Developing an Empirically Valid Function Acquisition Speed Test for Assessing Attitudes to and Predicting Real-world Behaviour

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

Isabella J. Lalor, B.A. (Hons)

October 2019

Research Supervisor: Dr. Bryan Roche

Head of Department: Prof. Andrew Coogan

Contents

List of Tables

List of Figures

Declaration

I, the undersigned, hereby certify that this material, which I now submit in fulfilment of a M.Sc. degree, has not been previously submitted as an exercise for a degree at this or any other University, and is, unless otherwise stated, entirely my own work.

Signed: ________________________________

Isabella Jardim Lalor

Student ID: 14387606

Date: 24th October 2019

Acknowledgements

This work would not have been possible without my supervisor Dr. Bryan Roche, who selflessly devoted time and energy towards helping me accomplish this thesis. I am eternally grateful for your guidance, patience, and willingness to provide support whenever needed.

I want to extend my sincere gratitude to Jamie Cummins, who selflessly took the time to express interest in my work, helping me make sense of my data and giving me advice throughout. This thesis would not be to the standard that it is without you, thank you.

I would also like to thank everyone in the Psychology department at NUI, Maynooth, for always being there to help in any way they can. I thank my fellow postgraduates, who have motivated me and given me advice along the way. In particular I would like to thank Alex, I would not have gotten through this past year without your countless messages of motivational support when needed the most.

Finally, I wish to thank my family and friends for their endless support. To Darragh, words cannot fully express how grateful I am for all your emotional support throughout. To my parents, Marise and Joseph Lalor, thank you for your love, patience, and encouragement, not only during this thesis, but in everything I have accomplished and everything I wish to pursue. Your love and support has most definitely not been taken for granted.

Abstract

The current research was designed to validate a newly developed Function Acquisition Speed Test (FAST; O'Reilly et al., 2012) developed to measure attitudes to abortion among the general population, and to retrospectively predict a real-world behaviour in relation to voting for or against the availability of abortion services, in a national referendum. Participants, consisting of those with openly expressed pro-life or pro-choice attitudes, were exposed to an Implicit Association Test (IAT), a FAST procedure, an explicit attitudes to abortion questionnaire, and a demographic survey. This procedure allowed for a comparison of FAST and IAT outcomes, inter-correlations between the explicit and implicit measures, as well as the predictive validity of the implicit measures of referendum voting behaviour and self-reported group affiliation. The first aim of this research was to assess the FAST's utility in providing converging test results to predict both group affiliation and voting behaviour in the 2018 referendum on the 8^{th} amendment concerning the legalisation of abortion. In seeking to validate the FAST, the gold standard IAT, as well as explicit test measures (i.e., the Attitudes to Abortion Questionnaire and Demographic survey), were administrated to assess convergent and predictive validity. The second aim of this research was to meaningfully compare the two implicit test measures, by employing common data scoring methods across the tests (i.e. using Rate-Fluency Differential or RFD scores, D scores and slope differential scores). Results showed that the FAST predicted group affiliation and voting behaviour satisfactorily. In particular, FAST RFD scores were found to be the optimal scoring metric when predicting voting behaviour, whereas FAST slope scores were found to be the optimal scoring metric when predicting group affiliation. Cut-off points were also provided for both scoring metrics, contributing to the production of a meaningful scoring system for the FAST. In terms of the implicit measures, the IAT D scores was found to have the best overall predictive validity. However, neither test outperformed the almost perfectly

predictive explicit test score, regardless of the scoring metric used. Nevertheless, the FAST offers a functional model of implicit attitudes that is wanting for the IAT, and additional research is needed in sensitive contexts, in order to increase the FAST's predictive validity up to and beyond that of explicit measures. Several methodological and conceptual issues are considered in an effort to accelerate fruitful research towards this end.

Chapter 1

General Introduction

1.1 Introduction

The current research seeks to contribute to the development of the Function Acquisition Speed Test (FAST) as a potentially reliable and valid implicit measure of stimulus relatedness, by employing a known group's paradigm approach to retrospectively predict real word behaviour using FAST scores. At present, there are two main schools of thought regarding the attitude concept and how it can be assessed to predict behaviour: the Social Cognition approach, and the Behaviour Analytic approach. The social-cognitive approach is the main stream approach, and is the most well-known approach in terms of developing alternative psychometric tests that potentially overcome the issues associated with explicit test measures. Critical analysis of its theoretical foundations and current popular forms of measurement, such as the Implicit Association Test (IAT) will be discussed. This approach will be contrasted throughout with the behaviour-analytic approach. A review of the relevant literature in Behaviour Analysis will be presented, as will one widely cited behaviour-analytic implicit test; the Implicit Relational Assessment Procedure (IRAP). This will lay the groundwork for outlining the origin and modus operandi of the Function Acquisition Speed Test. Abortion attitudes will be discussed, particularly in terms of how suitable they may be for assessment using implicit measures, and how this domain of attitude assessment may be a suitable one within which to assess the utility of an emerging implicit test. Finally, the current research programme to assess a FAST for distinguishing between known groups and predicting voting behaviour in relation to the legalisation of abortion services, will be outlined.

1.2 Explicit Test Approach

For decades, early psychological research was forced to rely primarily on self-report or explicit measures to assess participant attitudes or measure any given aspect of an individual's personal history, using surveys, questionnaires, and interviews (De Houwer, 2003; Gawronski, 2009). Ample evidence suggested that explicit measures satisfied important psychometric criteria such as usefulness and efficiency, where people's attitudes, beliefs, and personality characteristics were measured by simply asking them about their thoughts, feelings, and behaviours (Mischel, 1968). While pioneering scholars had formally presumed that an attitude could be accurately assessed using only a large set of direct questions that were selected via an elaborate procedure (e.g. Likert, 1932), this approach to attitude assessment is now considered questionable due to factors such as the introspection problem and social desirability bias (Jacoby, Lindsay, & Toth, 1992).

It was formerly assumed generally that people were aware of their attitudes, beliefs, and values that guided their behaviour, and that they were willing to reveal them if asked appropriately (Kihlstrom, 2004). However, these assumptions are not always valid, as people are sometimes unable or unwilling to reveal their opinions (Greenwald & Banaji, 1995). Explicit measures typically reference a target object in the participant's personal history (Jacoby, Lindsay, & Toth, 1992). Thus, they assume that the participant has already formed an opinion or is able to construct one on the spot, is aware of his or her attitude, and lastly, is willing to share it accurately with the researcher (Schwarz & Bohner, 2001). However, previously formed attitudes may not be easily accessible to the individual (Fazio, 1986; Fazio, Sanbonmatsu, Powell, & Kardes, 1986). From a behavioural point of view, this is because the individual cannot access the contingencies that would control their behaviour in the hypothetical scenarios often asked about in questionnaires. Thus, although individuals have a previously formed opinion, according to the social cognitive approach they report a newly created one. Even if participants can identify a prior experience, they may be unaware of its influence, hence providing inaccurate reports (Greenwald & Banaji, 1995). This limitation associated with self-report measures is known as the introspection problem, where individuals may not be consciously aware of the particular attitude being assessed by the researcher, and is therefore introspectively inaccessible (Cameron, Brown-Iannuzzi, & Payne, 2012). That is, participants may not be able to reflect on and express an attitude accurately on a self-report measure because they are unaware of their implicit attitudes (Dambrun & Guimond, 2004).

Often, participants are aware of their own attitudes or opinions, but are reluctant to share them with unfamiliar researchers. This is usually true for questions that are socially sensitive, where responses might be embarrassing or where people feel compelled to produce a socially acceptable answer (Kihlstrom, 2004), leading respondents to purposefully distort their answers for social-desirability or image-management reasons. For instance, the most blatant forms of prejudice and discrimination have decreased in many countries (Payne et al., 2010), where opinion polls indicate that fewer people express prejudiced attitudes towards particular groups than was the case decades ago. However, although fewer people express prejudiced attitudes, this does not necessarily mean people are actually less prejudice. Instead many people may intentionally hide their prejudices for fear of social consequences, expressing them only when they feel it is safe or socially appropriate (Holtgraves, 2004). Therefore, explicit measures assessing attitudes are prone to errors that greatly lower validity and reliability. As such, assessing socially sensitive topics such as race, gender and sexuality are more difficult to detect solely using explicit measures (Cullen & Barnes-Holmes, 2008).

Social desirability bias occurs when participants attempt to answer self-report questions in a more socially desirable way, which would appear more favourable to the researcher, as opposed to reporting their true beliefs or attitudes (Paulhus, 1984). In particular, research shows that when participants are given a standard self-report referring to socially sensitive information, results can be significantly distorted. This has been found in studies relating to sexuality and sexual practices (e.g. Fenton, Johnson, McManus, & Erens, 2001), mental health (e.g. Klassen, Anderson, & Hornstra, 1975), and regarding socially-favourable activities such as voting (Holbrook & Krosnick, 2010), where participants intentionally distorted their responses. For instance, in polls in Ireland in 1992, Sinn Féin, an extreme Republican party, received more votes in the general election than public opinion polls had predicted, likely due to socially desirable responding (Breen, 2000). Distortions are especially likely when an accurate response is seen as either violating social norms such as politeness or prejudice (Demo, 1985; Dovidio & Fazio, 1992), jeopardizing one's self-image (Dovidio & Fazio; Gaertner & Dovidio, 1986), or going against the stereotypical answer. In brief, no explicit measure can truly avoid the influence of respondents' control. This, in essence, is the main impetus behind the development of "implicit" measures.

1.3 Implicit Test Approach

Before the implicit test approach is outlined it is worth considering what it is meant by the term "implicit", at least within the social cognitive literature (see section 1.5 for a behavioural approach). Greenwald and Banaji (1995) note that attitudes, in addition to their conscious manifestations, might also operate in an indirect, unconscious, or implicit mode. Such implicit attitudes are activated automatically, not necessarily requiring conscious thought or attention. By nature, social cognitions and behaviours are implicit, unconscious and indirect, and are thusly inaccessible to traditional explicit measures. Although people may control explicit prejudices, or may honestly believe that they are not prejudiced, implicit prejudices may still guide their behaviour (De Houwer, Teige-Mocigemba, Spruyt, & Moors, 2009; Fazio, Jackson, Dunton, & Williams, 1995). Therefore, to overcome such complications with explicit measures, researchers have begun studying social biases using implicit measures, and research has shown that such measures produce results that often diverge from explicit self-report measures and appear to be more theoretically meaningful (e.g., they find biases among "known groups", who openly express sensitive and normally concealed attitudes; Fazio & Olson, 2003). A substantial portion of attitudes are shaped by cognitive processes that are outside conscious awareness and have influences that participants do not realize (Greenwald & Banaji, 1995). Procedures that measure these constructs at the implicit level are said to provide quantifiable insights into the underlying automatic associations that participants make. Recently, psychologists have shown renewed interest in this issue and have reclaimed the importance of the unconscious (Greenwald, 1992; Weinberger, 2000).

It has been suggested that an attitudinal construct, that is implicit attitudes, circumvents the problem of social desirability responding and introspective (Greenwald & Banaji, 1995). Implicit attitudes are described as automatic evaluations which influence an individual’s behaviour, but which are not necessarily accessible to that individual (De Houwer et al., 2009). The development of implicit measures was influenced by research carried out by Donders in the mid-19th Century, where he reported that simplistic mental tasks or computations lead to quicker response times and fewer errors by participants than that of more difficult or complex mental tasks (Lane, Banaji, Nosek, & Greenwald, 2007; Donders, 1868; 1969). This led to the widely influential idea that an individual’s inner state could potentially be inferred on the basis of their performance on some experimental task. This was demonstrated for the first time in a psychometric methodology, famously known as the Stroop paradigm (Stroop, 1935). Findings from Stroop experiments show that reaction times are consistently faster for the consistent (when the colour of the word matches the typed word) versus inconsistent (when the colour of the word does not match the typed word) tasks (Stroop, 1935). What the paradigm consistently shows is that, the automatic reading of a word once seen, interferes with the required naming of the ink colour and thus leads to the increased time required to respond in the inconsistent colour-word conditions. This test format laid the groundwork for creating a test format for use in psychological research that could potentially overcome the limitations of explicit measures (Rothermund & Wentura, 2004).

Drawing upon research in implicit memory, Greenwald and Banaji (1995) proposed a new theory of social behaviour they referred to as ‘Implicit Social Cognition’. It refers to how implicit cognitions are involved with, and how they interfere with, deliberate judgements. In other words, while attitudes are generally assumed to reflect cognitive associations between a concept and an attribute, implicit attitudes are proposed to reflect the associations that are outside of direct conscious awareness. In light of this theory, Greenwald & Banaji (1995) highlighted the need for a reliable laboratory measure that enables the efficient assessment of individual differences in implicit social cognition, or implicit attitudes. A number of such measures have emerged from within the field of Social Cognition that claim to assess implicit cognitive associations between concepts and thus underlying implicit attitudes. Most often, measures of implicit attitudes assess participant’s reaction times at cognitive tasks (Greenwald, McGhee, & Schwartz, 1998; Nosek, Hawkins, & Frazier, 2011). The most popular of these ‘associative’ implicit tests is the Implicit Association Test (Greenwald, McGhee & Schwartz, 1998).

1.4 A Social Cognitive Approach to the Attitude Concept

An attitude can be described as a positive or negative evaluative reaction towards a stimulus, such as a person, action, situation, object, or concept (Tesser & Shaffer, 1990). Our attitudes help define our identity, guide our actions, and influence how we judge others (Fazio & Roskos-Ewoldsen, 2005). Social psychologists have universally defined attitudes as individual cognitive structures; a focus on attitudes as cognitive representations that are acquired and possessed by individuals and which, to a great extent, are a part of human individuality (Bohner & Wänke, 2002; Fazio & Olson, 2003). For example, Gordon Allport wrote that attitudes are “a mental and neural state of readiness to respond, organized through experience, exerting a directive and/or dynamic influence upon the individual’s responses to all objects and situations” (Allport, 1935, pp. 810). Generally in the field of social cognition, an attitude can be summarised as an evaluative cognitive association (i.e. between a concept and an attribute) that has been established in the experiential history of an individual towards any given element of their environment, which will predispose a favourable or unfavourable behavioural response toward that element (Olsen & Zanna, 1993). Attitude researchers, therefore, focus more on psychological processes and structures of individuals than on the wider social context of attitude development and expression (Eagly & Chaiken, 1993, 2005; Prislin & Wood, 2005).

Although formal definitions of attitudes vary, most contemporary social psychologists agree that the characteristic attribute of an attitude is its evaluative (pro – con, pleasant – unpleasant) nature (Edwards, 1957; Eagly & Chaiken 1993). This view is strengthened by the fact that standard attitude scaling techniques result in a score that locates an individual on an evaluative dimension regarding the target attitude object (Green, 1954; Fishbein & Ajzen, 1975). Attitude measurements are frequently used by political scientists, sociologists, economists, and other academics to assess responses to social change, events and products, and to predict future behaviour (Bohner & Wänke, 2002). Attitudes are thought to be open to conscious inspection, although their expression often depends on their social desirability. In other words, traditional models assume that if you want to know someone’s beliefs, feelings, and behavioural tendencies toward an object, all you need to do is measure his or her attitude, provided that you are not investigating a socially sensitive area (Karpinski & Hilton, 2001). This is because self-reported attitudes toward socially sensitive topics (e.g. stereotypes) are particularly susceptible to social desirability (Hofmann, Gawronski, Gschwendner, Le, & Schmitt, 2005).

Social psychologists routinely measure attitudes, access how they change and attempt to understand their impact on overt behaviour (Krosnick, Judd, & Wittenbrink, 2005). In fact, the field of social psychology was originally defined as the scientific study of attitudes (Watson, 1925) because it was assumed that attitudes were the key to understanding human behaviour. Despite their general perceived utility, attitudes turned out to not correlate very well with future behaviour, leading to what was called the crisis of confidence in social psychology in the 1970s. Attitudes were usually found to be very poor predictors of actual behaviour, and many social psychologists began to worry about the utility of the attitude construct (e.g., Blumer, 1955; Campbell, 1963; Deutscher, 1966; Festinger, 1964). For example, LaPiere (1934) toured the USA with a young Chinese couple, stopping at 251 restaurants, hotels and other establishments. At the time, prejudice against Asians was widespread, yet the couple were refused service only once, and LaPiere judged their treatment to be above average in 40% of the restaurants visited. Later, LaPiere wrote to all of these establishments, asking if they would provide service to Chinese patrons. More than 90% of those who responded stated that they would not. Although the people who expressed negative attitudes in the survey were not guaranteed to be the same ones who, months earlier, had served the Chinese couple, the discrepancy between prejudicial attitudes and non-discriminatory behaviour seemed overwhelming and called the assumption of attitude-behaviour consistency into question.

Much of the impetus for the theory of reasoned action (TRA; e.g., Ajzen & Fishbein, 1980) was Fishbein’s realization that there was a problem with attitude measurement at the time. The TRA approach assumes that people’s behaviour follows reasonably from their beliefs, attitudes, and intentions. That is, attitudes determine subjective norms. According to TRA, our attitudes toward a particular behaviour are influenced by a combination of two related factors: our beliefs about the *outcome* of the behaviour (i.e., is the outcome likely or unlikely) and our *evaluation* of the potential outcome (i.e., is the outcome a good thing or a bad thing; Fishbein & Ajzen, 1975). One implication of this theoretical assumption is that researchers should be able to obtain strong correlations between attitudes and behavioural intentions (Fishbein & Ajzen, 2010). But this assumes, among other things, that a check mark on an attitude scale really indicates a person's attitude, and that a check mark on an intention scale really indicates a person's intention (Earp & Trafimow, 2015). The TRA has nothing to say about whether check marks on scales indicate attitudes or intentions; these are assumptions that are peripheral to the basic theory. Strong attitude-behaviour (or at least attitude-intention) correlations became the rule rather than the exception. They are known as *auxiliary assumptions* that researchers use to connect non-observational terms such as "attitude" and "intention" to observable phenomena such as check marks (Earp & Trafimow, 2015). Fishbein and Ajzen (1975) recognized this and took great pains to spell out, as well as possible, the auxiliary assumptions that best aid in measuring theoretically relevant variables (Ajzen & Fishbein, 1980). This is still a continuation of the social cognitive research agenda developing more non-functional accounts in an effort to save a non-functional idea that is attitudes.

The theoretical connection between attitudes and behaviours turned out to be correct (as far as we know) with the problem having been caused by incorrect auxiliary assumptions pertaining to attitude measurement. Decades of better controlled research indicate that attitudes can indeed predict behaviour (Fazio & Roskos-Ewoldsen, 2005). Three factors help explain why the attitude-behaviour relationship is strong in some cases but weak in others. First, attitudes influence behaviour more strongly when situational factors that contradict our attitudes are weak (Holt et al., 2015). For example, conformity pressures may lead us to behave in ways that are at odds with our inner convictions. According to the theory of planned behaviour and similar models (Ajzen, 1991), our intention to engage in a behaviour is strongest when we have a positive attitude towards that behaviour, when subjective norms (i.e., our perceptions of what other people think we should do) support our attitudes, and when we believe that the behaviour is under our control. Researchers have used this theory to predict successfully whether people will become smokers, exercise regularly, drive safely, donate blood and perform many other behaviours (Victoir, Eertmans, den Bergh, & den Broucke, 2005). Second, attitudes have a greater influence on behaviour when we are aware of them and when they are strongly held (Holt et al., 2015). Sometimes we seem to act without thinking, out of impulse or habit. Attitude-behaviour consistency increases when people consciously think about or are reminded of their attitudes before acting (White, Hogg, & Terry, 2002). Lastly, general attitudes best predict general classes of behaviour, and specific attitudes best predict specific behaviours. For example, Fishbein and Ajzen (1974) found almost no relation between people's general attitudes towards religion and 70 specific religious behaviours (such as frequency of attending services). But when they combined the 70 specific behaviours into a single global index of religious behaviour, the relation between general religious attitudes and overall religious behaviour was substantial.

From the behavioural perspective the lack of correlation between attitude and behaviour was not surprising, because the contingencies that control verbal behaviour are different to those that control socially embedded non-verbal action in a rich social context. There is no requirement that reports of attitudes and actions should correlate and of course there is a possibility that in a given instance an overt non-verbal action could lead to an attitude, rather than vice versa. For example, according to Festinger's (1957) theory of cognitive dissonance, people strive for consistency in their cognitions. When two or more cognitions contradict one another, the person experiences an uncomfortable state of tension, which Festinger calls *cognitive dissonance*, and becomes motivated to reduce this dissonance. The theory predicts that to reduce dissonance and restore a state of cognitive consistency, people will change one of their cognitions or add new cognitions. According to Bem's (1972) self-perception theory, we make inferences about our own attitudes by observing how we behave. In Bem's view, your attitude is not produced by a mysterious concept called *cognitive dissonance,* rather you simply observe how you have acted and infer how you must have felt to have behaved in this fashion. Both theories, however, agree that our behaviours can influence our attitudes.

Another part of the effort to save the concept of attitude appears to have been the realisation that methods could be developed to reduce social desirability. In recent years, the social cognitive approach to attitudes has gained favour because of the apparent success of new implicit testing tools, such as the Implicit Association Test (IAT), leaving more work for behavioural researchers who wish to build more functional models of such tests and procedures. The implicit test methods are interesting form a behavioural point of view because they use procedures remarkably close to the matching to sample and class formation procedures typical of research studies into stimulus equivalence and simple functional response classes. Before these linkages are outlined, and areas for cross fertilisation of ideas across the social cognitive and behavioural domains are presented, the method of the most common implicit test will be outlined.

1.4.1 Implicit Association Test

Currently, the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998), a general-purpose procedure for measuring strengths of automatic associations between concepts, is the most notable measure of implicit beliefs or attitudes. Building upon their theory of how implicit cognitive associations established in the individual's history mediate social feelings, thoughts, and actions, Greenwald and colleagues (1998) proposed a test that claims to measure the strength of socially significant cognitive associations or associative structures. The IAT is a computer-based categorization task that requires participants to respond under time pressure with differential key presses to stimuli that are presented on-screen, based on rules provided to the participant that vary across blocks (Randall, Rowe, Dong, Nock, & Colman, 2013). Many studies have now shown that most individuals will respond more rapidly and accurately when socially compatible stimuli share a response requirement, therefore revealing a particular bias (De Houwer, 2002). Is has also been suggested that this test format is relatively immune to faking efforts by the participant, even when they understand how the test works (Steffens, 2004).

