Evidence for online processing during causal learning

Pei-Pei Liu · Christian C. Luhmann

Published online: 9 December 2014
© Psychonomic Society, Inc. 2014

Abstract Many models of learning describe both the end product of learning (e.g., causal judgments) and the cognitive mechanisms that unfold on a trial-by-trial basis. However, the methods employed in the literature typically provide only indirect evidence about the unfolding cognitive processes. Here, we utilized a simultaneous secondary task to measure cognitive processing during a straightforward causal-learning task. The results from three experiments demonstrated that covariation information is not subject to uniform cognitive processing. Instead, we observed systematic variation in the processing dedicated to individual pieces of covariation information. In particular, observations that are inconsistent with previously presented covariation information appear to elicit greater cognitive processing than do observations that are consistent with previously presented covariation information. In addition, the degree of cognitive processing appears to be driven by learning per se, rather than by nonlearning processes such as memory and attention. Overall, these findings suggest that monitoring learning processes at a finer level may provide useful psychological insights into the nature of learning.

Keywords Causal learning · Human learning

When forming beliefs about causal relationships, people frequently rely on information about whether the potential cause tends to occur with its potential effect. For example, imagine a person becoming sick each time she eats a certain kind of nut, but not after eating other sorts of food. Eventually, she will likely believe that the nuts cause her sickness and avoid eating them in the future. Prevailing theories of causal learning (Busemeyer, 1991; Cheng, 1997; Einhorn & Hogarth, 1986;

Jenkins & Ward, 1965; Rescorla & Wagner, 1972; Schustack & Sternberg, 1981; White, 2002) have formalized such evidence using a covariation matrix such as that shown in Fig. 1. Assuming that events can be either present or absent, there are four types of observations, each represented by a different cell in the covariation matrix. For example, cell A of the matrix represents the frequency with which the cause and the effect are both present. Evidence from cells A and D (positive covariation) tends to suggest that the cause generates the outcome, whereas evidence from cells B and C (negative covariation) tends to suggest that the cause prevents the outcome (see Luhmann & Ahn, 2011, for systematic exceptions).

A variety of learning theories, both associative and nonassociative, describe the processes that operate to transform sequences of covariation information into causal judgments (Danks, Griffiths, & Tenenbaum, 2003; Hogarth & Einhorn, 1992; Lu, Rojas, Beckers, & Yuille, 2008; Luhmann & Ahn, 2007; Mackintosh, 1975; Pearce, 1994; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Wagner, 1981). These models specify what causal judgments should be made on the basis of any given set of covariation information, and provide details of the constituent processes that transform individual pieces of covariation information into the predicted judgments. The most prominent of these models have been derived from work on animal learning. For example, the model proposed by Rescorla and Wagner has been particularly influential. According to this model, the association between a potential cause and its potential effect is updated as individual pieces of covariation information are encountered, according to the extent to which learners' predictions are violated. The greater the violation (i.e., the prediction error), the greater the adjustment made to beliefs. According to the Rescorla–Wagner model, when first encountering a novel cause, the learner's predictions will be relatively inaccurate and will lead to large prediction errors. These prediction errors will, in turn, induce a rapid increase in the associative strength. As the learner

P.-P. Liu (✉) · C. C. Luhmann
Department of Psychology, Stony Brook University, Stony Brook,
NY 11794-25001, USA
e-mail: pei-pei.liu@stonybrook.edu

		Effect	
		Present	Absent
Cause	Present	A	B
	Absent	C	D

Fig. 1 Typical classification of covariation information. The contingency matrix classifies each observation into one of four categories, based on the presence and absence of the potential cause and the potential effect.

repeatedly encounters similar observations, the predictions will become more accurate, and the changes in the associative strength will become smaller. At some point, when the effect is fully expected to follow the cause, the presence of the outcome will produce neither prediction errors nor changes to the learner's beliefs.

For many years, Rescorla–Wagner and related associative models (Mackintosh, 1975; Pearce, 1994; Pearce & Hall, 1980; Wagner, 1981) were unique in their ability to make predictions about the learning process. However, a variety of process accounts have now been put forth within nonassociative frameworks (Danks et al., 2003; Hogarth & Einhorn, 1992; Lu et al., 2008; Luhmann & Ahn, 2007). For example, Hogarth and Einhorn proposed that learning could be viewed as an example of iterated anchoring and adjustment, in which individual pieces of evidence invoked adjustments, and requests for judgments prompted the creation of new anchors. Luhmann and Ahn (2007) described BUCKLE, a process model designed to explain people's apparent ability to make judgments about causes they have never directly observed. Like Rescorla–Wager, BUCKLE adjusts causal beliefs by computing prediction errors each time a piece of covariation information is encountered. Unlike Rescorla–Wager, however, BUCKLE uses Bayesian inference to “guess” when unobserved causes are present, and then learns as if this inferred information had actually been observed.

In the spirit of these Bayesian models, we have recently put forth a proposal (Luhmann & Ahn, 2011) suggesting that learners may attempt to generate causal explanations of individual pieces of covariation information as they are encountered, especially when learners encounter covariation information that conflicts with their prior causal beliefs. When such information is encountered, learners either use the contradictory information to revise the causal beliefs or construct explanations that allow the existing causal belief to be maintained (essentially, explaining away the contradictory covariation information).

