93) = 9.98, p < .001, \eta_p^2 = .24$, and contingency, $F(3, 93) = 4.91, p = .003, \eta_p^2 = .14$. These effects, however, should be interpreted in light of a significant interaction effect, $F(9, 279) = 2.04, p = .035, \eta_p^2 = .062$, indicating that the difference scores varied by type of evidence differently over the contingency conditions.⁸ Separate one-way repeated measures ANOVAs for ce , $\bar{c}\bar{e}$, $\bar{c}e$, and $c\bar{e}$ difference scores, with contingency as a factor, showed a significant main effect of contingency only in the ce evidence type, $F(3, 93) = 4.65, p = .004, \eta_p^2 = .13$, and not in $\bar{c}\bar{e}$ ($p = .21, \eta_p^2 = .047$), $\bar{c}e$ ($p = .22, \eta_p^2 = .046$), or $c\bar{e}$ ($p = .94, \eta_p^2 = .004$), consistent with Possibility 3 of Table 3. A within-subject contrast conducted for the ce evidence type difference scores found a significant linear trend across the contingency factor, $F(1, 31) = 11.02, p = .002, \eta_p^2 = .26$. This contrast was not significant for the $\bar{c}\bar{e}$ ($p = .61, \eta_p^2 = .008$), $\bar{c}e$ ($p = .16, \eta_p^2 = .061$), and $c\bar{e}$ ($p = .94, \eta_p^2 = .004$) difference scores.

Causal strength estimates differed across the conditions as would be predicted by the underlying contingency for the given condition (zero: $M = 28.3, SD = 26.3$; weak: $M = 36.2, SD = 27.5$; moderate: $M = 60.6, SD = 25.1$; perfect: $M = 81.0, SD = 25.9$), $F(3, 93) = 36.9, p < .001, \eta_p^2 = .54$. Participants were more confident in their causal strength estimates for the stronger contingencies (zero: $M = 4.67, SD = 1.10$; weak: $M = 4.61, SD = 1.09$; moderate: $M = 5.05, SD = 1.28$; perfect: $M = 5.62, SD = 1.28$), $F(3, 93) = 5.99, p = .001, \eta_p^2 = .16$.

Perceptual similarity posttest judgments—The perceptual post-test provides more evidence that the strength of the causal relationship for an existing hypothesis influences the causal assimilation process. (One participant's rating was missing in the moderate condition.) A one-way repeated measures ANOVA conducted for the posttest similarity judgments, with contingency as a factor showed a significant effect, $F(3, 90) = 4.38, p = .006, \eta_p^2 = .13$. Specific one sample t tests comparing the posttest judgments with the midpoint of the similarity scale (i.e., 5) found that A was rated as more similar to the end of the scale representing C than the scale's neutral midpoint in the perfect ($M = 6.19, SD = 1.86$) $t(31) = 3.60, p = .001, d = 0.64$ and moderate conditions ($M = 5.8, SD = 1.53$), $t(30) = 2.90, p = .007, d = 0.52$. This difference was not significant for the weak ($M = 4.92, SD = 1.63$; $p = .78, d = 0.050$) and the zero contingency conditions ($M = 5.25, SD = 1.40$; $p = .31, d = 0.18$).

To summarize, in Experiment 2, we found that as the causal relationship that a learner's hypothesis was built on became weaker, less ambiguous information was incorporated into frequency estimates. The type of evidence these ambiguous trials were grouped with remained the same, regardless of the strength of the causal relations: AE was assimilated as only ce trials, even when observed, well-defined trials contained a good number of $\bar{C}E$ trials. That is, instead of trying to mirror the observed, well-defined trials, the direction of assimilating ambiguous information during causal induction appears to be uniform.

⁸The sphericity assumption was violated for the interaction and was nonsignificant when Greenhouse–Geisser corrections were used: $F(4.7, 148.1) = 2.04, p = .079, \eta_p^2 = .062$.

Experiment 3: Influencing Causal Strength

Although Experiments 1 and 2 were not designed to test how causal assimilation could affect overall perceived causal strength, Experiment 3 used the design illustrated in Table 4 to overcome this limitation. Participants observed the AE or AĒ trials of Experiment 1, followed by a second phase in which the pairing of A with E or Ē was either consistent with the first phase (AE/AE and AĒ/AĒ conditions in Table 4) or opposite to the first phase (AE/AĒ and AĒ/AE conditions in Table 4). The question is how the causal assimilation of A, having taken place in the first phase, would influence the interpretation of A in the second phase and whether this could influence subsequent causal strength ratings.

Consider first the AE/AE and the AĒ/AĒ conditions in Table 4. As demonstrated in Experiment 1, participants in the AE/AE condition would interpret A as c in the first phase (as indicated in the brackets in Table 4), whereas participants in the AĒ/AĒ condition would interpret A as $\bar{c}$ in the first phase. Causal assimilation in the first phase alone would not result in any noticeable difference in causal strength judgments between the two conditions because in both conditions, ambiguous trials are assimilated as positive evidence. (For instance, assimilating all 20 ambiguous trials of the first phase would result in a ΔP of .85 for the first phase in both conditions.) Yet, maintaining these interpretations of A into the second phase would result in differences in causal strength judgments between the two conditions. In the AE/AE condition, AE trials in the second phase would be interpreted as ce trials (or evidence for a generative relationship) as A continues to be interpreted as c. However, in the AĒ/AĒ condition, the same AE trials would be interpreted as $\bar{c}e$ trials (or evidence against a generative relationship) as A continues to be interpreted as $\bar{c}$. These opposite interpretations would result in a noticeable difference in causal strength judgments; for the second phase, ΔP is .85 in the AE/AE condition but is 0.35 in the AĒ/AĒ condition when full assimilation occurs in line with the first phase interpretation (see Table 4). Thus, the current design can demonstrate apparently paradoxical phenomena; the AE/AE condition and the AĒ/AĒ condition would initially induce the same first phase causal strength judgments, and then, after observing identical second phase data (i.e., AE trials), the two conditions would end up with different causal strength estimates for the second phase. Such intriguing results would be obtained due to the carryover of the spontaneous interpretation of the ambiguous information that took place in the first phase.

Similar comparisons can be made between the AE/AĒ and the AĒ/AE conditions. After similar causal strength judgments in the first phase (due to positive assimilation in both conditions), observing identical second phase AĒ trials would result in a lower second phase causal strength judgment in the AE/AĒ condition than in the AĒ/AĒ condition. This is because AĒ trials in the second phase are interpreted as negative evidence ($c\bar{e}$) in the AE/AĒ condition but as positive evidence ($\bar{c}\bar{e}$) in the AĒ/AĒ condition.

Method

Participants—Twenty-eight Yale University undergraduates participated, either for pay or for partial fulfillment of an introductory psychology course requirement.

