UNIVERSITY OF AMSTERDAM

Avoidance Learning

A.M. Krypotos

Summary

The objective of this thesis was two-fold. First, we aimed to progress our understanding regarding the psychological mechanisms involved in avoidance learning.

Second, we attempted to underline the relevance of the insights gained by laboratory research for clinical practice.

For our objective, we completed an extensive literature review on avoidance learning, we conducted a series of experimental studies on avoidance tendencies and risk avoidance, and we presented a novel mathematical model for analysing approach-avoidance reaction time data. We end this dissertation by summarizing and discussing the main findings of the preceding chapters. In addition, we expand on the potential clinical implications of our results and suggest avenues for future research.

0.1 On Avoidance Learning

In Chapter 2, we reviewed traditional and contemporary models of avoidance learning, highlighting the theoretical propositions with the strongest experimental support. We ended our review by collating a set of principles underlying avoidance learning that are applicable to both experimental and clinical settings.

In general, the suggested principles fit modern theoretical approaches of avoidance learning (e.g., the Expectancy model; Lovibond, 2006), in that avoidance responses are mainly performed in order to prevent the encounter with an expected aversive outcome. At the same time, they incorporate the more traditional view that some avoidance responses are acquired more readily than others, due to the innate tendency of an organism to emit some types of responses (e.g., running) rather than others under a state of fear (Bolles, 1970). Lastly, the suggested principles differentiate between the acquisition of avoidance tendencies and overt avoidance, a distinction that fits dual-process models of behaviour (Strack & Deutsch, 2004).

Below we give an overview of the avoidance learning principles as described in Chapter 2. We also summarize the first experiment of Chapter 3, which provided support for some of our theoretical claims.

Summary of Avoidance Learning Principles

We argue (see Chapter 2) that the first step in the acquisition of an avoidance reaction is the pairing of a conditioned stimulus (CS+) with an aversive outcome (US), such that the presentation of the CS will evoke fear responses even when the US is omitted (i.e., Pavlovian conditioning). In partial contrast to traditional avoidance learning theories (e.g., Mowrer's two factor-theory) as well as contemporary cognitive reformulations (e.g., the Cognitive or the Expectancy model), according to which avoidance learning requires the operation of both Pavlovian and instrumental processes, we suggest that Pavlovian associations are suficient for evoking avoidance tendencies without the operation of instrumental processes. This suggestion is based on emotion theory according to which emotions, such as fear, are in essence action tendencies (Frijda, 1988, 2010; Lang, 1985). In other words, under a state of fear, someone will have the tendency to avoid the perceived threat (Beckers, Krypotos, Boddez, Efiting, & Kindt, 2013). As such, we assume that avoidance tendencies can be acquired by mere Pavlovian association, although these tendencies need not to be translated to overt action (Chapter 2).

We provided empirical support for the acquisition of avoidance tendencies via mere Pavlovian association in the first experiment of Chapter 3. In that experiment, our participants initially underwent a differential fear conditioning procedure, during which one neutral cue (e.g., a cube; CS+) was always paired with shock whereas another neutral cue (e.g., a cylinder; CS−) was never paired with shock. Then, our participants completed a symbolic approach-avoidance reaction task (AAT) during which they had to move a manikin figure towards and away from both CSs. Crucially, the shock electrodes were detached during the execution of the reaction time task, leaving little room for the operation of instrumental processes. Our data confirmed the expected pattern of responses, with participants being faster to avoid the CS+ and approach the CS− than the reversed. These findings suggest that in response to cues that elicit fear, avoidance tendencies are expressed, without the need for any instrumental reinforcement.

As already mentioned, acquired avoidance tendencies need not be translated into overt avoidance (Chapter 2). We suggest that a crucial factor deciding whether those tendencies will be expressed or not is the availability of cognitive resources during the perception of a potential danger in the environment. As such, we essentially treat avoidance tendencies as the output of an impulsive system similarly as described in dual-process models of behaviour (Strack & Deutsch, 2004). We propose that in case cognitive resources are depleted (e.g., due to high arousal or drug use), the acquired tendencies will readily be translated into overt behaviour. In other cases, however, those tendencies might be restrained by the operation of a reflective system, which executes cognitively demanding decision-making processes. In case avoidance tendencies are expressed, we assume they will take the form of an evolutionary relevant avoidance response (e.g., running; Species-Specific Defence Responses; SSDRs). As such, we partially agree with Bolles' (1970, 1971) proposition that under a state of fear, an organism is ready to emit evolutionary relevant avoidance responses without shaping by instrumental processes. On the other hand, we believe that instrumental learning does play a crucial role in maintaining the emitted SSDRs or in the learning of non-SSDRs. Specifically, we argue that if an SSDR leads to US omission, this SSDR will be negatively reinforced and performed again in similar threat situations. In the opposite case, non-SSDRs can be acquired, again via negative reinforcement stemming from US-omission. Of importance, and in line with Bolles' theory, we argue that non-SSDRs will be acquired after all SSDRs have been eliminated.

