Journal of Experimental Psychology: General
2011, Vol. 140, No. 4, 660–673

0096-3445/11/$12.00 DOI: 10.1037/a0023862

Temporal Contexts: Filling the Gap Between Episodic Memory and Associative Learning

Helena Matute
Universidad de Deusto

Ottmar V. Lipp
University of Queensland

Miguel A. Vadillo
Universidad de Deusto

Michael S. Humphreys
University of Queensland

People can create temporal contexts, or episodes, and stimuli that belong to the same context can later be used to retrieve the memory of other events that occurred at the same time. This can occur in the absence of direct contingency and contiguity between the events, which poses a challenge to associative theories of learning and memory. Because this is a learning and memory problem, we propose an integrated approach. Theories of temporal contexts developed in the memory tradition provide interesting predictions that we test using the methods of associative learning to assess their generality and applicability to different settings and dependent variables. In 4 experiments, the integration of these 2 areas allows us to show that (a) participants spontaneously create temporal contexts in the absence of explicit instructions; (b) cues can be used to retrieve an old temporal context and the information associated with other cues that were trained in that context; and (c) the memory of a retrieved temporal context can be updated with information from the current situation that does not fit well with the retrieved memory, thereby helping participants to best adapt their behavior to the future changes of the environment.

Keywords: associative learning, episodic memory, temporal contexts, paired associate, predictive learning

When people remember something that they experienced at a particular time of their life, they may also remember other things that occurred in the same epoch. For instance, over the course of a conversation, people may retrieve memories of childhood friends they made during a family holiday, and the revival of some episodes may also lead them to retrieve the memories of a historical event that also took place during that wonderful holiday they spent at the coast with their parents.

The phenomenon we just described may look at first glance like a simple associative effect. If the presentation of a given stimulus can activate the mental representation of another stimulus, this must mean that the two stimuli share an associative link. The challenge comes when we think about the way in which theories of learning and memory predict that links between stimuli are formed. Ideally, one stimulus must be consistently and contiguously followed by the other (contiguity) and should not occur in the absence of the other (contingency). Indeed, the most widely cited theories of associative learning (e.g., Rescorla & Wagner, 1972) assume that as contiguity or contingency are reduced, the strength of the association will weaken.[1] Theories of memory make a similar prediction, with some of the most popular views of memory assuming a buffer in which contiguous events become associated (e.g., search of associative memory; Raaijmakers & Shiffrin, 1980, 1981). Thus, most theories of learning and memory predict that direct associations should not form between isolated stimuli that are neither contiguous nor contingent on each other.

It appears quite obvious, however, that events that were experienced in the same temporal context can activate each other's mental representation. The question is how they do it. This is the focus of the present study. Because this is a problem of learning

---

Helena Matute and Miguel A. Vadillo, Departamento de Fundamentos y Métodos de la Psicología, Universidad de Deusto, Bilbao, Spain; Ottmar V. Lipp and Michael S. Humphreys, School of Psychology, University of Queensland, Brisbane, Queensland, Australia.

Support for this research was provided by Grant SEJ2007-63691/PSIC from the Spanish Ministry of Science and Innovation; Grant IT363-10 from the Departamento de Educación, Universidades e Investigación of the Basque Government; and Grant DP0770844 from the Australian Research Council. We thank M. Teresa Bajo, Mark E. Bouton, Pedro L. Cobos, Samuel D. Hannah, Marc W. Howard, Francisco J. López, Thomas Suddendorf, and Jason M. Tangen for their very helpful comments and suggestions on this research.

Correspondence concerning this article should be addressed to Helena Matute, Departamento de Psicología, Universidad de Deusto, Apartado 1, 48080 Bilbao, Spain. E-mail: matute@deusto.es

[1] There are exceptions to the contiguity principle, such as taste aversion learning, which occurs when the consumption of a given food is followed by illness. The association between food and illness is a very special one because poisonous foods can produce illness after several hours. Thus, to be effective, this type of learning needs to bridge very long delays (Garcia, Ervin, & Koelling, 1966). Nevertheless, even for taste aversion learning, the general contiguity principle still holds that the longer the interval, the weaker the association.

and memory, we first describe how this problem can best be addressed from each of those perspectives. After that, we suggest an integrated framework.

Within the associative learning tradition, Bouton's (1993, 1997) theory is the one most clearly devoted to understanding how contexts, physical or temporal, affect learning and behavior. It is well-known that the presentation of a cue can retrieve the mental representation of a second stimulus (i.e., an outcome) that has been associated with it. However, there are times in which the cue has been paired with different outcomes in different contexts. According to Bouton (1997), in those cases, the cue becomes ambiguous and the context can be used to disambiguate it. The association acquired second becomes context specific, which means that it will only be retrieved in the context in which it was acquired. If organisms are tested in the context in which the association acquired first was trained or in a new context, the association acquired first will prevail. A revision of this theory advanced by Rosas, Callejas-Aguilera, Ramos-Álvarez, and Fernández-Abad (2006) makes similar predictions, although it states that both the first- and the second-learned associations could, in principle, be specific to the context in which they were trained (see also Gawronski, Rydell, Vervliet, & De Houwer, 2010; Rosas & Callejas-Aguilera, 2006). According to this version of the theory, the context will play a role in the retrieval of the outcome if the organism has paid attention to the context. Attention may be devoted to the context not only because the cue is ambiguous, as suggested by Bouton (1993), but also because the participant has been instructed to attend to the context or for other reasons. Bouton's model, as well as Rosas et al.'s revision, has shown great heuristic value in guiding new research and in providing successful predictions in situations in which temporal contexts are involved, situations that are normally outside of the scope of simpler theories of learning (e.g., Rescorla & Wagner, 1972).

Consider, for example, an experiment by Brooks and Bouton (1993). Rats were first exposed to several pairings of a cue, A, with an outcome (e.g., A–O1). Then this learning was extinguished. That is, Cue A was now paired with nothing (i.e., A–O2). This normally produces responding appropriate for O2 if testing takes place immediately after A–O2 training (i.e., a recency effect; in this case, weak responding to A). So far, this can be simulated with very simple associative models of learning such as that of Rescorla and Wagner (1972). However, the presentation of A may also produce recovery of the response appropriate to O1 if testing takes place in a new physical context or after some time (i.e., in a new temporal context). This latter phenomenon is called *spontaneous recovery* (Pavlov, 1927). Because the Rescorla and Wagner model conceptualizes extinction as catastrophic forgetting of the association acquired first, the model cannot account for spontaneous recovery. Bouton's (1997) theory, however, does so by assuming that because the extinction phase contradicts what was learned in Phase 1, extinction becomes context specific. Thus, extinction-related behavior will only be retrieved if the organism is tested in the extinction context. Otherwise, what was learned during Phase 1, before the cue became ambiguous, will prevail. This explains the recovery of responding when testing occurs in new temporal (or physical) contexts. The main point of the study by Brooks and Bouton was to show that a retrieval cue for extinction presented after some time was able to induce behavior appropriate for extinction instead of the default spontaneous recovery that should have been observed at that time. That is, when rats were prompted to retrieve the context of extinction, behavior appropriate for the extinction phase was shown.

We know of no other theory of learning that could predict these and many other experimental results in which the retrieval of temporal contexts is involved. It must be noted, however, that in Bouton's model, the temporal context is not defined by explicit stimuli but by the mere passage of time. That is, even though external cues that are presented in a particular phase of training can be used as retrieval cues, it is the passage of time, rather than the changing external stimuli, that defines the temporal contexts. Despite their enormous heuristic value, neither Bouton's (1993, 1997) model nor Rosas et al.'s (2006) revision specifies how temporal contexts are represented, how the representation of old temporal contexts can be reactivated in memory, or how the presentation of a cue can reactivate a temporal context and the memories of other stimuli that were trained in that context. To our knowledge, these questions have not yet been addressed in the associative learning tradition.

These questions have, however, been thoroughly addressed in the memory tradition and, thus, both areas can complement each other. The temporal context model (TCM) of Howard and Kahana (2002), for instance, provides interesting insights into the mechanisms by which temporal contexts can be created and represented. According to TCM, stimuli become associated during training with the current state of a gradually changing representation of the temporal context. This context, in turn, gets associated with the stimuli. Later on, a presentation of the stimulus will enable the participant to retrieve this temporal context, and, by the same reasoning, the reactivation of the temporal context will also enable the participant to retrieve the representation of other stimuli that were linked to it. It is important to note that the temporal context is not an abstract concept here but a flexible and evolving representation that depends on, among other things, the stimuli (both internal and external to the organism) that are present at each moment. This means that stimuli that are present at a given time do indeed become part of the temporal context for other stimuli present at that time. Howard and Kahana's (2002) model focuses on simulating the evolution of the temporal contexts and the retrieval processes. Moreover, it provides a comprehensive theoretical framework of how a given cue can retrieve an entire episode or temporal context. The most recent version of this model (Sederberg, Howard, & Kahana, 2008) maintains the assumption of gradually evolving contexts while incorporating a new retrieval rule that is based on the leaky-accumulator decision model of Usher and McClelland (2001). In this new version, it is the relative activation of each of the different associations between context and stimuli that determine which memories are most readily retrieved at each time. The manner in which contexts are represented and retrieved is clearly specified in this model. Thus, the TCM can be used to complement the theories developed in the associative learning tradition.

In support of the TCM, researchers have shown, for example, that participants create temporal contexts that later guide their retrieval of the episodes. For instance, when studying paired associates, participants not only associate the required pairs (e.g., AB) but also form associations between these pairs and other pairs presented nearby in the list (e.g., Davis, Geller, Rizzuto, & Kahana, 2008). This also suggests that associations are formed not only among items but also between items and contexts and vice versa (see also Ranganath, 2010; Schwartz, Howard, Jing, & Kahana, 2005).

