Acquired equivalence and generalized suppression in a virtual reality environment

W. James Greville · Simon Dymond · Philip M. Newton ·
Bryan Roche

Abstract Acquired equivalence was investigated using a virtual reality conditioned suppression task administered in a first-person-shooter game. Two visual cues, A1 and B1, were followed by a tone (O1), and another two cues, A2 and B2, were followed by another tone (O2). During differential Pavlovian conditioning, A1 was paired with an instructed unconditioned stimulus (US) consisting of a flashing white screen, whereas A2 was not. All cues and outcomes were then presented at test, in the absence of the US, and suppression ratios were calculated for multiple response topographies (shots, hits, and breaks). Clear evidence of the suppression of shots was seen for A1 and B1, with no suppression being seen for either A2 or B2. Presentations of O1 and O2 resulted in significant suppression of shots and hits, whereas only O1 led to the suppression of breaks. The US expectancy ratings were consistent with these behavioral results. The findings are discussed in the light of differing accounts of acquired equivalence.

Keywords Acquired equivalence · Generalization · Conditioned suppression · Virtual reality environment

Acquired equivalence refers to the finding that stimuli sharing a common outcome are often shown to be equivalent, and that

W. J. Greville · S. Dymond (✉)
Department of Psychology, Swansea University, Singleton Park,
Swansea SA2 8PP, UK
e-mail: s.o.dymond@swansea.ac.uk

W. J. Greville · P. M. Newton
College of Medicine, Swansea University, Swansea, UK

B. Roche
Department of Psychology, National University of Ireland,
Maynooth, Ireland

changes applied to one stimulus may generalize to the other stimulus without further training (Hall, 1996; Honey & Hall, 1989). In a typical variant of common-outcome training, two stimuli predict different outcomes, which results in an increase in the discriminability of those stimuli; also, training in which two stimuli predict the same outcome results in a decrease in their discriminability. For instance, Ward-Robinson and Hall (1999) presented rats with two stimuli, A1 and A2, which were each followed by reward, whereas another, B, was not. Later, one stimulus (A1) was paired with shock, and it was found that fear generalized from A1 to A2, but not to B. Other studies with nonhumans (e.g., Bonardi, Rey, Richmond & Hall, 1993; Kaiser, Sherburne, Steirn & Zentall, 1997) adapted training designs from the early human discrimination literature to test acquired equivalence. Pigeons were first presented with two stimuli, A and B, which were followed by X, and a further two stimuli, C and D, which were followed by Y. Next, another discrimination was trained that was either consistent or inconsistent with the previous training. It was found that acquisition was faster when the reinforced stimuli were A and B and the nonreinforced stimuli were C and D, as compared to A and C and B and D, respectively, in the inconsistent condition. In this way, generalization was enhanced between stimuli that had shared a common event in the first stage of training (acquired equivalence).

Hall, Mitchell, Graham and Lavis (2003) found evidence for an associative-mediation account of acquired equivalence in experiments employing a multistage discrimination design. For example, initial training, in which no overt response was required, associated two shape stimuli (A and B) with the presentation of one outcome (the nonsense syllable *wug*), and another two shapes (C and D) with another (the nonsense syllable *zif*). In a second stage, groups of participants were trained in two different motor responses (press a key on the left or press a key on the right) in the presence of shapes that either shared a common outcome (consistent condition) or did not (inconsistent condition). Superior performance was observed in the consistent condition, indicating an acquired-

equivalence effect. According to the association-mediation account, Stage 1 training adds an associative link to the stimulus representations, such that subsequent retrieval of the associations during test facilitates generalized transfer.

Further studies, in which the influence of associative mediation was minimized, revealed evidence for an attentional-process account of acquired equivalence (Bonardi, Graham, Hall & Mitchell, 2005). The attentional-process, or feature salience (Meeter, Shohamy & Myers, 2009), account assumes that each of the four cues (A, B, C, and D) will share features with the others, such that A and B will have x in common, whereas C and D will share y in common. Training comes to establish feature x as being predictive of the outcome that follows A and B, whereas feature y is predictive of the common outcomes following C and D. Thus, discrimination is facilitated in the consistent training condition by paying increased attention to the predictive features and learning to make one response to both cues that share feature x, and another response to the cues that share feature y (see also Mackintosh, 1975). Bonardi et al. (2005) concluded that, with appropriately controlled designs, it is possible to demonstrate evidence for the attentional-process account, but that both attentional and associative processes likely operate in tandem in discrimination tasks of this kind.

