

A Question of Control? Examining the Role of Control Conditions in Experimental Psychopathology using the Example of Cognitive Bias Modification Research

Simon E. Blackwell¹, Marcella L. Woud¹ and Colin MacLeod²

¹ Ruhr-Universität Bochum (Germany)

² University of Western Australia (Australia)

Abstract. While control conditions are vitally important in research, selecting the optimal control condition can be challenging. Problems are likely to arise when the choice of control condition is not tightly guided by the specific question that a given study aims to address. Such problems have become increasingly apparent in experimental psychopathology research investigating the experimental modification of cognitive biases, particularly as the focus of this research has shifted from theoretical questions concerning mechanistic aspects of the association between cognitive bias and emotional vulnerability, to questions that instead concern the clinical efficacy of ‘cognitive bias modification’ (CBM) procedures. We discuss the kinds of control conditions that have typically been employed in CBM research, illustrating how difficulties can arise when changes in the types of research questions asked are not accompanied by changes in the control conditions employed. Crucially, claims made on the basis of comparing active and control conditions within CBM studies should be restricted to those conclusions allowed by the specific control condition employed. CBM studies aiming to establish clinical utility are likely to require quite different control conditions from CBM studies aiming to illuminate mechanisms. Further, conclusions concerning the clinical utility of CBM interventions cannot necessarily be drawn from studies in which the control condition has been chosen to answer questions concerning mechanisms. Appreciating the need to appropriately alter control conditions in the transition from basic mechanisms-focussed investigations to applied clinical research could greatly facilitate the translational process.

Received 28 February 2017; Revised 23 June 2017; Accepted 11 September 2017

Keywords: cognitive bias modification, experimental psychopathology, translational research.

Why Talk About Control Conditions? A Brief Introduction

“... and what should the control condition be?” is a ubiquitous yet often challenging question in experimental research. One area in which this question has become increasingly important in recent years is a line of experimental psychopathology research interested in the experimental modification of cognitive biases. In this domain of research, which has come to be known as the field of ‘Cognitive Bias Modification’ (CBM; Koster, Fox, & MacLeod, 2009), the initial development of procedures intended to manipulate cognitive biases was motivated by the goal of resolving theoretical questions concerning the causal role on anxiety of

negative biases in attention (MacLeod, Rutherford, Campbell, Ebsworthy, & Holker, 2002) and interpretation (Grey & Mathews, 2000; Mathews & MacLeod, 2002). However, in recent years there has been a proliferation of CBM research studies motivated by the quite different aim of determining the potential therapeutic effect of the bias modification procedures when delivered to clinical populations (see e.g., Woud & Becker, 2014). Although some of these more recent studies have led to optimism concerning the clinical promise of certain training paradigms (e.g., Wiers, Eberl, Rinck, Becker, & Lindenmeyer, 2011), across the broader field the pattern of findings has been quite mixed, highlighting the challenges associated with clinical translation (e.g., Fox, Mackintosh, & Holmes, 2014; Koster & Bernstein, 2015). Moreover, the design of these clinical studies, and the claims made on the basis of their findings, have not always taken adequate account of

Correspondence concerning this article should be addressed to Simon E. Blackwell, Mental Health Research and Treatment Center, Department of Psychology, Ruhr-Universität Bochum, Bochumer Fenster 3/05, Massenbergsstraße 9-13, 44787 Bochum, Germany.

E-mail: Simon.Blackwell@rub.de

Marcella L. Woud is supported by a Post Doctoral scholarship of the Daimler and Benz Foundation (32-12/4) and a grant of the Deutsche Forschungsgemeinschaft (DFG; WO2018/2-1). We would like to thank the many colleagues with whom we have discussed and debated the issues around control conditions; these discussions have been invaluable in developing the ideas elaborated within this paper.

How to cite this article:

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. Doi:10.1017/sjp.2017.41

the need to ensure that the conclusions drawn are limited to those allowed by the choice of control condition, and that this choice in turn is well-matched to the specific questions that the study aimed to address. With the growing number of clinical studies, the challenge posed by the choice of control condition has become increasingly recognised (e.g., Becker, Jostmann, & Holland, 2017; Blackwell et al., 2015; Hirsch, Meeten, Krahé, & Reeder, 2016; Kakoschke, Kemps, & Tiggemann, 2017).

While a control condition (or comparison group) is a vital element of any experimental or interventional research study, for various reasons the choice of this condition can cause difficulty. Anecdotally, selecting or designing an appropriate control condition can sometimes be more challenging than selecting or designing the 'active' condition of interest. Relatedly, interpreting the implications of finding (or failing to find) differences between the active and control conditions can be particularly challenging if the choice of control condition is suboptimal. Of course, these challenges are not unique to CBM research. However, the transitional nature of this particular field of research, which is currently going through the process of translation from experimental to clinical investigations, has brought to the fore the importance of ensuring that the chosen control condition(s) enables investigators to draw legitimate conclusions concerning the specific question(s) that their studies aim to answer. Thus, we hope that by discussing the factors that should bear upon the choice of control conditions in CBM research, we can elucidate issues relevant not only to the investigation of CBM but to the broader field of experimental psychopathology and translational research.

In this paper we reflect on the considerations that should guide the choice of control conditions within CBM research studies. The paper will not provide definitive answers as to which control conditions should be used, nor will it provide a complete overview of all the possible control conditions that could potentially be employed within this research field. Rather, we will use examples to illustrate the points of principle which, we argue, should serve to guide the selection of control conditions in CBM studies. These examples will be drawn mostly from CBM work that has focussed on the modification of biases in interpretation and attention, and we will assume some familiarity on the part of the reader with these training paradigms. However, the points of principle that we aim to communicate hold true not only for CBM studies targeting different types of cognitive bias, but for any translational research driven by advances in experimental psychopathology.

We will start with a brief consideration of why we need control conditions. As we will point out, although in all studies the inclusion of a control condition enables

investigators to contrast the impact of the active condition against this chosen comparison condition, the exact comparison required (and so the appropriate choice of control condition) will depend upon the precise question that the investigator wishes to answer. We then will discuss typical control conditions used in CBM research to date. In doing so, we will consider whether these have been well-suited to the changing nature of the questions that need to be addressed as a field negotiates the transition from theoretically motivated research in healthy participants to applied clinical research in patient populations. Finally, we will illustrate the alternative types of control conditions that can enable investigators to answer the quite different types of questions that can legitimately be asked across this broad and diverse domain of research. We hope that this paper will be useful for future work, stimulating thought and discussion about the critical issue of control conditions.

What Are Control Conditions and Why Do We Need Them?