In the first investigation of the IAT, a series of word pairs, such as ‘insect-bad’ and ‘flowers-good’ were flashed on a computer screen and participants were asked to press a corresponding computer key as quickly and accurately as possible, producing a reaction time (Greenwald et al., 1998). The “automatic association” between a concept (e.g., flowers) and an attribute (e.g., positive valence) was measured by observing the difference in speed between a condition in which flower names and pleasant-meaning words shared the same response key and a condition in which flower names and unpleasant-meaning words shared the same response key. The IAT assesses associations between the two target categories (e.g., flowers vs. insects) and two evaluative categories (e.g., good vs. bad) through recording response-times. Standardized reaction time scores based on the variance in response latencies across test blocks are of interest in the IAT, as it is assumed that placing participants under time pressure facilitates the capture of their implicit (or ‘automatic’) cognitions. The principle underlying the task is that participants will react more quickly when they perceive that the concepts are consistent with their attitudes than when the concepts are not. Thus, without conscious control, a participant will respond faster during a consistent block (e.g., insects-bad/ flowers-good) compared to an inconsistent block (e.g., insects-good/ flowers-bad), indicative of the participants attitudes (Greenwald & Banaji, 1995). These associative structures are believed to underlie implicit attitudes, which in turn mediate behaviours and judgements, and thus the relative strength of the associations should theoretically indicate the strength of the implicit attitude. In other words, there was a strong association between an image of an insect and the word ‘bad’ as opposed to the word ‘good’, and a weak association between an image of a flower and the word ‘bad’ as opposed to the word ‘good’. Here, the IAT is proposed to have revealed the pre-existing association that would underlie the individual’s implicit attitude towards insects and flowers (Greenwald et al., 1998).

The first investigations of the IAT (Greenwald et al., 1998) confirmed that the IAT could detect valence differences that were associated both with familiar non-social objects (flowers, musical instruments, insects, and weapons) and with significant social objects (Japanese and Korean ethnicity, and Black and White race). Greenwald et al. (1998) also demonstrated that IAT measures were stable across several procedural variations, including whether the pleasant category was assigned to a left-side or right-side response, the time interval between response to one stimulus and presentation of the next stimulus item (varied from 150 to 750ms), and whether concepts and attributes were represented by 5 or 25 items. Greenwald and colleagues (1998) suggested that the IAT may resist masking by self-preservation strategies, and may preclude the concealment of attitudes that an individual may either intentionally wish to conceal, or may not necessarily know they hold. The IAT is, therefore, one of the primary ways in which social cognitive researchers suggest circumventing the problems inherent in self-report measures. Given that it is believed to be superior to many explicit attitude tests in terms of being capable of overcoming experimental demand characteristics and precluding fake responses (Kim, 2003), the IAT has become widely utilised across many different fields of study within psychology, including forensic psychology (Brown, Grey & Snowdon, 2009), clinical psychology (Egloff & Schmukle, 2002) and health psychology (Wolff, von Hippel, Brener, & von Hippel, 2015).

1.4.2 Issues with the Implicit Association Test

Although the IAT is widely used to measure attitudes, many concerning methodological flaws have been uncovered. There have been concerns relating to the use of trial-by-trial feedback, arbitrary time penalties for error responses, score algorithms based on effect size and p values rather than experimental control, use of instructions rather than in-trial learning, and the use of arbitrary post-hoc response times rather than finite response windows as part of the procedure (O'Reilly, Roche, Ruiz, Tyndall, & Gavin, 2012). Overall, the most significant criticism of the IAT is that research into the core processes underlying its procedure is lacking (Rothermund & Wentura, 2004; Hughes & Barnes-Holmes, 2011). The IAT has also been found to have relatively low reliability (Kawakami & Dovidio, 2001). This was evident in meta-analysis conducted by Oswald and colleagues (2013) where weaker correlations between IAT scores and discriminatory behaviour were found compared with what Greenwald, Poehlman, Uhlmann, and Banaji found in a 2009 meta-analysis. Carlsson and Agerstrom (2016) suggest that there is little evidence to support that the IAT can even predict discrimination in a meta-analysis reviewing eleven IAT research studies. In effect, the IAT is poorly understood both conceptually and methodologically.

The developers of the IAT have been relatively conservative in attempting to precisely define the 'cognitive biases' the IAT is purported to measure (Greenwald & Nosek, 2008). As mentioned previously, the IAT measures a participants associations between concepts; where a score is typically given on a scale of -2.0 to 2.0, with anything above 0.65 or below -0.65 indicating a strong association. However, there is not enough research to validate that above and below that cut-off participants differ in any way based on that score (Blanton & Jaccard, 2006). There are queries surrounding what exactly the IAT technique is measuring (Fazio & Olson, 2003). Opinions on the IAT are mixed regarding its validity as a psychometric measure, and whether it actually assesses implicit attitudes or cultural knowledge (Olson & Fazio, 2004). The IAT is purported to measure relative strength of associations. However, some researchers have asserted that the IAT may instead be measuring constructs such as salience of attributes (Rothermund & Wentura, 2004) or cultural knowledge (Arkes & Tetlock, 2004). All in all, the IAT is considered noisy, in that there is no way of knowing whether it's measuring unconscious attitudes or simply associations picked up from the environment (Fazio & Olsen, 2003), and an answer to this has not been given by the developers themselves who are also unsure.

One of the major problems with the IAT is that it ultimately provides an indirect measure of beliefs as it focuses on associations rather than relations among stimuli (De Houwer, 2002). This measure is therefore controversial as implicit attitudes are relational in nature, not just associative (De Houwer, 2002).The IAT purportedly measures attitudes such that participants are not aware of what is being tested, do not have conscious access to the attitudes being tested, and/or have no control over the measurement outcome. De Houwer and colleagues (2009) criticizes the IAT for its relative lack of empirical research into those functional properties. A single study (Monteith, Voils, & Ashburn-Nardo, 2001) has suggested that many participants are aware of what a given IAT is examining, which calls into question the "implicitness" of the IAT in that sense, which has been suggested as one of the primary ways the IAT circumvents the issues in explicit measures (Greenwald et al., 1998). Some researchers have proposed that the IAT effect is heavily susceptible to faking, that is, it may be possible, even with the indirect format outlined above, to fake results on the IAT (e.g., Cvencek, Greenwald, Brown, Gray, & Snowden, 2010; Steffens, 2004). This effect is found to be particularly pronounced when participants have had previous experience with the IAT procedure (Fiedler & Bluemke, 2005).

From a behaviour-analytic approach, there are four main features of the IAT presentation format that has been scrutinised. Firstly is the imbalanced feedback, where participants are presented with negative feedback in the form of a red X on the screen only when an incorrect response is produced. This is done as a form of punishment for incorrect responses which has unknown effects on subsequent responses and response rates (O'Reilly, Roche, & Cartwright, 2015). In contrast, when a correct response is given, the participant is presented with no feedback whatsoever and instead is presented with the next trial. This imbalanced feedback method makes the emergence of IAT effects more likely. Specifically, such a procedure ensures that whichever of the two test blocks contains less errors will also be the block that involves less punishment, therefore possibly exaggerating the accuracy and reaction time differences across blocks in the expected direction. This exaggeration may occur because rapid responding in the consistent block is negatively reinforced by the removal of both interruptions to response fluency and the omission or reduction of negative feedback. In contrast, rapid responding in the inconsistent block is punished, thereby leading to slower responding. Previous research has shown that response caution during difficult IAT tasks can partially explain IAT effects (Klauer, Voss, Schmitz, & Teige-Mocigemba, 2007).

Additionally, concerns have been raised regarding the use of reaction times as an assessment criterion. Within the IAT, the response times are typically measured from the point of stimulus presentation to the end of the trial (i.e., when a correct response is made), for both the correct response and the incorrect response. However, as mentioned, feedback is only presented after an incorrect response, after which participants are required to respond correctly. Therefore, the recorded response times for incorrect trials are considerably larger than for correct trials but not because of the speed at which a participant responds under the stimulus control of the relational stimuli in a given trial, but because of the addition of an observation response time to the feedback to the overall response time recorded for that trial. This is intentional on the part of the developers of the IAT (Greenwald, McGhee, & Schwartz, 1998), and is designed to negate the need to impose an arbitrary time penalty added to recorded response times, a practice formerly used by IAT researchers to artificially increase reaction times for trials on which an error was made (Gavin, Roche, Ruiz, Hogan, & O'Reilly, 2012). However, some researchers have objected to this statistical technique, arguing that it does not accurately reflect either response rates or response time (see Gavin, Roche, & Ruiz, 2008, for a full review). In contrast, behaviour-analytic researchers have suggested that response accuracy or indeed fluency (combined accuracy and speed) rather than speed alone might be superior as an assessment of task difficulty within implicit measures (Gavin et al., 2012; O'Reilly et al., 2012). A move to fluency as the primary measure over response time also circumvents several problems regarding the manner in which response latencies (i.e., the time between stimulus presentation and the participant emitting a correct response) are calculated using the IAT (Ridgeway, Roche, Gavin, & Ruiz, 2010).

Lastly, the IAT relies on a method in which the raw IAT scores (i.e., participant error rates and reaction times) are transformed using what is known as the D-algorithm (Greenwald, Nosek & Banaji, 2003). This popular scoring technique involves a standardised reaction time measure (taken across a fixed trial block) as the core index of association strength. That is, by subtracting the mean response latency of the 4^{th} IAT test block from that of the 7^{th} IAT test block (or vice versa, depending on which is the 'consistent' block) and dividing this value by the standard deviation of response times across both critical blocks. It should be noted that the use of this d-score and test format is based primarily on a single paper from Greenwald, Nosek, and Banaji (2003), which compared the effectiveness of 10 different metrics on 7 different criteria (correlation with explicit measures, correlation with average latency, internal consistency, order effect correction, IAT experience correction, effect size, and implicit-explicit path in a confirmatory factor analysis; Richetin, Costantini, Perugini, & Schönbrodt, 2015). However, the trajectories of reaction times across trials are not assessed and response rates and changes in rate are usually irrelevant to the measures (Blanton & Jaccard, 2006). In line with behaviour analytic tradition, learning rates to predetermined criteria may be the better primary behavioural measure over standardized latency scores.

In a functional, Relational Frame Theory (RFT) driven approach, the IAT is viewed as a measure of an individual's verbal history and practices, by assessing their ability to form stimulus relations under time constraints, which may or may not in turn reflect personal attitudes or affective states and dispositions (O'Reilly, Roche, & Cartwright, 2015). IAT effects are conceived in terms of participants' fluency with the relevant verbal categories and their degree of experience at juxtaposing members of those verbal categories. For instance, an individual who has many dealings with people of a specific race, and has encountered both pleasant and unpleasant individuals from this racial group, will likely find it easy to juxtapose racial and evaluative terms in an IAT according to the test rules across the two test blocks. Such an individual will show no IAT effect (i.e., response time or accuracy differential across the text blocks). On the other hand, if they have experienced mostly unpleasant individuals from one racial group or other, the juxtaposition of response rules across the IAT blocks will likely expose a fluency differential across those two blocks (i.e., an IAT effect). This is the behavioural model of the IAT (see Roche, Ruiz, O'Riordan, & Hand, 2005), which later was developed into its own measure known as the Function Acquisition Speed Test (FAST; O'Reilly, Roche, Ruiz, Tyndall, & Gavin, 2012). In this sense, the behavioural phenomenon captured by the IAT is of interest. Indeed, behaviour-analytic researchers have attempted to account for the effects noted on implicit measures whilst simultaneously avoiding the issues of social-cognitive theory. Two main behaviour-analytic measures have been developed in this regard; the Implicit Relational Assessment Procedure and the Function Acquisition

Speed Test. Before these alternative test measures are outlined, the behavioural approach to the attitude concept will be discussed.

1.5 A Behavioural Approach to the Attitude Concept

The general social cognitive approach to the conceptualisation of attitudes is challenged by a behavioural account of the IAT, which would identify such social cognitive definitions as mentalistic. The Relational Frame Theory (RFT: Hayes, Barnes-Holmes & Roche, 2001) approach emerged from the behaviour-analytic literature around the same time that Implicit Social Cognition research began to take off, and has led to a multitude of advances in the behavioural understanding of complex human language and cognition. The RFT approach to language conceives the social-cognitive ‘attitude’ concept in terms of a network of trained and derived stimulus relations, established within an individual’s verbal and non-verbal interaction histories (Grey & Barnes, 1996). Behaviour-analytic measures work from within a functional and well-defined paradigm, potentially avoiding some of the theoretical and conceptual confusion that is associated with social cognitive measures (De Houwer, 2011; Hughes & Barnes-Holmes, 2011). The RFT approach is therefore believed to be a potentially powerful explanatory tool for research surrounding implicit attitudes (O’Reilly, Roche, & Cartwright, 2015).

Much of the research that lead to the current RFT approach has grown from Sidman’s (1971) investigations into the phenomenon of stimulus equivalence. The phenomenon of stimulus equivalence began as a way of teaching reading comprehension to participants who were severely developmentally disabled. Sidman used a Matching-to-Sample (MTS) procedure whereby a participant would learn a series of related conditional discriminations. The stimuli involved in those discriminations, usually presented in visual forms, often became related to each other in ways that were not explicitly trained. The simplest form of conditional discrimination involves four stimuli, two of which may be termed “samples,” and two “comparisons.” Conditional discrimination training normally consists of reinforcing the choice of one comparison (let’s call it B1) when presented with one of the samples (let’s call it A1), and reinforcing the choice of the alternative comparison (B2) when presented with the other sample (A2). Subsequent to such training, if a participant is presented with B1 as a sample, A1 will generally be chosen as a comparison, and if presented with B2 as a sample A2 will be chosen as a comparison. In effect, the two relations A1-B1 and A2-B2 are directly trained, but the relations B1- A1 and B2-A2 are derived without any further explicit training. If the individual concerned is then taught a second related conditional discrimination such as B1-C1/B2-C2, the number of relations that may be derived increase dramatically. In fact, it has been repeatedly demonstrated that training the two related conditional discriminations, A1- B1/A2-B2 and B1-C1/B2-C2, will result in the emergence of the following eight derived relations: B1-A1, B2-A2, C1-B1, C2-B2, A1-C1, A2-C2, C1-A1, and C2-A2. If these emergent or derived relational responses are observed, the stimuli involved are said to participate in equivalence relations (Dymond & Barnes, 1994; Sidman & Tailby, 1982).

Sidman’s conceptualization of stimulus equivalence defines it in terms of the mathematical relations of reflexivity, symmetry, and transitivity. Specifically, reflexivity is where the participant will match each stimulus with itself, i.e. when a participant is presented with A1, then A1 will be picked. Symmetry involves successfully reversing the direction of the trained relation, i.e. when a participant is presented with B1, A1 will be picked. Lastly, transitivity involves deriving the untrained identity relation between the stimuli that were never paired, i.e. when a participant is presented with C1, A1 will be selected. RFT considers equivalence as just one example of a diverse and extensive range of relations that fall under the phenomenon of Derived Relational Responding (DRR; Hayes et al., 2001; Törneke, 2010). RFT posits that DRR refers to the multitude of ways in which stimuli may be related due to an individual's verbal history of interaction within a social and non-social environment (Leigland, 1997). More specifically, in the same way that reinforcement contingences within a particular history may derive equivalence classes between stimuli (e.g., derived A-C relations), RFT proposes that there are a multitude of ways that different relations may be derived or abstracted between stimuli. Simply put, this essentially refers to the same core process as outlined above (i.e., A-C derived relations), but refers to the variety of different ways in which these relations can be derived (Hayes, 1994). The different forms of relational responding (difference, opposition, greater than, less than, etc.) can also be derived without explicit training (e.g., Dymond & Barnes, 1995; Lipkens, Hayes, & Hayes, 1993; Roche & Barnes, 1996). Therefore, DRR is perhaps the core behaviour underpinning language and cognition.

The pivotal role played by derived relational responding in language led researchers to investigate how histories of verbal behaviour might impact upon participants' formation of new equivalence classes. In their seminal study, Watt, Keenan, Barnes and Cairns (1991) used a simple stimulus equivalence paradigm in which participants were trained to relate stimuli with strong socially established functions in ways which were inconsistent with their social history. There was a tendency among Northern Irish individuals to categorize names on the basis of Protestant or Catholic religions (Cairns & Mercer, 1984), and so Watt and colleagues (1991) utilised stimuli representative of Catholic and Protestant names and symbols. Participants were exposed to a simple MTS procedure, whereby relations between Catholic family names and nonsense syllables (A-B relations) were trained alongside relations between the same nonsense syllables and Protestant symbols (B-C relations) with the aim of deriving AB-C equivalence classes. It was expected that there would be a resistance observed for the Northern Irish sample to the formation of derived equivalence classes between Catholic names and Protestant symbols relative to the English participants.

Indeed, subjects did not respond in accordance with new contingencies for equivalence (i.e., Catholic symbols with Protestant names) due to their previous and more extended learning histories with those stimuli (i.e., resistance to behavioural change was strong), with 12 out of the 19 Northern Irish participants incorrectly pairing a novel Protestant name rather than the expected Catholic name with the Protestant symbols. This procedure, therefore, provided clear experimental evidence that an individual's history of relational responding may interfere with the formation of novel relations (O'Reilly et al., 2012). In this way, the Watt et al. (1991) study can be regarded as an empirical example of the resistance to change observed when required to form novel equivalence classes (i.e., Catholic names — Protestant symbols) that oppose the existing verbal histories for an individual.

1.5.1 The Implicit Relational Assessment Procedure

The DRR paradigm has been used by behaviour-analytic researchers to develop more functional, inclusive, and discriminative investigative tools for ascertaining verbal histories of relational responding. The first of these measures, the Implicit Relational Assessment Procedure (IRAP; Barnes-Holmes et al., 2008) was developed to expand on the Watt et al. (1991) procedure, and to provide a measure capable of more comprehensively assessing derived relations. The IRAP provides a direct measure of implicit attitudes, postulating that complex stimulus relations, as opposed to simple associations, are fundamental to understanding attitudes. One of the key differences between the IAT and the IRAP is that the IAT produces a single metric of overall bias, whereas, the IRAP produces four separate metrics, one for each of the individual relational responses that are targeted by the IRAP (Hussey, Barnes-Holmes, & Barnes-Holmes, 2015). Additionally, the IRAP can measure the strength and directionality of relatively simple stimulus relations and more complex relational networks.

The IRAP is described as being a combination of the IAT and an earlier RFT based procedure known as the Relational Evaluation Procedure (REP; Hayes, Barnes-Holmes & Roche, 2001). One of the core assumptions of RFT is that the behavioural units of human language and thought may be defined in terms of derived stimulus relations and relational networks (Barnes-Holmes et al., 2005). According to RFT, the IRAP involves participants responding rapidly and accurately to stimuli considered both consistent and inconsistent with their prior verbal learning histories (Cullen, Barnes-Holmes, Barnes-Holmes, & Stewart, 2009). Similar to the REP, the IRAP involves presenting two specific relational terms (e.g., true, false, similar or opposite) so that the relations among the relevant stimuli can be measured and its properties assessed (Barnes-Holmes et al., 2008). Like the IAT, the IRAP requires participants to respond rapidly and accurately to multiple test blocks that are either consistent or inconsistent with the individual's suspected attitudes (Barnes-Holmes et al., 2008). However, due to the presence of contextual cues during trials (e.g., words such as true and false), the tasks are considerably more complex, and therefore more complex inferences can be drawn from response latencies and accuracies regarding relations between concept categories (i.e., attitudes). The basic hypothesis is that average response latencies should be faster and more accurate across blocks of consistent comparative to inconsistent trials (Barnes-Holmes et al., 2006).

The Relational Elaboration and Coherence (REC) model is a behaviour-analytic interpretation of the behaviours captured by the IRAP, which provides an explanation of how behaviour-analytic theory may be extended to the measurement of an individual's history of relational responding (Barnes-Holmes, Barnes-Holmes, Stewart, & Boles, 2010). Essentially, the REC model posits that the behaviours captured by indirect measures such as the IRAP reflect a relation within the individual's verbal and non-verbal histories between some targeted relation (e.g., 'Love'/ 'Hate' and 'Pleasant'/ 'Unpleasant') and a contextual cue (e.g., the relational terms ‘Similar’/ ‘Opposite’). Within the paradigm of the REC model, therefore, each trial on the IRAP is believed to elicit a brief and immediate relational response that is a function of participants’ learning histories. The most probable response is believed to be emitted first and in the shortest time, such that response latencies should be shorter for trials that are consistent with the individual’s learning histories. In contrast, if the trial is inconsistent with their verbal histories (that is, the required response is in opposition with their immediate or most probable response), the REC model posits that response latencies will be lower. In this way, measures such as the IRAP are believed to differ from traditional explicit measures primarily in terms of time pressure (Barnes-Holmes et al., 2010). That is, within the IRAP, participants are required to respond quickly and accurately within a particular time frame, and consequently, according to the REC model, the most immediate (i.e., the most probable) responses are emitted.

A number of published studies using the IRAP have examined social biases of adults for race (Power, Barnes-Holmes, Barnes-Holmes, & Stewart, 2009), gender (Drake, Primeaux, & Thomas, 2018), religion (Drake et al., 2010), age (Cullen et al., 2009), attractiveness (Murphy, MacCarthaigh, & Barnes-Holmes, 2014), as well as many other social biases (Hughes & Barnes-Holmes, 2011). Preliminary findings from these studies have also provided support for the IRAP in terms of reliability and validity (Drake, Seymour, & Habib, 2016). However, while the IRAP is making substantial progress in terms of providing a functionally understood implicit test methodology, and while it has spawned a very large amount of research worldwide, it too suffers from several methodological and conceptual limitations (Gavin, Roche, & Ruiz, 2008; Ridgeway et al., 2010). Generally speaking, because it is based heavily on the stimulus presentation and response quantification methods of the IAT, it has inherited many of the problems of the IAT as outlined above. In addition, it enforces strict inclusion criteria which can eliminate the majority of the participants from any given sample and takes up to 60 minutes to complete (e.g., Barnes-Holmes, Murphy, Barnes-Holmes, & Stewart, 2010). The length of time required to complete a typical IRAP may not lend itself as easily to widespread application within practical settings and, consequently, for practical purposes, the development of a shorter measure may be pragmatically beneficial.

