Accepted for publication in Journal of Experimental Psychology: General. © 2017, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors permission. The final article will be available, upon publication, via its DOI: 10.1037/xge0000341

Contemporary associative learning theory predicts failures to obtain blocking. Comment on Maes et al. (2016)

Fabian A. Soto
Department of Psychology, Florida International University

In a recent article, Maes et al. (2016) report the results from fifteen experiments in the blocking effect, all of which failed to replicate the basic phenomenon. While Maes et al. did not dispute the reality of the blocking effect, they concluded that the effect is more difficult to obtain than what could be assumed from the literature and that we lack insight into its boundary conditions. This conclusion is incorrect, as contemporary associative learning theory both agrees with the authors' conclusion that blocking is parameter-dependent, and it makes specific predictions about the experimental parameters likely to produce a small or no blocking effect. Ten out of the fifteen experiments presented by Maes et al. use exactly those parameters (same-modality stimuli for the compound AX), making their results completely unsurprising in the light of contemporary associative learning theory. The results from three other experiments are difficult to interpret due to a floor effect. A failure to replicate blocking in only two experiments is unsurprising and can be explained as the result of statistical variability or changes in experimental procedure.

Keywords: blocking, associative learning theory, replicability, configural and elemental processing

The discovery of stimulus competition effects in Pavlovian conditioning, back in the 1960s, completely changed our view of associative learning. A series of experiments involving the presentation of multiple conditioned stimuli (CSs) in a compound, followed by an unconditioned stimulus (US), revealed that CSs in the compound that are better predictors of the US acquire a stronger association with it (e.g., Kamin, 1969; Rescorla, 1968; Wagner et al., 1968). The best-known stimulus competition effect is Kamin's blocking effect (Kamin, 1969), in which pairings of a compound AX with a US result in low conditioning to X when A has already been paired with the US. The usual interpretation is that because A is a good predictor of the US, new learning about the uninformative stimulus X is unnecessary. This line of research culminated with the proposal of the Rescorla-Wagner model (Rescorla & Wagner, 1972) and similar models (e.g., Mackintosh, 1975; Pearce & Hall, 1980), which assume that learning about one cue is influenced by other cues concurrently presented during conditioning. These models are still hugely influential today.

Maes et al. (2016) have reported the results from fifteen experiments in the blocking effect, all of which failed to replicate the basic phenomenon. Maes et al. conclude that their results "raise doubts regarding the canonical nature of the blocking effect" and suggest that "blocking is a highly parameter-dependent phenomenon" (p. e58). While the authors "do not want to dispute that the blocking effect exists", they conclude that "a true blocking effect is more difficult to obtain than one might assume from the literature and that we lack insight into its boundary conditions" (p. e60). Due to the importance of the blocking effect for the development of error-driven associative learning models, these conclusions could be read as suggesting that such models are of questionable value, as they seem unable to explain under what conditions blocking is indeed observed. In this commentary, I will argue that contemporary associative learning theory does not really give blocking a "canonical" status. Rather, it agrees with the authors' conclusion that blocking is parameter-dependent, and it makes very specific predictions about the experimental parameters likely to produce a small or no blocking effect. As it turns out, most of the experiments presented by Maes et al. use exactly those parameters. Overall, the results described by Maes et al. are rather unsurprising, and for the most part can be explained in the light of contemporary associative learning theory.

Maes et al. recognize that models of associative learning propose very specific boundary conditions for the observation of a blocking effect. Although they seem to agree with the importance of checking that such boundary conditions have been met in their experiments, and they discuss some boundary conditions in their paper, their discussion is limited to models that are more than 30 years old: the models of Rescorla and Wagner (1972), Mackintosh (1975), and Pearce and Hall (1980). Such models were designed to explain the blocking effect and other stimulus competition phenomena, rather than to explain the specific circumstances in which blocking should and should not occur. Unsurprisingly, the models cannot explain the failures to obtain blocking reported by Maes et al.

However, we have learned much more about associative learning since 1980, and that knowledge has been incorporated into new models. Contemporary learning theory has been expanded beyond the Rescorla-Wagner and similar models, and it is important to understand how some of the results obtained by Maes et al. (2016) can be interpreted in the light of such work.

