

Systematic Explorations of Methodological Parameters of the Implicit Relational Assessment Procedure (IRAP)



*Thesis submitted to the Department of Psychology, Faculty of Science,
in fulfilment of the requirements for the degree of Masters of Science,
National University of Ireland, Maynooth.*

Colin Harte, B.A. (Hons.)

September 2015

Research Supervisor: Dr. Yvonne Barnes-Holmes

Head of Department: Dr. Andrew Coogan

Table of Contents

Acknowledgements	ii
List of Tables.....	iii
List of Figures	iv
List of Appendices	v
Abstract	vi
Chapter 1. <i>General Introduction:</i>	1
Chapter 2. <i>Experiment 1: Exploring the Use of Natural Language Statements</i>	18
Chapter 3. <i>Experiment 2: Manipulating Rule Order</i>	36
Chapter 4. <i>Experiments 3a and 3b: Manipulating Block Length and Block Number</i>	52
Chapter 5. <i>General Discussion</i>	76
References	90
Appendices	100

Acknowledgements

I'd like to thank Yvonne first and foremost for her endless support, wisdom and guidance, and the fantastic opportunity to work with her. To say it has been a pleasure and an amazing experience is definitely an understatement. You and Dermot are a constant source of inspiration and motivation.

I'd also like to thank Ciara McEnteggart and Ian Hussey for their support and endless help, whilst still somehow managing to juggle the countless other things that they had to deal with at the same time. In particular Ciara for your constant guidance. You never cease to amaze me (and I'm still not convinced you're actually human).

Aileen, Martin and Deirdre, the best team and friends anyone could ever be lucky to end up with. I fear to think where I'd be without the three of you! You're three of a kind and I'm so privileged to call you friends. Thanks for being there and being your crazy selves.

To my friends at home, particularly David, Ciaran, Caroline and Ellen. Thanks for your constant support in everything always.

And finally I'd like to thank my parents. Mam and Dad, your unwavering support in everything I decide to do (or not do) is a constant source of encouragement and inspiration in my life and this is no exception. You're truly amazing.

Thank you to everyone above and everyone else that have been involved in this project, directly or indirectly. Thanks for putting up with me; sure there'd be no thesis without you!

List of Tables

- Table 1: The stimuli used in the Typical and Natural Language IRAPs in Experiment 1, including labels, targets and response options
- Table 2: The stimuli used in the Self-esteem IRAP in Experiment 2, including labels, targets and response options
- Table 3: Mean scores with standard error values for the DASS sub-scales and the RSES used in Experiment 2
- Table 4: The stimuli used in the Spider Fear IRAP in Experiment 3a, including labels, targets and response options

List of Figures

- Figure 1: A comparison of a Typical IRAP and a Natural Language IRAP trial presented in Experiment 1
- Figure 2: Examples of the four trial-types in the Typical IRAP presented in Experiment 1
- Figure 3: Examples of the four trial-types in the Natural Language IRAP presented in Experiment 1
- Figure 4: Mean *D*-scores, with standard error bars, for the four trial-types on the Typical and Natural Language IRAPs in Experiment 1
- Figure 5: Examples of the four trial-types in the Self-esteem IRAP in Experiment 2
- Figure 6: Mean *D*-scores, with standard error bars, for the four trial-types in the Self-esteem IRAP in Experiment 2
- Figure 7: Examples of the four trial-types in the Spider Fear IRAP presented in Experiment 3a
- Figure 8: A brief illustration of the methodological differences between the Typical Switching IRAP and the Fast Switching IRAP in Experiment 3a
- Figure 9: Mean *D*-scores, with standard error bars, for the four trial-types in the Spider fear IRAPs in Experiment 3a
- Figure 10: A brief illustration of the methodological differences between the Typical Length IRAP and the Double Length IRAP in Experiment 3b
- Figure 11: Mean rolling *D*-scores, with standard error bars, for the four trial-types in the Spider fear IRAPs at the three, four, five and six block pair length in Experiment 3b

List of Appendices

- Appendix 1: National Adult Reading Test (NART) as employed in Experiment 1
- Appendix 2: Participant Consent Form
- Appendix 3: WAIS Full Scale, Verbal and Performance IQ predicted from number of NART errors as employed in Experiment 1
- Appendix 4: Depression, Anxiety and Stress SCALE (DASS 21) as employed in Experiment 2
- Appendix 5: Rosenberg Self Esteem Scale (RSES) as employed in Experiment 2
- Appendix 6: Fear of Spiders Questionnaire (FSQ) as employed in Experiments 3a and 3b
- Appendix 7: Behavioural Approach Task (BAT) scoring sheet as employed in Experiments 3a and 3b
- Appendix 8: Participant Debrief (Experiments 3a and 3b)

Abstract

Abstract

Having briefly reviewed implicit measures, including the Implicit Relational Assessment Procedure (IRAP), and examined the small number of published studies that have explored the possible role of methodological features of the task itself, the current thesis set about investigating other procedural parameters that may influence IRAP effects. The increasing use of the IRAP in clinically-relevant domains made this task seem even more pressing and necessary. At one level, we simply hoped to contribute to the small literature on the possible influence of these factors on IRAP outcomes. At another level, we hoped that these exploratory methodological analyses would help to pave the way for more robust use of the IRAP in applied and clinical domains.

The current thesis comprised of four studies (Experiments, 1, 2, 3a and 3b). In order to permit useful comparisons of the response patterns of the various procedural manipulations, there are strong overlaps in the experimental designs and analytic strategies employed across the four studies. Namely, all four studies involve a Typical IRAP as a control procedure and this is systematically compared to IRAPs with specific procedural modifications.

In Experiment 1 (Chapter 2), we compared a *Typical IRAP* with a *Natural Language IRAP* in which the label and target stimuli were combined, rather than presented separately as is usually the case. This manipulation was simply to determine whether the use of more complex label-target combinations (that would perhaps be better suited to clinical research) would produce different IRAP effects than those typically observed with separate label and target stimuli. Overall, patterns of responding were very similar across both IRAPs and the findings, therefore, from

Experiment 1 supported existing evidence which suggests that combining the label and target stimuli produces sound IRAP effects. However, such a manipulation involving the labels and targets appeared to have little or no direct impact on the outcomes.

In Experiment 2 (Chapter 3), we employed a *Self-esteem IRAP* and manipulated the order in which the self vs. others rules were presented. In the *Self-Rule First Condition*, participants were presented with the rules as follows: the self-positive rule followed by the others-negative rule, while in the *Others-Rule First Condition* participants were presented with the others-positive rule followed by the self-negative rule. This manipulation sought to determine whether the simple order in which the rules are presented influenced IRAP outcomes. The results indicated that patterns of responding appeared to differ somewhat, but not significantly, between conditions, especially on Self-Negative and Others-Positive trial-types. Overall, the findings from Experiment 2 supported some existing evidence which suggests that rules exert some influence on IRAP effects, and thus warrant empirical attention.

In Experiment 3a (Chapter 4), we employed the *Spider Fear IRAP* and manipulated the length of the test blocks presented between rules. To do so, we compared a *Fast Switching IRAP* and a *Typical Switching IRAP*. Specifically, the Fast Switching IRAP halved the length of the test blocks between rules, relative to the Typical Switching IRAP (and IRAPs generally), such that each block now contained only 16 trials, and not 32. Again, this manipulation sought to determine whether the typical rate at which blocks switch exerts its own influence on outcomes normally observed with the IRAP. Patterns of responding were somewhat similar across the two IRAPs. Stronger responding was recorded on the Fast Switching IRAP on the Fear-Spiders and Approach-Spiders trial-types, while the Typical Switching IRAP was

stronger on the Fear-Nature and Approach-Nature trial-types. The difference on Fear-Spiders was significant. While there were differences in the two outcomes, one could not argue that either condition yielded a better IRAP outcome than the other.

In Experiment 3b (Chapter 4), we again employed the Spider Fear IRAP and now manipulated the number of the test blocks presented. To do so, we compared a *Double Length IRAP* and a *Typical Length IRAP*. That is, instead of the typical 6 test blocks (i.e. 3 per rule), the Double Length IRAP presented 12 test blocks (6 per rule), although each test block contained the typical 32 trials. Overall, therefore, this modified IRAP presented a total of 384 trials -- that is double the number of trials presented in the Typical Length IRAP (i.e. 192). Again, this manipulation sought to determine whether the typical number of blocks exerts its own influence on outcomes normally observed with the IRAP. Patterns of responding were very similar across the Typical and Double Length IRAPs. In both cases, participants showed consistent responding with increasing numbers of block pairs, with no significant differences across each trial-type of increasing block pairs.

The current thesis, therefore, undertook a number of systematic manipulations of various procedural parameters of the IRAP. Overall, one must conclude that these exerted practically no significant influence on the observed IRAP effects, although some variations were detected. These changes are discussed in the context of the existing literature and particularly with regard to more extensive use of the IRAP in the future.

Chapter 1

General Introduction

Chapter 1

General Introduction

Research interest in implicit cognition has expanded at a considerable rate in the last 10-20 years and appears to be the result of two key factors: 1. Growing concerns about the limitations of explicit measures (in particular self-report procedures), especially their vulnerability to socially desirable responding (Power, Barnes-Holmes, Barnes-Holmes, & Stewart, 2009) and 2. increased belief in the role of implicit attitudes and bias in most aspects of higher cognition, particularly when trying to tap into information that is not accessible through introspective means (Greenwald & Banaji, 1995).

A number of measures of implicit attitudes have been developed in recent years. These include affective priming tasks (Fazio, Jackson, Dunton, & Williams, 1995), semantic priming tasks (Wittenbrink, Judd, & Park, 1997), the Go/No-go Association Task (GNAT; Nosek & Banaji, 2001), the Extrinsic Affective Simon Test (EAST; De Houwer, 2003), the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998) and the Implicit Relational Assessment Procedure (IRAP; Barnes-Holmes, Barnes-Holmes, Power, Hayden, Milne, & Stewart, 2006).

In short, affective priming tasks involve presenting a prime, such as the name of an attitude object, followed by a prompt for a quick evaluation of the prime via certain evaluative response options (e.g. good and bad). The response latency following being asked to make this evaluation is suggested to represent participants' attitudes toward the priming stimulus. Semantic priming tasks are almost identical to evaluative priming tasks, but differ in the response required. For example, the semantic decision task asks participants to quickly categorise words based on their perceived

semantic meaning, rather than evaluative meaning (e.g. words of professions followed by male or female pronouns; Banaji & Hardin, 1996).

Similar to typical priming tasks, the GNAT assesses the strength of associations between a pre-specified target category stimulus (e.g. insects) and some form of attribute, such as an evaluation (e.g. good or bad). It also differs from typical priming tasks in that the prime category under investigation (e.g. race) is generally made explicit to participants (De Houwer, Teige-Mocigemba, Spruyt, & Moors, 2009). In the GNAT, participants are required to respond in some cases by pairing the target category with one label response (e.g. insects-bad) and in other cases in the opposite manner (e.g. insects-good). In each of these conditions participants are also presented with distracting stimuli which they are required to ignore. It is hypothesised that distinguishing the targets and labels from the distractors should be easier in conditions in which the coherent implicit bias is held. The difference in accuracy between these conditions is also believed to represent an individual's implicit attitude.

The EAST operates on a similar premise, but differs in one key aspect. This procedure assesses performance by comparing trials within a single condition rather than a comparison of trials across two different conditions. On the EAST, responding is based on categorising a non-evaluative aspect of a valenced target word as positive or negative (e.g. respond "good" whenever a person-related word is presented and "bad" whenever an animal-related word is presented). It is argued that responding should be faster when the valence of the target word is coherent with the required response (e.g. responding "good" to the target word "friend" than responding "good" to the target word "enemy"; De Houwer et al., 2003). This difference in response times is believed to represent implicit attitudes.

The Implicit Association Test (IAT)

Undoubtedly, the most widely used and successful implicit measure is the IAT (e.g. Greenwald, Poehlman, Uhlmann, & Banaji, 2009; but see Gawronski and De Houwer, 2014, for a comprehensive review).

Similar to the other implicit measures, the IAT is an automated latency-based task in which participants pair words and/or images. In simple terms, the stimuli are comprised of, for example, target words (such as names of insects and flowers) and label words (such as various synonyms of ‘good’ and ‘bad’). Consistent with latency-based measures generally, the basic assumption of the IAT is that differences in reaction times in responding to the stimuli represent associations between these stimuli in memory. That is, participants will respond more quickly when associating stimuli that are consistent with their existing implicit beliefs (e.g. pairing “flowers” and “good” more quickly than “flowers” and “bad”). Consider one of the first IATs conducted by Greenwald and colleagues (1998) using flowers and insect stimuli. Participants were presented with flower names (e.g. “rose” and “tulip”) and insect names (e.g. “bee” and “wasp”) as target stimuli, and positively valenced (e.g. “happy”, “peace”) or negatively valenced words (e.g. “rotten”, “ugly”) as label stimuli. Participants were asked to associate or categorise these stimulus pairs in various ways that included pairing flowers and positive, insects and negative, flowers and negative, and insects and positive. As expected, participants more readily paired flowers with positive and insects with negative than flowers with negative and insects with positive. Crucially, this indicated that the flower-positive association was stronger than the flower-negative association, as was the insects-negative association relative to insects-positive. This difference is what is referred to as *the IAT effect*. The IAT data, from which these effects are observed, are analysed using the *D*-algorithm which converts

participant responses into *D*-scores (Greenwald, Nosek, & Banaji, 2003). This controls for any variations in participant response speed that may be the result of varying levels of cognitive ability. These variations may serve to confound the data when analysing differences between groups (Vahey, Barnes-Holmes, Barnes-Holmes, & Stewart, 2009).

In more recent years, the range of domains in which the IAT has been used has expanded rapidly and now includes, racial bias (McConnell & Leibold, 2001), self-esteem (Greenwald & Farnham, 2000), political voting (Greenwald, Smith, Sriram, Bar-Anan, & Nosek, 2009) and even predicting therapeutic outcomes (Gamer, Schmukle, Luka-Krausgrill, & Egloff, 2008). Consider the race study by McConnell and Leibold presented to white participants using target name stimuli typically associated with Black people (e.g. “Lamar”) and other names typically associated with White people (e.g. “Ian”), along with positive (e.g. “wonderful”) and negative labels (e.g. “disgusting”). Participants also completed a number of explicit measures of racial bias. As has been found repeatedly in race IATs, white participants more readily associated white related names with positive words, thus showing a pro-white bias, which also correlated with the explicit measures.

Among the more clinically-relevant phenomena assessed by the IAT are several studies of the implicit attitudes that may be involved in various patterns of psychological suffering, such as anxiety and depression. For example, Nock et al. (2010) studied suicidal behaviour in 157 participants in a clinical setting using a Death/Suicide IAT. Participants paired death/suicide words (e.g. “lifeless” and “suicide”) and life words (e.g. “survive” and “breathing”) with self (e.g. “me”) and other (e.g. “they”) words. The results indicated that participants who had previously attempted suicide more readily associated death/suicide with self than participants

with no suicidal history, but who were nonetheless clinically distressed. More importantly, this study showed the predictive validity of the IAT in that the effect for suicidal participants predicted risk of a suicide attempt in a six-month follow-up. Similar effects and predictive outcomes have been recorded with the IAT in other clinical domains (e.g. anxiety; see Egloff & Schmukle, 2002).

In spite of a considerable array of well-conducted and well-supported studies, and impressive findings overall, a number of authors have emphasised various conceptual and methodological weaknesses associated with the procedure (e.g. Barnes-Holmes et al., 2006; De Houwer, 2002). Central to the conceptual debate is the criticism that the IAT can only measure associations and thus fails to tap into the complexities of propositions or relations (De Houwer, 2002). That is, the format of the screen pairings does not permit access to the directionality or strength of the target associations, and thus the measure's ability to assess complex beliefs (especially those that control overt behaviour) is likely to be limited (Barnes-Holmes et al., 2006). While the development of the EAST represented an attempt to target propositions, the IRAP was specifically developed to target relations (although at a conceptual level, there is much overlap between a proposition and a relation, see De Houwer, 2014).

The Implicit Relational Assessment Procedure (IRAP)

It is perhaps surprising that, unlike most other implicit measures, the IRAP does not hail from the fields of cognitive or social psychology, but is instead based on a modern behavioural theory of human language and cognition known as Relational Frame Theory (RFT; Hayes, Barnes-Holmes, & Roche, 2001). While a detailed account of the theory is beyond the scope of the current thesis, it is useful to summarise the core tenets of RFT because it assists in understanding IRAP effects, and how they

can be most precisely and most meaningfully interpreted (for book-length reviews of RFT, see Dymond & Roche, 2013; Hayes et al.).

According to RFT, language and cognition emerge from our ability to relate stimuli in ways that are *not* based on their formal properties (e.g. shape or colour), but are based instead on arbitrarily applicable relations with other stimuli. For example, when offered one of two coins, a young child will likely select the larger, because of the superiority in physical size. However, if for example the coins were a 10 cent coin and a 5 cent coin, an older child will likely select the 10 cent coin, because even though it is physically smaller, the verbal community has established (arbitrarily) that it is greater in monetary value. Hence, the verbal functions of the smaller coin, as established by the social community, dominate the functions of physical superiority in terms of the child's responding. For RFT, it is this process (referred to as arbitrarily applicable relational responding) which is key to verbal sophistication and is involved in *all* verbal behaviour, as established through our developmental histories (Barnes-Holmes, Barnes-Holmes, Stewart, & Boles, 2010).

For RFT, this history is established by multiple exemplar training and natural language interactions. It is through this that individuals learn to relate stimuli and events that were never before related together, but one can do so based indirectly upon relevant histories of relating stimuli in similar ways. From a measurement perspective, RFT is interested in targeting these relational responses directly and exploring the types of verbal histories that give rise to these repertoires. And the IRAP was designed specifically for this purpose (Barnes-Holmes et al., 2010).

Consistent with the other automated latency-based implicit measures, the IRAP presents word and/or image pairings. In this case, responding to the stimulus pairs is controlled by one of two pre-specified rules given to participants prior to each

block of trials. One rule is deemed to be consistent with participants' general verbal histories (e.g. flowers are positive), while the other rule is inconsistent (e.g. flowers are negative). As with the other measures, an IRAP effect is defined by the latency difference between consistent and inconsistent responding. There are now over 50 published studies on the IRAP and the number of domains of interest has increased steadily.

Consider a study by Cullen, Barnes-Holmes, Barnes-Holmes, and Stewart (2009) that used an Age IRAP to measure the ageist beliefs of young adults. Across trials, participants were required to relate the label phrases "young people" and "old people" with either positive (e.g. "energetic") or negative (e.g. "weary") target words via the response options "Similar" and "Different". On consistent blocks, responding correctly involved relating: young people-positive-similar, young people-negative-different, old people-positive-different and old people-negative-similar. In contrast, correct responding on inconsistent blocks involved relating: young people-positive-different, young people-negative-similar, old people-positive-similar and old people-negative-different. In simple terms, consistent responding was based on the view that young participants would more readily evaluate young people as positive and old people as negative, and the results confirmed that this was the case.

