Individual variation in the vigor and form of Pavlovian conditioned responses: Analysis of a model system

Robert C. Honey*, Dominic M. Dwyer, Adela F. Iliescu

Cardiff University, UK

ABSTRACT

Pavlovian conditioning results in individual variation in the vigor and form of acquired behaviors. Here, we describe a general-process model of associative learning (HeiDI; How excitation and inhibition determine ideo-motion) that provides an analysis for such variation together with a range of other important group-level phenomena. The model takes as its starting point the idea that pairings of a conditioned stimulus (CS) and an unconditioned stimulus (US) result in the formation of reciprocal associations between their central representations. The asymptotic values of these associations and the rate at which these are reached are held to be influenced by the perceived salience of the CS (α_{CS}) and US (β_{US}). Importantly, whether this associative knowledge is exhibited in behavior that reflects the properties of the CS (e.g., sign-tracking) or US (e.g., goal-tracking) is also influenced by the relative values of α_{CS} and β_{US}. In this way, HeiDI provides an analysis for both quantitative and qualitative individual differences generated by Pavlovian conditioning procedures.

"*Give me a dozen healthy infants, well-formed, and my own specified world to bring them up in and I'll guarantee to take any one at random and train him to become any type of specialist I might select – doctor, lawyer, artist, merchant-chief and, yes, even beggar-man and thief, regardless of his talents, penchants, tendencies, abilities, vocations, and race of his ancestors. I am going beyond my facts, and I admit it, but so have the advocates to the contrary, and they have been doing so for many thousands of years.*" (p. 82; Watson, 1924).

Watson's central thesis might seem less controversial now than it did almost a century ago: An appropriately conducive environment – where different forms of training can be arranged – affords the development of selected paths in infants taken randomly from the normal population. His central thesis did not deny the existence of individual differences that have different origins (e.g., talents, penchants etc.), but it did suggest that training might free an individual from them. However, the fact that dramatic individual differences in acquired behavior can emerge in animals given identical training in a controlled environment is – if not antithetical – then certainly problematic from an empiricist perspective (e.g., Iliescu, Hall, Wilkinson, Dwyer, & Honey, 2018; Patitucci, Nelson, Dwyer, & Honey, 2016; but see, Byrom & Murphy, 2018). For example, such individual variation is beyond the scope of general-process theories of associative learning (e.g., Rescorla & Wagner, 1972; Mackintosh, 1975; Pearce & Hall, 1980; Wagner, 1981) in which the relationship between the strength of an association and performance is held to be monotonic: How could a single acquired property (like associative strength; V) be manifest in distinct ways across a set of rats? We have recently presented a model, HeiDI, which offers a potential answer to this question (Honey et al., 2020a, 2020b). Before we present that answer, we should first describe results that provided an important impetus for the development of HeiDI.

* Corresponding author at: School of Psychology, Cardiff University, Tower Building, Park Place, Cardiff, CF10 3AT, UK.
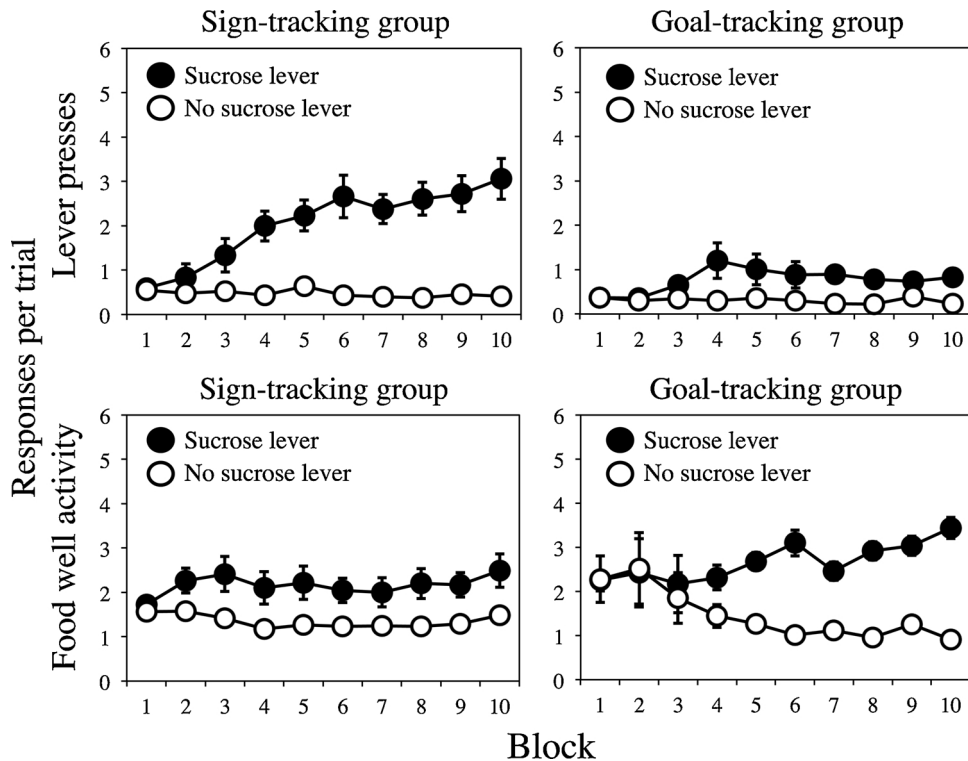 *E-mail address:* Honey@cardiff.ac.uk (R.C. Honey).

Fig. 1. The emergence of qualitative differences in conditioned behavior across 10 blocks of training. Mean (± SEM) levels of lever activity (sign-tracking) and food well activity (goal-tracking). Rats were divided into sign-trackers (left panels) and goal-trackers (right panels), with scores separated for the lever paired with sucrose and the lever that was not. Adapted from: Patitucci, E., Nelson, N., Dwyer, D.M., & Honey, R.C. (2016). The origins of individual differences in how learning is expressed in rats: A general-process perspective. *Journal of Experimental Psychology*: *Animal Learning and Cognition, 42,* 313-324.

1. Individual differences in conditioned behavior

What happens when a given set of healthy rats is food restricted and then receives identical training trials in a standard conditioning chamber? If the temporary insertion of a lever into the chamber serves as the conditioned stimulus (CS) and the delivery of sucrose (for example) serves as the unconditioned stimulus (US), then the CS will come to elicit conditioned responses (CRs): The rats develop a tendency to interact with the lever and to approach the well into which sucrose is about to be delivered. The fact that the CS can provoke multiple CRs is of interest in its own right: It suggests that the CR is not simply determined by the unconditioned responses provoked by the US (e.g., to approach the location where the US will be delivered; cf. Pavlov, 1927). More interesting, however, is the fact that the distribution of these two forms of CR differ across rats given this autoshaping procedure: with some being more likely to interact with the lever (called sign-tracking; e.g., Hearst & Jenkins, 1974) than to approach the food well (called goal-tracking; e.g., Boakes, 1977), and others being more likely to approach the food well than to interact with the lever (e.g., Iliescu et al., 2018; Patitucci et al., 2016). This variation is relatively continuous in nature and remarkably stable from one day to the next: These are *qualitative* individual differences in acquired behavior (see Fitzpatrick et al., 2013; see also, Matzel et al., 2003). Fig. 1 illustrates the results of a representative experiment in which rats received training where the insertion of one lever (for 10 s) was followed by the delivery of sucrose, while the insertion of another was not (Experiment 2; Patitucci et al., 2016). The rats have been separated into two groups based on their bias to sign-track and goal-track in the final block of training; with the bias scores calculated in the following way: (number goal-tracking responses – number of sign-tracking responses) / (number goal-tracking responses + number of sign-tracking responses). Using this measure of bias to split the rats into two groups enables the development of sign-tracking and goal-tracking behaviors to be illustrated. But as already noted, the individual differences are relatively continuous across the group of rats as a whole, as we will have cause to return to at a later point.