An important extension to past and contemporary theoretical models of avoidance learning is that Pavlovian and instrumental learning can be achieved by multiple pathways instead of as a result of direct experience only. Those pathways include vicarious learning (i.e., avoiding a stimulus after observing other people's avoidant reactions towards that stimulus), instructions (i.e., avoiding a stimulus after someone has provided information that this stimulus is dangerous), and symbolic generalization (i.e., generalization of avoidant responses towards stimuli that are symbolically equivalent to a stimulus that has been trained to elicit avoidance) (Dymond, Schlund, Roche, De Houwer, & Freegard, 2012; Rachman, 1977). The acquisition of avoidance via multiple pathways may better explain how maladaptive avoidance is acquired in psychological disorders, where often the onset of symptoms cannot be attributed to a patient's direct negative experience with the phobic object (Rachman, 1977).

Finally, the proposed principles (see Chapter 2) partially align with the Expectancy model in pointing to expectancy of a negative event as a key factor for both Pavlovian and instrumental learning in avoidance. Unlike the Expectancy model, however, we do not currently rule out the involvement of non-propositional processes in avoidance learning.

0.2 Manipulating Avoidance Tendencies

In the first experiment of Chapter 3, we provided evidence that avoidance tendencies can be acquired via mere Pavlovian association. In the second experiment of Chapter 3, we examined whether those tendencies are sensitive to standard Pavlovian manipulations, specifically to fear extinction and renewal. This experimental question is relevant for both theoretical and clinical reasons. Theoretically, sensitivity of avoidance tendencies to Pavlovian extinction and renewal would provide evidence that these responses follow similar rules as other conditioned fear responses (e.g., skin conductance). Clinically, finding ways to attenuate avoidance tendencies and prevent their return could prove helpful for future translational studies, aimed at countering maladaptive avoidance.

In our study, we included two groups of participants. The first group (i.e., ABB) received fear acquisition (see the first experiment of Chapter 3) in Context A (e.g., lights ofi), fear extinction in Context B (e.g., lights on) and a test of avoidance tendencies in Context B. The second group (i.e., ABA) underwent the same experimental procedure, with the crucial difference that avoidance tendencies were tested in the acquisition context. If avoidance tendencies follow the same principles as other conditioned responses (e.g., skin conductance) then it should be expected that participants in the ABB group should exhibit weaker avoidance tendencies compared to the ABA group. This was confirmed by the data, showing that after successful acquisition and extinction, the ABA group showed stronger avoidance tendencies compared to the ABB group.

These results may have both theoretical and clinical implications. In terms of theory, avoidance tendencies seem to follow the same learning principles as other commonly measured conditioned fear responses. The sensitivity of avoidance tendencies to fear extinction also differentiates them from overt avoidance responses that are, in general, resistant to extinction, at least in absence of response prevention (Solomon, Kamin, & Wynne, 1953). This differentiation provides another reason to distinguish between avoidance tendencies and overt avoidance (see Chapter 2). It may also suggest that different intervention strategies perhaps be used for attenuating maladaptive avoidance tendencies and behaviour.

The clinical implications of our results follow from the fact that fear extinction and renewal are regarded as laboratory models of exposure therapy and fear relapse, respectively (Bouton, 2000, 2002; Hermans, Craske, Mineka, & Lovibond, 2006). Our findings seem to suggest that although extinction-like therapies could be useful for attenuating avoidance tendencies, they do not prevent the return of those responses. Therefore, alternative procedures should be employed for improving the short and long-term effects of extinction-like therapies on avoidance tendencies.