Acquired equivalence may have relevance for understanding the generalization of conditioned emotional responses to cues indirectly related via common antecedents and/or outcomes (Hermans, Baeyens & Vervliet, 2013). For instance, a demonstration that fear-eliciting properties may generalize from a learned cue to related cues via acquired equivalence would mean that when the elements of this representational network are highly associated, activation of one element of the network (e.g., the perception of a dog or an odor) might be sufficient to activate the whole network and hence lead to the experience of fear or anxiety. From this perspective, generalization concerns the question of how stimuli that are related to the original CS are integrated into this associative network and thus acquire the potency to elicit the fear response. (Hermans et al., 2013, p. 127)

Generalization of conditioned fear such as this, as a consequence of acquired equivalence based on common outcomes, is a robust phenomenon in the nonhuman literature (e.g., Honey & Hall, 1989) but has not been well studied with humans. Although it is now accepted that the generalization of conditioned responses is a core component in the development, maintenance, and severity of psychopathology such as that seen in anxiety and other disorders (Dymond & Roche, 2009; Dymond, Roche & Bennett, 2013; Hermans et al., 2013; Lissek, Biggs, Rabin, Cornwell, Alvarez, Pine & Grillon, 2008), a demonstration of generalization via acquired equivalence involving a clinically relevant process such as conditioned suppression would be salutary.

In the present study, therefore, we employed a virtual reality environment "first-person-shooter" (FPS) videogame

(Greville, Newton, Roche & Dymond, 2013) to investigate the generalization of clinically relevant conditioned suppression via acquired equivalence based on common outcomes. Conditioned suppression is a model of the range and type of response disruption often seen in anxiety disorders, in which a fear-eliciting cue interferes with or otherwise suppresses ongoing instrumental behavior (Estes & Skinner, 1941). Although the human research that has been conducted to date on acquired equivalence has employed a range of innovative paradigms and designs (e.g., Bonardi et al., 2005; Hall et al., 2003; Hodder, George, Killcross & Honey, 2003; Meeter et al., 2009; Molet, Miller & Zentall, 2011; O'Reilly, Roche, Ruiz, Tyndall & Gavin, 2012; Preston, Shrager, Dudukovic & Gabrieli, 2004; Smyth, Barnes-Holmes & Barnes-Holmes, 2008), little is known about whether or not generalized suppression might be observed following a history of acquired equivalence based on common outcomes. Toward this end, virtual reality environments permit the fine-grained manipulation of different variables (such as common outcomes) and the measurement of multiple topographies of behavior in self-paced, motivating tasks that are well suited to the study of generalized suppression on the basis of acquired equivalence via common outcome training.

In the present study, using a within-subjects design, a three-stage training and testing procedure was arranged (see Table 1) in which presentations of two visual cues (A1 and B1) were followed by a common auditory stimulus (O1), whereas another two visual cues (A2 and B2) were followed by another auditory stimulus (O2). Next, during differential Pavlovian conditioning, A1 was followed by an instructed unconditioned stimulus (US), which was aversive in the context of the game, whereas A2 was not. During the crucial test phase, all cues and outcomes were presented, in the absence of the US, and suppression ratios were calculated for the multiple dependent measures afforded by the virtual reality task (Greville et al., 2013). If the common-outcome training was sufficient to create equivalent cues, then the conditioned suppressive effects associated with A1 should generalize to B1 (and O1), but not to B2 (and O2).

Table 1 Design of the experiment

Phase 1: Common-Outcome and Operant Training	
A1→O1	A2→O2
B1→O1	B2→O2
Phase 2: Pavlovian Conditioning	
A1+	A2−
Phase 3: Testing	
A1? B1? O1?	A2? B2? O2?

A1, A2, B1, and C2 represent visual cues, and O1 and O2 represent tones. "+", followed by the unconditioned stimulus (US); "−", not followed by the US; "?", test trials presented in the absence of the US.

Method

Participants

A group of 32 students and staff (17 women, 15 men) with a mean age of 27.7 years (SD = 6.84) were recruited from Swansea University and participated in return for either £5 or partial course credit.

Apparatus and stimuli

The experiment was conducted using a Dell Optiplex 755 PC running Microsoft Windows XP with a 27-in. Iiyama monitor as the visual display. A Logitech Rumblepad II wireless joy-pad controlled movement throughout the virtual environment (VE). Game sounds and the auditory stimuli used as two of the CSs were delivered via Grado SR60 headphones.

