

Institutt for psykologi

Eksamensoppgave i PSYPRO4412 – anvendt og klinisk kognitiv psykologi

Faglig kontakt under eksamen:

Tlf.:

Eksamensdato:

Eksamenstid (fra-til): 09:00-14:00

Hjelpemiddelkode/Tillatte hjelpemidler:

Målform/språk:

Antall sider (uten forside):

Antall sider vedlegg:

Informasjon om trykking av eksamensoppgave

Originalen er:

1-sidig 2-sidig

sort/hvit farger

skal ha flervalgskjema

Kontrollert av:

Dato

Sign

Langsvarsoppgaver – velg to av de tre oppgavene. Hver oppgave teller 40%.

1. Hva er Bayes teorem, hva er signaldetekteringsteori, hva er deres forbindelse, og når er de nyttige?
2. Hvor pålitelig er minnet, og hvilke mekanismer vet du som kunne produsere falske minner? (engelsk: How reliable is memory, and what mechanisms do you know that could produce false memories?)
3. How would you design a behavior modification programme? What theoretical considerations guide your design?

Kortsvarsoppgaver – velg en av de to oppgavene. Oppgaven teller 20%

1. What is the difference between actor-critic learning and Q-learning, and why does that matter?
2. What is the role of prediction error in reconsolidation?

Long questions (Choose two of three, each counts as 40% of the final grade)

Question 1: What is Bayes' theorem, what is signal detection theory, what is their connection, and when are they useful?

Hva er Bayes teorem, hva er signaldetekteringsteori, hva er deres forbindelse, og når er de nyttige?

Bayes' theorem combines previous knowledge, in the form of prior probabilities, with new evidence, in the form of detection rate and false alarm rate (the likelihoods), to produce updated knowledge, the posterior probabilities. An intuitive way to make this Bayesian calculation is by the tables advocated by Gigerenzer. When there are only two possible states of the world (for example, someone is either ill or healthy), and evidence either for or against each of them, that becomes a 2 x 2 table.

	positive	negative	
ill	85	15	100
healthy	495	9405	9900
			10000
p(ill positive test) =	0,146552		
p(healthy positive test)	= 0,853448		

$$\frac{p(\text{ill}|\text{positive test})}{p(\text{healthy}|\text{positive test})} = \text{posterior odds} = 0,171717$$

Base rate =	1,00 %
Detection rate =	85,00 %
False alarm rate =	5,00 %

Pick a base rate, detection rate false alarm rate, and a population size, calculate as Gigerenzer recommends.

In the example above, base rate = 0.01, or 1%, detection rate = 0.85 or 85%, and false alarm rate = 0.05 or 5%. Assume a population of 10000 people who divide exactly according to the probabilities. Then 100 out of those 10000 are ill, and 9900 are healthy. Of the 100 ill people, 85 are being picked up by the medical test, 15 have been missed. To calculate the probability of being ill given that the test says so, we must also know the number of false alarms. That is 5% of the people who are not ill (at least not with what is being tested for), so 5% of 9900, which is 495. The probability of being ill is then $85/(85+495) = 85/580$. That is a good enough answer. It is not required that the fraction is reduced or that a number in decimal notation is offered.

It is also possible to calculate using the odds form of Bayes' theorem, where posterior odds = prior odds * likelihood ratio. With a base rate of 1%, the prior odds of being ill are 1 to 99. The likelihood ratio is detection rate/false alarm rate = $85/5 = 17/1$. The posterior odds are then $17/99 = 0.171717\dots$

Bayes' theorem takes the detection rate and false alarm rate as given, as if they had been carved in stone and handed down by the deity. In cases where the world tidily divides into categorically different states, signal detection theory can tell us something about those likelihoods. Looking up z-scores, I find that if whatever the test measures is normally distributed, with equal standard deviations, but different averages in ill and healthy people, then a detection rate of 85% implies that the decision threshold lies 1.04 standard deviations below the average measurement of whatever it is (for example, antibodies for a virus binding to something in blood) ill people. A false alarm rate of 5% means that the decision threshold lies 1.64 standard deviations above the average measurement for healthy people. So the averages are 2.68 standard deviations apart. That distance is a measure of discriminability, and is called d' . If we decide that missing 15% of ill people is too high, that we can afford to miss at most 1% of ill people, we can make the decision criterion more liberal, meaning we accept a lower measurement as evidence of being ill. To get a miss rate of only 1%, the decision threshold must be 2.33 standard deviations below the average for ill people. The averages are only 2.68 standard deviations apart, so that puts the decision threshold only 0.35 standard deviations above the average for healthy people. In the table used in this course, the closest match is 0.36 standard deviations, which is associated with a false alarm rate of 36%. By reducing the misses, we have increased false alarms. Moving a decision criterion trades one error off against another. To see how that affects the posterior probability of someone actually being ill, given that the same test is applied with the new decision threshold, we can insert the new likelihoods into the table.