It takes a moment to fully appreciate the theoretical significance of the fact that rats given identical training exhibit what they have learnt in qualitatively different ways: There is no coherent mapping between learning and the two measures of performance (sign-tracking and goal-tracking). Focusing on lever-oriented behavior suggests that the sign-tracking group learnt more readily than the goal-tracking group, while focusing on food-well oriented behavior suggests just the opposite. This is true whether one construes learning as the development of associations between the representation of the CS (the lever) and the processes responsible for generating responding (e.g., Hull, 1943; Spence, 1937) or between the representations of the CS and the US (e.g., Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972). Indeed the problem is yet more ubiquitous: It is a challenge to any model that
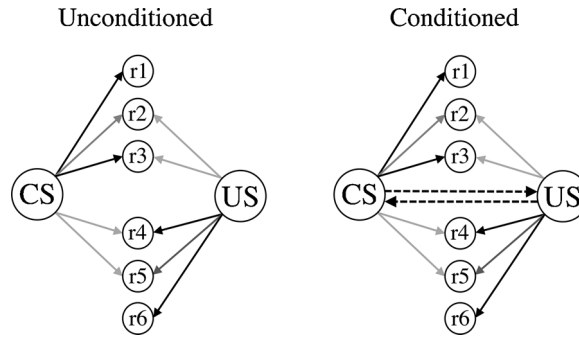
Unconditioned   Conditioned

Fig. 2. A schematic of the associative structures that are assumed to underpin Pavlovian learning and performance. The left-hand depicts the unconditioned structure (i.e., before conditioning), with the darkness of the links between the CS and r1-r6 and the US and r1-r6 indicating their strength; and the right-hand side depicts the conditioned structure (i.e., after conditioning). The reciprocal CS-US and US-CS associations are denoted by the dashed lines. Adapted from: Honey, R.C., Dwyer, D.M., & Iliescu, A.F. (2020). HeiDI: A model for Pavlovian learning and performance with reciprocal associations. *Psychological Review* (in press).

assumes that a single process underlies learning and its translation into performance (e.g., Gallistel & Gibbon, 2000; Stout & Miller, 2007). In the context of associative theories of learning, the problem reflects a surprising reluctance to specify how the strength of a CS-US association (i.e., V_{CS-US}) is translated into conditioned responding: with most models simply assuming that there is a monotonic relationship between V_{CS-US} and conditioned behavior. What is required is a general-process model that (i) specifies the mnemonic structures that underpin learning and performance, (ii) describes the rules governing how change occurs within those structures, and (iii) isolates potential sources of individual differences. HeiDI is such a model.

2. HeiDI

The mnemonic structures assumed to underpin learning and performance are depicted in Fig. 2. The left-hand side shows the pattern of unconditioned links between the CS, US and a set of response generating units (r1-r6). Some of these units are more strongly activated by the CS than the US (r1-r3), and some are more strongly activated by the US than the CS (r4-r6), with the darkness of the links denoting their strength. In this way, a general distinction is made between CS-oriented behaviors (r1-r3; e.g., orienting, lever approach, rearing) and US-oriented behaviors (r4-r6; e.g., food well approach, chewing, swallowing; see also, Holland, 1977, 1984). The right-hand side of Fig. 2 shows the conditioned structure, where standard Pavlovian conditioning trials – in which the CS immediately precedes the US – are assumed to result in the formation of reciprocal (CS-US and US-CS) associations (cf. Asratian, 1952, 1965; Pavlov, 1932; for a review, see Goremezano & Tait, 1976; but see, Konorski, 1948). In this way, a minimal functional cell assembly is formed: When the CS is presented activation propagates to the US, which is propagated back to the CS (e.g., Grossberg, 1980; Hebb, 1949). This structure contrasts with other trial-based models of associative learning (e.g., Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972) where the focus has been on the formation of CS-US associations, with their functional alignment to a process of prediction in spite of the fact that the models are trial based (i.e., the change in associative strength is assumed to occur upon presentation of the US). The reciprocal nature of the associations within HeiDI means that there is a basis for CS-oriented conditioned responses to change, through the influence of the US-CS association, independently of the influence of the CS-US association, which primarily affects US-oriented conditioned responses. The rules governing the formation of these associations are captured in Eqs. 1 and 2. Eq. 1 determines the development of the CS-US association, and is a simplification to the Rescorla-Wagner error-correcting learning rule, while Eq. 2 is the formally equivalent rule for the development of US-CS associations. These equations both reduce the number of free parameters from those in the Rescorla-Wagner model, and generate increased explanatory power, especially when coupled with formally equivalent equations for determining the formation of associations between two CSs (which we will come to shortly).

$$\Delta V_{CS\text{-}US} = \alpha_{CS}(c.\beta_{US} - \Sigma V_{TOTAL\text{-}US}) \tag{1}$$

$$\Delta V_{US\text{-}CS} = \beta_{US}(c.\alpha_{CS} - \Sigma V_{TOTAL\text{-}CS}) \tag{2}$$

The perceived salience of the US determines the asymptote for the CS-US association, and at the same time the perceived salience of the CS determines the asymptote for the US-CS association. In Eq. 1, the maximum strength of the CS-US association is 1 in units of V (denoted c) modulated by the value of the parameter β_{US}, which aligns to the perceived salience of the US (i.e., $c.\beta_{US}$). The value of α_{CS} determines the rate at which the CS-US association changes, and aligns to the perceived salience of the CS. In Eq. 2, the maximum strength of the US-CS association is again 1 in units of V (i.e., c), and is modulated by the value of α_{CS} (i.e., $c.\alpha_{CS}$). In this case, the value of β_{US} determines the rate at which the US-CS association changes. α_{CS} and β_{US} are confined to the unit interval $0 \leq \alpha_{CS}, \beta_{US} \leq 1$, and thus modulate both the rate of learning and the maximum strength of the reciprocal associations. In Eqs. 1 and 2, when the CS is absent, α_{CS} and $c.\alpha_{CS}$ are both set to 0 and when the US is absent, β_{US} and $c.\beta_{US}$ are both set to 0. This arrangement allows the strength of the CS-US association to change on trials on which the CS is present, but the US is no longer presented, and the US-CS association to change on trials on which the US is presented but the CS is absent (cf. Wagner & Rescorla, 1972; Wagner, Logan,

Haberlandt, & Price, 1968).

According to Eq. 1, during simple conditioning, the strength of the CS-US association (i.e., V_{CS-US}) converges asymptotically on c.β_{US}, with the change in associative strength of this association on a given trial (ΔV_{CS-US}) being determined by the error in the pooled error term for the US (c.β_{US} – $\Sigma V_{TOTAL-US}$). $\Sigma V_{TOTAL-US}$ denotes the total net associative strength of the CS with respect to the US, and excitatory learning stops when $\Sigma V_{TOTAL-US}$ = c.β_{US}, and the learning rate parameter α_{CS} affects the rate at which V_{CS-US} approaches c.β_{US}. Under conditions in which more than one CS (e.g., stimulus A and stimulus B) is paired with a US, the pooled error term means that the development of the A-US association will be influenced by the strength of the B-US association. That is, $\Sigma V_{TOTAL-US}$ is equal to the total or combined associative strengths of A and B with respect to the US.