The VE was designed using the Hammer Editor, part of the Source SDK package from Valve Software, with code written using Microsoft Visual C++ Express Edition and the Steam client used to launch the experiment (Greville et al., 2013). The VE consisted primarily of two buildings, with outdoor areas on each side and a pathway linking the two buildings. Walls were used to keep areas separate, such that participants could only follow one route through the VE, beginning in the initial outside area, then progressing through the first building, along the linking pathway, through the second building, and then to the final outside area. Participants were prevented from returning to earlier stages by the automatic closing and locking of doors. The two buildings comprised six identical interconnecting rooms consisting of the same basic layout, with a shelf along the west wall containing rows of crates and an archway leading to the next room on the north wall (see the supplemental materials for examples of screen shots).

Each room within the buildings was illuminated by a centrally positioned light source. The four CSs consisted of this light source changing from ambient white light to one of four colors: blue, green, red, or yellow. This ensured that the CS was always visible, no matter where participants were within the building. The colors and outcomes were partially counterbalanced across participants. The two outcomes were 4-s auditory tones, generated using Audacity and delivered at approximately 70 dB via the headphones. One outcome was a pure tone (simple sine wave) of 600 Hz. The second tone was an oscillating tone produced by adding a "phaser" effect (Stages 4, LFO Frequency 3.5, LFO Start phase 40, Depth 200, Feedback 0) to the 600-Hz pure tone. Pilot testing indicated that the two tones were discriminable.

The US was 3 s in duration and consisted of the on-screen display shaking violently, as if the building had sustained a severe impact. In addition, participants lost a small quantity of gold that they had collected (100 points) on each occasion (see the description of procedures below for details of the game play). If, however, participants made a response during the US, the screen flashed an additional time, as if there had been an explosion, and participants lost a larger quantity of gold (450–550 points). This element was analogous to the "Martians" paradigm (Arcediano, Ortega & Matute, 1996; Franssen, Clarysse, Beckers, van Vooren & Baeyens, 2010), in which the effects of the US are amplified if an instrumental response is made during the US (see the supplemental materials for a copy of all task instructions).

Design

Our paradigm involved six distinct CSs—labeled A1, A2, B1, B2, O1, and O2—presented at different points in a three-phase design (Table 1). During the common-outcome training, the stimuli were paired together, such that both A1 and B1 were always followed by O1, and A2 and B2 were always followed by O2. In Pavlovian conditioning, only A1 and A2 were presented, with the US being contingent on A1 but not on A2. Finally, during testing, all of the stimuli were presented.

Procedure

Participants were tested individually in a small, darkened experimental room. In the VE, they first found themselves in an outdoor area, facing the first building, which served as an initial orientation zone. During this initial familiarization, participants were advised via on-screen instructions to familiarize themselves with the controls and to practice moving around prior to entering the building, at which point the experiment commenced.

Phase 1: Common outcome and operant training Phase 1 took place in the first building. Participants were instructed that their primary aim was to find gold bars worth 100 points per set hidden in wooden crates. The crates were arranged along one wall in each of six adjoining rooms (24 crates per wall), and gold bars were placed in a randomly determined six of the 24 crates on each wall. In total, 144 crates were presented in Phase 1, with 36 gold bars being dispersed among these crates; this was sufficient to occupy participants for the entirety of Phase 1. During the instrumental (i.e., operant) training, participants learned to shoot crates in order to try to find gold, and thus to score points. Each crate took four shots to destroy; thus, every four on-target shots resulted in crate destruction (i.e., fixed ratio [FR] 4 schedule), and (on average) every four crates destroyed yielded gold bars (i.e., variable ratio [VR] 4 schedule). If a gold bar was uncovered, 100 points were added to the participant's score. Whenever the participant's score changed (gain or loss), the running total was displayed for 2 s in the top left corner of the screen.

While participants were engaged in this task, common-outcome training took place, which consisted of 32 trials. A1, B1, A2, and B2 were presented eight times each and were always followed by the respective outcome tone; either O1 (for A1 or B1) or O2 (for A2 or B2). Each trial was 10 s in duration, and the intertrial interval (ITI) was 10 s. Presentation of either A1, B1, A2, or B2 occurred at a random point within the first 2 s of the trial, for a 4-s duration, and was immediately followed by the appropriate outcome, also for 4 s (temporal variability was introduced in order to alleviate any impression that the outcome might be occurring on a fixed-interval schedule, rather than being contingent on the cue). Presentation of A1, B1, A2, and B2 occurred in a pseudorandom order, with the constraint that no more than two presentations of each could occur consecutively. This same restriction was also applied to stimulus presentation during the two following phases.