	positive	negative	
ill	99	1	100

healthy	3564	6336	9900
			10000
$p(\text{ill} \text{positive test}) =$	0,027027		
$p(\text{healthy} \text{positive test}) =$	0,972973		
$p(\text{ill} \text{positive test}) /$			
$p(\text{healthy} \text{positive test})$	0,027778		
= posterior odds =			

Base rate =	1,00 %	Pick a base rate, detection rate false alarm rate, and a population size, calculate as Gigerenzer recommends.
Detection rate =	99,00 %	
False alarm rate =	36,00 %	

The base rate is low. We gained only 14 extra correct diagnoses from the shift in decision criterion, from 85 to 99. And because the base rate is low, going from 5% of 9900 to 36% of 9900 took the *number* of false alarms from 495 to 3564, an increase of 3069. Consequently, the probability of actually being ill given a positive test result with the new, laxer decision threshold is only $99/3663 = 0.027$.

Calculating in odds form, the prior odds are still 1/99, but the likelihood ratio is now 99/36, giving us posterior odds of $1/36 = 0.027778$.

Correctly detecting most of the ill people means that a lot of healthy people got falsely diagnosed. That is not a problem if the consequence of that diagnosis is a better test, or a cheap treatment with few side effects. But if the decision based on the test is an expensive treatment, then that money is not available for other treatments, which might help more people with worse diseases, but cheaper treatments. Or if the decision based on the test is a medical treatment with severe side effects, is 3069 extra people suffering those side effects a good trade for treating 14 more people? If the side effects include death, it would be good to estimate whether more people might die from being unnecessarily treated than the extra 14 who are saved.

The shift in decision criterion would not increase the total number of false alarms so much if there were fewer healthy people, and the shift would pick up a larger number of ill people who would otherwise be missed. Repeat the calculations for a base rate of 20%. First, for the original, relatively demanding decision criterion that reduced false alarms more than misses.

	positive	negative	
ill	1700	300	2000
healthy	400	7600	8000
			10000
$p(\text{ill} \text{positive test}) =$	0,809524		
$p(\text{healthy} \text{positive test}) =$	0,190476		
$p(\text{ill} \text{positive test}) /$			
$p(\text{healthy} \text{positive test})$	4,25		
= posterior odds =			

Base rate =	20,00 %	Pick a base rate, detection rate false alarm rate, and a population size, calculate as Gigerenzer recommends.
-------------	---------	---

Detection rate = 85,00 %
 False alarm rate = 5,00 %

Out of 10000 people, 2000 are ill, the test picks up 1700 of them, and misses 300. The 5% false alarms out of 8000 are only 400, instead of 495 in the original example. The posterior probability of being ill given a positive test result is 0.81.

In odds form, prior odds = 1/4. The likelihood ratio is still 85/2 = 17/1. The posterior odds are then 1/4 * 17/1 = 17/4 = 4.25.

Now repeat for the laxer decision criterion.

	positive	negative	
ill	1980	20	2000
healthy	2880	5120	8000
			10000

$p(\text{ill}|\text{positive test}) = 0,407407$
 $p(\text{healthy}|\text{positive test}) = 0,592593$
 $\frac{p(\text{ill}|\text{positive test})}{p(\text{healthy}|\text{positive test})} = \text{posterior odds} = 0,6875$

Base rate = 20,00 %
 Detection rate = 99,00 %
 False alarm rate = 36,00 %

Pick a base rate, detection rate false alarm rate, and a population size, calculate as Gigerenzer recommends.

99% of 2000 ill people is 1980. Adopting the laxer decision threshold picked up an additional 280 people over the 1700 with the stricter threshold. The cost in terms of the *number* of false alarms is a growth from 400 with the strict threshold to 2880 with the laxer threshold, and increase of 2480. So with the 20% base rate, a shift in criterion gave us 280 extra ill people correctly diagnosed at the cost of an extra 2480 people wrongly diagnosed. The *same shift* in decision threshold when base rate was 1% gave an extra 14 people correctly diagnosed for an extra 3069 people falsely diagnosed.