One area of considerable work is related to the representation of stimuli in compounds. The original Rescorla-Wagner model assumed that stimuli presented in a compound (e.g., light and tone) are independently represented and associated with the US. This assumption has come to be known as "elemental" stimulus processing. On the basis of several experimental phenomena that could not be explained by traditional elemental models, Pearce (1987; 1994; 2002) proposed a model in which stimuli in compounds are represented as whole configurations, and it is this configural representation that gets associated with the unconditioned stimulus. The important difference between elemental and configural models is not really the kind of representation proposed (see Ghirlanda, 2015), but how much generalization they assume exists between different compound stimuli. For example, the Rescorla-Wagner model assumes that the associative strength acquired by A is fully transferred to the compound AX in a blocking experiment. For this reason, the US is perfectly predicted during the second phase of blocking and there is no new learning. On the other hand, Pearce's model assumes that the associative strength acquired by A is only partially transferred to the compound AX in a blocking experiment. In this case, the US is only imperfectly predicted and AX acquires some associative strength, which is then partially generalized back to the stimulus X. While the model predicts a blocking effect, the size of this blocking effect is smaller than in the case of the Rescorla-Wagner model. Furthermore, nothing prevents us from assuming even less generalization from A to AX in the blocking design, leading to a smaller blocking effect. In the extreme of full configural processing, learning about A does not transfer at all to the compound AX, and learning about AX does not transfer to X. With such low levels of compound generalization, no stimulus competition effects are observed. The take-home message is the following: more "elemental" stimulus processing produces a stronger blocking effect, while more "configural" stimulus processing produces a weaker blocking effect.

Neither traditional elemental models nor traditional configural models can explain the full range of experimental data (for reviews, see Melchers et al., 2008; Wagner, 2007). For this reason, a number of new models were proposed in the past fifteen years that allow for some level of flexibility in stimulus processing (Harris, 2006; Harris & Livesey, 2010; Kinder & Lachnit, 2003; McLaren & Mackintosh, 2002; Soto et al., 2014, 2015; Thorwart et al., 2012; Wagner, 2003, 2007). In addition, it has become clearer what experimental factors might produce more elemental processing and what factors produce more configural processing (see Melchers et al., 2008). For example, Kehoe et al. (1994) showed that the summation effect, which is indicative of elemental processing, is observed with stimuli from different modalities, but not with stimuli from the same modality. Many models (Harris, 2006; Harris & Livesey, 2010; McLaren & Mackintosh, 2002; Soto et al., 2014, 2015; Thorwart et al., 2012; Wagner, 2003, 2007) have explained this result by assuming that more similar stimuli, such as those coming from the same modality, produce more configural processing[1]. The logical consequence for blocking experiments is the following: contemporary associative learning theory predicts that more similar stimuli, and particularly those coming from the same modality, should produce a weaker blocking effect.

Figure 1 shows the predictions of Wagner's (2003; 2007) replaced elements model, a flexible version of Pearce's model proposed by Kinder and Lachnit (2003), and the latent causes model of Soto et al. (2014; 2015), for blocking experiments involving stimuli with different levels of similarity. These simulations have been adapted from Soto et al. (2014), so the interested reader can go to that original paper for a more detailed description of the models, the chosen parameter values and the simulation procedures. As it can be seen from the figure, all three models predict a stronger blocking effect with more dissimilar CSs, such as those coming from different modalities, and a weaker blocking effect with more similar CSs, such as those coming from the same modality. A point that is important to note is the following: the simulations shown in Figure 1 have been published in the literature (they are adapted from Figures 11 and 12 in Soto et al., 2014), and two of these models were proposed more than 10 years ago. This is not post-hoc storytelling. Also, both Wagner (2003) and Soto et al. (2014) explicitly associated "high similarity" leading to configural processing with stimuli from the same modality, and "low similarity" leading to elemental processing with stimuli from different modalities.

The exact mechanisms by which these models predict a smaller blocking effect with more similar stimuli differ across models. Wagner's replaced elements model

---

[1] Here it is important not to confuse *compound similarity,* which is similarity between compound stimuli as a function of their shared components (e.g., similarity between A and AX, AX and AY, AX and ABX, and so on), and *component similarity*, which is the similarity between the discrete stimulus components (e.g., the similarity between A and X). Different models of associative learning assume different levels of generalization as a function of *compound similarity*, with traditional configural models (e.g., Pearce's) assuming less generalization than traditional elemental models (e.g., Rescorla-Wagner). In addition, more recent models also assume that these generalization rules change as a function of *component similarity*, with lower levels of generalization when components are more similar. The specific mechanisms by which these models implement this hypothesis cannot be covered here, due to space limitations. For more details, interested readers should go to the papers by Wagner (2003; 2007), McLaren and Mackintosh (2002), Harris and colleagues (Harris & Livesey, 2010; Thorwart et al., 2012), and Soto and colleagues (Soto et al., 2014, 2015).