Eq. 2 is the complementary rule governing the formation of the US-CS association. Here, the change in the strength of this association (ΔV_{US-CS}) on a given trial is also determined by the error within the pooled error term for the CS (c.α_{CS} – $\Sigma V_{TOTAL-CS}$); with $\Sigma V_{TOTAL-CS}$ denoting the associative strength of the single US (in typical conditioning procedures) with respect to the CS. Learning ceases when $\Sigma V_{TOTAL-CS}$ = c.α_{CS}, and the learning rate parameter β_{US} affects the rate at which V_{US-CS} approaches c.α_{CS}. If we now consider what happens on a trial in which a compound of two stimuli (A and B) precedes a US, then the c.α_{CS} values for each CS in a compound (i.e., c.α_A and c.α_B) set independent asymptotes for the US-A and US-B associations; and B will compete with the US for association with A, and A will compete with the US for association with B (see Honey et al., 2020a, 2020b).

The associative structures depicted in Fig. 2, together with the learning (Eqs. 1–3) and performance (Eqs. 4 and 5) rules, are readily extended to the case where two CSs (A and B) become linked. In Eqs. 1 and 2, for example, α_{CS} and β_{US} can be replaced with α_A and α_B, respectively. In fact, group-level differences in such associations enable HeiDI to explain some of the phenomena that were beyond the scope of the model proposed by Rescorla and Wagner (1972); see Honey et al. (2020a, 2020b). But, individual differences in the perceived saliences of A and B could also constitute a basis for individual differences in learning and performance. These differences could be directly observed during pairings of two stimuli that evoke different unconditioned behaviors (e.g., a tone with a light; see Honey, Good, & Manser, 1998; Honey, Watt, & Good, 1998; see also, Narbutovich & Podkopayev, 1936; cited in Konorski, 1948, p. 91), or indirectly observed through a range of sensory preconditioning procedures (e.g., Pavlov, 1931/1932; cited in Kimmel, 1977; see also Brogden, 1939; Rescorla & Cunningham, 1978; Ward-Robinson & Hall, 1996). However, as far as we are aware, there have been no studies that have examined whether or not such individual differences are evident in the vigor or form of conditioned behavior acquired as a consequence of sensory-sensory pairings.

We have already noted that HeiDI assumes that α_{CS} and β_{US} are aligned to the perceived salience of the CS and US, respectively. When this assumption is combined with the learning rules (i.e., Eqs. 1 and 2) it is clear that individual differences in the perceived salience of the CS and US will affect both the asymptotic values of the CS-US and US-CS associations, and the rates at which these values are reached. That is, *quantitative* differences in learning are predicted to the extent that there are differences in the perceived salience of the CS and US. However, the individual differences in CS-oriented and US-oriented conditioned behaviors (e.g., sign-tracking and goal-tracking) are not only quantitative but also *qualitative* (see Fig. 1). This fact clearly requires that there is a more complex mapping of associative strength than a monotonic one. HeiDI first assumes that upon presentation of the CS, the associative strengths returned by Eqs. 1 and 2 are combined (V_{COMB}) in the way specified in Eq. 3. V_{COMB} represents the associative resonance within the CS-US assembly, and Eq. 3 weights the associative strength of the stimulus that is present (e.g., V_{CS-US}) more than the association involving an associatively activated node (e.g., V_{US-CS}; see Equation 3). In Eq. 3, V_{COMB} is in units of V, because V_{CS-US} is rendered dimensionless by multiplying it by 1/c. This combination rule captures the idea that the two stimuli function as an assembly, but one in which associative activity generated by the presentation of a stimulus (e.g., the CS) is subject to a process of dampening. That is, the US-CS association is only activated to the degree that the US itself is activated via the CS-US association.

$$V_{COMB} = V_{CS-US} + \left(\frac{1}{c}.\ V_{CS-US} \times V_{US-CS}\right)$$

(3)

Critically, HeiDI assumes that when the CS is presented, V_{COMB} is distributed into two components, which have different influences on performance: A CS-oriented component (which influences sign-tracking), and a US-oriented component (which influences goal-tracking). One way to do this, is according to the relative perceived saliences of the CS and US (i.e., α_{CS} and β_{US}): If $\alpha_{CS} > \beta_{US}$ then CS-oriented behavior would dominate US-oriented behavior, and if $\beta_{US} > \alpha_{CS}$ then the reverse is the case. However, this proposal is inadequate, because while the perceived salience of a CS will be available upon CS presentation, the perceived salience of the US will not: The US is not present. We have proposed, therefore, that the distribution of V_{COMB} is determined by the value of α_{CS} relative to V_{CS-US} (which reflects β_{US}). That is, HeiDI assumes that V_{COMB} is distributed into the CS-oriented and US-oriented components according to the perceived salience of the CS (α_{CS}) relative to the perceived salience of the retrieved US representation (i.e., V_{CS-US}); which is indirectly influenced by the perceived salience of the US (i.e., β_{US}; see Eq. 1). This idea is formally expressed in Eqs. 4 and 5 that generate two components, R_{CS} and R_{US}, which are held to affect the levels of CS-oriented and US-oriented responding, respectively. According to these equations, when $\alpha_{CS} > V_{CS-US}$ then $R_{CS} > R_{US}$, but the reverse is the case when $V_{CS-US} > \alpha_{CS}$. That is, the balance between CS- and US-oriented behavior is related to the individually perceived salience of the CS and US themselves. To address the fact that Eq. 1 (and Eq. 2) can return negative Vs, the absolute value of V_{CS-US} is used in Eqs. 4 and 5 to ensure that the proportions are < 1. As before, $|V_{CS-US}|$ is transformed into a dimensionless value by multiplying it by 1/c, which means that R_{CS} and R_{US} assume units of V. Clearly, Eqs. 4 and 5 are intimately related, being reciprocals of one another.