Methods and procedures—The stimuli and cover stories were the same as in Experiment 1. Each condition consisted of two phases. In two of the conditions (AE/AE and AE/AĒ conditions), the first phase was identical to the AE condition of Experiment 1. In the remaining two conditions (AĒ/AE and AĒ/AĒ conditions), the first phase was identical to the AĒ condition of Experiment 1. (See Table 4 for observed trial frequencies). After observing all the trials of the first phase, participants estimated the causal relationship between c and e (referred to as the first phase causal strength judgment) and provided a confidence estimate. The first phase contingencies were constructed such that even if participants had assimilated all As in line with the observed C–E contingency, causal strength estimates would not differ between

AE and A $\bar{E}$ versions. This design feature would ensure that any difference in the second phase causal strength judgments could not be due to differences in initial anchoring of the estimates (e.g., Tversky & Kahneman, 1974).

After making the first phase judgments, participants were instructed that they would observe another set of trials, after which they would assess the relationship between c and e for the second phase. The second phase depicted a shorter sequence of trials that again paired A with the presence or the absence of E. In the AE/AE and A $\bar{E}$ /AE conditions, A was paired with E, whereas in the AE/A $\bar{E}$ and A $\bar{E}$ /A $\bar{E}$ conditions, A was paired with $\bar{E}$. After observing these trials, participants were asked to estimate the causal strength between c and e “based on the second set of data you just viewed.” After the second phase judgment, participants made confidence judgments and posttest perceptual similarity judgments as in Experiment 1. Participants completed all four conditions (AE/AE, A $\bar{E}$ /AE, AE/A $\bar{E}$, and A $\bar{E}$ /A $\bar{E}$) in a counterbalanced order, with the four stimulus materials of Figure 1 being counterbalanced across conditions.

Results and Discussion

Causal strength estimates—Participants did not diverge on their first phase estimates (see Table 4 for the means). A one-way repeated measures ANOVA with condition (AE/AE, A $\bar{E}$ /AE, AE/A $\bar{E}$, A $\bar{E}$ /A $\bar{E}$) as a factor showed no significant differences ($p = .23$, $\eta_p^2 = .051$).

The second phase causal strength judgments diverged. A 2 (first phase: AE vs. A $\bar{E}$) $\times$ 2 (second phase: AE vs. A $\bar{E}$) repeated measures ANOVA on the second phase estimates found a significant interaction effect, $F(1, 27) = 10.3$, $p = .003$, $\eta_p^2 = .28$ (see Figure 4). None of the main effects were significant (first phase: $p = .57$, $\eta_p^2 = .012$; second phase: $p = .10$, $\eta_p^2 = .097$).

The causal strength judgments of the second phase were significantly higher in the AE/AE condition ($M = 60.7$) than in the A $\bar{E}$ /AE condition ($M = 48.3$), $t(27) = 2.45$, $p = .021$, $d = 0.46$ (see Figure 4), and they were significantly higher in the A $\bar{E}$ /A $\bar{E}$ condition ($M = 69.8$) than in the AE/A $\bar{E}$ condition ($M = 53.6$), $t(27) = 2.72$, $p = .011$, $d = 0.51$. That is, identical second phases led to different causal strength estimates, presumably because the interpretations of ambiguous information developed in the first phase affected causal strength judgments in the second phase.

Confidence and perceptual similarity posttest judgments—Changing the A pairing from the first phase to the second phase affected confidence estimates. (One participant’s confidence rating in the second phase A $\bar{E}$ /A $\bar{E}$ condition was missing.) A 2 (first phase: AE vs. A $\bar{E}$) $\times$ 2 (second phase: AE vs. A $\bar{E}$) repeated measures ANOVA on the first phase confidence estimates showed no significant main or interaction effects (first phase: $p = .76$, $\eta_p^2 = .003$; second phase: $p = .35$, $\eta_p^2 = .032$; interaction: $p = .94$, $\eta_p^2 = 0$). The same ANOVA on the second phase confidence estimates showed no significant main effects (first phase: $p = .74$, $\eta_p^2 = .004$; second phase: $p = .36$, $\eta_p^2 = .032$). However, a significant interaction was found, $F(1, 26) = 4.27$, $p = .049$, $\eta_p^2 = .14$, because participants’ confidence was marginally lower when the ambiguous trials changed in the second phase from AE to A $\bar{E}$ or vice versa (AE/A $\bar{E}$: $M = 4.59$, $SD = 1.22$; A $\bar{E}$ /AE: $M = 4.81$, $SD = 1.14$) than when the ambiguous trials remained the same (AE/AE test: $M = 5.18$, $SD = 1.07$; A $\bar{E}$ /A $\bar{E}$: $M = 5.06$, $SD = 1.40$).

The posttest perceptual similarity judgments were analyzed as in the previous experiments via one-sample t tests in which the ratings were compared against the midpoint value of 5. Again, higher ratings indicate greater similarity to C. Participants reported a significant shift in the

similarity of the ambiguous item only when it was consistently paired with the same form of E between the first and the second phases. In the AE/AE condition, A was seen as more similar to C than $\bar{C}$ ($M = 5.82$, $SD = 1.39$), $t(27) = 3.10$, $p = .004$, $d = 0.59$, and in the $A\bar{E}/A\bar{E}$ condition A was seen as more similar to $\bar{C}$ than C ($M = 4.14$, $SD = 1.48$), $t(27) = 3.05$, $p = .005$, $d = 0.58$. However, when A changed its pairing in the second phase, participants did not perceive A as significantly different from the midpoint of the similarity scale in the posttest ($A\bar{E}/AE$: $M = 5.33$, $SD = 1.43$, $p = .21$, $d = 0.24$; $AE/A\bar{E}$: $M = 4.95$, $SD = 1.71$, $p = .88$, $d = 0.028$). This last finding suggests that once information conflicting with the previous grouping of the ambiguous stimuli was provided, the ambiguous stimuli were seen as less like the items with which they were previously grouped. Overall, Experiment 3 showed that causal assimilation of ambiguous information could affect perceived causal strengths and could shift the perceived similarity of ambiguous information to other information.

Experiment 4: Causal Assimilation Over Temporal Sequences

The preceding experiments have all used discrete instances of the pairings of causes and effects. For example, C and $\bar{E}$ were presented side by side on a single computer screen, clearly indicating that no effect occurred following C. However, in real-life situations, pairings of cause and effect are often not clear-cut because many effects follow their causes by some delay. Recalling the example in the introduction, a person could be trying to decide whether a temporal delay between drinking caffeine and feeling alert was an instance of caffeine consumption causing alertness later (i.e., ce) or caffeine consumption with no alertness (i.e., $c\bar{e}$), followed by the absence of caffeine consumption and then alertness (i.e., $\bar{c}e$). Thus, temporal gaps between events provide another kind of ambiguity inherent in causal learning and a potential for assimilation, which is the focus of Experiment 4.