In that direction, we conducted a study (Chapter 4) in which we combined fear extinction with a task targeting action tendencies, namely a training AAT. In such a task, participants have to mainly approach one type of stimulus and avoid another type of stimulus. As a result, participants will typically exhibit attenuated avoidance tendencies to the first type of stimulus and attenuated approach tendencies towards the second type of stimulus. As noted in our introduction, there are some encouraging experimental findings showing that training AATs could result in

the modification of avoidant responding (e.g., Taylor & Amir, 2012). Therefore, we investigated whether the addition of a training AAT to an extinction procedure attenuated avoidance tendencies and/or prevented their return. Of importance, and in extension to our previous experiments (see Chapter 3), we also tested whether the modification of avoidance tendencies had generalized effects on subjective (e.g., expectancy of a US) and physiological fear responses (i.e., skin conductance and eye-blink startle).

Our results suggested that the addition of a training AAT to extinction led to weak modification of avoidance tendencies: between-group differences in avoidance tendencies arose, with participants in the experimental group tending to avoid the CS+ and approach the CS− after practising under the according contingencies and participants in the control group showing the opposite pattern. However, no between-group differences arose for any other measure of conditioned fear, be it subjective (e.g., US-expectancies) or physiological (e.g., skin conductance).

Given that experimental evidence shows only weak co-variation between different indices of emotion (Mauss, Levenson, McCarter, Wilhelm, & Gross, 2005; Mauss & Robinson, 2009), it need not be surprising that changes in avoidance tendencies were not accompanied by changes in other indices of conditioned fear. However, the fact that changes were observed in one out of three indices only makes it questionable that any change in avoidance tendencies will persist over time. As such, future studies should either attempt to strengthen the effects of training AATs or find alternative strategies that afiect fear more comprehensively (see \Limitations and Future Directions" for some suggestions).

0.3 Decomposing Performance on AATs

In Chapters 3, 4 and 6, we used AATs for assessing action tendencies. Such tasks are routinely employed in various branches of psychology for assessing behavioural urges towards stimuli with different hedonic value such as valenced words (De Houwer, Thomas, & Baeyens, 2001), pictures of spiders (Rinck & Becker, 2007), or alcoholic drinks (Wiers, Rinck, Dictus, & Van den Wildenberg, 2009).

Similarly to other response time tasks, statistical analyses of AATs often involve the comparison of response latencies across different experimental conditions (e.g., approach spiders versus avoid spiders) using repeated measures Analysis of Variance (ANOVA). Despite the commonness of this statistical approach, traditional reaction time analyses suffer from a series of shortcomings, such as the collapsing of reaction time distributions or the disregard of speed-accuracy trade-off (Heathcote, Popiel, & Mewhort, 1991; Wagenmakers, 2009). In order to overcome these limitations, we proposed a more accurate and insightful way to analyze AAT data in

Chapter 5. The advantages of using our Bayesian Hierarchical Drift Diffusion Model (Ratclifi & McKoon, 2008; Wiecki, 2013) for analysing AATs include: a) the consideration of the entire response time distribution in the data analysis, b) the accounting for the speed-accuracy trade off, c) the adequate consideration of individual differences, and d) a clear connection of the resulting parameters with underlying psychological processes (e.g., speed of information accumulation).

In order to demonstrate the strengths of our analytic approach, we fitted our model to two data sets; the first one (see Chapter 3, Experiment 1) concerned avoidance tendencies relevant for anxiety disorders and phobias , whereas the second data set (Van Gucht, Vansteenwegen, Van den Bergh, & Beckers, 2008; Experiment 1) concerned approach tendencies { mainly of interest in the addiction literature.

For our modelling strategy, we assumed that differences between conditions arise from differences in information accumulation (i.e., drift rate parameter). We explicitly note that the goal of our model was to provide a more comprehensive account of our data (\cognitive psychometrics"; Batchelder, 1998) and as such we refrained from any model comparison.

For both data-set 1 and 2, participants accumulated information faster in the congruent (i.e., approach CS− and avoid CS+ for data-set 1 and the reversed for data-set 2) compared to the incongruent (i.e., approach CS+ and avoid CS− for data-set 1 and the reversed for data set 2) condition. In addition, regarding dataset 2, the described between-condition differences were more pronounced for the ABA compared to the AAA group. Of importance, we were able to reach such conclusions despite the limited number of response times collected (i.e., 8 response times per condition) for each participant. This model feature makes our approach easily applicable to clinical research, where both the number of trials and the number of participants are often small.