The need for direct evidence

Each of the models described above outlines exactly how learners update their beliefs each time a single piece of covariation information is encountered. However, despite such widespread theoretical commitment to describing the processes that underlie learning, research on learning has a long tradition of focusing on the final learning product rather than on trial-by-trial processing. We believe this focus is partly due to the reliance on what we refer to as the *input–output paradigm*. The widely used input–output paradigm begins with the experimenter specifying a set of covariation information (the input) and then measuring learners' judgments regarding the relationship between events (the output). Each model under consideration is then evaluated with regard to its ability to account for the input–output relationship, in order to be considered empirically viable. Though this technique clearly represents a strong initial test of learning models, it has failed to generate direct evidence regarding the underlying processes that may be unfolding as learners encounter individual pieces of information. Researchers have claimed that specific variants of the input–output paradigm are able to demonstrate the existence of trial-by-trial processing (e.g., Chapman & Chapman, 1969; Dennis & Ahn, 2001; Katagiri, Kao, Simon, Castro, & Wasserman, 2007; Lopez, Shanks, Almaraz, & Fernandez, 1998; Luhmann & Ahn, 2011; Yates & Curley, 1986); however, subsequent work has demonstrated this not to be the case (e.g., Danks & Schwartz, 2006; Marsh & Ahn, 2006; Stout, Amundson, & Miller, 2005). Regardless, the very nature of the standard input–output paradigm clearly places an overwhelming emphasis on accurately predicting the transformation from covariation to judgment, with far less regard for the psychological mechanisms by which such a transformation comes about. More recently, researchers have started to investigate the trial-by-trial processing underlying learning, by employing nontraditional learning paradigms (De Houwer & Beckers, 2003; González-Martín, Cobos, Morís, & López, 2012; Le Pelley, Vadillo, & Luque, 2013; Vadillo & Matute, 2010; Waldmann & Walker, 2005), eyetracking (Beesley & Le Pelley, 2011; Kruschke, Kappenman, & Hetrick, 2005; Le Pelley, Beesley, & Griffiths, 2011; Wills, Lavric, Croft, & Hodgson, 2007), and electroencephalography (Luque, López, Marco-Pallares, Càmarà, & Rodríguez-Fornells, 2012; Morís, Luque, & Rodríguez-Fornells, 2013; Walsh & Anderson, 2011). These studies represent a refreshing alternative to the traditional input–output paradigm. Nonetheless, the amount of work directly focusing on process is still not comparable to the amount of work focused on the final learning products.

In the present study, we explored a novel paradigm in an attempt to examine the learning process more directly than is allowed by the standard input–output paradigm. Specifically, we wished to employ a method that would allow us to

investigate the cognitive processes engaged on a trial-by-trial basis. Accordingly, we had learners perform a simple secondary task that would allow us to selectively monitor learning-related processing. The secondary task that we chose was a relatively easy auditory classification task that has previously been employed as a measure of general cognitive processing. For example, Posner and Boies (1971) successfully employed a similar secondary task in order to distinguish between simple encoding processes and more complex comparative processes during a letter-matching task.

In a recent study (Liu & Luhmann, 2013), we demonstrated the value of secondary-task paradigms in elucidating the processes that underlie learning phenomena. We used a secondary tone discrimination task to monitor the cognitive processes unfolding in a learning task involving cue competition effects: when beliefs about the relationship between one cause and an effect are constrained by beliefs about the relationship between a second cause and that same effect (e.g., Kamin, 1969). During the learning task, participants had to discriminate and respond to auditory probes presented sporadically throughout the trial sequence. We found that responses to the probes were slower when a cause previously presented in isolation was first paired with a novel cause. This slowing was interpreted as reflecting the learning-related cognitive processes invoked by the presence of a novel, causally ambiguous cause. This interpretation was supported by the additional finding that individual differences in this reaction time (RT) effect were related to individual differences in learning; participants who exhibited greater blocking effects also had slower responses to the probes. Furthermore, the pattern of RTs was not epiphenomenal. A second study showed that blocking was disrupted when a cognitive load was imposed as novel causes were first encountered, but not when cognitive load was imposed at other times.

This prior application of the secondary-task paradigm demonstrated its value in providing insight into the operation and time course of the cognitive processes underlying cue competition effects. In the present study, we focused on more simplistic learning tasks that have been thoroughly studied using the standard input–output paradigm. In Experiment 1, we presented learners with relatively homogeneous sequences of covariation information. These sequences consisted of either a majority of positive covariation information (cells A and D) and a minority of negative covariation information (cells B and C), or vice versa. The auditory probes were then delivered during a small number of the majority trials and during a small number of the minority trials. In Experiment 2, we returned to the standard order-effect paradigm (e.g., Dennis & Ahn, 2001; Lopez et al., 1998), in which positive covariation information is presented early in a sequence and negative covariation information is presented late in a sequence, or vice versa. When a positive–negative sequence was presented, probes were delivered on positive trials in the first half and on

negative trials in the second half. In Experiment 3, we investigated whether the auditory probes were more closely related to a learner's subjective beliefs about causes or to the objective covariation information.

Experiment 1

In Experiment 1, participants received a sequence of trials either consisting of predominantly positive (cells A and D) covariation information intermixed with a small number of negative (cells B and C) trials or consisting of predominantly negative covariation information intermixed with a small number of positive trials. The majority trials were expected to strongly determine participants' expectations, and the minority trials served to occasionally violate these expectations. If participants develop expectancies on the basis of the more frequent majority trials and find the minority trials to be unexpected, then responses to the auditory probes might be different, depending on whether they are delivered during minority or majority trials. In contrast, if participants simply record the covariation for later processing, these different trial types might be processed in similar manners.

Method

Participants A group of 70 undergraduate students at Stony Brook University participated in the experiment for partial course credit. The participants were randomly assigned to one of the two conditions (positive/negative), with 35 in each condition. Two participants (one in each condition) missed all of the tones during minority trials and therefore were excluded from subsequent analyses.

Stimuli Participants were instructed that they would learn about whether a given medication was related to a hypothetical disease called Chromato. To do so, they were provided with a set of hypothetical medical records, each of which included the information about whether the patient had taken the medication and whether the patient had developed the disease. Each patient's record included information about whether or not the medication was taken and whether or not the patient had Chromato (Fig. 2). Medications of four different colors were used to represent the potential causes. When presented, the medications either were or were not crossed out, with the former representing the absence of the medications. Pictures of a sick-looking and a smiling face were used to represent the potential effects, with the sick face indicating the presence of the disease and the smiling face indicating the absence of the disease.

The auditory probes delivered throughout the experiment consisted of three different frequencies (high, medium, and low). The high tone had a frequency of 3520 Hz, the medium

Fig. 2 Example of a trial utilized in the present experiments. The novel medications (left) acted as potential causes of the fictional disease, Chromato (right). Each cause was a different color, to help differentiate them.

tone had a frequency of 880 Hz, and the low tone had a frequency of 220 Hz. Each tone lasted for 50 ms.