Participants in Experiment 4 first received training, during which pairs of events (e.g., a disk dropping into a toy followed by the toy playing a movie after a 3 s pause) were concluded with the presentation of an all black screen, so that participants could learn the absolute timing involved in cause and effect pairings. In the positive training condition, participants saw only CE and $\overline{CE}$ events, whereas in the negative training condition, participants saw only $\bar{C}E$ and $C\bar{E}$. Afterward, all participants received the same test sequences, which were continuous streams of trials presented without black screens separating event boundaries. Thus, during the test sequences, the sequences of $C\bar{E}$ followed by $\bar{C}E$ could be interpreted as $C\bar{E}$ followed by $\bar{C}E$ (if only the previously learned timing was taken into account) or as CE with a long delay between C and E. If causal assimilation takes place for temporal dimensions, the $C\bar{E}$ – $\bar{C}E$ sequences would be more likely to be interpreted as CE after the positive training than after the negative training.

As explained in the introduction to this article, however, this possibility is counterintuitive in light of the existing literature. First, temporal gaps on the order of more than 2 s (as were used in this experiment) destroy the perception of a causal relation (e.g., Shanks et al., 1989). Thus, even when one holds a positive causal hypothesis, people might be more predisposed to perceive a $C\bar{E}$ – $\bar{C}E$ sequence (or C followed by a temporal gap, followed by E) as the separate events of $c\bar{e}$ and $\bar{c}e$. Second, recent studies (e.g., Buehner & May, 2002; 2003; 2004; Buehner & McGregor, 2006; Hagmayer & Waldmann, 2002) provided results suggesting that data inconsistent with the learned, long temporal gap could be discounted or ignored. Thus, participants, once trained about the specific temporal gap durations for a certain causal relationship (e.g., 3 s), might be reluctant to perceive $C\bar{E}$ – $\bar{C}E$ sequence as a ce event because it violates the learned temporal relationship, even when one holds a positive causal hypothesis. That is, perceptual constraints may result in identical (and veridical) interpretations of the $C\bar{E}$ – $\bar{C}E$ sequence across the positive training and the negative training conditions. Alternatively,

and in line with Experiments 1–3, participants may use the hypotheses developed during the training phase to interpret the $C\bar{E}$ – $\bar{C}E$ sequences of the test phase differently.

Method

Participants—Twenty Yale University undergraduates participated, either for pay or for partial fulfillment of an introductory psychology course requirement.

Method and procedures—Two short animation clips were developed. One depicted an airport scanner that could flash a warning light when a piece of luggage labeled as a biohazard passed through it. The other depicted a child's toy that could play a movie when a disk was dropped into the toy (see Figure 5). In both sets of materials, the presentation of the causal stimulus was followed by a short temporal delay before the appropriate effect outcome would commence. The timing for the airport scanner materials was as follows: The luggage entered the scanner and interacted with the scanner for 2 s; then, there was no scanner activity for 4 s, resulting in the scanner displaying its warning light for 5 s. The toy materials used the following timing: The disk was dropped into the toy and interacted with the toy for 2 s; then, there was no toy activity for 3 s, resulting in the toy playing its movie for 5.5 s (see Figure 5). The $\bar{C}$ and $\bar{E}$ events were constructed by showing no activity in the apparatus for the same amount of time as the total length of the C and E events, respectively. This presentation equated the timing for a C and $\bar{C}$ event, as well as for an E and $\bar{E}$ event. Representing $\bar{C}$ and $\bar{E}$ as the absence of events made these events look the same as the temporal gap between C and E (see Figure 6 for an example).

At the beginning of each block, participants were told that the film clips they were about to watch depicted a new apparatus (e.g., the scanner or the toy) undergoing a series of tests of whether it reliably activated. Participants also were told how the apparatus was supposed to function (e.g., “dropping a red disk into a slot in the top of a toy causes the toy to play a short movie”). Participants were not told whether the apparatus worked correctly, nor were they told what estimates they would make after viewing the trials.

After these initial instructions, participants observed the training phase, which was designed to establish a positive or a negative causal hypothesis (i.e., c causes or does not cause e), as well as to teach participants the timing of the causal interactions. Six interactions of the potential cause and the apparatus were depicted with the beginning and end of each interaction being clearly demarcated by an all-black separating screen (B) presented for 2 s. In the positive training condition, the sequence was $CE, B, \bar{C}\bar{E}, B, CE, B, \bar{C}\bar{E}, B, CE, B, \bar{C}\bar{E}$. In the negative training condition, the training sequence was $C\bar{E}, B, \bar{C}E, B, C\bar{E}, B, \bar{C}E, B, C\bar{E}, B, \bar{C}E$.

Afterward, participants in both conditions were instructed that they were about to view a new series of tests conducted on the same apparatus as those they had just observed. It was purposely noted that it was the same apparatus, so that participants would rely on the specific timing of the events used during the training phase, which serves as a stronger test against the causal assimilation hypothesis. Participants were told that in the new sequence, the tests were conducted in a continuous manner so that there was no B when one trial began and another ended. Then, all participants saw 13 trials in the following order, as defined using the timing of the training phase: $CE, C\bar{E}, \bar{C}E, C\bar{E}, \bar{C}E, CE, C\bar{E}, \bar{C}E, CE, C\bar{E}, \bar{C}E, C\bar{E}, \bar{C}E$. Because these trials were presented without separating screens, there were five sets of $C\bar{E}$ followed by $\bar{C}E$, which could be collapsed as ce (see Figure 6). To provide a stronger test against causal assimilation, we also included standard CE trials in the test phase, using the timing of the training phase, which would have highlighted the difference between the $C\bar{E}$ – $\bar{C}E$ sequences and a normal CE event, possibly discouraging collapsing of $C\bar{E}$ – $\bar{C}E$ sequences into ce .

After viewing the test phase, participants estimated, in a random order, the frequencies of ce , $\overline{ce}$, $\bar{c}e$, and $c\bar{e}$ for the test phase alone. Participants made estimates for only the test phase because these trials were identical across conditions, allowing for the cleanest test of whether a formed causal hypothesis could be used to interpret ambiguous temporal events. Thus, participants were asked, for instance, “In the continuous stream of tests you just saw, how many times did the disk drop into the toy and, subsequently, the toy played a movie?” The instructions screen separating the training phase from the test phase emphasized the break between these phases and made it easier for participants to recall only the test phase in their estimates. (See a further discussion of this issue in the *Results and Discussion* section, below.) Following the frequency estimates, participants estimated the causal strength of the relationship between c and e for only the test phase, using a sliding scale of -100 (strongly inhibits) to 100 (strongly causes). Participants moved at their own pace through all stages of the experiment. The experiment was administered on eMac computers via the PsyScript experiment program (version 2.1.1).