The study by Cullen et al. (2009) summarised highlights an important feature of data analysis in the IRAP. That is, responding is divided according to four stimulus-combinations known as trial-types. In short, rather than calculating a single overall *D*-score as in the IAT (although this is possible and used in some analytical contexts), four trial-type *D*-scores are more typically calculated. Consider a study by Barnes-Holmes, Murphy, Barnes-Holmes, and Stewart (2010) that presented a Race IRAP to white male participants. Across trials, participants related images of white or black

males holding guns with positive (e.g. “safe”) or negative (e.g. “dangerous”) words. Hence, the four trial-types in this study, based on consistent and inconsistent blocks were: White-Safe; Black-Dangerous; White-Dangerous; and Black-Safe. On consistent blocks, responding correctly involved relating: white-safe-true, white-dangerous-false, black-safe-false and black-dangerous-true. In contrast, correct responding on inconsistent blocks involved relating: white-safe-false, white-dangerous-true, black-safe-true and black-dangerous-false. In simple terms, consistent responding was based on the view that white participants would more readily evaluate images of white males as safe and images of black males as dangerous. And the results confirmed that this was the case, thus showing a pro-white bias. It is important to emphasise that the four trial-types that typify the IRAP give more precision than the IAT and thus allow the researcher to determine whether a performance shows a pro-white bias, a non-pro-white bias, an anti-black bias or a non-anti-black bias.

Using the IRAP to Study Clinical Phenomena. Similar to the literature on the IAT, there is now a growing body of work using the IRAP to explore the implicit cognitions or relations that may underpin various social and clinical phenomena, including race (Barnes-Holmes, et al., 2010), depression (Hussey & Barnes-Holmes, 2012), obsessive compulsive disorder (OCD, Nicholson, Dempsey, & Barnes-Holmes, 2014) and drug dependence (Carpenter, Martinez, Vadhan, Barnes-Holmes, and Nunes, 2012). One such domain that has attracted considerable IRAP attention is spider fear. For example, Nicholson and Barnes-Holmes (2012) used a Spider Fear IRAP to compare responding to spiders versus scenes of nature, with target stimuli that specified fearful reactions versus approach responding. That is, participants paired words representing fear (e.g. “creeps me out”) or approach (e.g. “I could approach”) behaviours with pictures of spiders and scenes of nature, according to the rules

“Spiders scare me and I can approach nature” vs. “I can approach spiders and nature scares me”. The results indicated strong IRAP effects in terms of fear-spider relations among participants who had self-reported high spider fear on the Fear of Spiders Questionnaire (FSQ) and strong approach-nature relations overall. Responding on the former trial-type correlated significantly and positively with FSQ outcomes (i.e. strong fear-spider relations predicted high FSQ score) and negatively with responding on a behavioural approach task (BAT) involving a live tarantula (i.e. strong fear-spider relations predicted low approach behaviours in the BAT).

In a recently published meta-analysis, Vahey, Nicholson, and Barnes-Holmes (2015) reviewed the 15 studies ($N = 494$) published thus far that have employed clinically-focused IRAPs. When comparing their findings with similar reviews of IAT work in the clinical domain, the authors not only concluded that the IRAP compared favourably in this regard, but they also proposed that the procedure shows potential “as a tool for clinical assessment” (p.64).

Methodological Parameters of the IRAP. While IRAP research has, similar to the IAT, moved at considerable pace towards an ever-increasing range of domains, there remain methodological questions about parameters of the basic procedure itself which need to be addressed before researchers can be sure that clinical effects, for example, are genuine clinical effects and not spurious procedural artefacts. This has been a considerable (but often not primary) focus in IAT research. For example, Klauer and Mierke (2005) suggested that the length of IAT blocks has some influence on the effects (see also Back, Schmukle, & Egloff, 2005). Similarly, in terms of block order, Nosek, Greenwald, and Banaji (2007) suggested that this presentation variable exerted some influence over responding (see also De Houwer, 2001; Rothermund, Teige-Mocigemba, Gast, & Wentura, 2009).

Very few IRAP studies to date have asked similar methodological questions about the typical format in which the IRAP is presented. Among these, Campbell, Barnes-Holmes, Barnes-Holmes, and Stewart (2011) systematically manipulated the sequence of stimulus presentations and the locations of the response options in a simple pleasant-unpleasant IRAP. Consider a typical IRAP in which the label stimulus is randomised across trials (so that the same label does not appear on consecutive trials) and the response options also appear randomly on either the left- or right-hand sides of the screen. Campbell et al. compared the fixing (non-randomisation) of both of these features with the randomisation of both (e.g. in the Fixed-Fixed Condition, the sequence of the label stimuli was not random and the locations of the response options were fixed). The results indicated that some randomisation of the label stimuli (i.e. the Random [labels]-Fixed [response options] Condition) produced the strongest IRAP effects and the locations of the response options had limited influence.

Numerous published IRAP studies, including several targeting clinically-relevant phenomena, have also presented complex statements as stimuli on the task (e.g. Remue, De Houwer, Barnes-Holmes, Vanderhasselt, & De Raedt, 2013). And indeed, one of the core rationales for using the IRAP over an alternative implicit measure is that it is one of the few that can accommodate proposition-like statements (Gawronski & De Houwer, 2014). Consider the Future Thinking IRAP presented by Kosnes, Whelan, O'Donovan, and McHugh (2013) which contained the phrases "I expect/I don't expect" as label stimuli and the single word positive/negative targets (e.g. "Love" or "Worry"), presented to undergraduates with and without sub-clinical depression. The results indicated that the non-depressed sub-group showed stronger confirmation of the future-positive relation and denial of the future-negative relation, relative to those with self-reported sub-clinical depression. Parling, Cernvall, Stewart,

Barnes-Holmes, and Ghaderi (2012) similarly used phrases as labels (e.g. “I want to be/I don’t want to be”) and single word target stimuli (e.g. “Thin/Fat”) in a Body Image IRAP, presented to a sample with anorexia vs. controls. In this study, the group with anorexia showed stronger pro-fat-others and anti-fat-self responding than controls.

Several recent studies have also employed *target stimuli* that comprised more complex phrases than single words. Consider an IRAP on cocaine dependence presented by Carpenter et al. (2012). This presented the phrases “With cocaine/No cocaine” as labels and positive/negative phrases as targets (e.g. “I am friendly/I am alone”), presented to individuals in treatment for cocaine dependence. The results indicated that individuals with poorer treatment outcomes showed stronger confirmation of cocaine-positive and denial of cocaine-negative relations. Similarly complex target phrases were presented by Hooper, Vilatte, Neofotistou, and McHugh (2010) in an IRAP on experiential avoidance, and by Hussey and Barnes-Holmes (2012) in an IRAP on depression.

In another attempt to systematically explore presentation features of the IRAP, De Houwer, Heider, Spruyt, Roets, and Hughes (2015) created a variation of the IRAP, which they referred to as the Relational Responding Task (RRT). The key feature that differentiates the RRT from a typical IRAP is that the RRT combines the label and target stimuli into a single stimulus (e.g. instead of presenting them as two separate words, they appear together as a single sentence in the centre of the screen). Consider a typical IRAP in which the label stimulus appears at the top of the screen (e.g. “Flemish people”) and the target stimulus appears in the centre of the screen (e.g. “smart”). De Houwer et al. combined these two stimuli and presented them as a single phrase in the centre of the screen (e.g. “Flemish people are smarter than immigrants”).

Although the findings indicated the expected pattern of ethnic bias (i.e. pro-Flemish responding), it is difficult to determine whether the label-target combination influenced this effect because multiple procedural modifications were undertaken simultaneously, and these were not directly compared with a typical IRAP. However, the study does suggest that combining the label and target stimuli in a single sentence may not strongly influence outcomes relative to a typical IRAP.

In a similar vein, Barnes-Holmes et al. (2010) investigated the optimal response latency criterion across two Race IRAPs. In one study, the latency criterion remained at 3000ms as in a typical IRAP of that time, while the second study reduced this criterion to 2000ms. The results indicated that the stronger race effect (i.e. a pro-white bias) was recorded in the study with the shorter latency criterion. Although this effect was recorded with a Race IRAP, and various adjustments to the latency criterion are often needed with different IRAPs, it does highlight the influence of this variable on IRAP effects.

In a recently published study, O'Shea, Watson, and Brown (in press) systematically investigated the rules for IRAP responding. Participants were assigned to one of three conditions. The Standard Framing Condition was identical to a typical IRAP and participants were presented with the two response rules in the normal way. However, in the Positive Framing Condition, the positive response rule was always presented before the negative rule (e.g. thin person-positive followed by fat person-negative). Similarly, in the Negative Framing Condition, the negative response rule was always presented first (e.g. fat person-negative followed by thin person-positive). In a complex experimental design, the researchers compared these conditions across four IRAPs: a body weight IRAP, a nature IRAP, a social system IRAP and a nonsense syllable IRAP. The results indicated that the strongest IRAP effects were recorded in

the Positive Framing Condition, but only in the nonsense syllable and social systems IRAPs. These outcomes are interesting because they suggest that more established verbal relations, such as those involved in body weight and nature, are less sensitive to the way in which the rules are presented in the IRAP.

Another study has also explored the potential influence of response rules in the IRAP, in what have been described as faking manipulations. McKenna, Barnes-Holmes, Barnes-Holmes, and Stewart (2007) systematically compared three conditions on a simple pleasant-unpleasant IRAP. The No Faking Condition was identical to a typical IRAP. In the Faking with Strategy Condition, participants were encouraged to fake their responses and were given various instructions on how to do this (e.g. slowing down on consistent trials and speeding up on inconsistent trials). In the Faking with No Strategy Condition, participants were again encouraged to fake their responses, but were not given explicit instructions on how this might be done. The results indicated no differences among conditions, this suggesting little or no influence by the rule manipulations.

The Current Research

Having briefly reviewed implicit measures, including the IRAP, and examined the small number of published studies that have explored the possible role of methodological features of the task itself, the current thesis set about investigating other procedural task parameters that may influence IRAP effects. The increasing use of the IRAP in clinically-relevant domains made this task seem even more pressing and necessary. At one level, we simply hoped to contribute to the small literature on the possible influence of these factors on IRAP outcomes. At another level, we hoped

that these exploratory methodological analyses would help to pave the way for more robust use of the IRAP in applied and clinical domains.

The current thesis comprised of four studies (Experiments, 1, 2, 3a and 3b). In order to permit useful comparisons of the response patterns of the various procedural manipulations, there are strong overlaps in the experimental designs and analytic strategies employed across the four studies. Namely, all four studies involve a Typical IRAP as a control procedure and this is systematically compared to IRAPs with specific procedural modifications.

In Experiment 1 (Chapter 2), we compared a Typical IRAP with a Natural Language IRAP in which the label and target stimuli were combined rather than presented separately, as is usually the case. This manipulation was simply to determine whether the use of more complex label-target combinations (that would perhaps be better suited to clinical research) would produce different IRAP effects than those typically observed with separate label and target stimuli.

In Experiment 2 (Chapter 3), we employed a Self-esteem IRAP and manipulated the order in which the self vs. others rules were presented. In the Self-Rule First Condition, participants were presented with the rules as follows: the self-positive rule followed by the others-negative rule, while in the Others-Rule First Condition participants were presented with the others-positive rule followed by the self-negative rule. This manipulation sought to determine whether the simple order in which the rules are presented influenced IRAP outcomes.

In Experiment 3a (Chapter 4), we employed the Spider Fear IRAP and manipulated the length of the test blocks presented between rules. To do so, we compared a Fast Switching IRAP and a Typical Switching IRAP. Specifically, the Fast Switching IRAP halved the length of the test blocks between rules, relative to the

Typical Switching IRAP (and IRAPs generally), such that each block now contained only 16 trials, and not 32. Again, this manipulation sought to determine whether the typical rate at which blocks switch exerts its own influence on outcomes normally observed with the IRAP.

In Experiment 3b (Chapter 4), we again employed the Spider Fear IRAP and now manipulated the number of test blocks presented. To do so, we compared a Double Length IRAP and a Typical Length IRAP. That is, instead of the typical 6 test blocks (i.e. 3 per rule), the Double Length IRAP presented 12 test blocks (6 per rule), although each test block contained the typical 32 trials. Overall, therefore, this modified IRAP presented a total of 384 trials -- that is double the number of trials presented in the Typical Length IRAP (i.e. 192). Again, this manipulation sought to determine whether the typical number of blocks exerts its own influence on outcomes normally observed with the IRAP.

In all four studies, we employed Analyses of Variance (ANOVAs) to assess the extent to which trial-type performances varied across the IRAPs. All experiments also employed explicit measures relevant to each type of IRAP and these were primarily used to ensure that the response patterns observed in the IRAPs were representative of the typical population. Specifically, Experiment 1 employed the National Adult Reading Test (NART) to control for IQ. In Experiment 2, we employed the Depression, Anxiety and Stress Scales (DASS) to control for psychological distress and the Rosenberg Self-Esteem Scale (RSES) to control for self-esteem. In Experiments 3a and 3b, we employed the Fear of Spiders Questionnaire (FSQ) to control for spider fear. In these two final experiments, we also presented a Behavioural Approach Task (BAT). As well as ascertaining control, we conducted correlations among the explicit measures and the IRAPs to investigate the potential predictive

validity of the IRAPs across the various domains. One-sample t-tests were conducted to investigate the significance of the *D*-scores in each IRAP. Given that our experimental designs sought to systematically compare IRAPs, we also employed independent t-tests (there were dependent t-tests in Experiment 3b) to compare the trial-type outcomes across the various IRAPs. The similarities in strategy and analyses permitted optimal and most systematic comparisons across the four studies.

Chapter 2
Experiment 1

*Exploring the Use of
Natural Language Statements*

Chapter 2

Introduction

Very few published IRAP studies have asked specific methodological questions about the IRAP, although this has attracted considerable attention in IAT research (Greenwald et al., 2009). Among these, Campbell et al. (2011) systematically manipulated the sequence of stimulus presentations and the locations of the response options in a simple pleasant-unpleasant IRAP, and found that some randomisation of the label stimuli produced the strongest IRAP effects, while the locations of the response options had limited influence. Similarly, in their RRT variation of the IRAP, De Houwer et al. (in press) combined the label and target stimuli into a single stimulus, instead of presenting them as two separate words as is typically the case in the IRAP. While it is difficult to parse out the effects in the latter study, the results did suggest that combining the label and target stimuli may not strongly influence typical IRAP outcomes. Three conclusions could be drawn from this and related research. First, IRAPs will likely need to become more complex in presentation if more clinical domains are to be targeted. Second, given the multiple simultaneous manipulations of the RRT study, it is difficult to decipher the impact of increasing the complexity of the label and target stimuli, especially when they are combined. Third, impressive IRAP effects have been recorded by other researchers who have made the label or target stimuli more complex (e.g. Carpenter et al., 2012; Hooper et al., 2010; Kosnes et al., 2013; Parling et al., 2012; Remue et al., 2013).

Taken together, these studies influenced the design of Experiment 1, in which we compared a Typical IRAP with a *Natural Language IRAP* in which the label and target stimuli were combined rather than presented separately, as is usually the case.

This manipulation was simply to determine whether the use of more complex label-target combinations (that would perhaps be better suited to clinical research) would produce different IRAP effects than those typically observed with separate label and target stimuli.

It was difficult to make predictions about possible differences that may emerge between the Typical and Natural Language IRAPs because this manipulation had not been done before. However, given that previous evidence from similar manipulations and the fact that this was a very simplistic set of stimuli (i.e. responding to fruit and insects and positive or negative), we did not expect to find sizeable differences between performances on these IRAPs.

Method

Participants

Experiment 1 involved 24 undergraduate students from Maynooth University (15 female, 9 male) aged between 19 to 28 years ($M = 20.19$, $SD = 2.20$), all selected by random convenience sampling.

Setting

All aspects of Experiment 1 were conducted in an experimental cubicle at the Department of Psychology, Maynooth University. All participation was on an individual basis. The experimenter was only present during the instructional phase of the study and practice blocks of the IRAP, but vacated the room during completion of the explicit measures and the IRAP test blocks. Participation lasted approximately 30 minutes and breaks were scheduled, if requested, at specific points in the experimental sequence (e.g. between IRAP blocks).

Apparatus and Materials

All aspects of the experiment (including the questionnaire) were completed on a DELL desktop computer. The study involved an explicit (self-report) measure and an implicit measure (two IRAPs). The IRAP was delivered via PsychoPy (1.81.03) which controlled all aspects of stimulus presentation and recording of participant responses. The two IRAPs employed here are referred to as the *Typical IRAP* and the *Natural Language IRAP*. The explicit measure comprised of the National Adult Reading Test (NART: Nelson, 1982).

The National Adult Reading Test (NART). The NART consists of 50 words that are irregular in terms of their grapheme and phoneme correspondences (Nelson, 1982). Participants are required to read aloud and pronounce each word correctly, and accurate responding is based on the latter. The total score is calculated based on the number of errors, hence ranging from 0 to 50. Higher (error) scores indicate lower IQ. The NART has been shown to demonstrate high internal consistency (Crawford & Parker, 1989) and satisfactory inter-rater reliability (O'Carroll, 1987). See Appendix 1.

The Typical IRAP. The Typical IRAP employed here was referred to as such because its screen format was identical to that of almost all published IRAPs (e.g. Barnes-Holmes et al., 2010). That is, the label word stimulus was presented in the top centre of the screen, with the target word stimulus below, and the two response options on the bottom left- and right-hand sides of the screen. This typical screen format is illustrated in Figure 1 (left-hand side).




Figure 1. A comparison of a Typical IRAP and a Natural Language IRAP trial presented in Experiment 1.

The Typical IRAP presented a total of 12 label stimuli, adapted from Nosek and Mahzrin (2001). Six were the names of fruit (e.g. APRICOTS) and six were the names of insects (e.g. CENTIPEDES), see Table 1. Please note that stimuli that actually appeared on-screen in the IRAP are presented throughout the thesis in capital letters. This IRAP also presented a total of 12 target stimuli, adapted from Nosek and Banaji (2001). Six were positive words (e.g. SWEET) and six were negative words (e.g. ROTTEN). Two static response options (TRUE and FALSE) were presented on each trial at the bottom left- and right-hand corners of the screen, respectively. The label and target stimuli comprising the Typical IRAP are provided in Table 1.

Table 1

The stimuli used in the Typical and Natural Language IRAPs in Experiment 1, including labels, targets and response options

Fruit Label	Insects Label
APRICOTS	CENTIPEDES
PEACHES	COCKROACHES
RASBERRIES	MAGGOTS
WATERMELON	SPIDERS
GRAPES	WASPS
BLUEBERRIES	BEETLES
Positive Targets	Negative Targets
JUICY	ROTTEN
SWEET	NASTY
APPETIZING	TERRIBLE
TASTY	REVOLTING
DELICIOUS	FOUL
ENJOYABLE	DISGUSTING
Response Option 1	Response Option 2
TRUE	FALSE

The Typical IRAP comprised of four possible label-target combinations (trial-types): Fruit-Positive; Insects-Positive; Fruit-Negative; Insects-Negative, see Figure 2. Participants responded to these combinations by selecting TRUE or FALSE, assigned to the “d” and “k” keys, respectively.




Figure 2. Examples of the four trial-types in the Typical IRAP presented in Experiment 1. On each trial, a label stimulus (Fruit or Insects), a target stimulus (Positive or Negative) and two response options (TRUE and FALSE) appeared on-screen simultaneously. This generated four trial-types: Fruit-Positive; Fruit-Negative; Insects-Positive; Insects-Negative. Neither the phrases ‘Consistent’ or ‘Inconsistent’ nor the arrows were shown on-screen.