Convergent evidence for this proposal can be found in a study by Howard, Jing, Rao, Provyn, and Datey (2009), in which participants were presented with lists of paired associates (e.g., AB DE, BC EF) in a random order; the associations between items that were not presented closely together in time but were presented in similar contexts (i.e., with context in this case being the other member of the study pair such that A and C occur in the same context, namely, B) were subsequently examined. As should be expected from the TCM, participants were able to associate distant events that shared a common context. Similar effects were also evident when the stimuli were presented in separate lists or separated by several hundred seconds, that is, in the absence of contiguity (Howard, Youker, & Venkatadass, 2008). Associative-like effects similar to these have been demonstrated even for stimuli separated by distracters (Howard & Kahana, 1999). These findings are consistent with the idea that the presentation of a cue can retrieve the mental representation of the temporal context, which, in turn, can retrieve the representation of other stimuli that were also trained during the same temporal context.

It is important to note that if we are claiming that theories about temporal contexts developed in the memory tradition can be used to complement theories developed in the associative learning tradition, it should be shown that predictions derived from the memory literature can be verified in paradigms used in associative learning. Indeed, although the theories in the memory tradition have reached a greater level of development, they have been tested almost exclusively on paired associates and word lists. The associative learning methodology can provide convergent evidence using nonverbal behavior as the dependent variable (e.g., behavioral adaptation in simple videogames; see Arcediano, Ortega, & Matute, 1996; Costa & Boakes, 2011; Franssen, Clarysse, Beckers, van Vooren, & Baeyens, 2010; Lipp & Dal Santo, 2002). This, in turn, should speak to the generality of the effects, and new predictions should arise from the integration of both research areas. Using associative learning methods will also allow us to extend the findings from the memory tradition to new situations, such as those that lack a time interval within phases (something that is difficult to implement in the memory paradigms), or to situations that lack explicit instructions about the temporal contexts. Thus, the associative learning methods permit the investigation of how participants construct temporal contexts in the absence of external cues, such as changing learning lists or instructions.

Experiment 1

A recent experiment in the memory tradition by Humphreys, Murray, and Maguire (2009) is our starting point. Simplifying tremendously, participants first learned two different word lists, as in the AB, CD, interference paradigm. At test, participants were more accurate when a C cue followed another C cue than when it followed an A cue. This suggests that cues can retrieve temporal contexts, which, in turn, can retrieve other stimuli that occurred in that context. Humphreys et al. used both instructions and temporal grouping to establish the lists, or temporal contexts. They also used both instructions and a cue to reinstate the contexts at test. This result, we believe, resembles the findings by Brooks and Bouton (1993) described above, although, of course, Brooks and Bouton did not use instructions with their nonhuman participants. Nevertheless, Humphreys et al.'s study is more comprehensive: Information from both Phase 1 and Phase 2 was retrieved using retrieval cues, whereas Brooks and Bouton had argued that only the second phase learning (extinction) was subject to selective retrieval. Moreover, the possibility exists that Brooks and Bouton's results may be exclusive to a conditioning paradigm with nonhuman animals, which involves the use of biologically significant events such as unconditioned stimuli (see Gunther, Miller, & Matute, 1997). Our intent is therefore to follow up on Humphreys et al.'s study using the methods of human associative learning. Thus, we used symbolic rather than biologically significant events, and the dependent variable was predictive behavior. Predictive behavior and predictive judgments are often used in the human associative learning tradition as indices by which to determine that an association has been learned, and they are also regarded as analogues of conditioned responding in animals (e.g., Arcediano et al., 1996; Costa & Boakes, 2011; Franssen et al., 2010; Lipp & Dal Santo, 2002; Shanks & Dickinson, 1987; Vadillo, Bárcena, & Matute, 2006; Wasserman, 1990). According to this view, once an association has been acquired between a cue and an outcome, participants should be able to predict the outcome when the cue is presented, which means that anticipatory behavior (and judgments) appropriate to that particular outcome should occur in response to the cue.

In the current study, Target Cue X predicted outcomes that required different responses in different phases. X was associated with a negative outcome (O1) during Phase 1 and with a positive outcome (O2) during Phase 2, so that participants were trained not to respond to Cue X during Phase 1 and to respond to Cue X during Phase 2. Each phase of the study, including the test phase, was conducted immediately after termination of the previous phase and without interruption. Therefore, strong responding to X should be the default behavior observed at test, given that the temporal context of the test trial was identical to that of Phase 2 and very different from that of Phase 1. However, if we prime the temporal context of Phase 1 just before testing, then participants should retrieve the association trained during Phase 1 instead and should thus show weak responding to X at test, as if they were back in Phase 1. It should be noted that the current predictions go beyond those tested by Humphreys et al. (2009). First, we are testing whether a temporal context will be naturally formed with the different associations learned in different phases of the experiment. Second, we are testing whether a cue on its own can reinstate the temporal context of Phase 1.

Method

Participants and apparatus. Forty-one students from the University of Queensland, Australia, received course credit for their participation in the experiment. A computer program randomly assigned participants to one of two groups. This resulted in 20 participants in Group Retrieve 1 and 21 in Group Default. All of the materials were presented in an HTML document that included JavaScript functions to manage the presentation of the stimuli on the computer screen and to collect participants' responses.

Design and procedure. For this experiment, we used the *spy-radio task* (Matute, Vadillo, & Bárcena, 2007; Pineño, Ortega, & Matute, 2000).[2] In this task, participants are asked to imagine that they are soldiers ordered to rescue refugees that are hidden in a ramshackle building. On each trial, participants are given the opportunity to place a number of refugees on a truck and to take them to safety. Participants can place people on the truck by pressing the space bar repeatedly: The more they press the space bar, the more people they place. However, the refugees placed on the truck do not always arrive safely at their destination. In some trials, the road the truck has to drive on landmines that can explode. Participants can predict whether the road will be safe on a given trial by paying attention to the colored lights in a spy radio installed in the truck. In each trial, a cue (i.e., a colored light) is presented for 3 s during which the participant can respond. Certain colors in the spy radio predict that the road will be safe (and, therefore, indicate that participants should place as many refugees as possible on the truck while that cue is on), whereas other colors predict that the road will be mined (and, therefore, indicate that participants should avoid placing refugees on the truck during the presentation of those cues).

Immediately after the cue turned off, the outcome was presented: A message on the screen indicated whether the refugees had arrived safely at their destination; it also indicated the number of points earned or lost in that trial. Participants lost 1 point for each refugee placed on the truck on trials in which the road was mined (negative outcome; O1) and earned 1 point for each refugee placed on the truck on trials in which the road was safe (positive outcome; O2). They were not told which color predicted which outcome. Their main task was to learn this by paying attention to what happened during the learning trials. In each trial, the number of refugees placed on the truck during the 3-s interval for which the light was on was our dependent variable. This variable reflects the extent to which participants have learned that the cue (i.e., the color of the light) presented on a particular trial predicts that the road will be safe.

The design of the experiment is shown in Table 1. During Phase 1, two groups of participants received 10 trials with the target cue, X, which was always followed by O1, and 10 trials with Cue A, which was always followed by O2. These 20 trials were randomly intermixed. Then, during Phase 2, both groups received 10 additional trials with X, which was now followed by O2, and 10 trials with a new cue, B, which was followed by O1. The 20 trials of this phase were also randomly ordered. The colors that served as Cues A, B, and X were blue, yellow, and red, counterbalanced across participants. The duration of the intertrial interval was random, ranging from 3 to 7 s. During this time, all lights were off (i.e., grey) and pressing the space bar produced no effect.

Table 1
Design Summary of Experiment 1

Group	Training Phase 1	 Phase 2	 Phase 3	Test
Retrieve 1	X–O1/A–O2	X–O2/B–O1	A–O2	X
Default	X–O1/A–O2	X–O2/B–O1	B–O1	X

Note. Cues A, B, and X are color lights in the spy radio that can predict two different outcomes: O1 (the participants lose points if they perform the response) and O2 (participants earn points if they perform the response). In each phase, the different trial types were randomly intermixed. Phase 3 consists of just one trial. In the actual experiment, the different phases are not separated (i.e., a regular intertrial interval separates the last trial of one phase and the first trial of the next one).

There were no breaks between phases. Thus, if participants were to separate the continuous sequence of trials into different temporal contexts, they needed to spontaneously create temporal contexts themselves that coincided with those designed by the experimenters (see Table 1). At the end of the experiment, both groups received one test trial with the target cue, X. Before that, however, the critical manipulation occurred during Phase 3, which consisted of just one trial. This trial was presented immediately after the last trial of Phase 2 and immediately before the test trial (with a regular intertrial interval separating these trials). In Group Default, the Phase 3 trial consisted of a regular B–O1 trial (i.e., an additional Phase 2 trial). In Group Retrieve 1, however, this trial was an A–O2 trial (i.e., a Phase 1 trial). This manipulation should cue participants in Group Retrieve 1 to recall the context of Phase 1 just before testing. This, in turn, should prime them to respond to X at test with the response that was appropriate during Phase 1 (i.e., low responding to X) rather than with the response that was most recently acquired. That is, even though strong responding to X should be displayed by the end of Phase 2 (and this should be evident in Group Default), retrieving the temporal context of Phase 1 should lead participants in Group Retrieve 1 to respond weakly to X at test.

At this point, it is important to note that there was no clear contiguity or contingency between A and X in this experiment. Cues X and A were each contiguously followed by their respective outcomes, but the different X and A trials were separated from each other by the intertrial intervals and were randomly ordered within each phase. Therefore, the only relationship that existed between Cue A and the meaning that Cue X had during Phase 1 is that they were trained in the same phase. Indeed, participants were not even told that there were different phases in the experiment. However, we predicted that participants would be able to create the temporal context that corresponded to the phases planned by the experimenters, as the different cue–outcome associations presented in each phase remained constant for the duration of each phase and changed with each new phase.

Results and Discussion

The data selection criterion usually applied when using this and similar tasks in human associative learning is that at the end of training, there must be more responses to the positive cues than to the negative cues. This is a very lenient criterion that discards the data of participants who show no signs of learning during the training session. Such lack of learning can occur for a variety of reasons, such as color blindness or lack of attention. If the target associations were not acquired, it makes little sense to test for their retrieval. In the present experiment, this criterion requires that the number of responses given during the last positive trial of Phase 1 and Phase 2 be higher than the number of responses during the last

---

[2] A demo version of this program can be downloaded from http://www.labpsico.deusto.es/en/resources/

negative trials of Phases 1 and 2 (before the critical trial is introduced in Phase 3). Following this criterion, we eliminated two participants from Group Retrieve 1 and one participant from Group Default. This resulted in 18 participants in Group Retrieve 1 and 20 in Group Default.