Results and Discussion

Frequency estimates in the positive and negative training conditions showed that participants were more likely to collapse the $C\bar{E}$ and $\bar{C}E$ trials into ce in the positive training condition than in the negative training condition (see Table 5). A 2 (training: positive vs. negative) $\times 4$ (evidence type) repeated measures ANOVA showed a significant effect of training condition, $F(1, 19) = 4.60, p = .045, \eta_p^2 = .20$, a significant effect of evidence type, $F(3, 57) = 10.6, p < .001, \eta_p^2 = .36$, and a significant interaction, $F(3, 57) = 7.41, p < .001, \eta_p^2 = .28$. The positive training condition resulted in a higher frequency estimate of ce ($M = 6.40$) than did the negative training condition ($M = 4.90; t(19) = 2.36, p = .029, d = 0.53$). In contrast, the negative training condition had a significantly higher frequency estimate of $\bar{c}e$ ($M = 3.70$) than did the positive training condition ($M = 1.55, t(19) = 3.41, p = .003, d = 0.76$, and a significantly higher estimate of $c\bar{e}$ than did the positive training condition ($M = 3.40$ and 2.15 , respectively), $t(19) = 2.30, p = .033, d = 0.52$.

One potential concern with the significant differences between the positive and the negative training conditions reported earlier is that participants could have ignored the instructions to report frequency estimates only from the test phase and could have reported training phase trials as well. CE cases were present in the positive training phase but not in the negative training phase, and $C\bar{E}$ and $\bar{C}E$ cases were present in the negative training phase but not in the positive phase. Including these training trials in frequency estimates could have resulted in differences in frequency estimates across conditions roughly similar to those observed. Note, however, that incorporating the training phase trials would have resulted in higher $\overline{ce}$ estimates in the positive training condition than in the negative training condition. Yet, there was no difference between the two conditions for the $\overline{ce}$ estimates ($p = .12, d = 0.36$).

In addition, note that the total number of actual trials across the training and the test phases was identical between the two conditions (i.e., 6 trials in the training phase and 13 in the test phase). If causal assimilation did not occur in the test phase then the total number of estimated trials should therefore have been identical across conditions. In contrast, if causal assimilation is driving our results, the total number of estimated trials should be fewer in the positive training condition than in the negative training condition, regardless of whether the training phase trials were included in estimates. This is because each $C\bar{E} - \bar{C}E$ sequence of the test phase would have been combined into one event in the positive training condition, whereas those sequences would have been counted as two separate trials in the negative training condition (as outlined in the bottom of Figure 6). That is, because the same total number of trials was observed in the training phases of both conditions, participants including the training phase trials in their estimates would have resulted in the same increase in the total estimated trials across

conditions, and only assimilation in the test phase would account for fewer total numbers of reported trials in the positive training condition than in the negative training condition. Indeed, the positive training condition resulted, on average, in fewer total trials being reported (i.e., sum of ce , $\bar{ce}$, $\bar{ce}$, and $c\bar{e}$; $M = 11.1$, $SD = 3.09$) than did the negative training condition ($M = 13.9$, $SD = 5.64$), $t(19) = 2.14$, $p = .045$, $d = 0.48$. Although these findings cannot be used to definitively determine whether the training phase trials were incorporated into the estimates, it makes it more likely that the demonstrated differences in the frequency estimates between the two conditions must be at least in part due to causal assimilation. These results are echoed when comparing the condition totals with 13, the actual number of trials presented in the test phase alone. Specifically, the summed frequency estimates were significantly less than 13 in the positive training condition, $t(19) = 2.82$, $p = .011$, $d = 0.63$, whereas in the negative training condition they were not significantly different, $p = .50$, $d = 0.15$. These results are again predicted if causal assimilation of ambiguous evidence is occurring.

We also compared the reported frequencies of each event type to the corresponding observed trial frequencies in the test phase (see Table 5 for exact frequencies). For $\bar{ce}$ and $c\bar{e}$ estimates, we conducted one-sample t tests against 5 (i.e., the number of observed trials for each type), using a Bonferroni correction for four independent tests ($\alpha = .0125$). In accordance with a causal assimilation prediction, in the positive training condition both $\bar{ce}$ and $c\bar{e}$ were significantly underestimated, $t(19) = 6.83$, $p < .001$, $d = 1.53$; $t(19) = 5.01$, $p < .001$, $d = 1.12$, respectively. In the negative training condition, they were nonsignificant ($\bar{ce}$: $p = .06$, $d = 0.44$; $c\bar{e}$: $p = .018$, $d = 0.58$).

For the ce estimates, one-sample t tests against 3 (i.e., the number of observed trials) found that ce was overestimated both in the positive, $t(19) = 4.79$, $p < .001$, $d = 1.07$, and the negative training condition, $t(19) = 3.13$, $p = .005$, $d = 0.70$. That participants overestimated the ce frequency even in the negative training condition seems to be an interesting demonstration of the robustness of people's tendency to perceive events that follow sequentially in time as causally linked (e.g., Schlottmann & Shanks, 1992). Because our stimulus materials involve operations of machines, participants may have been more tolerant with longer temporal gaps, resulting in overestimation of ce events even in the negative training condition. Despite any natural perceptual bias, however, participants did report significantly more ce events in the positive training condition than in the negative training condition, as described above, which is more pertinent to the demonstration of causal assimilation.

To summarize, Experiment 4 demonstrated that a causal hypothesis established during the training phase altered the parsing of temporally unfolding events in the test phase. Specifically, when a participant believed C caused E, then a $C\bar{E} - \bar{C}E$ trial sequence was more likely to be viewed as ce . However, when participants believed C did not cause E, the same sequence was more likely to be interpreted as the events $C\bar{E}$ and $\bar{C}E$. The shifts in frequency estimates further resulted in higher overall causal strength estimates for the positive training condition ($M = 46.9$, $SD = 46.1$) than for the negative training condition ($M = 26.1$, $SD = 38.8$), $t(19) = 2.20$, $p = .040$, $d = 0.49$.

It is interesting to note that causal assimilation found in the positive training condition is consistent with recent findings that believing events to be causally related may shorten the perceived temporal duration between these events. For instance, Haggard, Clark, and Kalogeras (2002) found that participants judged the perceived time between a stimulus and an outcome to be smaller when the stimulus was a participant's self-generated causal action than when it was not (see Eagleman & Holcombe, 2002). In addition, Faro, Leclerc, and Hastie (2005) found that historical events that are causally related were judged to be closer together in time than were events that were not causally related. Similarly, our participants in the positive training condition could have perceived the long gap between C and E (resulting from the $\bar{E}$ and C

sections of the $C\bar{E} - CE$ trial sequence) as taking less time than did the timing they initially learned, due to their positive causal belief. Such changes in perception may be the driving force behind the causal assimilation of temporally ambiguous events.

General Discussion

We began this article with the following question: When learning new causal relations, how do people handle observations that are unclear as to the type of covariation evidence they represent between two events? In particular, we investigated whether people spontaneously categorize causal candidates into the four types of covariation evidence when the presence or absence of causal candidates was ambiguous due to the continuous nature of the candidate or due to its temporal properties.