The Natural Language IRAP. The Natural Language IRAP employed here was referred to as such because its screen format differed from a Typical IRAP. That is, on each trial the label word and target word were combined with the word “are” to form a short statement that was presented in the centre of the screen (see Figure 1 right-hand side). Consider the Typical IRAP in which the label “beetles” appeared above the target “delicious”. In contrast, in the Natural Language IRAP, these two stimuli were combined with “are” to form the statement “beetles are delicious”. All of

the stimuli used in the Natural Language IRAP were identical to the Typical IRAP (see Table 1).

The Natural Language IRAP comprised of four possible label-target combinations (trial-types): Fruit-Positive; Insects-Positive; Fruit-Negative; Insects-Negative, see Figure 3. It is important to note that the trial-types in the Natural Language IRAP were identical to the trial-types in the Typical IRAP. Participants responded to these combinations by selecting TRUE or FALSE, assigned to the “d” and “k” keys, respectively.




Figure 3. Examples of the four trial-types in the Natural Language IRAP presented in Experiment 1. On each trial, a label stimulus (Fruit or Insects), a target stimulus (Positive or Negative) and two response options (TRUE and FALSE) appeared on-screen simultaneously. This generated four trial-types: Fruit-Positive; Fruit-Negative; Insects-Positive; Insects-Negative. Neither the phrases ‘Consistent’ or ‘Inconsistent’ nor the arrows were shown on-screen.

Responses on trials from both IRAPs were defined as either consistent or inconsistent in terms of whether they were ‘pro-fruit’ or ‘pro-insects’. Consider, for example, a pro-insects block of trials in which participants are asked to respond positively to insects and negatively to fruit. When the trial presented an insect label (e.g. “CENTIPEDES) with a positive target (e.g. SWEET), TRUE was the correct response because it was consistent (i.e. pro-insect) and FALSE was inconsistent (not pro-insect). Now consider a pro-fruit block in which participants are asked to respond positively to fruit and negatively to insects. When the trial presented the fruit label with a positive target, TRUE was consistent and FALSE was inconsistent.

Procedure

The experimental sequence comprised of three stages. The presentation of the two IRAPs across Stages 1 and 2 was counterbalanced across participants, with participants being assigned to each IRAP semi-randomly. That is, half of the participants were presented with the Typical IRAP in Stage 1, followed by the Natural Language IRAP in Stage 2. For the remaining half of the participants, this sequence was reversed. All participants were presented with the NART in Stage 3. For illustrative purposes, the section below describes the experimental procedure presented to participants exposed to the Typical IRAP in Stage 1 and the Natural Language IRAP in Stage 2.

Stage 1: Typical IRAP. Prior to commencing the practice trials, participants were verbally instructed that each IRAP trial would present a word on the top of the screen, along with a word in the centre, and that their task was to respond with TRUE or FALSE to these pairings in accordance with the rule presented at the beginning of the block. Participants were informed that the rule would switch during the next block,

so that they would then be asked to respond in the opposite manner. That is, the ‘correct’ and ‘incorrect’ patterns of responding (depending upon the rule provided) alternated across blocks of trials. These instructions also highlighted two other key features of the task: the criterion for high levels of accurate responding (i.e. 80%) and the criterion for responding very quickly (<2,000 ms.).

Selecting the ‘correct’ response option removed all stimuli from the screen for a 400ms interval before the next trial appeared. Selecting the ‘incorrect’ response option was consequted with a red ‘X’ in the middle of the screen directly below the target. The next trial only appeared when a correct response was emitted.

As is always the case, IRAP blocks alternate between consistent and inconsistent responding (hence, participants must switch responding after each block), according to the rule specified at the beginning of each block. The Typical IRAP employed here comprised of two rules for responding. Rule A “Fruit taste good and insects taste bad” controlled consistent responding, while Rule B “Fruit taste bad and insects taste good” controlled inconsistent responding. Hence, correct responding involved switching from block to block (i.e. A, B, A, etc.) The actual order of the rules was counterbalanced across participants, such that half were exposed to a Rule A consistent block first, followed by Rule B block, while the reverse was the case for the remaining participants.

The IRAP always commenced with a minimum of two practice blocks and participants were required to achieve $\geq 80\%$ correct and a response latency ≤ 2000 ms for each. If they failed to achieve both performance criteria, they were provided with automated feedback and practice blocks continued to a maximum of eight blocks (i.e. four consistent blocks and four inconsistent blocks). Failing to meet the criteria after

eight practice blocks terminated participation. When the criteria were reached on a pair of practice blocks, participants proceeded automatically to a series of six test blocks (irrespective of performance, although participants continued to receive feedback on block performances). The program automatically recorded response accuracy (based on the first response emitted on each trial) and response latency (time in ms between trial onset and emission of correct response) for each participant on each trial. At the end of the last test block, the following message appeared on-screen: “Thank you. That is the end of this part of the experiment. Please contact the researcher”.

Stage 2: Natural Language IRAP. The Natural Language IRAP employed here was referred to as such because its screen format differed from a typical IRAP. That is, *on each trial the label word and target word were combined with the word “are” to form a short statement that was presented in the center of the screen*, see Figure 1 right-hand side. Consider the Typical IRAP in which the label “beetles” appeared above the target “delicious”. In contrast, in the Natural Language IRAP, these two stimuli were combined with “are” to form the statement “beetles are delicious”. All other procedural aspects of Stage 2 were the same as described above for the Typical IRAP in Stage 1.

Stage 3: Explicit Measures. In Stage 3, all participants completed the NART. They were thereafter fully debriefed as to the nature of the study and measures used and thanked for their participation.

IRAP Data Preparation

The IRAP response latency data were transformed according to the *D*-algorithm outlined in Barnes-Holmes et al. (2010) for both IRAPs conducted

presently. This algorithm is a derivation of the D algorithm developed for the IAT by Greenwald et al. (2003). The *D*-score has been shown to control for extraneous variables, such as general responding speed, age and cognitive ability (Greenwald et al., 2003). The various steps may be summarised as follows. 1. Latencies greater than 10,000ms were removed from the dataset. 2. Participants' performances on each test block were assessed against accuracy and latency mastery criteria (i.e. accuracy \geq 79% and median latency \leq 2000ms: see Barnes-Holmes et al.). 3. Participants who failed to maintain the mastery criteria on the practice blocks or subsequent test blocks had their data excluded from the analysis. 4. If data for a participant showed latencies $<$ 300ms on more than 10% of test blocks, the data were removed. 5. Twelve standard deviations were calculated for each pair of test blocks (four for each of the trial-types). 6. Twenty-four mean latencies were calculated, one for each of the four trial-types in each individual test block. 7. 12 difference scores were computed, one for each of the four trial-types in each pair of test blocks by subtracting the consistent 'Fruit-Positive' blocks from the corresponding inconsistent 'Fruit-Negative' blocks. 8. The difference scores were divided by their associated standard deviations from Step 5, resulting in 12 *D*-scores, one for each trial-type in each of the test block pairs. 9. Four final *D*-scores were calculated by averaging the three scores for each trial-type across the three test-block pairs. One *D*-score was calculated for each of the four trial-types in both IRAPs. The *D*-scores for trial-types 3 and 4 in each IRAP were reversed for ease of interpretation. Positive scores indicate "*positive*" or "*not-negative*" responding, whereas negative scores indicate "*negative*" or "*not-positive*" responding.

Ethical Considerations

All aspects of Experiment 1 adhered to the ethical guidelines outlined by the British Psychological Society (BPS, 2009) and the Psychological Society of Ireland (PSI, 2010), and received ethical approval from the Maynooth University Research Ethics Committee.¹ The key ethical issues pertaining to the current study may be summarised as follows. 1. Each participant completed a consent form, which highlighted specific ethically-relevant features, provided details on the nature and aims of the research, and outlined freedom to withdraw at any point (see Appendix 2). 2. All data were anonymised and analysed at group level. 3. The lack of distress or harm associated with the IRAP has been demonstrated by empirical evidence (personal communication with D. Barnes-Holmes). 4. Participants were fully debriefed and provided with researcher contact details. It is important to emphasise that no participant expressed any signs of distress prior to, during, or after involvement in the study.

Analytic Strategy

Experiment 1 sought to investigate a specific variation in the presentation of the label and target stimuli in the IRAP by comparing the statement-based Natural Language IRAP with a Typical IRAP (in which the label and target words remained separate). A very simple set of stimuli (positive and negative reactions to fruit and insects) were used for this purpose. One explicit measure (the NART) was used to control for participants IQ. The NART data were analysed in terms of relationship to a normative IQ. The IRAP data were analysed according to practices consistently used

¹ This study utilized a modified version of an IRAP originally developed for and run by Emer Long which had obtained full ethical approval.

in IRAP research (see below for a detailed description of IRAP *D*-score preparation). That is, the *D*-scores on each of the four trial-types were compared across the two IRAPs and a 4x2 repeated measures ANOVA was used to determine main effects for IRAP type or trial-type. One sample t-tests were then used to determine the significance of any of the *D*-scores. Independent t-tests determined whether performances on trial-types in one IRAP related to performances on the other. On this occasion, we also assessed whether the two IRAPs were comparable in terms of the time taken to complete them. Finally, a correlation matrix explored the potential relationships among the *D*-scores and the NART.

Results

NART Data

The mean score on the NART was 17.48 ($SD = 3.46$). This converted to a mean IQ of 109.14 ($SD = 4.46$) on the WAIS Full Scale (see Appendix 3). This indicated that participants overall were categorised within the ‘normal’ IQ range.

IRAP Data

One participant failed to meet the mastery criteria on the practice blocks and so the data were removed. Another two participants failed one or more test block pairs and these data were also removed. In total, therefore, three participants did not meet the pass criteria for one of the IRAPs, leaving $N = 21$.

The mean *D*-scores for each trial-type for each IRAP are illustrated in Figure 4. Both IRAPs recorded a strong positive responding on the Fruit-Positive trial-type, with only a marginally stronger effect on the Typical IRAP (that is, participants chose Fruit-Positive-True more quickly than Fruit-Positive-False). The IRAPs were again

similar on Fruit-Negative, but both effects were now weaker, and the larger of the two was recorded on the Natural Language IRAP. A very similar pattern was recorded on Insects-Positive, but the effects were somewhat weaker still and the stronger effect was again recorded on the Natural Language IRAP. The Insects-Negative trial-type showed almost no effect for either IRAP.

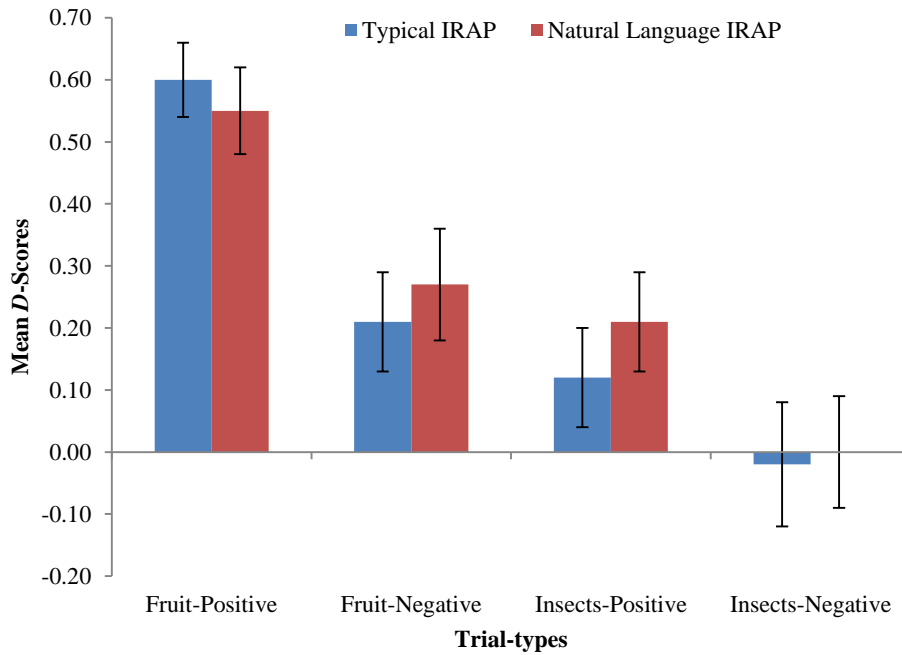


Figure 4. Mean *D*-scores, with standard error bars, for the four trial-types on the Typical and Natural Language IRAPs in Experiment 1. **Positive *D*-scores indicate positive responding and negative *D*-scores indicate negative responding.**

The *D*-scores on the two IRAPs were subjected to a 4×2 repeated measures ANOVA, with trial-type and IRAP type as within subject variables. As expected, there were no main effects and no interaction effect (all *p*'s >.05). One sample *t*-tests demonstrated that responding on Fruit-Positive was significant from zero on both IRAPs (Typical IRAP: $M = .60$, $SD = .29$, $p < .0001$; Natural Language IRAP: $M = .55$, $SD = .32$, $p < .0001$), as was responding on Fruit-Negative (Typical IRAP: $M = .21$, $SD = .37$, $p < .05$; Natural Language IRAP: $M = .27$, $SD = .42$, $p < .01$).

Responding on Insects-Positive was also significant on the Natural Language IRAP ($M = -.21$, $SD = .35$, $p = .01$; all other p 's $>.05$). Four independent t -tests were then conducted to determine if any of the trial-type D -scores differed significantly between each IRAP, but none did (all p 's $>.40$). Finally, we assessed whether the two IRAPs were comparable in terms of the time taken to complete them (Typical IRAP: $M = 9.84$ min, $SD = 3.21$; Natural Language IRAP: $M = 9.50$ min, $SD = 2.50$) and a dependent t -test confirmed no significant difference ($p = .57$).

Correlations between the IRAP and the Explicit Measures

A correlation matrix calculated the potential relationships among the trial-type D -scores on each IRAP and the NART. The only significant correlation reported with the NART was a moderate correlation on the Fruit-Positive trial-type on the Natural Language IRAP ($r = .47$, $p = .03$). To reduce the possible statistical error associated with running correlational matrices, Bonferroni tests were run for statistical integrity. This resulted in the p values above being corrected to $.003$. For the purpose of this thesis, the data before this correction will be mainly discussed. No other significant correlations were found (all r 's between $.03$ and $.2$, all p 's $> .29$).

Summary of Results

Patterns of responding were very similar across the Typical and Natural Language IRAPs. In both cases participants showed strong positive responding to Fruit-Positive and to Fruit-Negative. There was also strong positive responding to Insects-Positive on the Natural Language IRAP, with a weaker non-significant effect on the Natural Language IRAP. In both cases there was very weak responding on Insects-Negative. Responding on one IRAP did not differ with the other as shown by

independent t-tests, while responding did not correlate with the NART, which had indicated that the mean IQ was in the normal range. Overall, the outcomes on the IRAPs were similar and thus consistent with our provisional predictions in the current context.

Discussion

Few published IRAP studies have asked specific methodological questions about the procedure. From the data, one can conclude overall that there is some promise in presenting more complex stimuli (especially in clinical contexts), but this must be done carefully with regard to the various procedural parameters. In order to contribute to this small body of work, Experiment 1 compared a Typical IRAP with a Natural Language IRAP in order to determine the possible impact of combining the label and target stimuli. Our basic aim was to determine whether this would yield IRAP outcomes that were dissimilar to those typically observed and thus we employed a simple fruit-insects IRAP for this purpose. The data indicated very similar responding on both IRAPs that was consistent with what is typically seen in an IRAP, with perhaps a marginal superiority of effects on the Natural Language IRAP.

Overall, the findings from Experiment 1 supported existing evidence which suggests that combining the label and target stimuli produces sound IRAP effects, but does not have a direct impact on the outcomes or modify them in any noticeable manner. Given that our broader aim currently was to explore methodological parameters of the IRAP with a view to paving the way for the development of more complex and clinically-relevant IRAP studies, the findings from Experiment 1 go some way towards this. That said, there are numerous other procedural aspects of the

IRAP that remain untested in terms of their potential influence. Thus, in Experiment 2, we turned our attention to another such variable.

Chapter 3
Experiment 2

Manipulating Rule Order

Chapter 3

Introduction

The use of pre-specified, alternating rules for responding is one of the core tenets of the IRAP's structure. Opposing alternating rules lie at the centre of the IRAP's effects from which inferences about behaviour, based on participants' performances, are made. That is, they allow for calculation of the IRAP's *D*-scores that describe these effects (the difference in response times between responding according to one rule compared to responding according to the other). While this component of the procedure may be central to IRAP effects, there has been little research to date on this variable.

In a recent study, O'Shea et al. (in press) systematically investigated the rules for IRAP responding. They compared three different ways of presenting the rules to participants across a number of IRAPs. Their findings indicated that the strongest IRAP effects were recorded when the rules emphasised the positive dimension of the rule *before* the negative dimension (i.e. 'thin person-positive' followed by 'fat person-negative'). However, this effect was only found in two of the four IRAPs, namely those presenting nonsense syllables or stimuli relating to social systems, but the effect was not recorded in the IRAPs concerning bodyweight or nature. These outcomes suggest that more established verbal relations, such as those involved in body weight and nature, are possibly less sensitive to the ways in which the IRAP rules are provided.

In a similar vein, McKenna et al. (2007) explored the potential influence of response rules in the IRAP in what have been described as 'faking' manipulations. While various conditions attempted to decipher the impact of one type of faking rule

over another, or compared with no faking rule, the results did not differ from typical IRAP effects. This suggested little or no influence of the rule manipulations.

Taken together, the studies summarised above indicate that rules may or may not exert some influence over IRAP effects, and if they do, this may be observed only in certain contexts or in specific IRAPs. This suggests, therefore, that exploring rule manipulations in IRAP research would best be conducted in the context of an IRAP that has shown robust outcomes. One such IRAP is the Self-esteem IRAP which has yielded readily to various experimental manipulations and has demonstrated relatively strong replicability. Indeed, the IAT has also been used extensively to study self-esteem (e.g. De Raedt, Schacht, Franck, De Houwer, 2006).

The first study involving the Self-esteem IRAP was reported by Vahey et al. (2009). Across trials, participants related the labels “similar” or “opposite” with positive (e.g. “confident”) or negative (e.g. “worthless”) target words. Response options in this case were the participants’ own names and the participants name preceded by the word ‘Not’ (e.g. “Aileen” and “Not Aileen”). Hence, the four trial-types in this study, based on consistent and inconsistent blocks, were: Similar-Positive; Opposite-Negative; Similar-Negative; and Opposite-Positive. On consistent blocks, responding correctly involved relating: similar-positive-participant’s name, similar-negative-not participant’s name, opposite-negative-participant’s name and opposite-positive-not participant’s name. In contrast, correct responding on inconsistent blocks involved relating: similar-positive-not participant’s name, similar-negative-participant’s name, opposite-negative-not participant’s name and opposite-positive-participant’s name. In simple terms, consistent responding was based on the view that participants would more readily evaluate positive words as similar to themselves (i.e. self-positive) and negative words as opposite to themselves (i.e.

others-negative). And the results confirmed that this was the case, thus showing a self-positive bias. In addition, the outcomes on this IRAP correlated with standard explicit measures of self-esteem, such as the Rosenberg Self-Esteem Scale (RSES). Self-esteem IRAPs have also been used to identify ideal vs. actual self-esteem amongst individuals referred to as dysphoric or non-dysphoric (Remue et al., 2013) and among typically- vs. atypically-developing children (Scanlon, McEnteggart, Barnes-Holmes, & Barnes-Holmes, 2014). Again, the effects have been impressive and self-based relations did appear to differentiate the two groups in both studies.

