Mechanics of contingency-based Cognitive Bias Modification:
pre-existing bias affects potency of active training but not placebo conditions.

Anne-Wil Kruijt^{a*}

Per Carlbring^a

^a Department of Psychology
Stockholm University
SE-106 91 Stockholm
Sweden

Anne-Wil.Kruijt@psychology.su.se

Per.Carlbring@psychology.su.se

*corresponding author: dr. Anne-Wil Kruijt¹

¹. permanent adres: mail@awkruijt.nl

Abstract

Cognitive Bias Modification (CBM) refers to computerized training protocols aimed at modifying automatic information processing patterns associated with various forms of psychopathology. CBM protocols are widely studied as potential new treatments, the scope of application encompassing anxiety disorders, depression, substance abuse, disordered eating, pain perception, and insomnia. CBM often relies on (hidden) contingencies between stimulus content and desired responses. Active training conditions are typically contrasted with control training lacking the contingency. This paper focusses on the wide-spread, and intuitive, notion that pre-existing bias may affect the contingency experienced by an individual engaging in placebo training, which could render the intended placebo condition more potent. Applying probabilistic reasoning, we conclude that pre-existing bias likely does not increase the potency of placebo training. In contrast: lack of pre-existing bias may render an active training condition functionally similar to a placebo condition. In the second part of this paper we review literature with respect to our assumptions, and discuss the (im)possibility of empirically verifying the hypothesis that larger pre-existing bias associates with greater CBM efficacy in light of bias indices typically possessing low internal reliability.

Keywords: cognitive bias modification; task mechanics; placebo; measurement; reliability

A recent exchange of commentaries between colleagues in the field of food-related approach bias modification (CBM-A) caught our eye with an argument that rang familiar. Becker and colleagues forward the observation that, in a recent review, significant effects of food-related CBM-A appear to be restricted to studies using non-standard or no control conditions (Becker, Jostmann, & Holland, 2018). In contrast, they note, no significant effects were observed for studies utilizing a more standard ‘no-contingency’ control training. No-contingency control training conditions are commonly used across the different domains of contingency-based CBM and are known by various names such as ‘50/50 control’, ‘sham-training’, and ‘placebo condition’. Active CBM training conditions are typically designed to modify an individual’s response tendencies by means of a contingency hidden in the trial structure. These contingencies make the task easier to perform if an individual shifts from an undesired to a desired information processing strategy. Consequently, CBM placebo-control conditions are lauded as the perfect placebo condition, being identical to the active training task except for lacking the key feature of a hidden modifying contingency (yet see: Blackwell, Woud, & MacLeod, 2017). In their response, however, Kakoschke and colleagues argue that “*Although Becker et al. state that sham-training ‘should not change response biases in any specific direction’ (p. 3), this will clearly depend upon an individual’s initial level of approach bias toward the target stimulus category (here, unhealthy food). Most people are by definition likely to have a positive approach bias toward appetitive substances (otherwise there would be no point to approach bias re-training). Thus sham-training (with its 50/50 contingency) would logically serve to reduce such an approach bias. This should be particularly the case for more extreme or clinical samples who are likely to have elevated levels of initial approach bias (e.g., alcohol dependent inpatients or chocolate cravers)*” (Kakoschke, Kemps, & Tiggemann, 2018, p. 295). The authors offer examples of consistent result patterns in other CBM-A studies and conclude: “*Thus sham-training may actually*

function as an active condition, rather than the ‘neutral’ condition it is intended to be”

(Kakoschke et al., 2018, p. 295).

Although our background is in another subfield of CBM (Attention Bias Modification: ABM), we are familiar with this argument. The notion that the standard control conditions may act as (low-dose) active treatments, as a result of participant’s pre-existing bias, has occasionally popped up throughout the CBM subfields since nearly a decade. The argument was forwarded, for instance, in the discussion of one of the first larger randomized controlled trials for ABM. Discussing a finding of no differential effects for placebo and active ABM on symptoms in post-traumatic stress disorder patients, Schoorl and colleagues noted that *“Theoretically, this could mean that both treatments were effective, since we cannot be sure that the control condition acted as a real (neutral) placebo. For patients with anxiety disorders who have [bias] towards threat, one could speculate that training at 50/50 contingency may actually be a low-dose version of ABM”* (Schoorl, Putman, & Does, 2012). More recently, the argument featured in what is likely to prove an influential review for the attentional bias field: *“[...] the 50/50 control condition [...] may constitute a diluted ABM intervention for some individuals. Consider anxious participants prone to attend to threat 90% of the time. Exposure to a condition whereby probes follow threat ‘only’ 50% of the trials might diminish their [bias].”* (McNally, 2018, p. 5).

The idea of ‘unintentionally active placebo as a result of pre-existing bias’ was possibly first forwarded, tested, and rejected in a 2009 study on interpretation bias modification (CBM-I; Salemink, van den Hout, & Kindt, 2009). Observing higher post-training state anxiety in their placebo training group, Salemink and colleagues wrote: *“[state anxiety showed] an unexpected increase for the control group. Possibly, the control*

condition may have provoked anxiety, since almost half (38%) of their stories had a negative valence. The control condition might inadvertently have led to negative interpretations. Yet, based on both theoretical arguments as well as new empirical data this seems rather implausible. Theoretically, the argument that a negative bias was induced in combination with the present pattern of equal responses to positive and negative interpretations after control CBM-I, requires the assumption of a positive bias before CBM-I." (Salemink et al., 2009, p. 681/682). Because bias had been assessed post-training only, additional data was gathered from a third group of highly anxious individuals who did not undergo either intended active or placebo CBM-I. Observing that their scores did not differ from the control CBM-I group's scores post-training, the authors concluded: "[...] *control CBM-I condition did not induce a negative interpretive style and [...] functioned as a control condition, as intended.*" (p. 682). Nonetheless, the argument that 50% training may act as an active training kept hanging over the CBM-I domain too. For instance, Chan and colleagues noted in 2015 that "Similarly, *balanced training may not provide a suitable control if it significantly alters participants' usual interpretation style in either a positive or negative direction. Participants with predominantly negative biases may be more affected by the negative items than the 'balanced' amount of positive items. Alternatively, 50% of positive training may be sufficient to inducing positive interpretations. Hence, balanced training does not guarantee a neutral outcome.*" (Chan, Lau, & Reynolds, 2015, p. 1245)

Seeing the argument resurface in the 2018 exchange between Becker, Kakoschke, and colleagues, prompted us to dive into the mechanics common to contingency-based bias modification tasks to get to the bottom of the wide-spread, and intuitive, notion that the *experienced* contingency in a 50/50 control condition may be affected by the individuals pre-existing bias. Doing so, we arrived at the unforeseen conclusion that pre-existing bias affects

the experienced contingency for active rather than placebo training conditions. Opposite from the argument forwarded to date, pre-existing bias cannot increase the potency¹ of a 50/50 placebo condition yet it may render an active training condition functionally similar to a placebo condition. In the next section, we will develop this argument.