Figure 1. Predictions of conditioned responding to X in a blocking experiment from three models of associative learning. A larger difference between the black and white columns represents a stronger blocking effect. All three models predict that more similar stimuli, such as those coming from the same modality, should produce a weaker or no blocking effect than dissimilar stimuli, such as those coming from different modalities.

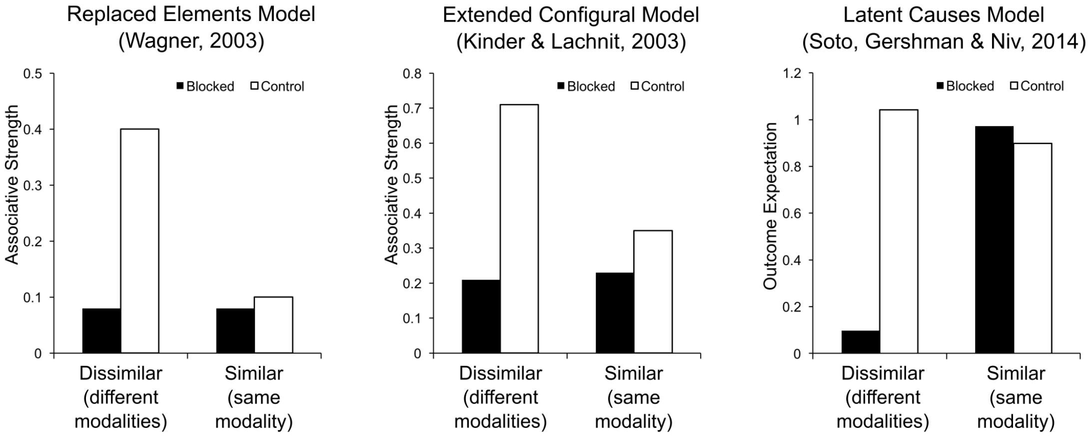

and the extended configural model propose that, with similar stimuli that foster configural processing, there is less generalization from A to the compound AX, and from AX to the test stimulus X. As indicated earlier, this should produce a smaller difference in responding to X between the blocked and control conditions. On the other hand, the latent causes model proposes that, with similar stimuli that foster configural processing, the same "configural" representation (a latent cause) is activated by A, X or AX. This configural representation is associated with the US during A+ and AX+ trials, and its associative strength is fully generalized to trials in which X is presented alone. Because A, X and AX share the exact same representation, both conditions of a blocking design (blocked and control) can be thought of as associating a single configural cue with the US and then testing its conditioned response. As seen in Figure 1, this leads to a prediction of high responding in both conditions.

Only extreme configural processing would lead to the complete absence of a blocking effect in the simulations presented in Figure 1. One might argue that such extreme configural processing is unlikely with any stimuli. Still, a reduction in the size of the blocking effect due to configural stimulus processing will necessarily reduce our ability to detect the effect in noisy data. With very similar stimuli, the effect might be too small to detect with the relatively small sample sizes used by Maes et al. (2016). In line with this idea, the results from several experiments (5, 7, 8, 11, 12 and 13) were (using the authors' words) numerically in line with the blocking effect, but without reaching statistical significance, and previous studies using stimuli similar to those used by Maes et al. (see original article for a discussion and references) and other same-modality stimuli (e.g., Dwyer et al., 2011; Jones & Haselgrove, 2013) did find evidence of a blocking effect. Thus, it is likely that the blocking effect is reduced rather than completely eliminated by using stimuli from the same modality, but the prediction from the literature is clear: such stimuli foster configural processing and reduce the likelihood of observing blocking.

Several other models in the literature are likely to make the same prediction as those shown in Figure 1: Harris (2006), Harris & Livesey (2010), and McLaren & Mackintosh (2002). To the best of my knowledge, specific predictions about the effect of stimulus similarity on the blocking effect have not been published for those models, but all of them propose less generalization from X to AX and from AX to A, and therefore a smaller blocking effect, when the two stimuli are similar (i.e., from the same modality, rather than from different modalities).