When an investigator examines the effect of an experimental manipulation on a dependent measure of interest, their motivation is usually to answer a specific question that preceded and guided the development of the study. For example, an investigator may be motivated to answer the following question: Does exposing people to the risk of contact with a spider serve to elevate their anxiety? In order to address this question, the investigator may deliver an active condition in which participants must comply with the instruction "put your hand into this box, which contains a spider", and measure whether their anxiety increased when they do so, compared to their anxiety level as assessed at the beginning of the experimental session. However, even if such an anxiety elevation were to be observed in this active condition, it would not be possible to conclude that this resulted from exposure to the risk of contact with a spider. There may be other factors, quite irrelevant to the theory that risk of contact with a spider will elevate anxiety, that could account for the observed change in anxiety between these two assessment points. Anxiety may have been increased, for example, by the general testing context, by repeated assessment, by the way in which the experimenter spoke, and so on. The investigator can only determine whether exposure to the risk of contact with a spider served to elevate anxiety by comparing the effects obtained in this active condition with the effects observed in a 'control condition', identical in all respects to the active condition except for the fact that the risk of contact with a spider is not introduced. For example, such a control condition could change the instruction to "put your hand into this box, which is quite empty", delivering this amended instruction in an identical manner in an identical

experimental context. Note that this *particular* control condition provides the benchmark for comparison that is required to answer the *particular* question that this investigator sought to resolve. Almost all experimental research uses control conditions of some kind or another, for this purpose of eliminating the potential influence of factors irrelevant to the precise question under consideration. This enables investigators to answer their specific questions with confidence by comparing the effects observed in the active and control conditions. Clearly, however, which factors are relevant or irrelevant to the question under scrutiny will depend upon the precise nature of this question. Thus, as the questions asked by the researchers change, so will the appropriate control condition.

Development of Cognitive Bias Modification Research and the Control Condition 'Problem'

The following section considers how problems associated with the choice of control conditions may have arisen in cognitive bias modification research, when changes in the types of questions asked by investigators have not been accompanied by corresponding changes in the types of control conditions employed. As we will discuss, this line of research initially developed to address theoretically motivated questions, concerning the causal nature of the relationship between cognitive bias and emotional vulnerability, but more recent extensions of CBM research have instead sought to address applied questions concerning whether CBM training procedures can deliver clinically meaningful therapeutic benefits to people with various forms of psychological dysfunction. The control conditions necessary to resolve the former type of questions may be quite different from the control conditions needed to address the latter type of questions. We will argue that problems have arisen when clinically motivated CBM studies have persisted with the use of control conditions better suited to address theoretical questions not directly relevant to the evaluation of therapeutic efficacy. This issue will then be elaborated upon, by considering from a clinical trials perspective the types of conclusions that can, and that cannot, legitimately be drawn from the kinds of control conditions most commonly used in clinical CBM studies to date.

Early CBM Research Investigating Causal Contribution of Cognitive Biases to Anxiety Vulnerability

Research in the field of cognitive bias modification (CBM) was originally motivated by a very particular theoretical question; specifically, does biased attentional and / or interpretive processing of emotional information causally contribute to variation in anxiety vulnerability? To address this question, early CBM studies first exposed

two groups of healthy volunteers to procedures intended to induce a transient group difference in either negative attentional bias (MacLeod et al., 2002) or negative interpretive bias (Grey & Mathews, 2000; Mathews & Mackintosh, 2000), and then assessed whether this also led to a corresponding group difference in anxiety vulnerability, as evidenced by e.g., anxiety reactivity to a laboratory stressor. Typically these studies were carried out in a single experimental session, and to maximise the prospect of inducing the required group difference in cognitive bias, investigators typically gave all participants an 'active' bias modification procedure, configured in one group to induce an increase in the target cognitive bias and configured in the other group to induce a decrease in this cognitive bias. Although such an experimental design technically precludes the designation of one condition as the 'control', the question of interest can be addressed only by comparing effects across the two CBM conditions. The contrasting of two opposing CBM conditions, used in early interpretation CBM studies (induction of negative versus benign interpretations; Mathews & Mackintosh, 2000) and in early attentional CBM studies (attend to threat versus avoid threat; MacLeod et al., 2002), continued to be widely used in experimental studies when the question of interest concerns issues such as whether the bias targeted in the CBM manipulation causally contributes to disorder-relevant symptomatology (e.g., Woud, Holmes, Postma, Dalgleish, & Mackintosh, 2012), or the potential causal impact of one type of bias on another (e.g., interpretation bias on memory; Tran, Hertel, & Joormann, 2011), or the duration and stability of the induced bias (e.g., Mackintosh, Mathews, Yiend, Ridgeway, & Cook, 2006).

Later CBM Research Investigating the Therapeutic Benefits of CBM-based Interventions

The capacity to directly modify biases in processes such as interpretation and attention using simple experimental manipulations, demonstrated in earlier CBM research addressing theoretical questions concerning causality, led clinical investigators to ask whether inducing benign or positive biases could alleviate symptoms of psychopathology in clinical populations. Studies designed to address this newer question have delivered CBM procedures to clinical (i.e., meeting diagnostic criteria for a disorder) or subclinical (i.e., scoring high on a questionnaire measure of psychopathology) samples. Many of these studies have used multiple training sessions across extended periods of time, although single-session studies have also been used when addressing this issue. Some investigators have asked the question of whether successfully attenuating a dysfunctional cognitive bias leads to corresponding reductions in clinically relevant symptoms, whereas

others have asked whether exposure to the procedures that aim to bring about such bias change leads to clinically relevant therapeutic benefits. The importance of distinguishing between these two kinds of questions has recently been highlighted by MacLeod and Grafton (2016), who refer to them respectively as questions concerning the therapeutic benefits of the bias change *process*, and questions concerning the therapeutic benefits of *procedures* intended to evoke this bias change process.

Emergence of the Control Condition 'Problem'?

In the earlier experimental research, the choice of the two comparison conditions within these studies can be seen to be clearly and directly dictated by the specific questions that the researchers wished to answer. To address the question of whether a given cognitive bias can make a causal contribution to anxiety vulnerability, it is desirable to induce the greatest possible group difference in this target cognitive bias, and so contrasting two conditions designed to modify this bias in opposing directions represents the optimal comparison. But if the aim of a study is to investigate whether administering a CBM procedure leads to clinically-relevant therapeutic benefits, a quite different kind of question, what in this case should the control condition be?

In one of the earliest papers reporting a study designed to investigate potential beneficial training effects of an interpretation training paradigm, there is an interesting discussion of the challenges associated with deciding on an appropriate comparison condition (Mathews, Ridgeway, Cook, & Yiend, 2007); in this case, the control condition chosen for comparison was a test-retest group. This paper, and others appearing in the literature at the time, pointed out that when CBM studies seek to determine whether positive CBM training delivers therapeutic benefits, then comparing such positive CBM training against negative CBM training is no longer appropriate (see also Becker et al., 2017).

As we will go on to discuss, in a study aiming to investigate possible therapeutic benefits of a CBM procedure there are many candidate control conditions that could be employed; the most appropriate choice will reflect the specific nature of the question asked. However, one particular kind of control condition has come to dominate this field of research. Specifically, researchers have used a control condition intended to match the active CBM procedure as closely as possible in every regard, while eliminating the specific contingency intended to bring about the reduction in the target cognitive bias.

