

The Rescorla-Wagner Model, Simplified

W. J. Wilson
Albion College

In 1972, Rescorla and Wagner proposed a mathematical model to explain the amount of learning that occurs on each trial of Pavlovian learning. The model recognized two important things:

1. Learning will occur if what happens on the trial does not match the expectation of the organism, and
2. The expectation on any given trial is based on the predictive value of all of the stimuli present.

Point 1 above says that SURPRISE is important: if you are surprised by a US that you do not expect, you will learn.

Point 2 says that what you expect on any given trial depends on what you have already learned. If you have no experience with a given CS, then it predicts nothing and you expect nothing; if the US occurs, you are surprised. You learn a lot about the CS' prediction of the US. If you have had many past experiences where the CS occurs then the US follows, then you have learned that the CS means the US is coming, and when the US comes you are not terribly surprised. You learn little more about the CS' prediction of the US.

Acquisition

The section above pretty much sums up what happens in acquisition. Early on, when the CS tells you little so you are surprised by the US, you learn a lot on each pairing of CS & US: the acquisition curve goes up quickly. Later in acquisition, the CS tells you a lot, so you learn

only a little on each trial. Eventually you stop learning more about the CS when the CS tells you with certainty that the US will come.

The Equation

$$\Delta V = \alpha\beta(\lambda - \Sigma V)$$

This is the Rescorla-Wagner equation. It specifies that the amount of learning (the change Δ in the predictive value of a stimulus V) depends on the amount of surprise (the difference between what actually happens, λ , and what you expect, ΣV). By convention, λ is usually set to a value of 1 when the US is present, and 0 when it is absent. A value other than 1 might be used if you want to model a larger or smaller US.

The other two terms, α and β , relate to the salience of the CS and the speed of learning for a given US, respectively. As far as Rescorla and Wagner are concerned, these parameters affect the rate of learning, but neither of them changes during learning; in most cases we can ignore α and β and focus solely on surprise to determine the extent to which learning will occur.

Blocking

One of the most important contributions made by the R-W model is that it predicts Blocking and Unblocking. Blocking occurs when a novel stimulus (because it is novel it has no predictive value) is presented together with a well-established CS (whose predictive value

is essentially equal to λ , that is, 1). Neither the well-established CS nor the novel stimulus will change their predictive value because no surprise occurs: the two stimuli together totally predict the US because their V s sum to λ .

$$\begin{aligned}\lambda - \Sigma V &= 1 - (1 + 0) \\ &= 1 - 1 \\ &= 0\end{aligned}$$

No surprise means no learning.

Now, if when the novel stimulus is added, the size of the US is increased, the predictive values of both the original CS and of the novel stimulus will increase, because surprise occurs: the two stimuli together do not predict the larger US; their V s sum to less than the now larger λ .

$$\begin{aligned}\lambda - \Sigma V &= 2 - (1 + 0) \\ &= 2 - 1 \\ &= 1\end{aligned}$$

and according to the R-W equation, the predictive values of the CS and the novel stimulus will both increase:

$$\begin{aligned}\Delta V &= \alpha\beta(\lambda - \Sigma V) \\ \Delta V_{CS} &= .2(2 - (1 + 0)) \\ &= .2(2 - 1) \\ &= .2(1) \\ &= .2 \\ \Delta V_{novel} &= .2(2 - (1 + 0)) \\ &= .2(2 - 1) \\ &= .2(1) \\ &= .2\end{aligned}$$

Because the equation specifies the *change* in the V , for the next trials the V s will be

$$\begin{aligned}V_{CS} &= V_{CS}OLD + \Delta V_{CS} \\ &= 1 + .2 \\ &= 1.2 \\ V_{Novel} &= V_{Novel}OLD + \Delta V_{Novel} \\ &= 0 + .2 \\ &= .2\end{aligned}$$

and on that next trial the learning will proceed as follow:

$$\begin{aligned}\Delta V &= \alpha\beta(\lambda - \Sigma V) \\ \Delta V_{CS} &= .2(2 - (1.2 + .2)) \\ &= .2(2 - 1.4) \\ &= .2(.6) \\ &= .12 \\ \Delta V_{novel} &= .2(2 - (1.2 + .2)) \\ &= .2(2 - 1.4) \\ &= .2(.6) \\ &= .12\end{aligned}$$

this results in V s on the next trial of

$$\begin{aligned}V_{CS} &= V_{CS}OLD + \Delta V_{CS} \\ &= 1.2 + .12 \\ &= 1.32 \\ V_{Novel} &= V_{Novel}OLD + \Delta V_{Novel} \\ &= .2 + .12 \\ &= .32\end{aligned}$$

The increase in the V s will continue until there is no more surprise, that is, until $\Sigma V = \lambda$. This will occur when $V_{CS} = 1.5$ and $V_{Novel} = .5$.

Now that the math is out of the way, think about it this way: In Blocking, when the novel stimulus is introduced and the US occurs, there is no surprise, because the CS predicts the US. In Unblocking when the novel stimulus is introduced *and* the US is larger there *is* surprise, because the CS predict US, not **US**. Therefore, learning occurs. Surprise leads to learning.