The dot-probe training as the simplest example of contingency-based CBM

The shared rationale among the vast majority of CBM training tasks is that, through repeated task trials, participants will adapt automatic tendencies/patterns in their cognitive processing due to a contingency hidden in the task's trial structure. In its most simple form there are two types of trials, which is the case in ABM's most studied procedure based on the dot-probe task (MacLeod, Rutherford, Campbell, Ebsworthy, & Holker, 2002). In each trial of this task, the individual is simultaneously exposed to two stimuli that differ in valence. Depending of the training objective, these can be negative and neutral faces, pictures of alcoholic and non-alcoholic beverages, healthy and unhealthy food, neutral words paired with words related to smoking, threat, pain, insomnia, etc. Although various implementations include neutral-neutral trials (which we will discuss later), the trials of interest always combine stimuli belonging to two separate classes of information. These will be referred to as the harmful and beneficial class stimuli (HCS and BCS). Note that stimuli in either class may be of neutral content/valence: these terms indicate whether the content/meaning of the stimuli in either class are to be considered harmful or beneficial relative to the other stimulus class. The objective of training is usually to shift preferential processing and/or automatic response tendencies away from the HCS and towards the BCS. In the dot-probe task, the display of two stimuli is followed by the appearance of a small response cue in either the location

¹ Throughout this manuscript we use the term 'potency' to indicate the capacity of a training contingency to elicit an effect.

previously taken by either the HCS or the BCS. The participant is instructed to respond to the cue as quickly as possible. The rationale behind the original dot-probe task (designed to measure, rather than modify, attentional bias) is that if an individual has a tendency to orient their focus of attention towards the HCS (e.g. biased attention towards threat), they will be able to respond relatively quickly (relative more often) on the trials in which the response cue appears on the location previously taken by the HCS (MacLeod, Mathews, & Tata, 1986). If, however, the individual responds equally quickly on trials in which the response cue appears in the location previously taken by either the BCS or HCS, this is taken as an indication of no bias. In a traditional dot-probe measurement task, the trial-structure is counterbalanced with respect to the number of trials in which the probe appears on the location previously taken by either kind of stimuli. In a dot-probe ABM training task, a hidden contingency is created by increasing the relative number of trials in which the response cue appears on the location initially taken by the BCS (up to 100%). The rationale is that the individual will adapt their attention orienting behaviour accordingly. This expectation appears to involve the assumption that a participant will have a more positive experience (speeded, easier, or otherwise more satisfactory) on trials in which they orient attention towards (the stimulus that is in) the location that will (subsequently) be taken by the probe. Thus, relative speed or ease of response that is the outcome in a measurement dot-probe is thought to be rewarding in the training variant. It appears that two crucial assumptions are that participants a) experience an easier or speeded response on trials in which their initial attention orienting is towards the probe location and b) experience this as rewarding (i.e. occasionally and anecdotally, some users report quick progression of the task stressful). Historically, it is thought that many of the biases targeted with CBM are automatic, i.e. that they occur in early automatic stages of information processing that are not under direct volitional control, and that therefore the CBM process of “learning”, or adapting to, the contingency should also happen implicitly,

i.e. without notifying participants of the contingency or the bias modifying objective of the training task, and ideally also without participants becoming aware during the training. We will write about CBM's implicitness assumptions quite extensively in the discussion. First, we focus on the likelihood that an individual's (implicit) experience of the hidden training contingency is affected by their pre-existing bias.

Quantifying the modifying effect of pre-existing bias on training potency

While in dot-probe parlance the trials in which the response cue appears on the location of the negative stimulus are typically labelled as congruent (with the bias to be measured or modified), we propose that for our current purposes it is more intuitive to label the two trial types according to congruency with the training objective. For the sake of developing a lexicon that applies across various CBM tasks, we define a term "Intended Training Ratio (incongruent/congruent)" (ITR(in/co)) to denote congruency with the training objective. Table 1 provides a glossary of the various terms and their acronyms used in this paper. For dot-probe tasks the ITR(in/co) is often the opposite of the traditional trial congruency, e.g. the ITR(in/co) for a 100% training task aimed at modifying bias away from HCS (negative/threatening/unhealthy – incongruent with the training objective) and towards the BCS (neutral/positive/healthy – congruent with the training objective) would be 0/100 expressing that 100% of the trials are congruent with the objective of reducing bias towards the HCS. For a placebo control condition, the $ITR(in/co) = 50/50$, which is identical to a dot-probe task used to assess (rather than modify) bias. Pre-existing bias can be expressed as the probability that, for a given trial, a participant's initial response tendency is the desired tendency (DT), which in a dot-probe is attention being oriented towards the BCS, or the undesired tendency (UT) which is attention being oriented towards the HCS.

Acronym	Term	Use
ITR(in/co)	Intended Training Ratio (incongruent/congruent)	ITR(in/co) refers to the task design and indicates the ratio of trials that are congruent/ incongruent with the intended training direction. If $ITR(in/co) = 0/100$, all trials are congruent with the training objective. If $ITR(in/co) = 15/85$, 85% of trials are congruent with the intended direction of training, and if $ITR(in/co) = 50/50$, no training contingency is present.
MER(co/we)	Maximum Experienced Ratio (consolidating / weakening)	The maximum experienced ratio of consolidating and weakening trials is a function of both the researcher-implemented ITR(in/co) and participant's pre-existing response tendencies.
HCS	harmful class stimuli	Stimuli representing the construct or class that is considered relatively negative/to be avoided or ignored as a result of training. May refer to negative, threatening, alcoholic, unhealthy, etc. etc., or neutral stimuli.
BCS	beneficial class stimuli	Stimuli representing the construct or class that is considered relatively positive/to be approached or attended to as a result of training. May refer to positive, neutral, non-alcoholic, healthy, etc. etc., or neutral stimuli.
DT & UT	desired tendency & undesired tendency	The response tendency that the training aims to increase/foster (DT) or decrease/suppress (UT). Could refer to directing of attention/gaze, pulling or pushing, a tendency to stop or go, or to interpret as positive or negative.
TC, TI & TN	training congruent, training incongruent & training neutral	Refers to a trial being congruent, incongruent or neutral with respect to the intended direction of training in a training task.
$P_{(weaken)}$, $P_{(consolidate)}$, $P_{(neutral)}$.		The probability that a given trial exerts a weakening, consolidating, or neutral effect on the participant's initial response tendency.

Table 1: glossary of terms and acronyms

Combining probabilities for response tendency and training congruency, we can compute the probability for each trial to be classified as either consolidating or weakening the individual's automatic response tendency for that trial. For each trial, the probability that the initial response tendency may be consolidated ($P_{(consolidate)}$) is the sum of the probability that the participant had the undesired tendency while the trial is incongruent with the training

objective ($P_{(UT)} * P_{(TI)}$), and the probability that the participant had the desired response tendency for a training congruent trial ($P_{(DT)} * P_{(TC)}$):

$$P(\text{consolidate}) = P(UT) * P(TI) + P(DT) * P(TC)$$

Imagine how a participant with a relatively strong bias towards HCS would tend to orient towards the HCS in 80% of the trials². For this participant $P_{(UT)} = .8$, and $P_{(DT)} = .2$. In a 0/100 ITR(in/co) task (in which all trials are congruent with the training objective), $P_{(TI)} = 0$ and $P_{(TC)} = 1$. Thus, $P_{(\text{consolidate})} = .8 * 0 + .2 * 1 = .2$. Therefore, the participant's initial response tendency will likely be consolidated on 20% of the trials, relative to 80% of trials in which they may experience that they might want to change something. The probability that for any given trial this participant's initial response tendency may be weakened is the sum of the probability that the participant had the desired response tendency on a trial incongruent with the training objective ($.2 * 0 = 0$) and the probability that the participant had the undesired tendency on a training congruent trial ($.8 * 1 = .8$). Thus $P_{(\text{weaken})} = .8$.

$$P(\text{weaken}) = P(DT) * P(TI) + P(UT) * P(TC)$$

Notice that such weakening would likely require the participant to experience (implicitly) that something can be changed in order to be rewarded, and therefore also that the relative speed on the other 20% of trials (the consolidating trials) is experienced as rewarding. Yet, we can't know with certainty whether the individual experiences reward or punishment as a result of the contingency on any of the trials. Therefore, we denote the ratio of consolidating and weakening trials the "Maximum Experienced Ratio (consolidating / weakening)" ($MER(\text{co/we})$). For the participant in our example, the $MER(\text{co/we})$ would be 20/80:

$$MER(\text{co/we}) = P(\text{consolidate}) * 100 / P(\text{weaken}) * 100$$

² We assume that the participant always orients attention towards one of the two stimuli at the moment these disappear (and the response cue appears), and not towards both stimuli simultaneously or anywhere else.

Assessing inertia of the placebo training condition

Now we return to the question of whether pre-existing bias may activate a 50/50 ITR(in/co) task into a more potent training task, which would be indicated by a MER(co/we) deviating from 50/50. For a 50/50 ITR(in/co) task, $P_{(TI)}$ and $P_{(TC)}$ are both .5. For an individual with no pre-existing bias $P_{(UT)}$ and $P_{(DT)}$ also are .5. Filling out the formulae results in $MER(co/we) = 50/50$. Ergo: in absence of bias, a 50/50 ITR(in/co) task is inert as intended.