The results are shown in Figure 1. They confirmed our predictions and replicated the basic findings of Humphreys et al. (2009) using a very different methodology and dependent variable. The training phases proceeded as expected, with the learning curves showing a gradual increase in the number of responses to the stimuli that were associated with gaining points and a decrease of responses to stimuli that were associated with losing points. This is hardly surprising, not only because the task was relatively easy but also because, as mentioned above, the data from those participants who did not learn to respond more to the positive than to the negative stimuli were discarded.

The critical results are those of the test phase. As expected, they show that participants in Group Retrieve 1 responded significantly less to X at test than did participants in Group Default, $t(36) = 5.07$, $p < .001$. As previously mentioned, the test trial is effectively an additional Phase 2 trial in Group Default and, thus, strong responding is observed in that group. The interesting result is that of Group Retrieve 1, in which participants, as expected, behaved differently. Presenting participants just before testing with the stimulus that had been trained along with X during Phase 1 made them respond to X as if they were back in Phase 1. These results suggest that memories of independent associations that were trained in the same temporal context can be used to retrieve each other even when there is no contiguity or contingency between them.

Experiment 2

Experiment 1 suggests that the presentation of a stimulus can be used to retrieve other stimuli or experiences that were trained at the same time even when there is no contiguity or contingency between the two stimuli. In the first experiment, it was important to use a condition in which strong responding to X was evident by the end of Phase 2 so that any changes in behavior shown at test could not be attributed to preasymptotic learning. In Experiment 2, however, to make certain that the results were not due to this particular aspect of the design, participants were exposed to the same treatment but with the order of phases reversed. That is, by the end of Phase 2, groups should now show little or no responding to the target cue, X (i.e., as in extinction). Like in Experiment 1, the test phase was conducted at the end of training and without interruption, which, in this case, was planned to minimize the likelihood of spontaneous recovery. Thus, testing should show extinction in Group Default. Group Retrieve 1, by contrast, was, just before testing, cued to retrieve the strong responding to X that in this experiment occurred during Phase 1. If our hypothesis is correct, Group Retrieve 1 should now show higher rather than lower responding than Group Default.

Method

Participants and apparatus. A total of 51 anonymous Internet users who visited our virtual laboratory (http://www.labpsico.com) volunteered for the study. The computer program randomly assigned participants to one of two groups. This resulted in 26 participants in Group Retrieve 1 and 25 in Group Default. To comply with ethical regulations in the conduct of Internet research, we did not request any personal data, nor did we use cookies or other software to obtain personal information without the participants' consent. All stimuli involved in the experiment were preloaded in the computer's memory before participants could start the experiment, so differences in the connection speed did not influence the pace of the experiment.

Although we did not record the computers' Internet addresses (i.e., IPs) in the present series of experiments, we have done so in the past and have verified that the rate of data sets coming from the same computer is negligible (around 2%), which confirms that multiple data submission is unlikely to pose a problem for the

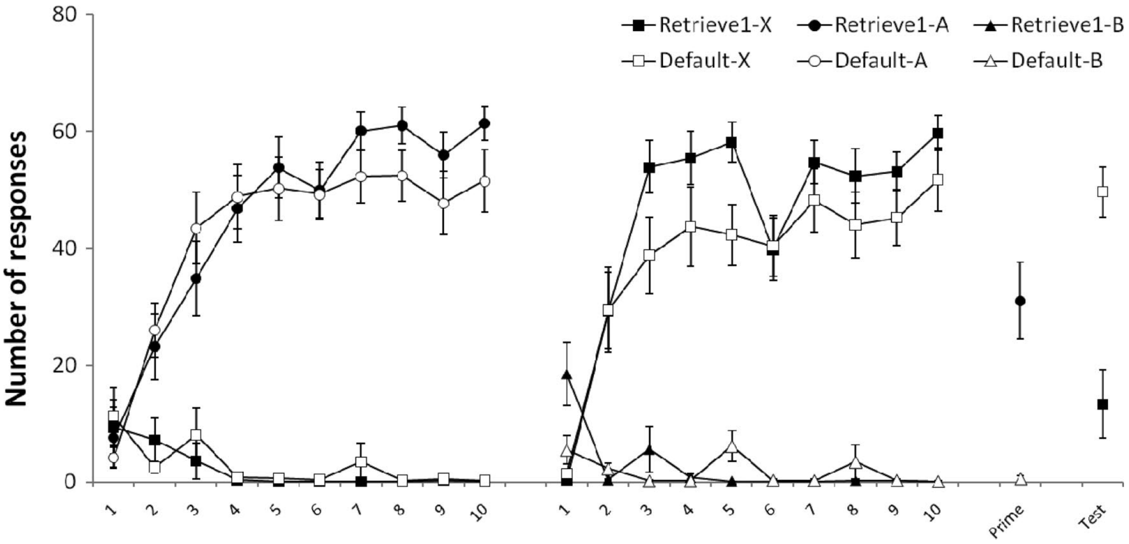

Figure 1. Mean number of responses during the different phases of Experiment 1 as a function of group and cue. Error bars represent the standard errors of the mean.

validity of these studies (see also Reips, 2002). Moreover, even those cases can represent different people using the same computer (e.g., family members often share the same computer, and different students in the same university can use the same computer at different moments). Thus, duplicate IPs do not necessarily prove that the same person is participating several times in the same experiment. Perhaps most important, previous experiments have shown that very similar results are normally obtained in the laboratory and through the Internet (for a review, see Kraut et al., 2004). This convergence between online and offline results has also been confirmed for experiments in associative learning that use procedures very similar to the ones we are using here (e.g., Matute et al., 2007; Vadillo et al., 2006; Vadillo & Matute, 2011). In any case and as a general rule of prudence, two of the experiments in this series were conducted through the Internet but the other two were conducted in the laboratory.

Design and procedure. As can be seen in Table 2, the only difference between this experiment and Experiment 1 is that the order of the first two phases was reversed. The critical manipulation in this study, which again took place during Phase 3, consisted of presenting an additional Phase 2 trial in Group Default and a Phase 1 trial in Group Retrieve 1. Therefore, Group Default should show weak responding to X at test. A comparatively stronger response should be shown in Group Retrieve 1 if, as suggested by Experiment 1, participants in this group behave at test as if they were back to the temporal context of Phase 1.

Results and Discussion

Applying the same data selection criterion used in Experiment 1, we eliminated three participants from Group Retrieve 1 and one from Group Default from the analyses. The results are shown in Figure 2. The training phase proceeded smoothly, with participants learning to respond more to the positive than to the negative stimuli. The critical results are those of the test phase. As expected, participants in Group Retrieve 1 responded to Cue X at test significantly more than did participants in Group Default, $t(45) = 3.62$, $p < .001$. The weak responding of Group Default at test is not surprising, given that responding to X had been extinguished by the end of Phase 2 and that the test trial was conducted in an identical temporal context. However, in line with the results observed in Experiment 1, participants in Group Retrieve 1 behaved differently. One presentation of the stimulus that had been trained along with X during Phase 1 just before test led these participants to behave at test as if they were back in Phase 1.

Table 2
Design Summary of Experiment 2

Group	Training			Test
	Phase 1	Phase 2	Phase 3	
Retrieve 1	X–O2/A–O1	X–O1/B–O2	A–O1	X
Default	X–O2/A–O1	X–O1/B–O2	B–O2	X

Note. Cues A, B, and X are color lights in the spy radio that can predict two different outcomes: O1 (the participants lose points if they perform the response) and O2 (participants earn points if they perform the response). In each phase, the different trial types were randomly intermixed. Phase 3 consists of just one trial. In the actual experiment, the different phases are not separated (i.e., a regular intertrial interval separates the last trial of one phase and the first trial of the next one).

Experiment 3

Experiments 1 and 2 used the strategy of showing in Group Default the behavior that should normally be expected at the end of Phase 2 and comparing this with behavior prompted by the presentation of a cue that reactivated the temporal context of Phase 1. However, it is possible that the presentation of a retrieval cue from Phase 1 at the end of Phase 2 had the additional (or alternative) effect of producing a disruption between Phase 2 and test. This disruption between training and testing was not present in Group Default, and it could, in principle, be responsible for the differential responding observed at test. Thus, in Experiment 3, we used two retrieve groups and omitted Group Default. The two groups used in this experiment were equally distracted at the end of Phase 2. Then, after both groups had been distracted, we cued either Phase 1 or Phase 2 for Group Retrieve 1 and Group Retrieve 2, respectively. A test trial with the target cue, X, was then presented to assess whether responding appropriate to Phase 1 or Phase 2 took place. If the disruption between Phase 2 and test was responsible for the observed effects, those effects should not be reproduced in the present experiment.

Therefore, in this experiment, the test phase took place well after Phase 1 and Phase 2 had finished. This normally produces spontaneous recovery of the association that was trained during Phase 1 (e.g., Bouton, 1997; Pavlov, 1927). Even though it could be argued that according to some theories (e.g., Bouton, 1993), excitatory associations should prevail at test, regardless of the phase in which they were trained, spontaneous recovery should nevertheless affect both groups to the same extent. That is, in our experiment, participants should either recover the excitatory association in both groups or the association trained first in both groups (i.e., retrieving X–O2 more strongly than X–O1 or vice versa). Our prediction, however, holds that the groups will behave differentially as a function of whether the contexts of Phase 1 or Phase 2 are primed for retrieval.

Method

Participants and apparatus. This experiment was conducted through the Internet with 105 anonymous volunteers. The computer program randomly allocated participants to Group Retrieve 1 (49 participants) and Group Retrieve 2 (56 participants).