Here I have focused on an explanation of the results reported by Maes et al. (2016) in terms of elemental vs. configural stimulus processing, because in this case quantitative models of associative learning offer clear predictions (see Figure 1). Still, it should be mentioned that other potential explanations for the failure to obtain blocking with stimuli from the same modality exist in the associative learning literature. In particular, stronger generalization between A and X should produce a smaller blocking effect, because the associative strength acquired by A is generalized to X. In addition, generalization from A to B in the control condition would make the detection of a blocking effect even harder. Such generalization is more likely with stimuli from the same modality, which can be highly similar, than with stimuli from different modalities, which are are dissimilar (note that in Experiments 5-14 A, X and B are all from the same modality). Furthermore, generalization between A and X can be supported by within-compound associations formed during AX+ training (Rescorla &

Durlach, 1981), and such associations are thought to be easier to develop between similar stimuli.

In sum, contemporary learning theories predict that using stimuli from the same modality in the compound AX reduces the likelihood of observing the blocking effect[2]. Ten out of the 15 experiments reported by Maes et al. (2016) used stimuli from the same modality for AX, known to foster configural processing of compounds (Experiments 5 to 14; see Appendix F in original article). That is, in the light of contemporary associative learning theory, most of the failures to obtain blocking obtained in this paper are unsurprising. Although the results from these experiments are compelling, they are not beyond the scope of our current understanding of associative learning.

What about the other 5 experiments?

Five of the experiments reported by Maes et al. (experiments 1-4 and 15) used stimuli from different modalities in the compound AX, which are thought to foster elemental rather than configural processing. The models discussed earlier would predict that a robust blocking effect should have been observed in all these experiments. On the other hand, problems with some of these experiments make their results extremely difficult to interpret.

In particular, as the authors indicate in page e57: "One might argue that the observation of blocking in Experiments 2, 3 and 15 was hampered by a floor effect–if the control group is hardly responding to X, lower responding in the experimental group cannot be expected." Thus, three of the five experiments in which we might expect a blocking effect have results indicative of a floor effect, a point admitted by the authors in their paper. The authors also perform a Bayesian meta-analysis of all other experiments, and conclude that "when excluding the potential influence of floor effects, we find substantial evidence in favor of the null hypothesis" (p. e57). This analysis and conclusion are quite convincing but, as indicated by the previous discussion, also completely expected by contemporary associative learning theory, which predicts a small or no blocking effect in 10 out of the 12 experiments included in this meta-analysis.

In conclusion, only two experiments (performed in mice) out of fifteen have results that might be considered unexpected from the point of view of contemporary associative learning theory. The failure to obtain blocking in two experiment is far less surprising than the failure to obtain blocking in fifteen experiments, and by itself it would be unlikely to generate the kind of attention that this article has generated in the field. There are many possible explanations for the results of these two experiments, but I would like to highlight that their post-hoc nature makes them different from the predictions from associative learning theory discussed earlier.

The simplest explanation is that statistically we expect some proportion of well-conducted blocking experiments to not produce a blocking effect, as the sample effect size in blocking experiments is a random variable. Given the large number of successful replications of blocking in the literature (as reviewed by Maes et al, 2016), including successful demonstrations of the phenomenon in mice (Bonardi et al., 2010; Sanderson et al., 2016; Yamada, 2010), two failed replications are, again, rather unsurprising. Still, the overall pattern of results observed in the mice experiments reported by Maes et al. (2016; Experiments 1-4), which can be observed in Figure 2, suggests an even simpler explanation: tones may have been much more salient events than lights for the mice in these experiments. This hypothesis would explain the floor effect observed in Experiments 2 and 3, in which X was a light that acquired no conditioned response even in the control condition, probably due to strong overshadowing by the more salient tone. In Experiments 1 and 4, this salient tone was the blocked stimulus X, which may have captured attention during compound trials, producing a strong external inhibition effect and thus a weakened generalization of the associative strength acquired by A to the compound AX. In line with this hypothesis, Hall et al. (1977) found that blocking is reduced when X is a more salient stimulus than A, and more specifically LoLordo et al. (1982) found that tones are difficult to block by lights in fear conditioning (as in Experiments 1-4) and concluded that this was due to the tones' higher salience.