Design Participants were randomly assigned to either the positive or the negative condition. Participants received four sequences of trials, with each sequence describing a different medication. In the positive condition, each sequence consisted of 12 trials from each of the positive cells of the covariation matrix (cells A and D) and three trials from each of the negative cells (cells B and C, $\Delta P = .6$). The negative condition was the mirror opposite, with cells B and C swapping roles with cells A and D ($\Delta P = -.6$). The order in which the trials were presented in each condition was pseudorandom, so as to ensure that the minority trials were spread out over the entire sequence. In addition, the first minority trial of a sequence could only appear after at least two majority trials.

In each sequence, an auditory probe was played during one trial from each cell of the covariation matrix—that is, two majority and two minority trials. At least four majority trials were presented before the first probe was played on a minority trial. The frequencies of the probes were randomly determined. Participants had to press one of three keys according to the pitch of the probe. The RTs to the tones were the main dependent measure.

Procedure Participants were first given an overview of their basic task. Before starting, they were familiarized with the auditory probes. The three different pitches were presented, and participants were required to simply identify whether the presented tone was high, medium, or low, using three keys on the computer keyboard. Feedback was presented after each response, and participants continued until they were able to discriminate the three tones (given how dissimilar the tones were, this portion of the task was short). Participants then completed a brief sequence of practice trials (involving a medication not used in the main task) to familiarize them with the process of making tone discrimination responses during the learning task and to allow them practice with making the causal judgments.

On each trial of the experiment, the presence or absence of the medication was first presented on the left side of the screen for 1,000 ms. The effect (the sick or happy face) was then presented on the right side of the screen (so that both the cause and the effect information were present on the screen) for 2,000 ms. On most trials, the screen was then cleared and a blank screen was presented for 2,500–3,500 ms before presenting the next trial. However, on those trials with an auditory probe, the probe was delivered 1,000 ms after the onset of the outcome. Participants then had up to 3 s to respond to the probe, and the information remained on the screen until participants responded. If they did not respond after 3 s, the trial ended and they moved on to the next trial.

After completing each sequence of trials, participants judged the causal relationship between the medication and the disease on a scale from –100 to 100, where 100 represented *definitely causes the disease* and –100 represented *definitely prevents the disease*. Participants made judgments by typing numbers into a dialog box. They completed four sequences of these trials (either four positive sequences or four negative sequences), learning about a new medication (i.e., a different color) in each sequence.

Results

The participants' judgments and RTs in the four sequences were averaged. Overall, they judged the cause to be strongly generative in the positive condition ($M = 72.05$, $SD = 26.74$), with judgments being significantly greater than zero [$t(33) = 15.71$, $p < .001$, $d = 2.69$]. Participants judged the cause to be strongly preventative in the negative condition ($M = -68.33$, $SD = 36.96$), with judgments being significantly less than zero [$t(33) = 10.78$, $p < .001$, $d = 1.85$]. Judgments in the two conditions were also significantly different from each other [$t(66) = 33.56$, $p < .001$, $d = 1.81$].

Figure 3 shows the mean RTs to the probes occurring during positive and negative trials in the two conditions. We performed a 2 (Condition: positive/negative) by 2 (Trial Type: positive/negative) analysis of variance (ANOVA) with repeated measures on the latter factor. The analysis yielded a significant interaction between condition and trial type [$F(1, 65) = 7.53$, $p < .01$, $\eta^2 = .10$]. Neither of the main effects was significant, $F_s < 1$. In the positive condition, responses to the probes were marginally faster during positive than during negative trials [$t(32) = 1.98$, $p = .057$, $d = 0.34$]. In the negative condition, responses to the probes were marginally faster during negative than during positive trials [$t(33) = 1.93$, $p = .063$, $d = 0.33$]. When we collapsed across conditions, RTs to the probes during majority trials and minority trials were significantly different [$t(66) = 2.77$, $p < .01$, $d = 0.34$]. That is, regardless of the condition, responses were faster to probes delivered during majority trials than to probes delivered during minority trials. Moreover, accuracy to tones on majority

Fig. 3 Results from Experiment 1: Reaction times were faster to condition-consistent information, and slower to condition-inconsistent information. Error bars indicate standard errors of the means.

($M = .84$, $SD = .30$) and minority ($M = .89$, $SD = .19$) trials did not differ [$t(66) = 1.18$, $p > .2$], suggesting that the RT difference between majority and minority trials was not the result of a speed–accuracy trade-off.

Discussion

Participants in Experiment 1 were faster to respond to auditory probes delivered during expected, majority trials than during unexpected, minority trials. This result suggests that observations encountered during the learning sequence were not necessarily subjected to equivalent processing. Instead, it appears that the unexpected nature of the infrequent minority trials demanded greater cognitive processing than did the expected, frequent majority trials. This has several immediate implications. First, it suggests that participants develop expectancies over the course of the learning sequence rather than simply recording observations passively. Second, it suggests that the processing of individual observations depends, in part, on the relationship between the current observation and previously encountered observations. Experiment 2 was designed to replicate and extend these findings.

Experiment 2

In Experiment 2, we explored the traditional order-effect paradigm, in which positive covariation information is presented early in the sequence and negative covariation information is presented later in the sequence, or vice versa. The auditory probes were delivered throughout the sequence, allowing us to investigate the time course of cognitive processing. If the results of Experiment 1 reflected processing evoked by violations of expectation, then secondary-task RTs should get faster as learners' expectations strengthen over the

course of the first half of the sequence. As they first enter the second half of the sequence, responses should slow due to the unexpected change in covariation information. To the extent that this new, contradictory information changes learners' beliefs, RTs might again drop late in the sequence.

Method

Participants A group of 24 undergraduate students at Stony Brook University participated in the experiment for partial course credit.

Materials and design The task and stimuli were identical to those used in Experiment 1. The only difference was in the design of the trial sequences. There were two orders of trial presentations. In the positive–negative sequence, the first half of the sequence consisted of seven trials from each of the positive cells of the covariation matrix (cells A and D) and one trial from each of the negative cells (cells B and C). The second half of the sequence consisted of seven trials from each of the negative cells and one from each of positive cells. The negative–positive sequence was the mirror opposite, with cells B and C swapping roles with cells A and D.