In the design of Experiment 2², we considered the existing but limited and mixed evidence on the impact of rules on IRAP effects, and believed that this warranted further attention, particularly (as noted previously) if the IRAP is to be used in more challenging clinically-relevant domains. The robust effects recorded across several Self-esteem IRAPs suggested that this might be a useful context in which to conduct rule-based manipulations. Of course, these manipulations may have been conducted in the context of a simple positive-negative IRAP, but then one may find limited generalisation to more complex IRAPs. We believed that the Self-esteem IRAP was somewhere along this continuum from basic to complex. Hence, Experiment 2 investigated a rule order manipulation in the context of a Self-esteem IRAP. Specifically, the study manipulated the order in which the self vs. others rules were presented. In the Self-Rule First Condition, participants were presented with the rules as follows: the self-positive rule followed by the others-negative rule, while in the Others-Rule First Condition participants were presented with the others-positive rule followed by the self-negative rule. This manipulation sought to determine whether the simple order in which the rules are presented influenced IRAP outcomes.

² It should be noted that the data in this experiment were collected in part by Ms. Sara Kenehan.

As was the case with Experiment 1, it was difficult to make predictions about possible differences that may emerge between the Self-Rule First Condition and the Others-Rule First Condition because this specific manipulation had not been done before and mixed findings have been recorded across previous manipulations of IRAP rules.

Method

Participants

Experiment 2 involved 41 undergraduate students from Maynooth University (21 females and 20 males) aged between 18 and 26 years ($M = 19.61$, $SD = 2.02$), all selected by random convenience sampling. The study comprised of two conditions, referred to as *Self-Rule First Condition* and *Other-Rule First Condition*, between which all participants were evenly and semi-randomly assigned.

Setting

All aspects of Experiment 2 were identical to Experiment 1.

Apparatus and Materials

All aspects of the experiment (including the questionnaires) were completed on a DELL desktop computer. The study involved both explicit (self-report) measures and an implicit measure (the IRAP). The IRAP was delivered via PsychoPy (1.81.03) which controlled all aspects of stimulus presentation and recording of participant responses. The IRAP employed here is referred to as the Self-esteem IRAP and was used previously by Vahey et al. (2009) and by Scanlon et al. (2014). The two explicit

measures comprised of the Depression, Anxiety Stress Scales (DASS-21: Lovibond & Lovibond, 1995) and the Rosenberg Self-Esteem Scale (RSES: Rosenberg, 1965).

The Depression, Anxiety Stress Scales (DASS). The short version of the DASS is a 21-item questionnaire that assesses depression, anxiety and stress (Lovibond & Lovibond, 1995). The measure contains three sub-scales (each with seven items) that indicate the presence and severity of emotional states in the past week. For example, the Depression sub-scale contains the statement “*I couldn't seem to experience any positive feeling at all*”. Responding on all sub-scales involves a 4-point Likert scale ranging from 0 (Did not apply to me at all) to 3 (Applied to me very much or most of the time). Scores from each sub-scale are summed and multiplied by two to generate a total score each for depression, anxiety and stress. Higher scores indicate higher levels of symptomatology and each sub-scale is sub-divided in terms of Normal (0-9 for depression; 0-7 for anxiety; 0-14 for stress), Mild (10-13 for depression; 8-9 for anxiety; 15-18 for stress), Moderate (14-20 for depression; 10-14 for anxiety; 19-25 for stress), Severe (21-27 for depression; 15-19 for anxiety; 26-33 for stress) and Extremely Severe (28+ for depression; 20+ for anxiety; 34+ for stress). The DASS-21 has demonstrated high internal consistency (Cronbach's α = .82, .90 and .93 for depression, anxiety and stress, respectively) and adequate construct validity in a large sample of normative participants (Henry & Crawford, 2005). See Appendix 4.

The Rosenberg Self-Esteem Scale (RSES). The RSES is a 10-item measure of self-esteem. All items (e.g. “*I take a positive attitude toward myself*”) are rated on a 4-point scale from 1 (Strongly agree) to 4 (Strongly disagree). The RSES yields an overall score with a maximum of 30 and a minimum of 0, with scores ranging from 15-25 indicating normal self-esteem, <15 indicating low self-esteem and >25

indicating high self-esteem. The RSES has demonstrated satisfactory test-retest reliability (.88), strong concurrent and predictive validity (Robins, Hendin, & Trzesniewski, 2001) and adequate construct validity in a large sample of students (Hagborg, 1993). See Appendix 5.

The Self-esteem IRAP. The Self-esteem IRAP presented only two label words, one from the category Self (i.e. I AM) and one from the category Others (i.e. OTHERS ARE). This IRAP also presented 12 target stimuli, six of which were positive words (e.g. KIND) and six of which negative words (e.g. MANIPULATIVE). Again, the response options TRUE and FALSE were present on each trial. The list of label stimuli, target stimuli and the response options comprising the Self-esteem IRAP are provided in Table 2.

Table 2

The stimuli used in the Self-esteem IRAP in Experiment 2, including labels, targets and response options

Self Label	Others Label
I AM	OTHERS ARE
Positive Targets	Negative Targets
LOYAL	MANIPULATIVE
TRUSTWORTHY	DISHONEST
KIND	CRUEL
MORAL	HORRIBLE
GENEROUS	SELFISH
FRIENDLY	HEARTLESS
Response Option 1	Response Option 2
TRUE	FALSE

The Self-esteem IRAP comprised of four possible trial-types: Self-Positive; Self-Negative; Others-Positive; and Others-Negative (see Figure 4). Again, participants responded to these by selecting TRUE or FALSE, assigned to the "d" and "k" keys, respectively.




Figure 5. Examples of the four trial-types in the Self-esteem IRAP presented in Experiment 2. On each trial, a label stimulus (Self or Others), a target stimulus (Positive or Negative) and two response options (TRUE and FALSE) appeared on-screen simultaneously. This generated four trial-types: Self-Positive; Self-Negative; Others-Positive; and Others-Negative. Neither the phrases 'Consistent' or 'Inconsistent' nor the arrows were shown on screen.

Responses on the IRAP trials were defined as either consistent or inconsistent in terms of whether they were 'pro-self' or 'pro-others'. Consider, for example, a pro-self block of trials in which participants are asked to respond positively to self and negatively to others. When the trial presented the Self label (I AM) with a positive

target (e.g. FRIENDLY), TRUE was the correct response because it was consistent (i.e. pro-self) and FALSE was inconsistent (not pro-self, see Figure 5). Now consider a pro-others block in which participants are asked to respond positively to others and negatively to self. When the trial presented the Self label with a negative target, TRUE was consistent and FALSE was inconsistent.

Procedure

The experimental sequence for Experiment 2 comprised of two stages that were identical for all participants, with participants being assigned to each condition semi-randomly. Stage 1 involved the Self-Esteem IRAP and Stage 2 involved the two explicit measures (the DASS and the RSES, always presented in that order). Unless otherwise stated, all procedural aspects of the Self-esteem IRAP were identical to the Typical IRAP as presented in Experiment 1.

Stage 1: Self-Esteem IRAP (Self-Rule First Condition and Others-Rule First Condition). In Stage 1, all participants completed the same Self-esteem IRAP, but the rule sequence with which they were provided distinguished the two conditions. That is, half of the participants were exposed to the Self-Rule First Condition in which the IRAP rules always specified self first and others next, while the other half were exposed to the Others-Rule First Condition, in which the IRAP rules always specified others first and self next.

Although all participants completed the same IRAP, and IRAP blocks always alternate between consistent and inconsistent responding (hence, participants must switch responding after each block), the two participant groups were differentiated in terms of the actual rule order specified on each block. That is, participants in the Self-Rule First Condition were always presented with a self-rule first. Specifically, half of

the blocks presented the rule “I am good and others are bad”, while the other blocks presented the rule “I am bad and others are good”. In short, for these participants every block of trial was preceded by a rule in which responding to ‘self’ was always specified first and responding to ‘others’ was always specified second. The opposite was the case for participants assigned to the Others-Rule First Condition. This group was always presented with an Others-rule first. Specifically, half of the blocks presented the rule “Others are good and I am bad”, while the other blocks presented the rule “Others are bad and I am good”. In short, for these participants every block of trial was preceded by a rule in which responding to others was always specified first and responding to the self was always specified second.

Stage 2: Explicit Measures. In Stage 2, all participants completed the DASS first, followed by the RSES. They were thereafter fully debriefed as to the nature of the study and measures used and thanked for their participation.

IRAP Data Preparation

All aspects of the data preparation for the Self-Esteem IRAP were identical to the previous study (see Experiment 1 for an in depth outline of the *D*-score calculation procedure). The *D*-scores for trial-types 3 and 4 in each IRAP were again reversed for ease of interpretation. Positive scores indicate “*positive*” or “*not-negative*” responding, whereas negative scores indicate “*negative*” or “*not-positive*” responding.

Ethical Considerations

All aspects of ethical consideration that applied to Experiment 2 were identical to Experiment 1 apart from one minor aspect. In this case, ethical approval was

attained at a departmental level as it was developed, in part, for the purpose of an undergraduate final year project.

Analytic Strategy

Experiment 2 sought to systematically compare responding on two rule order conditions, in which either the self-rule was presented first, followed by the others-rule, or the others-rule was presented first, followed by the self-rule. This manipulation was presented in the context of the Self-esteem IRAP. Two explicit measures (the DASS and RSES) were used to control for participants' levels of distress and self-esteem. The DASS and RSES data were analysed in terms of normative levels of distress and self-esteem. The *D*-scores on each of the four trial-types were compared across the two IRAP conditions and a 4x2 repeated measures ANOVA determined possible main effects for condition or trial-type. One sample t-tests determined the significance of any of the *D*-scores. Independent t-tests determined whether performance on trial-types in one condition related on the other. Finally, a correlation matrix explored the potential relationships among the *D*-scores and the explicit measures.

Results

Explicit Measures

The mean scores and standard errors were calculated for both explicit measures and are presented in Table 3.

Table 3

Mean scores with standard error values for the DASS sub-scales and the RSES used in Experiment 2

Explicit Measures	Mean (SD)
DASS Depression	12.34 (.64)
DASS Anxiety	10.98 (.55)
DASS Stress	13.27 (.61)
RSES	21.02 (.88)

The DASS Depression group mean was 12.34 and thus within the mild range (10-13). This mean reflected the scores of 13 participants who scored as normal, 14 who showed mild depression, 11 who were moderate and 3 who were severe. The Anxiety mean was 10.98 and moderate. This mean reflected the scores of 3 participants who scored as normal, 15 as mild anxious, 17 as moderate, 5 as severe and 1 as extremely severe. Finally, the Stress mean was 13.27 and normal. This mean reflected 31 participants who scored as normal, 3 as mild and 7 as moderate. The RSES group mean was 21.02 and thus categorised as normal in self-esteem. This mean reflected the scores of 28 participants who scored as normal, 6 who were low and 7 who were high. Four independent t-tests indicated that the participants in each condition did not differ significantly from each other on the DASS overall level of distress or on the RSES ($p > .05$ for all).

IRAP Data

All participants met the pass criteria for the Self-esteem IRAP, leaving $N = 41$.

The mean D -scores for each trial-type for each condition are illustrated in Figure 6. Both conditions showed strong positivity on the Self-Positive trial-type, although the effect was larger in the Self-First Rule Condition. On Self-Negative, no effect was recorded in the Others-First Rule Condition, while a marginal negative

response was observed in the Self-Rule First Condition. The opposite was the case for responding on the Others-Positive. There was no effect recorded for the Self-First Rule Condition, while a modest positive response was observed in the Others-Rule First Condition. On Others-Negative, the two conditions diverged considerably in that a marginal positive response was recorded in the Others-Rule First Condition, while a modest negative response was recorded in the Self-Rule First Condition.




Figure 6. Mean *D*-scores, with standard error bars, for the four trial-types in the Self-esteem IRAP in Experiment 2. **Positive *D*-scores indicate positive responding and negative *D*-scores indicate negative responding.**

The *D*-scores on the two IRAP conditions were subjected to a mixed 4×2 ANOVA, with trial-type as the within subjects variable and IRAP as the between subjects variable, but there were no main or interaction effects (all *p*'s >.05). Eight one-sample *t*-tests determined whether any of the *D*-scores differed significantly from zero. As expected, responding to Self-Positive did differ significantly in both

conditions (Self-Rule First: $M = .328$, $t(19) = 4.45$, $p = .0003$ and Other-Rule First: $M = .441$, $t(19) = 4.074$, $p = .006$). Four independent t-tests were then conducted to determine if any of the trial-type *D*-scores differed significantly between each IRAP, but none did (all p 's $> .05$).

Correlations between the IRAP and the Explicit Measures

A correlation matrix calculated potential relationships among the trial-type *D*-scores and the explicit measures. There were only two relationships that showed a moderate positive correlation between variables: the Depression sub-scale of the DASS ($p < .01$) and the RSES ($p < .01$) both correlated with responding on Others-Positive ($r = .5$). That is, greater implicit positivity towards others correlated with higher levels of self-esteem and depression. These correlations were only found in the Others-Rule First Condition. To reduce the possible statistical error associated with running correlational matrices, Bonferroni tests were run for statistical integrity. This resulted in the p values above being corrected to $.001$. For the purpose of this thesis, the data before this correction will be mainly discussed.

Summary of Results

Patterns of responding differed somewhat between the two conditions. In both, participants showed strong significant positive responding to Self-Positive, with the larger effect when the self-rule was presented first. There was no clear response on Self-Negative when the others-rule was presented first, but a modest negative response when the self-rule came first. The two conditions diverged even more on Others-Positive, with almost no response recorded when the self rule came first, but a modest positive response when the other rule came first. Divergence, albeit less, was also

recorded on Others-Negative with a marginal negative response when the self-rule came first, compared with a very marginal positive response when the others-rule came first. It was surprising, therefore, that the two conditions did not differ significantly on any trial-type. Responding on the Others-Rule First Condition correlated with depression and self-esteem, such that greater implicit positivity towards others correlated with higher self esteem and depression.

Discussion

The use of pre-specified, alternating rules for responding is one of the core tenets of the IRAP's structure. While this component of the procedure could be argued to be central to IRAP effects, there has been little research to date on this variable. The limited evidence base that is available suggests that rules may exert some influence over IRAP effects, but perhaps only in certain contexts or specific IRAPs. The robust effects recorded across several Self-esteem IRAPs suggested that this would be a useful context in which to conduct rule-based manipulations. Hence, Experiment 2 investigated a rule order manipulation in the context of a Self-esteem IRAP. Specifically, we manipulated the order in which the self vs. others rules were presented. In the Self-Rule First Condition, participants were presented with the rules as follows: the self-positive rule followed by the others-negative rule, while in the Others-Rule First Condition participants were presented with the others-positive rule followed by the self-negative rule. This manipulation sought to determine whether the simple order in which the rules are presented influenced IRAP outcomes.

The results indicated that while responding did not differ significantly between conditions, patterns of responding did differ somewhat between the two conditions. It should be noted that it is common practice in IRAP research for analyses and

speculation regarding non-significant findings to be noted in a manner that reflects the single subject focus of traditional behavioural research. The writer is conscious of this and it must be stressed that the discussion points generated as a result are speculative in nature as the conditions themselves did not differ significantly. Overall, one could perhaps conclude from the above data that the stronger IRAP effects were associated with the Other-Rule First Condition, which incidentally also correlated with depression and self-esteem.

Overall, the findings from Experiment 2 support some existing evidence which suggests that rules exert some influence on IRAP effects, and thus warrant empirical attention, especially in specific contexts. Again, our broader aim was to explore various methodological parameters of the IRAP with a view to paving the way for the development of more complex and clinically-relevant IRAP studies, and thus the findings from Experiment 2 again go some way towards this. Again, however, there remained other procedural aspects of the IRAP that were untested in terms of their potential influence. Thus, in Experiments 3a and 3b, we turned our attention to two other such variables.

Chapter 4

Experiment 3

*Manipulations of Block Length
and Block Number*

Chapter 4

Introduction

While IRAP research has moved at a considerable pace towards an ever-increasing range of domains, methodological questions, such as those investigated in Experiments 1 and 2, remain. While the explorations in the two previous studies of the label-target combinations in the Natural Language IRAP (Experiment 1) and rule ordering (Experiment 2) are important in terms of the presentation of the IRAP, it could be argued that deeper procedural questions need to be asked, especially if the IRAP is to meet the complex challenge of studying clinically-relevant relations. One such area of investigation that has concerned IAT researchers has involved the *length* of the test blocks. For example, Klauer and Mierke (2005) argued that the length of IAT blocks exerts some influence upon effects (see also Back et al., 2005). In a similar vein, Nosek et al. (2007) investigated block *order* and found that this variable exerted some influence over responding (see also De Houwer, 2001; Rothermund et al., 2009).

The rule order manipulation with the Self-esteem IRAP conducted in Experiment 2 added to only a handful of IRAP studies that have explored the possible impact of rules in IRAP performances. And indeed, the results from Experiment 2 indicated that, at least in that context, rules exerted only a marginal influence. In order to investigate the possible impact of rules further, we looked at the few studies that have investigated this issue with the IAT, and thus decided that similar questions should be examined in the context of the IRAP. Hence, Experiments 3a and 3b investigated the influence of task switching and block length, respectively. Again, rule manipulations in the IRAP would best be conducted in the context of an IRAP that has shown robust outcomes. But rather than using the Self-esteem IRAP again, we looked

toward another IRAP that has shown robust effects across several studies, namely the Spider Fear IRAP.

The first study involving the Spider Fear IRAP was reported by Nicholson and Barnes-Holmes (2012). Across trials, participants related images of spiders or scenes of nature with approach- (e.g. “I may approach”) or fear- (e.g. “Scares me”) related phrases. Hence, the four trial-types in this study, based on consistent and inconsistent blocks were: Fear-Spiders; Approach-Nature; Fear-Nature; and Approach-Spiders. On consistent blocks, responding correctly involved relating: fear-spiders-true, approach-spiders-false, fear-nature-false and approach-nature-true. In contrast, correct responding on inconsistent blocks involved relating: fear-spiders-false, approach-spiders-true, fear-nature-true and approach-nature-false. In simple terms, consistent responding was based on the view that participants would more readily evaluate images of nature as safe and images of spiders as fearful. And the results confirmed that this was the case, thus showing a fear-spiders bias. In addition, the outcomes on this IRAP correlated with the standard explicit measure of spider fear, namely the Fear of Spiders Questionnaire (FSQ).

In Experiment 3a³, we employed the Spider Fear IRAP and manipulated the length of the test blocks presented between rules. To do so, we compared a *Fast Switching IRAP* and a *Typical Switching IRAP*. Specifically, the Fast Switching IRAP halved the length of the test blocks between rules, relative to the Typical Switching IRAP (and IRAPs generally), such that each block now contained only 16 trials, and not 32. Again, this manipulation sought to determine whether the typical rate at which blocks switch exerts its own influence on outcomes normally observed with the IRAP. Once again, it was difficult to make predictions about possible differences that may

³ It should be noted that the data in experiments 3a and 3b were collected in part by Mr. Alan Kane.

emerge between the Fast Switching IRAP and the Typical Switching IRAP because this specific manipulation had not been done before.