For example, in many studies investigating the therapeutic impact of CBM procedures designed to modify attentional bias (CBM-A), the active CBM condition is a dot-probe procedure in which probes predominantly appear in locations distal to negative information, and

the control condition is a dot-probe procedure that is identical, except that probes now appear equally often in locations distal to and proximal to negative information (e.g., Beevers, Clasen, Enock, & Schnyer, 2015). In many study investigating the therapeutic impact of CBM procedures intended to reduce negative interpretation bias (CBM-I), the active CBM condition is a task that presents ambiguous stimuli under conditions that encourage participants to predominantly resolve this ambiguity in a positive or benign manner, and the control task is identical except that this training contingency is removed, such that participants are encouraged to resolve the ambiguity equally often in a positive/benign or negative manner (e.g., Salemink, van den Hout, & Kindt, 2009). Other variations of these control conditions have been developed in CBM-A and CBM-I studies, all based on the same intention of keeping the control condition identical to the active CBM condition, and specifically excising only the contingency intended to modify the target cognitive bias. For example, in some CBM-I studies researchers have employed a control condition in which the interpretation of the presented ambiguity is left unconstrained (Murphy, Hirsch, Mathews, Smith, & Clark, 2007); and in some recent CBM-A studies using eye-tracking approaches control participants have been 'yoked' to participants receiving the active CBM condition, such that they receive an identical experience but without the task contingency based on their responding, which in the active condition operates to modify attentional bias (Vazquez, Blanco, Sanchez, & McNally, 2016).

This type of approach to the development of control conditions has also been adopted in CBM studies designed to evaluate the therapeutic benefits of training changes in other types of dysfunctional bias. For example, an active CBM condition designed to induce behavioural avoidance of alcohol, by requiring participants to make avoidance movements to images of alcoholic beverages and approach movements to images of non-alcoholic beverages, may be compared to a control condition that presents identical images under identical experimental conditions, but requires an equal number of approach and avoidance movements to both classes of stimuli (Wiers et al., 2011). In all these examples, the general idea has been to use a closely-matched 'sham training' control that eliminates only the specific aspect of the training procedure that the investigator assumes to represent the 'active ingredient' of the CBM intervention. Viewing these studies, and these 'sham training' control conditions, within the broader frame of psychological treatment development research, the use of such closely-matched control conditions is perhaps unparalleled in terms of the resulting capacity it affords to isolate the impact

of a single component of a candidate therapeutic intervention (Clarke, Notebaert, & MacLeod, 2014). However, it does not follow from this that such a condition represents the appropriate control when the aim is to determine whether or not this candidate therapeutic procedure delivers clinically relevant benefits in terms of symptom reduction.

One problem is that, while this approach does enable investigators to draw conclusions concerning whether the experimental condition is preferable to the putatively 'sham' condition, comparison between the active and sham condition does not necessarily indicate whether or not the active condition delivers therapeutic benefits. In principle, the active condition may have significantly better outcomes than the sham condition, but both may be detrimental to well-being; alternatively, the active condition may not have significantly better outcomes than the sham condition, but both may be beneficial to well-being. Indeed, the latter possibility has been raised retrospectively by some investigators who have adopted this experimental approach, and have found that their intended sham condition has performed surprisingly well in clinical trials, in terms of associated symptom reduction (e.g., Blackwell et al., 2015). Although these studies have usually lacked the additional control conditions necessary to determine whether the intended sham condition was more effective than no intervention, or treatment as usual, or another particular established intervention, this does raise important questions concerning how findings from studies using 'sham training' control conditions should be interpreted, and whether these are the most helpful control conditions when the goal is to determine whether or not the active CBM condition yields therapeutic benefits. This issue compromises the conclusions that can be drawn from some recent meta-analyses of CBM studies (e.g., Cristea, Kok, & Cuijpers, 2015; Hallion & Ruscio, 2011), in which between-group effect sizes reflecting comparisons of emotional outcomes between tightly matched active and sham variants of CBM training have sometimes been interpreted as indexing the capacity of the active CBM condition to reduce dysfunctional symptoms, rather than testing hypothesis concerning active mechanisms by contrasting the outcomes of two closely-matched experimental manipulations.

The question therefore arises, if the aim of a study is to determine how effective CBM-based interventions may be in leading to clinically-relevant reductions in symptoms of psychopathology, is it appropriate to rely on comparison against this type of sham training control condition? We will argue, in the next section, that effects based on such comparisons may be of limited value in resolving this particular question.

Interpreting Effects when Contrasting Active and Sham CBM Conditions

Estimates of treatment efficacy always require that the outcomes associated with a treatment of interest are compared with the outcomes associated with something else. Consequently, whether or not the treatment of interest is found to be statistically significantly more effective than this 'something else', and the extent of this superiority as evidenced by magnitude of the effect size, will depend not only on the impact of the target treatment, but also on the nature and impact of the 'something else' against which it is compared (cf., Hitchcock, Werner-Seidler, Blackwell, & Dalgleish, 2017). How should we interpret the presence or absence of statistically significant differences, and their magnitude, when comparing a CBM intervention against a sham training condition? And is it appropriate to compare the magnitude of effect sizes resulting from this comparison against the magnitude of effect sizes observed when alternative psychological or pharmacological interventions are instead contrasted against very different control conditions, and thereby draw inferences concerning the relative therapeutic impact of CBM versus these alternative types of interventions?

Perhaps the most common assumption is that sham CBM training represents a placebo condition¹. Of course, representing a comparison condition as a 'placebo' presupposes knowledge of the therapeutic agent within the active condition. Thus, if we can justify restricting consideration of therapeutic efficacy exclusively to the impact of the training contingency within the active CBM paradigm (for example), we can argue that the commonly used sham training condition controls for the potential therapeutic impact of various other aspects of the CBM intervention that are not directly related to this training contingency, such as exposure to the emotional stimulus material, engaging in the computer task, researcher contact, repeated assessment, passage of time, and so on. We may then interpret the resulting between-group effect size as providing an estimate of the specific contribution of the putative 'active ingredient' to symptom change, over and above all the other elements of the CBM procedure and intervention setting. Representing sham CBM as a placebo condition has intuitive appeal, as the logic seems to parallel the approach taken to outcome evaluation in the pharmacotherapy literature, where drug versus placebo trials are common. However, closer consideration of this analogy reveals that there are problems associated with conceptualising sham CBM training as a placebo condition.

¹Note that even if the placebo analogy is not used, or is explicitly rejected, many of the difficulties in interpreting the between-group differences as elaborated in this section, e.g., the additivity assumption, still hold.

While the definition of placebo has been the focus of debate (see e.g., Howick, 2016; Maddocks, Kerry, Turner, & Howick, 2016), a common understanding is that a placebo should be identical to the 'active' treatment in every respect, except that the 'active ingredient' is removed, to make the placebo condition 'inert'². However, there is no compelling reason to suppose that sham CBM conditions that expose participants to emotional information and that deliver contingencies configured to drive 50/50 positive/negative resolutions of ambiguity, or that present probes equally often in locations proximal and distal to negative information, are necessarily 'inert'. Particularly when repeated across prolonged periods of time, exposure to such conditions may in fact alter patterns of processing and influence emotional responding. Indeed, exposure to emotional negative stimuli is a well-established intervention technique with demonstrated anxiolytic effects. Therefore, sham CBM conditions cannot be assumed to be 'inert'. Comparison between active CBM and this type of sham condition permits at best an estimate only of the specific contribution made by the training contingency to whatever overall therapeutic impact the CBM package does, or does not, exert on clinical symptoms. However, even this interpretation can be problematic. Although we may talk of 'removing' a specific aspect of the CBM paradigm, such as the training contingency, in practical terms this specific aspect is replaced by or substituted by something else (e.g., presentation of the same or similar stimuli in a different kind configuration), which may have 'specific' effects of its own. In some studies researchers have aimed to create something perhaps closer to a 'true' placebo by e.g., removing all emotional content (Bowler et al., 2017), or using a completely unrelated task with potential credibility as 'brain training' (Hoppitt et al., 2014). Such comparisons may result in a more accurate estimate of efficacy 'versus placebo', and may sometimes even provide a tougher comparison if they have greater credibility to participants than a more closely-matched sham training condition, but these comparisons still answer questions about specificity rather than questions of therapeutic utility.