We discussed several reasons why an active belief may not be able to help one categorize ambiguous information in a causal induction task. Previous research within the stereotype and attitude literature has shown that ambiguous data can be assimilated into a belief, if the belief is of high personal importance (e.g., Lord et al., 1979), and the categorization literature has shown that providing a categorization for a stimulus would lead to a biased interpretation of its ambiguous features (e.g., Medin et al., 1993). However, standard causal induction tasks do not involve causal beliefs of such personal importance, nor do they provide the preexisting categorization that could promote the assimilation of ambiguous information. In addition, assimilation during causal induction along temporal dimensions may be unlikely, given that people were shown to be predisposed to interpret temporal lags as negative information, regardless of the governing hypothesis (e.g., Shanks & Dickinson, 1988).

Nonetheless, we demonstrated the assimilation of ambiguous information during causal induction. In Experiment 1, people's frequency estimates of contingency data reflected that information ambiguous along a continuous quantity dimension was spontaneously categorized and assimilated, even in a cold, causal induction task. Experiment 2 showed that this causal assimilation process was moderated by the strength of the underlying causal relationship on which assimilation was based. Experiment 3 showed that causal assimilation could alter the overall perceived relationship between a cause and an effect. Finally, Experiment 4 showed that causal assimilation could occur over temporal sequences.

Although the current demonstrations of causal assimilation were limited to stimuli that are difficult to dichotomize, assimilation may take place even over values that are not inherently ambiguous. For example, assume that a learner has come to believe that caffeine ingestion (i.e., c) causes increased alertness (i.e., e). Subsequently, the learner observed that caffeine was ingested but a person did not become alert (event $c\bar{e}$). In the traditional view of covariation data, this event represents negative evidence. However, it does not have to represent such, if one postulates an alternative inhibitory cause that prevented the presence of e . (See Luhmann & Ahn, 2007, and Schulz & Somerville, 2006, for evidence that people will postulate the existence of inhibitory causes.) In this way, the $c\bar{e}$ event could still be consistent with the learner's original hypothesis by calling to the interference of an inhibitory cause. (For further demonstrations along this line, see Dennis & Ahn, 2001, and Marsh & Ahn, 2006.)

This type of dynamic, top-down, online interpretation of ambiguous covariation information is outside the boundary conditions of most current models of causal induction (e.g., Cheng, 1997; Rescorla & Wagner, 1972) because they all take as input events preclassified into presence or absence of a causal candidate, and it is not the goal of these models to account for how the events are initially classified. Yet, the results presented here suggest that a more complete model of human causal induction should take into consideration how input data changes in interpretation as a novel causal relation is learned because such changes in the

interpretation of data were shown to affect judgments of causal strengths (Experiments 3 and 4). (See Goodman, Mansinghka, & Tenenbaum, 2007, for a modeling approach that aims to describe how bottom-up information gathering processes can interact with causal model formation.)

The current study was conducted as one of the first empirical demonstrations of the assimilation phenomenon during causal induction, to bring attention to the importance of the dynamic interplay between top-down and bottom-up processes during causal induction (e.g., Wisniewski & Medin, 1994). Several further elements of the causal assimilation phenomenon need examination through much more extensive empirical investigations.

One important issue to be addressed in future research is whether causal assimilation occurs online when a reasoner is making individual observations or occurs at the end of observing a learning sequence. An online process of assimilation would proceed such that upon encountering each observation, a learner would rely on his or her hypothesis at that given point in time to determine what that particular observation means. Thus, if one's hypothesis changes during the causal learning process, the same ambiguous observations may be interpreted differently. In contrast, end-of-observation assimilation would entail all ambiguous observations being interpreted in the same way through the causal hypothesis that is most dominant at the end of observations. Understanding which of these accounts is applicable has interesting implications for causal models. For example, if assimilation occurs online, then models of causal reasoning that similarly update causal judgments online would need to account for how the two processes would coexist (e.g., Rescorla & Wagner, 1972). If assimilation occurs at the end of an observational sequence, then it becomes important to understand when these data are assimilated in relation to the time course of causal judgments, especially in relation to models that assess causality after all observations have been made (e.g., Cheng, 1997; Jenkins & Ward, 1965).

Another remaining question is how durable is the interpretation of ambiguous evidence made through assimilation. Experiment 3 demonstrated that participants maintained the initial interpretation of the ambiguous candidate despite a change in the pairing between A and E. This perseverance was great enough to modify their initial hypothesis about the causal strength of the relationship between c and e. This finding would seem to be consistent with the resilience of the interpretation of ambiguous stimuli once it has been made. However, the posttest similarity judgments for A suggest some possible attenuations of this phenomenon. When the form of E that A was paired with switched in the second phase (i.e., the AE/A $\bar{E}$ and A $\bar{E}$ /AE conditions of Experiment 3), the perception of A also started changing (that is, A did not seem more similar to C). Discovering under what conditions assimilation can persist in line with a first interpretation is an important step in understanding how ambiguous information can be incorporated into causal models.

Another unresolved issue is whether causal assimilation is a deliberate, reasoning-based process or a more automatic, implicit process. Assimilation observed in the current studies may be a form of rational judgment, in which an ambiguous observation is placed into the most likely category (C or $\bar{C}$) in an explicit fashion, based on all available information at that point. Using the example from Experiment 1, a participant may see a middle-height bacterium paired with an effect outcome and explicitly reason that because tall bacteria were what were normally paired with the effect and because the middle-height bacterium is paired with the effect, the middle-height bacterium must be an example of a tall bacterium. Alternatively, participants may have just come to perceive the ambiguous stimuli as more like one target category than another in a more implicit fashion.

The fact that we did not find any assimilation in the unknown conditions of Experiment 1 indirectly suggests that causal assimilation as an implicit process is a more likely possibility. In the unknown conditions, the ambiguous stimulus was stated to be completely unknown. If one were explicitly deciding about uncertain events then it would be reasonable to make the most educated guesses about the nature of the unknown stimuli, as one did with the perceptually ambiguous stimulus. In this way, unknown trials should have been incorporated in the same way as ambiguous trials. Thus, the lack of assimilation in the unknown conditions could suggest that this type of explicit reasoning is not occurring for these conditions. However, it is also possible that participants still explicitly reasoned about the unknown stimuli, as in the above description for ambiguous stimuli but did not feel sufficiently justified in making the same categorization decisions. Further experimentation is needed to describe the explicit or implicit nature of the assimilation phenomenon reported here.

One obvious limitation of the current experiments is that assimilation was investigated for cause candidates only. Clearly, more research is needed to investigate whether assimilation also takes place along effect variables. Given the current results on causal assimilation, it seems reasonable to expect that similar phenomena will obtain for effects. One interesting alternative, however, is that if both cause and effect variables are explicitly continuous, causal learning may become tantamount to function learning, or discovering a continuous relationship that exists between two variables. Because causal assimilation seems to take place in an attempt to group values in a categorical manner, assimilation may be unlikely to take place over such continuity.