Design and procedure. The design summary is shown in Table 3. The order of phases used in Experiment 1 was replicated because that experiment had been conducted only in the laboratory. Thus, an additional purpose of Experiment 3 was to make sure those results would replicate in the noisier domain of the Internet. The design was almost identical to that of Experiment 1 except that a new phase was added after Phase 2 to create a disruption between Phase 2 and test in both groups. In this new phase, 10 trials were provided in which a novel cue, C, was paired with O2. Cue C was a white light for all participants. Given that some time had elapsed since the end of Phase 2, there is no reason now to expect that the participants' behavior should continue reflecting the contingencies trained during Phase 2. Thus, we now tried to induce the retrieval of the temporal context of either Phase

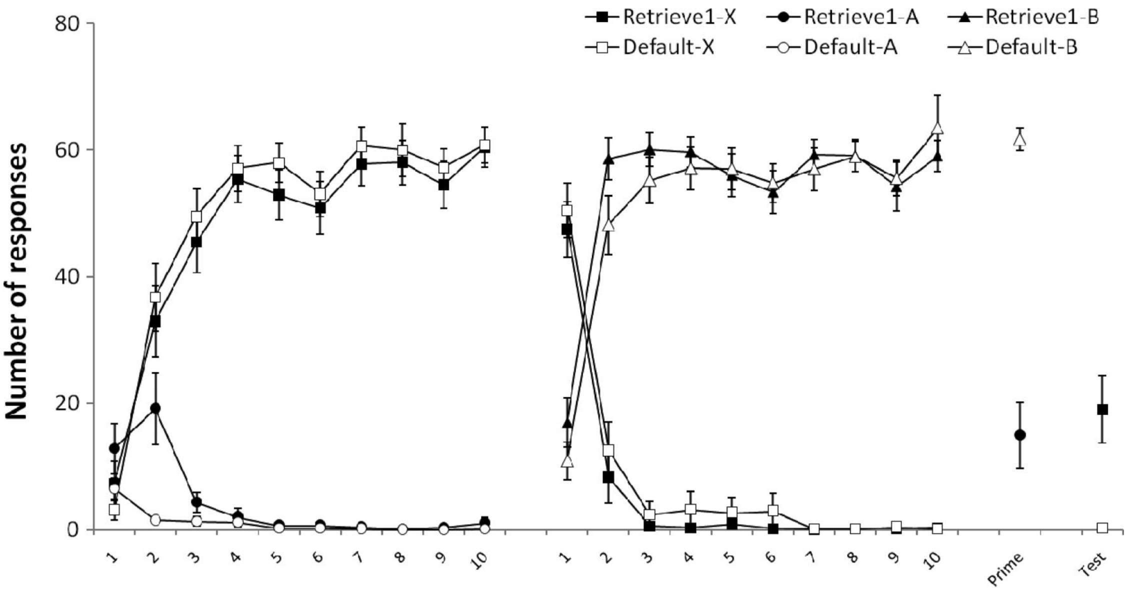

Figure 2. Mean number of responses during the different phases of Experiment 2 as a function of group and cue. Error bars represent the standard errors of the mean.

1 or Phase 2 in Group Retrieve 1 and Group Retrieve 2, respectively. To do so, we presented, just before testing, one trial with the stimulus (A or B) that had been trained along with X in Phase 1 or in Phase 2. Participants should now respond to the target cue X during the test trial as if they were back in Phase 1 or Phase 2, depending on whether they are being cued with A or B. All other procedural details were the same as in the previous experiments.

Results and Discussion

Two participants (one from each group) failed to meet the data selection criterion used in the previous experiments and were not included in subsequent analyses. The results are shown in Figure 3. As in the previous experiments, training proceeded smoothly, with the participants learning to respond more to the positive than to the negative stimuli by the end of training. Most important, the results show that participants in each group responded to X during test, as expected. Group Retrieve 1 responded significantly less than did Group Retrieve 2, $t(101) = 2.55$, $p < .05$. Thus, the cue that was presented before testing effectively activated the temporal context of one or the other phase.

Table 3
Design Summary of Experiment 3

	Training				
Group	Phase 1	Phase 2	Phase 3	Phase 4	Test
Retrieve 1	X–O1/A–O2	X–O2/B–O1	C–O2	A–O2	X
Retrieve 2	X–O1/A–O2	X–O2/B–O1	C–O2	B–O1	X

Note. Cues A, B, C, and X are color lights in the spy radio that can predict two different outcomes: O1 (the participants lose points if they perform the response) and O2 (participants earn points if they perform the response). In each phase, the different trial types were randomly intermixed. Phase 4 consists of just one trial. In the actual experiment, the different phases are not separated (i.e., a regular intertrial interval separates the last trial of one phase and the first trial of the next one).

Experiment 4

The results of Experiments 1–3 are compatible with the idea that people spontaneously construct temporal contexts using the stimuli that are present at a given time. They also show that people can use stimuli that are present in the current situation to cue older temporal contexts and the stimuli that had been associated with those contexts. Moreover, these experiments show that this cuing makes participants behave as if they were back in the old temporal context. But taking the idea of gradually evolving temporal contexts one step further, we can also make some new predictions.

Many memory researchers agree that the purpose of retrieving an old temporal context or episode and the memories associated with it is to guide adaptive behavior at present and to plan for future situations that might benefit from this knowledge (Schacter, Addis, & Buckne, 2007; Suddendorf & Corballis, 1997, 2007; Tulving, 2005). For this reason, when old episodes are retrieved, they are updated by integrating the current information into the old episode. Their function is not to provide an exact reproduction of the past but to provide a useful guide for the future. As such, we can assume quite straightforwardly that these episodes will become more useful on retrieval the more information they manage to integrate from the training situation.

The next experiment was a test of this idea. It was similar to the previous ones but tested the hypothesis that the mental representation of temporal contexts is updated with new information that is available when they are retrieved. To do so, we cued half of the participants in this experiment to recall the temporal context of Phase 1. However, we included a change so that the episode actually presented was different from what the participant had retrieved and was therefore expecting to occur. This should cause participants to update the episode and their representation of the

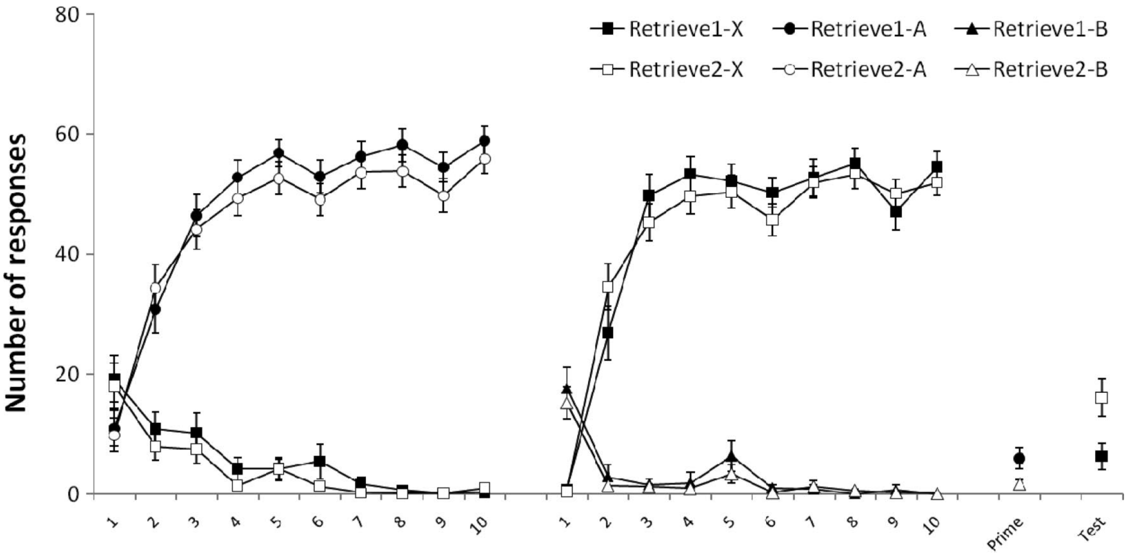

Figure 3. Mean number of responses during the different phases of Experiment 3 as a function of group and cue. Because of a programming error, the data from the distracting C–O2 trials in this experiment were not recorded. Error bars represent the standard errors of the mean.

Phase 1 temporal context. If this is to be transferred into behavior, as our previous experiments suggest, participants should no longer show the behavior appropriate to the temporal context of Phase 1 but a behavior that includes the update induced by our manipulation.

To this end, we changed our design to a miscuing paradigm (e.g., Lipp & Dal Santo, 2002; Lipp, Siddle, & Dall, 1993), which provides the ideal ingredients to test our prediction. In the standard miscuing procedure, the participants first receive one phase of training in which a given cue, X, is always followed by a given outcome (O1) and another cue, A, is followed by a second outcome (O2). Once this training is established, a single miscuing trial follows. In this trial, Cue A that was originally paired with O2 is now paired with O1, which is contrary to the participant's expectations. The effect of interest is that, as a consequence of this miscuing trial, participants update their expectations, and their behavior with respect to the other stimulus, X, changes.

The design is shown in Table 4. During Phase 1, participants learned that X produced a given outcome, O1, and that A produced O2. The X–O1 and A–O2 trials were randomly presented. During Phase 2, participants learned that a new target cue, Y, produced O1, whereas another cue, B, produced O2. Then, during Phase 3, participants received one trial with either Cue A or Cue B, the cues that had been trained as predictors of O2 during either Phase 1 or Phase 2. Therefore, this should retrieve the corresponding temporal context of either Phase 1 or Phase 2. It is important to note that Cues A and B were now followed by the other outcome, O1. Thus, A and B at this point were no longer mere reminders of an episode; they retrieved the episode and, because of their pairing with a different outcome, caused it to be updated. Thus, the old representation of the temporal context of Phase 1 should now be updated with this new evidence. What should participants do now when presented with Target Cues X and Y at test? If our reasoning is correct, participants should adapt their behavior so that the change in the predictive status of Cue A should also change how participants respond to the target cue that was originally trained in the same temporal context. In other words, the miscuing effect should affect selectively one target cue or the other, depending on which priming cue, A or B, was presented immediately before the test. These predictions go beyond those of any individual model of memory or learning that we are aware of.