Results from previous blocking experiments in mice also offer a possible explanation for the lack of blocking observed in Experiments 1 and 4. Sanderson et al. (2016) found that a visual cue could produce blocking of an auditory cue only if the visual cue was over-trained during the first phase of the experiment. Simply pairing the visual cue with the unconditioned stimulus until asymptotic responding was observed was not enough to observe a blocking effect. That is, some boundary conditions for blocking in mice have been empirically identified in the literature, and Experiments 1 and 4 were not designed to meet such boundary conditions (they were performed before the Sanderson et al. paper was published).

---

[2] To the best of my knowledge this prediction has not been directly tested, although there is much evidence in line with the assumption that stimuli from the same modality produce more "configural" processing, while stimuli from different modalities produce "elemental" processing (for reviews, see Melchers et al., 2008; Wagner, 2003, 2007). Importantly, while Maes et al. indicate that "theoretical accounts for blocking may offer clues regarding potential boundary conditions" (p. e58), their conclusion is not that the boundary conditions for blocking proposed by associative learning theory have not been directly tested. Instead, their conclusion is that "the blocking effect is indeed dependent on (a variety of) boundary conditions, the exact nature of which is yet to be determined" (p. e59, emphasis added) and that "a true blocking effect is more difficult to obtain than one might assume from the literature and that we lack insight into its boundary conditions" (p. e60, emphasis added).

The bigger picture

In their final conclusion, Maes et al. say that "blocking, rather than being a touchstone for our theories of elementary learning, should be the subject of further investigation" (p. e60). It seems important to clarify that, regardless of what introductory textbooks might say, it was not blocking alone that drove the development of the Rescorla-Wagner model and similar theories, but stimulus competition phenomena in general. This includes not only blocking, but also unblocking (Kamin, 1969), the relative validity effect (Wagner et al., 1968), the contingency effect (Rescorla, 1968), overshadowing (Mackintosh, 1976; Pavlov, 1927), procedures to produce conditioned inhibition (Rescorla, 1969), and phenomena that were first predicted by such models, such as overexpectation (Rescorla, 1970) and superconditioning (Rescorla, 1971). It is the accumulation of evidence for stimulus competition what is the touchstone for traditional theories of elementary learning, and not blocking alone.

Furthermore, the assumption that a global error signal drives learning has been kept by most contemporary models, mostly because it has received support from a variety of lines of research, including not only conditioning research, but also the study of complex forms of learning (for an early review, see Siegel & Allan, 1996) like category learning (Gluck & Bower, 1988; Soto & Wasserman, 2010a,b), object recognition (Soto & Wasserman, 2012; Soto et al., 2012) and causal and contingency learning (e.g., Dickinson, 2001; Shanks, 1985), and research on the neurobiological substrates of learning (e.g., Waelti et al., 2001; Kim et al., 1998; McNally et al., 2011). Much of that research has involved variations of the original blocking design.

Conclusion

In sum, the results from ten out of the fifteen experiments described by Maes et al. are rather unsurprising in the light of contemporary associative learning theory. The conclusion that "a true blocking effect is more difficult to obtain than one might assume from the literature and that we lack insight into its boundary conditions" (p. e60) seems far-fetched. The results from three other experiments are uninterpretable, due to an extremely low level of conditioned responding to X in the control group, pointing to a floor effect. The results from the other two experiments can be explained in a number of ways, and a failure to replicate blocking in only two experiments is unsurprising. I would recommend that we do not take blocking out of the handbooks of psychology just yet, but instead update those books to reflect our current understanding of associative learning processes.

References

Bonardi, C., Bartle, C., Bowles, K., de Pulford, F., & Jennings, D. J. (2010). Some appetitive procedures for examining associative learning in the mouse: Implications for psychopathology. *Behavioural Brain Research*, *211*(2), 240–247. Retrieved 2017-04-07, from `http://www.sciencedirect.com/science/article/pii/S0166432810002457`

Dickinson, A. (2001, February). The 28th Bartlett Memorial Lecture Causal learning: An associative analysis. *Quarterly Journal of Experimental Psychology*, *54B*(1), 3–25.