Experiment 3a

Method

Participants

Experiment 3a involved 79 undergraduate students from Maynooth University (47 females and 32 males) aged between 18 and 42 years ($M=21.91$, $SD=4.54$), all selected by random convenience sampling. The study involved of two IRAPs, referred to as the *Fast Switching IRAP* and *Typical Switching IRAP*, between which all participants were evenly and semi-randomly assigned.

Setting

All aspects of the setting in Experiment 3a were identical to the previous studies, with the exception that a tarantula molt was now used as part of the Behavioural Approach Task (BAT). During the BAT, the researcher remained inside the experimental cubicle and interacted directly with each participant.

Apparatus and Materials

All aspects of Experiment 3a (including the questionnaire) were completed on a DELL desktop computer. The study involved an explicit (self-report) measure, the IRAP and a Behavioural Approach Task (BAT). The IRAP was again delivered via PsychoPy (1.81.03) which controlled all aspects of stimulus presentation and recording of participant responses. The IRAP employed here is referred to as the



Spider Fear IRAP and was used previously by Nicholson and Barnes-Holmes (2012). *It is important to note that this IRAP differed specifically from the IRAPs presented in Experiments 1 and 2 in that both pictures and words were presented here (only words were presented previously).* The explicit measure comprised of the Fear of Spiders Questionnaire (FSQ; Szymanski & O'Donohue, 1995). The BAT involved approaching a tarantula molt, a variation of which involving a live tarantula was used previously by Nicholson and Barnes-Holmes in conjunction with the Spider Fear IRAP.

Fear of Spiders Questionnaire (FSQ). The FSQ is an 18-item scale that measures fear of spiders. All items (e.g. *"I now would do anything to try to avoid a spider"*) are rated on a 7-point scale from 0 (Strongly disagree) to 6 (Strongly agree). The FSQ yields an overall score with a maximum of 108 indicating high spider fear and a minimum of 0 indicating low spider fear. The FSQ has been shown to have high internal consistency (Cronbach's α = .92), test-retest reliability (.97) and split-half reliability (.89; Szymanski & O'Donohue, 1995). While this was presented to participants on-screen, a paper based version of the FSQ is provided in Appendix 6.

The Spider Fear IRAP. The IRAP presented a total of eight pictures, four of which presented spiders and four of which presented scenes of nature. The IRAP also presented three positive target label statements (e.g. I COULD APPROACH) and three negative target label statements (e.g. FRIGHTENS ME), as well as the response options TRUE and FALSE on each trial. The label stimuli, target stimuli and response options comprising the Spider Fear IRAP are provided in Table 4.

Table 4

The stimuli used in the Spider Fear IRAP in Experiment 3a, including labels, targets and response options

Fear Label	Approach Label
SCARES ME	I COULD APPROACH
DISGUSTS ME	I MAY APPROACH
CREEPS ME OUT	I CAN APPROACH
Spider Targets	Nature Targets
	
Response Option 1	Response Option 2
TRUE	FALSE

The Spider Fear IRAP comprised of four trial-types: Fear-Spiders; Fear-Nature; Approach-Spiders; Approach-Nature (see Figure 7). Participants responded to these by selecting TRUE or FALSE, assigned to the “d” and “k” keys, respectively.

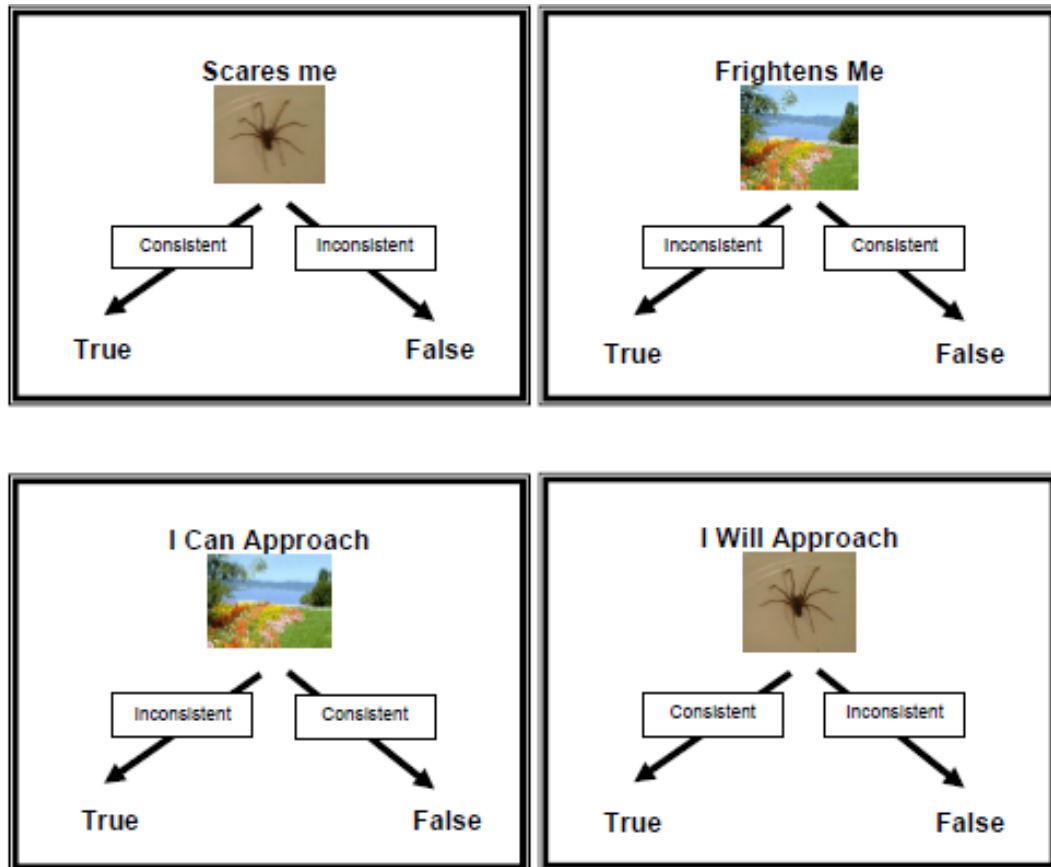


Figure 7. Examples of the four trial-types in the Spider Fear IRAP presented in Experiment 3a. On each trial, a label stimulus (e.g. SCARES ME or FRIGHTENS ME), a target stimulus (spiders or nature images) and two response options (TRUE and FALSE) appeared on-screen simultaneously. This generated four trial-types: Fear-Spiders; Approach-Spiders; Fear-Nature; Approach-Nature. Neither the phrases ‘Consistent’ or ‘Inconsistent’ nor the arrows were shown on screen (figure taken from Nicholson & Barnes-Holmes, 2012).

Responses on the IRAP trials were defined as either consistent or inconsistent in terms of whether they were ‘pro-spiders’ or ‘pro-nature’. Consider, for example, a pro-spider block of trials in which participants are asked to respond positively to spiders and negatively to nature. When the trial presented an approach label (e.g. I CAN APPROACH) with a target picture of a spider, TRUE was the correct response because it was consistent (i.e. pro-spider) and FALSE was inconsistent (not pro-spider, see Figure 6). Now consider a pro-nature block in which participants are asked to respond positively to scenes of nature and negatively to pictures of spiders. When the

trial presented the fear label (e.g. SCARES ME) with a target picture of a spider, TRUE was consistent and FALSE was inconsistent.

The Behavioural Approach Task (BAT). The BAT comprised of a tarantula molt that was approximately 4 inches in width and placed in a plastic container.

Procedure

The experimental sequence for Experiment 3a comprised of three stages that were identical for all participants. Stage 1 involved the Spider-Fear IRAP, Stage 2 involved the explicit measure (the FSQ), and Stage 3 presented the BAT. Unless otherwise stated, all procedural aspects of the Spider Fear IRAP were identical to the Typical IRAP as presented in Experiment 1.

Stage 1: Spider Fear IRAPs (Typical Switching IRAP and Fast Switching IRAP). In Stage 1, participants completed one of the two IRAPs, either the Typical Switching IRAP or the Fast Switching IRAP. *The fast switching IRAP halved the length of the test blocks between rules, relative to the Typical Switching IRAP (and IRAPs generally). That is, each block now contained only 16 trials, and not 32, such that each rule controlled a 16-trial test block and not the standard 32-trial test block.* See Figure 8. However, in order to ensure that this IRAP maintained the total minimum number of test trials that comprise the Typical IRAP, the current version now presented 24 test blocks (12 per rule). This matched the total number of test trials (192) that are presented in the Typical IRAP. Half of the participants were exposed to the Typical Switching IRAP, while the other half were exposed to the Fast Switching IRAP.




Figure 8. A brief illustration of the methodological differences between the Typical Switching IRAP and the Fast Switching IRAP in Experiment 3a.

Stage 2: FSQ. In Stage 2, all participants completed the FSQ.

Stage 3: BAT. In Stage 3, all participants completed the BAT. They were instructed initially that this aspect of the study was designed to assess their willingness to approach a spider. Participants were not made explicitly aware that the spider was a molt and not a live spider. Participants were instructed to complete as many steps of the BAT as were comfortable to them. The BAT comprised of six possible steps as follows: Step 1 (completing no steps); Step 2 (opening the door of the room and looking inside the room); Step 3 (entering the room); Step 4 (touching the spider box); Step 5 (opening the box); and Step 6 (touching the spider for 10 seconds). When participants indicated the final BAT step at which they felt comfortable, participation in the experiment was complete (see Appendix 7). They were thereafter fully debriefed and thanked for their participation (see Appendix 8).

IRAP Data Preparation

While most aspects of the data preparation for the Spider Fear IRAP were identical to the two previous studies (see Experiment 1 for in depth outline of the general *D*-score calculation procedure), a few minor modifications were necessary for

the analyses of data from the Fast Switching IRAP. Specifically, 24 (rather than 12) standard deviations were calculated for each pair of test blocks (four for each of the trial-types; this is twice that of the Typical Switching IRAP data as it consisted of twice the number of test blocks but half the length). Also, all *D*-scores were averaged across *pairs of pairs* of test blocks, so that three *D*-scores (rather than six), each consisting of 32 trials, were produced.

Ethical Considerations

All aspects of ethical consideration that applied to Experiment 3a again were almost identical to Experiment 1. Ethical approval for this project was confirmed by personal communication with Maynooth University's Social Research Ethics Subcommittee (SRESC) on the 20 May 2015. This can be confirmed by request. However, extra caution was taken in Experiment 3a because of the use of the Spider Fear IRAP and the tarantula molt. Specifically, participants were exposed to potentially stress-related stimuli for a brief period, most likely fear generated by the IRAP's spider words and pictures, as well as the spider molt used in the BAT. Indeed, the behavioural approach task with the spider molt is designed to elicit fear and so contained the possibility to cause minimal psychological discomfort to participants. However, each participant was made fully aware that they were not obliged to do any of the tasks if they did not wish to, and had full control over the extent of their participation in these tasks. Specifically, participants were never required to approach the molt and could choose not to do so at any point in the approach task. Evidence from the literature which has utilised similar tasks did not find any participants to be psychologically stressed by participation (see Teachman & Saporito, 2009; Nicholson & Barnes-Holmes, 2012). As a result, the likelihood of risk to any participant was

extremely low. Participants were reminded that they were free to withdraw from the study at any point and could decide at what point they wish to stop the experiment. In the potential case that any participant felt distress following the experiment, information of the counselling services provided by the University was on hand. It is important to emphasise that no participant expressed any signs of distress prior to, during, or after involvement in the study.

Analytic Strategy

Experiment 3a sought to systematically compare responding on two IRAPs, in which the point at which participants were required to switch responding was manipulated via presentation of the Typical Switching IRAP versus the Fast Switching IRAP (i.e. switching after every 32 trials vs. switching after every 16 trials, respectively). The Spider Fear IRAP was presented in both cases, along with the FSQ as an explicit measure and the BAT as a behavioural measure. The FSQ data were analysed in terms of relationship to a normative level of spider fear. The *D*-scores on each trial-type were compared across the two IRAPs and a mixed 4x2 ANOVA determined possible main effects for IRAP type or trial-type. One sample t-tests determined the significance of any of the *D*-scores. Independent t-tests determined whether performance on trial-types in one IRAP related to performances on the other. Finally, a correlation matrix explored the potential relationships among the *D*-scores, the FSQ and the BAT.

Results

FSQ Data

The mean on the FSQ was 52.37 ($SD = 29.96$). This indicated relatively normal levels of fear, given that the scale ranged from 0 to 108.

IRAP Data

Three participants did not meet the pass criteria for the Typical Switching IRAP, although all passed the Fast Switching IRAP, hence the total number of participants who completed both was $N = 77$.

The mean D -scores for each trial-type for each IRAP are illustrated in Figure 9. Both IRAPs showed fear responding on the Fear-Spiders trial-type, with a small effect on the Typical Switching IRAP and a strong effect on the Fast Switching IRAP. On the Approach-Spiders trial-type, both IRAPs showed modest approach responding, although again, the effect was larger in the Fast Switching IRAP. The opposite was the case for responding on Fear-Nature. There was limited fear responding on the Fast Switching IRAP and a more modest response on the Typical Switching IRAP. On Approach-Nature, both IRAPs showed moderately strong approach responding, with the larger effect recorded on the Typical Switching IRAP.

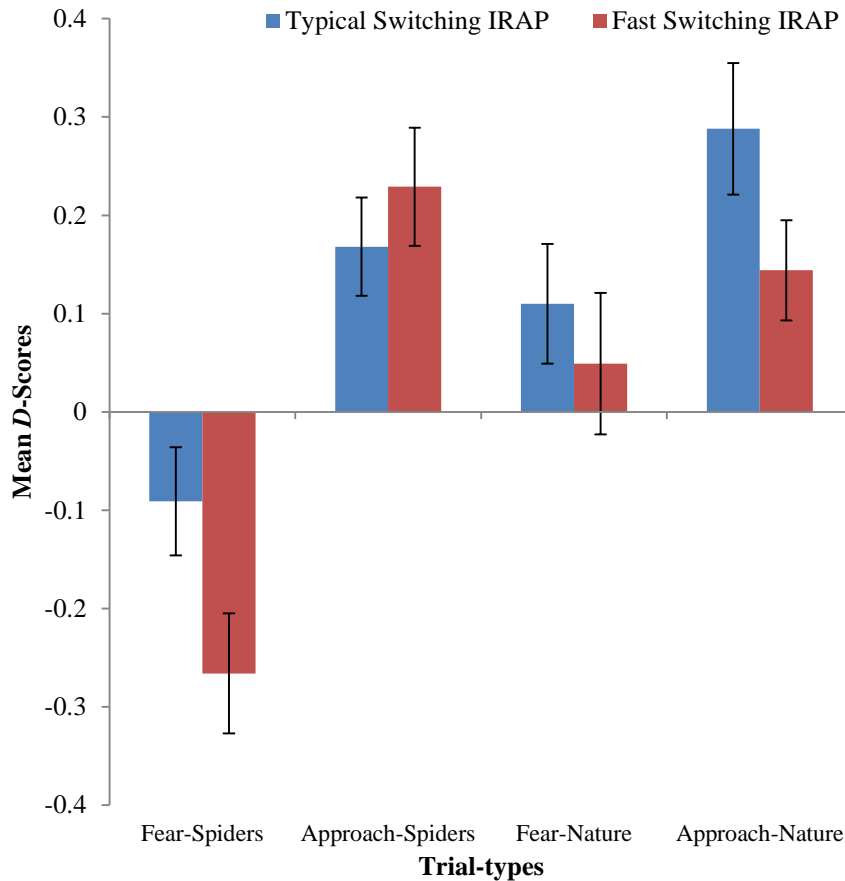


Figure 9. Mean *D*-scores, with standard error bars, for the four trial-types in the Spider fear IRAPs in Experiment 3a. **Positive *D*-scores indicate approach responding and negative *D*-scores indicate fear responding.**

The *D*-scores on the two IRAPs were subjected to a mixed 4x2 ANOVA, with trial-type as the within subjects variable and IRAP type as the between subjects variable. This found a main effect for trial-type ($f(3, 74)=13.56, p<.001$) and an interaction effect ($f(3, 222)=3.08, p=.03$). Two one way ANOVAs (one per IRAP) indicated a significant main effect for trial-type in both IRAPs (the Fast switching IRAP ($f(3, 38)=5.963, p < .0001$) and the Typical Switching IRAP ($f(3, 36)=11.854, p < .0001$)).

Eight one-sample *t*-tests determined whether the trial-types in either IRAP differed significantly from zero. Approach-Spiders (Typical Switching: $M = .168, t(36) = 3.34, p = .002$; Fast Switching: $M = .229, t(38) = 3.796, p = .0005$) and Approach Nature (Typical Switching: $M = .288, t(36) = 4.268, p = .0001$; Fast

Switching: $M = .144$, $t(38) = 2.848$, $p = .0071$) were significant on both IRAPs. Fear-Spiders was also significant on the Fast Switching IRAP ($M = .266$, $t(38) = 4.365$, $p < .0001$) but, not for the Typical Switching IRAP ($p = .11$, all other p 's). Four independent t -tests were then conducted to determine whether any of the D -scores differed significantly between the two IRAPs. Responding on Fear-Spiders was significant ($M = .175$, $t(74) = 2.14$, $p = .04$) and Approach-Nature approached significance ($M = -.14$, $t(74) = -1.712$, $p = .09$; all other p 's $> .44$).

Correlations among IRAPs and Other Measures. A correlation matrix calculated the potential relationships among the D -scores on both IRAPs, the FSQ and the BAT. The FSQ showed a positive correlation with responding on Fear-Nature in the Fast Switching IRAP ($r = .414$, $p = .008$) and this IRAP effect also showed a negative correlation that approached significance with the BAT ($r = -.293$, $p = .07$). That is, greater implicit fear of nature correlated with higher levels of spider fear and less approach behaviour on the BAT, but only on the Fast Switching IRAP. To reduce the possible statistical error associated with running correlational matrices, Bonferroni tests were run for statistical integrity. This resulted in p values being corrected to .0005 and .005 for the Fear-Nature trial type correlations with the FSQ and the BAT mentioned above. For the purpose of this thesis, the data before this correction will be mainly discussed.

Summary of Results

Patterns of responding were somewhat similar across the two IRAPs. In both cases, participants showed negative responding to Fear-Spiders, with a much larger effect in the Fast Switching IRAP. Both IRAPs showed positive responding on all three remaining trial-types. Specifically, there was strong significant positive

responding to Approach-Spiders, again with a larger effect in the Fast Switching IRAP. Responding was more modest on Fear-Nature, with a larger effect now on the Typical Switching IRAP. Stronger significant responding was recorded on Approach-Nature, with a much larger effect again recorded in the Typical Switching IRAP. Given these results, it was not surprising that the analyses yielded a significant main effect for trial-type in both IRAPs. Furthermore, responding on the two IRAPs differed significantly on the Fear-Spiders trial-type, while the difference between the two groups approached significance on the Approach-Nature trial-type. Responding on Fear-Nature in the Fast Switching IRAP was the only trial-type that correlated with any other measure. A significant positive correlation was recorded for the FSQ while a negative correlation for the BAT approached significance. That is, greater implicit fear of nature correlated with higher spider fear and less approach behaviour on the BAT.