Our model assumptions (e.g., that between-condition differences result from differences in drift rate) are based on the structure of a typical AAT trial, in which participants encounter a stimulus with a specific hedonic value (e.g., positive or negative) and simultaneously have to decide whether to approach or avoid it based on a specific stimulus feature (e.g., frame orientation; see Experiment 2 of Chapter 3). Future studies could potentially modify this trial sequence in order to disentangle the role of the diffusion model parameters operating during AATs. For example, participants may initially encounter the conditioned fear cue, which is supposed to evoke the automatic tendency for either approaching or avoiding it. This tendency for either response could be reected in the a-priori bias parameter of the diffusion model. Subsequently, following a set time interval after presentation of the CS, participants may encounter the stimulus feature according

to which they should emit an approach or avoidance response (see Gladwin, Mohr, & Wiers, 2014, for an example in that direction). It could be expected that in presence of this stimulus feature, participants will have to accumulate information for the type of response to be performed, with the rate of information accumulation being reected in the drift rate parameter of the diffusion model. Such a trial sequence, in which the different psychological parameters addressed by the diffusion model are expected to be influenced by the separate trial events, could be useful in disentangling the role of the diffusion model parameters involved in AATs.

0.4 Risk Decision-Making

As mentioned in our introduction, there has been a growing interest across the fields of psychology and economics in risk avoidance, which refers to the tendency to refrain from engaging in risk situations. Of importance, there is good evidence that individual characteristics can predict either exaggerated or restrained risk taking (Zuckerman & Kuhlman, 2000). Past studies, for example, have shown that trait anxiety, a personality characteristic associated with a generalized tendency to react with feelings of worry and apprehension across a range of situations (Endler & Kocovski, 2001; Spielberger, Gorcuch, & Lushene, 1970) is positively correlated with risk avoidance. In Chapter 6, we aimed to explore whether that correlation is modulated by the conditioned threat or safety value of stimuli presented during risk decision-making. Our experimental hypothesis was based on findings showing that individuals with high trait anxiety attend safety and threat information more readily than individuals with low levels of trait anxiety (Derryberry & Reed, 2002; Hainaut, Monfort, & Bolmont, 2006). As such, we hypothesized that safety and threat cues could modulate the correlation between trait anxiety and risk avoidance, with a stronger correlation when risk decision-making is performed in presence of threat cues than in presence of safety cues.

In order to test our hypotheses, we combined a standard fear conditioning task (see Chapters 3 and 4) with a Balloon Analogue Risk Task (BART), which is a reliable measure of risk taking (Lejuez et al., 2002). During the BART, our participants had to make risk decisions in one block of trials while the CS+ (i.e., threat stimulus) was present and in another block of trials during which the CS− (i.e., safe stimulus) was present. Results showed that differential BART performance was correlated with trait anxiety, with stronger correlations between trait anxiety and risk avoidance reported during the CS+ compared to the CS− block.

Our results could have important implications for the risk decision-making literature. Traditionally, decision-making was regarded as a \cold process", free of emotional influence. That idea has long been contradicted by data providing

strong evidence that emotions can influence decision processes (Newell, Lagnado, & Shanks, 2007; Zajonc, 1980). Our data move this evidence one step further as they also suggest that potential insights on risk avoidance can be gained by testing how the interaction of emotions with individual differences factors influences risk decision-making. For example, we have previously referred to studies showing that trait anxiety positively correlates with risk avoidance (Eisenberg, Baron, & Seligman, 1998; Peng, Xiao, Yang, Wu, & Miao, 2014). Our results indicate that this may not always be the case, as the emotional value of the stimuli (e.g., safety cues) presented during risk taking could efiectively nullify the above correlation. Extending on those results, it can be hypothesized that high trait anxious individuals will exhibit weaker risk avoidance if stimuli associated with safety are present in the environment. This proposition could be helpful for real life risk situations such as in police shootings or when making financial decisions over one's household. In such situations, individuals have to make a decision in a risk situation while being potentially influenced by other emotional factors unrelated to the decision at hand. By finding ways to mute the influence of those factors, potentially more beneficial decisions can be reached. However, although the presentation of safety stimuli during risk decision-making could prove helpful in that direction, they do not necessarily result in long-term changes as any reduction in risk avoidance levels due to the presentation of safety cues would be eliminated by removing these stimuli.