While the IRAP is making substantial progress in terms of providing a functionally understood implicit test methodology, it has been suggested that one slight limitation of the Watt and colleagues (1991) procedure and the IRAP measure is the fact that the procedural implicitness of the tests may, in a similar manner to the IAT, be compromised by presenting the specific relation under analysis on-screen during the procedure (Roche, O'Reilly, Gavin, Ruiz, & Arancibia, 2012). That is, while the exact nature of the test may not be immediately discernible to participants, it is unclear whether or not having the target stimuli on-screen alongside the relational term alerts the participant to the issue under investigation. While one study has directly investigated the fake-ability of the IRAP (i.e., the degree to which the IRAP effect can be diminished by conscious control) and revealed no significant faking effects (McKenna, Barnes-Holmes, Barnes-Holmes, & Stewart, 2007), an additional measure adopting a different approach would undoubtedly contribute to the growing literature on behaviour analytic implicit measures. Lastly, as previously discussed, some researchers have suggested that relying on response accuracy or fluency rather than reaction time (as in the IAT and IRAP methodologies) may provide a more accurate representation of participant difficulty across blocks (Gavin et al., 2012).

1.5.2 The Multi-Dimensional Multi-Level Framework

From the IRAP and REC Model, a more recent conceptual work on RFT has led to the development of a multidimensional, multilevel (MDML) conceptual framework for analysing the dynamics of arbitrarily applicable relational responding (Barnes-Holmes, Barnes-Holmes, Hussey, & Luciano, 2015). The MDML constitutes a step beyond the REC model, where it readily accommodates dynamic interactions in arbitrarily applicable relational responding across levels of relational development. The framework currently consists of four dimensions; coherence, complexity, derivation, and flexibility; and five levels of relational development; mutual entailing, relational framing, relational networking, relating relations, and relating relational networks (Barnes-Holmes, Barnes-Holmes, Luciano & McEnteggart, 2017). Within the MDML, each of the dimensions intersects with each of the levels, yielding 20 potential units of behavioural analysis, defined as functional-analytic abstractive relational quanta. And, perhaps most importantly of all, each of these units may be conceptualized as a verbal or relational response class that may enter into a discriminated operant, thus allowing for direct manipulation via appropriately arranged environmental contingencies of reinforcement. Within the MDML, therefore, there is no simplistic dichotomy between verbal and nonverbal behaviour based on whether a response is derived versus controlled by direct acting contingencies. Rather, all units of analysis within the MDML remain verbal response classes even when they are impacted upon by direct contingencies of reinforcement (O'Hora, Barnes-Holmes, & Stewart, 2014). In effect, the MDML serves to highlight the intensely operant but wholly verbal nature of RFT, with a focus on the impact of direct acting contingencies on its (verbal) operant units of analysis, from the most simple or basic relational responses to the most complex contextually controlled interactions among complex relational networks (Barnes-Holmes et al., 2016).

The MDML has also thoroughly examined a systematic analysis of the behavioural dynamics involved in the IRAP, but not as a measure of implicit cognition as such (Barnes-Holmes et al., 2015). Questions about the various patterns of behaviour observed on the IRAP, and the variables that appear to be important in producing those patterns, have been interpreted in terms of this novel MDML framework (Barnes-Holmes et al., 2017). For example, two patterns of arbitrarily applicable relational responding are required on a Flowers-Pleasant trial-type. Recall that a relatively large IRAP effect would be expected in a history-consistent direction (i.e., responding “True” more quickly than “False”). Such an IRAP effect may be conceptualized within the MDML as arbitrarily applicable relational responding for which the functional-analytic abstractive relational quanta’s are high in coherence, low in complexity, low in derivation, and low in flexibility during consistent trials, but during inconsistent trials coherence reduces, and complexity, derivation, and flexibility increase. In other words, confirming that flowers are pleasant is likely to cohere with the reinforcement contingencies operating in the wider verbal community. The required relational response is not particularly complex (i.e., confirming a mutually entailed relation), will have been derived numerous times in many contexts (e.g., when buying or receiving flowers as a gift), and may be relatively inflexible (assuming that there are very few, if any instances, in which flowers have acquired aversive functions in the natural environment; Barnes-Holmes et al., 2017). In summary, the MDML model offers a novel framework for RFT, as opposed to being a new or alternative theory, focusing on those features of the theory that appear to be the most important currently but have remained understated in much of the early work on RFT (Hayes, Barnes-Holmes, & Wilson, 2012).

1.5.3 Function Acquisition Speed Test

One behaviour-analytic alternative to the IAT and IRAP is the Function Acquisition Speed Test (FAST; O’Reilly et al., 2012), a bottom-up approach to test development based entirely on well understood learning principles. Almost every feature of its presentation and scoring format is based on systematic experimentation and so it is a methodologically transparent approach to measuring associations between concepts. It does not reply upon top down or post-hic theoretical approach (such as the MMDL), but relies instead on empirically controlled effects to drive behaviour-analytic explanations in terms of behavioural principles.

The FAST is an effective tool for the quantification of stimulus relatedness, which simply refers to the degree to which stimuli have been trained (either directly or through derivation) as being paired in the previous history of the individual (Cartwright, Roche, Gogarty, O'Reilly, & Stewart, 2016; O'Reilly et al., 2012; Roche, O'Reilly, et al., 2012; O'Reilly, Roche, Gavin, Ruiz, Ryan, & Campion, 2013). The FAST has a presentation format that at first glance is similar to that of the IAT but it is understood in a conceptually different way, and differs methodologically in empirically valid ways in almost every respect. The most obvious difference between the FAST and alternative tests is that it does not measure differences in response times across two key test blocks but rather the raw (not standardised) difference in the slope of the learning curves (i.e., the trial-by-trial rate of increase in speed and accuracy across successive trials) measured for each of two key test blocks. This is a sensitive measure that has been shown in fully controlled laboratory experiments, using entirely laboratory created relational networks to measure verbal relations reliably (O'Reilly et al., 2012; 2013).

The FAST shares a common overarching goal with the IRAP, in that it aims to assess whether natural verbal relations impact laboratory-induced patterns of relational responding. However, where both the IAT and IRAP train participants to respond quickly and accurately with a series of practice/categorization blocks prior to testing, the FAST focuses on the rate of acquisition of those patterns themselves (O'Reilly et al., 2012). That is, rather than investigating whether verbal histories facilitate or impede the speed at which classes of stimuli may be related together in a novel context, the FAST is concerned with the impact of history upon the differential rate at which these relations may be acquired. In brief, the FAST procedure establishes functional equivalence classes between sets of stimuli by requiring participants to learn, via corrective feedback, a common response (e.g., a left or right key-press) when presented with individual stimuli on-screen. Participants are presented with two contrasting test blocks (one that establishes culturally-consistent response classes, and another that establishes culturally-inconsistent classes), with the primary metric being the difference in learning rate across the two blocks. Each block is comprised of four stimulus classes, with two classes sharing a left key-press response and two sharing a right key-press response (i.e., they are spatially and formally equivalent within the task). In contrast to the IAT and IRAP, instructions on how to respond correctly are not given prior to the procedure, and no category/attribute labels or rules are provided at any point in the task. Rather, the establishment of a functional equivalence class is achieved exclusively through reinforcement (i.e., via verbal feedback). Latency pressure is applied by restricting the response window to a relatively narrow time frame (usually 3000ms), after which participants are presented with a timeout stimulus (O'Reilly et al., 2012).

The parameters of the FAST can be modified to suit different populations or experimental requirements. To date, the FAST paradigm has demonstrated utility in the assessment of both directly trained and derived equivalence relations established in the laboratory (O'Reilly et al., 2013; O'Reilly et al., 2012). The FAST has also shown to be a useful tool to assess sensitive sexual behaviour-related topics. For example, findings from a precursor of the FAST studied sexual interests among a normal population that indicated a negative effect for prepubescent female images and sexual imagery (i.e., these pairs of stimuli were not related), and indicated a positive FAST effect when the images were of female adults (i.e., these stimuli pairs were related). These findings suggest that subject's histories enabled the formation of a response class for adult female images and generic sexual images, but inhibited the formation of such a response class for prepubescent female images and generic sexual images (Roche et al., 2012). These findings can be supported by O'Reilly and colleagues (2013), who also found that on average participants had a quicker rate of acquisition for response classes that were consistent with a prior learning history compared to those which were inconsistent. This evidence proves that the FAST may be a useful tool for quantifying stimulus relatedness (Cummins, Roche, Tyndall, & Cartwright, 2018; Roche et al., 2012), and has since been used as a measure of verbal stimulus relations in the context of condom use (Cummins, Tyndall, Curtis, & Roche, 2019), and in assessing implicit gender stereotypes (Cartwright et al., 2016).

1.6 Predicting Behaviour

One of the main issues that has emerged within the literature on implicit attitudes is their predictive validity. If implicit attitudes do not differ significantly from explicit attitudes, and if they do not provide additional information for predicting behaviour, then the usefulness of measuring implicit attitudes is limited. In general, research has found that explicit attitudes predict deliberate behaviour and implicit attitudes predict spontaneous behaviour (e.g., Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997; Dovidio, Kawakami, & Gaertner, 2002; Fazio, et al., 1995; Spalding & Hardin, 1999). Fazio's (1990) MODE (Motivation and Opportunity as Determinants of processing) model provided a framework to explain when behaviour would be spontaneous, and best predicted by implicit attitudes, and when behaviour would be deliberate, and best predicted by explicit attitudes. A meta-analysis of the predictive validity of the IAT found that the IAT was a good predictor of spontaneous behaviour. However, the IAT was also effective in predicting some behaviour that was controllable, such as brand related choices and voting behaviour (Greenwald et al., 2009).

Recently, political psychologists have proposed the use of the IAT to measure implicit attitudes toward political parties and candidates and to predict subsequent voting behaviour. For example, researchers found that implicit race attitudes as measured by the IAT correctly predicted US voter intention in the 2008 US presidential election (Greenwald, Smith, Sriram, Bar-Anan & Nosek, 2009), highlighting the validity of the IAT in determining implicit attitudes to race that may impact participant behaviour. Previous research has demonstrated that although people may be aware of their implicit attitudes, they may not be aware of the ways in which their implicit preferences may impact their behaviour (Gawronski, Hofmann, & Wilbur, 2006). Hence, in the political realm, individuals' implicit attitudes toward parties or candidates may affect their eventual voting behaviour beyond their explicit attitudes. In fact, IATs were shown to provide predictive validity over and above explicit measures in a number of studies conducted ahead of political elections (e.g. Friese, Bluemke, & Wanke, 2007; Arcuri, Castelli, Galdi, Zogmaister, & Amadori, 2008). Therefore, evidence strongly suggests that political researchers should take into account both explicit and implicit measures when concerned with the prediction of voting behaviour.

Friese and colleagues (2007) used IATs for the first time in the lead-up to the 2002 parliamentary elections in Germany. They found incremental predictive validity of Single-Target IATs (ST-IATs), a modification of the IAT measure that assesses the strength of evaluative associations with a single attitude object (Bluemke & Friese, 2007), for the five major political parties. Although self-reported party preferences were stronger predictors of voting intention and actual voting behaviour, the implicit measures added significantly to the explanation of voting behaviour. Recent research suggests that implicit attitudes may be more useful for predicting the vote of undecided voters, particularly in the case of specific political issues through referendums, rather than elections (Raccuia, 2016). IATs may be more useful for the prediction of voting behaviour in the case of specific political issues rather than political elections, as a result of less elaborated attitudes toward these issues (Galdi, Arcuri, Gawronski, 2008; Friese, Smith, Plischke, Bluemke, Nosek, 2012). In contrast to political elections where voting behaviour is typically determined by long-lasting party affiliations, voting on referendums represents a much more complex situation. If this is true, then the IAT should be a useful tool for the prediction of political outcomes in Ireland, particularly in the past abortion issue-related referendum.

The IRAP also appears to be equal to or better than direct measures in predicting real-life behaviour. This is shown in a study conducted by Carpenter and colleagues (2012), who administered the IRAP along with a series of self-report questionnaires to a group of participants who were cocaine dependent before the commencement of a 6-month outpatient treatment program. The study found that, unlike the participants self-reports, participants IRAP performances before the treatment program correctly predicted their likelihood of not only attending the program, but also whether they adhered to the program during the first 12 weeks of treatment. Similarly, the IRAP has also been found to predict clinically relevant outcomes with respect to aversive responding in spider fearful individuals. Consistent with past work using the IAT (e.g. Teachman, Gregg, & Woody, 2001) the IRAP successfully distinguished between high and low spider fearful participants. Additionally, the IRAP successfully predicted how likely participants were to approach and interact with a live spider following the procedure. Juarascio and colleagues (2011) reported that an IRAP measure designed to assess implicit idealisation of thinness, was prospectively predictive of changes in weight, disordered eating, and body dissatisfaction over the subsequent year. The IRAP measure demonstrated incremental predictive validity above applied explicit measures. Other studies have used the IRAP to predict sexual orientation identification (Rönspies, 2015) and sexual offender status (Dawson, Barnes-Holmes, Gresswell, Hart, & Gore, 2009).

1.7 Known Groups Paradigm Approach

If a test is “valid,” one criterion could be that test scores should discriminate across groups that theoretically are expected to be different on the trait measured. If this is so, there is evidence of the usefulness of the test as a decision-making instrument and evidence that it can be generalized on a meaningful psychological trait across different samples of people (Hattie & Cooksey, 1984). Cronbach and Meehl (1955) first described this when presenting four methods that examine construct validity, one of them being to test for known group differences (e.g., a depression measure should differentiate between those who have been diagnosed with depression and those who have not). This “known-groups” approach to validity argues that a good measure should reliably distinguish between members of different groups, based on a prior predictions or knowledge about those groups.

Numerous IAT studies have utilised a known-groups approach to testing validity, beginning with prior expectations about the groups where differences in attitude is expected, and comparing that with their IAT scores (Lane et al., 2007). For instance, when entomologists show a smaller effect on the same test, compared to a control group (Lane et al., 2007) that is, they demonstrate a weaker relative preference for flowers over insects, we can conclude that the IAT reveals individual differences in strength of association, using a known group’s validation approach. That is, we begin with prior expectations of a group to show differences in attitude, and test this expected difference using the IAT measure. Furthermore, if individual differences in strength of association represent meaningful differences in attitude, the test should predict other behaviours, such that those who associate insects with positive traits are more likely to spend time with insects and to act in favourable ways toward insects, such as feeding and caring for them and being invested in their survival (Lane et al., 2007). Taken together, these findings suggest that groups differ in strength of associations in predicted ways, that one’s personal experience alters the magnitude of those associations, and that those associations relate reliably to other individual judgments and behaviours.

The IAT has indeed demonstrated theoretically predicted patterns of known-group membership. Previous research examining the validity of the IAT using the known-groups method (i.e., comparing groups expected to differ on implicit attitudes), has shown that adult smokers' implicit attitudes were less negative than that of non-smokers (Swanson, Rudman, Greenwald, 2001; Huijding, de Jong, Wiers, & Verkooijen, 2005), that Japanese Americans exhibited strong preference for their group relative to Korean Americans, whereas Korean Americans showed the opposite pattern (Greenwald et al., 1998), that East and West Germans each exhibited preference for their in-group (Kuhnen et al., 2001), and even members of groups artificially created in the laboratory showed preference for their in-groups (Ashburn-Nardo, Voils, & Monteith, 2001).

It can be concluded that a known-groups paradigm approach is essential in validating implicit test measures, especially that of the newly developed FAST procedure. However, it is of equal importance to find an attitude that is clear and unambiguous in order to reduce variance. Examining an attitude that clearly distinguishes groups will test the FASTs ability in predicting known group membership. Here, the main concern is to firstly assess the FAST where behaviour is known and long standing unambiguous attitudes are known, and to secondly compare its validity to that of the IAT. Therefore, the abortion attitude would be an ideal attitude to examine in that it tends to be unambiguous, where many people are either strongly against abortion or strongly for abortion.

1.8 Attitudes to Abortion

If any issue attitude could play the role of "irresistible force," it is abortion (Killian & Wilcox, 2008). The issue has transformed ordinary people into extraordinary activists (Maxwell, 2002). It has produced collective action of unusual intensity and duration, inspiring mass demonstrations and passionate debates. Abortion attitudes can even influence vote choice in presidential, local and general elections; where liberals are essentially in favour of prochoice ideology, and conservatives are in favour of pro-life ideology

(Abramowitz, 1995; Cook, Jelen, & Wilcox, 1994). Liberals are characterized by the belief that freedom should be exercised by the individual; whereas, conservatives believe freedom should be exercised by institutions, such as communities, church, businesses, and families (Saletan, 2003). Yet, the majority of the public holds an ambivalent and collectively moderate position on abortion (Cook, Jelen, & Wilcox, 1992). This ambivalence comes from conflicting values; many people value both an emergent foetal life and women's moral autonomy (Alvarez & Brehm, 1995). To the question of whether abortion should be allowed by law, the median voter replies "it depends" on the reasons the woman is seeking the abortion, on the timing of the abortion, and even on the procedures used (Norrander & Wilcox, 1999).

1.8.1 History of Abortion in Ireland

Abortion laws in the Republic of Ireland were among the most restrictive in the world. For decades, women living in Ireland who had a pregnancy they did not want or did not feel they could continue to have had three options; travel abroad to obtain a safe, legal termination of pregnancy, self-induce using an unsafe method, or remain pregnant (Aiken, Digol, Trussell, & Gomperts, 2017). Between 1970 and 2015, an estimated 180,797 women from Ireland travelled, most often in secrecy, to access abortion services in Great Britain (Aiken, Gomperts, & Trussell, 2017). Cultural and religious stigma, along with social pressures had a massive influence in regards to access to abortion care (Pinter, 2002). The stigma surrounding abortion in Ireland stemmed from the social attitudes to sexual activity cultivated by the Catholic social teaching which deemed contraceptive use, sexual intercourse outside marriage, and abortion to be sinful and immoral (Inglis, 1998).

Catholicism had become ingrained in Irish society and culture in the twentieth century, particularly in the area of sexual morality, and the Church and State worked together to create their desired image of a puritanical Irish community (Coakley & Gallagher, 2018).The alliance between Church and State played a huge role in the emergence of moral purity as a national trademark, one that sought to conceal social problems and represent Ireland as a beacon of goodness and chastity (Smith, 2007). Within the first five years of the establishment of the Free State in 1922, the Church and State had already collaborated on legislation that had censored films and books and outlawed divorce (Beaumont, 1997). In 1930, the Carrigan Committee was formed to address issues such as Ireland's legal age of consent and the rise of juvenile prostitution. However, the Carrigan Report portrayed Ireland in an unflattering light, so it was decided on direction from the Department of Justice that it would not be made public (Kennedy, 2000). It was determined that wide circulation of the report would encourage public discussion about morality in Ireland, and that doing so could undermine the Catholic Church (Smith, 2004). This marked the beginning of stigmatization surrounding illegitimacy and contributed directly to the perpetuation of sexual oppressive in Ireland (Smith, 2004).

Social attitudes to sexual activity and contraceptive use in Ireland was shaped by the Catholic social teaching which deemed sexual intercourse outside marriage to be sinful and immoral (Inglis, 1998). Many parents therefore, presumed that their children would follow such teaching, so sexual intimacy and the use of contraceptives were not discussed within families (Mahon, Conlon & Dillon, 1998). The lack of knowledge about sex could be considered a major contributing factor to social problems such as illegitimacy, abortion and infanticide. Sweetman (1979) found considerable naiveté and ignorance about sex amongst the Irish population, where one of the women she interviewed admitted to becoming pregnant with her first child without even knowing how. This naiveté stemmed from a lack of adequate sex education from schools and parents (Inglis, 1998), where there was a strategy of silence and prudence about all matters concerning sex as it was perceived as being something shameful to talk about (Elias, 1978). The central feature of the Catholic Church's approach to sex education was the constitution of human sexuality and reproduction within the natural law of god (Inglis, 2005). Human birth was taught as the miracle of life, where pregnancy and birth was something which happened to women as a part of God's plan. Human reproduction became wrapped in mystery, vague, romantic imagery (Elias, 1978).

Ireland's preoccupation with sexual immorality was mainly due to the number of children being born outside of marriage (Ferriter, 2012). However, religious teachings alone failed to prevent many thousands of single young women from becoming pregnant (Nolan, 1989). The State failed to adequately inform and prepare students, deferring to religious instruction to steer clear of taboo subjects in order to prevent moral impurity. This was the main reason for high levels of illegitimacy throughout the 19th and 20th centuries (Inglis, 1998). Between 1922 and 1979 over 100,000 illegitimate children were born in Ireland, giving a cause for concern on the part of both the Irish government and the Catholic Church who were fixated over these births (Nolan, 1989). During the twentieth century, there was a conscious attempt by the Irish government to move away, in some respect from what was considered British (Miller, 1975). Ireland was promoted as a Catholic and morally pure country, thus those who transgressed this moral code were frowned upon and, if possible, were to be somewhat hidden from society. Unmarried mothers and their children were one such group (Rossiter, 2009). Therefore, sexuality and perceptions of it were directly concerned with unmarried motherhood (Ferriter, 2012).