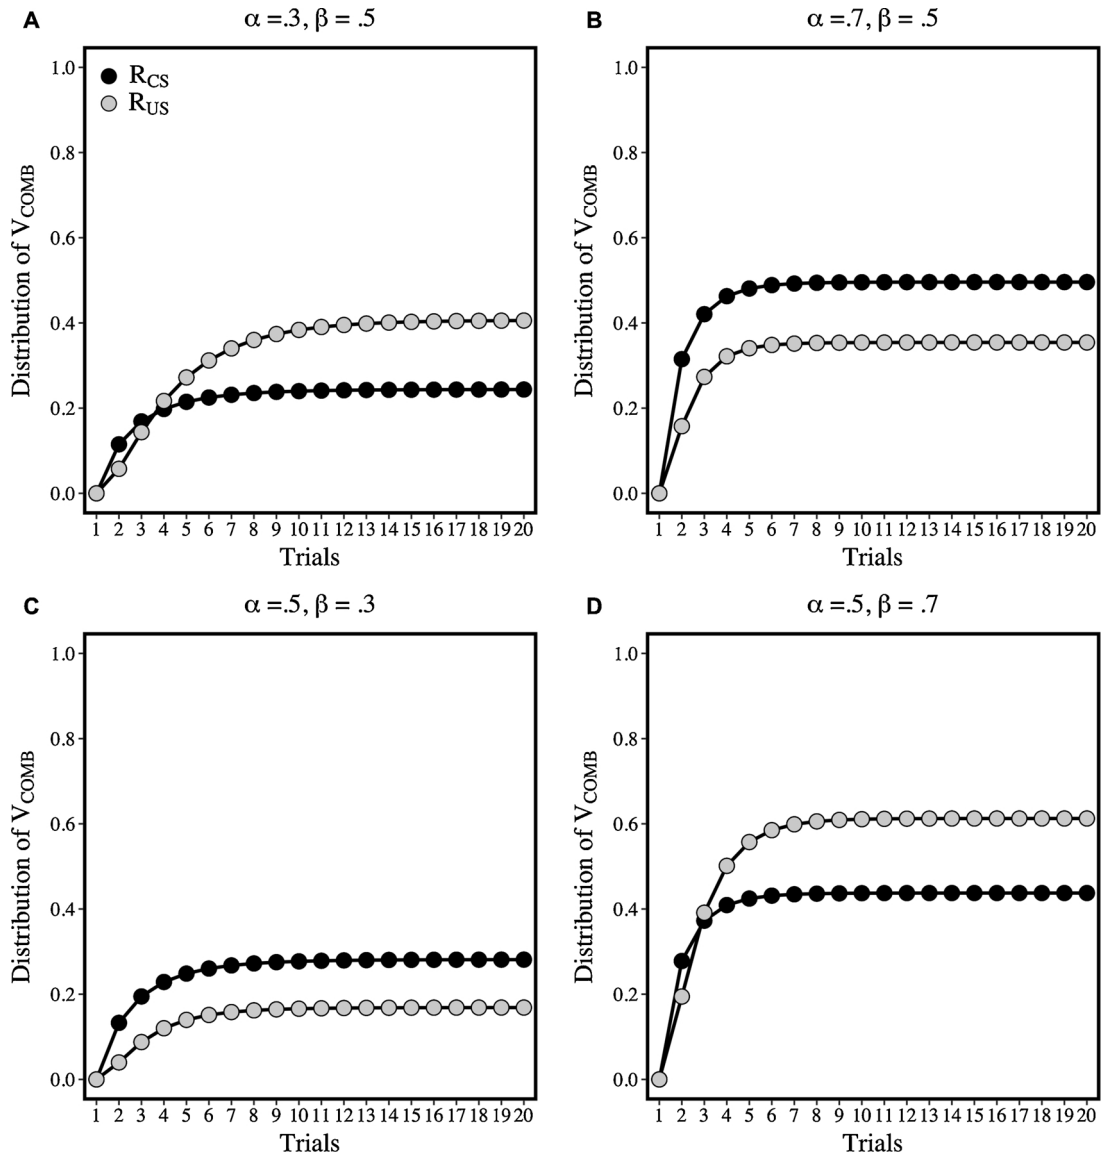$$R_{CS} = \frac{\alpha_{CS}}{\alpha_{CS} + \frac{1}{c}.\ |V_{CS-US}|}V_{COMB}$$

(4)

Fig. 3. Simulations of the distribution of V_{COMB} into R_{CS} and R_{US} across 20 conditioning trials. R_{CS} (black symbols) and R_{US} (grey symbols) outputs were generated when the following values of α_{CS} and β_{US} were used in Eqs. 1–5. Panels A and B: α_{CS} was either .30 (A) or .70 (B) and β_{US} was fixed at .50. Panels C and D: α_{CS} was fixed at .50 and β_{US} was either .30 (C) or .70 (D). Adapted from: Honey, R.C., Dwyer, D.M., & Iliescu, A.F. (2020). HeiDI: A model for Pavlovian learning and performance with reciprocal associations. *Psychological Review* (in press).

$$R_{US} = \frac{\frac{1}{c} \cdot |V_{CS\text{-}US}|}{\alpha_{CS} + \frac{1}{c} \cdot |V_{CS\text{-}US}|} V_{COMB}$$

(5)

$$r1 = \left(\frac{1}{c} \cdot R_{CS} \times V_{CS\text{-}r1} \right) + \left(\frac{1}{c} \cdot R_{US} \times V_{US\text{-}r1} \right)$$

(6)

Simulations confirm that changing the value of α_{CS} relative to β_{US} results in changes in R_{CS} relative R_{US}: Fig. 3 shows that when $\alpha_{CS} > \beta_{US}$ then $R_{CS} > R_{US}$ (see panels B and C), and when $\alpha_{CS} < \beta_{US}$ then $R_{CS} < R_{US}$ (see panels A and D). If these differences are multiplied by the strengths of the links (see Fig. 2) between the CS and r1-r6 (e.g., $V_{CS\text{-}r1}$) and the US and r1-6 (e.g., $V_{US\text{-}r1}$), according to Eq. 6, then we have the basis for the translation of associative strength into different forms of behavior. We assume that the value of r1-r6, which takes units of V, is reflected in the vigor of their corresponding response forms.

HeiDI provides a simple associative analysis for a broad range of group-level phenomena that have proven to be a challenge over a protracted period, and does so without the need to appeal to some tendentious assumptions or additional (e.g., attentional) processes

(see Honey et al., 2020a, 2020b). However, for the present purposes, it is most relevant to consider a series of findings that provide support for the analysis of individual differences in the form of responding that HeiDI provides; findings that are also beyond extant associative models. One finding that is consistent with the idea that β_{US} affects performance (through its impact on $V_{CS\text{-}US}$) was reported by Patitucci et al. (Experiment 2, 2016). In this experiment, blocks of training in which rats received pairings of the presentation of a lever with the delivery of a sucrose reinforcer, were interposed with individual sessions where the affective responses of the rats to sucrose was assessed through an analysis of microstructure of the licking (see Dwyer, 2012). Patitucci et al. observed that individual differences in goal-tracking a positively correlated with the affective value of sucrose, whereas sign-tracking had - if anything - a negative relationship with the affective value of sucrose (see also, Morrison, Bamkole, & Nicola, 2015). This is precisely what would be expected if the affective value of the US, as measured by the microstructure of licking, reflected β_{US}. Patitucci et al. (Experiment 1, 2016) also observed that if the presentation of one lever was followed by one US (e.g., sucrose) while the presentation of a second lever was followed by a different US (e.g., food pellets) then the bias to either sign-tracking or goal-tracking on one lever was uncorrelated with the bias observed on the other lever. This observation is consistent with the interpretation that different USs have different β_{US} values across a group of rats. This interpretation received direct support from the observation that when the same US (food or sucrose) is paired with both levers then there was now a strong correlation between the biases shown during presentations of the two levers (Iliescu et al., 2018). Finally, HeiDI assumes that the current value of $V_{CS\text{-}US}$ affects how V_{COMB} is distributed into R_{CS} and R_{US} (see Eqs. 4 and 5): when $V_{CS\text{-}US}$ is low relative to α_{CS} then $R_{US} < R_{CS}$. This assumption predicts that a procedure involving experimental extinction, which should result in a reduction in net $V_{CS\text{-}US}$ while leaving α_{CS} the same, should result in the preferential distribution of V_{COMB} to the R_{CS} component (see Eq. 4) rather than to the R_{US} component (see Eq. 5); and thus sign-tracking being more resistant to extinction than goal-tracking. This is precisely what Iliescu et al. (Experiment 1, 2018) observed in an experiment in which rats first received a discrimination in which one lever was reinforced and a second was not, and then these contingencies were reversed (see also, Ahrens, Singer, Fitzpatrick, Morrow, & Robinson, 2016): The reduction in goal-tracking to the formerly reinforced lever was more rapid than the reduction in sign-tracking. This observation was apparent irrespective of whether the rats had an original bias to goal-track or sign-track (see also, Anselme, Robinson, & Berridge, 2012). The key results from the study by Iliescu et al. (2018) are depicted in Fig. 4.