0.5 Clinical Implications

Historically, experimental research has provided the basis for the development of some of the most successful interventions for mental disorders (Emmelkamp et al., 2014). Indeed, laboratory procedures often provide deep knowledge of the psychological processes underlying mental disorders, and as such can establish a firm basis for the development of targeted clinical interventions.

Throughout this thesis, we have addressed the clinical implications of our findings, providing concrete suggestions for future interventions. In this section we

attempt to summarize and extend the insights provided in the separate chapters.
We propose that intervention protocols for maladaptive avoidance should target both overt avoidance and avoidance tendencies. Current procedures for maladaptive avoidance, such as those used in behaviour therapy (Beck, 2011), are primarily focused on challenging irrational beliefs about the relation between the phobic stimulus and the probability of a noxious outcome. As already mentioned in our review (see Chapter 2), we believe that such intervention techniques are valuable for fighting maladaptive avoidance. However, the high relapse of avoidance symptomatology after essentially successful treatment, points to the limited long-term effects of these interventions.

A potential explanation for the return of avoidance behaviour is the disregard of impulsive processes during therapeutic interventions. For example, when challenging the irrational thoughts of phobic patients, it is assumed that individuals will always behave in a rule-base manner and always recognize that their excessive avoidance towards the phobic stimulus is not justified. However, it is quite common that individuals will behave in a reexive manner and as such make a behavioural decision in absence of strong cognitive control.

Therefore, we have suggested that tasks aiming on retraining impulsive avoidance tendencies could be added to common therapeutic interventions for maladaptive avoidance. Although training AATs are argued to be successful in that direction, both our experiment (Chapter 4) as well as other studies (e.g., Asnaani, Rinck, Becker, & Hofmann, 2014) have failed to provide solid support for that claim. As such, perhaps alternative tasks or learning procedures may provide more fruitful directions (see \Limitations and Future Directions" section).

Finally, our review stressed the importance of applying therapeutic protocols that combine exposure with response prevention as well as removal of any safety behaviours. By not allowing avoidance responses to be performed (response prevention) individuals can learn that the phobic stimulus is not followed by a noxious outcome. In order to reach that goal, patients should also be prohibited from performing any safety behaviours, as it has been shown that such behaviours often mitigate the effects of exposure therapy (Salkovskis, Clark, Hackmann, Wells, & Gelder, 1999).

0.6 Limitations and Future Directions

We were unable to find a robust way to counter conditioned avoidance tendencies. Although the training AAT efiect seemed to generalize to an alternative task measuring avoidance tendencies, the transfer efiect was weak and did not afiect other emotion indices. Future studies could potentially enhance the effects by training action tendencies in multiple sessions. This is in line with a recent study showing that training AAT used to counter approach tendencies towards alcohol should be administrated six times for optimal results (Eberl et al., 2014).

Future studies may also use alternative methods to attenuate avoidance tendencies and prevent their return. An exciting possibility in that direction is to use experimental procedures that target memory reconsolidation. This proposition is based on the notion that although fear extinction may attenuate fear responses, it does not directly modify the initial fear memory (Bouton, 2002). As such, the initial fear memory is prone to return, something that could explain the relapse of anxiety symptomatology after seemingly successful therapy. However, experimental findings show that a way to modify the initial fear memory is by targeting its reconsolidation. Specifically, it has been argued that after being retrieved, fear memories enter a liable, protein-synthesis dependent state (reconsolidation phase), during which they are sensitive to modification (Alberini, 2005). Recent studies have suggested that retrieved fear memories can be modified by simple behavioural manipulations during reconsolidation (e.g., fear extinction). As a result, fear responses towards threat cues are suppressed (Monfils, Cowansage, Klann, & LeDoux, 2009; Schiller et al., 2009). However, several studies have failed to replicate this memory updating efiect (Chan, Leung, Westbrook, & McNally, 2010; Kindt & Soeter, 2013; Soeter & Kindt, 2011). Still, the fact that fear memories can become labile again and in need of reconsolidation upon retrieval would seem to open up a window of opportunity to modify them. Perhaps an alternative to conducting extinction training during reconsolidation might be to perform approach training, using a training AAT, during reconsolidation. Training AATs are thought to target action tendencies, which are argued to be a key aspect of emotions (Frijda, 1988; Lang, 1985). Perhaps then, changing action tendencies and thereby the emotion elicited by a CS during memory updating could serve the permanent modification of the fear memory.