Discussion

Experiment 3a investigated the influence of task switching on outcomes on the Spider Fear IRAP. In short, we manipulated the length of the test blocks presented between rules by comparing the Fast Switching IRAP with the Typical Switching IRAP. In other words, this manipulation sought to determine whether the typical rate at which blocks switch exerts its own influence on outcomes normally observed with the IRAP.

Patterns of responding were somewhat similar across the two IRAPs. Stronger responding was recorded on the Fast Switching IRAP on Fear-Spiders and Approach-Spiders, while the Typical Switching IRAP was stronger on Fear-Nature and Approach-Nature. The difference on Fear-Spiders was significant. Overall, therefore,

while there were differences between the IRAPs and some of these differences were even significant one could not argue that either IRAP yielded a better outcome, at least in the context of Spider Fear. Hence, in Experiment 3b, we turned our attention to a related potential source of influence and looked specifically at the possible role of block number.

Experiment 3b

Introduction

Given the data from Experiment 3a, we employed the Spider Fear IRAP again in Experiment 3b, but now manipulated the number of test blocks presented. To do so, we compared a *Double Length IRAP* and a *Typical Length IRAP*. That is, instead of the typical 6 test blocks (i.e. 3 per rule), the Double Length IRAP presented 12 test blocks (6 per rule), although each test block contained the typical 32 trials. Overall, therefore, this modified IRAP presented a total of 384 trials -- that is double the number of trials presented in the Typical Length IRAP (i.e. 192). Again, this manipulation sought to determine whether the typical number of blocks exerts its own influence on outcomes normally observed with the IRAP, and it was difficult to predict whether any differences in IRAP performances across the two would be observed.

Method

Participants

Experiment 3b involved 40 undergraduate students from Maynooth University (24 females and 16 males) aged between 18 and 42 years ($M = 22.78$, $SD = 5.32$), all selected by random convenience sampling. The study comprised of two IRAPs, referred to as the *Typical Length IRAP* and the *Double Length IRAP*, between which all participants were evenly and semi randomly assigned.

Setting

All aspects of the setting in Experiment 3b were identical to the previous study.

Apparatus and Materials

Most aspects of Experiment 3b were identical to Experiment 3a (i.e. the Spider Fear IRAP was employed, along with the FSQ and BAT; see Experiment 3a for outlines of these materials), except that the length of the IRAP to which participants was exposed was manipulated through the distinction between the Typical Length IRAP and the Double Length IRAP.

Experiment 3b involved two variations of the same Spider Fear IRAP from the previous study (see Figure 9). The Typical Length IRAP was identical to the Typical Switching IRAP (and thus identical to a standard IRAP) presented in Experiment 3a, while the Double Length IRAP *differed from a typical IRAP only with regard to the number of test blocks.*

Procedure

The experimental sequence of Experiment 3b comprised of three stages that were identical for all participants. Stage 1 involved the IRAP, Stage 2 involved the explicit measure (the FSQ) and Stage 3 presented the BAT. Unless otherwise stated, all procedural aspects of the Spider Fear IRAP were identical to the Typical IRAP as presented in Experiment 1.

Stage 1: Spider-Fear IRAPs (Typical Length IRAP and Double Length IRAP). In Stage 1, participants completed one of the two IRAPs, either the Typical Length IRAP or the Double Length IRAP. *The double length IRAP differed from a typical IRAP only with regard to the number of test blocks. That is, instead of the typical 6 test blocks (i.e. 3 per rule), the double length IRAP presented 12 test blocks (6 per rule), although each test block contained the typical 32 trials. Overall, therefore, this modified IRAP presented a total of 384 trials -- that is double the*

number of trials presented in the Typical Length IRAP (i.e. 192). See Figure 10. Half of the participants were exposed to the Typical Length IRAP, while the other half were exposed to the Double Length IRAP.

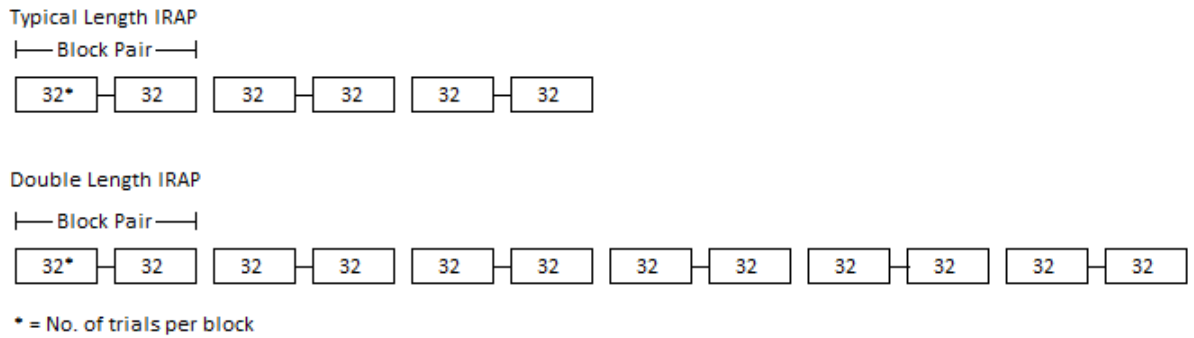


Figure 10. A brief illustration of the methodological differences between the Typical Length IRAP and the Double Length IRAP in Experiment 3b.

Stage 2: FSQ. In Stage 2, all participants completed the FSQ.

Stage 3: BAT. In Stage 3, all participants completed the BAT. They were thereafter fully debriefed and thanked for their participation (see Appendix 8).

IRAP Data Preparation

While most aspects of the data preparation for the Spider Fear IRAP were identical to the previous study (see Experiment 1 for in depth outline of the general *D*-score calculation procedure), a few minor modifications were necessary for the analyses of data from the Double Length IRAP. Specifically, *D*-scores were averaged across each pair of test blocks to produce six sets of *D*-scores (not three as is typical due to the double length nature of the IRAP), with each set consisting of four *D*-scores, one for each trial-type. Following this, instead of averaging *D*-scores across test blocks which is the usual convention, sequential averages were created across test block pairs to calculate a *rolling D-score*. This resulted in 6 overall rolling *D*-scores. This helped to assess the relative utility of incremental numbers of test block pairs (i.e. block pair

1, average of block pairs 1 + 2, average of block pairs 1 + 2 + 3, etc). Although these scores do not inform us about the ways in which performance may change across individual test block pairs, they do show how adding additional test blocks may change the overall performance on the task up to that point. Three participants did not meet the pass criteria for the double length IRAP leaving $N = 37$.

Ethical Considerations

All aspects of ethical consideration that applied to Experiment 3a also applied to the current study (see Ethical Considerations for Experiment 3a).

Analytic Strategy

Experiment 3b sought to systematically compare responding on two IRAPs, in which the total number of test blocks was manipulated via presentation of the Typical Length IRAP versus the Double Length IRAP (i.e. 192 trials vs. 384 trials, respectively). The Spider Fear IRAP was presented in both cases, along with the FSQ and the BAT. Again, the FSQ data were analysed in terms of normative level of spider fear. The *D*-scores on each trial-type were compared across the two IRAPs and a 4x2 repeated measures ANOVA determined possible main effects for IRAP type or trial-type. Instead of averaging *D*-scores across test blocks which is the usual convention, sequential averages were created across test block pairs to calculate a *rolling D-score* to assess the relative utility of incremental numbers of test block pairs. One sample t-tests determined the significance of any of the *D*-scores. Dependent t-tests determined whether there were any differences in participant scores between a Typical Length IRAP of three test block pairs compared to one consisting of four, five and six test

block pairs respectively across each trial-type. Finally, a correlation matrix explored the potential relationships among the *D*-scores, the FSQ and the BAT.

Results

FSQ Data

The mean on the FSQ was 53.14 ($SD = 29.28$). This indicated relatively normal levels of fear, given that the maximum possible score on the scale is 108.

IRAP Data

The mean rolling *D*-scores for each trial-type for each IRAP are illustrated in Figure 11. That is, scores up to three block pairs represent the Typical Length IRAP while those from four to six block pairs inclusive represent the Double Length IRAP. For reasons of parsimony it should be noted that to compare the Typical Length IRAP and the Double Length IRAP as presented procedurally, the blue (block pairs 1-3 or Typical Length IRAP) and purple (block pairs 1-6 or Double Length IRAP) respectively, should be compared. To assess the way in which adding additional test blocks may change the overall performance on the task up to that point as stated above, participant responses were divided in a manner that resembled four IRAPs of increasing block length (additional red and green bars below), but this was done on a post-hoc basis. Responding on all four block pairs showed modest fear responding on the Fear-Spiders trial-type and on the Fear-Nature trial-type. There was moderate approach responding on the Approach-Spiders trial-type again in all four comparisons. Similarly, on Approach-Nature, all response patterns showed strong approach responding.




Figure 11. Mean rolling *D*-scores, with standard error bars, for the four trial-types in the Spider fear IRAPs at the three, four, five and six block pair length in Experiment 3a. **Positive *D*-scores indicate approach responding and negative *D*-scores indicate fear responding.**

The *D*-scores on all four block pairs were subjected to four mixed 4x2 ANOVAs, with trial-type and block pair (IRAP) as the between subjects variables. Main effects were found for each block pair (Block Pairs 1-3/Typical Length IRAP: $f(3, 36) = 11.854, p < .0001$; Block Pairs 1-4: $f(3, 36) = 16.322, p < .0001$; Block Pairs 1-5: $f(3, 36) = 16.717, p < .0001$; Block Pairs 1-6/Double Length IRAP: $f(3, 36) = 19.076, p < .0001$).

Sixteen one-sample *t*-tests determined whether the trial-types in each block pair differed significantly from zero. Approach-Spiders (Block Pairs 1-3/Typical Length: $M = .168, t(36) = 3.340, p = .002$; Block Pairs 1-4: $M = .156, t(36) = 3.194, p = .003$; Block Pairs 1-5: $M = .132, t(36) = 2.793, p = .008$; Block Pairs 1-6/Double

Length: $M = .148$, $t(36) = 3.232$, $p = .003$) and Approach Nature (Block Pairs 1-3/Typical Length: $M = .288$, $t(36) = 4.268$, $p = .0001$; Block Pairs 1-4: $M = .305$, $t(36) = 5.019$, $p < .0001$; Block Pairs 1-5: $M = .294$, $t(36) = 5.119$, $p < .0001$; Block Pairs 1-6/Double Length: $M = .296$, $t(36) = 5.842$, $p < .0001$) were significant on all IRAPs. Fear-Spiders was significant on the Block Pairs 1-4, Block Pairs 1-5, and Block Pairs 1-6/Double Length IRAPs (Block Pairs 1-4: $M = .115$, $t(36) = 2.464$, $p = .02$; Block Pairs 1-5: $M = .104$, $t(36) = 2.402$, $p = .02$; Block pairs 1-6/Double Length: $M = .132$, $t(36) = 3.448$, $p = .002$), while Fear-Nature ($M = -.110$, $t(36) = -1.793$, $p = .054$) seemed to approach significance at Block Pairs 1-4 (all other p 's $> .08$).

Three dependent t-tests were also conducted to determine if there was any difference in participant scores between a Typical Length IRAP (Block Pairs 1-3) and all other block pairs including Block Pairs 1-6/Double Length IRAP across each trial-type. There were no significant differences (all p 's $> .24$).

Correlations Among the IRAPs and Other Measures. A correlation matrix calculated the potential relationships among the trial-type rolling D -scores on the IRAPs and the other measures. Only the fear-spiders trial-type in the Double Length IRAP (Block Pairs 1-6) approached significance in a moderate correlation with the FSQ ($r = .313$, $p = .0548$). Again, as was the case in Experiment 3b, to reduce the possible statistical error associated with running correlational matrices, the Bonferroni correction were run for statistical integrity. This resulted in the p value for this correlation being corrected to .002. For the purpose of this thesis, the data before this correction will be mainly discussed.

Summary of Results

Patterns of responding were very similar across the Typical and Double Length IRAPs and indeed across all block pair comparisons. Responding on each IRAP did not correlate with the BAT or FSQ, although the Fear-Spiders trial-type on the Double Length IRAP did appear to approach significance.

Discussion

Given the data from Experiment 3a, we employed the Spider Fear IRAP again in Experiment 3b, but now manipulated the number of test blocks presented by comparing the Double Length IRAP with the Typical Length IRAP. Again, our aim was to determine whether the typical number of blocks exerts its own influence on outcomes normally observed with the IRAP. The results indicated that responding was very similar on both IRAPs and thus block number, in and of itself, exerted little influence, at least in the context of the Spider Fear IRAP. The use of the rolling *D*-scores was particularly useful in this context

Chapter 5

General Discussion

Chapter 5

General Discussion

The current thesis sought to investigate the possible role of various methodological features of the IRAP on performances, as typically observed. This seemed timely given that the IRAP is increasingly used to present more complex stimuli in a broader range of domains. We specifically targeted four features: label-target stimulus combinations, rule order, task switching and block length. The current chapter is divided as follows. In the first section, the rationale and findings from Experiment 1 are summarised. This is followed accordingly by Experiment 2, and Experiments 3a and b thereafter. In each case, the data recorded here are compared with those from the existing literature, and any methodological or conceptual issues pertaining directly to any of the experiments are also discussed in this context.

Experiment 1 (Chapter 2): Findings and Issues

Very little published research has asked specific methodological questions about the IRAP, despite considerable attention to this issue in IAT research (Greenwald et al., 2009). From the limited body of evidence available thus far, such as work by Campbell et al. (2009), one can conclude overall that there is some promise in presenting more complex stimuli within an IRAP (especially in clinical contexts), but this must be done carefully with regard to the various procedural parameters. In order to contribute to this small body of work, Experiment 1 compared a Typical IRAP with a Natural Language IRAP using a simple fruit-insects comparison to determine the possible impact of combining the label and target stimuli. Our basic aim was to

determine whether this would yield IRAP outcomes that were dissimilar to those typically observed.

The findings from Experiment 1 showed patterns of responding that were very similar across the Typical and Natural Language IRAPs. In both cases, participants showed strong positive responding to Fruit-Positive and Fruit-Negative. There was also strong positive responding to Insects-Positive on the Natural Language IRAP, with a weaker non-significant effect on the Natural Language IRAP. In both cases, there was very weak responding on Insects-Negative. Overall, responding on one IRAP did not differ *significantly* from the other, although one might generously conclude that, if anything, the stronger effect overall emerged with the Natural Language IRAP. Fruit-Positive in this IRAP was the only trial-type that correlated with the NART, although this explicit measure was only used to control for IQ, which was overall in the normal range. Overall, the outcomes on the IRAPs were broadly similar and thus consistent with our provisional predictions.

In terms of recording basic effects on a simple fruit-insects IRAP, the findings from Experiment 1 are similar to the existing evidence base with regard to the basic effects in IRAPs that presented relatively basic stimuli (e.g. Campbell et al., 2011; McKenna et al., 2009). All published studies that have employed an IRAP of this nature have reported strong responding on the positive trial-type described as consistent, and this was also observed currently (e.g. ‘fruit-positive-true’ and ‘pleasant-love-similar’). There can often be variability observed in the other trial-types due to the different procedural variations and different samples that have been employed across the various studies which can make direct comparability somewhat difficult. However, in terms of the basic trends observed with IRAPs of this kind, the current study’s effects are generally similar.

In the existing literature, we identified only seven studies to which Experiment 1 could be directly compared. Three of these (Carpenter et al., 2012; Hussey et al., 2012; Hooper et al., 2010) had previously used target stimuli that were more complex than single words, while two (Kosnes et al., 2013; Parling et al., 2012) used more complex label stimuli. Campbell et al. (2009) manipulated the locations of the label stimuli, and De Houwer et al. (in press) effected a series of manipulations in the RRT. With the exception of the work by Campbell et al., none of the above systematically manipulated IRAP features in terms of comparing changes with a typical IRAP, hence it is difficult to draw firm conclusions about the possible impact of these changes. However, it is the case that the IRAP effects observed overall did not appear to differ greatly from those typically recorded in IRAP studies. As a result, one can conclude that adding these types of complexity to the various stimulus presentations does not appear to alter the IRAP effects one would expect to see. And the data from Experiment 1 support this conclusion. That is, outcomes on the Natural Language and Typical IRAPs did not differ markedly. While this is a positive sign for further use of complex stimuli within the IRAP, this move is perhaps even more promising when one considers that Experiment 1 indicated, albeit very modestly, that the effects from the Natural Language IRAP were even stronger than the Typical IRAP.

It has been argued that the IRAP is difficult and/or lengthy to complete, at least relative to other implicit measures (De Houwer et al., in press). However, each IRAP here took on average less than 10 minutes to complete in its entirety, including instructions, and this is not unusual. Furthermore, the total attrition rate was 8.3%, which is again consistent with other IRAP studies, particularly those in which the procedure is delivered by a trained researcher. It is important to emphasise, therefore, that the typical effects observed in the Natural Language IRAP did come at a cost of

more time by participants or the researcher. In all practical respects, it did not differ from the Typical IRAP. This suggests that the Natural Language IRAP offers researchers the possibility of presenting more complex relations, but within a procedure that remains highly user-friendly. Again, the latter is particularly important if the IRAP is to be used more often to study clinical phenomena and especially if it is to be presented to clinical populations.

In terms of the wider psychological literature, the importance of studying the link between environmental events, derived relations, and the physiological activity that takes place in the brain and the central nervous system has been noted and discussed by a number of authors (e.g. Hayes & Bisst, 1998, Barnes-Holmes et al., 2004). In this respect, and specifically with regard to the natural language statement manipulation above, the cognitive neuroscientific literature has examined the effects of sentence manipulation on semantic processing via event related potentials (ERPs; i.e. time-locked electrical signals locked to a stimulus or set of stimuli that are presented repeatedly [Barnes-Holmes et al., 2004]). This seems particularly warranted given the study conducted by Barnes-Holmes et al. (2004) which suggested that semantic and derived stimulus relations may encompass a functional overlap, sensitively measured by ERPs on an IAT. Much research conducted in the cognitive neuroscientific literature has examined the manipulation of sentences as measured by ERPs. Studies of this nature have typically found that an N400 (a 400ms negatively peaked ERP) will be elicited when participants encounter an incongruent, or also merely an unexpected, end to a sentence when they are asked to silently read a sentence presented to them one word at a time (Kutas, Neville, & Holcomb, 1987). It has been suggested that the degree of N400 is a function of how easy the word coheres with the given context (e.g. Van Berkum et al., 1999; in Diaz & Swaab, 2007). This, alongside

the fact that this anomaly occurs even when the end to a sentence is congruent but merely unexpected, suggests possible avenues of interest relating to the natural language statements used in the IRAP and potential ERP analysis. These findings and the general area of research may prove to be beneficial in terms of considering future research into IRAP development and response measure when considering semantic stimulus presentation.