Ireland always had a reputation for being ignorant, incompetent, shy and awkward when it came to sex (Ferriter, 2004). After years of sexual oppression, or what James Smith coined in 2007 (pp. 2) as "Ireland's architecture of containment", sex was often associated with feelings of guilt, shame and embarrassment. There was a general belief that sex was a sin and that a young girl who had sex outside of marriage was tainted (Inglis, 2005). In an interview conducted by Richards (1981), one interviewee gave insight into the stigma associated with illegitimacy, stating that, “For as long as I could remember the worst thing that could happen to anyone’s daughter in Ireland was to have a baby and not be married” (Richards, 1981, pp. 53). Girls who did become pregnant before marriage were often brought to the Magdalene laundries or mother and baby homes by their families or their local clergy (Smith, 2007). This was done so that the respectability and good name of the family would not be tarnished by any scandal (McCarthy, 2010). This is depicted in the Irish film ‘Sex in a Cold Climate’ (Humphries, 1998), where one of the women interviewed who was brought to the Magdalene Laundries accounted that, “The biggest sin in Ireland, well apart from having a baby in them days without being married, was to talk… You never let the neighbours know.” Diarmaid Ferriter rightly called this “a preoccupation with what was seen, not what was suffered” (Ferriter, 2012; pp. 546). These institutions served to contain socially marginalised women and the reality of extra-marital sex in Ireland in this period (McCarthy, 2010). “Not in Ireland” (Ferriter, 2012, pp. 398) became the nation’s unofficial mantra. It wasn’t just the Catholic Church, but rather the wider deference to the Catholic Church by both government and society, which resulted in a culture where shame was used as a tool and silence was seen as the only solution (Keating, 2004).

Gradually, from the late 1960s, the Church’s dominance in the field of sexuality began to be challenged, primarily through the media where censorship began to be relaxed (Inglis, 1998). Over the next thirty years, progressive and liberal viewpoints about sexuality existed side by side with the traditional perspective of the Catholic Church (Inglis, 1998). However, women with unwanted pregnancies soon realised that, despite the messages received through the media, the revolution had not started in their family or community (Rossiter, 2009). They were on their own, and left with feelings of shame, remorse and guilt. Irish women often travelled to England for abortions, where the numbers increased rapidly each year (Mahon, Conlon & Dillon, 1998). Mahon and colleagues (1998) indicated that the main reason Irish women had abortions was due to the shame attached to women's pre-marital sexual activity, and due to their fear of disclosing their pregnancies to their parents. They had witnessed other lone mothers in their community being subjected to stigma and they were concerned about the stigma that their family would be subjected to if they were to "come out" and disclose their pregnancy (Mahon, Conlon & Dillon, 1998, pp.530). The route travelled between Ireland and England for abortions became known as the abortion trail and the people who travel it Ireland's "invisible export" (Rossiter, 2009, pp. 36).

It was not until the mid-nineties that a formal sexual education program was introduced in the Republic (Mayock & Kitching, 2007). It is said that local events contributed to this shift in the educational discourse around sex, such as the death of pregnant teenager Ann Lovett and her child (Maguire, 2001). Lovett was 15 years old when she was found clinging to life in a grotto of the Blessed Virgin outside Granard, Ireland in the winter of 1984. The body of the baby boy she was carrying, which no one admitted to knowing about, was found on a stone beneath the Blessed Virgin statue, and Ann herself died in the hospital shortly afterwards (Ferriter, 2009). Ann's death marked the death of denial in Ireland; no longer could people claim ignorance or hide behind the impenetrable cloak of the Catholic Church (Maguire, 2001). The growing phenomenon of infanticide, as well as global tragedies, such as the HIV crisis and the abuse scandals of the Catholic Church, put pressure on the Irish government to introduce sexual education into schools (Wilentz, 2016). This was met with strong opposition from the Catholic Church, and it would take over 10 years for Ireland to introduce the Social, Personal and Health Education (SPHE) curriculum, which included Relationships and Sexuality Education (RSE) (Wilentz, 2016).

1.8.2 History of Abortion Laws

The first official document criminalizing abortion in Ireland was passed in 1861, where procuring a termination of pregnancy, or assisting a woman in such a procedure, had been a criminal offence in the UK and Ireland, under the Offenses Against the Person Act. These laws remained in place even after Irish independence in 1922. It would be over 120 years after the Offences Against the Person Act before abortion laws would once again be addressed in Ireland (Ferriter, 2009). After the 1973 Roe v. Wade decision in the U.S., which granted a woman the right to have an abortion under the Fourteenth Amendment of the Constitution, people and politicians in Ireland feared a similar judicial ruling that would allow women to seek abortion privately (Fletcher, 2000). The Catholic conservative lobby was strong in the 1980s in Ireland, where several activist groups such as the Council of Social Concern, the League of Decency and the Irish Family League, united to form the Pro-Life Amendment Campaign (PLAC). They campaigned for the introduction of a bill that would insert a "pro-life" amendment into the Irish Constitution that would make it unequivocal in its ban on abortion (Rhinehart, 2013). They successfully campaigned to have a constitutional ban on abortion introduced in 1983. The Eighth Amendment of the Irish Constitution (Article 40.3.3) was passed by the Irish public by a 67 percent majority, which effectively banned abortion in all situations in Ireland, guaranteeing to protect the life of the unborn and its equal right to life to the mother (Schweppe, 2008).

In 1992, the refusal to allow a 14-year-old rape victim who was suicidal to travel abroad for a termination of pregnancy by a High Court injunction, led to the well-known "X Case" (Smyth, 1998). The young girl, known as X, expressed suicidal thoughts as a result of the pregnancy to her mother, and the family decided to travel abroad to terminate the pregnancy. Before they left, the family asked the police if DNA could be used from the aborted foetus to help convict the perpetrator, who denied his involvement. Hearing that X planned to have an abortion, the local Attorney sought an injunction against her under the Eighth Amendment, preventing the family from following through with their plans for termination (Smyth, 1998). The X case was first brought to the Irish High Court, who granted a temporary and then permanent injunction against X, essentially barring her from leaving Ireland until the birth of her child. The parents of X then appealed the High Court's decision to the Supreme Court, who overturned it and ruled that suicide was a valid reason for obtaining an abortion under the clause of Article 40.3.3 requiring the State to have "due regard to the equal life of the mother" (Whitty, 1993). The resultant common law and two subsequent referenda in 1992 and 2002 upheld the right to abortion services in Ireland if there was a real and substantial risk to the life of the mother, including the risk of suicide. A constitutional right for Irish women to travel abroad to access abortion services was established (Whitty, 1993).

The decision in Open Door and Well Woman vs. Ireland also had bearing in another piece of legislation passed in Ireland's ongoing abortion discussion (Fletcher, 2000). The applicants were two non-profit organizations, neither of which advocated or encouraged abortion, but provided information about pregnancy-related options. The Ireland Supreme Court imposed restrictions on the two applicant companies and prohibited them from providing information to pregnant women as to the location or identity of, or method of communication with, abortion clinics in Great Britain (European Court of Human Rights, 1992). The European Convention on Human Rights held that the Ireland Supreme Court's injunction restraining counselling agencies from providing pregnant women with information concerning abortion facilities abroad violated Article 10 of the Convention (European Court of Human Rights, 1992). In response to the success of this litigation, and prompted by the controversy on the X Case, a referendum was held in 1992 on the Fourteenth Amendment, which passed. The Fourteenth Amendment (1992) specified that the prohibition of abortion would not limit the right to distribute information about abortion services in foreign countries. This was governed by the Regulation of Information Act, passed in 1995, giving legal permission to doctors and medical groups to disseminate information on abortion services abroad to patients soliciting this information, but preventing direct referrals for abortion services in the UK.

The Irish government faced increasing pressure to legislate for situations when a doctor could terminate a pregnancy if the life of a mother was in danger, both externally from the European Court of Human Rights, and internally from the political fallout from the death of Savita Halappanavar (Murray, 2016). The Protection of Life During Pregnancy Act (PLDPA) was passed in 2013 permitting Irish women to obtain an abortion in Ireland in very limited circumstances, if deemed appropriate by medical professionals. This legislation came in response to Savita Halappanavars death during an inevitable miscarriage, six days after her husband first requested a termination of pregnancy due to her severe pain and rapidly declining health (Felzmann, 2014). The PLDPA stated that termination can occur legally in three circumstances; if the Mother's life is at risk from physical illness, if the Mother's life is at risk during an emergency, or if there is a clear risk of suicide. To access abortion services under mental health grounds, the PLDPA required that one obstetrician and two psychiatrists must jointly certify that there is a substantial risk to the woman's life. Subsequent guidance was introduced in 2014 to help healthcare professionals interpret this legislation. Criticisms of both the PLDPA and guidance document came from the UN Human Rights Committee and doctors who articulated that excessive scrutiny would be placed on women by medical professionals and that women seeking abortions due to rape, incest, fatal foetal abnormality or serious risks to health were denied (Murray, 2016). Twenty-six terminations were permitted under the PLDPA in 2015; 14 arising from a risk of physical illness, 9 from emergencies arising from physical illness and 3 arising from risk of suicide. This contrasts to the 3735 women with an Irish address who travelled to the UK for an abortion in 2014 (Aitken, Patek, & Murphy, 2017).

The Oireachtas Committee endorsed a full repeal, asking that the Eighth Amendment not be replaced, but be removed from the Irish Constitution completely (Report of the Joint Committee on the Eighth Amendment of the Constitution, 2017). Upon taking office in 2017, new Taoiseach Leo Varadkar immediately announced that a referendum on the Eighth Amendment would occur in 2018 (McNamee & Power, 2017). In 2018, an abortion rights campaign group in Ireland Together for Yes (TFY), consisting of over 70 diverse civil society bodies, including the Abortion Rights Campaign, the Irish Family Planning Association, as well as many political parties within the Irish government was founded. TFY was the national civil society campaign to remove the 8th Amendment from the Irish Constitution, strengthening their visibility and determination, and bringing the issue of abortion to the forefront of Irish politics. Several pro-life groups campaigned for the right of the unborn child, including the Pro-Life Campaign, One Day More, Youth Defence, the Love Both Project, Precious Life, Iona Institute, as well as many Pro-Life University Societies. On the 25th of May 2018 Ireland's electorate voted in a referendum to repeal the Eighth Amendment to the Constitution by 66.4% to 33.6%, allowing the Dáil to legislate for abortion provision. President Michael D. Higgins signed the Regulation of Termination of Pregnancy Bill in December 2018, legalising abortion in Ireland.

1.9 The Current Study

The current research is concerned with validating the novel implicit Functional Acquisition Speed Test (FAST) attitude measurement tool using *a known-groups paradigm* approach. That is, members of two activist groups with known and openly expressed pro-life and pro-choice attitudes will be recruited. The aim is not to measure the attitudes of group members for its own sake, but to take their overtly expressed attitudes as a given in assessing the FAST's utility in providing converging test results to predict this group membership. Therefore, my first study will employ *a known group's paradigm* approach to validate the FAST's measurement of concept associations through correlating it with group membership and self-reported attitudes (through a questionnaire and demographic survey). The study will also record how the participants voted in the 2018 referendum on the 8th amendment (a proxy for the individuals' attitudes towards abortion), and attempt to retrospectively predict the voting behaviour based on FAST scores using binary logistic regressions. As well as assessing its usefulness in distinguishing known groups and predicting their related voting behaviour, it is necessary to validate the FAST against other popular implicit tests and to begin the long effort of calibrating its scoring scale. An IAT will also be employed, using the same stimuli, for comparative purposes.

The first study will adopt a recently proposed response rate differential score (Cummins et al., 2018) as the primary metric for the FAST, for purely inductive reasons. While slope differential scores have been employed to index a FAST effect in several recently published studies, there is clear agnosticism regarding the optimal scoring method in several publications (see Cummins et al., 2018), until one is arrived at through iterative and published research. This is a process that has not occurred with either the IRAP or the IAT and has led to accusations of opaqueness in the methodology (O'Reilly et al., 2012; Ridgeway et al., 2010). Thus, while a slope differential score has great intuitive appeal, there may be alternatives that prove to more sensitive to capturing the dynamics of a complex verbal learning history. However, in a detailed data analysis methodology chapter, slope scores will be considered to allow for a systematic comparison of the relative utilities of these indices.

An area-under-curve (AUC) analyses will also be employed in the second study to assess predictive validity of the IAT, the self-reported measures, and the various FAST methodologies under investigation. In particular, the AUC method will be used to calculate the FAST's sensitivity to known group membership of participants (i.e., being a pro-life or pro-choice advocate) as well as in predicting voting behaviour, and comparing that with AUC results for both the IAT and explicit measures. It is predicted that the FAST will predict group affiliation and voting behaviour, and that cut-off RFD scores for pro-choice and pro-life groups will be identified providing a useful scale. This would provide evidence of the FASTs ability to assess and predict voting behaviour in future research for both decided and undecided voters.

Chapter 2

Assessing the Predictive Validity of a Function Acquisition Speed Test for Attitudes towards Abortion

2.1 Introduction

The current experiment sets to employ a known group's paradigm approach to validate the novel implicit Functional Acquisition Speed Test (FAST) attitude measurement tool through correlating a first-build abortion attitudes FAST test's scores with; 1) pro-life or pro-choice activist group membership, 2) self-reported levels of religiosity, conservativeness and approval of the availability of abortion services, 3) scores on an established implicit attitude test called the Implicit Association Test, 4) scores on one established scale of abortion attitudes, and responses on a bespoke set of self-report demographic and abortion attitude questions, and 5) real world referendum voting behaviour in the 2018 referendum on the 8th amendment to the Irish constitution. Members of two activist groups with known and openly expressed pro-life and pro-choice attitudes were recruited. The aim was not to measure the attitudes of group members for its own sake, but to take their overtly expressed attitudes as a given in assessing the FAST's utility to predict this group membership. FAST scores could then be validated by the extent to which scores converge with those on other tests and scales administered as well as past referendum voting behaviour. Good convergent validity would support the FAST's potential later use with individuals who may not so openly express their pro-choice or pro-life views.

The FAST will be used to help in identifying a pre-existing pro-life or pro-choice attitude based primarily upon rate-fluency differential (RFD) scores (see Cummins et al. 2018 and Results section), and to help in predicting subsequent referendum voting behaviour. These scores should co-vary with known group membership, confirmed by using an Attitudes to Abortion questionnaire and a demographic questionnaire, as well as with voting behaviour (or abstinence, where FAST scores are moderate). As well as assessing its usefulness in distinguishing known groups and predicting their related voting behaviour, it is necessary to validate the FAST against other popular implicit tests. Thus, an IAT will also be employed, using the same stimuli, for comparative (validation) purposes. The FAST and IAT will be assessed for (retrospective) predictive validity. It is predicted that both the FAST and IAT will predict group affiliation and voting behaviour.

2.2 Method

2.2.1 Participants

Sixty participants volunteered to take part in the current study, primarily consisting of Caucasian, Irish undergraduate Maynooth University students who would be considered to have strong pro-life or pro-choice views. Participants were active members of pro-choice, pro-life, feminist and/or religious societies. Participants were also recruited from other nearby Universities, and from pro-choice and pro-life organisations based in Ireland. Participants received no remuneration for their participation. All participants were adults between the ages of 19 and 35 (M = 22.47, SD = 3.34), where 24 identified as male and 36 identified as female. Of the 60 participants, 42 identified as being pro-choice and 18 identified as being pro-life. All participants had normal or corrected vision, were fluent in English, and did not suffer from any condition that would interfere with their ability to perform routine cognitive tasks. Participants had no prior experience with implicit testing, and were asked to self-exclude from the study if they had any conditions (physical or mental) that might be exacerbated as a result of this procedure. Informed consent was obtained in writing from all participants.

2.2.2 Ethical Considerations

The current experiment received approval by the Maynooth University Ethics Committee prior to beginning data collection. Identities of participants were not recorded at any stage, and participants were represented by ID number for both the computer tasks and the questionnaires. There was a small risk that the topic of abortion could lead to discomfort for the participants. However, participants were aware of the socially sensitive nature of this attitude research before volunteering to participate in the study, and a thorough screening process was implemented. Participants were informed through briefing and through an information sheet (Appendix I) given before participation of their right to self-exclude if they had any conditions, physical or mental, that may be exacerbated as a result of this procedure, or to remove themselves at any stage during the procedure. This way, participants did not have to reveal reasons for self-exclusion, but were provided with ample opportunity before any participation has taken place to do so. All participants were informed of the true nature of the study after participation through a full debriefing, and were given the opportunity to withdraw their data at that stage. In order to reduce risk arising from providing an interpretation of test scores to participants, they were advised that their data would be processed at a group level, and that individual scores could not be meaningfully interpreted.

The Attitudes to Abortion questionnaire was strictly utilized to assess general abortion attitudes, and did not ask sensitive or overly personal questions, and participants were reminded of their right to skip any questions that they did not wish to answer. All of the word stimuli used for the FAST and IAT consisted of everyday English words and phrases that were not unusually evocative and did not constitute offensive or shocking terms. Relevant contact details were provided to the participants along with the information sheet if he/she had any questions or queries about the study after participation, along with contact details for relevant crises pregnancy helplines. All participants were presented with and signed a consent form before proceeding to the first phase of the experiment (See Appendix II).

2.2.3 Apparatus

2.2.3.1 Attitudes to Abortion and Demographic Questionnaire

The current experiment aimed to compare the scores on the FAST and IAT in identifying two groups of individuals with known and different pre-existing attitudes towards abortion. Due to its reported high levels of reliability and internal consistency, Hess and Rueb's (2005) *Survey of Student Attitudes toward Abortion* was utilised for identifying and indexing the strength of these pre-existing attitudes. Hess and Rueb's (2005) survey contained 40 questions addressing abortion opinions, religious beliefs, and other factors potentially contributing to pro-choice and pro-life stances. The survey contained two sections; the first section examined opinions regarding the legal and moral status of abortion under various conditions (the Attitudes to Abortion questionnaire), and the second section addressed other factors, such as religious and political beliefs, and other relevant demographic variables (the demographic questionnaire). The Attitudes to Abortion questionnaire utilized a 5-point Likert scale (1 –Strongly Disagree, 5–Strongly Agree) to measure the level of agreement of test takers with a range of statements regarding abortion. The demographic questionnaire asked participants to report, using simple scales and check boxes, their age (to the nearest 6 months), gender, religion, religiosity, feelings towards abortion, and political leanings.

For the purpose of this study, some questions from Hess and Rueb's (2005) survey were omitted or modified, as they were either irrelevant to an Irish population or regarded as too personal for a participant to be asked to answer in this study. For example, one of the questions in Hess and Rueb's (2005) survey asked participants to tick what region of the United States they were originally from, this was modified for the current study to what region of Ireland the participant was originally from. Furthermore, the original survey asked participants to answer either yes or no to the statement "I had an abortion". This was omitted from the questionnaire presented in the current study. Therefore, the questionnaire used for the current study had a combined total of 32 questions. This had no implications on the scoring of the Attitudes to Abortion questionnaire, as the questions omitted were used as filler items in the original Hess and Rueb's (2005) survey, and not used in calculating the total score. The questionnaire used in this study also included asking how participants voted or intended to vote in the recent 2018 referendum on the 8th amendment to the Irish constitution. The modified Attitude to Abortion scale and the bespoke demographic questionnaire are presented in Appendix III.

2.2.3.2 The FAST and IAT Software Package and Stimulus Set

The experimental procedure was conducted in a small, quiet experimental room (1.5m x 1.5m approximately) in the Department of Psychology in Maynooth University. Procedures were carried out by all participants on a 15.6 inch screen (resolution of 1366 x 768 pixels) *laptop* computer. The FAST and the IAT were developed with the computer programming software *Inquisit* (version: 5.0.14). All responses consisted of pressing one of two keyboard buttons, and all responses and their timings were recorded by the software program.

The FAST and IAT tests were similar from the point of view of the participant, with the FAST consisting of two 50-trial task blocks, one consistent and one inconsistent, and the IAT consisting of two main 40-trial task blocks (consistent and inconsistent), with two shorter 20-trial task blocks (consistent and inconsistent) presented before each of the main task blocks as practice, and three sorting 20-trial training blocks (7 blocks in total). Both tests worked by exposing participants to brief blocks of learning in a computer-based task. Participants were required to simply learn whether to press the left key or the right key on the computer keyboard depending upon which of four category exemplar words were displayed on screen. The current abortion attitudes FAST and IAT were designed to assess fluency in associating terms related to abortion with positive terms compared to negative terms. Participants who already associated, for example, abortion related terms with negative terms (i.e., those who identify as being pro-life) would find it difficult to learn the task if abortion-related and positive terms required the same shared directional response (i.e., it is difficult to learn common responses to incompatible word stimuli). Whereas participants who would easily associate abortion related terms with positive terms (i.e., those who identify as being pro-choice) would find it difficult to learn the task if abortion-related and negative terms required the same directional response.

The FAST and IAT employed verbal stimuli as exemplars of the categories of interest, with four exemplars per stimulus category. The real word stimuli consisted of 16 words in total, 8 related to the topic of abortion, and 8 consisted of positive and negative terms listed below (Table 1). For the FAST and IAT to be conceptually valid and reliable, it is crucial to ensure that the stimuli selected are optimally representative of the nominal categories according to which participants will categorise target stimuli (Gawronski & Payne, 2010). Stimuli for the target categories were chosen carefully based on their perceived semantic representativeness of the categories of interest. For the stimuli representing the categories 'Anti-Abortion' and 'Pro-Abortion', descriptive words were carefully selected by the author from those used by pro-life and pro-choice campaigners that were non-intentionally offensive while still being representative of that category. For the 'Positive' and 'Negative' attribute categories, unambiguously positive and negative 'feeling' words were selected from the Project Implicit website (Retrieved from implicit.harvard.edu), which were also used in a study conducted by Nosek, Banaji and Greenwald in 2002. The word stimuli presented during the FAST were identical to those presented during the IAT (See section 2.3.3 for procedural details).