The ‘50/50 condition possibly not inert in the presence of bias’ argument represents the intuitive notion that individuals with a strong pre-existing bias would receive unequal exposure to the two trial types due to their preferential processing of one of the two trial types. Using the above formulae, we can test this assumption for an individual with absolute bias towards the HCS: their $P_{(UT)} = 1$, and $P_{(DT)} = 0$. If this individual engages in a 50/50 ITR(in/co) task, their $P_{(consolidate)} = 1*.5 + 0*.5 = .5$, their $P_{(weaken)} = .5$, and their $MER(co/we) = 50/50$. Similarly, for a person with a strong beneficial bias causing them to orient towards the BCS in 80% of the trials. Their $P_{(DT)} = .8$ and $P_{(UT)} = .2$. In a 50/50 ITR(in/co) task their $P_{(consolidate)} = .2*.5 + .8*.5 = .5$, their $P_{(weaken)} = .8*.5 + .2*.5 = .5$, resulting in a $MER(co/we)$ of 50/50. Thus it is demonstrated that pre-existing bias does not modify a 50/50 ITR(in/co) task into a training task.

Assessing potency of active training conditions

In our very first example, however, we have seen a $MER(co/we)$ that deviated from the ITR(in/co) as a result of pre-existing bias: in that example the pre-existing bias was 80% towards HCS, which for a 0/100 ITR(in/co) task resulted in a $MER(co/we)$ of 20/80. It turns out that the intuitive notion that pre-existing bias modifies exposure to a hidden training contingency holds true *in the presence of such a contingency*. When pre-existing harmful bias

is absolute and $ITR(in/co)$ is 0/100, the $MER(co/we)$ is 0/100 as intended. If, however, an individual has an absolute pre-existing beneficial bias (away from the HCS and towards the BCS), the $MER(co/we)$ is the inverse of the $ITR(in/co)$. For such an individual and a 0/100 $ITR(in/co)$ task, the $MER(co/we) = 100/0$, i.e. all trials serve to consolidate their pre-existing bias. Similarly, $MER(co/we) = 80/20$ when someone with an absolute *beneficial* bias engages in an 20/80 $ITR(in/co)$ task. More problematically, for someone with a moderate harmful bias orienting towards the HCS in 70% of the trials, the potency of an $ITR(in/co) = 0/100$ training will be reduced to a $MER(co/we)$ of 30/70 rather than the intended 0/100. Figure 1 illustrates how the bias modifying potency ($P_{(weaken)}$) of contingency-based CBM training is related to both the intended training ratio ($ITR(in/co)$) and pre-existing bias.

Figure 1. Probability that a trial weakens an initial response tendency (y-axis) as a function of pre-existing bias (x-axis), and intended training ratio incongruent congruent ($ITR(in/co)$; lines).

Expansion to more complex CBM tasks:

The above demonstration used a ‘two trial types’ dot-probe task as the simplest form of CBM in terms of number of trial types. Yet, most contingency-based CBM procedures employ more than two trial types. In this next section we will explore how the above reasoning applies to several other tasks and variants.

First, we will look at a dot-probe based training with three trial types. One commonly used ABM procedure employs 120 neutral-threat trials plus 40 neutral-neutral trials for a total of 160 trials per session (available from: people.socsci.tau.ac.il/mu/anxietytrauma/research). In the training version of this task, the response cue always appears on the location previously taken by the threatening stimulus, i.e., the ITR(in/co) is 0/100 within the subset of 120 ‘training relevant’ neutral-threat trials. Thus, for an individual with an absolute bias towards threat, the MER(co/we) on the neutral-threat trials will be 0/100. Yet, assuming that the participant first has to become aware (even if only to a limited extent) that there is something that they can change in order to perform faster, it may be more realistic to assume a MER(co/we) encompassing all 160 trials. To do so we modify the formulae by adding a ‘training neutral’ term ($+ .5 * P_{(TN)}$) to both $P(\text{weaken})$ and $P(\text{consolidate})$ and adjust the values of $P_{(TC)}$, $P_{(TI)}$ and $P_{(TN)}$ to .75, 0, and .25 (reflecting the proportion of trials for each trial type). For an individual with absolute bias towards threat the resulting overall $\text{MER}(\text{co/we}) = 12.5/87.5$. Thus, adding neutral-neutral trials induces a discrepancy between ITR(in/co) and (overall) MER(co/we) that will be highest at the most extreme values of pre-existing bias. For an individual with no bias engaging in either the active or placebo version of this task, the MER(co/we) remains 50/50 as intended.

In dot-probe based CBM participants typically receive no other instructions than to respond to the response cue as quickly as possible, and it is hoped that they will (subconsciously) adapt to the contingency, namely the response cue appearing relatively often on the location of either BCS or HCS. In approach bias modification training, however, there is only one stimulus presented in each trial (HCS or BCS) and it is common practice to add a 'foil dimension' which participants are instructed to respond to. For instance, in the seminal procedure for modifying (alcohol) approach tendencies, participants receive training that is intended to make it easier to push a joystick in order to 'push away' pictures of alcoholic beverages (HCS) and/or to pull the joystick to 'bring closer' pictures of non-alcoholic beverages (BCS). Yet participants receive explicit instructions to move the joystick based on the orientation of the pictures: push if landscape and pull if portrait. The rationale is that this foil dimension ensures that differences in response time between the two trial types is driven by automatic/implicit evaluation of the stimulus content (Wiers, Rinck, Kordts, Houben, & Strack, 2010). In addition, the foil dimension is thought to conceal the contingency and the purpose of training, allowing the intended change in approach tendency/bias to take place at a more implicit level (yet see: Van Dessel, De Houwer, & Gast, 2016; Wiers et al., 2010). The foil and target dimensions are mapped within the trial structure so that in a measurement task (and in a 50/50 placebo condition), pictures of alcoholic and non-alcoholic beverages appear equally often in portrait and landscape format. In a maximum training contingency task, non-alcoholic beverages always appear in portrait format (pull), and non-alcoholic beverages in landscape (push). Thus, this type of training has four rather than two trial types (BCS/HCS * portrait/landscape). These four trial types can, however, be divided in two sets of two trial types: one set for HCS and one for BCS. Within either set we can apply the formulae defined above. For the HCS set (trials with alcoholic pictures), the DT is to push, and the UT to pull. For the BCS set this is the other way around (DT is pull and

UT is push). Often, mirrored ITR(in/co) are used for the HCS and BCS trial sets, such that the mapping of portrait/landscape across the trials works out to instruct participants to push all HCS and pull all BCS. Let's fill out the above formulae for a task with an HCS-ITR(in/co) of 0/100, a BCS-ITR(in/co) of 0/100, and a participant with an absolute tendency to pull alcohol and push soda pictures (UT = 1, DT = 0). First the HCS part: $P(\text{weaken}) = 0 * 0 + 1 * 1 = 1$, and the resulting HCS-MER(co/we) = 0/100. Then for the BCS part: $P(\text{weaken}) = 0 * 0 + 1 * 1 = 1$, and BS-MER(co/we) = 0/100. Thus, for a participant with absolute (and complementary) biases towards alcohol and away from soda pictures, MER(co/we) is 0/100 (for both HCS and BCS) as intended. As we have seen before, however, if a participant has no pre-existing biases, MER(co/we)s drop to 50/50.

Yet, it is not inconceivable that a participant has 'asymmetrical' pre-existing biases for HCS and BCS. For instance, a strong tendency to pull pictures of unhealthy food paired with no noticeable preference to push or pull pictures of healthy food. In that case the HCS-MER(co/we) will be 0/100, the BCS-MER(co/we) 50/50, and the overall MER(co/we) = 25/75 instead of the intended 0/100. For an individual with a tendency to pull on 80% of HCS, and to push on 60% of BCS, the HCS-MER(co/we) = 20/80, the BCS-MER(co/we) = 40/60 and the overall MER(co/we) = 30/70 (assuming equal numbers of BCS and HCS trials). Finally, we look at approach bias modification for an individual with a tendency to pull the joystick in 70% of the trials, regardless of the content of the stimulus picture (anecdotal evidence has it that many find joysticks easier to pull than the push). Calculating MER(co/we)s for a task with ITR(in/co) = 0/100 for both HCS and BCS, the resulting HCS-MER(co/we) would be 70/30, the BCS-MER(co/we) = 30/70 and the overall MER(co/we) = 50/50. Thus, if the participants response tendency is largely determined by factors other than the content of the stimulus pictures, the intended training contingency will likely not be

experienced by the participant. This reminds us that MER(co/we) is related to the probability that an (implicit) connection is made between the content of the stimulus pictures and the required response, even if the participant has been explicitly instructed to respond according to the foil dimension (picture orientation, colour of the picture border, etc). In this type of CBM, the reward that could be experienced is not so much ease or speed of responding in itself, but rather the individual knowing what the correct response is and experiencing increased ease in giving the correct response if they manage to overcome a pre-existing tendency to pull the joystick when presented with an HCS in the format for which the explicit instructions tell them to push it.