Finally, it is also not clear whether the demonstrated assimilations are limited to causal learning paradigms. It is possible that the causal assimilation phenomenon is characteristic of a more general covariation detection bias in which events that show similar co-occurrence patterns are grouped together. Alternatively, the assimilation phenomenon as shown here may be specific to causal relations alone. That is, perhaps knowing, alone, that the relation between a middle-height bacteria and an outcome is correlational may not be enough impetus to justify assimilation, and instead, an explicit causal belief must be held to induce assimilation. Further experimentation is needed to choose between these possibilities.

Summary

The findings in the current experiments provide a crucial initial step in understanding how individual causal theories affect the way in which people evaluate and interpret information in a causal reasoning paradigm. In short, causal beliefs can trigger spontaneous categorizations of ambiguous information that would otherwise have no easy parsing according to models of causal induction. Further experimentation is needed to determine what the prevalence of the assimilation phenomenon in causal induction is and what it means for theories of causal learning.

Acknowledgments

We thank the National Institutes of Health for Grant MH57737 awarded to Woo-kyoung Ahn and Yale University for the University Dissertation Fellowship and Psychology Department graduate research funds awarded to Jesseca K. Marsh, all of which helped fund this project. We thank the following dissertation committee members for their comments and insight into the reported experiments: Susan Nolen-Hoeksema, Marcia Johnson, Brian Scholl, and Teresa Treat. We also thank Miriam Bassok and Marc Buehner for their helpful comments on drafts of this article.

References

- Ahn, W.; Marsh, J.K.; Luhmann, C.C. Dynamic interpretations of covariation data. In: Gopnik, A.; Schulz, L., editors. *Causal learning: Psychology, philosophy, and computation*. Oxford, England: Oxford University Press; 2007. p. 280–293.
- Allan L, Tangen J, Wood R, Shah T. Temporal contiguity and contingency judgments: A Pavlovian analogue. *Integrative Psychological and Behavioral Science* 2003;38:214–229.
- Anderson CA. Implicit personality theories and empirical data: Biased assimilation, belief perseverance and change, and covariation detection sensitivity. *Social Cognition* 1995;13:25–48.
- Archambault A, O'Donnell C, Schyns P. Blind to object changes: When learning the same object at different levels of categorization modifies its perception. *Psychological Science* 1999;10:249–255.
- Buehner M, May J. Knowledge mediates the timeframe of covariation assessment in human causal induction. *Thinking and Reasoning* 2002;8:269–295.
- Buehner M, May J. Rethinking temporal contiguity and the judgment of causality: Effects of prior knowledge, experience, and reinforcement procedure. *The Quarterly Journal of Experimental Psychology: Section A* 2003;56:865–890.
- Buehner M, May J. Abolishing the effect of reinforcement delay on human causal learning. *The Quarterly Journal of Experimental Psychology: Section B* 2004;57:179–191.
- Buehner M, McGregor S. Temporal delays can facilitate causal attribution: Towards a general timeframe bias in causal induction. *Thinking and Reasoning* 2006;12:353–378.
- Cheng PW. From covariation to causation: A causal power theory. *Psychological Review* 1997;104:367–405.
- Cheng PW, Novick LR. A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology* 1990;58:545–567. [PubMed: 2348358]
- Collins DJ, Shanks DR. Momentary and integrative response strategies in causal judgment. *Memory and Cognition* 2002;30:1138–1147.
- Corneille O, Huat J, Becquart E, Brédart S. When memory shifts toward more typical category exemplars: Accentuation effects in the recollection of ethnically ambiguous faces. *Journal of Personality and Social Psychology* 2004;86:236–250. [PubMed: 14769081]
- Corneille O, Leyens J, Yzerbyt V, Walther E. Judgeability concerns: The interplay of information, applicability, and accountability in the overattribution bias. *Journal of Personality and Social Psychology* 1999;76:377–387.
- Dennis MJ, Ahn W. Primacy in causal strength judgments: The effect of initial evidence for generative versus inhibitory relationships. *Memory and Cognition* 2001;29:152–164.
- Dunning D, Sherman DA. Stereotypes and tacit inference. *Journal of Personality and Social Psychology* 1997;73:459–471. [PubMed: 9294897]
- Eagleman D, Holcombe A. Causality and the perception of time. *Trends in Cognitive Sciences* 2002;6:323–325. [PubMed: 12140076]
- Eberhardt J, Dasgupta N, Banaszynski T. Believing is seeing: The effects of racial labels and implicit beliefs on face perception. *Personality and Social Psychology Bulletin* 2003;29:360. [PubMed: 15273013]
- Faro D, Leclerc F, Hastie R. Perceived causality as a cue to temporal distance. *Psychological Science* 2005;16:673–677. [PubMed: 16137251]
- Gawronski B, Geschke D, Banse R. Implicit bias in impression formation: Associations influence the construal of individuating information. *European Journal of Social Psychology* 2003;33:573–589.
- Goldstone RL, Lippa Y, Shiffrin RM. Altering object representations through category learning. *Cognition* 2001;78:27–43. [PubMed: 11062321]
- Goodman, N.D.; Mansinghka, V.K.; Tenenbaum, J.B. In: McNamara, D.S.; Trafton, J.G., editors. *Learning grounded causal models; Proceedings of the 29th annual conference of the Cognitive Science Society*; Austin, TX: Cognitive Science Society; 2007. p. 305–310.
- Greville WJ, Buehner MJ. The influence of temporal distributions on causal induction from tabular data. *Memory and Cognition* 2007;35:444–453.