The decision threshold, and its associated detection rate and false alarm rate, are not carved in stone and handed down by divine fiat. Where that threshold should be depends on an evaluation of the costs of the mistakes that are traded off against each other, the benefits of the correct decisions traded off against each other, and on the base rate. Bayes' theorem doesn't deal with that. Signal detection theory does, for cases where the world can be divided into neat, mutually exclusive categories, or where it approximates that simplicity closely enough.

This example has been phrased in terms of medical diagnosis. It is good if a student knows that the same principle can be applied to other situations, for example how should locking up criminals who harm others be balanced against the chance of locking up innocent people? How should the chance of not giving welfare support to someone who needs it be balanced against the chance of paying money to someone who plays the system? How should the chance of seducing someone be balanced against

the chance of committing sexual harassment? And so on. Bayes' theorem can be applied to a wide range of decisions, but the optimal placement of the decision threshold depends on each individual's subjective evaluation of the possible consequences. Different evaluations leads to different opinions on where the decision threshold should be. Students may even discuss how signal detection theory could explain persecutory delusions.

Bonus points for a student who realises that things get more complicated when the world is not so tidy, as for example with autism, which does not form a categorically separate population. Scores on the tree core traits of autism are *unimodally* distributed, and any calculation of detection rates and false alarm rates gets a lot more complicated, and is well beyond the scope of this course.

Question 2: How reliable is memory, and what mechanisms do you know that could produce false memories? Hvor pålitelig er minnet, og hvilke mekanismer vet du som kunne produsere falske minner?

Sensorveiledning: Geraerts et al. divided memories of childhood sexual abuse into continuously remembered (there was never a time when the person reporting abuse had forgotten the abuse) and discontinuous, and they divided the discontinuous memories into those remembered within or outside of therapy. Geraerts et al. do not mention whether the therapy in question was specifically recovered memory therapy or a broader range of therapies. In this study, *none* of the memories recovered in therapy was corroborated by independent evidence, while memories recovered out of therapy were as likely to corroborated as continuous memories (those that had never been forgotten). The best predictor of the existence of corroborating evidence was being surprised by the memory. Memories that were corroborated also tended to be recovered suddenly, rather than gradually.

Shaw and Porter interviewed subjects three times, asking them both about an emotionally salient event that (according to the subjects' parents) they really had experienced, and about a fictitious crime. Shaw and Porter were able to persuade 70% of subjects that they had committed a crime serious enough that police had gotten involved. Their interviews used tactics that had elicited false confessions in previous research, such as presenting false evidence, social pressure, building rapport, suggestive retrieval techniques including guided imagery, leaving pauses for the subjects to fill, and implying that experimenters knew further details.

The finding that corroborated discontinuous memories were recovered more suddenly and were more surprising than uncorroborated (and therefore more likely false) memories suggests that false memories are built up slowly. There are three candidate mechanisms within the penum. One is reconsolidation. Hupbach et al demonstrated that episodic memory can be altered through reconsolidation. They showed people a list of objects on a Monday. On Wednesday, the reconsolidation group was asked to remember what they had done on Monday. Then they were shown a second set of objects. One control group saw the same second set of objects, but without being reminded of the previous set. Those two groups returned on Friday, joining the second control group. The experimental group encoded more Wednesday objects along with the Monday objects of which they had been reminded than the control group without the reminder. The Monday & Friday only group provided a baseline of spontaneous intrusions without exposure to the Wednesday objects. The intrusions were asymmetric. The experimental group did not include Monday items in their recall of the Wednesday items. The increased

intrusions were of Wednesday items seen at a time when the Monday items were re-encoded after the reminder. Reconsolidation thus provides a mechanism by which memories may be altered, rather than merely new memories being established that may compete with the old ones. That matters in the context of false memories. Depending on the circumstances in which a false memory is established, the original information may be overwritten.

An alternative method of altering an existing memory is to change the probabilistic relationship between a conditioned and an unconditioned stimulus gradually. The theory behind this approach is that a sudden change in the relationship between conditioned and unconditioned stimulus acts like a context change and triggers new learning, represented in a separate, parallel association. Gradual change in the relationship targets the original association. It is unclear whether this would generalise from conditioning to episodic or semantic memory.

A sudden change in the CS-US relationship is more likely to generate new, parallel associations. In the case of conditioning, there would be no explicit information regarding the source of the information, and it would be impossible to tell directly which is the original memory.