We propose that the formulae that we provided can be applied to virtually all hidden contingency-based CBM tasks, which apart from the tasks discussed, include training variants of the Posner single cueing tasks, some tasks for interpretation bias modification (CBM-I), and Go/NoGo training if interpreted according to the ‘behaviour-stimulus interaction’-account (see Veling et al., 2017).

For a more interactive demonstration of the here described interactions between pre-existing bias and training contingency, we created an *R* shiny app which can be accessed at <https://awkruijt.shinyapps.io/mechanics-of-CBM-app>. It generates and combines random sequences of participant’s response tendencies and trial types based on user settings of pre-existing bias, ITR(in/co), task type, and numbers of trials. These random sequences are then combined to a simulated trial sequence and the resultant MER(co/we)s are shown.

Discussion

The first part of this paper presented a probabilistic approach to quantify the maximum experienced contingency in contingency-based CBM tasks, as a function of individuals' pre-existing bias and the intended training contingency implemented by the researcher. The conclusions are two-fold: a) the regularly forwarded suggestion that placebo training may act as an active training in the presence of a pre-existing bias is likely incorrect and b) the intuitive notion that pre-existing bias affects exposure to a hidden contingency likely holds true, but only when such a hidden contingency is present, i.e. in active CBM training conditions. This suggests that pre-existing bias may be a prerequisite for successful modification of bias through contingency-based CBM. In the second part of the paper we discuss a) non-contingency based CBM, b) the hidden-contingency assumption, and c) barriers to obtaining empirical evidence of pre-existing bias moderating CBM effects.

Non-contingency-based CBM?

The reader may note that we speak of contingency-based CBM, raising the question if there also exists non-contingency-based CBM. Indeed, at least one form of CBM-I does not rely on a contingency hidden in the trial structure. In imagery-focused (components of) CBM-I (Blackwell et al., 2015; Blackwell & Holmes, 2010), participants receive instruction on how to generate mental imagery (e.g., using first person perspective) and are subsequently exposed to a series of vignettes which resolve in either a positive or a negative manner (usually always positive in an active training condition). The participant's task is to generate mental imagery according to the instructions and for each vignette rate how well they were able to do so. Thus, there exists no contingency in the sense of participants having to display one of two possible responses or behaviours for two or more different trial types with the trial structure being manipulated by the researcher to create a contingency where one tendency is

favourable for (most) trials of one type, and the other tendency for (most) trials of the other type. Word-sentence-association CBM-I, on the other hand, can be understood as a contingency-based CBM (Amir, Bomyea, & Beard, 2010; Beard & Amir, 2008). In this CBM-I form, each trial starts with a cue-word that is either positive or negative, followed by an initially ambiguous vignette which is resolved in the final sentence (sometimes the participant is required to fill in missing letters on the disambiguating word). The participants task is to indicate whether or not the vignette is related to the cue-word. Feedback is provided and manipulated according to a contingency: in an ITR(in/co)=0/100 training. Feedback on trials with a negative cue-word is always positive if the participant indicates that the word is unrelated to the interpretation of the vignette, and always negative if the participant indicates that the vignette and cue-word are related (and vice-versa for positive cue-word trials). Visual search ABM is a CBM form for which it is not immediately obvious whether it is contingency-based (Dandeneau, Baldwin, Baccus, Sakellaropoulo, & Pruessner, 2007; De Voogd, Wiers, Prins, & Salemink, 2014; Kruijt, Putman, & Van der Does, 2013b; Platt, Murphy, & Lau, 2015). When used in CBM context, it typically features just one trial type for which participants are instructed to locate a happy face among negative faces as quickly as possible. The typical control condition is identical except that stimuli are pictures of 7- and 9- petaled flowers. Notice, however, that this type of training can be reframed as an ITR(in/co) = 0/100 version of a ‘find the face that is different from the others’ visual search task, paired with a ‘neutral class stimuli’ version as the control condition. Yet, visual search ABM differs from other CBM forms most crucially in one easily overlooked feature, namely the explicit instructions: find the happy face.

The hidden contingency assumption

The reader may have felt that, throughout this paper, we have been quite mysterious regarding the implicit nature of CBM training, the hidden-ness of the training contingency, and the presumed unawareness of participating individuals. Many CBM researchers (including ourselves) believe that unawareness of training could be relevant for the success of training. Most of us would like these training techniques to amount to more than simply instructing patients to ignore bad information, not in the least because we tend to believe behaviour to be affected by a plethora of implicit processes (Wiers, Teachman, & De Houwer, 2007). In the upcoming section we will discuss the hidden contingency assumption in more detail.

To our knowledge, the overarching defining features of CBM methods have only ever been tentatively formally defined, namely in the introduction to a special section in the journal of abnormal psychology that may be regarded as the formal starting point of the CBM field (Koster, Fox, & MacLeod, 2009)³. Koster and colleagues explain that their selection of studies as examples of CBM work was guided by two key features. The first being that the target of change is a cognitive bias, for which they add that “*Commonly, it has been the case that the targeted bias represents a pattern of processing selectivity that appears to operate automatically, in the sense of proceeding swiftly without intention, and so is not readily amenable to volitional control.*” (p. 3). The second key feature is that “[...] *the method of manipulating the target cognitive bias has not principally relied on instruction but instead has involved extensive practice on a cognitive task designed to encourage and facilitate the desired cognitive change. Thus, for example, this task may incorporate a certain contingency that renders performance easier or more successful to the extent that the desired pattern of*

³ although the very first use of the term ‘Cognitive Bias Modification’ likely dates back to a 2007 paper on imagery rescripting (Holmes, Arntz, & Smucker, 2007)

processing selectivity is acquired” (Koster, Fox, Macleod, 2009, p.3). Although it was not explicitly stated that unawareness of training is required, the assumption that implicit training (i.e. training relying on hidden contingencies) may produce better and/or better generalizing effects became widely held in the CBM field (see for instance: MacLeod & Mathews, 2012). Reports for several CBM studies describe how care is taken not to communicate the aim and/or the contingency to participants and it is not uncommon that post-training awareness interviews or questionnaires are taken as manipulation checks (for instance in: Chan et al., 2015; Kruijt, Putman, & Van der Does, 2013a; Salemink et al., 2009; Schoenmakers, Wiers, Jones, Bruce, & Jansen, 2007). In addition, concealing the contingency is often given as a rationale for using ITR(in/co)s lower than 0/100 (Schumacher, Kemps, & Tiggesmann, 2016; Wells & Beevers, 2010; Wiers et al., 2010), or for testing features such as adding a working-memory load (Clarke et al., 2017). Several studies specifically tested differential effects of implicit versus explicit instructions for dot-probe ABM (Grafton, Mackintosh, Vujic, & MacLeod, 2013; Krebs, Hirsch, & Mathews, 2010; Lazarov, Abend, Seidner, Pine, & Bar-Haim, 2017; Nishiguchi, Takano, & Tanno, 2015). Most of these reported larger (initial) bias change in the explicit-instruction conditions, but also found that these additional gains did not necessarily generalize to other computerized tasks (Krebs et al., 2010; Lazarov et al., 2017; Nishiguchi et al., 2015). One study examined effects on stress reactivity and found that, while bias change in the explicit-instruction condition was comparable to bias change in the hidden-contingency ABM condition, “*providing explicit information about the training contingency, and engaging participants in active practice of the target attentional change, eliminates the desired impact of attentional change on anxiety vulnerability*” (Grafton et al., 2013, p. 185). In a later paper, this finding has been interpreted as “[...] *consistent with the possibility that explicit contingency awareness may contribute to more fragile bias change, while conditions that discourage explicit processing of the contingency may more effectively contribute to bias*