Each participant completed four sequences of trials (two positive–negative and two negative–positive), each employing a different medication. The order of the four sequences was counterbalanced across participants. In each sequence, two trials from each half of the sequence were probed by auditory tones. Each type of trial (cells A, B, C, and D) was probed once. For example, in the positive-first sequence, one cell A trial and one cell D trial were probed in the first half of the sequence, whereas one cell B trial and one cell C trial were probed in the second half of the sequence. In each half of the sequence, one probe was delivered early in the half, and the other probe was delivered late in the half.

Results

Participants' judgments for positive–negative sequences were averaged, as were their judgments for negative–positive sequences. The judgments elicited by the two orders were significantly different [$t(22) = 5.47$, $p < .001$, $d = 1.14$]. Participants judged the medication to be generative in the positive–negative sequences ($M = 19.91$, $SD = 15.85$) [$t(22) = 3.84$, $p < .001$, $d = 1.26$] and preventative in the negative–positive sequences ($M = -21.97$, $SD = 28.53$) [$t(22) = -3.69$, $p < .001$, $d = 0.77$].

As can be seen in Fig. 4, responses to the probes over the course of learning showed similar patterns in the positive–negative and negative–positive sequences. A 4 (Position of the Probes: early in the first half, late in the first half, early in the second half, late in the second half) by 2 (Sequence Type: positive–negative vs. negative–positive) repeated measures

Fig. 4 Results of Experiment 2. Reaction times decreased over the course of the first half of the sequence. When encountering the initial trials of the inconsistent second half, reaction times increased, but over the course of the second half, reaction times again decreased. Error bars indicate standard errors of the means.

ANOVA was conducted. The main effect of probe position was significant [$F(3, 66) = 10.12, p < .001, \eta^2 = .55$]. Neither the main effect of sequence type nor the interaction between probe position and sequence type was significant (both $F_s < 1$). Since RTs to the probes did not differ between the positive-negative and negative-positive orders ($p_s > .2$), we collapsed over this factor in subsequent analyses. In the first half of the sequence, RTs were faster to the second probe ($M = 1.30, SD = 0.05$) than to the first probe ($M = 1.48, SD = 0.06$) [$t(23) = 4.49, p < .001, d = 0.92$]. RTs to the first probe in the second half of the sequence ($M = 1.45, SD = 0.06$) were slower than those to the second tone in the first half of the sequence [$t(23) = -4.38, p < .001, d = 0.90$]. Finally, RTs to the second tone of the second half of the sequence ($M = 1.39, SD = 0.06$) were faster than those to the first tone in the second half of the sequence [$t(23) = 2.14, p < .05, d = 0.44$]. Again, the accuracies to the tones in the first half ($M = .91, SD = .16$) and the second half ($M = .91, SD = .14$) did not differ [$t(23) < .01$], nor did any of the four trial types ($t_s < 0.54, p_s < .6$), suggesting that the observed RT effects were not the result of speed-accuracy trade-offs.

Discussion

In Experiment 2, participants became faster to respond to the auditory probes as they were presented with homogeneous trials conveying either positive or negative covariation information. As participants first entered the second half of the sequence, in which the trials conveyed new, contradictory covariation information, secondary-task responses slowed. Then, as participants continued to encounter this new, contradictory information, responses became faster again.

The results of Experiment 2 provide further evidence that learners subject covariation information to differential processing as it is encountered. Furthermore, the pattern of results

over the trial sequence is consistent with the idea that participants are developing expectations about the underlying causal relationships. As these expectations became stronger, expectation-consistent information was processed more easily. When these expectations were violated, as when learners encountered the second half of the sequence, greater cognitive resources were dedicated to processing the covariation information.

Experiment 3

Our strategy in Experiments 1 and 2 had been to manipulate the statistical properties of the trial sequences in ways known to affect learning, and to investigate whether these manipulations also led to changes in online processing. Because of this strategy, however, the RT effects that we have demonstrated so far could be straightforwardly explained as a consequence of the frequencies of the various trial types (or other simplistic, stimulus-bound processes). Thus, one could argue that the relatively slow RTs associated with the minority trials were not due to any learning-related cognitive processes per se, but directly to the fact that these trials were objectively less frequent.

In Experiment 3, we attempted to further examine the RT effects that we have observed so far, specifically testing whether they reflected learning-related processes or the more basic regularities in the task. If the RTs that we observed do indeed reflect violations of learned expectations, then the more consistent a participant's belief is with the underlying contingency, the stronger the violations on the minority trials would be, which in turn should lead to slower RTs, regardless of the objective contingencies. In contrast, if a participant failed to learn the contingency objectively describing a particular sequence, the information presented on the minority trials should not be particularly unexpected, and hence responses to probes should not be particularly slow. That is, if the RT effects do reflect learning-related processes, we should be able to see that the pattern of RTs would be more strongly associated with participants' own causal strength judgments than with the underlying statistical properties of the learning sequence. On the other hand, if the RT effects simply reflect the objective frequency of the trial types, then variability in RT effects should not be associated with variability in subjective causal beliefs regarding the underlying contingencies.

In order to test these alternatives, we employed a modified version of Experiment 1. Modifications were made because the causal strength judgments observed in Experiment 1 suffered from a lack of variability (e.g., in the positive condition of Exp. 1, 50 % of judgments were between 83 and 100) that would minimize our ability to detect a relationship between these judgments and RTs. To remedy this, in Experiment 3 we

required learners to simultaneously learn about three different causes (i.e., within the same trial sequence). Prior research has shown that simultaneously learning about more than one cause tends to increase cognitive load (Marsh & Ahn, 2006), which suggests that this modification should slow the rate at which participants learn the contingencies, and thus reduce the likelihood that causal judgments would be near ceiling.

Method

Participants A group of 44 undergraduate students at Stony Brook University participated in the experiment for partial course credits.