Conditioned Inhibition

If a novel stimulus is presented along with a well-established CS and the US does not come then surprise occurs.

$$\begin{aligned}\lambda - \Sigma V &= 0 - (1 + 0) \\ &= -1\end{aligned}$$

The negative surprise means that both stimuli must lose predictive value, and because the novel stimulus is starting at $V = 0$ its V becomes negative; this novel stimulus becomes a conditioned inhibitor (CI). Repeated presentations of CS and CI together with no US, and CS with US, will result in the CS returning to a $V = 1$, serving as a perfect predictor of the US, and the CI reaching a $V = -1$, becoming a perfect predictor of the absence of the US.

Protection From Extinction

The Rescorla-Wagner Model makes some unexpected predictions, and many of them have been demonstrated in real learning situations. One example of this is protection from extinction. Normally a well-established CS presented in the absence of a US will undergo extinction; that is, it will produce smaller and smaller CRs until there is no longer a CR at all. R-W explains this extinction as a loss in predictive value (a decrease in the size of V until it eventually reaches 0); more on this later. If that

same CS is presented without the US but accompanied by a well-established Conditioned Inhibitor (CI), that is, a stimulus that predicts the absence of a US (in R-W terms, a stimulus with a predictive value of -1) then R-W predicts that the CS will *not* undergo extinction (its V will not decrease in size).

Think about this in terms of surprise. In extinction, when the CS is presented and the US does not occur, you are surprised; the CS predicted the US and it did not appear. This is negative surprise ($\lambda - \Sigma V = -1$) so the change in the V for the CS is negative: it gets smaller. In protection from extinction, although the CS predicts the US, the CI predicts the absence of the US, and so there is no surprise ($\lambda - \Sigma V = 0 - (1 + -1) = 0 - 0 = 0$); in the absence of surprise there is no learning, so the V s do not change.

Overexpectation

Another unexpected prediction that came out of the Rescorla-Wagner model is that if there are two CSs, each by itself fully predicting a US (that is, each with a $V = 1$), and they are presented together, they will predict a larger US ($\Sigma V = 1 + 1 = 2$). If these two CSs together are followed by the original US, there will be surprise: you expected **US**, but you got US. The surprise is negative, so the V s for both CSs must decline. So the CSs lose predictive value even though they are still paired with the same US that they had come to predict.

Even stranger is the result when these two CSs plus a novel stimulus are presented together and followed by the original US. Together they all predict a larger US, but what occurs is the same old US, so there is surprise. The surprise is negative, so the V s for all three stimuli must decline. The two CSs go from V s of 1 to values less than 1, but the novel stimulus goes from a $V = 0$ to a V that is neg-

ative; in other words this novel stimulus becomes inhibitory even though it was presented along with the US. It acts like a conditioned inhibitor.

Problems

Extinction Issues

Extinction poses some problems for the R-W model. Specifically, the model deals adequately with the reduction in the size of the CR that occurs when a well-established CS is presented without the US, but the model does not predict spontaneous recovery or rapid reacquisition.

Spontaneous Recovery. If a CR has diminished as a result of presentations of a CS alone, the CR will recur when that CS is presented after a rest period. R-W explains extinction as a reduction in the predictive value V of the CS; there is nothing in the model that allows V to become larger after a rest period. If the CR is gone after extinction because $V = 0$ and the CS no longer predicts the US, a rest period cannot make the CS predict the US again.

Rapid Reacquisition. CRs to a CS that has undergone extinction will be re-established quickly once that CS is again followed by the US; this reacquisition will occur more rapidly than initial acquisition of CRs to that CS, or than acquisition of CRs to a novel CS. Because R-W suggests that extinction drives V back down to 0, that is, it causes the CS to have no predictive value, reacquisition should proceed at the same rate as it did originally.

Latent Inhibition

Latent inhibition is a well-established phenomenon in classical conditioning. A stimulus that has been presented alone prior to its pairing with a US will be slower to acquire CRs

than will a stimulus that was not presented by itself. According to R-W, the presentation of the stimulus by itself should do nothing to its predictive value, as there is no surprise. The stimulus predicts nothing, and nothing occurs, so there is no surprise and thus no learning.

$$\begin{aligned}\lambda - \Sigma V &= 0 - 0 \\ &= 0\end{aligned}$$

However, when the stimulus is later followed by a US, acquisition is slower, as if the V was driven below 0 by the prior experience with the stimulus. The Rescorla-Wagner model suggests no change to the predictive value V of a stimulus if what follows it is not surprising, and it is not especially surprising if nothing occurs following a novel stimulus.

Conclusion

The Rescorla-Wagner model does a great job of explaining many important phenomena of classical conditioning, and even predicts some unexpected results. However, it fails to model some very basic phenomena such as spontaneous recovery, rapid reacquisition, and latent inhibition. It stimulated much research, drove the theoretical analysis of Pavlovian conditioning forward, and ultimately led to the development of more complete approaches.

References

- 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*, pp. 64–99. Appleton-Century-Crofts.