An additional purpose of this experiment was to ensure that the observed effects were not specific to the particular task that we used in Experiments 1–3. Thus, we used an alternative task that resembles a Martians videogame for this experiment (see footnote 2). Finally, the experiment was aimed at verifying that the present results cannot only be observed in experiments run via the Internet, as in Experiments 2 and 3, but also in the laboratory, as in Experiment 1.

Table 4
Design Summary of Experiment 4

	Training			
Group	Phase 1	Phase 2	Phase 3	Test
M1T1	X–O1/A–O2	Y–O1/B–O2	A–O1	X
M1T2	X–O1/A–O2	Y–O1/B–O2	A–O1	Y
M2T1	X–O1/A–O2	Y–O1/B–O2	B–O1	X
M2T2	X–O1/A–O2	Y–O1/B–O2	B–O1	Y

Note. Cues A–Y were background colors in the Martians task and could predict two different outcomes: O1 (a Martian invasion) and O2 (no invasion). In each phase, the different trial types were randomly intermixed. Phase 3 consists of just one trial. There was no separation between phases. M1 and M2 refer to whether Phase 1 or Phase 2 was miscued; T1 and T2 indicate whether testing took place with the target cue from Phase 1 or from Phase 2.

Method

Participants and apparatus. Sixty-four undergraduate students from the University of Queensland volunteered to participate for course credit. On arrival at the laboratory, participants were randomly assigned to one of four experimental groups.

Design and procedure. Participants were run individually using a Martians videogame (see footnote 2) that resembles conditioning and is often used to study associative learning with humans (Arcediano et al., 1996; Costa & Boakes, 2011; Franssen et al., 2010; Lipp & Dal Santo, 2002). The aim of participants in this game is to prevent an invasion of Martians. Martians appear on the screen and can be destroyed by firing a laser gun (i.e., pressing the space bar) just before they appear. A successful shot is indicated by the appearance of an explosion. Martians appear at a steady rate on the screen of approximately four to five per second. Participants are first trained to press the space bar regularly to destroy as many Martians as possible. Thus, the purpose of this preliminary phase is simply to produce stable baseline behavior of about four to five responses per second, against which the focal behavioral changes will be assessed during the subsequent experiment.

Once this behavior is well established, participants receive a new instructional set and the experiment proper starts. They are now informed that the Martians have developed an antilaser shield and that, if they fire the gun while the shield (a white flashing screen) is connected, hundreds of Martians will invade their screen immediately. The flashing screen lasts just for 0.5 s, which means that participants need to find a way to predict its activation to be able to cease responding before the flashing occurs. Otherwise, they will suffer an invasion. They are explicitly encouraged to learn to predict when the Martians are about to connect the shield, although they are not told how to do it. Thus, in this paradigm, participants are responding regularly to Martians during the intertrial intervals and a trial begins when a cue is presented. The cues are changes of the screen background color. Some of these cues signal the activation of the shield and, thus, a potential invasion (O1). Others are followed by nothing (O2).[3] The dependent measure is the degree to which ongoing behavior is suppressed when a cue is presented, which is indicative of the participant's expectation that O1 will follow.

In this experiment, cues were presented for 1 s during training with the exception of the last presentation of each cue, which lasted for 2 s, and the test trial presentations, which lasted for 3 s each. The reason for doing so is that 1 s is normally not enough to properly assess the degree of suppression. However, if participants are always given 3 s for each cue, they learn this temporal pattern and keep responding as much as they can during the presentation of the cue (Arcediano et al., 1996). In this case, no assessment of the dependent variable can take place either. The solution suggested by Arcediano et al. (1996) was to use a longer presentation of cues only in trials in which assessment of the dependent variable was most critical and keep the other trials to a shorter duration. In this way, the procedure is most sensitive. Given that the test trial and the latest training trial were the most critical in this experiment, we opted for using longer cues that would allow assessment only in those trials.

During Phase 1, all participants received eight trials with Cue X, which always signaled danger (i.e., O1), and 10 trials with Cue A, which signaled nothing (O2). Then, during Phase 2, they received 10 trials with Cue Y, which always signaled danger (i.e., O1), and 10 trials with Cue B, which signaled nothing (i.e., O2). The order with which the different types of trials were presented within each phase was randomized, as was the duration of the intertrial intervals, which ranged from 6 to 12 s. The colors that served as Cues A, B, X, and Y were yellow, blue, light blue, and red, all counterbalanced. The background color of the screen was black during the intertrial intervals.

Phase 3 consisted of just one trial. As in the preceding experiments and because we used no breaks between phases, this Phase 3 trial was actually the last trial of Phase 2 and the one immediately preceding test. This critical Phase 3 manipulation consisted of a miscuing trial in which the Martian invasion was now signaled by one of the cues (i.e., A or B) that in a previous temporal context (Phase 1 or Phase 2) had predicted nothing. This miscuing trial had the purpose of (a) cuing the temporal context of either Phase 1 or Phase 2 and (b) making participants integrate the current information with the representation of the older temporal context, as the cue that they would expect to predict nothing was now a predictor for danger. If our hypothesis is correct and participants do update the representation of an old temporal context as they get new information, this update should be reflected in a corresponding behavioral change at test that should affect other stimuli trained in that context (in this case, Cues X or Y). Thus, during a subsequent test trial, participants should no longer respond to the target cues, X and Y, as they had learned previously. Instead, they should respond to them in just the opposite way. It is important to note that this change should occur only if the cue presented in Phase 3 had been trained in the same temporal context as the target cue. Thus, half of the participants in each group were tested with the target cue from Phase 1, X, whereas the other half were tested with the target cue from Phase 2, Y. In total, there were four groups (M1T1, M1T2, M2T1, and M2T2) as a function of whether they received miscuing with a cue from Phase 1 (M1) or Phase 2 (M2) and as a function of whether they were tested with a target cue from Phase 1 (T1) or Phase 2 (T2).

Results and Discussion

In associative learning studies in which the target behavior consists of suppressing ongoing behavior, a suppression ratio is normally calculated so that the number of responses (i.e., in this case, presses of the space bar) that occur during the cue presentation can be compared with the base rate of responses immediately preceding the presentation of the cue. The dependent variable is thus computed as the ratio of (responses during cue)/(responses during cue + responses during an equivalent time period preceding the cue).

The critical result that we expect in this experiment is not a mere difference between the groups at test, as in the previous experiments. Instead, we expect the miscuing trial to change the manner in which participants respond to the target cue relative to the last

---

[3] It should be noted that in this experiment, O2 is not a positive outcome, as in the previous experiments. However, given that the participants' motivation throughout this experiment is to maintain a stable and high rate of responding, the fact that some cues are followed by nothing means that participants should keep responding during those cues.

training trial. Thus, the result of interest is the degree to which responding at test differs from responding during the last training trial, which, as stated above, is the reason why these two trials were given longer durations than the other ones, thereby allowing for the critical suppression ratios to be assessed. Figure 4 shows these results.

As can be seen in Figure 4, a reduction in suppression from the last training trial to test was evident when the stimuli presented during the miscuing trial and during test had been trained together in the same phase but not when the stimuli were trained in different phases. This impression was confirmed in a 2 (last training trial vs. test trial) × 2 (miscuing cue from Phase 1 vs. miscuing cue from Phase 2) × 2 (target cue from Phase 1 vs. target cue from Phase 2) factorial analysis of variance, which yielded a main effect for trial, $F(1, 60) = 18.09$, $p < .001$, and Miscuing Cue × Test Cue, $F(1, 60) = 13.18$, $p < .001$, and Trial × Miscuing Cue × Test Cue interactions, $F(1, 60) = 5.45$, $p < .05$. Follow-up tests found less suppression during test than during the last training trial in participants presented with miscuing and test cues that had been trained together in Phase 1, $t(60) = 4.33$, or Phase 2, $t(60) = 2.80$. There was no effect of miscuing on suppression during test if the miscuing cue and the test cue had been trained in different phases, both ts < 2.1. This suggests that the miscuing trial led to the updating of other memories and predictions associated with cues trained in the same temporal context but left behavior dependent on cues trained in different temporal contexts untouched.

General Discussion

The results of the present experiments show that events that are trained in the same phase of the experiment can activate the mental representation of each other even when there is no direct contiguity or contingency between them. The most widely accepted theories of associative learning and memory do not anticipate this result (e.g., Raaijmakers & Shiffrin, 1980; Rescorla & Wagner, 1972).

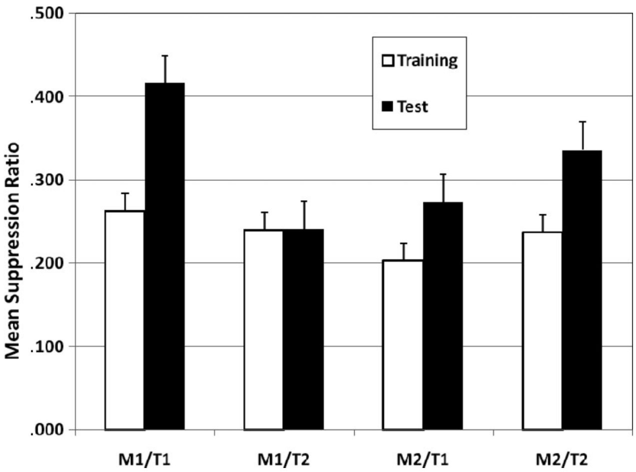

Figure 4. Mean suppression ratios during the last trial of training and the test trial of Experiment 4 as a function of whether the miscuing and test cue had been trained in the same temporal context (M1/T1, M2/T2) or in different temporal contexts (M1/T2, M2/T1; error bars represent the standard error of the mean). Note that this figure needs to be read in the opposite way as the previous ones. Values closer to zero indicate stronger suppression (with zero being complete suppression of ongoing behavior in response to the stimulus, therefore, a very strong response). In contrast, values closer to .500 indicate that the number of responses during the cue differs little from the number of responses during an equivalent baseline period. Thus, .500 indicates null responding to the stimulus.