This phase was intended simultaneously to train the operant response of shooting crates and to establish acquired equivalence between the stimuli and the outcomes; consistent with previous research on this topic (e.g., Hall et al., 2003), we did not conduct a separate, formal test for the acquired equivalence of the cues, other than that presented during the final test phase.

Phase 2: Pavlovian conditioning On completing Phase 1, participants left the first building and proceeded to the central pathway, where they were provided with on-screen instructions. Phase 2 commenced upon entering the second building, which had a layout identical to the first. A delay-conditioning procedure was used, with a 4-s delay between CS+ onset and US onset (i.e., CS duration = 4 s) and no trace interval separating CS+ and US. That is, if the US was scheduled, it was delivered immediately following termination of the CS+ (i.e., A1). Participants were exposed to ten presentations each of the CS+ and CS– in a pseudorandom order, with the US following the CS+ on eight of those ten trials (i.e., CS+/US contingency of .8). The US never followed the CS– (i.e., A2). Trials were again 10 s in duration, with a 10-s ITI and with the cue being presented at a random point within the first 2 s of the trial. Any loss of points due to the US was immediately displayed on screen.

Phase 3: Testing We tested for the generalization of suppression via acquired equivalence during this phase, which consisted of 12 trials: two presentations of each of the six CSs in a pseudorandom order, with no US presentations. The trial length was the same as we described previously. At the conclusion of Phase 3, participants were directed to exit the building and proceed to the final outdoor area, where they were informed that their mission was complete, presented with their total score (on screen), and asked to contact the experimenter.

The experimenter then reentered the experimental room and undertook a brief manipulation check to identify retrospectively the extent to which participants expected the US following each color ("To what extent did you expect an explosion to occur after the [colored] light? 1 = *not at all*, and 5 = *all the time*."). Participants were then debriefed and compensated.

Data analysis

Suppression ratios of the total number of responses (shots), shots hitting the targets (hits), and crates destroyed (breaks) during Phase 3 were calculated: $X/(X+Y)$, where X is the total number during the CS, and Y is the total number immediately (4 s) prior to the CS. All trials were included in the analysis; these data were not normally distributed according to a D'Agostino and Pearson (1973) omnibus normality test, and thus were compared using Wilcoxon signed-rank tests (Streiner & Norman, 2011).

Results

Significant suppression of shots was found when participants were presented with A1, B1, and O1, but not with A2, B2, and O2 (Fig. 1a). As expected, the greatest suppression was observed for A1, $W = -588.0$, $p < .0001$, which had been directly paired with the US in Phase 2. A2, which was never paired with the US, showed no significant suppression, $W = -143.0$, $p = .1246$. The B1 cue, which shared a common outcome (O1) with A1, showed significant suppression, $W = -377.0$, $p = .0013$, whereas the extent of suppression to B2, which shared a common outcome (O2) with A2, was not significant, $W = -133.0$, $p = .271$. This suggests that A1 and B1, and A2 and B2, respectively, were functionally equivalent by virtue of their common outcomes, and that the conditioning applied to one stimulus (A1) successfully generalized to B1, but not to B2. Significant suppression of shots was found for both of the outcomes: O1, $W = -281.0$, $p = .0040$, and O2, $W = -291.0$, $p = .0029$. Similar effects were found for hits (Fig. 1b) and breaks, although no significant suppression of breaks was observed for O2 (Fig. 1c).

Overall, suppression was significantly greater during A1 than during A2 for all measures (all ps < .001). Likewise, suppression was significantly greater during B1 than during B2 for breaks ($p = .0195$), and marginally significant for hits ($W = -238$, $p = .0523$) and shots ($W = -195$, $p = .0831$). We observed no significant difference between O1 and O2 for shots ($W = -40$, $p = .75$), hits ($W = -102$, $p = .35$) or breaks ($W = -134$, $p = .09$).