Dwyer, D. M., Haselgrove, M., & Jones, P. M. (2011). Cue interactions in flavor preference learning: a configural analysis. *Journal of Experimental Psychology: Animal Behavior Processes*, *37*(1), 41. Retrieved 2017-04-11, from `http://psycnet.apa.org/journals/xan/37/1/41/`

Ghirlanda, S. (2015, February). On elemental and configural models of associative learning. *Journal of Mathematical Psychology*, *64–65*, 8–16. doi: 10.1016/j.jmp.2014.11.003

Gluck, M. A., & Bower, G. H. (1988, September). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*(3), 227–247.

Hall, G., Mackintosh, N. J., Goodall, G., & Martello, M. D. (1977, May). Loss of control by a less valid or by a less salient stimulus compounded with a better predictor of reinforcement. *Learning and Motivation*, *8*(2), 145–158. Retrieved 2017-05-10, from `http://www.sciencedirect.com/science/article/pii/0023969077900017` doi: 10.1016/0023-9690(77)90001-7

Harris, J. A. (2006, July). Elemental representations of stimuli in associative learning. *Psychological Review*, *113*(3), 584–605.

Harris, J. A., & Livesey, E. J. (2010, January). An attention-modulated associative network. *Learning & Behavior*, *38*(1), 1–26. doi: 10.3758/LB.38.1.1

Jones, P. M., & Haselgrove, M. (2013). Blocking and associability change. *Journal of Experimental Psychology: Animal Behavior Processes*, *39*(3), 249. Retrieved 2017-04-11, from `http://psycnet.apa.org/journals/xan/39/3/249/`

Kamin, L. J. (1969). Selective association and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental Issues in Associative Learning* (pp. 42–64). Halifax: Dalhousie University Press.

Kehoe, E. J., Horne, A. J., Horne, P. S., & Macrae, M. (1994). Summation and configuration between and within sensory modalities in classical conditioning of the rabbit. *Animal Learning and Behavior*, *22*(1), 19–26.

Kim, J. J., Krupa, D. J., & Thompson, R. F. (1998). Inhibitory cerebello-olivary projections and blocking effect in classical conditioning. *Science*, *279*(5350), 570–573. Retrieved 2017-04-12, from `http://science.sciencemag.org/content/279/5350/570.short`

Kinder, A., & Lachnit, H. (2003, March). Similarity and Discrimination in Human Pavlovian Conditioning. *Psychophysiology*, *40*(2), 226–234. doi: 10.1111/1469-8986.00024

LoLordo, V. M., Jacobs, W. J., & Foree, D. D. (1982). Failure to block control by a relevant stimulus. *Learning & Behavior*, *10*(2), 183–192. Retrieved 2017-04-12, from `http://www.springerlink.com/index/A4G1885Q616307W4.pdf`

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*, 276–298.

Mackintosh, N. J. (1976). Overshadowing and stimulus intensity. *Animal Learning & Behavior*, *4*(2), 186–192.

Maes, E., Boddez, Y., Alfei, J. M., Krypotos, A. M., D'Hooge, R., De Houwer, J., & Beckers, T. (2016). The elusive nature of the blocking effect: 15 failures to replicate. *Journal of Experimental Psychology: General*, *145*(9), e49–e71.

McLaren, I. P. L., & Mackintosh, N. J. (2002, August). Associative learning and elemental representation: II. Generalization and discrimination. *Animal Learning & Behavior*, *30*(3), 177–200.

McNally, G. P., Johansen, J. P., & Blair, H. T. (2011, May). Placing prediction into the fear circuit. *Trends in Neurosciences*. Retrieved 2011-06-02, from `http://www.cell.com/trends/neurosciences/abstract/S0166-2236(11)00055-5` doi: 10.1016/j.tins.2011.03.005

Melchers, K. G., Shanks, D. R., & Lachnit, H. (2008, March). Stimulus coding in human associative learning: Flexible representations of parts and wholes. *Behavioural Processes*, *77*(3), 413–427.

Pavlov, I. P. (1927). *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*. London: Oxford University Press.

Pearce, J. M. (1987, January). A model for stimulus generalization in Pavlovian conditioning. *Psychological Review*, *94*(1), 61–73.

Pearce, J. M. (1994, October). Similarity and discrimination: A selective review and a connectionist model. *Psychological Review*, *101*(4), 587–607.

Pearce, J. M. (2002, May). Evaluation and development of a connectionist theory of configural learning. *Animal Learning & Behavior*, *30*(2), 73–95.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, *87*(6), 532–552.