Materials and design The task and stimuli were identical to those used in Experiment 1. The information presented was also similar to that in Experiment 1, except for the following changes. We included three medications in each sequence, with one embodying a positive contingency, one embodying a negative contingency, and a filler cause embodying a contingency of zero. For the positive and negative medications, the proportions of trials from each cell of the covariation matrix were identical to those in the positive ($\Delta P = .6$) and negative ($\Delta P = -.6$) conditions in Experiment 1. For the filler medication, there were seven trials from each cell of the covariation matrix. The trials involving the different medications were randomly intermixed in the presentation sequence. Note that, as in Experiments 1 and 2, a picture of the medication crossed out was shown to indicate its absence. For the positive and negative medications, auditory probes were delivered on two trials of each type (e.g., two cell A trials, two cell B trials, etc.). For the filler medications, auditory probes were delivered on one trial of each type. The probe timing was identical to that in Experiment 1. Participants completed three sequences of trials, learning about three different medications in each sequence (i.e., each sequence included one positive, one negative, and one filler medication). Causal strength judgments for each medication were made at the end of each sequence.

Results

Overall, the positive causes were judged to be generative ($M = 40.04$, $SD = 30.13$), with judgments being significantly greater than zero [$t(43) = 8.81$, $p < .001$, $d = 1.33$]. The negative causes were judged to be preventative ($M = -38.99$, $SD = 37.77$), with judgments being significantly less than zero [$t(43) = 6.85$, $p < .001$, $d = 1.03$]. Judgments in the two types of causes were also significantly different from each other [$t(43) = 9.39$, $p < .001$, $d = 1.42$]. Although these gross measures suggest that participants' judgments were largely consistent with the underlying contingencies, aspects of the results also suggest that our attempt to increase the complexity

of the learning task was successful. First, despite presenting identical objective contingency information, these causal judgments were substantially less extreme (i.e., farther away from the end points of the scale) than were those observed in Experiment 1. That is, subjectively speaking, the positive causes in Experiment 3 were viewed as being less positive than those in Experiment 1 (judgments of 40.04 vs. 72.05, respectively), and the negative causes in Experiment 3 were viewed as being less negative than those in Experiment 1 (judgments of -38.99 vs. -68.33 , respectively). Second, judgments of both the positive ($SD = 47.20$) and negative ($SD = 53.84$) causes were more variable in Experiment 3 than in Experiment 1 ($SDs = 26.74$ and 36.9 , respectively). If the patterns of RTs that we observed in Experiments 1 and 2 were the consequences of the objective trial frequencies, these cross-experimental differences in causal judgments should be irrelevant when considering the magnitude of the RT effects. However, if the RT effects instead reflect learning-related processes, then we should expect the RT effects to be significantly attenuated relative to those observed in Experiment 1.

Indeed, analyses revealed that RTs on the majority trials ($M = 1.31$, $SD = 0.27$) were not different ($p > .9$) from those on the minority trials ($M = 1.31$, $SD = 0.27$). The fact that RTs were not significantly modulated by the statistical properties of the sequence (i.e., the effect found in Exps. 1 and 2) is consistent with the idea that learning about multiple medications simultaneously increased the complexity of the learning task (Marsh & Ahn, 2006) and reduced the magnitude of the RT effects observed in Experiments 1 and 2. Moreover, accuracies to the tones on majority ($M = .91$, $SD = .09$) and minority ($M = .92$, $SD = .10$) trials were not different. Given that RTs were not strongly related to the objective trial frequencies (i.e., majority vs. minority), the question of interest, then, is whether they were instead related to learners' subjective beliefs and expectations.

In order to investigate the relationship between participants' own idiosyncratic causal beliefs and their RTs, we constructed 44 individual linear regression models, one for each participant. Each of these regression models allowed the individual participant's subjective causal beliefs an opportunity to predict that individual's RT differences between majority and minority trials.

We first computed the difference between the participant's average RT on the minority trials and the participant's average RT on the majority trials ($RT_{\min} - RT_{\max}$). This calculation was performed separately for each cause so that each participant had six difference scores, one for each cause (three positive, three negative). We then quantified participants' causal beliefs. For the positive causes, we simply took participants' raw causal judgments, because *larger* judgments (i.e., more positive) were expected to predict larger RT differences (i.e., $RT_{\min} >> RT_{\max}$). For the negative causes, the causal

judgments were multiplied by -1 , because *smaller* judgments (i.e., more negative) were also expected to predict larger RT differences (i.e., $RT_{\min} \gg RT_{\max}$). These two constructed variables were then entered into linear regression models for each participant, with the participant's RT effect as the criterion variable and the participant's causal beliefs being the predictor variable. Each regression yielded a beta coefficient representing the degree to which that participant's RTs were predicted by that participant's causal beliefs. These beta values were then pooled ($M = 52.52$, $SD = 160.24$) and found to be significantly greater than zero [$t(43) = 2.17$, $p < .05$, $d = 0.33$]. This suggests that, on average, an individual participant's own causal beliefs predicted that participant's own RTs, whereas the objective contingencies did not.

Discussion

Experiment 3 was designed to investigate whether the RT effects observed in Experiments 1 and 2 reflected the objective statistics or learners' subjective causal beliefs of the underlying contingencies. To induce greater variability into participants' causal judgments, we simultaneously presented covariation information about three causes (rather than a single cause, as was done in Exp. 1). The results suggest that this modification had its intended consequences. The relationship between the objective contingencies and participants' causal judgments appears to have been weakened; causal judgments were both less extreme (i.e., closer to zero) and more variable than those observed in Experiment 1. The relationship between the objective contingencies and participants' RTs appears to have been similarly weakened; RTs on the majority and minority trials were not significantly different from each other. Taken together, these results suggest that having participants simultaneously learn about three causes weakened the relationship between the objective contingencies and the participants' behavior.