It is perhaps useful to provide a cautionary note regarding possible uses of the Natural Language IRAP. In short, the fact that the findings here indicate that using short natural language statements does not yield different outcomes from separate labels and targets does not mean that caution should not be exercised when researchers construct these stimulus arrangements. For example, adapting items from questionnaires is notoriously difficult in the IRAP (even if the questionnaire itself is psychometrically sound) and it often makes little sense to participants when reading the screen. A key issue here appears to be the importance of the researcher identifying clearly which relations are to be targeted. For example, one might be tempted to take items, such as “*My painful experiences and memories make it difficult for me to live a life that I would value*”, from the Acceptance and Action Questionnaire (Bond et al., 2011) and use them in an IRAP. However, one must then ask about the construction of the three remaining trial-types and about what the researcher is trying to achieve with this statement. In short, the IRAP is not merely a questionnaire that is completed under accuracy and time constraints, rather it is a tool with which to conduct functional analyses of relational responding (see Hussey, Barnes-Holmes, & Barnes-Holmes, 2015).

Experiment 2 (Chapter 3): Findings and Issues

The use of pre-specified, alternating rules for responding is one of the core tenets of the IRAP's structure and yet there has been little research to date on this variable. The limited evidence base that is available suggests that rules may exert some influence over IRAP effects, but perhaps only in certain contexts or specific IRAPs (e.g. McKenna et al., 2009; O'Shea et al., 2015). The robust effects recorded across several Self-esteem IRAPs (e.g. Scanlon et al., 2014; Vahey et al., 2009) suggested that this would be a useful context in which to conduct rule-based manipulations. Hence, Experiment 2 investigated a rule order manipulation in the context of a Self-esteem IRAP. Specifically, we manipulated the order in which the self- vs. others-rules were presented. This manipulation simply sought to determine whether the simple order in which the rules were presented influenced IRAP effects. The findings from Experiment 2 showed patterns of responding that differed somewhat between the two conditions, in a manner that was much more distinct than that recorded between the two IRAPs in Experiment 1. In both conditions in Experiment 2, participants showed strong significant positive responding to Self-Positive, with the larger effect when the self rule was presented first. There was no clear response on Self-Negative when the other rule was presented first, but a modest negative response when the self rule came first. The two conditions diverged even more on Others-Positive, with almost no response recorded when the self rule came first, but a modest positive response when the other rule came first. Divergence, albeit less, was also recorded on Others-Negative with a marginal negative response when the self rule came first, compared with a very marginal positive response when the other rule came first. It was surprising, therefore, that the two conditions did not differ significantly on any trial-type. Furthermore, responding on the Others-Rule First Condition correlated

with depression and self-esteem, such that greater implicit positivity towards others correlated with higher self esteem and depression. It should be once again noted that it is common practice in IRAP research for analyses and speculation regarding non-significant findings to be noted in a manner that reflects the single subject focus of traditional behavioural research. The writer is conscious of this and it must be stressed that the resulting discussion points are only speculative as the conditions themselves did not differ significantly.

In terms of recording effects on the Self-esteem IRAP, the findings from Experiment 2 are similar to the existing evidence base with regard to the basic IRAP effects (e.g. Scanlon et al., 2014; Remue et al., 2013; Vahey et al., 2009). All published studies that have employed a variation of the Self-esteem IRAP have reported strong self-positive responding, as demonstrated currently. Similarly, the variability in the Others trial-types which we recorded also matches variability in this regard as recorded in previous studies, although admittedly different procedural variations and different samples have been employed across the various studies.

In terms of determining the influence of the rule manipulation, the key study with which the findings from Experiment 2 can be directly compared is the work by O'Shea et al. (in press). In that study, the more established verbal relations (on the weight IRAP and the nature IRAP) were less sensitive to rule framing than less established relations (on a nonsense syllable IRAP and a social systems IRAP). If one looks closely at the Self-esteem IRAP in Experiment 2, one can find a pattern of results that supports this outcome. That is, both conditions produced significant positive performances on the Self-Positive trial-type, which clearly is a well-established relation (participants were overall categorised as normal on the DASS and RSES). Hence, the rule manipulation had little or no influence on this effect. However, the

small *D*-scores and divergent patterns of responding recorded on the other trial-types, as well as the differences that emerged between conditions, suggest perhaps that these involved less well-established relations (e.g. Others-Negative) and are thus more sensitive to the rule manipulation.

Experiments 3a and 3b (Chapter 4): Findings and Issues

While IRAP research has moved at a considerable pace towards an ever-increasing range of domains, methodological questions, such as those investigated in Experiments 1 and 2, remain. While the explorations in the two previous studies of the label-target combinations in the Natural Language IRAP (Experiment 1) and rule ordering (Experiment 2) are important in terms of the presentation of the IRAP, it could be argued that deeper procedural questions need to be asked, especially if the IRAP is to meet the complex challenge of studying clinically-relevant relations. One such area of investigation that has concerned IAT researchers involved the *length* of the test blocks (e.g. Back et al., 2005; Klauer & Mierke, 2005; Nosek, et al., 2007). This was the focus of Experiment 3a.

Experiment 3a. In Experiment 3a, we employed the Spider Fear IRAP and manipulated the length of the test blocks presented between rules. To do so, we compared a Fast Switching IRAP and a Typical Switching IRAP. Specifically, the Fast Switching IRAP halved the length of the test blocks between rules, relative to the Typical Switching IRAP, such that each block now contained only 16 trials, and not 32.

The findings from Experiment 3a, in terms of the overlap between the two IRAPs, fell somewhere between Experiment 1's and Experiment 2's outcomes. In both cases, participants showed negative responding to Fear-Spiders, with a larger effect in

the Fast Switching IRAP. Both showed positive responding on all three remaining trial-types. There was strong significant positive responding to Approach-Spiders, again stronger in the Fast Switching IRAP. Responding was more modest on Fear-Nature, but now stronger with a in the Typical Switching IRAP. Stronger significant responding was recorded on Approach-Nature, with the larger effect also in the Typical Switching IRAP. Indeed, responding on the two IRAPs differed significantly on Fear-Spiders, while the difference on Approach-Nature approached significance. Responding on Fear-Nature in the Fast Switching IRAP correlated positively with the FSQ and a negative correlation with the BAT approached significance.

In summary, Experiment 3a indicated that the rate of block switching appeared to exert some influence over the outcomes. While it is difficult to specify the nature of this influence precisely, closer inspection of Figure 9 indicates that the Fast Switching IRAP produced stronger effects than the Typical IRAP on both of the spider trial-types, while the Typical Switching IRAP produced stronger effects on both of the nature trial-types. One could argue, therefore, that the spiders trial-types that are most critical to the Spider Fear IRAP (the nature trial-types are simply the contrast categories) generate stronger *D*-scores when participants are forced to switch more quickly than is typically the case in an IRAP. This might suggest, perhaps, that IRAPs targeting more socially sensitive or clinically-relevant domains should consider increasing the rate at which switching is forced.

In terms of recording effects on the Spider Fear IRAP, the findings from Experiment 3a are similar with Nicholson and Barnes-Holmes (2012). Similar to this original study, participants currently showed a significant fear of spiders, but only on the Fast Switching IRAP. Both studies also recorded no fear responding to nature and strong approach nature responses. However, while Nicholson and Barnes-Holmes

reported little or no effect for approaching spiders (even for participants low in the FSQ), both IRAPs here showed significant spider approach responding.

In terms of determining the influence of the block length manipulation, there were no existing published IRAP studies with which we could compare the findings from Experiment 3a, and most of the IAT studies, have focused on the relationship between block switching and cognitive models, than on the possible influence of switching as a variable in IAT outcomes. For example, Klauer and Mierke's (2005) task switching cost model has been tested extensively within the IAT. The findings from various studies (e.g. Back et al., 2001; Rubinstein et al., 2001) suggest that a higher rate of task switching should lead to a shallower response pattern. However, the findings from the current study show overall *strong* responding in the Fast Switching IRAP, with three of the four trial-types producing significant *D*-scores. Nonetheless, in a similar vein, Sohn and Anderson (2001) argued that familiarity and exposure can reduce interference from task switching and thus reduce the cost. Hence, one could argue that that was the case here and that strong performances were observed on the Fast Switching IRAP because of the very common verbal relations (spiders and nature) that were targeted.

With regard to the wider cognitive and neuropsychological literature, it may also merit some discussion to consider the IRAP, and in particular its task switching components, in the context of executive function. Executive function in this domain can be thought of as the planning and organisation of secondary tasks performed by the cortical and subcortical areas of the brain, and is thought to be one of the primary functions of the prefrontal cortex (Aron, Robbins, & Poldrack, 2004). Within this broader literature, the use of other measures involving a task switching element have often suggested that performance on these measures are mediated by executive

function to some degree, of which task switching is thought to be a primary component (Banich, 2009). For example, research involving the Wisconsin Card Sorting Task, a task of which the core tenet is task switching and designed to study executive control, has suggested that this task can in some cases be used as a direct measure of frontal lobe damage (Strauss, Sherman, & Spreen, 2006, in Nyhus & Barcelo, 2009). In light of this, future research using the IRAP, and particularly in a similar methodological vein to the current thesis, should aim to explore its possible sensitivity to frontal dysfunction and the possible related confounds and implications.

Experiment 3b. Given the data from Experiment 3a and the influence of manipulating block length, we employed the Spider Fear IRAP again in Experiment 3b in order to manipulate number of test blocks. To do this, we compared a Double Length IRAP (12 test blocks) and a Typical Length IRAP (6 test blocks). Again, this manipulation sought to determine whether the typical number of blocks exerts its own influence on outcomes normally observed with the IRAP.

Unlike the findings from Experiment 3a, the manipulation of block number appeared to exert little or no influence on the IRAP outcomes. That is, the patterns of responding were very similar across the Typical and Double Length IRAPs. The data showed negative responding to Fear-Spiders and positive responding to the three remaining trial-types, all irrespective of number of block pairs. All block comparisons on Approach-Spiders and Approach-Nature were significant, as were the comparisons on Fear-Spiders, with the exception of Block Pairs 1-3 (Typical Length IRAP). None of the effects on Fear-Nature were significant. There were basically no correlations with the FSQ or BAT.

In terms again of recording effects on the Spider Fear IRAP, the findings from Experiment 3b are similar with Nicholson and Barnes-Holmes (2012). Similar to this

original study and Experiment 3a, participants in Experiment 3b showed a significant fear of spiders, but this effect became non-significant at the typical length. All three studies also recorded no fear responding to nature and strong approach nature responses. And as above, while Nicholson and Barnes-Holmes reported little or no effect for approaching spiders, Experiment 3a showed significant spider approach responding. And this effect was recorded again in Experiment 3b. Hence, with the exception of approaching spiders, both Experiments 3a and 3b recorded similar effects to the original Spider Fear IRAP study by Nicholson and Barnes-Holmes.

In terms of determining the influence of the block number manipulation, there were again no existing published IRAP studies with which we could compare the findings from Experiment 3b and indeed little or no IAT studies.

The current analyses employing rolling *D*-scores not only illustrated most clearly the comparison between the Typical length IRAP (Block Pairs 1-3) and the Double Length IRAP (Block Pairs 1-6), and all remaining block pairs, but allowed a more fine-grained analysis of all block pairs that is usually analysed in IRAP research. Interestingly, this detailed analysis showed very little fluency in responding across blocks, suggesting a lack of practice effects. It was also interesting that this applied equally to the Double Length IRAP, as to the Typical IRAP. These findings lend some support to the view that the IRAP is targeting participants' pre-experimental verbal histories, rather than generating specific patterns in vivo.

Concluding Comments

The current thesis was brief and succinct in ambition and design, but placed the importance of systematic comparisons at its core. It may be deemed somewhat alarming that the IRAP is being used increasingly to study clinical phenomena and even to try to distinguish between clinical and non-clinical groups, when many of its basic procedural parameters have not been tested. As a result, any effects recorded with complex IRAPs cannot necessarily be assumed genuine because they may simply reflect procedural artefacts. The current thesis was designed to alert or perhaps even reassure IRAP researchers of the need and utility of various task features. It was surprising overall that the manipulations we conducted exerted almost no significant influence on the outcomes. Of course, other manipulations to the IRAP more general in nature to the task than those presented in the current thesis are also of potential interest to future research. These may include handedness and the use of one-handed responding as opposed to the two involved in the IRAP currently for example. However as far as the current thesis is concerned, in short, one can conclude from that many of the IRAP effects that have been published are likely to be robust and genuine IRAP effects and not merely methodological artefacts. At one level this is good news for the published body off IRAP evidence. And, at another level the findings from the current studies are good news for more extensive and complex uses of the IRAP in the future.

References

References

- Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, 8(4), 170-177.
- Back, M. D., Schmukle, S. C., & Egloff, B. (2005). Measuring task switching ability in the Implicit Association Test. *Experimental Psychology*, 52(3), 167-179.
- Banaji, M. R., & Hardin, C. D. (1996). Automatic stereotyping. *Psychological Science*, 7, 136-141.
- Banich, M. T. (2009). Executive function: The search for an integrated account. *Current Directions in Psychological Science*, 18(2), 89-94.
- Barnes-Holmes, D., Barnes-Holmes, Y., Stewart, I., & Boles, S. (2010). A sketch of the Implicit Relational Assessment Procedure (IRAP) and the Relational Elaboration and Coherence (REC) model. *The Psychological Record*, 60, 527-542.
- Barnes-Holmes, D., Barnes-Holmes, Y., Power, P., Hayden, E., Milne, R., & Stewart, I. (2006). Do you really know what you believe? Developing the Implicit Relational Assessment Procedure (IRAP) as a direct measure of implicit beliefs. *The Irish Psychologist*, 32, 169-177.
- Barnes-Holmes, D., Murphy, A., Barnes-Holmes, Y., & Stewart, I. (2010). The Implicit Relational Assessment Procedure (IRAP): Exploring the impact of private versus public contexts and the response latency criterion on pro-white and anti-black stereotyping among white Irish individuals. *The Psychological Record*, 60, 57-66.
- Barnes-Holmes, D., Staunton, C., Barnes-Holmes, Y., Whelan, R., Stewart, I., Commins, S., Walsh, D., Smeets, P., & Dymond, S. (2004). Interfacing relational frame theory with cognitive neuroscience: Semantic Priming, the

- Implicit Association, and event related potentials. *International Journal of Psychology and Psychological Therapy*, 4(2), 215-240.
- Bond, F. W., Hayes, S. C., Baer, R. A., Carpenter, K. M., Guenole, N., Orcutt, H. K., Waltz, T., & Zettle, R. D. (2011). Preliminary psychometric properties of the Acceptance and Action Questionnaire-II: a revised measure of psychological inflexibility and experiential avoidance. *Behavior Therapy*, 42(4), 676–688.
- Campbell, C., Barnes-Holmes, Y., Barnes-Holmes, D., & Stewart, I. (2011). Exploring screen presentations in the Implicit Relational Assessment Procedure (IRAP). *International Journal of Psychology and Psychological Therapy*, 11(3), 377-388.
- Carpenter, K. M., Martinez, D., Vadhan, N. L., Barnes-Holmes, D., & Nunes, V. E. (2012). Performance based measures of attentional bias and relational responding are associated with behavioural treatment outcome for cocaine dependence. *The American Journal of Drug and Alcohol Abuse*, 38, 146-154.
- Crawford, J. R. & Parker, D. M. (1989). Prediction of WAIS IQ with the National Adult Reading Test: Cross validation and extension. *British Journal of Clinical Psychology*, 28, 267-273.
- Cullen, C., Barnes-Holmes, D., Barnes-Holmes, Y., Stewart, I. (2009). The implicit relational assessment procedure (IRAP) and the malleability of ageist attitudes. *The Psychological Record*, 59, 591-620.
- Diaz, M. T. & Swaab, T. Y. (2007). Electrophysiological differentiation of phonological and semantic integration in word and sentence contexts. *Brain Research*, 1146, 85-100.
- De Houwer, J. (2001). A structural and process analysis of the Implicit Association Test. *Journal of Experimental Social Psychology*, 37, 443-451.

- De Houwer, J. (2002). The Implicit Association Test as a tool for studying dysfunctional associations in psychopathology: Strengths and limitations. *Journal of Behavior Therapy and Experimental Psychiatry*, 33, 115-133.
- De Houwer, J. (2003). The Extrinsic Affective Simon Task. *Experimental Psychology*, 50(2), 77-85.
- De Houwer, J. (2014). A propositional model of implicit evaluation. *Social and Personality Psychology Compass*, 8(7), 342-353.
- De Houwer, J., Heider, N., Spruyt, A., Roets, A., & Hughes, S. (2015). The relational responding task; toward a new implicit measure of beliefs. *Frontiers in Psychology*, 6, 1-9.
- De Houwer, J., Teige-Mocigemba, S., Spruyt, A., & Moors, A. (2009). Implicit measures: A normative analysis and review. *Psychological Bulletin*, 135(3), 347-368.
- De Raedt, R., Schacht, R., Franck, E., & De Houwer, J. (2006). Self-esteem and depression revisited: Implicit positive self-esteem in depressed patients. *Behaviour Research and Therapy*, 44(7), 1017-1028.
- Dymond, S. & Roche, B. (2013). *Advances in Relational Frame Theory & Contextual Behavioral Science: Research & Application*. CA, USA: New Harbinger Publications.
- Egloff, B. & Schmukle, S. C. (2002). Predictive validity of an Implicit Association Test for assessing anxiety. *Journal of personality and Social Psychology*, 83(6), 1441-1455.
- Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, 69, 1013-1027.

- Gamer, J., Schmukle, S. C., Luka-Krausgrill, U., Egloff, B. (2008). Examining the dynamics of the implicit and explicit self-concept in social anxiety: Changes in the Implicit Association Test-Anxiety and the Social Phobia Anxiety Inventory following treatment. *Journal of Personality Assessment*, 90(5), 476-480.
- Gawronski, B., & De Houwer, J. (2014). Implicit measures in social and personality psychology. In H. T. Reis, & C. M. Judd (2nd Ed), *Handbook of research methods in social and personality psychology* (pp. 283-310). New York, NY: Cambridge University Press.
- Greenwald, A. G. & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, 102(1), 4-27.
- Greenwald, A. G. & Farnham, S. G. (2000). Using the Implicit Association Test to measure self-esteem and self-concept. *Journal of Personality and Social Psychology*, 79(6), 1022-1038.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, 74(6), 1464.
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the implicit association test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, 85(2), 197-216.
- Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*, 97(1), 17-41.