- Haggard P, Clark S, Kalogeras J. Voluntary action and conscious awareness. *Nature Neuroscience* 2002;5:382–385.
- Hagmayer Y, Waldmann M. How temporal assumptions influence causal judgments. *Memory and Cognition* 2002;30:1128–1137.
- Huttenlocher J, Hedges L, Vevea J. Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General* 2000;129:220–241. [PubMed: 10868335]
- Jenkins HM, Ward WC. Judgment of contingency between responses and outcomes. *Psychological Monographs: General & Applied* 1965;79:17.
- Kunda Z, Sherman-Williams B. Stereotypes and the construal of individuating information. *Personality and Social Psychology Bulletin* 1993;19:90–99.
- Lagnado D, Sloman S. Time as a guide to cause. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2006;32:451–460.
- Lien Y, Cheng PW. Distinguishing genuine from spurious causes: A coherence hypothesis. *Cognitive Psychology* 2000;40:97–137.
- Livingston KR, Andrews JK, Harnad S. Categorical perception effects induced by category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1998;24:732–753.
- Lord CG, Ross L, Lepper MR. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology* 1979;37:2098–2109.
- Luhmann CC, Ahn W. BUCKLE: A model of unobserved cause learning. *Psychological Review* 2007;114:657–677. [PubMed: 17638500]
- Marsh, JK.; Ahn, W. In: Alterman, R.; Kirsh, D., editors. Interpretation of ambiguous information in causal induction; Proceedings of the 25th annual conference of the Cognitive Science Society; Mahwah, NJ: Erlbaum; 2003. p. 775-780.
- Marsh JK, Ahn W. Order effects in contingency learning: The role of task complexity. *Memory and Cognition* 2006;34:568–576.
- Medin DL, Goldstone RL, Gentner D. Respects for similarity. *Psychological Review* 1993;100:254–278.
- Michotte, A. The perception of causality. Miles, TR., translator. London: Methuen; 1946.
- Munro GD, Ditto PH. Biased assimilation, attitude polarization, and affect in reactions to stereotyped-relevant scientific information. *Personality and Social Psychology Bulletin* 1997;23:636–653.
- Munro GD, Ditto PH, Lockhart LK, Fagerlin A, Gready M, Peterson E. Biased assimilation of sociopolitical arguments: Evaluating the 1996 US presidential debate. *Basic and Applied Social Psychology* 2002;24:15–26.
- Murphy GL. Comprehending complex concepts. *Cognitive Science* 1988;12:529–562.
- Quinn P, Schyns P, Goldstone R. The interplay between perceptual organization and categorization in the representation of complex visual patterns by young infants. *Journal of Experimental Child Psychology* 2006;95:117–127. [PubMed: 16797025]
- Rescorla, R.; Wagner, A. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black, AH.; Prokasy, WF., editors. Classical conditioning II: Current research and theory. New York: Appleton-Century-Crofts; 1972. p. 64-99.
- Schlottmann A, Shanks D. Evidence for a distinction between judged and perceived causality. *The Quarterly Journal of Experimental Psychology: Section A* 1992;44:321–342.
- Schulz LE, Sommerville J. God does not play dice: Causal determinism and preschoolers' causal inferences. *Child Development* 2006;77:427–442. [PubMed: 16611182]
- Schyns P, Goldstone RL, Thibaut JP. The development of features in object concepts. *Behavioral and Brain Sciences* 1998;21:1–54. [PubMed: 10097010]
- Schyns P, Oliva A. Dr. Angry and Mr. Smile: When categorization flexibly modifies the perception of faces in rapid visual presentations. *Cognition* 1999;69:243–265. [PubMed: 10193048]
- Schyns PG, Rodet L. Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1997;23:681–696.
- Sedikides C, Anderson CA. Causal explanations of defection: A knowledge structure approach. *Personality and Social Psychology Bulletin* 1992;18:420–429.

- Shanks, D.; Dickinson, A. Associative accounts of causality judgment. In: Bower, GH., editor. *The psychology of learning and motivation*. Vol. 21. San Diego, CA: Academic Press; 1988. p. 229-261.
- Shanks, D.; Lopez, F.; Darby, R.; Dickinson, A. Distinguishing associative and probabilistic contrast theories of human contingency judgment. In: Shanks, DR.; Medin, DL.; Holyoak, KJ., editors. *The psychology of learning and motivation*. Vol. 34. San Diego, CA: Academic Press; 1996. p. 265-311.
- Shanks D, Pearson S, Dickinson A. Temporal contiguity and the judgment of causality by human subjects. *The Quarterly Journal of Experimental Psychology: Section B* 1989;41:139–159.
- Tversky A, Kahneman D. Judgment and uncertainty: Heuristics and biases. *Science* 1974 September 27;185:1124–1131. [PubMed: 17835457]
- Waldmann MR, Hagmayer Y. Categories and causality: The neglected direction. *Cognitive Psychology* 2006;53:27–58. [PubMed: 16497289]
- Williams P, Tarr MJ. RSVP: Experimental control software for MacOS. 2001 Retrieved June 15, 2000.
- Wisniewski EJ, Medin DL. On the interaction of theory and data in concept learning. *Cognitive Science* 1994;18:221–281.
- Yzerbyt V, Leyens J, Schadron G. Social judgeability and the dilution of stereotypes: The impact of the nature and sequence of information. *Personality and Social Psychology Bulletin* 1997;23:1312.
- Yzerbyt V, Schadron G, Leyens J, Rocher S. Social judgeability: The impact of meta-informational cues on the use of stereotypes. *Journal of Personality and Social Psychology* 1994;66:48–55.

Figure 1.

Pretested materials used as stimuli in Experiments 1 through 3. Each set's far left and far right items represent the ends of the spectrum used in the pretest for that stimulus set. To the right of each set is the mean (M) and standard deviation (SD) for the ratings of the A form of the stimuli from the pretest. Five was the midpoint of the rating scale. Below this information are presented the values for the variations of C (candidate cause), A (ambiguous candidate cause), and $\bar{C}$ (polar opposite of candidate cause) used in the actual experiments. The units used to describe these variations are parameter settings used in the Canvas image creation software program (version 5), and they are presented here to illustrate the distributions of the variations within each material set.

Figure 2.

Mean perceptual similarity posttest ratings for Experiment 1. The dotted line indicates the midpoint of the similarity scale. The two bars on the left represent the conditions where A (ambiguous candidate cause) was present, and the two bars on the right indicate the two conditions where the unknown causal candidate (U) was presented. Ratings close to 9 represent high similarity to C (candidate cause), and ratings close to 1 represent high similarity to $\bar{C}$ (polar opposite of the causal candidate). Error bars represent standard error.

Figure 3.

Difference between observed and reported trial frequencies across conditions of Experiment 2. Error bars represent standard error. ce = cause present, effect present; $c\bar{e}$ = cause absent, effect present; $\bar{c}e$ = cause present, effect absent; $\bar{c}\bar{e}$ = cause absent, effect absent (see text for details).

Figure 4.

Mean second phase causal strength ratings for Experiment 3. Error bars represent standard error. AE = phase pairing ambiguous cause and the presence of the effect; AĒ = phase pairing ambiguous cause and the absence of the effect.

Figure 5.

Timing for the child's toy materials of Experiment 4. CE = cause present, effect present; $\overline{CE}$ = cause absent, effect absent; $\overline{CE}$ = cause absent, effect present; $\overline{CE}$ = cause present, effect absent.

Figure 6.

Test phase sequence and timing for the toy materials of Experiment 4. Top: Shows details of events and the timing that occurred for $\bar{C}E$ and CE trials. Bottom: Dotted lines between observed trials represent onset of a new trial from the experimenter's perspective; there were no demarcations of trials for participants.

Table 1

The Four Types of Information Used to Evaluate a Causal Relationship

	Candidate	Candidate
Effect	cause present	cause absent
Effect present	ce	$\bar{c}e$
Effect absent	$c\bar{e}$	$c\bar{e}$

Note. The bar over the letters represents the absence of an event. c = candidate cause; e = effect.