change.” (Clarke et al., 2017, p. 26). As mentioned before, it is common practice in the ABM-A domain to use foil-dimensions to conceal the purpose of training. Nonetheless, there exists discussion on the necessity to make participants focus on ‘irrelevant features’ (such as the orientation of the stimulus picture) rather than the content of the stimuli. In one commentary it has been suggested that tasks relying on implicit and explicit instructions tap into different constructs: *“It could at least theoretically be the case that the kind of avoidance associations as measured by a relevant-feature SRC cause an increased chance of relapse, while alcohol-avoidance associations retrained and assessed by an irrelevant-feature AAT decrease the chance of relapse.”* (Wiers, Gladwin, & Rinck, 2013, p. 1). On the other hand, in papers reporting training awareness-checks (throughout the CBM spectrum), it is not always clear whether non-awareness is believed to be crucial to the procedure’s efficacy or whether precautions and checks are intended to guard the intended (double-)blind nature of the study design.

Another aspect of the unawareness assumption is the question whether we believe that the (emotional/palatable) valence of stimuli in a hidden-contingency task reliably influences initial response tendencies. Paradoxically, the tasks that we use to measure whether bias exist may only function in the manner that we intend them to function if we assume that bias exists. One of us (AWK) once performed 30 funnelled awareness interviews as part of a dot-probe ABM study (Kruijt et al., 2013a). It was observed that a considerable portion of the participants appeared to have paid little attention to the content of the stimulus pictures. When prompted, several did not recognize that they had seen negative and neutral (or positive and neutral) faces during a total of 1000 trials. Moreover, the fourth and final session had consisted of 200 bias training trials, immediately followed, without a warning or break, by 100 bias assessment trials. Because different stimulus sets were used during the training

and the assessment trials, ‘the faces changed’ at the same time that the ITR(in/co) changed from 15/85 to 50/50. When asked if they had noticed anything else changing when the faces changed, a large proportion of participants indicated not to have noticed the change in stimulus sets at all. While, of course, this is a mere anecdotal observation, it may merit some thought. It seems not entirely inconceivable that a participant wishing to speed up their performance on a(n often rather boring and repetitive) training task, may adopt a strategy of ignoring the stimulus content entirely, focussing exclusively on features relevant to their response: the response cue in a dot-probe task, the picture orientation in an approach-bias training task, etc. Of course, this could be taken as being in line with the intention of keeping the contingency hidden, but only to the extent that we believe the stimulus content to be implicitly registered after all, and therefore of bias (in the form of unvolitional preferential processing) to occur. In contrast, the feature-specific attention allocation account (FSAA) proposes that task-irrelevant emotional features may not be processed when participants are focussing on, for instance, non-affective semantic stimulus features. Most FSAA work focusses on priming tasks, yet one study has suggested that emotional Stroop and dot-probe tasks are similarly affected (Everaert, Spruyt, & Houwer, 2013).

In the approach-avoidance domain, data supporting the notion that stimulus content may be more effectively ignored than might be desirable comes from a study comparing two sets of stimulus-response-compatibility (SRC, also known as the mannikin task) and approach-avoidance tasks (Kersbergen, Woud, & Field, 2015). Both approach-avoidance and SRC tasks are regularly used to measure alcohol related action tendencies. In relevant feature versions of each task, participants were instructed to pull/approach alcohol and push/avoid stationery pictures. In irrelevant feature versions, participants were instructed to respond to the orientation of the pictures (tilted left/right). Picture orientation was counterbalanced over

stimulus-content of the pictures, so that these tasks were identical to ITR(in/co) 50/50 tasks. Internal reliability for the relevant feature tasks was moderate-to-high ($r = .70$ and $.58$), but dropped to essentially zero and negative for the irrelevant feature versions ($r = -.05$ and $-.77$). Bias estimates derived from the irrelevant feature versions also did not associate with self-reported drinking behaviour (Kersbergen et al., 2015). A possible explanation for this task-reliability pattern is that the stimulus content simply did not affect response times in the irrelevant feature task versions, thereby reducing the index of alcohol-stationery bias to essentially measurement error.

Only in interpretation bias modification does it seem truly necessary for the participant to engage with the stimulus content in order to advance the task. It is therefore perhaps not surprising that some of the highest contingency awareness ratings have been reported for CBM-I (e.g., 94% in Chan et al., 2015), that no differential effects were observed for implicit and explicit instruction CBM-I training (Mobini et al., 2014), and that contingency awareness appears to be historically more accepted in this CBM subfield (Browning, Holmes, Charles, Cowen, & Harmer, 2012; Holmes, Arntz, & Smucker, 2007; Salemink, van den Hout, & Kindt, 2007).

All in all, it seems that there exist different lines of reasoning, both between and within CBM subfields, on the necessity of keeping participants unaware of the training contingency. When the CBM field was first defined, the implicit nature of training was forwarded as a defining feature of CBM. Yet, some task traditions allow awareness and use explicit instructions (imagery based CBM-I, visual-search ABM). For other tasks, the consensus seems to be that non-awareness is preferable and may even be necessary for generalization of effects to emotional vulnerability (dot-probe ABM), while in the approach

CBM tradition foil response dimensions to conceal the contingency are integral to the standard task designs and it is suggested that these may tap into a different construct than non-foil task versions. The question of (un)awareness falls apart in (at least) two dimensions, namely whether unawareness is required to achieve the intended CBM effects, and whether unawareness is required to achieve full blinding, and reduce demand effects, in (double) blind study designs. Two further awareness-related unresolved questions, specific to evaluation of clinical potential of CBM procedures, are a) whether participants in clinical trials reporting not to have become aware of the contingency/purpose of the training may reflect low ‘treatment credibility’ (Kuckertz et al., 2018; Wiers, Gladwin, Hofmann, Salemink, & Ridderinkhof, 2013), and b) the medical-ethical implications if it would turn out that unawareness is required for training efficacy: we cannot disseminate a treatment for which the mechanism needs to be kept ‘secret’ (Kruijt, 2014). We suggest that, regardless of the reasons to aim for unawareness, our calculations and conclusions apply to any contingency-based CBM procedure in which the presence of a contingency is purposefully hidden.

Empirical evidence of pre-existing bias moderating CBM effects?

When we argue that active CBM may act as placebo in the absence of bias, our argument is purely task-mechanical, based on the probabilistic line of reasoning outlined above, which suggests that a hidden contingency is likely not experienced by participants who do not have a strong pre-existing bias. In this final section we outline a major barrier to empirically verifying the hypothesis that more extreme pre-existing bias results in more effective bias modification, namely low reliability of bias measures in general, and the associated regression towards the mean in specific.

Regression towards the mean refers to the statistical artefact that the observed association between any two measures will be reduced by the involvement of random processes, and that this reduction is most pronounced at the tail-ends of either measure's distribution. In the context of experimental research, regression towards the mean is often understood to be a function of measurement error (or conversely: measurement error is defined as all variance due to random processes). In practice, regression towards the mean causes change over time to be negatively correlated with initial levels. When (a proportion of) the individuals with more extreme pre-existing bias scores, scored more extremely than their "true bias score" as a result of redistribution of (randomly distributed) error variance, those participants are likely to score more moderately on a second assessment, (also) in the absence of actual change in the construct of interest. This mechanism results in spurious correlations between a baseline and a change score involving that same baseline (or any measure correlating with that baseline measure), as well as false positive results when assessing change in subgroups characterized by more extreme baseline scores (Bland & Altman, 1994).