- Greenwald, A. G., Smith, C. T., Sriram, N., Bar-Anan, Y., & Nosek, B. A. (2009). Implicit race attitudes predicted vote in the U.S. presidential election. *Analyses of Social Issues and Public Policy*, 9(1), 241-253.
- Hagborg, W. J. (1993). The Rosenberg Self-Esteem Scale and Harter's Self-Perception Profile for adolescents: A concurrent validity study. *Psychology in the Schools*, 30, 132-136.
- Hayes, S. C., Barnes-Holmes, D., Roche, B. (2001). *Relational frame theory: A post-Skinnerian account of human language and cognition*. New York, MA: Kluwer Academic Publishers.
- Hayes, S. C. & Bisset, R. T. (1998). Derived stimulus relations produce mediated and episodic priming. *The Psychological Record*, 48, 617-360.
- Henry, J. D., & Crawford, J. R. (2005). The short-form version of the Depression Anxiety Stress Scales (DASS-21): Construct validity and normative data in a large non-clinical sample. *British Journal of Clinical Psychology*, 45, 227-239.
- Hooper, N., Villatte, M., Neofotistou, E., & McHugh, L. (2010). An implicit versus explicit measure of experiential avoidance. *International Journal of Behavior Consultation and Therapy*, 6(3), 233-244.
- Hussey, I., & Barnes-Holmes, D. (2012). The Implicit Relational Assessment Procedure as a measure of implicit depression and the role of psychological flexibility. *Cognitive and Behavioral Practice*, 19, 573-582.
- Hussey, I., Barnes-Holmes, D., & Barnes-Holmes, Y. (2015). From relational frame theory to implicit attitudes and back again: Clarifying the link between RFT and IRAP research. *Current Opinion in Psychology*, 2, 11–15.
- Klauer, K. C. & Mierke, J. (2005). Task-set inertia, attitude accessibility, and compatibility-order effects: New evidence for a task-set switching account of

- the Implicit Association Test effect. *Personality and Social Psychology Bulletin*, 31(2), 208-217.
- Kosnes, L., Whelan, R., O'Donovan, A., McHugh, L. A. (2013). Implicit measurement of positive and negative future thinking as a predictor of depressive symptoms and hopelessness. *Consciousness and Cognition*, 22, 898-912.
- Kutas, M., Neville, H. J., & Holcomb, P. J. (1987). A preliminary comparison of the N400 response to semantic anomalies during reading, listening and signing. In R.J. Ellingson, N. M. F. Murray & A. M. Haliday (Eds.), *The London Symposia* (EEG suppl. 39)(p. 325-330). Elsevier Science Publishers.
- Lovibond, S. H., & Lovibond, P. F. (1995). *Manual for the Depression Anxiety Stress Scales* (2nd ed.). Sydney: The Psychology Foundation of Australia.
- McConnell, A. R. & Leibold, J. M. (2001). Relations among the Implicit Associations Test, discriminatory behaviour, and explicit measures of racial attitudes. *Journal of Experimental Social Psychology*, 37, 435-442.
- McKenna, I., Barnes-Holmes, D., Barnes-Holmes, Y., & Stewart, I. (2007). Testing the fake-ability of the Implicit Relational Assessment Procedure (IRAP): The first study. *International Journal of Psychology and Psychological Therapy*, 7(2), 253-268.
- Nelson, H. E. (1982). *National Adult Reading Test*. Windsor, UK: NFER-Nelson.
- Nicholson, E., & Barnes-Holmes, D. (2012). The Implicit Relational Assessment Procedure (IRAP) as a measure of spider fear. *The Psychological Record*, 62, 263-278.
- Nicholson, E., Dempsey, K., & Barnes-Holmes, D. (2014). The role of responsibility and threat appraisals in contamination fear and obsessive-compulsive

- tendencies at the implicit level. *Journal of Contextual Behavioral Science*, 3(1), 31-37.
- Nock, M., K., Park, J. M., Finn, C. T., Deliberto, T. L., Dour, H. J., & Banaji, M. R. (2010). Measuring the suicidal mind: Implicit cognition predicts suicidal behaviour. *Psychological Science*, 21, 511-517.
- Nosek, B. A. & Banaji, M. R. (2001). The go/no-go association task. *Social Cognition*, 19(6), 625-664.
- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2007). The Implicit Association Test at age 7: A methodological and conceptual review (pp. 265–292). In J. A. Bargh (Ed.), *Automatic Processes in Social Thinking and Behavior*. Psychology Press.
- Nyhus, E. & Barcelo, F. (2009). The Wisconsin Card Sorting Test and the cognitive assessment of prefrontal executive functions: A critical update. *Brain and Cognition*, 71(3), 437-451.
- O'Carroll, R. E. (1987). The inter-rater reliability of the National Adult Reading Test (NART): a pilot study. *British Journal of Clinical Psychology*, 26, 229–230.
- O'Shea, D., Watson, D. G., & Brown, G. D. A. (in press). Measuring implicit attitudes: A positive framing bias flaw in the Implicit Relational Assessment Procedure (IRAP). *Psychological Assessment*.
- Parling, T., Cernvall, M., Stewart, I., Barnes-Holmes, D., & Ghaderi, A. (2012). Using the Implicit Relational Assessment Procedure to compare implicit pro-thin/anti-fat attitudes of patients with anorexia nervosa and non-clinical controls. *Eating Disorders: The Journal of Treatment and Prevention*, 20(2), 127-143.

- Power, P., Barnes-Holmes, D., Barnes-Holmes, Y., Stewart, I. (2009). The Implicit Relational Assessment Procedure (IRAP) as a measure of implicit relative preferences: A first study. *The Psychological Record*, 59, 621-640.
- Remue, J., De Houwer, J., Barnes-Holmes, D., Vanderhasselt, A., & De Raedt, R. (2013). Self-esteem revisited: Performance on implicit relational assessment procedure as a measure of self-versus ideal self-related cognitions in dysphoria. *Cognition and Emotion*, 27(8), 1441-1449.
- Robins, R. W., Hendin, H. M., & Trezenieski, K. H. (2001). Measuring global self-esteem: Construct validation of a single-item measure and the Rosenberg Self-Esteem Scale. *Personality and Social Psychology Bulletin*, 27(2), 151-161.
- Rosenberg, M. (1965). *Society and the adolescent self-image*. Princeton, NJ: Princeton University Press.
- Rothermund, K., Teige-Mocigemba, S., Gast, A., & Wentura, D. (2009). Eliminating the influence of recoding in the Implicit Association Test: The Recoding-Free Implicit Association Test (IAT-RF). *Quarterly Journal of Experimental Psychology*, 62 (1), 84-98.
- Rubinstein, J. S., Meyer, D. E., & Evans, J. E. (2001). Executive control of cognitive processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance*, 27(4), 763.
- Scanlon, G., McEnteggart, C., Barnes-Holmes, Y., & Barnes-Holmes, D. (2014). Using the Implicit Relational Assessment Procedure (IRAP) to assess implicit gender bias and self-esteem in typically-developing children and children and children with ADHD and with dyslexia. *Behavioral Development Bulletin*, 19 (2), 48-58.

- Sohn, M. H., & Anderson, J. R. (2001). Task preparation and task repetition: two-component model of task switching. *Journal of Experimental Psychology: General*, 130(4), 764.
- Szymanski, J. & O'Donohue, W. (1995). Fear of spiders questionnaire. *Journal of Behavior Therapy and Experimental Psychiatry*, 26(1), 31-34.
- Teachman, B. A. & Saporito, J. (2009). I am going to gag: Disgust cognitions in spider and blood-injury-injection fears. *Cognition and Emotion*, 23, 399-414.
- Vahey, N. A., Barnes-Holmes, D., Barnes-Holmes, Y., & Stewart, I. (2009). A first test of the implicit relational assessment procedure as a measure of self-esteem: Irish prisoner groups and university students. *The Psychological Record*, 59, 371-388.
- Vahey, N. A., Nicholson, E., & Barnes-Holmes, D. (2015). A meta-analysis of criterion effects for the Implicit Relational Assessment Procedure (IRAP) in the clinical domain. *Journal of Behavior Therapy and Experimental Psychiatry*, 48, 59-65.
- Wittenbrink, B., Judd, C. M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology*, 72(2), 262-274.

Appendices

Appendix 1: National Adult Reading Test

FOR EXPERIMENTER'S USE

NART pronunciation and definitions

Word	Say	Definition
Ache	<i>Rhymes with take</i>	Any dull, continuous pain
Debt	Det	Anything which one owes to another
Psalm	Sahm	A sacred song or hymn
Depot	Deppo (or deepo)	A place where things are kept or stored
Chord	Kord	<ol style="list-style-type: none"> 1. <i>Maths</i>: a straight line segment joining two points on a curve. 2. a string on a musical instrument 3. <i>Music</i>: a group of three or more notes played together in harmony
Bouquet	Bo- kay or bookay	<ol style="list-style-type: none"> 1. a bunch of flowers 2. the characteristic smell of wines or liqueurs
Deny	De- nigh	<ol style="list-style-type: none"> 1. to declare as untrue 2. to refuse to believe or acknowledge 3. to refuse to grant
Capon	Kay -pon	A domestic cock which has been castrated to improve its flesh for eating
Heir	Air	<ol style="list-style-type: none"> 1. a person who inherits, or will inherit, money, property, title, etc. 2. a person, group or society to which something such as tradition, ideas, etc. Is passed on
Aisle	Ile	Any passage between blocks of seats, as in a theatre
Subtle	Sutt 'l	Fine, slight or delicate, so as to be difficult to detect, etc.
Nausea	Nawsia	<ol style="list-style-type: none"> 1. a feeling of sickness in the stomach, often followed by vomiting 2. a feeling of extreme disgust or loathing
Equivocal	Ikkwivvi -k'l	Ambiguous or unclear
Naïve	Nie- eev	Unaffected or unsophisticatedly simple and artless (free from deceit or cunning)
Thyme	Time	A low shrub with fragrant leaves used in cooking
Courteous	Kertius	Polite and well-mannered
Gaoled	Jaled	Also spelt jail : a building where convicted criminals are kept
Procreate	Pro -kree-ate	To produce offspring
Quadruped	Kwodroo -ped	Any animal with four feet
Catacomb	Katta -koom or Katta -kome	(usually plural) an underground cemetery consisting of tunnels with recesses for graves
Superfluous	Soo- per floo-us	More than is needed
Radix	Ray -diks	<i>Maths</i> : a number used as the base of a system of numbers, logarithms, etc.
Assignate	Ass -ignate	Arrnage to meet in private, arrange rendezvous

Ache	_____	Simile	_____
Aeon	_____	Debt	_____
Cellist	_____	Psalm	_____
Zealot	_____	Depot	_____
Abstemious	_____	Chord	_____
Gouge	_____	Bouquet	_____
Placebo	_____	Deny	_____
Façade	_____	Capon	_____
Aver	_____	Heir	_____
Leviathan	_____	Aisle	_____
Chagrin	_____	Subtle	_____
Détente	_____	Nausea	_____
Gauche	_____	Equivocal	_____
Drachm	_____	Naïve	_____
Idyll	_____	Thyme	_____
Beatify	_____	Courteous	_____
Banal	_____	Gaoled	_____
Sidereal	_____	Procreate	_____
Puerperal	_____	Quadruped	_____
Topiary	_____	Catacomb	_____
Demesne	_____	Superfluous	_____
Labile	_____	Radix	_____
Phlegm	_____	Assignate	_____
Syncope	_____	Gist	_____
Prelate	_____	Hiatus	_____

Appendix 2: Participant Consent Form

Consent Form

Participant

I consent to participate in an experimental psychology study being run by Colin Harte, a postgraduate student at Maynooth University. I understand and consent to the following:

- The experiment will last a maximum time of one hour.
- I understand that the experiment has three parts including a questionnaire, a computer task, and a behavioural approach task. I understand that I may not have to complete all parts of the experiment.
- I am free to terminate my participation in the study at any time and may withdraw the data obtained from my participation, if I so wish.
- I understand that I participate under my own volition and that my participation will not have any effect on my subsequent academic results. I also understand that no monetary remuneration will result from participation.
- I understand that the data collected will be safeguarded in a code protected computer system, and any raw data will be locked in a cabinet, for a period of ten years, after which it will be destroyed. Up until publication of the research project, I may have access to the data collected.
- I understand that the data will be combined, analysed, and may be presented at International Conferences, or submitted to international journals for publication.
- I confirm that I am not under the age of 18.
- I consent to the retention and further use of the data collected from my participation in this study for possible use in future research and analyses (It should be noted that keys to identifying individual participant data will be destroyed following the publication of the research project leaving only fully anonymised data): Yes___ No___

I have received this information in an understandable way. All my questions have been answered.

Please print and sign your name below if you are willing to abide fully by the conditions stated above.

Name: _____ **(Please print in block capitals)**

Signature: _____

Date: _____

Experimenter

I, Colin Harte, can confirm that all the necessary safety precautions have been taken.

Signature of experimenter: _____

Date: _____

Contact Details

Researcher: Colin Harte, Department of Psychology, 2nd Floor, John Hume Building, Maynooth University, Maynooth, Co. Kildare

Tel: +353 85 133 9736

Email: colin.harte.2012@mumail.ie

Supervisor: Dr. Yvonne Barnes-Holmes, Department of Psychology, 2nd Floor, John Hume Building, Maynooth University, Maynooth, Co. Kildare.

Tel: +3531 708 6080

Email: yvonne.barnes-holmes@nuim.ie

If during your participation in this study you feel the information and guidelines that you were given have been neglected or disregarded in any way, or if you are unhappy about the process, please contact the Secretary of the National University of Ireland Maynooth Ethics Committee at research.ethics@nuim.ie or +353 (0)1 708 6019. Please be assured that your concerns will be dealt with in a sensitive manner.

Appendix 3: WAIS Full Scale, Verbal and Performance IQ predicted from number of NART errors

NART Errors	Predicted Full Scale IQ	Predicted Verbal IQ	Predicted Performance IQ
0	131	127	128
1	129	126	127
2	128	125	126
3	127	124	125
4	126	123	123
5	124	122	122
6	123	121	121
7	122	119	120
8	121	118	119
9	120	117	118
10	118	116	117
11	117	115	116
12	116	114	115
13	115	113	114
14	113	111	112
15	112	110	111
16	111	109	110
17	110	108	109
18	108	107	108
19	107	106	107
20	106	105	106
21	105	103	105
22	103	102	104
23	102	101	102
24	101	100	101

NART Errors	Predicted Full Scale IQ	Predicted Verbal IQ	Predicted Performance IQ
25	100	99	100
26	98	98	99
27	97	97	98
28	96	95	97
29	95	94	96
30	94	93	95
31	92	92	94
32	91	91	93
33	90	90	91
34	89	89	90
35	87	87	89
36	86	86	88
37	85	85	87
38	84	84	86
39	82	83	85
40	81	82	84
41	80	81	83
42	79	80	82
43	77	78	80
44	76	77	79
45	75	76	78
46	74	75	77
47	73	74	76
48	71	73	75
49	70	72	74
50	69	70	73

Appendix 4: Depression Anxiety and Stress Scales

DASS 21

Please read each statement and circle a number 0, 1, 2 or 3 which indicates how much the statement applied to you **over the past week**

0	1	2	3		
Did not apply to me at all	Applied to me to some degree, or some of the time	Applied to me to a considerable degree, or a good part of time	Applied to me very much, or most of the time		
1	I found it hard to wind down	0	1	2	3
2	I was aware of dryness of my mouth	0	1	2	3
3	I couldn't seem to experience any positive feeling at all	0	1	2	3
4	I experienced breathing difficulty (eg, excessively rapid breathing, breathlessness in the absence of physical exertion)	0	1	2	3
5	I found it difficult to work up the initiative to do things	0	1	2	3
6	I tended to over-react to situations	0	1	2	3
7	I experienced trembling (eg, in the hands)	0	1	2	3
8	I felt that I was using a lot of nervous energy	0	1	2	3
9	I was worried about situations in which I might panic and make a fool of myself	0	1	2	3
10	I felt that I had nothing to look forward to	0	1	2	3
11	I found myself getting agitated	0	1	2	3
12	I found it difficult to relax	0	1	2	3
13	I felt down-hearted and blue	0	1	2	3
14	I was intolerant of anything that kept me from getting on with what I was doing	0	1	2	3
15	I felt I was close to panic	0	1	2	3
16	I was unable to become enthusiastic about anything	0	1	2	3
17	I felt I wasn't worth much as a person	0	1	2	3
18	I felt that I was rather touchy	0	1	2	3
19	I was aware of the action of my heart in the absence of physical exertion (eg, sense of heart rate increase, heart missing a beat)	0	1	2	3
20	I felt scared without any good reason	0	1	2	3
21	I felt that life was meaningless	0	1	2	3

Appendix 5: Rosenberg Self Esteem Scale

RSES

Below is a list of statements dealing with your general feelings about yourself. Please rate how much you agree with each statement by circling a number next to it. Use the scale below to make your choice.

Strongly Agree	Agree	Disagree	Strongly Disagree
1	2	3	4

1.	I feel that I am a person of worth, at least on an equal plane with others.	1	2	3	4
2.	I feel that I have a number of good qualities.	1	2	3	4
3.	All in all, I am inclined to feel that I am a failure.	1	2	3	4
4.	I am able to do things as well as most other people.	1	2	3	4
5.	I feel I do not have much to be proud of.	1	2	3	4
6.	I take a positive attitude toward myself.	1	2	3	4
7.	On the whole, I am satisfied with myself.	1	2	3	4
8.	I wish I could have more respect for myself.	1	2	3	4
9.	I certainly feel useless at times.	1	2	3	4
10.	At times I think I am no good at all.	1	2	3	4

Appendix 6: Fear of Spiders Questionnaire

Please rate the following statements from 1 (strongly disagree) to 7 (strongly agree)

1. If I came across a spider now, I would get help from someone else to remove it.
2. Currently, I am sometimes on the lookout for spiders.
3. If I saw a spider now, I would think it will harm me.
4. I now think a lot about spiders.
5. I would be somewhat afraid to enter a room now, where I have seen a spider before.
6. I now would do anything to try to avoid a spider.
7. Currently, I sometimes think about getting bit by a spider.
8. If I encountered a spider now, I wouldn't be able to deal effectively with it.
9. If I encountered a spider now, it would take a long time to get it out of my mind.
10. If I came across a spider now, I would leave the room.
11. If I saw a spider now, I would think it will try to jump on me.
12. If I saw a spider now, I would ask someone else to kill it.
13. If I encountered a spider now, I would have images of it trying to get me.
14. If I saw a spider now I would be afraid of it.
15. If I saw a spider now, I would feel very panicky.
16. Spiders are one of my worst fears.
17. I would feel very nervous if I saw a spider now.
18. If I saw a spider now I would probably break out in a sweat and my heart would beat faster.

Seven-point Likert-scale

1 – Strongly disagree

2 – Disagree

3 – Somewhat disagree

4 – Neither agree nor disagree

5 – Somewhat agree

6 – Agree

7 – Strongly agree

Appendix 7: Behavioural Approach Task Scoring Sheet

Prompt

The following task is designed to assess how willing you are to approach a spider. I am going to ask you if you are willing to complete a number of tasks and if you are willing, I will then ask you to complete this task.

There is a spider in a container on the table in the next room. Would you be.....

Scoring

Enter the individuals BAT score below, where:

0 = completed no steps

1 = willing to open the door and look

2 = willing to enter the room

3 = willing to touch the box

4 = willing to open the box

5 = willing to touch the spider for 10 seconds

Appendix 8: Participant Debrief (Experiments 3a and 3b)

While participants will not receive a physical debriefing sheet, the following forms the information that each participant was debriefed with and the procedure for doing so.

Following the experiment, participants will be thanked for their contribution to the study. They will be told that the current study is investigating two procedural parameters of the computer based task they just completed (the IRAP). Specifically they will be told that it was investigating whether effects generated from the task are produced by the rate at which they switch between rules of responding (that is, the amount of times they respond true and false before being asked to change rule), and also whether effects are possibly mediated by the amount of blocks that they complete. They will be told that the use of spiders and nature stimuli are not directly relevant to the study but were simply used as a proxy for the methodological questions at hand due to its previous successful use in predicting behaviour and accurately measuring spider fear. They will be told that the spider involved in the final part of the study was a molt and was not alive. Any questions that they have regarding the nature of the study will be addressed again and they will be reminded that they are free to withdraw their data until publication has been reached.

Regardless of whether they experience any after effects or stress following the study, contact details for the counselling service in the University will be provided.