The upper panels of Fig. 4 depict the relationships between lever presses on training blocks 5 and 6 (left panel) and between food-well entries on the same blocks (right panel). Each circle represents a given rat. The filled and open circles denote those classified as sign-trackers and goal-trackers, respectively, on the basis of their bias score on the final training block. There was a clear correlation between lever presses on blocks 5 and 6 and between food-well entries on these blocks. Moreover, the upper panels indicate that the rats designated sign-trackers showed higher levels of lever presses than those designated as goal-trackers, and rats designated goal-trackers showed higher levers of food-well entries than those designated sign-trackers. The lower panels depict the relationships between the two forms of responding on the final block of training and the first block in which the contingencies were changed (i.e., reversed). It is clear that while lever pressing remained relatively unchanged by this manipulation, goal-tracking no longer reflected the previous contingencies. In fact, over the course of reversal training, goal-tracking came to rapidly match the changed contingencies while sign-tracking changed less rapidly (see Iliescu et al., 2018).

The evidence presented thus far has implicated the perceived salience of the US in the distribution of US-oriented and CS-oriented behavior. HeiDI also supposes that the perceived salience of the CS should have an effect: a CS with higher perceived salience should result in more CS-oriented and less US-oriented behavior. We do not have any direct evidence that bears on this prediction. However, if one assumed that the perceived salience of the CS is high at the outset of a stimulus and declines across its duration (e.g., due to habituation), then this should be evident as a reduction in sign-tracking and an increase in goal-tracking across the duration of a CS. That is, a plausible additional assumption predicts that while some indices of conditioning should decline across a CS, others should increase (i.e., show inhibition of delay; Pavlov, 1927). It is worth noting that a similar idea was briefly entertained by Mackintosh (p. 62, 1974; see also, Pavlov, 1927, p. 104; Staddon, 2005; Staddon & Higa, 1999; Wagner, 1981), who argued that a decay process might enable the early and later parts of a CS to be discriminated, and thereby provide a basis for inhibition of delay. In any case, formal simulations of HeiDI together with an analysis of archival data confirmed the accuracy of the predictions identified above (see Figs. 5 and 6; Iliescu, Dwyer, & Honey, 2020).

The simulations were conducted with the α_{CS} values set from the outset of the CS at either .30 or .70 and β_{US} fixed at .50 (panels A ad B of Fig. 5), or with the α_{CS} value set from the outset of the CS at .50 and β_{US} set to .30 or .70 (panels C and D of Fig. 5). Across the 4 epochs of each trial, these α_{CS} values were subject to exponential decay: $\alpha_{CS}(1 - 0.10)^2$. The α_{CS} value in the fourth epoch was used to calculate the change in associative strength in Eqs. 1 and 2. R_{CS} and R_{US} were calculated by entering the decaying value of α_{CS} into Eqs. 4 and 5; and α_{CS} was reset to its starting value after a given trial. It is clear that irrespective of whether the values for the critical parameters, α_{CS} and β_{US}, were set to generate higher values of R_{US} relative to R_{CS} (panels A and D) or higher values of R_{CS} relative to R_{US} (panels B an C), R_{US} increased across epochs within a trial while R_{CS} decreased.

Fig. 6 shows how the lever presses and food-well entries change across training blocks (T1-T6) and across four 2.5-s epochs within the 10-s trials, with the upper panes depicting a group of rats designated sign-trackers (group ST) and the lower panes depicting a group of rats designated goal-trackers (group GT). z-transformed scores were used to put the different responses (lever presses and food-well entries) on the same scale. Inspection of Fig. 6 shows that the lever-pressing bias in group ST was most evident at the start of trials, whereas food-well entry bias in group GT was most evident at the end of the trials (see also, Derman, Schneider, Juarez, & Delamater, 2018; but see, Lee et al., 2018). These patterns of results match those predicted by HeiDI.
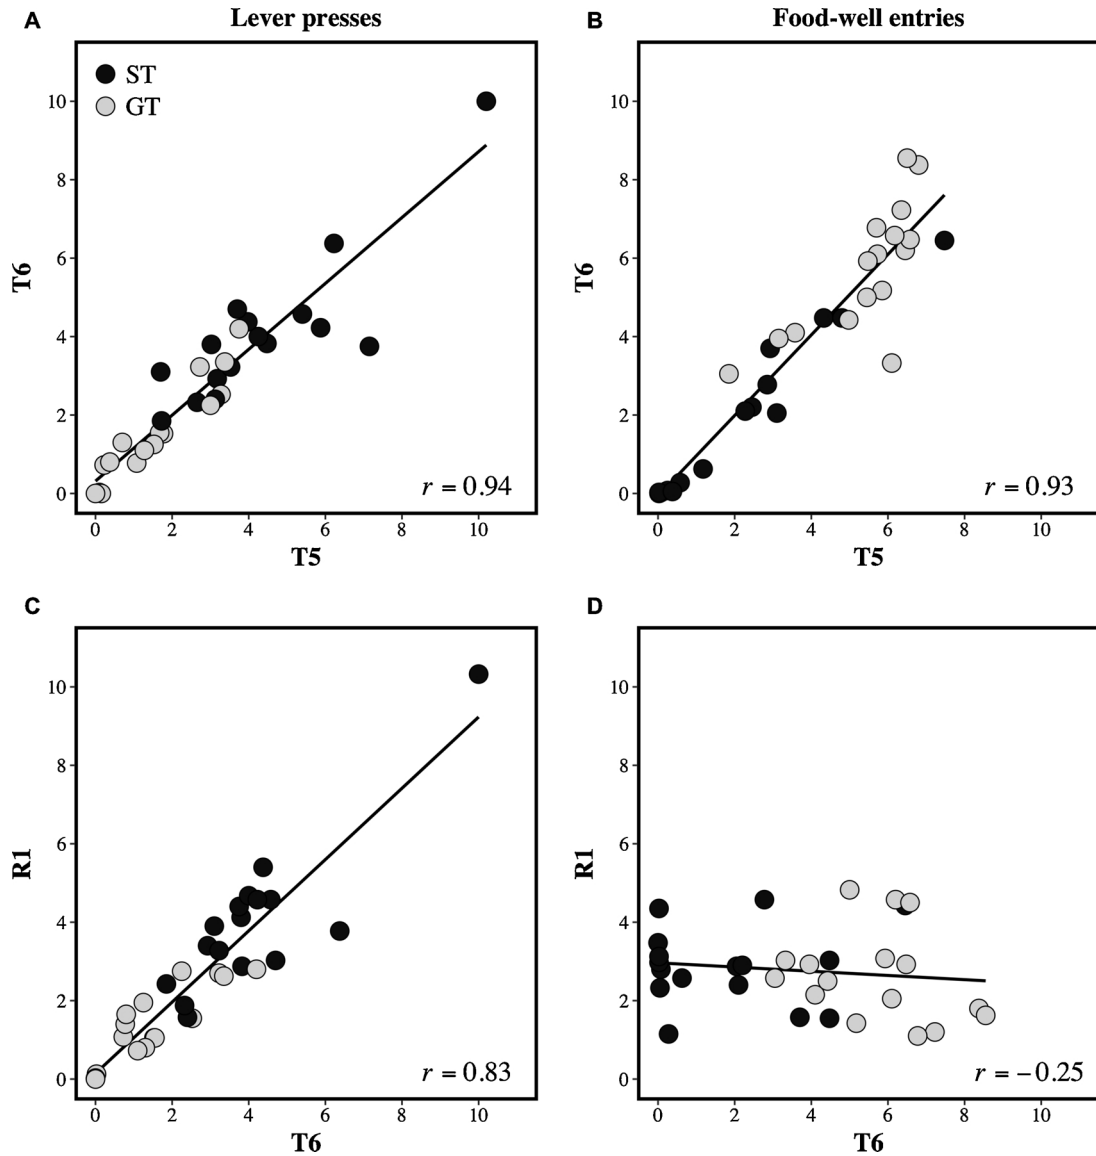
Fig. 4. The upper panels show the relationship between the mean number of responses per reinforced trial for lever presses (left-hand panel) and for food-well entries (right-hand panel) on training blocks 5 (T5) and 6 (T6). The lower panels show the relationship between the final block of training (T6) and the first block on which the formerly reinforced lever was no longer reinforced (R1). The black symbols correspond to rats classified as sign-trackers (i.e., group ST) and the grey symbols to those classified as goal-trackers (i.e., group GT). Adapted from: Iliescu, A.F., Hall, J., Wilkinson, L., Dwyer, D.M., & Honey, R.C. (2018). The nature of phenotypic variation in Pavlovian conditioning. *Journal of Experimental Psychology*: *Animal Learning and Cognition*, *44*, 358-369.