The pensum in this course does not contain any studies that attempt to establish how much each of these mechanisms may contribute to any particular false memory. All that matters here is that students understand the distinct mechanisms, that two of them modify an original memory, while one establishes a new memory that is stored in parallel with the old, and that papers in the pensum only provide evidence for reconsolidation affecting episodic memory.

Question 3: How would you design a behavior modification programme? What theoretical considerations guide your design?

Sensorveiledning: The basic problem with inducing behavior change is relapse, and one important cause of relapse is that behavior modification training often does not affect the original learning that causes undesirable behavior, but causes new learning. That means there are competing and conflicting pieces of information in memory. The new learning is often context-specific, and relapse may occur when returning to the context of original learning. That would be an example of ABA renewal, where A is the original context and B is the context in which the learning occurs that is supposed to change behavior. Students may choose to define and explain AAB, ABA and ABC renewal as well as spontaneous recovery, reinstatement, rapid reacquisition and resurgence, or they may choose just to mention these problems and focus on how to address them. There are several approaches.

One possibility is to accept that new learning is context-specific, but try to make it as effective as possible, and to get around context specificity to some degree by training in multiple contexts. Further, if extinction proceeds very gradually, for example in extinguishing fear conditioning by starting with the least scary situations, there is the risk that reduced fear is based mostly on the non-associative process of habituation. If instead mildly and quite scary situations are mixed, then prediction error is kept high enough to increase learning rate and speed up new learning. It also helps if subjects are encouraged to make very specific predictions so that it is clear when those predictions fail. Prediction errors may also be increased by first extinguishing responses to several conditioned stimuli separately, then combining those stimuli, by removing safety signals (when treating phobias, cognitive treatments that reduce anxiety would also reduce prediction error and should therefore only be given *after*

extinction training), and by occasional CS-US pairings. Generalisation of the new learning is encouraged by varying the conditioned stimuli and the contexts in which they are presented.

An alternative view is that a sudden change in the relationship between conditioned and unconditioned stimulus acts like a context change and that this is what triggers new learning. If instead the probabilistic relationship is changed more gradually (though presumably not so gradually to risk relying on habituation alone), then the original association is modified. This avoids the context specificity of new inhibitory learning and so avoids renewal and reinstatement. This approach seems to contradict the idea of maximizing prediction error.

Finally, reconsolidation may be employed. Within a time window more than 10 minutes but less than 6 hours after retrieval, a memory may be labile and thus modifiable. The memory may then be modified through extinction training within that time window, or a range pharmacological interventions, most commonly administration of cortisol. As in the case of gradual change of the CS-US relationship, targeting the original memory avoids the effects of context specificity such as renewal, spontaneous recovery and resurgence.

Short questions (Choose one of two, each answer counts as 20% of the final grade)

Question 1: What is the difference between actor-critic learning and Q-learning, and why does that matter?

Sensorveiledning: In actor-critic (model-free, stimulus-response, habitual) cognition, you choose whatever response has been best rewarded in the past in this situation, while this stimulus was present.

For example, if you learn to find your way around by actor-critic learning, you just know that in the presence of *this* stimulus (this place), you choose the response of going in *this* direction, without knowing where that will take you, what you will find at the end of this chain of stimulus-response associations, or how long it takes to get there. Or you might habitually prepare a food that you don't like any longer because the last time you ate it, you became violently ill. Stimulus-response learning does not represent the outcome, and so it is not immediately sensitive to changes in the value of the outcome. Only new (and slow) learning can change the response.

In Q learning (model-based, response-outcome or stimulus-(response-outcome) or goal-directed cognition), experience is compiled into a (possibly hierarchical) generative model of the world—a mechanistic, causal understanding of the causes and consequences of actions and events. When faced with a particular situation, this model can be searched, and the quality of various behaviours deduced—even if they have never been tried or experienced. As this involves somehow simulating or inferring future possibilities, it can have high computational costs.