Table 1: Verbal stimuli employed in the FAST and IAT

Positive	Negative	Anti-Abortion	Pro-Abortion
Happy	Nasty	Continue Pregnancy	Terminate Pregnancy
Wonderful	Awful	Do Carry to Term	Don't Carry to Term
Glorious	Terrible	Keep Foetus	Abort Foetus
Peace	Horrible	Pro-Life	Pro-Choice

2.3 Procedure

2.3.1 Informed Consent

Prior to any involvement all participants were informed that they are taking part in an experiment assessing a new measure of attitudes, specifically focusing on abortion attitudes, that will require them to fill out a short abortion attitude and demographic questionnaire, and to complete two computer-based learning tasks, during which a variety of stimuli related to the topic of abortion will be presented to them on the screen. Participants were presented with an information sheet detailing the format of the task (i.e., that a word would be presented and that a response would be required in the form of a button press, as instructed by the computer program), and of the likely duration of the experiment (approximately 30 minutes). Participants were then required to sign a consent form and were informed that they were free to withdraw from participating in the research even after having signed the form.

2.3.2 Attitudes to Abortion and Demographic Questionnaire

Upon supplying written consent, participants were asked to complete the Attitudes to Abortion questionnaire and the demographic questionnaire. As participants were aware of the general nature of the test, administrating the self-report measure before or after the implicit measures would exert little influence over the outcomes of the test according to Nosek, Greenwald and Banaji (2005). However, performing the implicit measure first might cause reactivity that is reflected in subsequent responses to an explicit questionnaire. On the other hand, implicit measures are less likely to be consciously controlled or to suffer from reactive effects (Rudman, 2011). Because the attitude being assessed here was a sensitive one, it was considered prudent to assess explicit attitudes first.

Participants were left on their own in the experimental room until they had completed the explicit measure so as not to feel pressured to answer all of the questions and to mitigate the effects of social desirability bias (Nederhof, 1985). Upon completion, the participants turned in their surveys face down on the table before continuing the experiment. The standard questionnaire required approximately 10 minutes to complete, and the demographic questionnaire required approximately 5 minutes to complete.

2.3.3 Implicit Measures

Following the participants completion of the explicit measures, both of the computer-based implicit tests were administrated. This involved firstly ensuring that the participant was comfortable, that the environment was quiet, and that there were limited external distraction. The participant was then directed to the laptop computer where they began with either the IAT or the FAST, the order of which was counter-balanced across participants. Participants were also given the option to take a short break in between the two tasks.

2.3.3.1 Function Acquisition Speed Test

The FAST involved the formation of functional responses (i.e., pressing the “Z” or “M” key) for particular stimuli presented in sequence on the computer screen. The FAST consisted of two 50-trial learning task blocks, one being the *consistent* block, and the other being the *inconsistent* block. Within each block, relations were established between two stimuli or sets thereof via explicitly trained shared response functions (i.e., a common key-press response). The order of block presentation on each FAST was randomized across participants based on a random number generator built into the software. For participants who identified themselves as being pro-choice, the consistent block involved the reinforcement of common responses to stimuli from both positive and pro-choice stimuli, as well as stimuli from pro-life and negative verbal categories (see Fig. 1 for a schematic of the functions for the stimuli in each of the two FAST blocks). The consistent block involved the reinforcement of responses consistent with verbal category compatibilities that are reflective of certain beliefs, whereas the inconsistent block involved the opposite response patterns being reinforced. The consistent and inconsistent block is defined arbitrarily here from the point of view of a pro-choicer, that is, Anti-Abortion/Positive and Pro-Abortion/Negative shared response functions for the consistent block, and Anti-Abortion/Negative and Pro-Abortion/Positive shared response functions for the inconsistent block (see Figure 1). From some participant's point of view, the inconsistent block will yield greater response fluency and the consistent block reduced fluency. This difference in performance across participants with different attitudes to abortion is precisely what will underlie the FAST's ability to identify these attitudes.

The FAST commenced with a series of instructions on the screen informing the participants of the general nature of the task. The following text was presented to participants at the beginning of each block:

> In this task, you will need to use the 'Z' and 'M' keys on the computer keyboard. Please locate these keys now. When you next press the 'M' key, positive and negative words, and words related to anti-abortion and pro-abortion attitudes, will begin to flash up on the screen, one at a time. The goal of the task is to learn to press either the 'Z' or the 'M' key depending on what word appears on the screen, based on the feedback that you are given. Try to respond AS QUICKLY AND AS ACCURATELY AS POSSIBLE. When you're ready, press the 'M' key to begin.

These instructions remained on-screen until participants felt confident they understood the task, after which they pressed the “M” key on the keyboard to begin, as instructed. Following the key press, the first inter-trial interval (ITI) of 500ms was presented (i.e., a blank white screen) before participants began a randomized block. A trial began with the presentation of one stimulus from the four sets of word stimuli (i.e., positive stimuli, negative stimuli, anti-abortion stimuli, or pro-abortion stimuli). Stimuli were presented in black 48-type Ariel font in the centre of a white background screen. Each trial was limited to a 3000ms response window in which participants had to respond by pressing either the “Z” or “M” key. On every trial upon the production of a response, corrective feedback was presented whereby the participant was informed whether their response was correct or incorrect by the immediate display of the words ‘Correct’ or ‘Wrong’ in red 48-point font in the centre of the screen for 1500ms. If no response was emitted within the 3000ms response window (i.e., no key was pressed), a time-out response (‘Wrong’) was presented in the centre of the screen. The screen would then clear and the next trial would begin after the 500ms ITI. Thus, the reinforcement contingency employed could be described as a FR1 with a 3s limited hold. Stimuli were presented in a quasi-random order, with each stimulus presented once in a given cycle of four trials. In each block, participants were required to learn, through corrective feedback, a common response to members of two sets of stimuli and a different response to the members of another two sets.

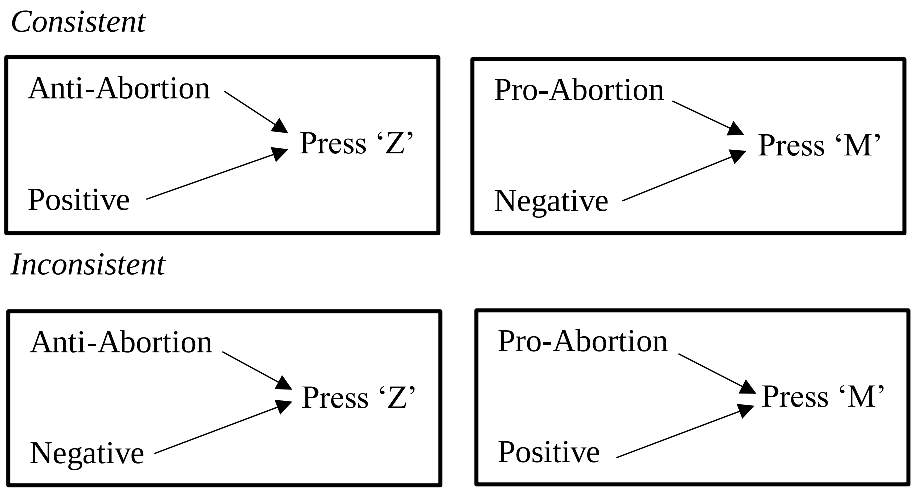

Figure 1: The trained response functions for FAST stimuli across both the (arbitrarily defined) consistent and inconsistent blocks (see section 2.3.3.1).

After the second and final block of the task, the computer screen displayed a message informing the participant that this part of experiment is now concluded. For all participants the FAST took no more than 5 minutes to complete. For participants who began with the FAST after completing the explicit measure, they were given the option to take a short break or to continue on and begin the IAT. Participants who finished with the FAST were then debriefed.

2.3.3.2 The Implicit Association Test

The IAT was ostensibly very like the FAST from the participant's perspective, but with some key and well noted differences (see Roche et al., 2015). During the IAT participants are asked to quickly and accurately categorize pairs of stimuli in terms of spatial responses (i.e., by pressing the "E" key for exemplars from either of two verbal categories and the "I" key for exemplars of the other two verbal categories).

Stimuli were presented on the centre of a black background screen and remained until a correct response was given (i.e., no time limit). Positive and negative evaluative stimuli were presented in green 28-type Arial font and stimulus exemplars from the anti-abortion and pro-abortion categories were presented in white 28-type Arial font. All stimuli were identical to those used in the FAST. If the correct key was pressed, a new trial would begin. If the incorrect key was pressed a red X appeared below the stimuli, indicating that the participant needed to press the correct key before the next trial could begin (i.e., trial remained in progress and response time was recorded from trial onset to first *correct* response). Participants were reminded of this requirement by a message that was displayed at the bottom of the screen for the duration of all training and test blocks in grey 18-type Arial font; "If you make an error, a red X will appear. Press the other key to continue."

The current IAT consisted of seven blocks as per the original IAT (Greenwald et al, 1998), with instructions presented before each block in white Arial font and summary instructions remaining on screen at all times in the top left and top right of the screen (see Fig. X below). The first two blocks consisted of the target category (pro and anti-abortion) sorting task (20-trails) and the attribute stimuli sorting task (20-trials). The target category sorting task block entailed participants sorting abortion related terms into one of two categories. Specifically, if the category "Anti-Abortion" was labelled on the top left of the screen, and one of the anti-abortion word stimuli appeared on the centre of the screen (e.g., *continue pregnancy*), participants would be required to press the "E" key (i.e., left). In the attribute sorting task block the participants were required to sort words related to the attribute (e.g., positive and negative). Similarly, if the category "Positive" was labelled on the top left of the screen, and a positive word evaluative word stimulus appeared on the screen (e.g., happy), participants would be expected to again press the "E" key.

The following instructions, presented along with the categories for that particular on the top left and right corners, is an example of what was presented to participants before each of these first two blocks:

> Put your left finger on the ‘E’ response key for items that belong to the category ‘Anti-Abortion’. Put your right finger on the ‘I’ response key for items that belong to the category ‘Pro-Abortion’. Items will appear one-by-one in the middle of the screen. If you make an error, a red X will appear – to continue, press the other response key. Go as fast as you can while making as few errors as possible. Press the spacebar to begin.

The third and fourth blocks of the IAT involved presenting stimuli from all four categories and participants were required to respond spatially to both target and attribute words presented in an inter-mixed way within a single block (see Figure 2). The third block was a brief version of 20 trials, where response requirement may have been more or less relationally consistent with the verbal history of the participant, intended to serve as a practice block before the critical task block to follow. The fourth block (i.e. the critical task block) consisted of 40 trials, where the response requirements remained the same as the third block. The following instructions are an example of what was presented to participants before both the third and fourth block:

> Press the left ‘E’ key for ‘Anti-Abortion’ and ‘Positive’. Press the right ‘I’ key for ‘Pro-Abortion’ and ‘Negative’. Each item belongs to only one category. If you make an error, a red X will appear – to continue, press the other response key. Go as fast as you can while making as few errors as possible. Press the spacebar to begin.

In the fifth block the target category sorting task block was presented again but with the targets switching sides, and the evaluative categories remaining in place. The following instructions is an example of what was displayed to participants before the fifth block, where the block began once they pressed the space bar:

> Attention! The labels have changed sides. Press the left ‘E’ key for ‘Pro-Abortion’. Press the right ‘I’ key for ‘Anti-Abortion’. Go as fast as you can while making as few errors as possible. Press the spacebar to begin.

In the sixth and seventh test blocks of the IAT, the categories were combined in a way that was orthogonal to that used in the third and fourth blocks. For example, if the categories Anti-Abortion/Positive and Pro-Abortion/Negative previously shared a positional response, now Pro-Abortion/Positive and Anti-Abortion/Negative share a positional response. The sixth block was a brief version of 20 trials, intended to serve as a practice block before the critical task block to follow. The fourth block (i.e. the critical task block) consisted of 40 trials, where the response requirements remained the same as the sixth block. The order in which blocks three and four and blocks six and seven were presented was counterbalanced per participant, where participants either began with two consistent blocks (a practice task block and critical task block) or two inconsistent blocks.

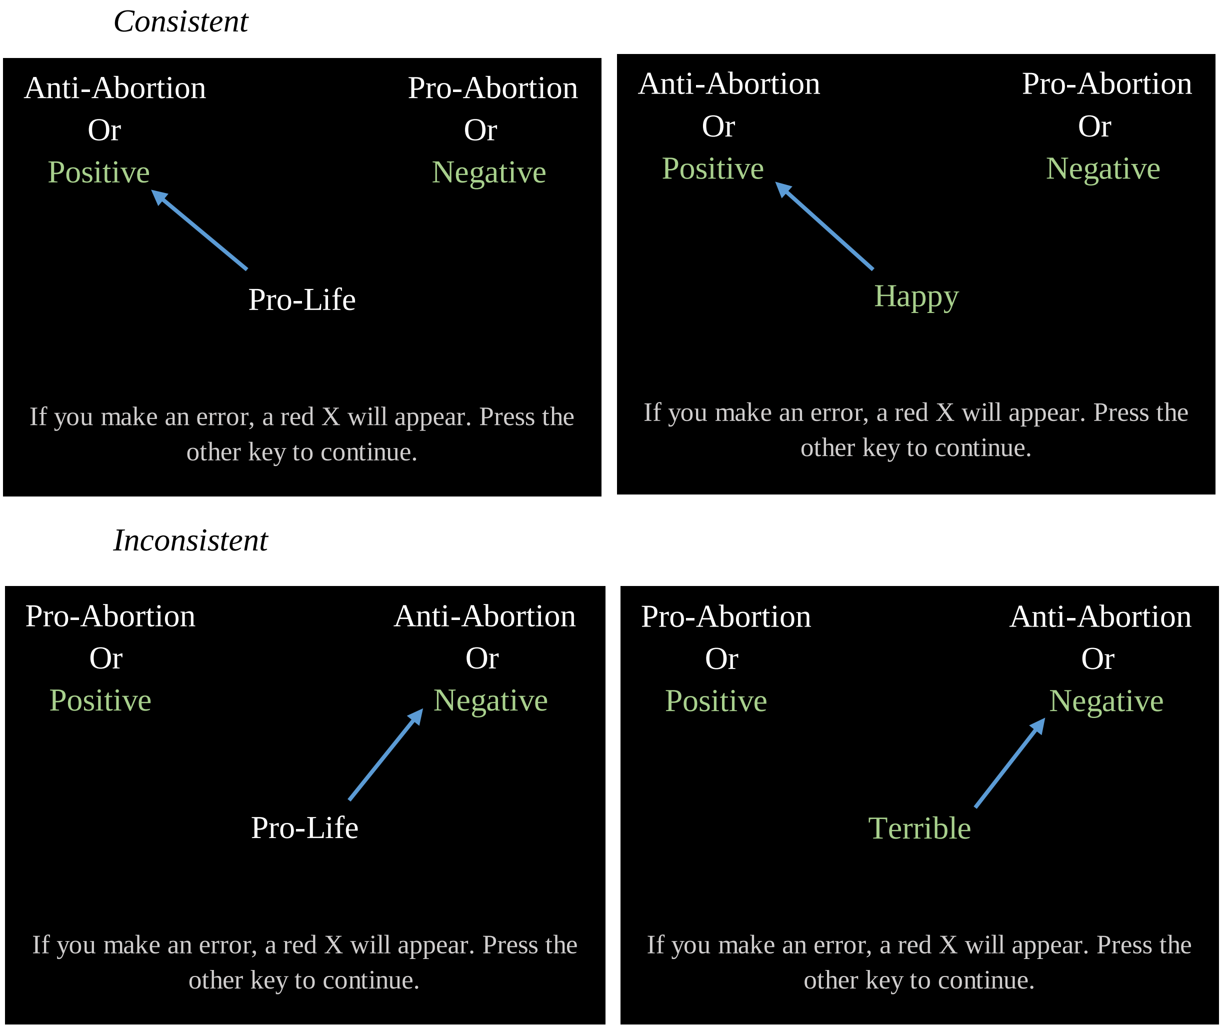

Figure 2: Representation of consistent and inconsistent trials. The superimposed arrows indicate what was defined here as a correct response for that particular block (left arrow for “E” key and right arrow for “I” key; the arrows are for illustration purposes only and did not appear on screen).

2.3.4 Debriefing

Participants were fully debriefed regarding the nature of the experiment immediately upon completion. Data was available on request, but participants were informed that data could not be interpreted. Participants were assured that all data analyses were to be conducted at a group level, where no personally identifying information relating to them individually would be stored at any time.

2.4 Results

2.4.1 Attitudes to Abortion and Demographic Questionnaire

2.4.1.1 Scoring

To analyse the data, all Likert responses to the pro-life slanted questions (1, 8, 15, and 18) were reverse scaled. These revised scores and the original scores for items 1, 2, 4, 7-15, and 18 combined to form a participant's total score ranging from 13-65. Hence, a low score indicated pro-life attitudes, with the lowest possible score being 13. A high score indicated a pro-choice attitude, with the highest score being 65. Several items were essentially filler items. One of these filler items included asking participants, on a scale of 1-5, how strongly they agreed or disagreed to the statements "I am pro-choice" and "I am pro-life". To determine if the total score construct was valid, correlations were run between the total score of the questionnaire and the participant's 1-5 scale scores regarding their prevailing pro-choice and pro-life attitudes. Both sets of correlations were significant ($p < .001$), as investigated using Pearson product-moment correlation coefficient. Specifically, the correlation between the scale scores obtained from the statement "I am pro-choice" and the total Attitudes to Abortion total score was $r = .794$; whereas, the correlation between the scale scores obtained from the statement "I am pro-life" and the total Attitudes to Abortion scale score was $r = - .919$. These correlations suggest that the Attitudes to Abortion questionnaires total score construct is valid, where scores are accurately depicting self-reported known group views.

Total scores on the Attitudes to Abortion questionnaire for those who identified as pro-life ranged from 13 to 60 (M = 27.06, SD = 11.5), where scores tended to cluster around the lower end of the scale with mostly extremes and one outlier (participant 27; see Table 2). Total scores for those who identified as pro-choice ranged from 25 to 64 (M = 48.24, SD =

8.3), with scores clustering around the higher end of the scale and with mostly extremes (Table 2). With all data combined, 41.88 was the overall average total score (*SD* = 13.48), with total scores of 13 to 24 considered to be the extreme pro-life score range, in that these scores were 100% consistent with participants who strongly agreed to being pro-life and strongly disagreed to being pro-choice. Total scores of 43 to 64 were considered to be the extreme pro-choice score range, in that these scores were 100% consistent with participants who strongly agreed to being pro-choice and strongly disagreed to being pro-life.

Table 2: Total Score on Attitudes to Abortion Questionnaire for Group One (Pro-Choice) and Group Two (Pro-Life)

Participant	Group	Total Score	Participant	Group	Total Score
1	2	27	**31**	1	25
2	2	24	**32**	1	38
3	2	27	**33**	1	36
4	1	39	**34**	1	56
5	1	52	**35**	1	59
6	2	36	**36**	1	41
7	1	51	**37**	1	54
8	1	49	**38**	1	54
9	2	13	**39**	1	51
10	1	57	**40**	1	40
11	1	60	**41**	2	38
12	1	64	**42**	1	48
13	2	43	**43**	2	27
14	1	53	**44**	1	43
15	1	45	**45**	1	48
16	2	24	**46**	1	47
17	2	22	**47**	1	42
18	2	18	**48**	1	55
19	1	41	**49**	1	47
20	1	41	**50**	1	56
21	1	28	**51**	1	50
22	2	19	**52**	1	52
23	1	60	**53**	1	53
24	1	47	**54**	2	21
25	1	50	**55**	1	55
26	1	46	**56**	1	56
27	2	60	**57**	1	45
28	2	15	**58**	2	21
29	1	51	**59**	1	41
30	2	19	**60**	2	33

2.4.1.2 Group Affiliation

An independent sample t-test was conducted to compare the Attitudes to Abortion questionnaire scores across self-identified pro-choice and pro-life participants. Recall that in the demographic questionnaire, participants rated from 1 to 5 in relation to whether they were pro-life or pro-choice. All participants provided a response, with 41 classifying themselves as pro-choice (i.e., a 4 or 5 rating), 17 reporting to being pro-life (i.e., a 4 or 5 rating), and two reporting that they were in the middle (i.e., rated 3). The independent sample t-test found a significant difference in the questionnaires total score between pro-life and pro-choice groups ($t(58) = 8.042$, $p < 0.001$). The magnitude of the differences in the means (mean difference = 21.18, 95% CI: 15.9 to 26.45) was very large (eta squared 0.54). Thus, the attitudes to abortion questionnaire is showing strong convergent validity with self-reports on attitudes.

2.4.1.3 Voting Behaviour

An independent sample t-test was conducted to compare the Attitude to Abortion scale scores of participants who voted differently in the abortion referendum. In the demographic questionnaire, participants were asked how they voted in the past abortion referendum, where they were given the option of three boxes to tick; either yes, no or undecided. All participants answered, with 47 reporting a Yes vote (pro-abortion choice), 13 to a No vote (anti-abortion), and none were undecided. The independent sample t-test found a significant difference in attitude scale scores across yes and no voters ($t(58) = -9.47$, $p < 0.001$). The magnitude of the differences in the means (mean difference = -25.28, 95% CI: -30.6 to -19.9) was very large (eta squared 0.6). Thus, the attitude scale scores displayed good convergent validity with self-reports of voting behaviour.

2.4.2 Function Acquisition Speed Test

2.4.2.1 Scoring

Scores on the FAST were calculated automatically by the Inquisit software. The primary index used here as a FAST score was suggested by Cummins et al. (2018). It is based on a Skinnerian interpretation of the notion of 'fluency' (Skinner, 1937), as well as elements of assessment used within the precision teaching literature from applied behaviour analysis (Binder, 1996; Calkin, 2005). This method involves calculating two rates of responding for both the consistent and inconsistent blocks: the rate of correct responses per minute (CRPM), and the rate of incorrect responses per minute (IRPM). These rates were found by dividing the number of correct (for CRPM) or incorrect (for IRPM) responses by the total time the participant took to complete the FAST block, and multiplying this value by 60,000 (i.e., to convert the value to 'per minute'). Both inter-trial interval and feedback presentation times are included within the total time taken to complete the procedure. The rationale for their inclusion is based on the fact that the measure is attempting to describe the true rate of the behaviour within the procedure. Latent periods (such as I and feedback presentation) in which the participant cannot respond necessarily constrain the overall rate of responding in a given block, and so are included. The rate of IRPM is subtracted from the rate of CRPM to provide a composite index of learning rate, and this provides a response rate differential score (RRD; or fluency) for the block.