Participants' postexperimental US expectancy ratings are shown in Fig. 1d. As predicted, the highest level of expectancy was for A1; with a mean rating of 4.417, this was
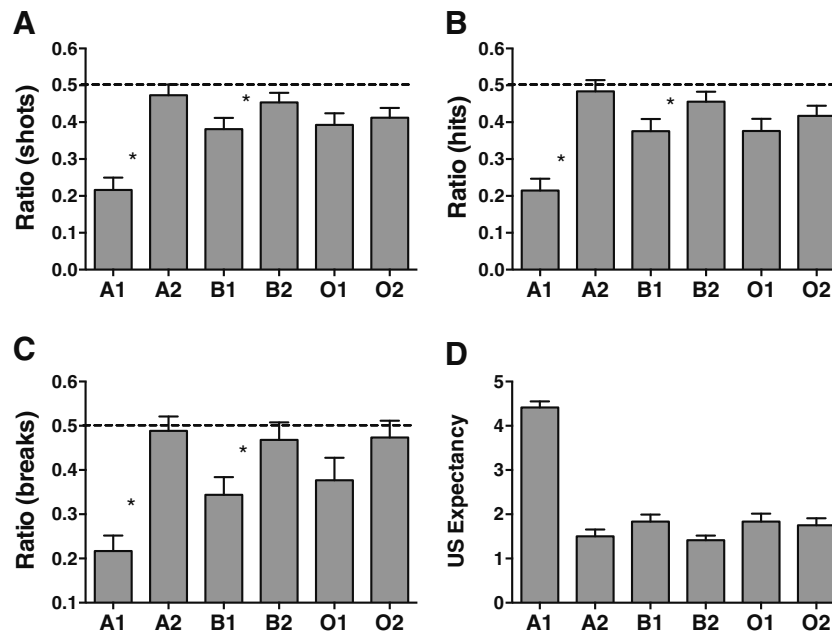
Fig. 1 Suppression ratios for shots (**a**), hits (**b**), and breaks (**c**), as well as the mean unconditioned-stimulus (US) expectancy ratings (**d**), for all cues and outcomes presented during the test phase. Error bars indicate standard errors of the means. *Significantly different from the paired stimulus (i.e., A1 vs. A2, B1 vs. B2)

significantly higher than the expectancy rating of 1.5 for A2, $W = -586$, $p < .001$. The ratings made to B1 (1.833) and B2 (1.417) differed ($W = -75, p < .01$), whereas those made to O1 (1.833) and O2 (1.75) did not differ significantly ($W = -12$, $p = .6543$). Finally, the ratings of O1 and O2 were not correlated with any of the behavioral suppression measures (O1: p values between .311 and .930; O2: p values between .113 and .951).

Discussion

The present study was designed to test the predictions of an account based on acquired equivalence: that the suppressive effects of Pavlovian training would generalize to stimuli sharing a common outcome, and not to other stimuli sharing another outcome. Using a novel virtual reality conditioned suppression task (Greville et al., 2013), participants underwent a three-stage training and testing procedure in which two cues (A1 and B1) were first followed by a common outcome (O1), whereas two other cues (A2 and B2) were followed by a different outcome (O2). During Pavlovian training, one cue (A1) was followed by a flashing white screen US, which functioned as an aversive stimulus, whereas A2 was not, and in the final test phase, all cues and outcomes were presented in the absence of the US. To the extent that the first stage of training was sufficient to create equivalent cues, the suppressive effects of A1 should generalize to B1 and not to B2 (or A2). Both outcomes should increase discrimination

between the cues, such that greater suppression should also be demonstrated for O1 over O2. These findings revealed clear evidence of suppression of shots for A1, which was directly paired with the US, and for B1, which shared common outcome O1. No suppression was observed for either A2, which was never directly paired with the US, or with B2, with which A2 shared the common outcome O2. Presentations of O1 and O2 resulted in significant suppression of shots and hits, whereas only O1 led to suppression of breaks. Overall, we found significant suppression of A1 over A2 for all measures obtained from the virtual reality task, with similar suppressive effects being noted for B1 over B2 for breaks and, at marginally significant levels, for hits and shots. Although O1 and O2 did not differ on the three behavioral measures (shots, hits, and breaks), the findings from the US expectancy ratings for these outcome stimuli were consistent with our behavioral results for the A and B cues (see Fig. 1). Unlike the A and B cues, the difference in suppression between O1 and O2 was not significant, suggesting that participants did not learn to discriminate between the outcomes very effectively.

In the introduction, we also considered the predictions made by an attentional-process account of acquired equivalence (e.g., Bonardi et al., 2005). Applied to the present findings, an attentional account would presuppose that the A1 and B1 cues shared feature x in common, and A2 and B2 shared feature y in common. Feature x comes to predict the common outcome (O1) shared by A1 and B1, whereas feature y is predictive of the common outcome (O2) shared by A2 and

B2. Following Pavlovian training, changes applied to A1 and A2 are also assumed to share the individual features x and y in common, respectively, which is likely to further facilitate discrimination between the cues and enhance generalization of conditioned suppression from A1 to B1 (and O1), and of nonsuppression from A2 to B2 (and O2).