The results of our Experiments 1–3 were predicted by the TCM model of Howard and Kahana (2002; Sederberg et al., 2008). They were obtained using a task that is very different from those that have traditionally been used to test this model. More specifically, we used a methodology that comes from the associative learning tradition, which speaks to the notion of a common framework for episodic memory and associative learning. In essence, the TCM explains the results of Experiments 1–3 by assuming that people are continuously forming and updating temporal contexts. To do so, people use all stimuli that are available at any given time. Each stimulus becomes associatively linked to these contexts, and then, as the contexts evolve with time and get linked to other stimuli, the older stimuli, if presented again, can be used to cue the old context along with all of the information and memories from the original training episode. The results of Experiments 1–3 support this view. Priming a given temporal context at test with one of the stimuli that had been trained in that context made participants retrieve the meaning of the target cue trained in that context. Thus, responding to the target cue at test was dependent on the retrieved context information.

Results of Experiment 4 support that view as well but go beyond the predictions of the TCM. Here, we used a cue from a particular temporal context to retrieve the memory of that context. Once participants had retrieved that memory and were, therefore, expecting that a given outcome followed the cue, we presented the other outcome. This led participants to update the representation of the old temporal context. As a consequence, they also changed their behavior when the second stimulus from that context was presented at test. This behavioral change did not transfer to stimuli that had been trained in a different temporal context. Note that the TCM predicts only part of this. It does predict the update of the temporal context representation, but it does not specify whether and how this should affect the behavior to the other stimulus. Indeed, these results suggest that the episodes were created by the sets of stimuli that were trained in a temporal window. Moreover, the episode is encoded in such a way that retrieving it and altering what traditional theories of associative learning would regard as a single independent association trained in that episode can also alter other associations trained independently in the same episode.

We are not claiming that this has to be an automatic process of episodic restructuring, but, functionally, it works like it. The possibility that deliberate processes might have caused this result instead should be subject to future research. Indeed, it could be argued that our participants were creating different rules for each phase and then going back to these rules on retrieval of one or the other context so that they could then flexibly use or modify these rules as needed. Note that this rule-based account does not depart from the main point we are making. That is, the rules that apply during a given time period can also create a context or episode. What is clear, however, is that regardless of whether this process is deliberate or automatic, the results show that there is a psychological reality to the contexts in Phase 1 and Phase 2. Otherwise, the disconfirmation of one expectation during miscuing should have changed all other expectations in the experiment. The fact that it changed the expectation only for the second item that had occurred in that phase but not the expectation for the items that occurred in the other phase supports this psychological reality. This result was not predicted by any of the standard models of learning or memory and is a good example of how an integration of the learning and memory traditions can produce fruitful advances.

Our experiments show that people can update their representation of a retrieved temporal context as a function of new information that is available in the current situation. Experiment 4 demonstrated that retrieving an old temporal context and updating it with the changes that occurred in the current situation helped participants adjust their future behavior, even though this adaptation required responding to the target cue in just the opposite way as was done during the acquisition stage. This suggests the existence of a very direct link between remembering the past and planning for the future, as has already been suggested by many other researchers (e.g., Schacter et al., 2007; Suddendorf & Corballis, 1997, 2007; Tulving, 2005). That is, memory provides people with knowledge and experiences that feed learning, and this integration of memory and learning helps people plan for the future and adapt future behavior to the most likely changes in the environment.

Bouton's theory of contexts in associative learning cannot explain our results. It predicts that only the associations acquired second (Bouton, 1997; or only the inhibitory associations; see Bouton, 1993) become context specific. This means that in our experiments, only the associations acquired second (or the inhibitory associations) should have been subject to selective retrieval. This was not the case. Rosas et al.'s (2006) revision of Bouton's theory might perhaps be better equipped to assume that either the associations acquired first or the associations acquired second may be subject to selective retrieval. This theory, however, is as silent as Bouton's concerning how temporal contexts may be created and represented.

It should be noted, however, that several components of the TCM were foreshadowed in the associative learning literature. Bouton's (1993, 1997) theory assumes time-dependent changes in context (although the temporal context according to Bouton is provided by time itself rather than explicit stimuli that occur within a period of time) and Capaldi's (1994) sequential theory of extinction assumes that trials can provide contexts for other trials in a series (see also Bouton, Woods, & Pineño, 2004). From this perspective, it could be argued that what was manipulated in the present experiments was not time but explicit stimuli and trials. Admittedly, we cannot manipulate time per se and travel back in time to a previous phase. For this reason, we manipulated stimuli and trials instead, under the assumption that, as suggested by the TCM, the stimuli that are present at a given time are what define that temporal context. This, we believe, is a powerful approach. It allows specific predictions concerning (backward and forward) mental time travel. It makes a clear prediction about how contexts should be represented, how they should change with time, and how their representation might be accessed. The TCM also makes a specific prediction about associations being formed between the stimuli that occur in a given temporal context. Our results clearly confirm this prediction of the TCM.

However, it may be possible that some other mechanism different from those proposed by the TCM be involved. One alternative would be to assume that Cues A and B defined two distinct physical contexts. Physical contexts, however, are generally defined by cues that are continuously present and that do not change from one trial to the next. Indeed, physical contexts are assumed to get compound, simultaneous training with the target cues according to mainstream theories of learning (e.g., Bouton, 1993; Miller & Matzel, 1988; Pearce, 1994; Rescorla & Wagner, 1972). This was not the case in our experiments. The notion of temporal (rather than physical) contexts seems to be, at least in principle, a more parsimonious explanation for our results.

A second alternative is to assume that the majority of the stimuli that are trained in the same phase are more contiguous (and contingent) to each other than to the majority of the stimuli trained in a different phase. This would allow the formation of direct associations between the stimuli that were trained in the same phase, on the basis of their relative contiguity and contingency to each other as compared with stimuli trained in different phases. That is, contiguity and contingency could probably be relative to the temporal window that the participant is using at a given time.

It should be noted, however, that the design that we used explicitly assured that the training of, say, Target Cue X with Outcome O1 was independent of the training of the second cue that was trained in the same phase, say, Target Cue A with Outcome O2. The order in which those two types of trials occurred in each phase was randomized. Sometimes the X–O1 pairings preceded the pairings of A–O2, sometimes this order was reversed, yet at other times several trials of X–O1 or A–O2 occurred in succession. In essence, the design that we used can be thought of as identical to that of a typical discrimination learning design in which a given cue is followed by an outcome and a different cue is followed by another outcome, with the order of these two trial types randomized (e.g., Arcediano et al., 1996; Chiszar & Spear, 1969; Vila, Romero, & Rosas, 2002). It is also identical to many within-subject designs that train several independent conditions, such as an experimental condition and a control condition, by randomly intermixing different types of learning trials during the same phase. In those cases, traditional learning theories would hold that no contiguity or contingency exists between the different trial types (i.e., between the experimental and control conditions) and that no associations are formed between them.

Nevertheless, some might argue that even such weak and relative contiguity and contingency might be sufficient to create associations. This view would require specification of the conditions under which relative contiguity would be sufficient to produce associations among stimuli and how those associations would develop. For instance, why should the latest trials of Phase 1 become associated with other trials of Phase 1 rather than with the more contiguous trials presented at the beginning of Phase 2? In essence, developing such a proposal would be very similar to suggesting that participants created temporal contexts that coincided with those designed by the experimenters. The advantage that we see in applying the TCM is that it has already formalized those ideas and can be used to derive testable predictions.

Although we have used two different tasks that differed in the dependent variable and several other details, it is also true that both tasks involved the gain or loss of points in a computer game. It would be interesting to know whether similar results can be observed in other situations that do not require the gain or loss of points. Indeed, there is a tradition within the associative learning literature that assumes that when different stimuli are trained in the same way or are associated with the same set of objects, they become functionally equivalent. This is called *mediated generalization* and it has been used to explain concept formation and different aspects of language acquisition (Keller & Schoenfeld, 1950; Wasserman, DeVolder, & Coppage, 1992). This account could perhaps be used to explain the results of our Experiment 4. Given that A and X were trained in the same temporal context (here playing the role of the common set of objects), this kind of mediated generalization might explain why the new information about A generalizes to X but not to other cues trained in other temporal contexts. Our view, however, is that this is a description of the results more than an explanation. Moreover, this account could also accommodate the opposite result: If whatever is learned about A is generalized to X, we should expect less miscuing in Experiment 4, which is contrary to the results of the experiment.

Our results suggest that participants were able to construct different temporal windows, one for each of the phases designed by the experimenters. Exploring how they do it seems to be a very exciting venue for future research. TCM provides a possible explanation for temporal context formation, and experiments from the associative tradition can shed light on how this process takes place. Indeed, the spontaneous separation of two experimental phases into separate temporal contexts has often been observed in human learning experiments (e.g., Alvarado, Jara, Vila, & Rosas, 2006; Collins & Shanks, 2002; Lipp & Purkis, 2006; Matute, Vegas, & De Marez, 2002; Vadillo, Vegas, & Matute, 2004). It is important to note that although those experiments were not explicitly designed to test this particular question, many of them showed that participants can use subtle, nonphysical cues to construct efficient and flexible temporal contexts (Collins & Shanks, 2002; Lipp & Purkis, 2006; Matute et al., 2002; Vadillo et al., 2004). For instance, Vadillo et al. (2004) used a contingency judgment task in which the cues and outcomes were fictitious medicines and allergic reactions, respectively. These tasks are very different from the ones used in the present article (e.g., the dependent variable is the number of presses of the space bar in the present research, whereas it is a subjective judgment of causality or prediction in contingency judgment studies). Participants were required to provide their judgment either in every trial or at the end of training. The (physical) context in which these judgments occurred was manipulated orthogonally.

It is interesting that the results of those experiments showed that changing the frequency with which the judgments were requested during testing with respect to training (e.g., changing from no request during training to a request at test) produced an effect that was similar to changing the physical context between training and testing. That is, participants created flexible temporal contexts as a function of when and how often they were requested to provide their judgments. When the physical contexts did not change and participants were asked to provide only one judgment at the end of training, participants assumed that the relevant temporal context to consider was that of the entire experiment, that is, the one including the two learning phases. When the physical contexts did not change but participants were asked to provide frequent judgments through the training phases, participants assumed that the relevant temporal context to consider at test was that of the most recent phase. These results suggest that participants used the differences in the frequency with which they were asked to provide their judgments in the different phases of the experiment as a subtle clue that helped them to decide when exactly to segment the continuous flow of information, so that they could create the different temporal contexts or episodes (see also Kurby & Zacks, 2008, for discussion on how segmentation of events may take place).