Table 2

Observed and Estimated Trial Frequencies for Experiment 1

Trial type	Conditions			
	AE	UE	AĒ	UĒ
Observed frequencies				
CE	18	18	18	18
CE	18	18	18	18
CE	2	2	2	2
CĒ	2	2	2	2
Ambiguous or unknown trials paired with the effect	20	20	0	0
Ambiguous or unknown trials paired without the effect	0	0	20	20
Possible estimates if fully assimilating ambiguous trials ^a				
CE	38	18	18	18
CE	18	18	38	18
CE	2	2	2	2
CĒ	2	2	2	2
Participants' frequency estimates				
ce	27.5 ± 9.5	20.0 ± 5.1	19.3 ± 7.1	18.1 ± 4.0
ce	20.7 ± 7.3	19.2 ± 4.6	28.8 ± 9.5	17.8 ± 4.1
cē	3.30 ± 3.6	2.30 ± 1.5	3.70 ± 6.3	2.75 ± 1.3
cē	3.35 ± 3.6	3.50 ± 4.1	3.35 ± 4.5	2.40 ± 1.1

Note. Bold numbers indicates cells that may absorb A trials if assimilation occurs. Upper case lettering denotes observed trials and lower case lettering denotes reported trials. CE and ce = cause present, effect present; $\overline{CE}$ and $\overline{CE}$ = cause absent, effect absent; $\overline{CE}$ and $\overline{ce}$ = cause absent, effect present; $\overline{CE}$ and $\overline{cē}$ = cause present, effect absent; AE = ambiguous cause present, effect present; $\overline{AE}$ = ambiguous cause present, effect absent; UE = unknown cause present, effect present; $\overline{UE}$ = unknown cause present, effect absent.

^aTable 2 illustrates an extreme version of possible estimates in which all As get assimilated. It is not expected that all 20 instances of A would be incorporated as shown in this possibility.

Table 3

Observed and Reported Trial Frequencies for Experiment 2

Trial type	Condition			
	Zero	Weak	Moderate	Perfect
Observed frequencies				
CE	10	13	16	20
CE	10	13	16	20
CE	10	7	4	0
CĒ	10	7	4	0
AE	20	20	20	20
AĒ	0	0	0	0
Possibility 1: Incorporating all AE trials positively				
ce	10	13 + 20	16 + 20	20 + 20
CE	10	13	16	20
cē	10	7	4	0
cĒ	10	7	4	0
Possibility 2: Incorporating AE trials as positive and negative evidence				
ce	10	13 + 13	16 + 16	20 + 20
ce	10	13	16	20
cē	10	7 + 7	4 + 4	0 + 0
cĒ	10	7	4	0
Possibility 3: Incorporating AE trials proportionately				
ce	10	13 + 13	16 + 16	20 + 20
ce	10	13	16	20
cē	10	7	4	0
cĒ	10	7	4	0
Participants' frequency estimates				
ce	13.2 ± 5.3	15.3 ± 7.7	21.5 ± 8.4	28.3 ± 9.8
ce	11.8 ± 5.2	12.4 ± 4.9	15.6 ± 6.8	20.9 ± 6.1
cē	11.5 ± 4.9	8.3 ± 3.0	6.9 ± 6.4	3.6 ± 7.6
cĒ	11.2 ± 4.3	8.7 ± 3.8	5.5 ± 3.0	1.5 ± 4.3

Note. The numbers in bold represent the numbers of ambiguous trials that would be incorporated into that frequency estimate. Upper case lettering denotes observed trials and lower case lettering denotes reported trials. CE and ce = cause present, effect present; **CE** and **ce** = cause absent, effect present; CĒ and cē = cause absent, effect present; CĒ and cĒ = cause present, effect absent; AE = ambiguous cause present, effect present; AĒ = ambiguous cause present, effect absent; UE = unknown cause present, effect present; UĒ = unknown cause present, effect absent.

Table 4

Number of Observed Trials in the First and Second Phases of Experiment 3 and Mean Causal Strength Estimates at the End of Each Phase

Trial type	Conditions			
	AE/AE	AĒ/AE	AE/AĒ	AĒ/AĒ
Observed contingencies in first phase				
CE	18	18	18	18
CĒ	18	18	18	18
CE	2	2	2	2
CĒ	2	2	2	2
AE	20 [ce]	0	20 [ce]	0
AĒ	0	20 [cē]	0	20 [cē]
Causal strength estimates of the first phase	63.1 ± 32.6	57.0 ± 33.2	61.3 ± 32.6	66.9 ± 26.2
Observed contingencies in second phase				
CE	9	9	9	9
CĒ	9	9	9	9
CE	1	1	1	1
CĒ	1	1	1	1
AE	10	10	0	0
AĒ	0	0	10	10
Hypothetical frequency estimates of the second phase when full assimilation occurs in line with the first phase assimilation				
ce	9 + 10	9	9	9
cē	9	9	9	9 + 10
cē	1	1 + 10	1	1
cē	1	1	1 + 10	1
ΔP	0.85	0.35	0.35	0.85
Causal strength estimates of the second phase	60.7 ± 32.6	48.3 ± 31.5	53.6 ± 29.3	69.8 ± 27.9

Note. ce or $\overline{c\bar{e}}$ in brackets indicates predicted interpretation of ambiguous trials. Bold numbers indicates evidence that would be interpreted as negative if assimilation occurs. The +10 notation indicates where second phase ambiguous trials would be assimilated. Upper case lettering denotes observed trials and lower case lettering denotes reported trials. CE and ce = cause present, effect present; $\overline{C\bar{E}}$ and $\overline{c\bar{e}}$ =cause absent, effect absent; $\bar{C}E$ and $\bar{c}e$ = cause absent, effect present; $\bar{C}\bar{E}$ and $\bar{c}\bar{e}$ = cause present, effect absent; AE = ambiguous cause present, effect present; AĒ = ambiguous cause present, effect absent; UE = unknown cause present, effect present; UĒ = unknown cause present, effect absent.

Table 5

Observed and Reported Trial Frequencies for Experiment 4

Trial type	Training condition	
	Positive training	Negative training
Observed trials in the training phase		
CE	3	0
CE	3	0
CE	0	3
CĒ	0	3
Observed trials in the test phase		
CE	3	3
CE	0	0
CE	5	5
CĒ	5	5
Participants' frequency estimates		
Ce	6.40 ± 3.17	4.90 ± 2.71
ce	0.95 ± 1.90	1.85 ± 2.80
cē	1.55 ± 2.26	3.70 ± 2.94
cĕ	2.15 ± 2.54	3.40 ± 2.76

Note. Bold numbers represents predicted site of assimilation. Upper case lettering denotes observed trials and lower case lettering denotes reported trials. CE and ce = cause present, effect present; $\overline{CE}$ and $\overline{ce}$ = cause absent, effect absent; CĒ and cē = cause absent, effect present; CE and cĕ = cause present, effect absent; AE = ambiguous cause present, effect present; AĒ = ambiguous cause present, effect absent; UE = unknown cause present, effect present; UĒ = unknown cause present, effect absent.