The weakened relationship between the objective contingencies and participants' behavior allowed us an opportunity to investigate whether the RT effects observed in Experiments 1 and 2 were driven by the participants' own idiosyncratic causal beliefs. To evaluate this possibility, we evaluated whether individual participants' own causal judgments predicted their own individual RT effects. The results indicate that these factors were significantly related. So, if a participant judged a particular cause to be "strong" (either strongly generative or strongly preventative), this participant also tended to respond more slowly on minority than on majority trials. In contrast, if this participant judged another cause to be weak (despite this cause being described by the same objective contingency as the "strong" cause), this participant would tend to respond just as quickly on minority as on majority trials.

Though these results are obviously correlational, they are consistent with the idea that the patterns of RTs observed in

Experiments 1 and 2 do not simply reflect the event frequencies embodied by the trial sequence (i.e., the contingency information). Instead, the present results are consistent with the idea that these RT effects reflect cognitively demanding processes that are related to learning, such as the violation of learners' subjective expectations. Just as in Experiment 1, when these expectations were strong (i.e., when causal strength judgments were far from zero), expectation-inconsistent information appeared to demand greater processing than expectation-consistent information. However, when expectations were weaker, despite identical covariation, objectively consistent and objectively inconsistent information evoked similar cognitive processing.

General discussion

Three experiments utilized an auditory discrimination task to measure cognitive processing during causal learning. The results from each of these experiments demonstrated that covariation information is not subject to uniform cognitive processing. Instead, we observed systematic variation in the processing dedicated to individual pieces of covariation information. In Experiments 1 and 2, responses to the tone discrimination task were slower when learners encountered covariation information that was inconsistent with earlier portions of the sequence. Experiment 3 showed that variability in discrimination responses was related to learners' subjective beliefs about the causes. When the results are taken together, the present study suggests that individual pieces of covariation information that are inconsistent with learners' causal beliefs elicit greater cognitive processing than do individual pieces of covariation information that are consistent with learners' causal beliefs.

Online processing

As we mentioned earlier, a critical component of extant process models is a comparative process describing how beliefs are revised in light of new evidence. Indeed, many process models suggest that learning relies critically on the explicit computation of prediction error (e.g., Catena, Maldonado, & Candido, 1998; Danks et al., 2003; Hogarth & Einhorn, 1992; Luhmann & Ahn, 2007; Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972). In the present experiments, prediction error should have been relatively small on majority trials in Experiments 1 and 3, and on trials near the end of the first half of the sequence in Experiment 2; all of these trials were associated with relatively fast response times. In contrast, prediction error should have been relatively large on the minority trials in the learning task in Experiments 1 and 3, and on the trials early in the second half of the sequence in

Experiment 2; all of these trials were associated with relatively slow RTs. This suggests that the processing dedicated to an individual piece of covariation information is modulated by the relationship between that piece of information and the learner's current causal beliefs. Thus, the systematic pattern of secondary-task RTs in the present experiments is consistent with the notion that individual pieces of covariation information engage one or more learning-related cognitive processes, the nature of which scales with prediction errors. That being said, we expect that many learning-related processes meet the criteria.

One possibility is that the patterns in participants' secondary-task RTs reflect the calculation of prediction error itself. As we described above, this calculation is a critical part of many process models. That being said, traditional associative models emphasize the automatic aspects of learning processes that require little or no cognitive resources (Evans, 2003; Evans & Over, 1996; Kahneman, 2003; Sloman, 1996; Stanovich & West, 2000). Such accounts would appear to have difficulty explaining the present findings. Of course, a variety of associative models also permit more deliberative processing (e.g., Wagner, 1981; Wagner & Brandon, 1989), as do several nonassociative models that rely on prediction error (Catena et al., 1998; Danks et al., 2003; Hogarth & Einhorn, 1992; Lu et al., 2008; Luhmann & Ahn, 2007). Any of these models could presumably be consistent with the idea of cognitively demanding computation of prediction error.

Alternatively, prediction errors could be computed quite automatically, but drive a cascade of cognitively demanding consequences aimed at altering both internal representations (e.g., beliefs about causal relationships) and future behavior (e.g., causal interventions; Lagnado & Sloman, 2004). For example, we have recently suggested (Liu & Luhmann, 2013) that prediction errors may drive shifts in attentional weights (Kruschke, 2001; Mackintosh, 1975), and that such attentional shifts may be cognitively demanding. This proposal represents a hybrid account, with automatic processes driving learning per se, but being followed by auxiliary, learning-related processes that might be cognitively demanding.

Luhmann and Ahn (2011) suggested that even more deliberative processes could be invoked when expectancies are violated. Specifically, when learners encounter information that contradicts their hypotheses, they may try to revise the existing hypothesis, or they may seek to dismiss the contradictory information. Luhmann and Ahn (2011) suggested that the latter strategy is both favored by learners and more cognitively demanding. In the present study, the cognitively demanding processing invoked by unexpected outcomes may reflect learner's efforts to engage in either or both of these two strategies.

Converging measures of learning

With the demonstration that secondary-task RTs reflect the expectedness of individual pieces of information during the course of learning, the present results also point to the potential benefits of directly scrutinizing the learning process itself. Here, we briefly highlight prior studies that have employed divergent methods, simply to begin illustrating how the potential diversity of methods that we have at our disposal can be leveraged to better understand processes like learning.

Initially, researchers believed that the standard input–output paradigm could be employed to assess the existence and nature of trial-by-trial learning processes. For example, substantial work has focused on the influence of presentation order (e.g., Dennis & Ahn, 2001; Katagiri et al., 2007; Lopez et al., 1998; Luhmann & Ahn, 2011; Yates & Curley, 1986), because it was believed that certain models made unique predictions that were tied to specific sequences of covariation information. For example, the Rescorla–Wagner model exhibits recency effects, in which recent covariation information exerts a disproportionate influence on associative strengths. However, research showed that order effects could be driven by memory-related processes (Stout et al., 2005), by changes in learning rates (Danks & Schwartz, 2005), or by the availability of working memory resources (Marsh & Ahn, 2006). Another approach asked learners to make trial-by-trial contingency judgments throughout the learning sequence (for a review, see Hogarth & Einhorn, 1992). However, analyses showed that the frequency with which judgments were elicited influenced the judgments themselves, suggesting that the technique was too invasive to provide clean measurements of learning.