To produce a FAST score, the response rate differential for the inconsistent block is subtracted from the response rate differential of the consistent block, and the difference in this value signifies the difference in the fluency of responding for that participant's FAST, referred to as the rate-fluency differential (RFD) score (see equation 1 below for an expression of this formula. A higher RFD score indicates that participants were more fluent on pro-abortion positive blocks and anti-abortion negative blocks (consistent with a pro-choice stance), while a lower RFD score indicates that participants were more fluent on anti-abortion positive blocks and pro-abortion negative blocks (consistent with a pro-life stance). This can be seen in Figure 3 showing the trend of the RFD scores across the two groups.

$$RFD = \left(\left(\frac{TC_C - TI_C}{TT_C}\right)\right) - \left(\left(\frac{TC_I - TI_I}{TT_I}\right)\right) \text{ x } 60000$$

Equation 1: Formula to calculate the rate-fluency differential score. TC = Total Correct, TI = Total Incorrect, TT = Total Time, and where a 'C' subscript indicates the consistent block, and an 'I' subscript indicates the inconsistent block.

2.4.2.2 Group Affiliation

Total RFD scores for those who identified as pro-life ranged from -16.82 to 11.92 (*M* = -6.84, *SD* = 6.88), where scores tended to cluster around the lower end of the scale with a mode of -16.82, and one outlier (participant 30). Total scores for those who identified as pro-choice ranged from -15.48 to 25.02 (*M* = -1.66, *SD* = 7.39), with scores clustering around the higher end of the scale, with two very extreme scores (participants 20 and 34; see Appendix IV), and with one outlier (participant 36). Both groups tended to therefore have an anti-abortion attitude insofar as abortion related terms were more easily established as common members of a functional response class with negative terms than with positive terms for both groups, on average (See Figure 3). That is, both groups are showing an overall FAST effect in the same direction. However, they are still distinguishable in this regard in terms of the extent of the negative effect.

An independent samples t-test was conducted to compare FAST RFD scores of pro-life and pro-choice groups. A significant difference was found in RFD scores between pro-life and pro-choice members (t (58) = 2.53, p = .014). The magnitude of the differences in the means was relatively large, where eta squared = .1.

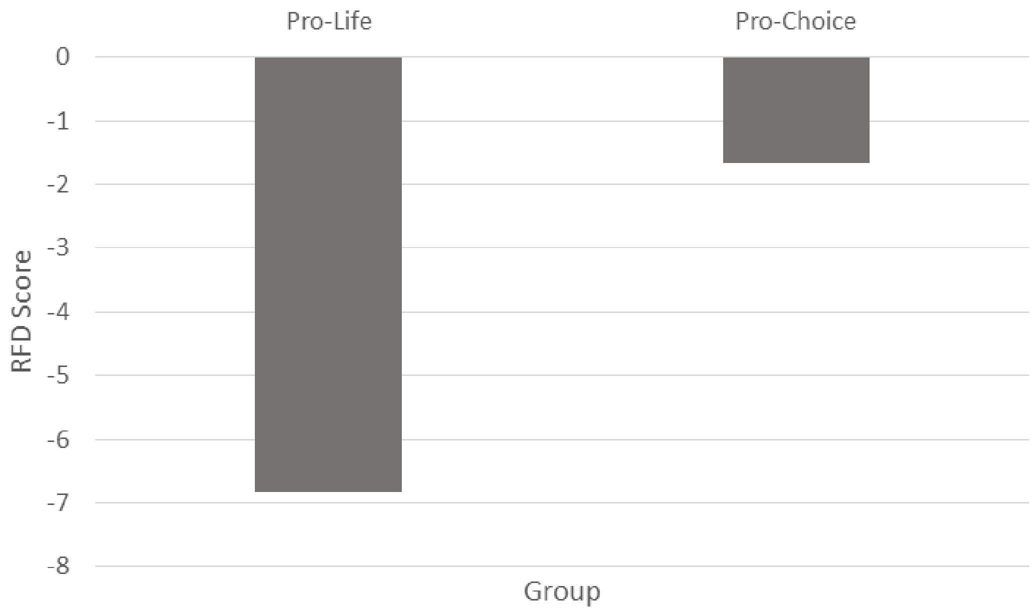

Figure 3: Figure showing the difference of RFD scores (mean) as a function of the two group affiliations.

2.4.2.3 Voting Behaviour

An independent samples t-test was also conducted to compare the RFD scores of participants who voted yes and no in the 2018 Irish referendum on abortion. There was a significant difference in RFD scores between yes voters and no voters (t (58) = -3, p = .004), where the magnitude of the differences in the means was large (eta squared = 0.13).

2.4.2.4 Relationship between FAST RFD scores and Explicit Scores

The relationship between the FAST RFD scores and the Attitudes to Abortion questionnaire total scores were investigated using a Pearson product-moment correlation. There was a significantly moderate, positive correlation between the two variables, r = .405, p < .001, with higher RFD scores in the FAST (higher scores indicative of pro-choice attitudes) associated with higher total scores in the explicit measure (also indicative of pro-choice attitudes). This indicates that scores obtained in the FAST reflected those obtained in the Attitudes to Abortion questionnaire. Thus, in effect, correlations were in the predicted direction.

2.4.3 Implicit Association Test

2.4.3.1 Scoring

The IAT score is based on how long it takes a person, on average, to produce correct positional responses to the words in the third and fourth block of the IAT versus the sixth and seventh block of the IAT. The strength of an association between concepts is measured by the standardized mean difference score of the 'hypothesis-inconsistent' pairings and 'hypothesis-consistent' pairings (D-score; Greenwald, Nosek, & Banaji, 2003). In general, the higher the D-score the stronger is the association between the 'hypothesis-consistent' pairings. In contrast, negative D-scores suggest a stronger association between the 'hypothesis-inconsistent' pairings (decided by the researcher). The Inquisit software calculated d-scores automatically using the scoring algorithm recommended by Greenwald and colleagues (2003).

The IAT effect (D-score measures) is found by computing the difference between the mean response times during the consistent and inconsistent blocks, and then dividing this difference by the standard deviation of all the individual response times during those main test blocks. That is, the standard deviation for all of the trials at blocks 3 and 6, and the standard deviation for all of the trials at block 4 and 7 were calculated (as all other blocks were practice and training blocks). The difference scores between Block 3 and Block 6 as well as the difference scores between block 4 and block 7 were then computed. Both the response latencies and the error rates are taken into account in this algorithm. These two sets of D-scores were divided by their corresponding standard deviations. Finally, the averages of the two sets of scores were computed. A positive score indicated that participants completed blocks where anti-abortion and positive, and pro-abortion and negative were established as common members of the functional positional response class more quickly, indicating that they were compatible prior to experimentation (i.e., participant has a pro-life stance); whereas, a negative score indicated that participants completed blocks where pro-abortion and positive, and pro-life and negative stimuli were more easily formed into distinct functional response classes (i.e., participant has a pro-choice stance).

2.4.3.2 Group Affiliation

Total D-scores for those who identified as pro-life ranged from -0.26 to 1.45 (*M* = .721, *SD* = .433), where scores tended to cluster around the higher end of the scale with one outlier (participant 27). Total scores for those who identified as pro-choice ranged from -.64 to 1.07 (*M* = -.017, *SD* = .409), with scores clustering around the lower end of the scale, with a mode of -.639.

An independent samples t-test was conducted to compare D-scores of pro-life and pro-choice groups. A significant difference was found in D-scores of pro-life and pro-choice members (t (58) = -6.3, p = .000). The magnitude of the differences in the means was large, where 40% of the variance in RFD scores is explained by group membership (eta squared = .4). Pro-Life groups tended to therefore have an anti-abortion attitude insofar as abortion related terms were more easily established as common members of a functional response class with negative terms than with positive terms on average. Pro-choice groups tended to have a pro-abortion attitude insofar as abortion related terms were more easily established as common members of a functional response class with positive terms than with negative terms on average (See Figure 4). That is, both groups are showing an overall IAT effect in different directions.

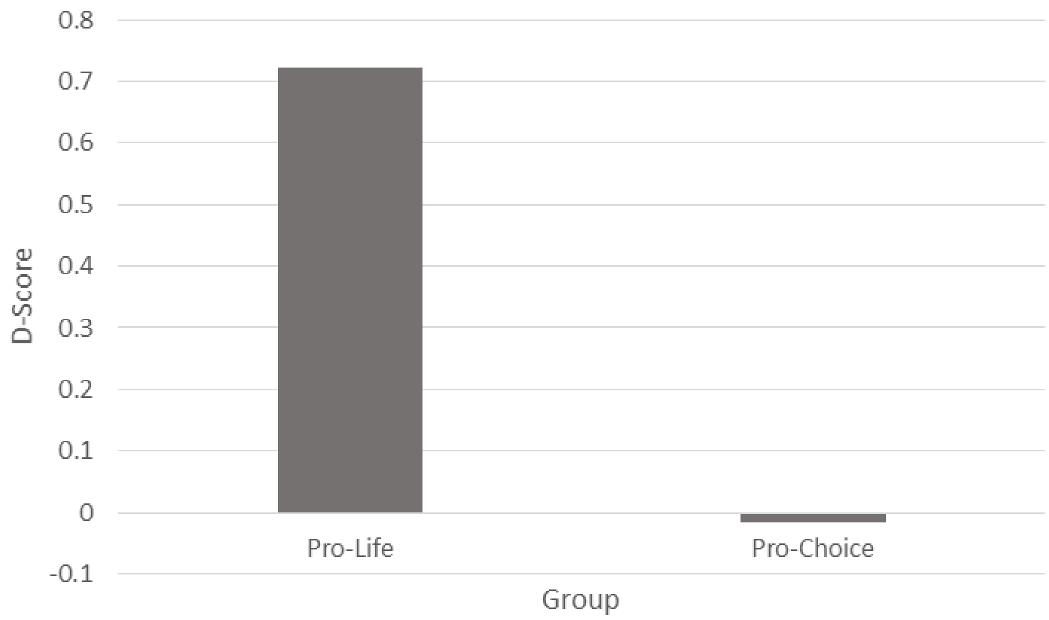

Figure 4: Figure showing the trend of D-scores (mean) as a function of the two group affiliations.

2.4.3.3 Voting Behaviour

An independent samples t-test was also conducted to compare D-scores of participants who voted yes or no in the 2018 Irish referendum on abortion. There was also a significant difference in D-scores between yes voters and no voters ($t (58) = 6.9$, $p = .000$), where the magnitude of the differences in the means was very large (eta squared = 0.45).

2.4.3.4 Relationship between IAT D-scores and Explicit Scores

The relationship between IAT D-scores and the Attitudes to Abortion scale score was investigated using Pearson product-moment correlation coefficient. There was a strong, negative correlation between the two variables, $r = -.647$, $p < .001$, with high scores in the IAT (indicative of pro-life attitudes) associated with lower total scores in the Attitudes to

Abortion scale (also indicative of pro-life attitudes). This indicates that scores obtained in IAT reflected those obtained in Attitudes to Abortion questionnaire. Thus, in effect, correlations were in the predicted direction.

2.4.4 Predicting Voting Behaviour

Participants who had a FAST RFD score and an IAT D score in the same direction as their own voting behaviour and reported attitude on abortion was tallied (for validity purposes). The percentage of accuracy in detecting the leaning of the participant based on FAST RFD scores was 35.7% for those who identified as pro-choice (i.e. scored above 0) and 88.8% for those who identified as pro-life (i.e. scored below 0). The percentage of accuracy in detecting the leaning of the participant based on IAT D scores was 54.7% for those who identified as pro-choice (i.e. scored below 0) and 94.4% for those who identified as pro-life (i.e. scored above 0). Of those who identified as pro-choice, 97.6% reported that they voted Yes in the abortion referendum (pro-abortion choice), and of those who identified as pro-life, 66.6% reported that they voted No in the abortion referendum (pro-life choice).

2.4.4.1 Binary Logistic Regression

A logistic regression was performed to ascertain the relationship between the FAST RFD score, the IAT D-score, and the attitudes to abortion scale score, on voting behaviour (the likelihood that participants voted yes or no in the referendum). Aside from the suitability of a logistic regression over a multiple linear regression for this data set in which the dependent variable (voting) is binary, it also helps to avoid confounding effects by analysing the association of all variables simultaneously rather than sequentially.

Table 2 below present's results from binary logistic regression models for the prediction of voting behaviour among yes and no voters. Implicit RFD (model 1a) scores and explicit (model 1b) scores both predicted voting behaviour, adding significantly to the model/prediction. In fact, the FAST allowed for the correct classification (% CCC) of 83.3 percent of yes or no voters. However, the explicit measure outperformed the implicit measure, correctly classifying 93.3% percent of cases. Controlling for the explicit measure (model 2) only slightly increased Nagelkerke's R^2 by 2 percentage points while the %CCC increased by 1.7 percentage points. However, looking at the incremental validity of the implicit measure, the FAST did not significantly predict voting behaviour for the abortion referendum beyond the explicit measure.

Table 2. Results of logistic regression for prediction of voting behaviour for RFD and explicit scores.

Step	Variable	B	*SE*	Wald	p	Exp(B)	Nagel-kerke's R²	%CCC
1a	Constant	2.21	.553	15.992	<.001	9.113	.24	83.3%
	RFD	.171	.062	7.676	.006	1.186		
1b	Constant	-9	2.9	9.677	.002	.000	.826	93.3%
	Explicit	.316	.101	9.857	.002	1.372		
2	Constant	-8.6	3.04	7.97	.005	.000	.846	95%
	RFD	.134	.101	1.755	.185	1.143		
	Explicit	.331	.115	8.313	.004	1.393		

Notes: B = regression weight; SE = standard error of the regression weight; Wald = Wald criterion; Exp (B) = Odds ratio, the relative amount by which the odds increase (Exp (B) > 1.0) or decrease (Exp (B) < 1.0) when the value of the predictor is increased by 1 unit; CCC = correctly classified cases.

Table 3 depicts the results for the prediction of voting behaviour, showing that the implicit-only IAT (model 1a) scores added significantly to the model/prediction. The implicit IAT measure allowed for the correct classification of 90 percent of cases compared to the higher 93.3 percent of cases that were correctly predicted by explicit attitudes (model 1b) measured by the abortion attitudes survey. Adding the IAT to the explicit measure (model 2) led to an increase in Nagelkerke's R^2 by 4.2 percentage points and an increase in the %CCC of 3.4 percentage points. In line with these descriptive results, IAT scores did not significantly improve model fit for predicting voting behaviour once attitude scale scores were controlled for.

Table 3. Results of logistic regression for prediction of voting behaviour for IAT and explicit scores.

Step	Variable	B	*SE*	Wald	p	Exp(B)	Nagel-kerke's R^2	%CCC
1a	Constant	4.08	1.16	12.424	<.001	58.939	.66	90%
	IAT	-5.7	1.72	11.133	.001	.003		
1b	Constant	-9	2.9	9.677	.002	.000	.826	93.3%
	Explicit	.316	.101	9.857	.002	1.372		
2	Constant	-5.5	3.54	2.401	.121	.004	.868	96.7%
	IAT	-4.4	2.86	2.359	.125	.012		
	Explicit	.297	.145	4.199	.04	1.345		

2.5 Discussion

The current experiment sought to develop and assess the FAST's utility in the real world context of predicting self-reported abortion attitudes and relevant voting behaviour. From using a known-groups paradigm approach, consisting of pro-choice and pro-life groups, it was expected that explicit self-report measures would correlate highly with both group affiliation and voting behaviour. Once confirming that participants had polarised pro-choice and pro-life views using these explicit measures, the FAST implicit measure could then be used to retrospectively predict group membership and voting behaviour, and its utility in doing so could be compared to that of the IAT.

Significant correlations were found between the total score of the Attitudes to Abortion questionnaire and the responses participants gave as to whether they were pro-choice (r = .794) or pro-life (r = -.919). These correlations suggest that the explicit scoring construct is valid. An independent samples t-test showed that the explicit scores between pro-life and pro-choice groups were significantly different ($t = 8.042$, $p < 0.001$). Not only did differences exist between the two groups, the overall ratings were very much polarised. 58.8% of participants who identified as pro-life on the self-report survey had extreme pro-life scores (ranging from 13 to 24) and 75.6% of participants who identified as pro-choice had extreme pro-choice scores (ranging from 43 to 64). This signified that this sample was suitable for assessing the validity of the FAST in a known group's paradigm.

In identifying polarising groups with strong and somewhat extreme pro-choice and pro-life views through self-reported explicit measures, this may undermine the need for implicit tests in this context. While self-reports may function at least as well or better in predicting behaviours consistent with attitudes towards abortion, the first and foremost concern is to assess the FAST where behaviour was known. It was also important that long standing unambiguous attitudes were known before being assessed through the FAST, and that an IAT was also used to compare its validity in assessing these attitudes.

For the FAST, both mean RFD scores for the pro-choice group and the pro-life group were below zero (as seen in Figure 3), indicating essentially that all participants are against abortion. This indicates that those in the pro-choice group are not pro-abortion, as even they showed some negative effect towards abortion. However, the pro-life group were even more against abortion. So, while the configuration of the FAST may have been such that it better suited an attitude phenomenon in which people vary equally around a zero point, it was nevertheless apparently still sensitive to varying degrees of aversion to abortion across the two self-identified groups.

A significant difference was found in RFD scores between pro-life and pro-choice members ($t(58) = 2.53$, $p = .014$). Moreover, the FAST was found to significantly predict group affiliation and subsequent voting behaviour in a binary logistic regression, with an impressive 83.3% correct classification. However, it was outperformed by the IAT in this regard. Although it was found that the IAT may be a better predictor of voting behaviour than the FAST, it could be suggested that the two measures cannot be compared meaningfully as they use very different scoring. It is important to understand that of course the metrics for scoring are part and parcel of the measure and in that regard no further comparison is necessary. However, from a behavioural and functional point of view, and as argued in several publications, the underlying behavioural processes of both tests is surely the same, although several aspects of the IAT obscure or conflate these processes with other unknown processes (e.g., RTs over 10s produced within potentially infinite response windows are post-hoc eliminated as if they did not occur, but the functional effect of these very slow responses on subsequent trial performance is complex and unknown). To more meaningfully compare the two tests, therefore, they could be scored in the same way (i.e. either using a D-scores algorithm for both measure or using RFD scores for both measures).

In addition, the FAST scores examined here were calculated using the novel RFD score method (Cummins et al. 2018). It is not yet known if this measure has more utility than the previous learning slope differential score method (Cartwright et al., 2016, Cummins et al., 2018) and examining this is beyond the scope of the current chapter. In the next chapter, therefore, an extended exploratory analysis of scoring systems will be employed across the IAT and FAST. This will allow for thorough analyses of the effects considered here but using various scoring metrics across the IAT and FAST. In addition, it will allow for a more in-depth analysis of the predictive validity of the FAST and IAT using more advanced statistical methods.

Chapter 3

Function Acquisition Speed Test Scoring Methodologies and their Relation to Predictive Validity

3.1 Introduction

The previous chapter outlined the experimental procedures and main analyses that could address the central questions of concern in this thesis. More specifically, that chapter examined whether or not FAST scores (and IAT scores) differed across, and therefore could distinguish between, known groups and, using a regression analysis, the extent to which these measures, and a self-report attitude questionnaire could retrospectively predict voting behaviour. While important conclusions can be drawn from that analysis regarding the potential utility of the FAST, some more in-depth and interrogative analyses of that rich data set are possible. In particular, a more direct and meaningful comparison of the IAT and FAST is possible using the current data set, as is a more focused analysis of their predictive utility, using more advanced statistical methods, such as Area under the Curve (AUC) analysis. Thus, rather than conduct further experimental investigations to better understand the relative predictive utilities of the FAST and IAT, it was decided that an *ex post facto* research methodology would be adopted for the current chapter, and existing data re-coded in such a way as to allow further important questions to be answered. In particular, the FAST and IAT will be rescored using each other's scoring algorithms and also in novel ways (i.e., learning slope differential scores), so that meaningful comparisons can be made between the two measures. In effect, in this chapter, the first systematic analysis of the FAST in relation to the IAT will be conducted while holding scoring methods constant across both tests. That is, analyses conducted using the original scoring methods used for the FAST (RFD scores) and the IAT (D scores), will be complimented by analyses conducted on these test outcomes using D scores, RFD scores, and also learning slope differential scores. The slope score is of particular interest as it is arguably more comprehensive than the D score, in that it accounts for both accuracy and response times, and avoids the conceptual issues associated with the use of post-hoc manipulation of the data (see Ridgeway, Roche, Gavin, & Ruiz, 2010). As such, this chapter constitutes *ex post facto* research and is exploratory in nature.

3.2 The Receiver Operating Characteristic (ROC) Curve

In the previous chapter a binary logistic regression was conducted, which is a test that can be viewed as assessing whether a model is well calibrated or not. The model is said to be well calibrated if the observed risk matches the predicted risk (i.e. the probability; Dankers, Traverso, Wee, & van Kuijk, 2019). This model assigns every observation the same predicted probability. However, the model isn't considered very useful as it doesn't *discriminate* between those observations at high risk and those at low risk. As well as being well calibrated, the model would need to have high discrimination ability (Dankers et al., 2019). In the binary outcome context, this means that observations with y = 1 ought to be predicted high probabilities, and those with y = 0 ought to be assigned low probabilities. Such a model allows us to discriminate between low and high risk observations.

One approach to assessing the discrimination of a fitted logistic model would be to use the Receiver Operating Characteristic (ROC) curve (Fawcett, 2006). By evaluating different thresholds for rounding our model predictions, we can determine many sensitivity and specificity pairs. If we plot the sensitivity versus (1 – specificity) for all these pairs (i.e. the true positive rate versus the false positive rate), we obtain the ROC curve. Therefore, the ROC curve is created by plotting the true positive rate against the false positive rate at various threshold settings. ROC curves are used to observe how any predictive model can distinguish between the true positives and negatives. The true-positive rate is also known as *sensitivity* (i.e., the probability of predicting a real positive will be a positive), and the false-positive rate is also known as *1-specificity* (i.e., the probability of predicting a real negative will be a negative). The sensitivity is defined as the probability of the prediction rule or model predicting an observation as 'positive' given that in truth (y = 1). In other words, the sensitivity is the proportion of truly positive observations which is classified as such by the model or test. Conversely the specificity is the probability of the model predicting 'negative' given that the observation is 'negative' (y = 0).