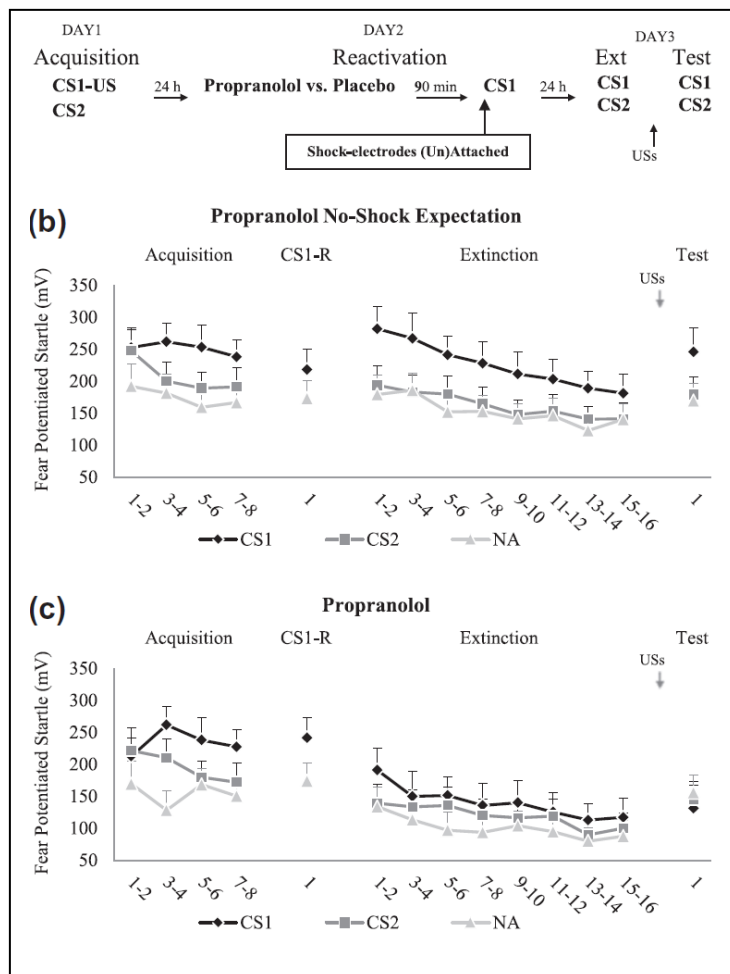
For example, if you learn to find your way around by associating a stimulus with a response-outcome association, meaning you represent the *value* or *quality* of the outcome, you can plan your path. If you are in the presence of *this* stimulus (you are *here*), you can say that if you perform *this* response (go in *this* direction) you will experience *this* outcome (end up in *this* place). Then you can take this imagined outcome as the stimulus for a new stimulus-(response-outcome) association, and so simulate a whole chain of actions.

One reason why the difference matters is that Q-learning, which represents the outcome, is immediately sensitive to changes in the value of the outcome. Therefore persuading client that an outcome is not desirable should quickly change behavior. The actor-critic learning system, however, will need to be retrained by actively monitoring habits (mindfulness) and substituting good new habits for the old, bad habits until the new habits stick.

Further, specific impairments have distinct outcomes. Impairments in representing positive outcomes in Q-learning is associated with negative symptoms in schizophrenia, while giving excessive weight to positive outcomes is associated with one form of impulsivity.

Question 2: What is the role of prediction error in reconsolidation?