One alternative approach for learning focuses on testing whether higher-order processes are critical to learning, by manipulating the availability of either time (González-Martín et al., 2012; Vadillo & Matute, 2010) or cognitive resources (De Houwer & Beckers, 2003; Waldmann & Walker, 2005). Though the introduction of such methodologies is refreshing, these specific findings yielded only ambiguous evidence for cognitively demanding processes in learning. For example, De Houwer and Beckers failed to find an influence of cognitive load on learning per se, only finding an effect when load was imposed during the process of making judgments themselves (i.e., after the trial-by-trial sequence had concluded). Waldmann and Walker only observed an influence of cognitive load when learners were engaged in diagnostic learning (determining whether some cues were caused by an event) involving cue interaction effects; no effect of load was found when learning in the predictive direction (determining whether some cues caused an event). On the other hand, González-Martín et al. and Vadillo and Matute found that certain learning effects (e.g., interference) were only observed when

participants were under time pressure, suggesting that automatic processing rather than analytical processing is crucial.

Le Pelley, Vadillo, and Luque (2013) used a secondary dot-probe task to measure the processes underlying learning of cue predictiveness. In their study, participants first learned that some cues were predictive of category memberships, and other cues were not. Later, cues of different predictivenesses were used to prime different spatial locations in which a dot could appear. Detection was faster when dots appeared in locations primed by predictive cues than when dots appeared in locations primed by nonpredictive cues, suggesting an attentional bias related to the learned predictiveness of the cues. Moreover, Le Pelley et al. administered the dot-probe task at different points of the learning task and found that such attentional bias increases as learning progressed.

Several studies utilizing eyetracking have explored cue competition effects (Beesley & Le Pelley, 2011; Kruschke et al., 2005; Le Pelley et al., 2011; Wills et al., 2007). These studies were particularly aimed at evaluating an attentional account of learning, which suggests that people learn not only causal strengths, but also attentional weights for individual causes. Although there are more traditional approaches to tackling this question (e.g., Kruschke & Blair, 2000), these studies utilized eyetracking in order to gain a direct measure of spatial attention as learners completed a traditional causal-learning task. Their results demonstrate that learners acquired significant attentional biases during the task. They further demonstrated a strong relationship between individual variability in attentional biases and learning, suggesting that the observed attentional effects were critically involved in cue competition. This finding, like the present results and those of Hastie and Kumar (1979), suggests that attentional and memory processes may, in large part, constitute the learning process itself.

More recently, studies have utilized electroencephalography to track neural processing during learning (Luque et al., 2012; Moris, Luque, & Rodríguez-Fornells, 2013; Walsh & Anderson, 2011). These studies focused on event-related potential (ERP) components that were either related to outcome evaluations (i.e., feedback-related negativity, FRN) or related to outcomes expectations (i.e., the stimulus-preceding negativity). Overall, these studies revealed that the magnitude of these learning-related ERP components change as learning progresses. For example, Walsh and Anderson found that expected and unexpected outcomes elicited FRNs of different magnitudes, and this difference grew as learners learned to better anticipate outcomes. This finding suggests that the FRN reflects the computation of prediction errors (Holroyd & Coles, 2002) and that expectations become stronger as the relationships between events are learned.

These studies go some way toward illustrating the diversity of methods that are applicable to the study of the processes

underlying causal learning. We expect that these are not the only, or even the best, methods available, but they should point the way for others to develop newer, more innovative approaches. Our belief is that a diversity of approaches is not only informative, but will also prevent the field from wasting time chasing after ever-more-nuanced predictions in an attempt to differentiate ever-more-similar theories (see Machery, 2009, for a depressing review of just such a situation).

References

- Beesley, T., & Le Pelley, M. E. (2011). The influence of blocking on overt attention and associability in human learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 37, 114–120.
- Bussemeyer, J. R. (1991). Intuitive statistical estimation. In N. H. Anderson (Ed.), *Contributions to information integration theory* (Vol. 1, pp. 189–215). Hillsdale, NJ: Erlbaum.
- Catena, A., Maldonado, A., & Cándido, A. (1998). The effect of the frequency of judgment and the type of trials on covariation learning. *Journal of Experimental Psychology: Human Perception and Performance*, 24, 481–495.
- Chapman, L., & Chapman, J. (1969). Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of Abnormal Psychology*, 74, 271–280.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367–405. doi:10.1037/0033-295X.104.2.367
- Danks, D., Griffiths, T. L., & Tenenbaum, J. B. (2003). Dynamical causal learning. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems 15* (pp. 67–74). Cambridge, MA: MIT Press.
- Danks, D., & Schwartz, S. (2005). Causal learning from biased sequences. In B. G. Bara, L. Barsalou, & M. Bucciarelli (Eds.), *Proceedings of the 27th Annual Meeting of the Cognitive Science Society* (pp. 542–547). Mahwah, NJ: Erlbaum.
- Danks, D., & Schwartz, S. (2006). Effects of causal strength on learning from biased sequences. In R. Sun & N. Miyake (Eds.), *Proceedings of the 28th Annual Meeting of the Cognitive Science Society* (pp. 1180–1185). Mahwah, NJ: Erlbaum.
- De Houwer, J., & Beckers, T. (2003). Secondary task difficulty modulates forward blocking in human contingency learning. *Quarterly Journal of Experimental Psychology*, 56B, 345–357. doi:10.1080/02724990244000296
- Dennis, M. J., & Ahn, W.-K. (2001). Primacy in causal strength judgments: The effect of initial evidence for generative versus inhibitory relationships. *Memory & Cognition*, 29, 152–164. doi:10.3758/BF03195749
- Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause. *Psychological Bulletin*, 99, 3–19. doi:10.1037/0033-2909.99.1.3
- Evans, J. S. B. T. (2003). In two minds: Dual process accounts of reasoning. *Trends in Cognitive Sciences*, 7, 454–459. doi:10.1016/j.tics.2003.08.012
- Evans, J. S. B. T., & Over, D. E. (1996). *Rationality and reasoning*. Hove, UK: Psychology Press.
- González-Martín, E., Cobos, P. L., Moris, J., & López, F. J. (2012). Interference between outcomes, spontaneous recovery, and context