Even if we are primarily interested in the therapeutic contribution made specifically by the training contingency alone, or assume that we have successfully created a 'placebo' version of the training, the interpretation of the effect size observed when comparing active and sham CBM can be problematic for a number of reasons

(e.g., Rutherford & Roose, 2013). One such reason is that this approach assumes an additive model of treatment effects, presupposing that the individual contributions to the overall outcome in a treatment arm can be isolated and quantified by observing their addition or subtraction in the control arm. Rutherford and Roose (2013) challenge the validity of this additive assumption (in the context of pharmacological interventions), pointing out that there is little or no evidence for additiveness of medication and placebo responses (and see Berna et al., 2017 for evidence against this additivity assumption, in the context of pain). One key message is that the contribution of the 'active ingredient' to symptom change can vary according to the size of the 'non-specific' effects in the same treatment arm (and vice versa). Thus, for studies in which non-specific effects are kept to a minimum, an active treatment has the opportunity to show a much greater 'specific' effect compared to studies in which non-specific effects are maximised across both the active and placebo conditions. Further, in most study samples, there will be an upper limit on how much symptom improvement can be shown, due to e.g., inherent refractoriness and heterogeneity in the sample. Thus, an intervention with a very potent specific effect may nevertheless perform no better than a placebo when 'non-specific' effects are maximised, even though symptom improvement may be carried primarily by the specific treatment effect in the treatment arm, but by 'non-specific' factors in the placebo condition. However, it would be incorrect to conclude from this lack of differential efficacy that the 'active ingredient' has no effect. Rather, it could be concluded that the active ingredient is effective, but that when participants are already exposed to a rich array of non-specific therapeutic factors then the inclusion of this active ingredient does not enhance outcomes (cf., Rutherford & Roose, 2013). Thus, whether or not active CBM training outperforms sham CBM training, and the exact size of the difference (if present) cannot readily be translated into confident generalizable conclusions concerning the potency of the training contingency in driving therapeutic change.

Of course, there are other potential complications associated with the use of this type of sham CBM condition as a control condition in a treatment trial. For example, within studies testing the therapeutic benefits of CBM procedures, participants' expectancies may influence how they interact with the training materials, during and after sessions, within the sham condition as well as the active condition (Blackwell et al., 2015). Participants may develop their own theory of how they should change their processing to reap the expected benefits of training, and take action accordingly even when they are in the sham condition. Thus, participants

²In fact, this conceptualisation of a placebo fits well with recent re-considerations of the long-standing debate about how to define a placebo (Howick, 2016), one version of which is summarised by Maddocks et al. (2016): a placebo should contain all of the 'incidental' features of the treatment, none of the 'characteristic' features, and nothing else.

in this control condition may afford more weight, or may pay greater attention to, the positive interpretations of ambiguity that they encounter within this sham condition, perhaps rehearsing these between sessions, or endeavouring to impose such positive resolutions themselves on unresolved ambiguity (Blackwell et al., 2015). Alternatively, a sham training that is essentially a 'watered-down' version of the active training may not necessarily remain equally credible to participants (particularly when participants know that they have a 50% chance of receiving a 'sham' training), thus no longer adequately controlling for expectancy (see e.g., Maddocks et al., 2016 for an interesting parallel discussion concerning trying to develop a placebo version of an exercise intervention). Therefore, although there may be a clear and precise distinction between the *procedures* that participants are exposed to when they receive active or sham CBM conditions, this procedural distinction may not reliably translate into a difference in the degree to which bias change *processes* are elicited by these two conditions. Importantly, it has been pointed out that when studies have failed to observe any differing therapeutic impact of active vs sham CBM procedures, then almost invariably they have also failed to obtain evidence that these differing procedures succeeded in elicited the differing patterns of cognitive bias that they were intended to induce in recipients (MacLeod & Grafton, 2016).

A final problem we will note in relation to a reliance on active versus 'placebo' contrasts in treatment evaluation research concerns the difficulties associated with inferring clinical utility from the observed outcomes. There has been increasing recognition within the pharmacotherapy literature of the inadequacy of placebo-controlled trials as a means of determining the clinical value of new drugs (e.g., Nunn, 2009). While a placebo-controlled trial may be appropriate to answer questions concerning the specificity of treatment effects relating to target mechanisms of interest, the resulting between-group effect size may provide little information about the clinical utility of the new treatment. One reason for this is that the size of this effect is likely to be heavily influenced by factors relating to the specific context of the trial, such as aspects of the treatment trial design that inflated or suppressed placebo responding, or that influenced the credibility of the interventions under evaluation. Moreover, the knowledge that one is in a placebo-controlled trial can affect expectancy, influencing response to both the placebo and the active treatment. Thus, the generalization of the placebo-controlled effect size is problematic, as it is unlikely to accurately reflect the therapeutic change that will occur in the real world setting, where participants generally know that they are receiving an active treatment (cf., Rutherford et al., 2017).

Overall, while sham training control conditions can be extremely valuable for a number of quite specific purposes, the between-group effect size relating to the impact on symptoms of a CBM intervention and a sham training control does not readily permit conclusions concerning the clinical utility of the CBM intervention. However, it has been common to base such conclusions on this type of comparison. For example, meta-analyses have tended to draw conclusions concerning the efficacy of CBM procedures as interventions for mental health problems on the basis of the findings obtained when active vs sham CBM procedures have been compared (e.g., Cristea et al., 2015; Hallion & Ruscio, 2011). Indeed, meta-analyses have at times made claims concerning the relative efficacy of CBM approaches and other psychological interventions, by comparing effect sizes obtained when active vs sham CBM conditions have been contrasted to effect sizes obtained when alternative psychological interventions (e.g., cognitive behavioural therapy) have been contrasted against very different types of control conditions that are very much weaker in terms eliminating the impact of non-specific factors (e.g., no treatment or waitlist controls). A failure to recognize the importance of the control condition when interpreting the meaning of observed effects compromises the validity of conclusions concerning the potential utility of CBM-based interventions in comparison to other kinds of interventions. The problem is not that sham training controls are inappropriate – in fact they can be extremely valuable – but rather that their use enables researchers only to answer a particular type of question (as will be discussed below). Importantly, this question does not concern the clinical utility of CBM-based interventions. Whenever researchers seek to answer a particular question about a candidate clinical intervention, then it is vitally important that they contrast that candidate intervention against the specific control condition that does in fact serve to resolve the researcher's exact question.