To some extent, it is difficult to reconcile the present design and findings with the attentional account, for several reasons. First, the predictive features x and y remain unspecified. Although it is possible that pairs of stimuli could share populations of stimulus elements in common, it is at present difficult to determine what role they might have played in capturing and retaining attentional resources. Second, all of the cues were colored lights that did not differ along any other physical dimensions—more precisely, they had salient features equally in common. Cues of the same illumination (as far as can be assumed, in the absence of individual psychophysics), size, and duration were presented in similar, equally predictable contexts within the game, and thus it is difficult to determine how the potential salience of the cues might have been unintentionally varied within the present procedures (cf. Bonardi et al., 2005). Third, the attentional account may struggle to explain the mediated-conditioning findings (i.e., O1 > O2). It does, however, allow for the possibility that mediated conditioning processes may operate in tandem with attentional processes, which would reconcile the account with the present findings. Finally, we employed a within-subjects design that did not require manipulating the consistency of any subsequent training involving cues that did or did not share a common outcome (as has been done in the majority of research conducted to date on acquired equivalence). Of course, it is readily possible to propose a variant of the present design in which groups would receive an additional training stage after common-outcome training in which a particular response (e.g., pressing a marked key on the left) would be taught for A1 and B1, and another response (e.g., pressing a marked key on the right) would be taught for A2 and B2 (consistent condition), whereas in the inconsistent condition, one response might be required for A1 and B2, and another for A2 and B1. Given this arrangement, we would predict findings similar to those of the present study, although the attentional account would still need to address exactly which salient features came to be predictive, given the aforementioned issues. In conclusion, although it is beyond the remit of the present article to evaluate the relative merits of the associative-mediation and attentional accounts, it is likely that either a combination of associative and attentional processes or additional (e.g., configural) processes underlie acquired-equivalence effects like those seen here (for alternative accounts, see Honey, Close & Lin, 2010; Honey & Ward-Robinson, 2002).

An associative-mediation account (Hall et al., 2003) of the present findings would assume that initial common-outcome training adds an associative link, such that presentations of A1 and B1 evoke representations of each other and of O1; that A2 and B2 evoke representations of each other and of O2; and that test performance is facilitated by retrieval of these associatively linked representations. This account may partly explain our findings, with the exception that suppression was observed to both O1 and O2, and not just O1. One potential factor in explaining the present findings may be the within-group design employed, which involved a multistage training and test sequence rather than a consistent/inconsistent between-group manipulation during the second training stage, when responses were trained in the presence of stimuli that did or did not share a common outcome (Hall et al., 2003). However, a number of other possible procedural explanations could account for the unexpected suppression to O2.

First, the outcomes were presented in different modalities than the A and B cues (i.e., auditory tones rather than colored lights). Although pilot testing indicated that these tones were discriminable, they may have been less so than the light cues, and hence, generalization from one to the other may have occurred in ways that would be difficult to predict and control when the generalization gradient was unknown. Second, the discriminability of the outcome stimuli is also questionable, given the fact that participants were not required to attend to outcomes, as they were for the A stimuli. In other words, the outcome stimuli were not themselves discriminative for outcomes, whereas the A and B stimuli were. As a result, we would not expect the O stimuli to have distinct and clearly acquired stimulus functions. Consequently, any generalized transfer from the A to the O stimuli, and from one O stimulus to the other, would be difficult to predict. Controlling such an outcome was not an expressed purpose of this study, although it is now interesting to note that when they were probed for, the outcome stimuli did not appear to acquire the same functions as the discriminated conditioned stimuli using the present preparation. Third, it is not necessarily a valid assumption that the outcome stimuli should acquire the same functions as the A and B cues, since O1 and O2 always *followed* presentations of the A and B cues, rather than predicting them. It would be unreasonable to expect in practice for the outcomes to acquire the same stimulus properties as the A and B stimuli via backward-conditioning effects (see Hall, 1996). For these reasons, the purpose of checking for emergent outcome stimulus functions was largely inductive.

Taking all of these arguments together, it may have been the case that, having no expectation either way about the functions of O1 and O2, that participants erred on the side of caution with regard to O1 and O2, in reducing response rates in the presence of both. Although this explanation is certainly plausible, it raises yet another query: If nothing was learned about O1 and O2, and participants merely erred on the side of caution, then given that similar levels of suppression were exhibited for B1 as for O1 and O2, can we be confident that

participants learned that A1 and B1 were equivalent via common-outcome training, or were they simply erring on the side of caution with regard to B1 also? We would suggest that because we found significant suppression to B1 but not B2, and because the differences between these cues were significant for all measures, we can be satisfied that acquired equivalence was demonstrated. Meanwhile, with regard to the outcomes, we can tentatively conclude that a degree of learning took place, but that this was impaired by the different modalities of the stimuli, perhaps by their poor discriminability, and by the fact that they were outcomes rather than cues. If nothing else, even if we were to assume that any novel stimulus might have shown the same suppressive effects as B1, O1, and O2, we can conclude that participants learned that the A2 stimulus was "safe" and that this knowledge successfully transferred to B2, which likewise did not show significant suppression.