ROCs are often used in this way to assess the efficiency of a diagnostic tests or to compare the diagnostic performance of two or more laboratory or diagnostic tests (Griner et al., 1981). ROC and related methods are the engine for statistically appraising a test’s performance at classifying cases into groups correctly, such as those with versus without mood disorder (Youngstrom, 2014). These methods also provide a statistical process for comparing different tests and deciding whether one is superior for making these classification decisions. Although these methods are often presented in the context of diagnostic decisions, they could be used with any dichotomous variable, and can extend to scenarios with multiple categories (Robin et al., 2011) or continuous dependent variables (Kruschke, 2011).

The Receiver Operating Characteristic (ROC) curve will be used throughout this chapter to compare scoring methodologies, firstly in relation to predicting voting behaviour (i.e., a yes or no abortion referendum vote), and secondly in relation to predicting group affiliation (i.e. pro-choice or pro-life). Scoring metrics that will be assessed and compared across both the FAST and IAT tests include RFD scoring, D-algorithm scoring, and Slope scoring. These scoring metrics for both the FAST and IAT tests will also be compared against the explicit test measure used in the Attitudes to Abortion Questionnaire.

The Area under the ROC curve (AUC) will be used to assess cut-off points for the more newly developed RFD scoring for the FAST in relation to both voting behaviour and group affiliation. The AUC is a measure of how well a parameter can distinguish between two groups, where an AUC of 1 would signify a perfect test, making correct predictions for every participant, and an AUC of 0.5 would signify a worthless test, giving random prediction results for participants. The AUC is the probability that if you were to take a random pair of observations, one with y = 1 and one with y = 0, the observation with y = 1 has a higher predicted probability than the other. The AUC thus gives the probability that the model correctly ranks such pairs of observations.

3.3 ROC Analysis of Voting Behaviour

3.3.1 Original Data

The ROC curves of participants RFD scores from the FAST, D scores from the IAT, and total scores from the Attitudes to Abortion Questionnaire (as reported in Chapter 2) in relation to voting behaviour, are displayed in Figure 5. The graph is constructed with sensitivity on the vertical axis and 1-specificity on the horizontal. Here, the true-positive rate (y = 1) represents a Yes vote (consistent with pro-choice views) and the false-positive rate (y = 0) represents a No vote (consistent with pro-life views), in relation to FAST RFD scores, IAT D scores and the Attitude to Abortion Questionnaires (AAQ) total scores. The diagonal reference line, from (0, 0) to (1, 1), is indicative of an independent variable that discriminates no different from guessing (50/50 chance). However, the ROC curve for FAST RFD scores, IAT D scores and the Attitudes to Abortion Questionnaire total scores in Figure 5 are well above this line. In fact, the area under the ROC curve (AUC) for FAST RFD scores was 0.809, for IAT D scores was 0.939, and for total scores from the Attitudes to Abortion Questionnaire was 0.983 (see section 3.3.1.1 for further detail), as compared to that of the diagonal reference line which is always 0.5 (half the graph).

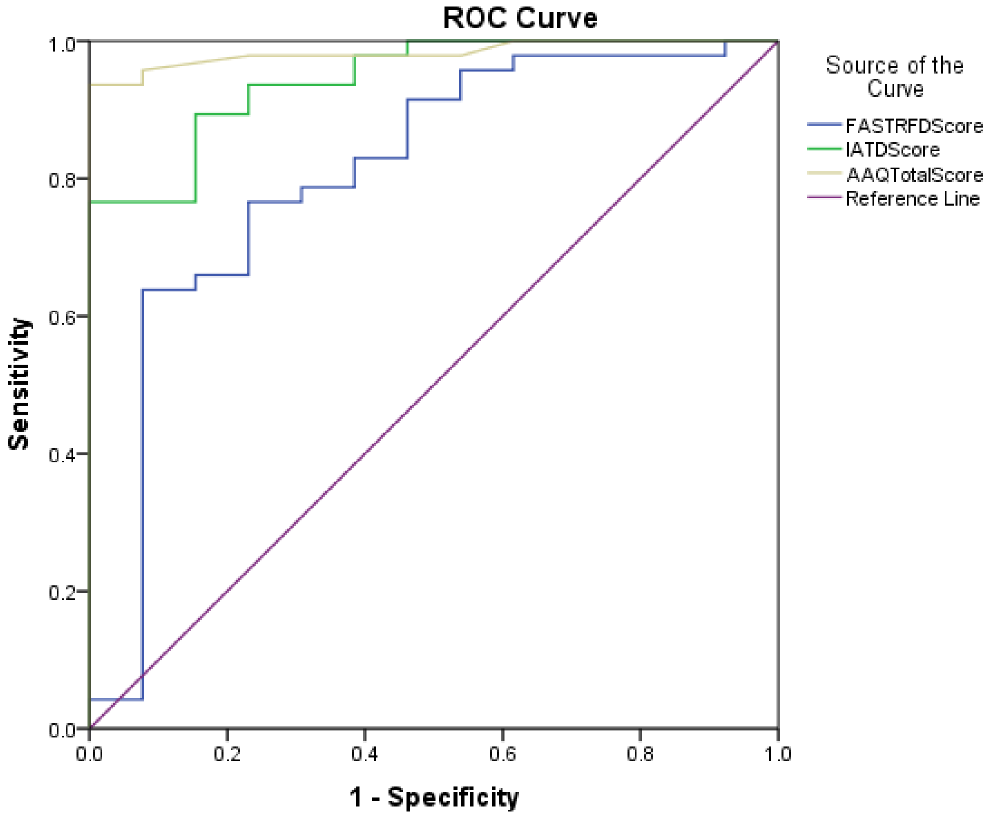

Figure 5: ROC curve with the test variables being FAST RFD scores (blue curve), IAT D scores (green curve) and Attitudes to Abortion Questionnaire (AAQ) total score (grey curve) with the state variable being voting behaviour where y = 1 represents yes votes and y = 0 represents no votes. The diagonal reference line is represented in purple.

3.3.1.1 Area Under the Receiver Operating Characteristic Curve

Powerful models have ROC curves that approach the upper left corner, which indicates that the model achieves near the maximum of 100% sensitivity and 100% specificity simultaneously. Conversely, a poor model with no predictive value will have an ROC curve close to the y = x or 45 degree line. This has led to the use of the Area Under the Curve (AUC) of the ROC curve, as a widely used metric for interpreting individual model performance but also for comparing between models (Green & Swets, 1966). Strong performing models have higher ROC curves and thus larger AUC values.

Through conducting an AUC analysis it was found that FAST RFD scores correctly classified randomly drawn pairs of yes/no voters 80.9% of the time, IAT D scores correctly classified randomly drawn pairs of yes/no voters 93.9% of the time, and explicit scores correctly classified randomly drawn pairs of yes/no voters 98.3% of the time, all of which translate to an excellent diagnostic power. More specifically, the AUC for FAST RFD scores of 0.809, for example, means that of all possible yes/no vote pairs created from the current sample, the model with RFD scores had (in 80.9% of the pairs) a higher predicted probability of being yes to participants who identified as pro-choice than it did to those who identified as pro-life. Table 4 displays significance tests for the AUC analysis, showing results for FAST RFD scores, IAT D scores and explicit scores form the Attitudes to Abortion questionnaire.

Table 4: Area Under the Curve significance tests for the FAST RFD scores, IAT D scores and Total scores from the explicit measure in relation to voting behaviour.

Test	Area	Std. Error	Asymptotic Sig	Asymptotic 95% Confidence Interval	
				Lower Bound	Upper Bound
FAST RFD scores	.809	.076	.001	.659	.958
IAT D scores	.939	.031	.000	.878	1.000
Explicit scores	.983	.014	.000	.955	1.000

3.3.1.2 Cut-off points for FAST RFD Scores

In a ROC curve the true-positive rate (sensitivity) is plotted in function of the false-positive rate (1-specificity) for different cut-off points. Each point on the ROC curve represents a sensitivity/specificity pair corresponding to a particular decision threshold. A test with perfect discrimination (no overlap in the two distributions) has a ROC curve that passes through the upper left corner. That is, the model or prediction rule is perfect at classifying observations if it has 100% sensitivity and 100% specificity. Therefore the closer the ROC curve is to the upper left corner, the higher the overall accuracy of the test (Zweig & Campbell, 1993). Therefore, the area in the upper left region of the ROC graph presented in Figure 5 provides the most useful discrimination in terms of identifying cut-off scores for the FAST RFD scoring metric (Streiner & Cairney, 2007).

For each observation, the fitted model in Figure 5 can be used to calculate the fitted probabilities. One way to create such a classification rule is to choose a cut-off point, and classify those observations with a fitted probability above the cut-off point as positive and those at or below it as negative. For this particular cut-off, we can estimate the sensitivity by the proportion of observations with $y = 1$ which have a predicted probability above this cut-off point, and similarly we can estimate specificity by the proportion of $y = 0$ observations with a predicted probability at or below the cut-off point. A cut-off RFD score for yes and no voters was calculated by analysing the sensitivity and specificity values on the ROC curve in this manner. However, as the cut-off point increases, fewer observations will be predicted as positive. This will mean that fewer of the $y = 1$ observations will be predicted as positive (reduced sensitivity), but more of the $y = 0$ observations will be predicted as negative (increased specificity). Thus, there is an intrinsic trade-off between sensitivity and specificity when selecting and deciding between different cut-off points.

The optimum cut-off point selected from this data set for FAST RFD scores was -4.85087, taking both sensitivity and specificity into account equally (see Appendix V for the FAST RFD scores sensitivity and specificity coordinate points for the ROC curve in relation to voting behaviour). Theoretically, participants who scored above this score were more likely to have voted yes, and participants who scored below this were more likely to have voted no. The cut off score was determined by selecting the point on the graph closest to 1 on the y-axis (presented in Figure 5). For this particular data set, 47.6% of participants who scored below -4.85087 voted no, and 92.3% of participants who scored above -4.85087 voted yes.

3.3.2 Converting FAST RFD Scores to D Scores

Greenwald and colleagues (2003) suggested improved scoring procedures for the IAT, the D measures, which were optimized with regard to the IAT's psychometric criteria (e.g., increased internal consistency, higher correlations with explicit measures, and resistance to some extraneous procedural influences). D measures differ from the conventional algorithms previously used for implicit tests in several aspects, including modified upper and lower tail treatment of latencies, inclusion of both correct and incorrect responses, with incorrect response latencies being increased by an error penalty, and an individual standardization similar to that in Cohen's effect size measure d (see Greenwald et al., 2003; Nosek et al., 2007). Since Greenwald and colleagues (2003) identified the D score (difference between the critical blocks means divided by the inclusive standard deviation) as the best way to compute the IAT score, most of these tests of psychometric validity of the IAT have been performed only using this particular scoring method (see Section 2.4.3.1 for a detailed account of the D score measure). In other words, research has mainly considered the psychometric properties of one type of IAT score. In accordance with Greenwald and colleagues (2003) findings and recommendations, it seems that the best approach is to calculate the difference between the two critical blocks is the D score. This method is one way to reduce the effect of the heavy tails of the distributions of each individual affecting both means and standard deviation.

Much of the IAT's popularity in terms of the D measures scoring method, may be attributable to its comparatively satisfactory reliability estimates. Therefore, D scores are now widely used for all IATs. Because of this, it is of interest to convert the original RFD scores from the FAST to the golden standard D score measure, so that the FAST and IAT can be more meaningfully compared. Scores were converted for the FAST using the D scoring metric as outlined in section 2.4.3.1 of this thesis. An ROC analysis presented in Figure 6 compares the predictive validity in terms of voting behaviour for the FAST and IAT, both using the D score metric. The graph is constructed with sensitivity on the vertical axis and 1-specificity on the horizontal. Here, the true-positive rate (y = 1) represents a Yes vote (consistent with pro-choice views) and the false-positive rate (y = 0) represents a No vote (consistent with pro-life views), for both FAST D scores and IAT D scores.

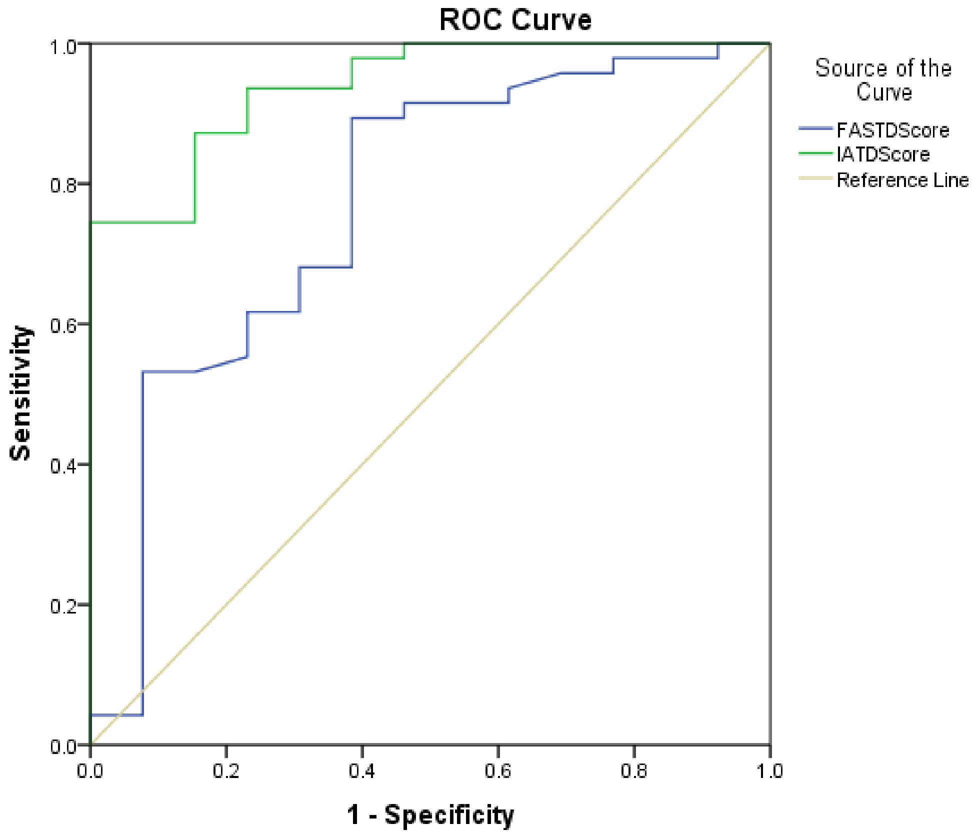

Figure 6**:** ROC curve with the test variables being FAST D scores (blue curve) and IAT D scores (green curve) with the state variable being voting behaviour where y = 1 represents yes votes and y = 0 represents no votes. The diagonal reference line at (0, 0) is represented in grey.

3.3.2.1 Area Under the Receiver Operating Characteristic Curve

The AUC for FAST D scores was 0.769, and for IAT D scores was 0.939, as compared to that of the diagonal reference line which is always 0.5 (half the graph). That is, through conducting an AUC analysis it was found that FAST D scores correctly classified randomly drawn pairs of yes/no voters 76.9% of the time, and that IAT D scores correctly classified randomly drawn pairs of yes/no voters 93.9% of the time, all of which signify an excellent diagnostic power. Therefore, the IAT D score had a higher predicted probability for distinguishing between voting behaviour (with 93.9% of the pairs) compared to that of the FAST D score. Table 6 displays significance tests for the AUC analysis, showing results for FAST D scores and IAT D scores in relation to voting behaviour.

Table 5: Area Under the Curve significance tests for the FAST D scores and IAT D scores in relation to voting behaviour.

Test	Area	Std. Error	Asymptotic Sig	Asymptotic 95% Confidence Interval	
				Lower Bound	Upper Bound
FAST D score	.769	.081	.003	.611	.928
IAT D score	.939	.031	.000	.878	1.000

3.3.2.2 Cut-off points for FAST D Scores

The optimum cut-off point selected from this data set for FAST D scores was -.665, taking both sensitivity and specificity into account equally (see Appendix VI for the FAST D scores sensitivity and specificity coordinate points for the ROC curve in relation to voting behaviour). Theoretically, participants who scored above this score were more likely to have voted yes, and participants who scored below this were more likely to have voted no. The cut off score was determined by selecting the point on the graph closest to 1 on the y-axis (presented in Figure 6). For this particular data set, 61.5% of participants who scored below -.655 voted no, and 89.3% of participants who scored above -.655 voted yes.

3.3.3 Converting FAST RFD Scores and IAT D Scores to Slope Scores

The RFD measures of the FAST is a novel scoring metric. When first used, RFD scores proved to be the most sensitive measure to changes in experimentally controlled inter-class stimulus relatedness in the FAST, and it consistently achieved greater statistical power than the pre-existing standard measure of slope scores for the FAST (Cummins, 2017). However, FAST learning slope differential scores could be seen as conceptually more appropriate as a measure of learning rate differentials. Therefore, as the slope score was that standard measure used for several published FAST studies (although see O'Reilly et al., 2012, 2013, and Cummins et al., 2019 for an alternative scoring system based on a different FAST training and scoring format), it is important to see here if slope scores may more accurately predict voting behaviour than that of RFD scores and D scores for both the FAST and IAT.

The FAST slope score is a measure that combines both response time and accuracy (i.e., fluency; Cartwright et al., 2016). Specifically, to find the slope score, cumulative learning curves are produced for each block and are fitted with regression lines, which are then contrasted. The linear regression lines are created according to the cumulative record generated by responses on each block. Higher FAST slope scores are represented by a steeper regression line, which indicates a combination of less overall time taken (i.e., quicker response times) and more accurate responses in a given block. The slope of the regression line recorded for the inconsistent block is subtracted from that recorded for the consistent block, and this value is that of the standard FAST score, the slope score. The size and direction of difference between FAST slope scores is taken to indicate the degree to which one pair of response classes is acquired more easily than the other, referred to as the FAST effect. A positive FAST slope score indicates a steeper learning rate on the consistent block compared to that of the inconsistent block, which would suggest that the functional classes established in the consistent block would be compatible with the stimulus equivalence training and testing contingencies. A negative FAST score indicates the opposite (i.e., superiority of response performance on the inconsistent block relative to the consistent block). This provides an indirect or "implicit" measure of the stimulus relation configurations that exist pre-experimentally in the repertoire of the subject (see O'Reilly et al., 2015 for a conceptual analysis of this effect, and Cartwright et al., 2016 for a recent empirical analysis using real-world stimuli).

This slope scoring methodology used in the quantification of implicit measure effects as seen in the FAST, overcomes a significant issue associated with the typical response time scoring of other implicit measures. Specifically, the use of means and standard deviations to analyse response times (as employed by the D algorithm), which assumes that the values are generally normally distributed. Treating response time data as normally distributed can then affect overall effect sizes, and consequently lead to complications associated with striking differences in response times (Whelan, 2008). Given that the slope score doesn't parameterize the response time data, this scoring metric avoids this confound. Therefore, IAT D scores were converted to slope scores as described above, and compared to that of FAST slope scores (also converted from RFD scores to slope scores as described above) using an ROC analysis as presented in Figure 7. The graph is constructed with sensitivity on the vertical axis and 1-specificity on the horizontal. Here, the true-positive rate (y = 1) represents a Yes vote (consistent with pro-choice views) and the false-positive rate (y = 0) represents a No vote (consistent with pro-life views), in relation to FAST slope scores and IAT slope scores.

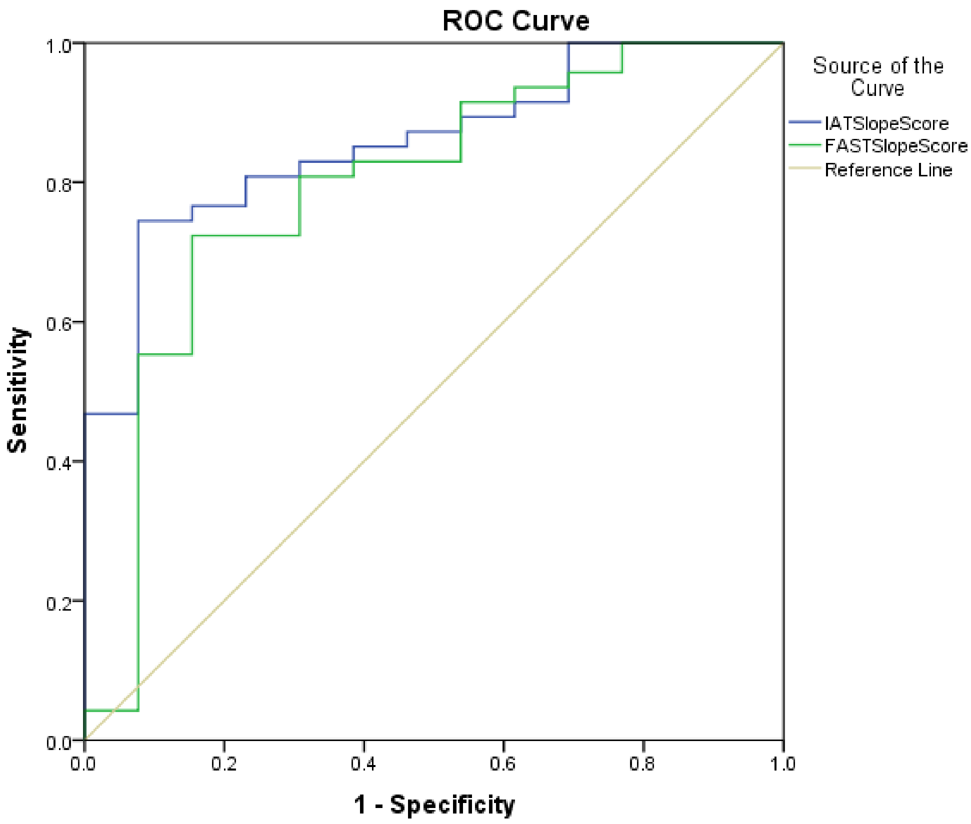

Figure 7: ROC curve with the test variables being IAT slope scores (blue curve) and FAST slope scores (green curve) with the state variable being voting behaviour where y = 1 represents yes votes and y = 0 represents no votes. The diagonal reference line at (0, 0) is represented in grey.