Choosing Control Conditions Based on the Questions that Studies Seek to Address

It may seem obvious that the choice of control condition should be driven by the particular question that a study seeks to address. It should be similarly clear that the conclusions we can legitimately draw from any previous study should concern only those particular issues that can be resolved in light of the specific control condition adopted in that study. Nevertheless, as a field develops and research progresses, conventions can easily emerge, such that certain methodological features well-suited to resolving early questions in the field become so familiar that they are carried into new contexts, where different questions are asked that these

familiar control condition are no longer well-placed to answer. This may explain why the use of sham CBM control conditions, well-suited to answering particular theoretical questions that motivated development of the CBM field, has persisted into recent extensions of CBM research intended to address very different applied questions concerning the clinical utility of CBM approaches. In many cases, the use of sham CBM as the chosen control condition appears to be based on habit, rather than consideration of the study's aim. Thus, instead of investigators asking themselves, 'Given our specific research question about the CBM procedure we are evaluating, what is the best control group?' they may jump straight to 'How do we create a sham version of the CBM procedure we are evaluating?' Unfortunately, such a lack of reflection can be counterproductive to the field. There are many different kinds of questions that can be addressed via a study using a CBM procedure, and answering each of these questions is likely to require a different kind of control comparison. In the following sections we consider two broad kinds of questions that CBM researchers have been motivated to address, and we discuss how the type of control condition required differs in each case. First, we consider studies that have sought to address questions concerning whether variation in a particular cognitive bias causally contributes to emotional symptomatology of interest. Second, we consider studies addressing questions concerning the clinical utility of CBM approaches. As will be seen, quite different types of comparison conditions are best suited to successfully answer these quite different types of question.

Studies Asking whether Cognitive Biases Contribute to Emotional Variability of Interest

As discussed earlier, the initial work that drove the development of CBM procedures, and much ongoing research, sought to answer questions concerning whether variability in particular cognitive biases causally contributes to variability in emotional vulnerability (or to other types of behavioural or cognitive variability). In general, experimental manipulations are required to address questions of causality (e.g., Kraemer et al., 1997). To answer questions concerning the causal impact of a cognitive bias, the main requirement is that an intended CBM manipulation must directly induce a group difference in the target cognitive bias, in order to determine whether this gives rise to a concomitant group difference in the emotional vulnerability (or behavioural or cognitive measures) of interest. Unless the experimenter succeeds in inducing this group difference in the target cognitive bias it is not possible to determine the causal contribution of this bias to other variability of interest (MacLeod & Grafton, 2016; see also Kraemer et al., 1997 for a more general

elaboration). The prospect of successfully inducing this required group difference in cognitive bias, in order to answer the causal questions under scrutiny, can be maximised by contrasting one CBM condition configured to attenuate the target cognitive bias against another CBM condition configured instead to amplify this bias. It is important to note that this contrast does not reveal whether either CBM condition is 'beneficial' in an absolute sense, only whether inducing a discrepancy in this bias induces a discrepancy in the emotional measure thought to be potentially causally influenced by this bias. Other experimental contrasts may be used to investigate the role of different parameters of a bias-training procedure (e.g., provision of explicit instructions as to the training contingency, Grafton, Mackintosh, Vujic, & MacLeod, 2014; or imagery versus verbal processing, Holmes & Mathews, 2005). Again, in this case there will generally be two contrasting experimental conditions, and it could be argued that neither is a 'control' condition as such.

While transiently inducing a cognitive bias thought to potentially make a causal contribution to individual differences in emotional vulnerability may be acceptable when working with healthy participants, the results of such CBM research cannot determine whether cognitive biases causally contribute to dysfunctional symptoms of clinical relevance. Answering this question requires researchers to study cohorts of clinical or sub-clinical participants who report experiencing such dysfunctional symptoms. While it remains appropriate to deliver to these participants a CBM procedure configured to reduce the bias thought to potentially contribute to their dysfunctional symptoms, it becomes ethically dubious to expose them to a CBM procedure configured instead to increase this bias. Consequently, researchers asking whether a given cognitive bias causally contributes to clinically relevant symptoms most often instead employ a variant of the 'sham training' control condition that we already have discussed. Once again, however, the goal is simply to induce a group difference in this target cognitive bias after the CBM manipulation, in order to determine whether this gives rise to a corresponding group difference in the dysfunctional symptom of interest, as predicted by the causal hypothesis under test. Again, the capacity of the study to answer the casual question addressed by the investigators depends on the two groups showing different biases at post-training, and so care should be taken in piloting the CBM protocol to maximize the likelihood of this occurrence. If post-training there is no bias difference between the two groups, then no conclusions can be drawn about the potential impact of the induced bias on the dysfunctional symptom of interest (cf., Clarke et al., 2014; MacLeod & Grafton, 2016). If such a group difference in the cognitive bias is successfully induced,

and this is found to be accompanied by a corresponding group differences in a laboratory measure of the dysfunctional symptom of interest, then this will confirm that the cognitive bias can functionally contribute to this type of symptom. However, such a finding does not in itself permit conclusions concerning the clinical utility of a CBM intervention targeting this cognitive bias.

Of course, the finding that a CBM manipulation inducing a transient group difference in a target cognitive bias serves to evoke a corresponding group difference in some laboratory measure of a dysfunctional symptom cannot permit the conclusion that *enduring* individual differences in this cognitive bias causally contribute to such dysfunctional symptomatology in a real world setting. Some researchers who have taken CBM outside of the laboratory setting have been motivated to address this more sophisticated causal question. Again, clinical or subclinical participants are selected who report experiencing the dysfunctional symptom of interest. Again, the goal is still to directly induce a group difference in the target bias, though now the aim is to have this induced group difference in bias persist outside the laboratory, and so multiple CBM sessions are usually delivered. But the same two comparison conditions remain appropriate. That is, it is appropriate to elicit this group difference by giving one group of participants extended exposure to a CBM condition configured to attenuate the target bias, and the other group equivalent exposure to a sham CBM condition. If an enduring group difference in the target cognitive bias is successfully induced in this manner, then the finding that a corresponding group difference in real world dysfunction symptomatology also becomes evident will warrant an affirmative answer to the question under test. That is, such results indicate that this cognitive bias does causally contribute to the occurrence of the dysfunctional symptom within the naturalistic setting. But for the reasons already discussed, this finding will not in itself permit conclusions concerning whether or not the CBM intervention has clinical utility, as the sham condition represents a suboptimal control condition for answering this quite different question.

While these studies, involving extended delivery of active and sham CBM procedures to participants reporting clinical or subclinical dysfunction, may not serve to adequately determine the clinical utility of these CBM procedures, they do provide ‘proof-of-principle’ that it may be possible to deploy CBM procedures in appropriately designed treatment trials. An essential part of the translational process, when moving from lab to clinic, is to demonstrate specificity of the proposed intervention, by showing that its influence stems from the putative ‘active ingredient’, and not simply from non-specific factors such as expectancy, generic

aspects of the intervention, or the way it is delivered such as contact with a researcher. A ‘sham training’ control condition, or an appropriate ‘placebo’ variant is ideal for serving such a purpose. The requirement to demonstrate specificity within a clinical sample, prior to conducting formal treatment trials, helps prevent premature execution of large-scale RCTs, which are not only expensive in terms of patient and investigator time, money, and other resources, but in the absence of adequate preparatory work are likely to fail, potentially compromising the development of promising new intervention approaches.