3. General discussion

Individual variation in the form of Pavlovian conditioned behavior has been largely neglected by general process models of associative learning. Here, we have illustrated how one model, HeiDI, provides a simple analysis for such variations. The analysis relies on a novel specification of the associative structures that underpin learning and performance (see Fig. 2), and the rules governing how changes occur within those structures (Eqs. 1 and 2). These rules have two free parameters, α_{CS} and β_{US}, which are aligned to the perceived salience of the CS and US, respectively. These parameters are assumed to be fixed for a given rat, but to vary across a set of rats. Variation in the two parameters thereby provides the basis for individual differences in the vigor and form of conditioned behavior. Returning to Watson (1924), we now have a general process model of Pavlovian conditioning that provides a basis for both quantitative and qualitative individual differences in acquired behavior – changes that emerge as a consequence of giving the same training to a group of rats. What is more, this model also provides an analysis for group-level phenomena that have
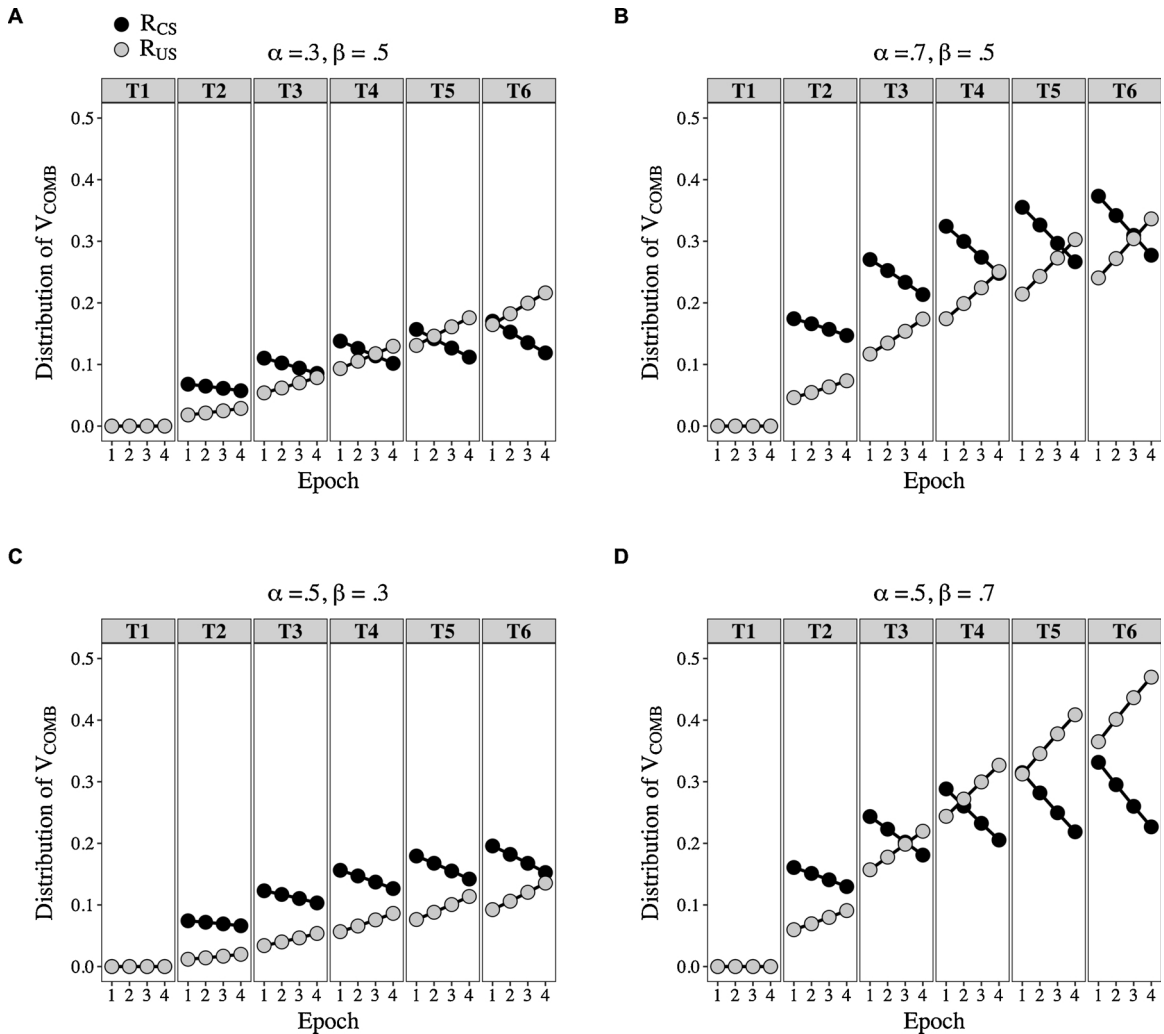
Fig. 5. Simulations of R_{CS} and R_{US} across training blocks (T1-T6) and trial epoch (1-4). The V_{COMB} outputs used in Eqs. 4 and 5 to calculate R_{CS} (black symbols) and R_{US} (grey symbols) were generated using Eqs. 1–3. In panels A and B, at the start of a CS presentation α_{CS} was either .3 (A) or .7 (B) and β_{US} was fixed at .5; and in panels C and D, at the start of the CS presentation α_{CS} was fixed at .5 and β_{US} was either .3 (C) or .7 (D). Across the 4 epochs of a trial, the value of α_{CS} was subject to exponential decay ($\alpha_{CS}(1 - 0.10)^2$) and was reset at the start of each trial. The terminal values of α_{CS} were used in Eqs. 1–3, while the (within-trial) decaying values were used in Eqs. 4 and 5. Adapted from: Iliescu, A.F., Dwyer, D.M., & Honey, R.C. (2020). Individual differences in the nature of conditioned behavior across a conditioned stimulus: Development and application of a model. *Journal of Experimental Psychology: Animal Learning and Cognition* (in press).

proven resistant to alternative associative models (see Honey et al., 2020a, 2020b). While HeiDI was developed in the context of findings from a model system (i.e., sign- and goal-tracking in rats), it has very broad explanatory power, only some of which derives from its adaptation of the pooled error term proposed by Rescorla and Wagner (1972).