The present results should also be linked to the experimental tradition of associative interference in paired associate learning. There is substantial evidence that in the AB AC paradigm, the first list association can coexist with the second list association. When participants are asked to recall the items associated with A in the first list and in the second list, the recall probabilities appear to be independent (Martin, 1971). Thus, there is little or no evidence that there is any destructive interference (Dyne, Humphreys, Bain, & Pike, 1990), and it has been shown that this pattern of results can be replicated by a model that depends heavily on context to keep the two responses separate (Chappell & Humphreys, 1994). This lack of evidence for destructive interference is compatible with the account presented in this article and with associative learning accounts of the effect of extinction training and contexts on acquired associations (Bouton, 1993, 1997).

In summary, our experiments add to the convergent evidence originating from findings in the memory tradition using paired associate paradigms and word lists (e.g., Humphreys et al., 2009; Humphreys, Bain, & Pike, 1989; Sederberg et al., 2008) and from the associative learning tradition (Bouton, 1993, 1997; Capaldi, 1994) by providing evidence (a) for the retrieval of the representation of a temporal context through a cue trained in that context (as opposed to retrieving it using instructions; Humphreys et al., 2009) and (b) for the spontaneous creation of a context or an episode through the stimulus contingencies found in a temporal window. Also, using nonverbal behavior as the dependent variable provides interesting confirming evidence that speaks of the generality of these effects across different tasks and measures. These results indicate the promise of our approach to integrate concepts and procedures from the associative learning and episodic memory traditions to advance the understanding of the basic processes involved in using past experiences for adaptive behavior in the present and future.

References

Alvarado, A., Jara, E., Vila, J., & Rosas, J. M. (2006). Time and order effects on causal learning. *Learning and Motivation, 37,* 324–345. doi:10.1016/j.lmot.2005.11.001

Arcediano, F., Ortega, N., & Matute, H. (1996). A behavioural preparation for the study of human Pavlovian conditioning. *The Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology, 49*(B), 270–283.

Bouton, M. E. (1993). Context, time, and memory retrieval in the interference paradigms of Pavlovian conditioning. *Psychological Bulletin, 114,* 80–99. doi:10.1037/0033-2909.114.1.80

Bouton, M. E. (1997). Signals for whether versus when an event will occur. In M. E. Bouton & M. S. Fanselow (Eds.), *Learning, motivation, and cognition: The functional behaviorism of Robert C. Bolles* (pp. 385–409). Washington, DC: American Psychological Association. doi:10.1037/10223-019

Bouton, M. E., Woods, A. M., & Pineño, O. (2004). Occasional reinforced trials during extinction can slow the rate of rapid reacquisition. *Learning and Motivation, 35,* 371–390. doi:10.1016/j.lmot.2004.05.001

Brooks, D. C., & Bouton, M. E. (1993). A retrieval cue for extinction attenuates spontaneous recovery. *Journal of Experimental Psychology: Animal Behavior Processes, 19,* 77–89. doi:10.1037/0097-7403.19.1.77

Capaldi, E. (1994). The sequential view: From rapidly fading stimulus traces to the organization of memory and the abstract concept of number. *Psychonomic Bulletin & Review, 1,* 156–181. doi:10.3758/BF03200771

Chappell, M., & Humphreys, M. S. (1994). An auto-associative neural network for sparse representations: Analysis and application to models of recognition and cued recall. *Psychological Review, 101,* 103–128. doi:10.1037/0033-295X.101.1.103

Chiszar, D. A., & Spear, N. E. (1969). Stimulus change reversal learning, and retention in the rat. *Journal of Comparative and Physiological Psychology, 69,* 190–195. doi:10.1037/h0027947

Collins, D. J., & Shanks, D. R. (2002). Momentary and integrative response strategies in causal judgment. *Memory & Cognition, 30,* 1138–1147. doi:10.3758/BF03194331

Costa, D. S. J., & Boakes, R. A. (2011). Varying temporal contiguity and interference in a human avoidance task. *Journal of Experimental Psychology: Animal Behavior Processes, 37,* 71–78. doi:10.1037/a0021192

Davis, O. C., Geller, A. S., Rizzuto, D. S., & Kahana, M. J. (2008). Temporal associative processes revealed by intrusions in paired-associate recall. *Psychonomic Bulletin & Review, 15,* 64–69. doi:10.3758/PBR.15.1.64

Dyne, A. M., Humphreys, M. S., Bain, J. D., & Pike, R. (1990). Associative interference effects in recognition and recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 813–824. doi:10.1037/0278-7393.16.5.813

Franssen, M., Clarysse, J., Beckers, T., van Vooren, P., & Baeyens, F. (2010). A free software package for a human online-conditioned suppression preparation. *Behavior Research Methods, 42,* 311–317. doi:10.3758/BRM.42.1.311

Garcia, J., Ervin, F. R., & Koelling, R. A. (1966). Learning with prolonged delay of reinforcement. *Psychonomic Science, 5,* 121–122.

Gawronski, B., Rydell, R. J., Vervliet, B., & De Houwer, J. (2010). Generalization versus contextualization in automatic evaluation. *Journal of Experimental Psychology: General, 139,* 683–701. doi:10.1037/a0020315

Gunther, L. M., Miller, R. R., & Matute, H. (1997). CSs and USs: What's the difference? *Journal of Experimental Psychology: Animal Behavior Processes, 23,* 15–30. doi:10.1037/0097-7403.23.1.15

Howard, M. W., Jing, B., Rao, V. A., Provyn, J. P., & Datey, A. V. (2009). Bridging the gap: Transitive associations between items presented in similar temporal contexts. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35,* 391–407. doi:10.1037/a0015002

Howard, M. W., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25,* 923–941. doi:10.1037/0278-7393.25.4.923

Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology, 46,* 269–299. doi:10.1006/jmps.2001.1388

Howard, M. W., Youker, T. E., & Venkatadass, V. (2008). The persistence of memory: Contiguity effects across several minutes. *Psychonomic Bulletin & Review, 15,* 58–63. doi:10.3758/PBR.15.1.58

Humphreys, M. S., Bain, J. D., & Pike, R. (1989). Different ways to cue a coherent memory system: A theory for episodic, semantic and procedural tasks. *Psychological Review, 96,* 208–233. doi:10.1037/0033-295X.96.2.208

Humphreys, M. S., Murray, K. L., & Maguire, A. M. (2009). Context and control operations used in accessing list-specific, generalized, and semantic memories. *Cognitive Psychology, 58,* 311–337. doi:10.1016/j.cogpsych.2008.10.001

Keller, F. S., & Schoenfeld, W. N. (1950*). Principles of psychology.* New York, NY: Appleton-Century-Crofts.

Kraut, R., Olson, J., Banaji, M., Bruckman, A., Cohen, J., & Couper, M. (2004). Psychological research online: Report of Board of Scientific Affairs' Advisory Group on the conduct of research on the Internet. *American Psychologist, 59,* 105–117. doi:10.1037/0003-066X.59.2.105

Kurby, C. A., & Zacks, J. M. (2008). Segmentation in the perception and memory of events. *Trends in Cognitive Sciences, 12,* 72–79. doi:10.1016/j.tics.2007.11.004

Lipp, O. V., & Dal Santo, L. A. (2002). Cue competition between elementary trained stimuli: US miscuing, interference, and US omission. *Learning and Motivation, 33,* 327–346. doi:10.1016/S0023-9690(02)00001-2

Lipp, O. V., & Purkis, H. M. (2006). The effects of assessment type on verbal ratings of conditional stimulus valence and contingency judgments: Implications for the extinction of evaluative learning. *Journal of Experimental Psychology: Animal Behavior Processes, 32,* 431–440. doi:10.1037/0097-7403.32.4.431

Lipp, O. V., Siddle, D. A. T., & Dall, P. J. (1993). Effects of miscuing on Pavlovian conditioned responding and on probe reaction time. *Australian Journal of Psychology, 45,* 161–167. doi:10.1080/00049539308259134

Martin, E. (1971). Verbal learning theory and independent retrieval phenomena. *Psychological Review, 78,* 314–332. doi:10.1037/h0031030

Matute, H., Vadillo, M. A., & Bárcena, R. (2007). Web-based experiment control software for research and teaching on human learning. *Behavior Research Methods, 39,* 689–693. doi:10.3758/BF03193041

Matute, H., Vegas, S., & De Marez, P. J. (2002). Flexible use of recent information in causal and predictive judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28,* 714–725. doi:10.1037/0278-7393.28.4.714

Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 51–92). San Diego, CA: Academic Press.

Pavlov, I. P. (1927). *Conditioned reflexes.* London, England: Clarendon Press.

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

Pineño, O., Ortega, N., & Matute, H. (2000). The relative activation of the associations modulates interference between elementally-trained cues. *Learning and Motivation, 31,* 128–152. doi:10.1006/lmot.1999.1047

Raaijmakers, J. G. W., & Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 14, pp. 207–262). New York, NY: Academic Press.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review, 88,* 93–134. doi:10.1037/0033-295X.88.2.93

Ranganath, C. (2010). Binding items and contexts: The cognitive neuroscience of episodic memory. *Current Directions in Psychological Science, 19,* 131–137. doi:10.1177/0963721410368805

Reips, U.-D. (2002). Internet-based psychological experimenting: Five dos and five don'ts. *Social Science Computer Review, 20,* 241–249.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning: II. Current research and theory* (pp. 64–99). New York, NY: Appleton-Century-Crofts.

Rosas, J. M., & Callejas-Aguilera, J. E. (2006). Context switch effects on acquisition and extinction in human predictive learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 461–474. doi:10.1037/0278-7393.32.3.461

Rosas, J. M., Callejas-Aguilera, J. E., Ramos-Álvarez, M. M., & Fernández-Abad, M. J. (2006). Revision of retrieval theory of forgetting: What does make information context-specific? *International Journal of Psychology and Psychological Therapy, 6,* 147–166.