The extent to which regression towards the mean occurs is directly and inversely related to the measure's reliability, defined as the proportion of true variance relative to the total variance ($\text{true variance} / (\text{true} + \text{error variance})$). Low reliability indicates that a relatively large proportion of the variance is due to measurement error, and the larger the proportion of error variance, the more pronounced the effect of 'random error variance (re)distribution' (i.e. regression towards the mean). Low (and unacceptably low) measurement reliability poses an ever more clearly emerging problem across the CBM research field. Specifically: while the component scores for the bias indices that are the main outcome for many CBM tasks often possess outstanding internal reliability (e.g. in the $> .95$

range for response times on either congruent or incongruent trials in the dot probe task, indicating a very low proportion of error variance), the difference score (bias index) often possess minimal to no internal reliability (tentatively indicating that these scores reflect mostly the small amount of (measurement) error variance in response time on one trial type that remains unexplained by the other). Low reliability of difference scores derived from highly correlating component scores is a statistical phenomenon leading to extensive debate since decades (See for instance: Cronbach & Furby, 1970; Edwards, 2001; Rogosa & Willett, 1983; Trafimow, 2015; Williams & Zimmerman, 1977, 1996). Although it can be argued that the problem of low reliability might be not so grave under certain circumstances (for instance when a difference score reflects change in a construct that, in itself, is the construct of interest and is measured with reasonable reliability), the implications are especially poignant for constructs that can be operationalized only as a difference score, as is the case for virtually all cognitive bias indices. Using the dot probe assessed bias as an example once more: the actual measurement (i.e. the response time to either congruent or incongruent trials) often possesses outstanding reliability, yet the average response time for a single trial type is uninformative with respect to the question of whether an individual (or a group) responded relatively fast to one trial type compared to the other (i.e. bias).

The practical implication of low difference/bias score reliability is that participants who, within a sample, attained the highest bias score on a first assessment, may be the lowest scoring participants on a repeated assessment. I.e. the rank ordering of bias scores observed within a sample is largely random. Thus, while the suggestion to test CBM in individuals pre-selected for having a consistent bias observed over multiple baseline assessments has been repeatedly raised (most often at conferences and other ‘debate not in print’), the recurring answer is that assessment of bias is not reliable enough to detect presence of bias at the

individual level (Bar-Haim, 2010; Elfeddali, de Vries, Bolman, Pronk, & Wiers, 2016; Wiers, Gladwin, Hofmann, et al., 2013). Nonetheless, several studies and one meta-analysis reported findings suggestive of pre-existing bias moderating CBM induced change in bias (Eberl et al., 2013; Fox, Zougkou, Ashwin, & Cahill, 2015; Mogoase, David, & Koster, 2014, yet see: Heeren, Mogoase, Philippot, & McNally, 2015), or symptoms (Amir, Taylor, & Donohue, 2011; Fox et al., 2015; Kuckertz et al., 2014; Price et al., 2016). These findings could have been taken as supportive of the hypothesis forwarded in this paper (that pre-existing bias is necessary for CBM training to function as originally intended), where it not the case that the typical low internal reliability of bias indices a) impedes us from reliably identifying participants with higher baseline bias, and b) raises the likelihood that any finding of baseline bias associating with subsequent change in bias is (partly or entirely) due to regression towards the mean.

Discussions on the issue of low reliability for difference scores often reference a phenomenon that is currently most well-known as the ‘reliability paradox’ (Hedge, Powell, & Sumner, 2018): measures with low levels of variance within groups will be characterized by low internal reliability, yet between-group effects will be more readily detected by statistical tests. The classical (non-emotional) Stroop task is an often-cited example of a task for which data tends to show this variance pattern. Large parts of cognitive bias research can be thought of as derived from the classical Stroop task, as it gave rise to the emotional Stroop task, which in turn was the inspiration for the dot-probe task and the subsequent development of ABM. The reliability paradox can easily be misinterpreted to suggest that while low reliability may lead to reduced correlations between measures (such as the correlations reported by Miloff, Savva, & Carlbring, 2015), assessment at group level is unaffected by measure reliability, i.e. that biases may be reliably observed as a heightened average bias

within a group of patients (Staugaard, 2009). Group-level comparisons are, of course, not unaffected by reliability, even though the relationship between reliability and power is somewhat complex and moderated by the within/between distribution of variance (Parsons, Kruijt, & Fox, 2018, fig 1; Zimmerman & Zumbo, 2015). In addition, typical dot-probe data, for instance, is not at all characterized by low within-group variance (nor particularly high between group variance) and therefore does not fit the data-pattern associated with the reliability paradox. Moreover, a recent meta-analysis of baseline dot-probe bias in ABM RCTs enrolling a total of 1005 clinically anxious adults found that the group level baseline attention bias commonly assumed to be associated with anxiety either does not exist or otherwise is not detected by the dot-probe task (Kruijt, Parsons, & Fox, 2018). Thus, it appears that this specific bias cannot be reliably detected at individual nor group level, casting doubt on its existence as well as the wisdom of developing treatments based on this task. A conclusion to which the conclusions drawn in the first part of this manuscript add an additional layer of doubts.

While this section focused mostly on reliability issues described for anxiety-related attention bias (modification), similar reliability issues may be expected to be at play for other bias assessment tasks due to the reliance on difference scores (see also: Brown, Eley, & Broeren, 2014; Kersbergen et al., 2015; Parsons et al., 2018; Reinecke, Becker, & Rinck, 2010). Apart from the wider implications, the practical implication for the hypothesis generated by the probabilistic arguments presented in the first part of this paper (that pre-existing bias is required for CBM procedures to exert effects in the manner that they were originally designed for) is that these cannot be empirically verified for most of the existing CBM-related bias measures because the effect of interest would almost certainly be contaminated with regression towards the mean.

Conclusion

In contingency-based CBM, level of pre-existing bias cannot increase the potency of a 50/50 placebo CBM training. Lack of pre-existing bias, however, may render an intended active training task functionally similar to a placebo training. Regression towards the mean, and the ever more clearly emerging problems with low task-reliability in the CBM domain, form a major barrier to empirically verifying these conclusions.

Role of Funding Sources

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Author contributions

AWK conceived the manuscript idea, coded the illustrating app, and wrote the first manuscript draft. Both authors contributed to, and have approved, the final manuscript.

Declaration of Interest

The authors declare that no conflicts of interest exist.

Acknowledgements

In the course of writing this manuscript, the authors have sought input on their assumptions and ideas from various colleagues. We'd like to thank dr. Inge Kersbergen for her input on the subconscious is better assumption, dr. Marcella Woud for help on delineating different forms of CBM-I, dr Simon Blackwell for input on the preprint, and members of the UWA CARE lab for their persistent insistence that the 'placebo may be active for those with a strong bias'-argument is incorrect.

References

- Amir, N., Bomyea, J., & Beard, C. (2010). The effect of single-session interpretation modification on attention bias in socially anxious individuals. *Journal of Anxiety Disorders, 24*(2), 178–182. <https://doi.org/10.1016/j.janxdis.2009.10.005>
- Amir, N., Taylor, C. T., & Donohue, M. C. (2011). Predictors of response to an attention modification program in generalized social phobia. *Journal of Consulting and Clinical Psychology, 79*(4), 533–541. <https://doi.org/10.1037/a0023808>
- Bar-Haim, Y. (2010). Research review: Attention bias modification (ABM): A novel treatment for anxiety disorders. *Journal of Child Psychology and Psychiatry, 51*(8), 859–870. <https://doi.org/10.1111/j.1469-7610.2010.02251.x>
- Beard, C., & Amir, N. (2008). A multi-session interpretation modification program: Changes in interpretation and social anxiety symptoms. *Behaviour Research and Therapy, 46*(10), 1135–1141. <https://doi.org/10.1016/J.BRAT.2008.05.012>
- Becker, D., Jostmann, N. B., & Holland, R. W. (2018). Does approach bias modification really work in the eating domain? A commentary on Kakoschke et al. (2017). *Addictive Behaviors, 77*, 293–294. <https://doi.org/10.1016/J.ADDBEH.2017.02.025>
- Blackwell, S. E., Browning, M., Mathews, A., Pictet, A., Welch, J., Davies, J., ... Holmes, E. (2015). Positive imagery-based cognitive bias modification as a web-based treatment tool for depressed adults: A randomized controlled trial. *Clinical Psychological Science, 3*(1), 91–111. <https://doi.org/10.1177/2167702614560746>
- Blackwell, S. E., & Holmes, E. A. (2010). Modifying interpretation and imagination in clinical depression: A single case series using cognitive bias modification. *Applied Cognitive Psychology, 24*(3), 338–350. <https://doi.org/10.1002/acp.1680>
- Blackwell, S. E., Woud, M. L., & MacLeod, C. (2017). A Question of Control? Examining the Role of Control Conditions in Experimental Psychopathology using the Example