The analysis offered by HeiDI for individual differences in Pavlovian conditioned responding rests on a single learning process – the development of reciprocal associations between the CS ad US – the behavioral sequelae of which are influenced by the relative perceived saliences of the CS (i.e., α_{CS}) and the US (i.e., β_{US}; as given by V_{CS-US}). However, others have explicitly argued that sign-tracking and goal-tracking, in particular, originate in two process of learning: with sign-tracking reflecting the fact that a lever CS gains incentive salience through its association with the US, while goal-tracking reflects the predictive value of the lever CS (see Berridge & Robinson, 2016; see also, Flagel, Akil, & Robinson, 2009; Lesaint, Sigaud, Flagel, Robinson, & Khamassi, 2014). Similarly, it has been suggested that while sign-tracking reflects a stimulus-response association (i.e., a lever-response association), goal-tracking reflects a stimulus-stimulus association (i.e., lever-outcome association; see Iliescu et al., 2018; Patitucci et al., 2016). These less parsimonious, dual-process accounts predict some of the results that are consistent with HeiDI, particularly the fact that sign-tracking is more resistant to extinction than is goal-tracking (e.g., Iliescu et al., 2018), and the related fact that a partial reinforcement schedule (or uncertainty) maintains higher levels of sign-tracking than does a continuous reinforcement schedule (see Robinson, Anselme, Fischer, & Berridge, 2014). However, these accounts are rather less consistent with other aspects of the behavioral results that we have summarized here, which provide support for HeiDI. For example, they provide no very clear basis for the observation
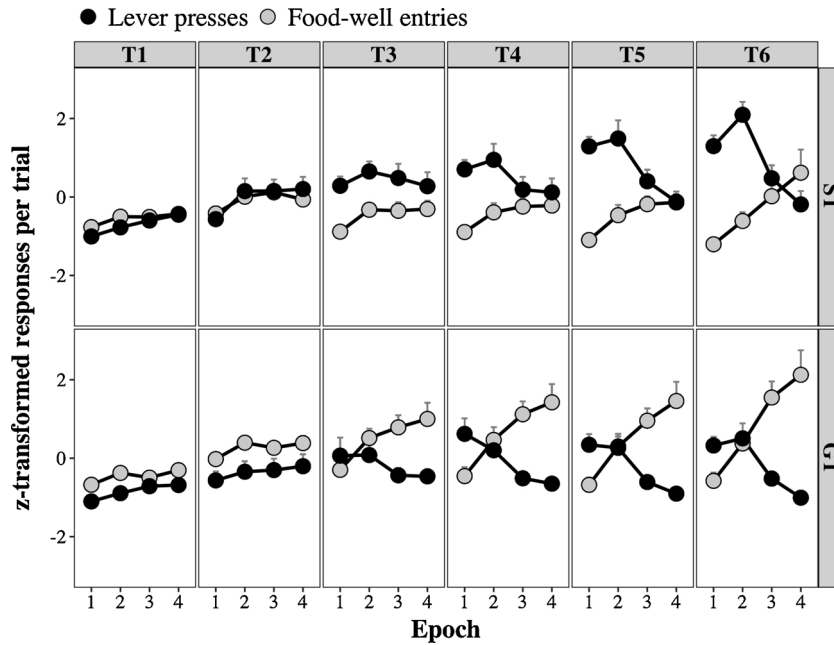
Fig. 6. Mean (+SEM) z-transformed lever presses (black symbols) and food-well entries (grey symbols) during reinforced lever presentations in groups ST (upper panels) and GT (lower panels). The results from the 10-s lever presentations are broken down into 4 successive 2.5-s epochs for each of the 6 blocks of training (T1-T6). Adapted from: Iliescu, A.F., Dwyer, D.M., & Honey, R.C. (2020). Individual differences in the nature of conditioned behavior across a conditioned stimulus: Development and application of a model. *Journal of Experimental Psychology*: *Animal Learning and Cognition* (in press).

that the sign-tracking phenotype is most evident at the outset of a CS while the goal-tracking phenotype is most evident at the end of the stimulus (see Iliescu et al., 2020). Moreover, and as we have already noted, HeiDI provides an analysis of a very broad range of group-level phenomena that are simply beyond the scope of models developed in the context of one model system: Autoshaping in the rat.[1]

If we accept the analysis provided by HeiDI for variation in the vigor and form of Pavlovian conditioned behavior, then a natural question is what are the origins of differences in the perceived salience of the CS (i.e., α_{CS}) and US (i.e., β_{US})? The simple answer to this question is that we do not know. But, one issue that might well constrain our capacity to address this question at behavioral, computation and neural levels is the fact that the sign-tracking and goal-tracking responses are (necessarily) quite different in a standard rat autoshaping procedure: A lever enables a set of responses that is simply very different from the set of responses directed towards the food well. This fact makes it easy to identify the two forms of responding, but it does not allow their ready comparison. To circumvent this issue, it would be advantageous to develop procedures where the sign-tracking and goal-tracking responses are measured (more) equivalently. For example, one could have adjacent food wells in which the (internal) illumination of one food well signals the impending arrival of food in the other food well. In this way, the measures of sign- and goal-tracking would be much more similar: Sign-tracking would be evident as entering the illuminated food well and goal-tracking by entering the well in which food was delivered. It remains to be seen whether this approach is a viable one: Will it yield individual differences in the form of conditioned responding that are evident in the conventional rat autoshaping procedure.

To summarize: The development of HeiDI was inspired by the marked quantitative but especially qualitative individual differences in the form of conditioned responding observed in a Pavlovian conditioning procedure. These differences have not been dealt with by extant general process models of Pavlovian conditioning, for which associative strength (or some other unitary construct) is assumed to have a monotonic relationship to conditioned behavior (e.g., Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972; see also, Gallistel & Gibbon, 2000; Stout & Miller, 2007). It should be noted that the rat autoshaping procedure in which the qualitative differences are observed is not one that is routine: While the appetitive US is standard, the lever CS is not. However, many of the behavioral phenomena that provide converging support for HeiDI from this procedure have clear counterparts in other conditioning procedures (see Honey et al., 2020a, 2020b). Under these conditions, it seems entirely appropriate that HeiDI is formulated as a general process model, providing a relatively simple analysis for an extensive array of findings.

---

[1] It should be apparent, however, that there are other potential sources of variation within HeiDI that could also contribute to differences in the form of responding: The pre-existing differences in the connections between the CS, US and r1-r6; and the obvious potential for these connections to change (see Fig. 2; Honey, Dwyer, & Iliescu, 2020).

Author note

We dedicate this article to our dear friend and unique colleague, William J. Macken (Bill), who passed in February 2020. Correspondence about this article should be addressed to: R.C. Honey; email: honey@cardiff.ac.uk.

Author statement

All three authors contributed to the ideas presented in this manuscript, and to its preparation for publication.

References

Ahrens, A. M., Singer, B. F., Fitzpatrick, C. J., Morrow, J. D., & Robinson, T. E. (2016). Rats that sign-track are resistant to Pavlovian but not instrumental extinction. *Behavioural Brain Research, 296*, 418–430.

Anselme, P., Robinson, M. J., & Berridge, K. C. (2012). Reward uncertainty enhances incentive salience attribution as sign-tracking. *Behavioural Brain Research, 238*, 53–61.

Asratian, E. A. (1952). On the physiology of temporary connections. *Proceedings of the XV Conference of Higher Nervous Activity*.

Asratian, E. A. (1965). *Compensatory adaptations, reflex activity, and the brain.* Oxford: Pergamon Press.

Berridge, K. C., & Robinson, T. E. (2016). Liking, wanting, and the incentive-sensitization theory of addictions. *The American Psychologist, 71*, 670–679.

Boakes, R. A. (1977). Performance on learning to associate a stimulus with positive reinforcement. In H. Davis, & H. M. B. Hurwitz (Eds.). *Operant-pavlovian interactions* (pp. 67–97). NJ: Lawrence Erlbaum Associates.