In Chapter 2 we argued that avoidance can be acquired via other pathways than direct experience. It could be of interest to test whether avoidance can be attenuated via multiple pathways as well. A recent study by Golkar, Selbing, Flygare, Ohman, and Olsson (2013) showed that vicarious extinction, which entails the observation of an individual undergoing a standard fear extinction procedure, resulted in stronger extinction learning and stronger prevention of reinstatement than a standard extinction procedure. Both theoretically and clinically it would be relevant to test whether similar vicarious learning procedures can result in attenuated avoidance tendencies.

In all our experiments (see Chapters 3, 4 and 6), we assessed action tendencies through AATs (De Houwer, Crombez, Baeyens, & Hermans, 2001; Krieglmeyer & Deutsch, 2010). Given the nature of those tasks, in which multiple trials are needed in order to draw valid conclusions, we were able to assess action tendencies only at the end, rather than during conditioning. As such, although we could evaluate how our manipulation influenced action tendencies, we were not able to track action tendencies during each learning phase, as we could for expectancy ratings (see Chapters 3, 4 and 6)or physiological responses (skin conductance and startle reex; see Chapter 4). By testing how avoidance tendencies are influenced by the different CS-(no-)US contingencies, potentially greater insights could be gained in the nature of action tendencies. For example, in the first experiment of Chapter 3, the pattern of results suggested that participants were slower to avoid the CScompared to the other three combinations (i.e., approach and avoid the CS+ and approach the CS−), which were largely similar. As we also point out in Appendix A, it is hard to reliably interpret this pattern of results as it could have arisen in multiple ways (e.g., by participants becoming faster in some conditions while their responses to the other conditions remain the same as prior to the fear acquisition phase). This dificulty in interpreting AAT effects could possibly be surpassed by testing action tendencies during the entire fear acquisition phase and acquiring learning curves for approaching and avoiding each separate stimulus.

Alternative experimental procedures could be used for overcoming the above shortcoming. An interesting possibility in that direction is the assessment of body posture during conditioning. It has recently been found that body posture can serve as an index of approach-avoidance tendencies (Eerland, Guadalupe, Franken, & Zwaan, 2012). Therefore, in a potential future experiment, participants could undergo the different conditioning procedures while their body posture is measured, with potential changes in body posture during the conditioning trials serving as proxies of action tendencies.

The development of a way to measure avoidance tendencies online could also help the investigation of emotion coherence during avoidance learning, which refers to \the coordination, or association, of a person's experiential, behavioural, and physiological responses as the emotion unfolds over time." (Mauss et al., 2005, p. 175). Although from a layman's perspective, emotions are expected to be characterized by coordinate changes across subjective, physiological and behavioural measures, experimental evidence suggests very modest correlations across measures (Mauss et al., 2005; Mauss & Robinson, 2009). A common argument for the lack of correlation is that different emotion indices are measured at different time points, not allowing a moment-to-moment comparison. In terms of fear conditioning, measuring action tendencies concurrently with physiological responses and subjective apprehension measures, could help the moment-to-moment assessment of fear across different fear indices, and as such possibly enable to reliably evaluate co-variation of different fear indices.

Lastly, we argue that future research can benefit from the application of mathematical modelling techniques to behavioural data. We have given an example in that direction in Chapter 5, where we applied a simple mathematical model to AAT data. This model allowed us to quantify evidence regarding different decisionmaking parameters (e.g., speed of information accumulation), something impossible with traditional statistical techniques. Those models could prove valuable also when applied to conditioning data or avoidance learning data (e.g., in Maia, 2010). One of the main advantages of those models is that they can decompose performance in parameters which refer directly to latent psychological processes, such as in the diffusion model in which the reaction time distributions for correct and error responses are decomposed to meaningful variables of decision-making variables. Of importance, apart from behaviour, mathematical modeling can be used to infer the neurobiological causes of the exhibited behaviour (Wiecki, Poland, & Frank, in press). As such, we suggest that the wealth of parameters provided in mathematical models and the clear connection of those parameters with psychological latent variables, could prove valuable for the deeper understanding of avoidance behaviour.

All in all, we have provided a fresh view on how adaptive and maladaptive avoidance is acquired and expressed. We have also translated the insights gained from our work to clinical practice. We hope that the present thesis will serve as further impetus for the revival of avoidance learning research.