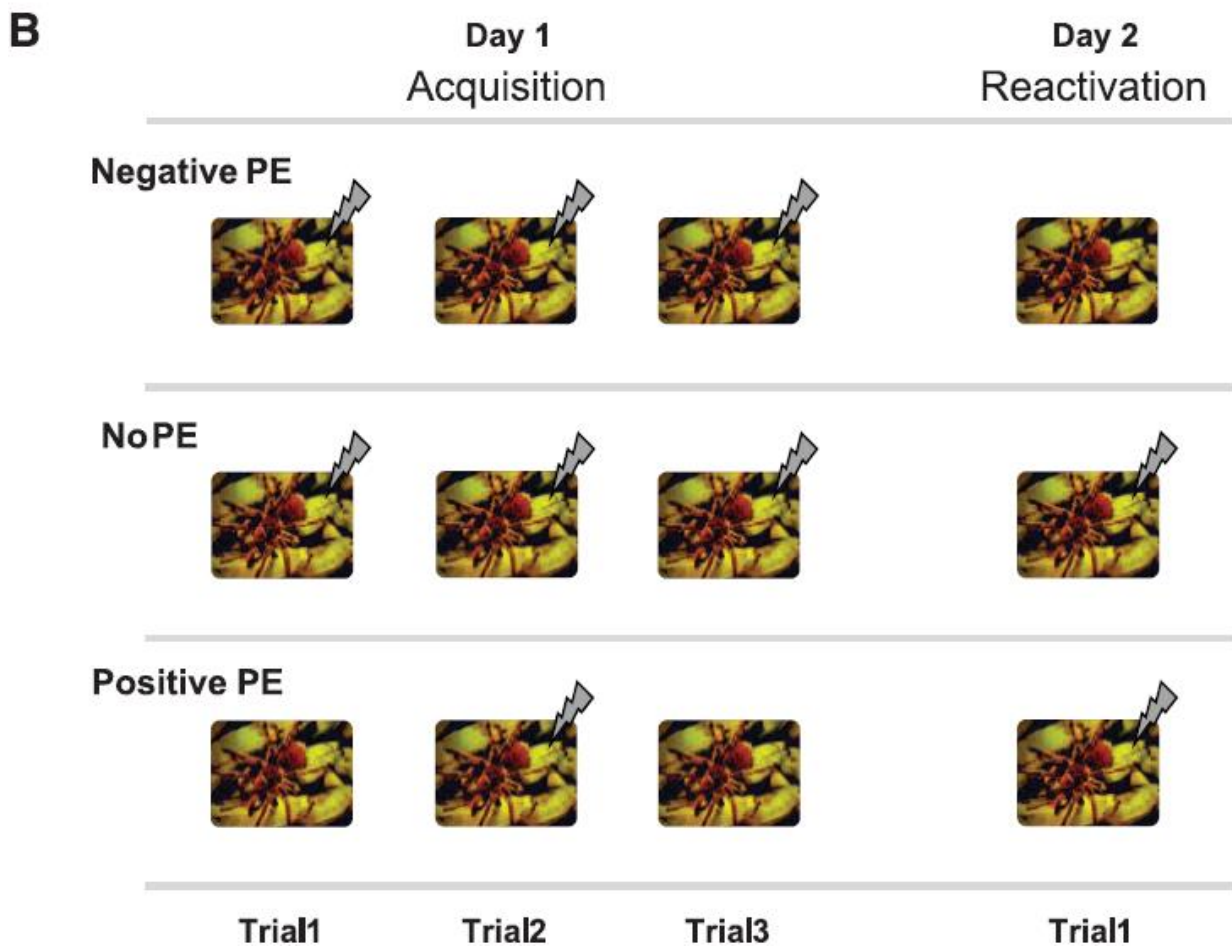

Sensorveiledning: Reconsolidation occurs when being reminded of some information in memory makes that information easier to change during a reconsolidation time window. It occurs in multiple different memory systems. In associative learning, there is evidence that a reminder alone is not enough to trigger reconsolidation, that the association only becomes labile if the conditioned stimulus is not followed by its usual consequence during the reminder. Agren’s review provides a verbal description of Sevenster et al.’s (2012) experimental design. In fear conditioning, administration of propranolol combined with a reminder was previously found to selectively reduce the emotional response to a CS. In Sevenster et al.’s study, this was true only for participants who were connected to the electrodes that had previously delivered electric shock. People who were not connected, and who had no reason to expect any shock, and who therefore experienced no prediction error when shock was absent during the reminder, behaved like the placebo group. Their fear potentiated startle did not decrease from the end of acquisition to the beginning of extinction. More importantly, the group



did not decrease from the end of acquisition to the beginning of extinction. More importantly, the group

without shock expectation and thus no prediction error showed reinstatement after being shocked again, which indicates that the original association still existed, and had only been masked by an inhibitory association created during extinction training.

A second study by Sevenster et al. (2013) is not described in Agren’s review, but students may remember it from the lecture. Initial fear conditioning involved three presentations of a spider picture, with either one or three shocks in those three trials. A no prediction error group received a shock during the reminder trial, which was expected given the 100% reinforcement during acquisition. Negative prediction error came from the omission of an expected shock, positive prediction error came from a shock when past reinforcement history (spider picture followed by shock only one time out of three) indicated that no shock was more likely.



Although explicit expectation of shock depended on its past frequency, fear potentiated startle, the measure of emotional response, depended only on whether there was prediction error, not whether it

was positive or negative. In the presence of propranolol, either prediction error led to an immediate decrease of fear potentiated startle after the reminder trial at the beginning of extinction, and after extinction training no reinstatement after being exposed to shock.

There is good evidence for prediction error being necessary for a reminder to make an association labile. It is not clear what prediction error even means in episodic memory. There, a reminder of a previously learned list can make the memory of that list vulnerable to intrusion errors from a newly learned list, consistent with the idea that the remembered list is being re-encoded while the new list is being learned, and so new list items are added to the memory of the old list. But what would be the prediction error in this situation? So the relevance of prediction error has been demonstrated when reconsolidating associative learning, so far specifically fear conditioning, but it is unclear whether prediction error is relevant to reconsolidation in all memory systems.