Brogden, W. J. (1939). Sensory pre-conditioning. *Journal of Experimental Psychology, 25*, 323–332.

Byrom, N. C., & Murphy, R. A. (2018). Individual differences are more than a Gene × Environment interaction: The role of learning. *Journal of Experimental Psychology: Animal Learning and Cognition, 44*, 36–55.

Derman, R. C., Schneider, K., Juarez, S., & Delamater, A. R. (2018). Sign-tracking is an expectancy-mediated behavior that relies on prediction error mechanisms. *Learning & Memory,* 550–563.

Dwyer, D. M. (2012). Licking and liking: The assessment of hedonic responses in rodents. *The Quarterly Journal of Experimental Psychology, 65*, 371–394.

Fitzpatrick, C. J., Gopalakrishnan, S., Cogan, E. S., Yager, L. M., Meyer, P. J., Lovic, V., et al. (2013). Variation in the form of Pavlovian conditioned approach behavior among outbred male Sprague-Dawley rats from different vendors and colonies: Sign-tracking vs. goal-tracking. *PloS One.* https://doi.org/10.1371/journal.pone.0075042.

Flagel, S. B., Akil, H., & Robinson, T. E. (2009). Individual differences in the attribution of incentive salience to reward-related cues: Implications for addiction. *Neuropharmacology, 56*, 139–148.

Gallistel, C. R., & Gibbon, J. (2000). Time, rate and conditioning. *Psychological Review, 107*, 289–344.

Goremezano, I., & Tait, R. W. (1976). The Pavlovian analysis of instrumental conditioning. *Integrative Psychological and Behavioral Science, 11*, 37–55.

Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review, 87*, 1–51.

Hearst, E., & Jenkins, H. M. (1974). *Sign tracking: The stimulus-reinforcer relation and directed action.* Austin: TX: Psychonomic Society.

Hebb, D. O. (1949). *The Organization of Behavior.* New York: Wiley & Sons.

Holland, P. C. (1977). Conditioned stimulus as a determinant of the form of the Pavlovian conditioned response. *Journal of Experimental Psychology: Animal Behavior Processes, 3*, 77–104.

Holland, P. C. (1984). Origins of behavior in Pavlovian conditioning. *The Psychology of Learning and Motivation, 18*, 129–174.

Honey, R. C., Dwyer, D. M., & Iliescu, A. F. (2020a). HeiDI: A model for Pavlovian learning and performance with reciprocal associations. *Psychological Review* (in press).

Honey, R. C., Dwyer, D. M., & Iliescu, A. F. (2020b). Elaboration of a model of Pavlovian learning and performance: HeiDI. *Journal of Experimental Psychology: Animal Learning and Cognition* (in press).

Honey, R. C., Good, M., & Manser, K. L. (1998). Negative priming in associative learning: Evidence from a serial-habituation procedure. *Journal of Experimental Psychology: Animal Behavior Processes, 24*, 229–237.

Honey, R. C., Watt, A., & Good, M. (1998). Hippocampal lesions disrupt an associative mismatch process. *Journal of Neuroscience, 18*, 2226–2230.

Hull, C. L. (1943). *Principles of behavior.* NY: Appleton-Century-Crofts.

Iliescu, A. F., Dwyer, D. M., & Honey, R. C. (2020). Individual differences in the nature of conditioned behavior across a conditioned stimulus: Development and application of a model. *Journal of Experimental Psychology: Animal Learning and Cognition* (in press).

Iliescu, A. F., Hall, J., Wilkinson, L., Dwyer, D. M., & Honey, R. C. (2018). The nature of phenotypic variation in Pavlovian conditioning. *Journal of Experimental Psychology: Animal Learning and Cognition, 44*, 358–369.

Kimmel, H. D. (1977). Notes from "Pavlov's wednesdays": Sensory preconditioning. *The American Journal of Psychology, 90*, 319–321.

Konorski, J. (1948). *Conditioned reflexes and neuron organization.* Cambridge: Cambridge University Press.

Lee, B., Gentry, R. N., Bissonette, G. B., Herman, R. J., Mallon, J. J., Bryden, D. W., et al. (2018). Manipulating the revision of reward value during the intertrial interval increases sign tracking and dopamine release. *PLoS Biology, 16*, Article e2004015.

Lesaint, F., Sigaud, O., Flagel, S. B., Robinson, T. E., & Khamassi, M. (2014). Modelling individual differences in the form of Pavlovian conditioned approach responses: A dual learning systems approach with factored representations. *PLoS Computational Biology, 10* e1003466.

Mackintosh, N. J. (1974). *The psychology of animal learning.* London: Academic Press.

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

Matzel, L. D., Han, Y. R., Grossman, H., Karnik, M. S., Patel, D., Scott, N., et al. (2003). Individual differences in the expression of a "general" learning ability in mice. *Journal of Neuroscience, 23*, 6423–6433.

Morrison, S. E., Bamkole, M. A., & Nicola, S. M. (2015). Sign tracking, but not goal tracking, is resistant to outcome devaluation. *Frontiers in Neuroscience, 9*(December), 1–12.

Patitucci, E., Nelson, A. J. D., Dwyer, D. M., & Honey, R. C. (2016). The origins of individual differences in how learning is expressed in rats: A general-process perspective. *Journal of Experimental Psychology: Animal Learning and Cognition, 42*, 313–324.

Pavlov, I. P. (1927). *Conditioned Reflexes: An investigation of the physiological activity of the cerebral cortex.* London: Oxford University Press.

Pavlov, I. P. (1932). The reply of a physiologist to psychologists. *Psychological Review, 39*, 91–127.

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

Rescorla, R. A., & Cunningham, C. L. (1978). Within-compound flavor associations. *Journal of Experimental Psychology: Animal Behavior Processes, 4*, 267–275.

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: Appleton-Century-Crofts.

Robinson, M. J. F., Anselme, P., Fischer, A. M., & Berridge, K. C. (2014). Initial uncertainty in Pavlovian reward prediction persistently elevates incentive salience and extends sign-tracking to normally unattractive cues. *Behavioural Brain Research, 266*, 119–130.

Spence, K. W. (1937). The differential response in animals to stimuli varying within a single dimension. *Psychological Review, 44*, 430–444.

Staddon, J. E. (2005). Interval timing: Memory, not a clock. *Trends in Cognitive Sciences, 9*, 312–314.

Staddon, J. E. R., & Higa, J. J. (1999). Time and memory: Towards a pacemaker-free theory of interval timing. *Journal of the Experimental Analysis of Behavior, 71*, 215–251.

Stout, S. C., & Miller, R. R. (2007). Sometimes-competing retrieval (SOCR): A formalization of the comparator hypothesis. *Psychological Review, 114*, 759–783.

Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear, & R. R. Miller (Eds.). *Information processing in animals: Memory mechanisms* (pp. 5–148). Hillsdale, NJ: Erlbaum.

Wagner, A. R., & Rescorla, R. A. (1921). Inhibition in Pavlovian conditioning: Applications of a theory. In R. A. Boakese, & M. S. Halliday (Eds.). *Inhibition and learning* (pp. 301–336). New York: Academic Press.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology, 76*, 171–180.

Watson, J. B. (1924). *Behaviorism.* New York, NY: W.W. Norton & Company Inc.

Ward-Robinson, J., & Hall, G. (1996). Backward sensory preconditioning. *Journal of Experimental Psychology: Animal Behavior Processes, 22*, 395–404.