- effects as measured by a cued response reaction time task: Evidence for associative retrieval models. *Journal of Experimental Psychology: Animal Behavior Processes*, 38, 419–432. doi:10.1037/a0029517
- Hastie, R., & Kumar, P. A. (1979). Person memory: Personality traits as organizing principles in memory for behaviors. *Journal of Personality and Social Psychology*, 37, 25–38.
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24, 1–55.
- Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, 109, 679–709. doi:10.1037/0033-295X.109.4.679
- Jenkins, H. M., & Ward, W. C. (1965). Judgement of contingency between responses and outcomes. *Psychological Monographs*, 79, 1–17.
- Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58, 697–720. doi:10.1037/0003-066X.58.9.697
- Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279–296). New York, NY: Appleton-Century-Crofts.
- Katagiri, M., Kao, S. F., Simon, A. M., Castro, L., & Wasserman, E. A. (2007). Judgments of causal efficacy under constant and changing interevent contingencies. *Behavioural Processes*, 74, 251–264. doi:10.1016/j.beproc.2006.09.001
- Kruschke, J. K., & Blair, N. J. (2000). Blocking and backward blocking involve learned inattention. *Psychonomic Bulletin and Review*, 7, 636–645.
- Kruschke, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology*, 45, 812–863.
- Kruschke, J. K., Kappenman, E. S., & Hetrick, W. P. (2005). Eye gaze and individual differences consistent with learned attention in associative blocking and highlighting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 830–845. doi:10.1037/0278-7393.31.5.830
- Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 856–876. doi:10.1037/0278-7393.30.4.856
- Le Pelley, M. E., Beesley, T., & Griffiths, O. (2011). Overt attention and predictiveness in human contingency learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 37, 220–229.
- Le Pelley, M. E., Vadillo, M., & Luque, D. (2013). Learned predictiveness influences rapid attentional capture: Evidence from the dot probe task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39, 1888–1900.
- Liu, P., & Luhmann, C. C. (2013). Evidence that a transient, but cognitively demanding process underlies forward blocking. *Quarterly Journal of Experimental Psychology*, 66, 744–766.
- Lopez, F. J., Shanks, D. R., Almaraz, J., & Fernandez, P. (1998). Effects of trial order on contingency judgments: A comparison of associative and probabilistic contrast accounts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 672–694.
- Lu, H., Rojas, R. R., Beckers, T., & Yuille, A. (2008). Sequential causal learning in humans and rats. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 185–190). Austin, TX: Cognitive Science Society.
- Luhmann, C. C., & Ahn, W.-K. (2007). BUCKLE: A model of unobserved cause learning. *Psychological Review*, 114, 657–677. doi:10.1037/0033-295X.114.3.657
- Luhmann, C. C., & Ahn, W.-K. (2011). Order effects during learning: Expectations and interpretations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37, 568–587.
- Luque, D., López, F. J., Marco-Pallares, J., Cámara, E., & Rodríguez-Fornells, A. (2012). Feedback-related brain potential activity complies with basic assumptions of associative learning theory. *Journal of Cognitive Neuroscience*, 24, 794–808. doi:10.1162/jocn_a.00145
- Machery, E. (2009). *Doing without concepts*. Oxford, UK: Oxford University Press.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276–298. doi:10.1037/h0076778
- Marsh, J. K., & Ahn, W.-K. (2006). Order effects in contingency learning: The role of task complexity. *Memory & Cognition*, 34, 568–576. doi:10.3758/BF03193580
- Moris, J., Luque, D., & Rodríguez-Fornells, A. (2013). Learning-induced modulations of the stimulus-preceding negativity. *Psychophysiology*, 50, 931–939. doi:10.1111/psyp.12073
- Pearce, J. M. (1994). Similarity and discrimination: A selective review and a connectionist model. *Psychological Review*, 101, 587–607. doi:10.1037/0033-295X.101.4.587
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian conditioning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87, 532–552. doi:10.1037/0033-295X.87.6.532
- Posner, M. I., & Boies, S. (1971). Components of attention. *Psychological Review*, 78, 391–408. doi:10.1037/h0031333
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York, NY: Appleton-Century-Crofts.
- Schustack, M. W., & Sternberg, R. J. (1981). Evaluation of evidence in causal inference. *Journal of Experimental Psychology: General*, 110, 101–120.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3–22. doi:10.1037/0033-2909.119.1.3
- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate. *Behavioral and Brain Sciences*, 23, 645–665. disc. 665–726.
- Stout, S. C., Amundson, J. C., & Miller, R. R. (2005). Trial order and retention interval in human predictive judgment. *Memory & Cognition*, 33, 1368–1376. doi:10.3758/BF03193369
- Vadillo, M. A., & Matute, H. (2010). Augmentation in contingency learning under time pressure. *British Journal of Psychology*, 101, 579–589.
- Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5–47). Hillsdale, NJ: Erlbaum.
- Wagner, A. R., & Brandon, S. E. (1989). Evolution of a structured connectionist model of Pavlovian conditioning (ÆSOP). In S. B. Klein & R. R. Mowrer (Eds.), *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theories* (pp. 149–189). Hillsdale, NJ: Erlbaum.
- Waldmann, M. R., & Walker, J. M. (2005). Competence and performance in causal learning. *Learning & Behavior*, 33, 211–229.
- Walsh, M. M., & Anderson, J. R. (2011). Modulation of the feedback-related negativity by instruction and experience. *Proceedings of the National Academy of Sciences*, 108, 19048–19053.
- White, P. A. (2002). Causal attribution from covariation information: The evidential evaluation model. *European Journal of Social Psychology*, 32, 667–684.
- Wills, A. J., Lavric, A., Croft, G. S., & Hodgson, T. L. (2007). Predictive learning, prediction errors, and attention: Evidence from event-related potentials and eye tracking. *Journal of Cognitive Neuroscience*, 19, 843–854. doi:10.1162/jocn.2007.19.5.843
- Yates, J., & Curley, S. (1986). Contingency judgment: Primacy effects and attention decrement. *Acta Psychologica*, 62, 293–302.