A potential limitation of the present findings concerns the fact that participants' US expectancy ratings were obtained after a period of extinction. Although the expectancy ratings data were generally consistent with the observed behavioral effects, it would be useful to conduct a further study in which concurrent ratings were obtained within the conditioning task. This would allow for the simultaneous tracking of expectancies and behavior across the phases of the task.

The conceptual paradigm within which this work was conducted may offer a further possible explanatory mechanism for the effects observed. More specifically, the present research was conducted within a behavior-analytic paradigm in which explanations for the types of experimental effects reported here might be proffered in terms of stimulus control alone (Dougher & Markham, 1996; Saunders & Williams, 1998). That is, from our functional perspective, our experimental preparations alone constitute a form of explanation for the present effects, insofar as they allowed for both prediction of and influence over the experimental outcomes (Chiesa, 1994; Hayes & Brownstein, 1986; Skinner, 1953). Thus, whereas attentional and associative-mediation accounts of the observed effects may serve as useful explanatory heuristics, they may also burden the behavioral researcher with an obligation to devote undue research attention to the effort to ratify one or the other account, rather than focus resources on increasing the predictive and controlling features of the relevant experimental preparations. Adopting a functional approach to the analysis of behavior may ultimately provide the more complete and coherent account of any laboratory-created effects. Indeed, the functional, pragmatic, and parsimonious approach is increasingly viewed as one deserving of consideration by cognitive researchers (De Houwer, 2011; De Houwer, Barnes-Holmes & Moors, 2013). In conclusion, we believe that the present findings extend our understanding of acquired equivalence and offer an innovative methodology for studying both such effects and conditioned suppression, regardless of a researcher's theoretical perspective.

References

Arcediano, F., Ortega, N., & Matute, H. (1996). A behavioural preparation for the study of human Pavlovian conditioning. *Quarterly Journal of Experimental Psychology, 49B,* 270–283. doi:10.1080/713932633

Bonardi, C., Graham, S., Hall, G., & Mitchell, C. (2005). Acquired distinctiveness and equivalence in human discrimination learning: Evidence for an attentional process. *Psychonomic Bulletin & Review, 12,* 88–92. doi:10.3758/BF03196351

Bonardi, C., Rey, V., Richmond, M., & Hall, G. (1993). Acquired equivalence of cues in pigeon autoshaping: Effects of training with common consequences and with common antecedents. *Animal Learning & Behavior, 21,* 369–376. doi:10.3758/BF03198003

Chiesa, M. (1994). *Radical behaviorism: The philosophy and the science.* Boston: Authors' Cooperative.

D'Agostino, R. B., & Pearson, E. S. (1973). Tests of departure from normality. Empirical results for the distribution of b^2 and $\sqrt{b^1}$. *Biometrika, 60,* 613–622.

De Houwer, J. (2011). Why the cognitive approach in psychology would profit from a functional approach and vice versa. *Perspectives on Psychological Science, 6,* 202–209.

De Houwer, J., Barnes-Holmes, D., & Moors, A. (2013). What is learning? On the nature and merits of a functional definition of learning. *Psychonomic Bulletin & Review, 20,* 631–642. doi:10.3758/s13423-013-0386-3

Dougher, M. J., & Markham, M. R. (1996). Stimulus classes and the untrained acquisition of stimulus functions. In T. R. Zentall & P. M. Smeets (Eds.), *Stimulus class formation in humans and animals* (pp. 137–152). Amsterdam: Elsevier.

Dymond, S., & Roche, B. (2009). A contemporary behavior analysis of anxiety and avoidance. *Behavior Analyst, 32,* 7–28.

Dymond, S., Roche, B., & Bennett, M. (2013). Relational frame theory and experimental psychopathology. In S. Dymond & B. Roche (Eds.), *Advances in relational frame theory: Research and application* (pp. 199–218). Oakland: New Harbinger.