Schacter, D. L., Addis, D. R., & Buckne, R. L. (2007). Remembering the past to imagine the future: The prospective brain. *Nature Reviews Neuroscience, 8,* 657–661. doi:10.1038/nrn2213

Schwartz, G., Howard, M. W., Jing, B., & Kahana, M. J. (2005). Shadows of the past: Temporal retrieval effects in recognition memory. *Psychological Science, 16,* 898–904. doi:10.1111/j.1467-9280.2005.01634.x

Sederberg, P. B., Howard, M. W., & Kahana, M. J. (2008). A context-based theory of recency and contiguity in free recall. *Psychological Review, 115,* 893–912. doi:10.1037/a0013396

Shanks, D. R., & Dickinson, A. (1987). Associative accounts of causality judgment. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 21, pp. 229–261). San Diego, CA: Academic Press.

Suddendorf, T., & Corballis, M. C. (1997). Mental time travel and the evolution of the human mind. *Genetic, Social, and General Psychology Monographs, 123,* 133–167. Retrieved from http://cogprints.org/725/

Suddendorf, T., & Corballis, M. C. (2007). The evolution of foresight: What is mental time travel and is it unique to humans? *Behavioral and Brain Sciences, 30,* 299–313. doi:10.1017/S0140525X07001975

Tulving, E. (2005). Episodic memory and autonoesis: Uniquely human? In H. S. Terrace & J. Metcalfe (Eds.), *The missing link in cognition: Origins of self-reflective consciousness* (pp. 3–56). New York, NY: Oxford University Press.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review, 108,* 550–592. doi:10.1037/0033-295X.108.3.550

Vadillo, M. A., Bárcena, R., & Matute, H. (2006). The Internet as a research tool in the study of associative learning: An example from overshadowing. *Behavioural Processes, 73,* 36–40. doi:10.1016/j.beproc.2006.01.014

Vadillo, M. A., & Matute, H. (2011). Further evidence on the validity of Web-based research on associative learning: Augmentation in a predictive learning task. *Computers in Human Behavior, 27,* 750–754. doi:10.1016/j.chb.2010.10.020

Vadillo, M. A., Vegas, S., & Matute, H. (2004). Frequency of judgment as a context-like determinant of predictive judgments. *Memory & Cognition, 32,* 1065–1075. doi:10.3758/BF03196882

Vila, J., Romero, M., & Rosas, J. M. (2002). Retroactive interference after discrimination reversal decreases following temporal and physical context changes in human subjects. *Behavioural Processes, 59,* 47–54. doi:10.1016/S0376-6357(02)00063-3

Wasserman, E. A. (1990). Detecting response–outcome relations: Toward an understanding of the causal texture of the environment. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 26, pp. 27–82). San Diego, CA: Academic Press.

Wasserman, E. A., DeVolder, C. L., & Coppage, D. J. (1992). Non-similarity-based conceptualization in pigeons via secondary or mediated generalization. *Psychological Science, 3,* 374–379. doi:10.1111/j.1467-9280.1992.tb00050.x

Received August 31, 2010
Revision received March 30, 2011
Accepted April 1, 2011 ■

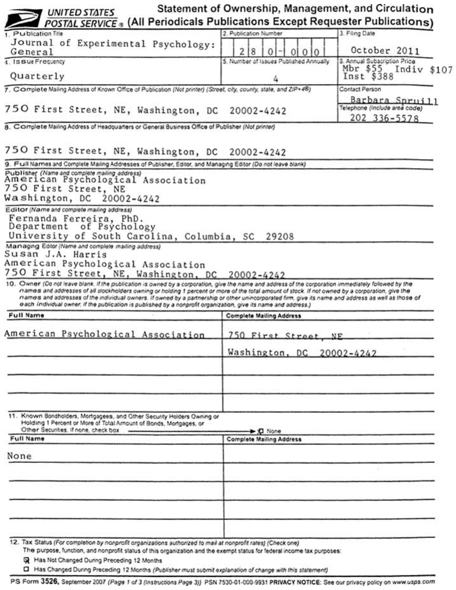

UNITED STATES POSTAL SERVICE®

Statement of Ownership, Management, and Circulation (All Periodicals Publications Except Requester Publications)

1. Publication Title: Journal of Experimental Psychology: General
2. Publication Number: 2 8 0 - 0 0 0
3. Filing Date: October 2011
4. Issue Frequency: Quarterly
5. Number of Issues Published Annually: 4
6. Annual Subscription Price: Mbr $55 Indiv $107 Inst $388
7. Complete Mailing Address of Known Office of Publication *(Not printer) (Street, city, county, state, and ZIP+4®)*: 750 First Street, NE, Washington, DC 20002-4242
Contact Person: Barbara Spruill
Telephone *(Include area code)*: 202 336-5578
8. Complete Mailing Address of Headquarters or General Business Office of Publisher *(Not printer)*: 750 First Street, NE, Washington, DC 20002-4242
9. Full Names and Complete Mailing Addresses of Publisher, Editor, and Managing Editor *(Do not leave blank)*

Publisher *(Name and complete mailing address)*
American Psychological Association
750 First Street, NE
Washington, DC 20002-4242

Editor *(Name and complete mailing address)*
Fernanda Ferreira, PhD.
Department of Psychology
University of South Carolina, Columbia, SC 29208

Managing Editor *(Name and complete mailing address)*
Susan J.A. Harris
American Psychological Association
750 First Street, NE, Washington, DC 20002-4242

10. Owner *(Do not leave blank. If the publication is owned by a corporation, give the name and address of the corporation immediately followed by the names and addresses of all stockholders owning or holding 1 percent or more of the total amount of stock. If not owned by a corporation, give the names and addresses of the individual owners. If owned by a partnership or other unincorporated firm, give its name and address as well as those of each individual owner. If the publication is published by a nonprofit organization, give its name and address.)*

Full Name	Complete Mailing Address
American Psychological Association	750 First Street, NE Washington, DC 20002-4242

11. Known Bondholders, Mortgagees, and Other Security Holders Owning or Holding 1 Percent or More of Total Amount of Bonds, Mortgages, or Other Securities. If none, check box → ☒ None

Full Name	Complete Mailing Address
None	

12. Tax Status *(For completion by nonprofit organizations authorized to mail at nonprofit rates) (Check one)*
The purpose, function, and nonprofit status of this organization and the exempt status for federal income tax purposes:
☒ Has Not Changed During Preceding 12 Months
☐ Has Changed During Preceding 12 Months *(Publisher must submit explanation of change with this statement)*

PS Form **3526**, September 2007 *(Page 1 of 3 (Instructions Page 3))* PSN 7530-01-000-9931 **PRIVACY NOTICE**: See our privacy policy on *www.usps.com*

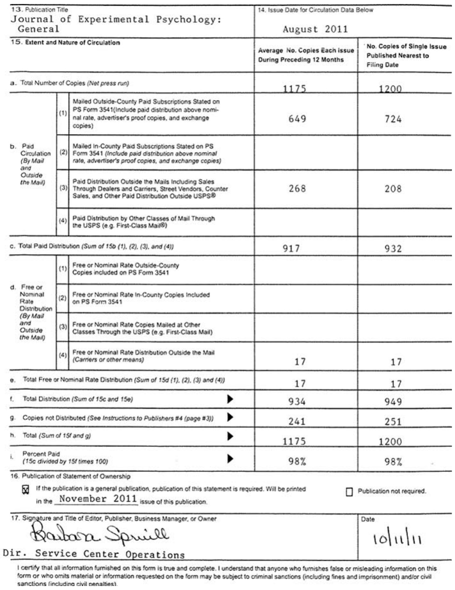

13. Publication Title: Journal of Experimental Psychology: General
14. Issue Date for Circulation Data Below: August 2011

15. Extent and Nature of Circulation		Average No. Copies Each Issue During Preceding 12 Months	No. Copies of Single Issue Published Nearest to Filing Date
a. Total Number of Copies *(Net press run)*		1175	1200
b. Paid Circulation *(By Mail and Outside the Mail)*	(1) Mailed Outside-County Paid Subscriptions Stated on PS Form 3541(Include paid distribution above nominal rate, advertiser's proof copies, and exchange copies)	649	724
	(2) Mailed In-County Paid Subscriptions Stated on PS Form 3541 *(Include paid distribution above nominal rate, advertiser's proof copies, and exchange copies)*		
	(3) Paid Distribution Outside the Mails Including Sales Through Dealers and Carriers, Street Vendors, Counter Sales, and Other Paid Distribution Outside USPS®	268	208
	(4) Paid Distribution by Other Classes of Mail Through the USPS (e.g. First-Class Mail®)		
c. Total Paid Distribution *(Sum of 15b (1), (2), (3), and (4))*		917	932
d. Free or Nominal Rate Distribution *(By Mail and Outside the Mail)*	(1) Free or Nominal Rate Outside-County Copies included on PS Form 3541		
	(2) Free or Nominal Rate In-County Copies Included on PS Form 3541		
	(3) Free or Nominal Rate Copies Mailed at Other Classes Through the USPS (e.g. First-Class Mail)		
	(4) Free or Nominal Rate Distribution Outside the Mail *(Carriers or other means)*	17	17
e. Total Free or Nominal Rate Distribution *(Sum of 15d (1), (2), (3) and (4))*		17	17
f. Total Distribution *(Sum of 15c and 15e)* ▶		934	949
g. Copies not Distributed *(See Instructions to Publishers #4 (page #3))* ▶		241	251
h. Total *(Sum of 15f and g)* ▶		1175	1200
i. Percent Paid *(15c divided by 15f times 100)* ▶		98%	98%

16. Publication of Statement of Ownership
☒ If the publication is a general publication, publication of this statement is required. Will be printed in the November 2011 issue of this publication.
☐ Publication not required.

17. Signature and Title of Editor, Publisher, Business Manager, or Owner
Barbara Spruill
Dir. Service Center Operations
Date: 10/11/11

I certify that all information furnished on this form is true and complete. I understand that anyone who furnishes false or misleading information on this form or who omits material or information requested on the form may be subject to criminal sanctions (including fines and imprisonment) and/or civil sanctions (including civil penalties).