The nature of the control condition is not the only design aspect of such ‘proof-of-principle’ studies that may limit the conclusions that can be drawn from them regarding clinical utility. Researchers will naturally want to maximise power, that is, the chance of finding an effect if it is indeed there. One well-discussed contributor to power is of course sample size (itself perhaps a problem in much CBM research, e.g., Rinck, 2017), but other design features that may increase power are to select a relatively homogeneous sample (for example by restricting participant variability on characteristics such as severity, medication use etc.), to minimise the potential influence of non-specific factors such as expectancy (for example by avoiding advertising the study as a trial of an exciting new treatment), and to minimise variability due to natural remission or ‘life events’, for example by adopting a short time-frame for training and follow-up. All of these are positive design features that can maximise the chances of detecting ‘specific’ training effects, which is appropriate for the particular questions that may be addressed in such a ‘proof-of-principle’ study context. However, they compromise the degree to which inferences can be drawn concerning the clinical utility of suitably designed CBM variants when appropriately configured and delivered in a real-world treatment application; in such a real-world application the participants would likely be more heterogeneous, both the intervention and follow-up periods may be more extended, and it may be appropriate to adaptively harness non-specific factors likely to enhance the success of the CBM intervention, rather than to systematically eliminate them. Most importantly, for the purpose of our current discussion, a study intending to address questions of treatment utility in a ‘real-world’ application may also require quite different kinds of control conditions.

Studies asking whether CBM Procedures have Clinical Utility as Interventions

There are many different questions that can reasonably be asked about the clinical utility of a new candidate intervention approach, each of which can usefully inform the decisions that clinicians or service-providers make

concerning whether or not to employ this intervention approach with patients under their care. Importantly, within a clinical setting a service provider will rarely be making a choice between offering a treatment or a placebo (or even less 'sham training'), and thus comparisons against control conditions intended to demonstrate specificity of effects may have limited relevance to clinical decision-making. Rather, clinicians are likely to want answers to questions such as "are my patients likely to improve if I give them this treatment?", "would adding this to the treatment I usually provide lead to better outcomes?" or "would this be better than the treatment I usually offer?" As we will go on to discuss, in order to answer these different type of questions concerning clinical utility, differing comparison conditions are required.

When choosing a control condition, investigators seeking to evaluate the clinical utility of a CBM intervention could usefully ask themselves: Why are we bothering to develop this particular CBM procedure into a potential intervention in the first place? How are we expecting it to be used, by whom and in what kind of context, and what kind of effects do we expect it to have? Once these issues have been thought through (perhaps with the help of a framework such as 'Patient-Intervention-Comparison-Outcome', or PICO; Schardt, Adams, Owens, Keitz, & Fontelo, 2007), designing a trial that can determine the clinical utility of this CBM procedure, with respect to its intended purpose, becomes much more straightforward. We will illustrate this general idea with a few examples.

Let us imagine that our intended purpose is to deliver a particular CBM intervention to assist with symptom relief while patients are on a waiting list for face-to-face therapy. This purpose leads us to ask whether patients on the waitlist who receive the CBM intervention experience greater symptom relief during the waiting period than do patients on the waitlist who do not receive the CBM intervention. Thus, a clinically meaningful comparison would be to compare the CBM intervention to simply being on the waitlist, as this would tell us the relative advantage (and cost-effectiveness) of adding in the intervention versus not doing so. Of course, there are many other things patients on a waiting list could potentially be offered, such as access to self-help materials (whether internet or paper-based), or another CBM or cognitive training intervention. So, we could also ask whether the CBM procedure is superior to these alternatives in terms of alleviating symptoms while patients await face-to-face treatment. This quite different question is also clinically relevant, but answering it would instead require us to compare the CBM intervention to both the waitlist and one or more alternative low-intensity intervention that could readily be made available to waitlist patients. The outcome of such comparisons would serve to guide

clinical decision making. More complex questions, for example concerning whether particular subsets of individuals on the waitlist benefit more or less from each of these alternative interventions, would require more complex comparisons involving participant subgroups. However, by constructing the comparisons to resolve specific questions of clinical relevance, we increase the chances that our findings will be clinically informative.

Perhaps, however, we wish to determine whether a particular CBM intervention could augment the effectiveness of some already validated treatment, such as internet-delivered CBT (e.g., Williams, Blackwell, Mackenzie, Holmes, & Andrews, 2013), pharmacotherapy, or a more complex care package such as inpatient treatment or an existing multi-modal psychosocial intervention (e.g., Ferrari, Becker, Smit, Rinck, & Spijker, 2016; Wiers et al., 2011). Again, answering a specific question about how much additional benefit is gained from adding in the CBM procedure (and how cost-effective this is) requires a particular comparison, this time between the validated intervention alone and the validated intervention with the CBM module added. The outcome is then directly informative as to the clinical utility of adding in the CBM procedure in this context. And, as with the waitlist example, there could be many different variants of this question, each requiring a distinctive type of comparison to be carried out in order to provide an answer. For example, the intention may be to add a module to the validated intervention, and the question could be whether it would be better for this module to be a particular CBM procedure, or instead to comprise some alternative candidate therapeutic component such as another cognitive or behavioural intervention. Answering this would require comparing the original intervention alone against variants that include either the additional CBM component, or the additional alternative component. The findings revealed by the chosen contrasts should be clinically informative, so long as these are appropriately chosen to answer specific questions of clinical relevance.

Yet another possible reason for seeking to evaluate the therapeutic impact of a CBM procedure would be to determine whether this intervention approach may be more beneficial (or more cost-effective) than the intervention approach routinely made available to patients when they approach a particular clinical service (e.g., a GP surgery). To answer this question, it would be appropriate to contrast the CBM intervention against 'treatment as usual' (that is, the treatment patients routinely received prior to commencement of the study). For example, if in normal circumstances everyone approaching a particular clinical service with a certain disorder was offered computerized CBT, then this trial would randomly assign future patients to either computerized CBT or to the CBM intervention. The resulting comparison between these

groups, in terms of symptom improvement, should be clinically meaningful, because it serves to answer a clinically meaningful question.

Questions regarding the relative efficacy, cost-effectiveness, or utility as measured by another index (e.g., accessibility, scalability), of different active treatments, and the subsequent comparison conditions these entail, may be particularly valuable from a clinical perspective. Moreover studies in which comparison conditions are active treatments avoid the methodological problems that arise when comparing treatments conditions only against placebo, sham training, no-treatment, or waitlist conditions, in which the knowledge of potentially being in a placebo/sham training condition, or definitely being in a no-treatment or waitlist condition, will influence participants' expectancies of benefits (or lack of them). Finally, taking the perspective of clinical utility allows researchers to move towards addressing important questions concerning which treatments may be most useful for which subsets of patients. Knowing whether one treatment on average performs better than, or equivalently to, another in itself goes a useful step beyond simply proliferating multiple treatment approaches all of which perform better than waitlist or their 'placebo' variant. However, in comparing two active treatments, by carefully distinguishing participants on potential dimensions of relevance (such as disorder subtype or symptom profile) it becomes possible to start trying to answer critically important questions such as 'which of these alternative treatments is likely to be most effective for this particular individual', thereby matching patients to treatments. While such patient stratification is seen as a valuable goal by many, it is in fact only really possible if treatments are directly compared. The exciting promise of research that illuminates the specific cognitive mechanisms through which different variants of CBM deliver their therapeutic benefits, is that this enables us to move beyond relatively blunt traditional candidate moderators of treatment outcome (such as age, gender, or symptom severity), and instead identify the precise patterns of cognitive bias (or other indices of emotional processing) that indicate whether a given individual is likely to benefit from one or more alternative variants of CBM interventions. For example, it may be that emotional response to brief 'sample' of a CBM intervention can be used to identify those for whom a full series of CBM sessions is, or is not, likely to be useful (Blackwell et al., 2015). There is still some way to go before such targeted identification of likely responders becomes a reality, but it is an exciting and productive route to explore.