Rescorla, R. A. (1968, August). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, *66*(1), 1–5. doi: 10.1037/h0025984

Rescorla, R. A. (1969). Pavlovian conditioned inhibition. *Psychological Bulletin*, *72*(2), 77–94. doi: doi:10.1037/h0027760

Rescorla, R. A. (1970). Reduction in the effectiveness of reinforcement after prior excitatory conditioning. *Learning and Motivation*, *1*(4), 372–381.

Rescorla, R. A. (1971). Variation in the effectiveness of reinforcement and nonreinforcement following prior inhibitory conditioning. *Learning and Motivation*, *2*(1), 113–123.

Rescorla, R. A., & Durlach, P. J. (1981). Within-event learning in Pavlovian conditioning. In N. E. Spear & R. R. Miller (Eds.), *Information Processing in Animals: Memory Mechanisms* (pp. 83–111). Hillsdale, NJ: Erlbaum.

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 theory and research* (pp. 64–99). New York: Appleton-Century-Crofts.

Sanderson, D. J., Jones, W. S., & Austen, J. M. (2016, January). The effect of the amount of blocking cue training on blocking of appetitive conditioning in mice. *Behavioural Processes*, *122*, 36–42. Retrieved 2017-04-07, from `http://www.sciencedirect.com/science/article/pii/S0376635715300747` doi: 10.1016/j.beproc.2015.11.007

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgement. *The Quarterly Journal of Experimental Psychology*, *37B*(1), 1–21.

Siegel, S., & Allan, L. G. (1996, September). The widespread influence of the Rescorla-Wagner model. *Psychonomic Bulletin & Review*, *3*(3), 314–321.

Soto, F. A., Gershman, S. J., & Niv, Y. (2014). Explaining compound generalization in associative and causal learning through rational principles of dimensional generalization. *Psychological Review*, *121*(3), 526–558.

Soto, F. A., Quintana, G. R., Pérez-Acosta, A. M., Ponce, F. P., & Vogel, E. H. (2015, October). Why are some dimensions integral? Testing two hypotheses through causal learning experiments. *Cognition*, *143*, 163–177. doi: 10.1016/j.cognition.2015.07.001

Soto, F. A., Siow, J. Y. M., & Wasserman, E. A. (2012, June). View-invariance learning in object recognition by pigeons depends on error-driven associative learning processes. *Vision Research*, *62*, 148–161. Retrieved 2012-05-04, from `http://www.sciencedirect.com/science/article/pii/S0042698912001101` doi: 10.1016/j.visres.2012.04.004

Soto, F. A., & Wasserman, E. A. (2010a). Error-driven learning in visual categorization and object recognition: A common elements model. *Psychological Review*, *117*(2), 349–381.

Soto, F. A., & Wasserman, E. A. (2010b). Missing the forest for the trees: Object discrimination learning blocks categorization learning. *Psychological Science*, *21*(10), 1510–1517. doi: 10.1177/0956797610382125

Soto, F. A., & Wasserman, E. A. (2012). Visual object categorization in birds and primates: Integrating behavioral, neurobiological, and computational evidence within a "general process" framework. *Cognitive, Affective, & Behavioral Neuroscience*, *12*(1), 220–240. Retrieved 2011-11-22, from `http://www.springerlink.com/content/475240w080x15129/` doi: 10.3758/s13415-011-0070-x

Thorwart, A., Livesey, E., & Harris, J. (2012). Normalization between stimulus elements in a model of Pavlovian conditioning: Showjumping on an elemental horse. *Learning & Behavior*, *40*(3), 334–346. doi: 10.3758/s13420-012-0073-7

Waelti, P., Dickinson, A., & Schultz, W. (2001). Dopamine responses comply with basic assumptions of formal learning theory. *Nature*, *412*(6842), 43–48.

Wagner, A. R. (2003, February). Context-sensitive elemental theory. *Quarterly Journal of Experimental Psychology*, *56B*(1), 7–29.

Wagner, A. R. (2007). Evolution of an elemental theory of Pavlovian conditioning. *Learning & Behavior*, *36*(3), 253–265.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology*, *76*(2), 171–80.

Yamada, K. (2010, November). Strain differences of selective attention in mice: Effect of Kamin blocking on classical fear conditioning. *Behavioural Brain Research*, *213*(1), 126–129. Retrieved 2017-04-07, from `http://www.sciencedirect.com/science/article/pii/S0166432810003189` doi: 10.1016/j.bbr.2010.04.037