- of Cognitive Bias Modification Research. *The Spanish Journal of Psychology*, 20, E54–E54. <https://doi.org/10.1017/sjp.2017.41>
- Bland, J. M., & Altman, D. G. (1994). Statistics Notes: Some examples of regression towards the mean. *BMJ*, 309(6957), 780–780. <https://doi.org/10.1136/bmj.309.6957.780>
- Brown, H., Eley, T., & Broeren, S. (2014). Psychometric properties of reaction time based experimental paradigms measuring anxiety-related information-processing biases in children. *Journal of Anxiety Disorders*, 28, 97–107. <https://doi.org/10.1016/j.janxdis.2013.11.004>
- Browning, M., Holmes, E. A., Charles, M., Cowen, P. J., & Harmer, C. J. (2012). Using attentional bias modification as a cognitive vaccine against depression. *Biological Psychiatry*, 72(7), 572–579. <https://doi.org/10.1016/j.biopsych.2012.04.014>
- Chan, S. W. Y., Lau, J. Y. F., & Reynolds, S. A. (2015). Is cognitive bias modification training truly beneficial for adolescents? *Journal of Child Psychology and Psychiatry*, 56(11), 1239–1248. <https://doi.org/10.1111/jcpp.12368>
- Clarke, P. J. F., Branson, S., Chen, N. T. M., Van Bockstaele, B., Salemink, E., MacLeod, C., & Notebaert, L. (2017). Attention bias modification training under working memory load increases the magnitude of change in attentional bias. *Journal of Behavior Therapy and Experimental Psychiatry*, 57, 25–31. <https://doi.org/10.1016/J.JBTEP.2017.02.003>
- Cronbach, L. J., & Furby, L. (1970). How we should measure “change”: Or should we? *Psychological Bulletin*, 74(1), 68–80. <https://doi.org/10.1037/h0029382>
- Dandeneau, S. D., Baldwin, M. W., Baccus, J. R., Sakellaropoulou, M., & Pruessner, J. C. (2007). Cutting Stress Off at the Pass: Reducing Vigilance and Responsiveness to Social Threat by Manipulating Attention. *Journal of Personality and Social Psychology*, 93(4), 651–666. <https://doi.org/10.1037/0022-3514.93.4.651>

- De Voogd, E. L., Wiers, R. W., Prins, P. J. M., & Salemink, E. (2014). Visual search attentional bias modification reduced social phobia in adolescents. *Journal of Behavior Therapy and Experimental Psychiatry*, 45(2), 252–259.
<https://doi.org/10.1016/j.jbtep.2013.11.006>
- Eberl, C., Wiers, R. W., Pawelczack, S., Rinck, M., Becker, E. S., & Lindenmeyer, J. (2013). Approach bias modification in alcohol dependence: Do clinical effects replicate and for whom does it work best? *Developmental Cognitive Neuroscience*, 4, 38–51.
<https://doi.org/10.1016/J.DCN.2012.11.002>
- Edwards, J. R. (2001). Ten Difference Score Myths. *Organizational Research Methods*, 4(3), 265–287. <https://doi.org/10.1177/109442810143005>
- Elfeddali, I., de Vries, H., Bolman, C., Pronk, T., & Wiers, R. W. (2016). A randomized controlled trial of Web-based Attentional Bias Modification to help smokers quit. *Health Psychology*, 35(8), 870–880. <https://doi.org/10.1037/hea0000346>
- Everaert, T., Spruyt, A., & Houwer, J. De. (2013). *On the malleability of automatic attentional biases: Effects of feature-specific attention allocation On the malleability of automatic attentional biases: Effects of feature-specific attention allocation*. (April), 37–41.
- Fox, E., Zougkou, K., Ashwin, C., & Cahill, S. (2015). Investigating the efficacy of attention bias modification in reducing high spider fear: The role of individual differences in initial bias. *Journal of Behavior Therapy and Experimental Psychiatry*, 49, 84–93.
<https://doi.org/10.1016/J.JBTEP.2015.05.001>
- Grafton, B., Mackintosh, B., Vujic, T., & MacLeod, C. (2013). When Ignorance is Bliss: Explicit Instruction and the Efficacy of CBM-A for Anxiety. *Cognitive Therapy and Research*. <https://doi.org/10.1007/s10608-013-9579-3>

- Hedge, C., Powell, G., & Sumner, P. (2018). The reliability paradox: Why robust cognitive tasks do not produce reliable individual differences. *Behavior Research Methods*, 50(3), 1166–1186. <https://doi.org/10.3758/s13428-017-0935-1>
- Heeren, A., Mogoase, C., Philippot, P., & McNally, R. J. (2015). Attention bias modification for social anxiety: A systematic review and meta-analysis. *Clinical Psychology Review*, 40, 76–90. <https://doi.org/10.1016/j.cpr.2015.06.001>
- Holmes, E. A., Arntz, A., & Smucker, M. R. (2007). Imagery rescripting in cognitive behaviour therapy: Images, treatment techniques and outcomes. *Journal of Behavior Therapy and Experimental Psychiatry*, 38, 297–305. <https://doi.org/10.1016/j.jbtep.2007.10.007>
- Kakoschke, N., Kemps, E., & Tiggemann, M. (2018). What is the appropriate control condition for approach bias modification? A response to commentary by Becker et al. (2017). *Addictive Behaviors*, 77, 295–296. <https://doi.org/10.1016/J.ADDBEH.2017.02.024>
- Kersbergen, I., Woud, M. L., & Field, M. (2015). The validity of different measures of automatic alcohol action tendencies. *Psychology of Addictive Behaviors*, 29(1), 225–230. <https://doi.org/10.1037/adb0000009>
- Koster, E. H. W., Fox, E., & MacLeod, C. M. (2009). Introduction to the Special Section on Cognitive Bias Modification in Emotional Disorders. *Journal of Abnormal Psychology*, 118(1), 1–4. <https://doi.org/10.1037/a0014379>
- Krebs, G., Hirsch, C. R., & Mathews, A. (2010). The effect of attention modification with explicit vs. Minimal instructions on worry. *Behaviour Research and Therapy*, 48(3), 251–256. <https://doi.org/10.1016/j.brat.2009.10.009>

- Kruijt, A.-W. (2014). General discussion. In *Depression vulnerability: Studying components of cognitive models* (Doctoral dissertation). Retrieved from <https://openaccess.leidenuniv.nl/bitstream/handle/1887/28632/06.pdf?sequence=11>
- Kruijt, A.-W., Parsons, S., & Fox, E. (2018). *No evidence for attention bias towards threat in clinical anxiety and PTSD: a meta-analysis of baseline dot probe bias in attention bias modification RCTs*. <https://doi.org/10.31234/OSF.IO/RFJUP>
- Kruijt, A.-W., Putman, P., & Van der Does, W. (2013a). A Multiple Case Series Analysis of Six Variants of Attentional Bias Modification for Depression. *ISRN Psychiatry*, 2013, 1–8. <https://doi.org/10.1155/2013/414170>
- Kruijt, A.-W., Putman, P., & Van der Does, W. (2013b). The effects of a visual search attentional bias modification paradigm on attentional bias in dysphoric individuals. *Journal of Behavior Therapy and Experimental Psychiatry*, 44(2), 248–254. <https://doi.org/10.1016/j.jbtep.2012.11.003>
- Kuckertz, J. M., Gildebrant, E., Liliequist, B., Karlström, P., Våppling, C., Bodlund, O., ... Carlbring, P. (2014). Moderation and mediation of the effect of attention training in social anxiety disorder. *Behaviour Research and Therapy*, 53(1), 30–40. <https://doi.org/10.1016/j.brat.2013.12.003>
- Kuckertz, J. M., Schofield, C. A., Clerkin, E. M., Primack, J., Boettcher, H., Weisberg, R. B., ... Beard, C. (2018). Attentional Bias Modification for Social Anxiety Disorder: What do Patients Think and Why does it Matter? *Behavioural and Cognitive Psychotherapy*, 1–23. <https://doi.org/10.1017/S1352465818000231>
- Lazarov, A., Abend, R., Seidner, S., Pine, D. S., & Bar-Haim, Y. (2017). The Effects of Training Contingency Awareness During Attention Bias Modification on Learning and Stress Reactivity. *Behavior Therapy*, 48(5), 638–650. <https://doi.org/10.1016/J.BETH.2017.03.002>

MacLeod, C., & Mathews, A. (2012). Cognitive Bias Modification Approaches to Anxiety.