Conclusions

Research investigating whether the direct modification of cognitive biases serves to alter the tendency to

experience psychological dysfunction has given rise to a substantial body of work, and this now represents a significant field of research. The potential to experimentally manipulate key cognitive processes and observe the effects on indices of psychopathology offers a wide array of opportunities to advance theoretical understanding, while also developing new types of potential clinical interventions. As CBM research has extended into the clinical domain, the issues addressed by researchers have changed from being questions of a purely theoretical nature concerning the causal role of biased cognition in the determination of dysfunctional symptomatology, to now also include important applied questions concerning the clinical utility of CBM procedures in the therapeutic remediation of such dysfunction. However, despite this major change in the types of questions CBM studies have sought to address, the control conditions they have most often employed remain largely the same. This has constrained the capacity of more recent CBM studies to provide clear answers to clinically pertinent questions.

In this manuscript we have reviewed the control conditions commonly employed in CBM studies, considered their strengths and limitations, and have suggested a way forwards. Specifically, we propose that the choice of control conditions within future CBM studies should be closely guided by the specific question that each such study is conducted to address. Consistent with this advice, we argue that the conclusions drawn from the findings obtained in any CBM study should be constrained to those that concern the specific question that the choice of control condition permits that study to answer. We point out that, although comparing an active CBM condition with a sham version of that same CBM procedure can provide answers to important questions concerning the causal impact of the target bias on symptoms of interest, and to questions concerning the mechanisms through which this impact is delivered, comparisons involving this particular control condition are unlikely to provide a directly meaningful estimate of clinical utility. Such comparisons have led to confusion between estimates of treatment specificity and estimates of clinical efficacy, and the use of the resulting effect sizes to make inferences about clinical utility can lead to unwarranted conclusions. We caution that, when an estimate of clinical utility is the desired outcome of a CBM study then a 'sham training' control condition may represent a poor choice. As we have illustrated through example, the ideal control condition under these circumstances will be determined by the particular question concerning clinical utility that the investigator wishes to answer.

The rapid evolution of CBM research across the past decade, from laboratory studies investigating theoretical issues to field studies intended to examine

therapeutic efficacy, has required many investigators to navigate a steep learning curve concerning the methodological challenges associated with the clinical translational process. However, the issues addressed in this article around the impact of decisions concerning control conditions, and the suggestions we offer concerning how these choices should be made, are by no means restricted to CBM research. Rather, they have broad applicability across many areas of translational research, in which investigators seek to develop experimental designs used to address theoretical questions by evaluating the impact of manipulations delivered in the laboratory, into methodologies capable of determining whether these manipulations can deliver meaningful therapeutic benefits in the clinical setting. We believe that engaging with these important issues, and addressing them successfully, will facilitate the translational process in ways that further enhance the contribution that such work makes to clinical progress.

While new candidate interventions are always welcomed, sometimes with great enthusiasm, control conditions seldom generate much interest or excitement, with the consequence that the control condition can feel like the 'poor cousin' of the active training condition in treatment studies. And yet, the value of such studies is critically dependent on the selection and construction of the appropriate control condition, and this affords many opportunities for creativity and ingenuity in study design. Although much of our discussion has focussed on problems and challenges, we hope that by placing control conditions at centre-stage this paper also provides an opportunity to celebrate this often overlooked, but vitally important, component of our scientific studies, which warrants more consideration and deserves greater credit than it sometimes receives.

References

- Becker D., Jostmann N. B., & Holland R. W. (2017). Does approach bias modification really work in the eating domain? A commentary on Kakoschke et al. (2017). *Addictive Behaviors*, Available online. <https://doi.org/10.1016/j.addbeh.2017.02.025>
- Beevers C. G., Clasen P. C., Enock P. M., & Schnyer D. M. (2015). Attention bias modification for major depressive disorder: Effects on attention bias, resting state connectivity, and symptom change. *Journal of Abnormal Psychology*, *124*, 463–475. <https://doi.org/10.1037/abn0000049>
- Berna C., Kirsch I., Zion S. R., Lee Y. C., Jensen K. B., Sadler P., ... Edwards R. R. (2017). Side effects can enhance treatment response through expectancy effects: an experimental analgesic randomized controlled trial. *Pain*, *158*, 1014–1020. <https://doi.org/10.1097/j.pain.0000000000000870>
- Blackwell S. E., Browning M., Mathews A., Pictet A., Welch J., Davies J., ... Holmes E. A. (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>
- Bowler J. O., Hoppitt L., Illingworth J., Dalgleish T., Ononaiye M., Perez-Olivas G., & Mackintosh B. (2017). Asymmetrical transfer effects of cognitive bias modification: Modifying attention to threat influences interpretation of emotional ambiguity, but not vice versa. *Journal of Behavior Therapy and Experimental Psychiatry*, *54*, 239–246. <https://doi.org/10.1016/j.jbtep.2016.08.011>
- Clarke P. J. F., Notebaert L., & MacLeod C. (2014). Absence of evidence or evidence of absence: Reflecting on therapeutic implementations of attentional bias modification. *BMC Psychiatry*, *14*, 8. <https://doi.org/10.1186/1471-244X-14-8>
- Cristea I. A., Kok R. N., & Cuijpers P. (2015). Efficacy of cognitive bias modification interventions in anxiety and depression: Meta-analysis. *The British Journal of Psychiatry*, *206*(1), 7–16. <https://doi.org/10.1192/bjp.bp.114.146761>
- Ferrari G. R. A., Becker E. S., Smit F., Rinck M., & Spijker J. (2016). Investigating the (cost-) effectiveness of attention bias modification (ABM) for outpatients with major depressive disorder (MDD): A randomized controlled trial protocol. *BMC Psychiatry*, *16*(1), 370. <https://doi.org/10.1186/s12888-016-1085-1>
- Fox E., Mackintosh B., & Holmes E. A. (2014). Travellers' tales in cognitive bias modification research: A commentary on the Special Issue. *Cognitive Therapy and Research*, *38*, 239–247. <https://doi.org/10.1007/s10608-014-9604-1>
- Grafton B., Mackintosh B., Vujic T., & MacLeod C. (2014). When ignorance is bliss: Explicit instruction and the efficacy of CBM-A for anxiety. *Cognitive Therapy and Research*, *38*, 172–188. <https://doi.org/10.1007/s10608-013-9579-3>
- Grey S., & Mathews A. (2000). Effects of training on interpretation of emotional ambiguity. *The Quarterly Journal of Experimental Psychology Section A*, *53*, 1143–1162. <https://doi.org/10.1080/713755937>
- Hallion L. S., & Ruscio A. M. (2011). A meta-analysis of the effect of cognitive bias modification on anxiety and depression. *Psychological Bulletin*, *137*, 940–958. <https://doi.org/10.1037/a0024355>
- Hirsch C. R., Meeten F., Krahé C., & Reeder C. (2016). Resolving ambiguity in emotional disorders: The nature and role of interpretation biases. *Annual Review of Clinical Psychology*, *12*(1), 281–305. <https://doi.org/10.1146/annurev-clinpsy-021815-093436>
- Hitchcock C., Werner-Seidler A., Blackwell S. E., & Dalgleish T. (2017). Autobiographical episodic memory-based training for the treatment of mood, anxiety and stress-related disorders: A systematic review and meta-analysis. *Clinical Psychology Review*, *52*, 92–107. <https://doi.org/10.1016/j.cpr.2016.12.003>
- Holmes E. A., & Mathews A. (2005). Mental imagery and emotion: A special relationship? *Emotion*, *5*, 489–497. <https://doi.org/10.1037/1528-3542.5.4.489>
- Hoppitt L., Illingworth J. L., MacLeod C., Hampshire A., Dunn B. D., & Mackintosh B. (2014). Modifying social anxiety related to a real-life stressor using online cognitive