Estes, W. K., & Skinner, B. F. (1941). Some quantitative properties of anxiety. *Journal of Experimental Psychology, 29,* 390–400. doi:10.1037/h0062283

Franssen, M., Clarysse, J., Beckers, T., van Vooren, P. R., & 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

Greville, W. J. S., Newton, P. M., Roche, B., & Dymond, S. (2013). Conditioned suppression in a virtual environment. *Computers in Human Behavior, 29,* 552–558. doi:10.1016/j.chb.2012.11.016

Hall, G. (1996). Learning about associatively activated stimulus representations: Implications for acquired equivalence and perceptual learning. *Animal Learning & Behavior, 24,* 233–255. doi:10.3758/BF03198973

Hall, G., Mitchell, C., Graham, S., & Lavis, Y. (2003). Acquired equivalence and distinctiveness in human discrimination learning: Evidence for associative mediation. *Journal of Experimental Psychology: General, 132,* 266–276. doi:10.1037/0096-3445.132.2.266

Hayes, S. C., & Brownstein, A. J. (1986). Mentalism, behavior-behavior relations, and a behavior-analytic view of the purposes of science. *Behavior Analyst, 9,* 175–190.

Hermans, D., Baeyens, F., & Vervliet, B. (2013). Generalization of acquired emotional responses. In M. D. Robinson, E. Watkins, &

E. Harmon-Jones (Eds.), *Handbook of cognition and emotion* (pp. 117–134). London: Guilford Press.

Hodder, K. I., George, D. N., Killcross, A. S., & Honey, R. C. (2003). Representational blending in human conditional learning: Implications for associative theory. *Quarterly Journal of Experimental Psychology, 56B,* 223–238.

Honey, R. C., Close, J., & Lin, T. C. E. (2010). Acquired distinctiveness and equivalence: A synthesis. In C. Mitchell & M. E. Le Pelley (Eds.), *Attention and associative learning: From brain to behavior* (pp. 159–186). Oxford: Oxford University Press.

Honey, R. C., & Hall, G. (1989). Acquired equivalence and distinctiveness of cues. *Journal of Experimental Psychology: Animal Behavior Processes, 16,* 178–184.

Honey, R. C., & Ward-Robinson, J. (2002). Acquired equivalence and distinctiveness of cues: I. Exploring a neural network approach. *Journal of Experimental Psychology: Animal Behavior Processes, 28,* 378–387.

Kaiser, D. H., Sherburne, L. M., Steirn, J. N., & Zentall, T. R. (1997). Perceptual learning in pigeons: Decreased ability to discriminate samples mapped onto the same comparison in many-to-one matching. *Psychonomic Bulletin & Review, 4,* 378–381. doi:10.3758/BF03210796

Lissek, S., Biggs, A. L., Rabin, S., Cornwell, B. R., Alvarez, R. P., Pine, D. S., & Grillon, C. (2008). Generalization of conditioned fear-potentiated startle in humans: Experimental validation and clinical relevance. *Behaviour Research and Therapy, 46,* 678–687.

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

Meeter, M., Shohamy, D., & Myers, C. E. (2009). Acquired equivalence changes stimulus representation. *Journal of the Experimental Analysis of Behavior, 91,* 127–141.

Molet, M., Miller, H. C., & Zentall, T. R. (2011). Acquired equivalence between stimuli trained in the same context. *Psychonomic Bulletin & Review, 18,* 618–623. doi:10.3758/s13423-011-0087-8

O'Reilly, A., Roche, B., Ruiz, M. R., Tyndall, I., & Gavin, A. (2012). The Function Acquisition Speed Test (FAST): A behavior-analytic implicit test for assessing stimulus relations. *Psychological Record, 62,* 507–528.

Preston, A. R., Shrager, Y., Dudukovic, N. M., & Gabrieli, J. D. E. (2004). Hippocampal contribution to the novel use of relational information in declarative memory. *Hippocampus, 14,* 148–152.

Saunders, K. J., & Williams, D. C. (1998). Stimulus-control procedures. In K. A. Lattal & M. Perone (Eds.), *Handbook of research methods in human operant behavior* (pp. 193–228). New York: Plenum Press.

Skinner, B. F. (1953). *Science and human behavior*. New York: Macmillan.

Smyth, S., Barnes-Holmes, D., & Barnes-Holmes, Y. (2008). Acquired equivalence in human discrimination learning: The role of propositional knowledge. *Journal of Experimental Psychology: Animal Behavior Processes, 34,* 167–177. doi:10.1037/0097-7403.34.1.167

Streiner, D. L., & Norman, G. R. (2011). Correction for multiple testing: Is there a resolution? *Chest, 140,* 16–18.

Ward-Robinson, J., & Hall, G. (1999). The role of mediated conditioning in acquired equivalence. *Quarterly Journal of Experimental Psychology, 52B,* 335–350.