Annual Review of Clinical Psychology, 8(1), 189–217.

<https://doi.org/10.1146/annurev-clinpsy-032511-143052>

MacLeod, C., Mathews, A., & Tata, P. (1986). Attentional bias in emotional disorders.

Journal of Abnormal Psychology, 95(1), 15–20.

MacLeod, C., Rutherford, E., Campbell, L., Ebsworthy, G., & Holker, L. (2002). Selective attention and emotional vulnerability: Assessing the causal basis of their association through the experimental manipulation of attentional bias. *Journal of Abnormal*

Psychology, 111(1), 107–123. <https://doi.org/10.1037//0021-843X.111.1.107>

McNally, R. J. (2018). Attentional bias for threat: Crisis or opportunity? *Clinical Psychology Review*. <https://doi.org/10.1016/j.cpr.2018.05.005>

Miloff, A., Savva, A., & Carlbring, P. (2015). Cognitive bias measurement and social anxiety disorder: Correlating self-report data and attentional bias. *Internet Interventions*, 2(3), 227–234. <https://doi.org/10.1016/j.invent.2015.03.006>

Mobini, S., Mackintosh, B., Illingworth, J., Gega, L., Langdon, P., & Hoppitt, L. (2014).

Effects of standard and explicit cognitive bias modification and computer-administered cognitive-behaviour therapy on cognitive biases and social anxiety.

Journal of Behavior Therapy and Experimental Psychiatry, 45(2), 272–279.

<https://doi.org/10.1016/j.jbtep.2013.12.002>

Mogoşe, C., David, D., & Koster, E. H. W. (2014). Clinical Efficacy of Attentional Bias Modification Procedures: An Updated Meta-Analysis. *Journal of Clinical*

Psychology, 70(12), 1133–1157. <https://doi.org/10.1002/jclp.22081>

Nishiguchi, Y., Takano, K., & Tanno, Y. (2015). *Explicitly Guided Attentional Bias Modification Promotes Attentional Disengagement From Negative Stimuli*.

<https://doi.org/10.1037/a0039117>

- Parsons, S., Kruijt, A.-W., & Fox, E. (2018). Psychological Science needs a standard practice of reporting the reliability of cognitive behavioural measurements. *Psyarxiv*.
<https://doi.org/10.17605/OSF.IO/6KA9Z>
- Platt, B., Murphy, S. E., & Lau, J. Y. F. (2015). The association between negative attention biases and symptoms of depression in a community sample of adolescents. *PeerJ*, 3, e1372. <https://doi.org/10.7717/peerj.1372>
- Price, R. B., Wallace, M., Kuckertz, J. M., Amir, N., Graur, S., Cummings, L., ... Bar-Haim, Y. (2016). Pooled patient-level meta-analysis of children and adults completing a computer-based anxiety intervention targeting attentional bias. *Clinical Psychology Review*, 50, 37–49. <https://doi.org/10.1016/j.cpr.2016.09.009>
- Reinecke, A., Becker, E. S., & Rinck, M. (2010). Three Indirect Tasks Assessing Implicit Threat Associations and Behavioral Response Tendencies. *Zeitschrift Für Psychologie / Journal of Psychology*, 218(1), 4–11. <https://doi.org/10.1027/0044-3409/a000002>
- Rogosa, D. R., & Willett, J. B. (1983). Demonstrating the Reliability of the Difference Score in the Measurement of Change. *Journal of Educational Measurement*, 20(4), 335–343. Retrieved from JSTOR.
- Salemink, E., van den Hout, M., & Kindt, M. (2007). Trained interpretive bias and anxiety. *Behaviour Research and Therapy*, 45(2), 329–340.
<https://doi.org/10.1016/J.BRAT.2006.03.011>
- Salemink, E., van den Hout, M., & Kindt, M. (2009). Effects of positive interpretive bias modification in highly anxious individuals. *Journal of Anxiety Disorders*, 23(5), 676–683. <https://doi.org/10.1016/j.janxdis.2009.02.006>
- Schoenmakers, T., Wiers, R. W., Jones, B. T., Bruce, G., & Jansen, A. T. M. (2007). Attentional re-training decreases attentional bias in heavy drinkers without

- generalization. *Addiction*, 102(3), 399–405. <https://doi.org/10.1111/j.1360-0443.2006.01718.x>
- Schoorl, M., Putman, P., & Does, W. Van Der. (2012). Attentional bias modification in posttraumatic stress disorder: A randomized controlled trial. *Psychotherapy and*
- Schumacher, S. E., Kemps, E., & Tiggemann, M. (2016). Bias modification training can alter approach bias and chocolate consumption. *Appetite*, 96, 219–224. <https://doi.org/10.1016/j.appet.2015.09.014>
- Staugaard, S. R. (2009). Reliability of two versions of the dot-probe task using photographic faces. *Psychology Science Quarterly*, 51(3), 339–350.
- Trafimow, D. (2015). A defense against the alleged unreliability of difference scores. *Cogent Mathematics*, 2(1). <https://doi.org/10.1080/23311835.2015.1064626>
- Van Dessel, P., De Houwer, J., & Gast, A. (2016). Approach–Avoidance Training Effects Are Moderated by Awareness of Stimulus–Action Contingencies. *Personality and Social Psychology Bulletin*, 42(1), 81–93. <https://doi.org/10.1177/0146167215615335>
- Wells, T. T., & Beevers, C. G. (2010). Biased attention and dysphoria: Manipulating selective attention reduces subsequent depressive symptoms. *Cognition & Emotion*, 24(4), 719–728. <https://doi.org/10.1080/02699930802652388>
- Wiers, R. W., Gladwin, T. E., Hofmann, W., Salemink, E., & Ridderinkhof, K. R. (2013). Cognitive Bias Modification and Cognitive Control Training in Addiction and Related Psychopathology: Mechanisms, Clinical Perspectives, and Ways Forward. *Clinical Psychological Science*, 1(2), 192–212. <https://doi.org/10.1177/2167702612466547>
- Wiers, R. W., Gladwin, T. E., & Rinck, M. (2013). Should we train alcohol-dependent patients to avoid alcohol? *Frontiers in Psychiatry*, 4, 33. <https://doi.org/10.3389/fpsy.2013.00033>

- Wiers, R. W., Rinck, M., Kordts, R., Houben, K., & Strack, F. (2010). Retraining automatic action-tendencies to approach alcohol in hazardous drinkers. *Addiction, 105*(2), 279–287. <https://doi.org/10.1111/j.1360-0443.2009.02775.x>
- Wiers, R. W., Teachman, B. A., & De Houwer, J. (2007). Implicit cognitive processes in psychopathology: An introduction. *Journal of Behavior Therapy and Experimental Psychiatry, 38*(2), 95–104. <https://doi.org/10.1016/J.JBTEP.2006.10.002>
- Williams, R. H., & Zimmerman, D. W. (1977). The Reliability of Difference Scores When Errors Are Correlated. *Educational and Psychological Measurement, 37*(3), 679–689. <https://doi.org/10.1177/001316447703700310>
- Williams, R. H., & Zimmerman, D. W. (1996). Are Simple Gain Scores Obsolete? *Applied Psychological Measurement, 20*(1), 59–69. <https://doi.org/10.1177/014662169602000106>
- Zimmerman, D., & Zumbo, B. (2015). Resolving the Issue of How Reliability is Related to Statistical Power: Adhering to Mathematical Definitions. *Journal of Modern Applied Statistical Methods, 14*(2). <https://doi.org/10.22237/jmasm/1446350640>