- bias modification for interpretation. *Behaviour Research and Therapy*, 52, 45–52. <https://doi.org/10.1016/j.brat.2013.10.008>
- Howick J.** (2016). The relativity of ‘placebos’: Defending a modified version of Grünbaum’s definition. *Synthese*, 194, 1363–1396. <https://doi.org/10.1007/s11229-015-1001-0>
- Kakoschke N., Kemps E., & Tiggemann M.** (2017). What is the appropriate control condition for approach bias modification? A response to commentary by Becker et al. (2017). *Addictive Behaviors*, Available online. <https://doi.org/10.1016/j.addbeh.2017.02.024>
- Koster E. H. W., & Bernstein A.** (2015). Introduction to the special issue on cognitive bias modification: Taking a step back to move forward? *Journal of Behavior Therapy and Experimental Psychiatry*, 49, 1–4. <https://doi.org/10.1016/j.jbtep.2015.05.006>
- Koster E. H. W., Fox E., & MacLeod C.** (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>
- Kraemer H. C., Kazdin A. E., Offord D. R., Kessler R. C., Jensen P. S., & Kupfer D. J.** (1997). Coming to terms with the terms of risk. *Archives of General Psychiatry*, 54, 337–343. <https://doi.org/10.1001/archpsyc.1997.01830160065009>
- Mackintosh B., Mathews A., Yiend J., Ridgeway V., & Cook E.** (2006). Induced biases in emotional interpretation influence stress vulnerability and endure despite changes in context. *Behavior Therapy*, 37, 209–222. <https://doi.org/10.1016/j.beth.2006.03.001>
- MacLeod C., & Grafton B.** (2016). Anxiety-linked attentional bias and its modification: Illustrating the importance of distinguishing processes and procedures in experimental psychopathology research. *Behaviour Research and Therapy*, 86, 68–86. <https://doi.org/10.1016/j.brat.2016.07.005>
- 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>
- Maddocks M., Kerry R., Turner A., & Howick J.** (2016). Proleptic placebos in physical therapy trials. *Journal of Evaluation in Clinical Practice*, 22, 598–602. <https://doi.org/10.1111/jep.12582>
- Mathews A., & Mackintosh B.** (2000). Induced emotional interpretation bias and anxiety. *Journal of Abnormal Psychology*, 109, 602–615. <https://doi.org/10.1037/0021-843X.109.4.602>
- Mathews A., & MacLeod C.** (2002). Induced processing biases have causal effects on anxiety. *Cognition & Emotion*, 16, 331–354. <https://doi.org/10.1080/02699930143000518>
- Mathews A., Ridgeway V., Cook E., & Yiend J.** (2007). Inducing a benign interpretational bias reduces trait anxiety. *Journal of Behavior Therapy and Experimental Psychiatry*, 38, 225–236. <https://doi.org/10.1016/j.jbtep.2006.10.011>
- Murphy R., Hirsch C. R., Mathews A., Smith K., & Clark D. M.** (2007). Facilitating a benign interpretation bias in a high socially anxious population. *Behaviour Research and Therapy*, 45, 1517–1529. <https://doi.org/10.1016/j.brat.2007.01.007>
- Nunn R.** (2009). It’s time to put the placebo out of our misery. *The British Medical Journal*, 338, b1568. <https://doi.org/10.1136/bmj.b1568>
- Rinck M.** (2017). CBM research needs more power: Commentary on the special issue on cognitive bias modification. *Journal of Behavior Therapy and Experimental Psychiatry*, 57, 215. <https://doi.org/10.1016/j.jbtep.2016.03.001>
- Rutherford B. R., & Roose S. P.** (2013). A model of placebo response in antidepressant clinical trials. *American Journal of Psychiatry*, 170, 723–733. <https://doi.org/10.1176/appi.ajp.2012.12040474>
- Rutherford B. R., Wall M. M., Brown P. J., Choo T.-H., Wager T. D., Peterson B. S., ... Roose S. P.** (2017). Patient expectancy as a mediator of placebo effects in antidepressant clinical trials. *The American Journal of Psychiatry*, 174, 135–142. <https://doi.org/10.1176/appi.ajp.2016.16020225>
- Salemink E., van den Hout M., & Kindt M.** (2009). Effects of positive interpretive bias modification in highly anxious individuals. *Journal of Anxiety Disorders*, 23, 676–683. <https://doi.org/10.1016/j.janxdis.2009.02.006>
- Schardt C., Adams M. B., Owens T., Keitz S., & Fontelo P.** (2007). Utilization of the PICO framework to improve searching PubMed for clinical questions. *BMC Medical Informatics and Decision Making*, 7(1), 16. <https://doi.org/10.1186/1472-6947-7-16>
- Tran T. B., Hertel P. T., & Joormann J.** (2011). Cognitive bias modification: Induced interpretive biases affect memory. *Emotion*, 11(1), 145–152. <https://doi.org/10.1037/a0021754>
- Vazquez C., Blanco I., Sanchez A., & McNally R. J.** (2016). Attentional bias modification in depression through gaze contingencies and regulatory control using a new eye-tracking intervention paradigm: Study protocol for a placebo-controlled trial. *BMC Psychiatry*, 16(1), 439. <https://doi.org/10.1186/s12888-016-1150-9>
- Wiers R. W., Eberl C., Rinck M., Becker E. S., & Lindenmeyer J.** (2011). Retraining automatic action tendencies changes alcoholic patients’ approach bias for alcohol and improves treatment outcome. *Psychological Science*, 22, 490–497. <https://doi.org/10.1177/0956797611400615>
- Williams A. D., Blackwell S. E., Mackenzie A., Holmes E. A., & Andrews G.** (2013). Combining imagination and reason in the treatment of depression: A randomized controlled trial of internet-based cognitive bias modification and internet-CBT for depression. *Journal of Consulting and Clinical Psychology*, 81, 793–799. <https://doi.org/10.1037/a0033247>
- Woud M. L., & Becker E. S.** (2014). Editorial for the special issue on cognitive bias modification techniques: An introduction to a time traveller’s tale. *Cognitive Therapy and Research*, 38, 83–88. <https://doi.org/10.1007/s10608-014-9605-0>
- Woud M. L., Holmes E. A., Postma P., Dalgleish T., & Mackintosh B.** (2012). Ameliorating intrusive memories of distressing experiences using computerized reappraisal training. *Emotion*, 12, 778–784. <https://doi.org/10.1037/a0024992>