3.3.3.1 Area Under the Receiver Operating Characteristic Curve

The AUC for FAST slope scores was 0.794, and for IAT slope scores was 0.858, as compared to that of the diagonal reference line which is always 0.5 (half the graph). That is, through conducting an AUC analysis it was found that FAST slope scores correctly classified randomly drawn pairs of yes/no voters 79.4% of the time, and that IAT slope scores correctly classified randomly drawn pairs of yes/no voters 85.8% of the time, all of which signify an excellent diagnostic power. Therefore, the IAT slope score had a higher predicted probability for distinguishing between voting behaviour (with 85.8% of the pairs) compared to that of the

FAST slope score. Table 6 displays significance tests for the AUC analysis, showing results for FAST slope scores and IAT slope scores in relation to voting behaviour.

Table 6: Area Under the Curve significance tests for the FAST slope scores and IAT slope scores in relation to voting behaviour.

Test	Area	Std. Error	Asymptotic Sig	Asymptotic 95% Confidence Interval	
				Lower Bound	Upper Bound
FAST Slope	.794	.077	.001	.643	.944
IAT Slope	.858	.052	.000	.756	.959

3.3.3.2 Cut-off points for FAST Slope Scores

The optimum cut-off point selected from this data set for FAST slope scores was -.0513, taking both sensitivity and specificity into account equally (see Appendix VII for the FAST slope scores sensitivity and specificity coordinate points for the ROC curve in relation to voting behaviour). Theoretically, participants who scored above this score were more likely to have voted yes, and participants who scored below this were more likely to have voted no. The cut off score was determined by selecting the point on the graph closest to 1 on the y-axis (presented in Figure 7). For this particular data set, 44% of participants who scored below -.0513 voted no, and 94.3% of participants who scored above -.0513 voted yes.

3.3.4 Converting IAT D Scores to RFD Scores

As mentioned, scoring metrics used for the FAST procedure (i.e. slope scores and RFD scores) overcomes issues associated with that of the typical response time scoring of other implicit measures specifically that of the D algorithm used commonly for the IAT. Similarly to that of the slope score, the RFD score doesn't parameterize the response time data, making it an interesting scoring metric to explore when comparing the FAST and IAT. Therefore, IAT D scores were converted to the newly developed RFD score metric (see section 2.4.2.1 for how RFD scores were calculated), and compared to that of FAST RFD scores using an ROC analysis as presented in Figure 8. The graph is constructed with sensitivity on the vertical axis and 1-specificity on the horizontal. Here, the true-positive rate (y = 1) represents a Yes vote (consistent with pro-choice views) and the false-positive rate (y = 0) represents a No vote (consistent with pro-life views), for both FAST slope scores and IAT slope scores.

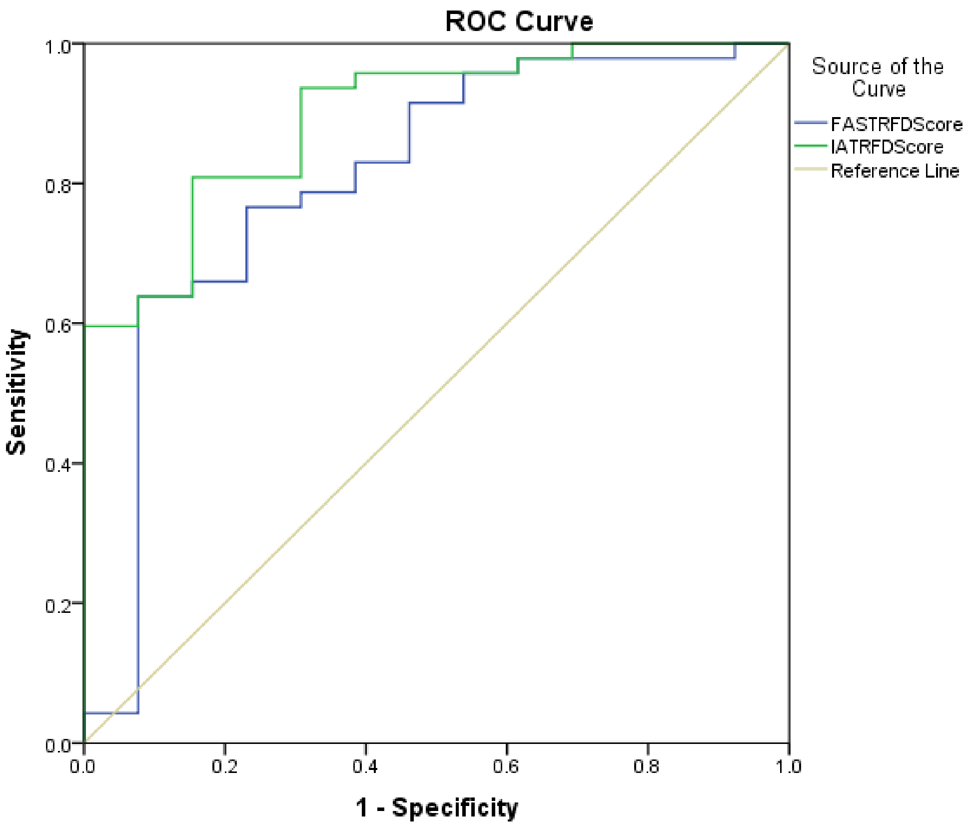

Figure 8: ROC curve with the test variables being FAST RFD scores (blue curve) and IAT RFD scores (green curve) with the state variable being voting behaviour where y = 1 represents yes votes and y = 0 represents no votes. The diagonal reference line at (0, 0) is represented in grey.

3.3.4.1 Area Under the Receiver Operating Characteristic Curve

The AUC for FAST RFD scores was 0.809, and for IAT RFD scores was 0.895, as compared to that of the diagonal reference line which is always 0.5 (half the graph). That is, through conducting an AUC analysis it was found that FAST RFD scores correctly classified randomly drawn pairs of yes/no voters 80.9% of the time, and that IAT RFD scores correctly classified randomly drawn pairs of yes/no voters 89.5% of the time, all of which signify an excellent diagnostic power. Therefore, the IAT RFD score had a higher predicted probability for distinguishing between voting behaviour (with 89.5% of the pairs) compared to that of the FAST RFD scores. Table 7 displays significance tests for the AUC analysis, showing results for FAST RFD scores and IAT RFD scores in relation to voting behaviour.

Table 7: Area Under the Curve significance tests for the FAST RFD scores and IAT RFD scores in relation to voting behaviour.

Test	Area	Std. Error	Asymptotic Sig	Asymptotic 95% Confidence Interval	
				Lower Bound	Upper Bound
FAST RFD	.809	.076	.001	.659	.958
IAT RFD	.895	.045	.000	.807	.984

3.3.5 Summary of Results

3.3.5.1 Area Under the Receiver Operating Characteristic Curve

The area under the (ROC) curve analysis has allowed for the comparison of different implicit test scoring metrics for two different implicit tests, the FAST and the IAT, in relation to participants voting behaviour. A summary of the results can be seen in Table 8. Here, the FAST has the best predictive validity in terms of voting behaviour when using the RFD scoring metric, whereas the IAT has the best predictive validity in terms of voting behaviour when using the D scoring metric. When comparing the FAST and IAT side by side, where both tests use the same scoring metric, the IAT outperforms the FAST in all scoring metrics (RFD scoring, Slope scoring and D scoring). However, neither tests outperform the results found for the explicit measure from the Attitudes to Abortion questionnaire at an AUC of .983. The AUC for the IAT D score of .939 is the closest result to that of the explicit measure.

Table 8: Summary Table of AUC analysis for all three scoring metrics for both the FAST and IAT in relation to voting behaviour, as well as an AUC analysis of the explicit measure.

Test Measure	Method of Scoring	Area under the Curve
FAST	RFD Score	.809
	Slope Score	.794
	D Score	.769
IAT	RFD Score	.895
	Slope Score	.858
	D Score	.939
Attitudes to Abortion Questionnaire	Total Score	.983

3.4 ROC Analysis of Group Affiliation

3.4.1 Original Data

The ROC curves in relation to participant's group affiliation for RFD scores from the FAST, D scores from the IAT, and total scores from the Attitudes to Abortion Questionnaire, as collected from the previous chapter, are displayed in Figure 9. The graph is constructed with sensitivity on the vertical axis and 1-specificity on the horizontal. Here, the true-positive rate (y = 1) represents participants who identified as pro-choice and the false-positive rate (y

= 0) represents participants who identified as pro-life, in relation to FAST RFD scores, IAT D scores and the Attitude to Abortion questionnaires total scores. The diagonal reference line, from (0, 0) to (1, 1), is indicative of an independent variable that discriminates no different from guessing (50/50 chance). However, the ROC curve for RFD FAST scores, IAT D scores and the Attitudes to Abortion Questionnaire total scores in Figure 9 are well above this line. In fact, the AUC for FAST RFD scores was 0.742, for IAT D scores was 0.889, and for total scores from the Attitudes to Abortion Questionnaire was 0.917 (see section 3.4.1.1 below), as compared to that of the diagonal reference line which is always 0.5 (half the graph).

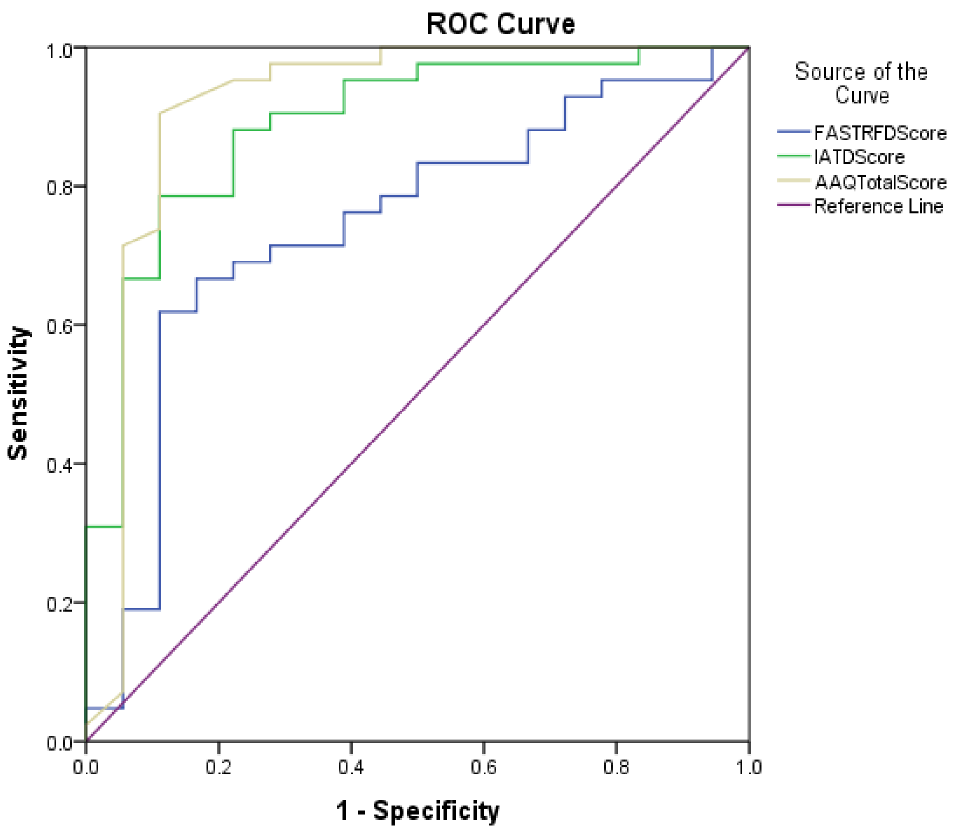

Figure 9: ROC curve with the test variables being FAST RFD scores (blue curve), IAT D scores (green curve) and Attitudes to Abortion questionnaire (AAQ) total score (grey curve) with the state variable being group affiliation where y = 1 represents participants who identified as pro-choice and y = 0 represents participants who identified as pro-life. The diagonal reference line is represented in purple.

3.4.1.1 Area Under the Receiver Operating Characteristic Curve

Through conducting an AUC analysis it was found that RFD scores correctly classified randomly drawn pairs of pro-choice/pro-life participants 74.2% of the time, IAT scores correctly classified randomly drawn pairs of pro-choice/pro-life participants 88.9% of the time, and explicit scores correctly classified randomly drawn pairs of pro-choice/pro-life participants 91.7% of the time, all of which translate to an excellent diagnostic power. Table 9 displays significance tests for the AUC analysis, showing results for FAST RFD scores, IAT D scores and explicit scores form the Attitudes to Abortion questionnaire.

Table 9: Area Under the Curve significance tests for the FAST RFD scores, IAT D scores and Total scores from the explicit measure in relation to group affiliation.

Test	Area	Std. Error	Asymptotic Sig	Asymptotic 95% Confidence Interval	
				Lower Bound	Upper Bound
FAST RFD score	.742	.072	.003	.602	.882
IAT D score	.889	.047	.000	.796	.981
Explicit score	.917	.053	.000	.813	1.000

3.4.1.2 Cut-off points for FAST RFD Scores

An optimum cut-off point for RFD scores was calculated by conducting a ROC curve (see graph 8) and analysing the sensitivity and specificity values, where a cut off scale value of -1.9421 was determined (see Appendix VIII for the FAST RFD scores sensitivity and specificity coordinate points for the ROC curve in relation to group affiliation). For this particular data set, participants who had an RFD score above -1.9421 were 92.8% likely to be pro-choice. Participants who had an RFD score below -1.9421 were 50% likely to be pro-life.

3.4.2 Converting FAST RFD Scores to D Scores

Scores were converted for the FAST from RFD scores to D scores using the D scoring metric as outlined in section 2.4.3.1 of this thesis. An ROC analysis presented in Figure 10 compares the predictive validity in terms of group affiliation for the FAST and IAT, both using the D score metric. The graph is constructed with sensitivity on the vertical axis and 1-specificity on the horizontal. Here, the true-positive rate (y = 1) represents participants who identified as pro-choice and the false-positive rate (y = 0) represents a participants who identified as pro-choice in relation to FAST D scores and IAT D scores.

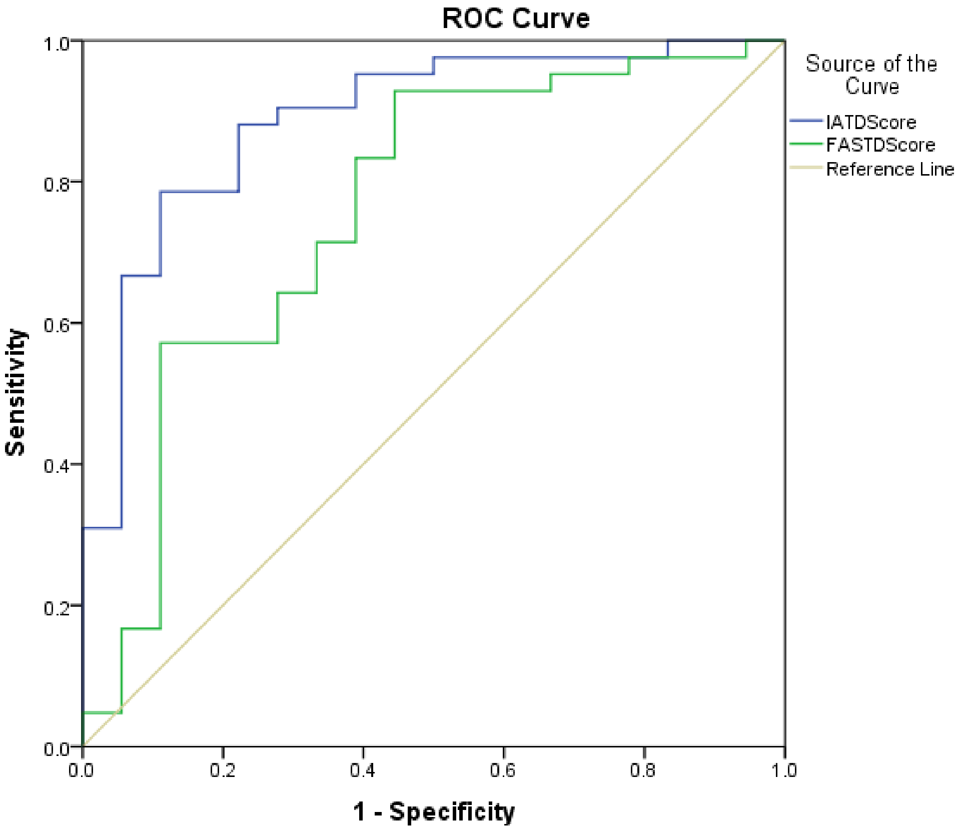

Figure 10: ROC curve with the test variables being IAT D scores (blue curve) and FAST D scores (green curve) with the state variable being group affiliation where y = 1 represents pro-choice groups and y = 0 represents pro-life groups. The diagonal reference line at (0, 0) is represented in grey.

3.4.2.1 Area Under the Receiver Operating Characteristic Curve

The AUC for FAST D scores was 0.759, and for IAT D scores was 0.889, as compared to that of the diagonal reference line which is always 0.5 (half the graph). That is, through conducting an AUC analysis it was found that FAST D scores correctly classified randomly drawn pairs of pro-choice/pro-life participants 75.9% of the time, and that IAT D scores correctly classified randomly drawn pairs of pro-choice/pro-life participants 88.9% of the time, all of which signify an excellent diagnostic power. Therefore, the IAT D score had a higher predicted probability for distinguishing between pro-life and pro-choice groups (with 88.9% of the pairs) compared to that of the FAST D score. Table 10 displays significance tests for the AUC analysis, showing results for FAST D scores and IAT D scores in relation to group affiliation.

Table 10: Area Under the Curve significance tests for the FAST D scores and IAT D scores in relation to group affiliation.

Test	Area	Std. Error	Asymptotic Sig	Asymptotic 95% Confidence Interval	
				Lower Bound	Upper Bound
IAT D score	.889	.047	.000	.796	.981
FAST D score	.759	.073	.002	.616	.903

3.4.2.2 Cut-off points for FAST D Scores

An optimum cut-off point for D scores was calculated by conducting a ROC curve (see graph 9) and analysing the sensitivity and specificity values, where a cut off scale value of –0.64 was determined (see Appendix IX for the FAST D scores sensitivity and specificity coordinate points for the ROC curve in relation to group affiliation). For this particular data set, participants who had an RFD score above -0.64 were 82.9% likely to be pro-choice. Participants who had an RFD score below -0.64 were 76.9% likely to be pro-life.

3.4.3 Converting FAST RFD Scores and IAT D Scores to Slope Scores

IAT D scores and FAST RFD scores were converted to slope scores as described in section 3.3.3, and compared using an ROC analysis as presented in Figure 11. The graph is constructed with sensitivity on the vertical axis and 1-specificity on the horizontal. Here, the true-positive rate (y = 1) represents pro-choice views and the false-positive rate (y = 0) represents pro-life views, in relation to FAST slope scores and IAT slope scores.

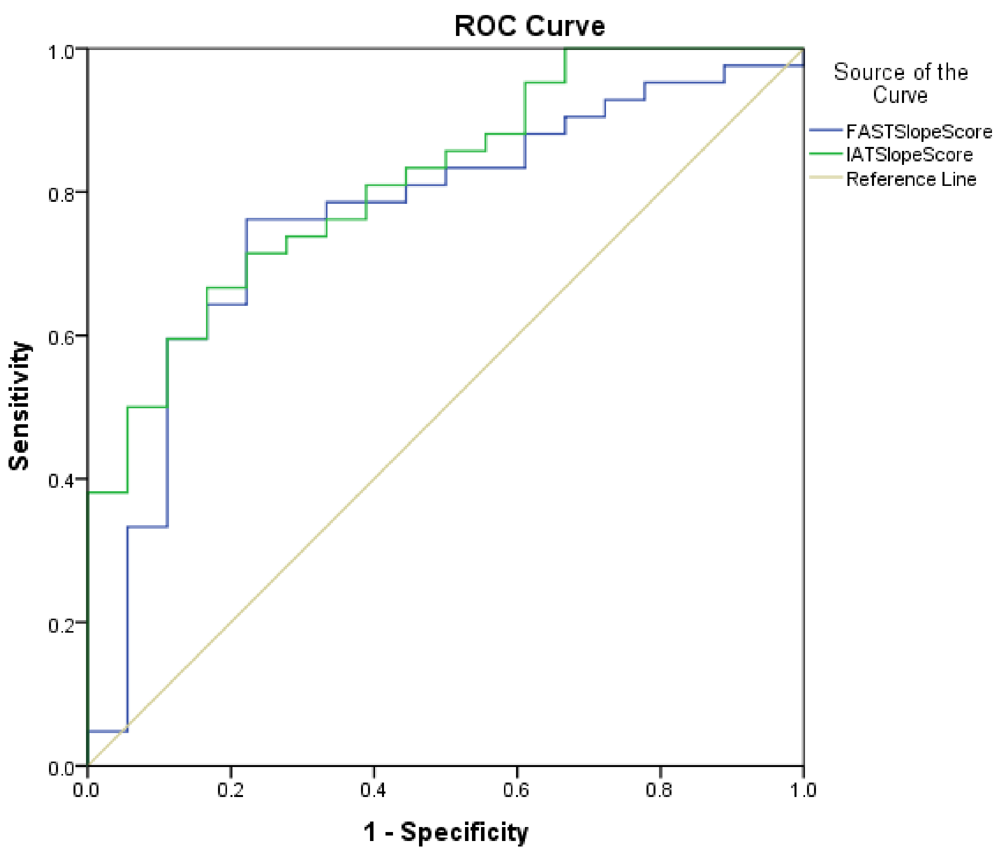

Figure 11: ROC curve with the test variables being FAST slope scores (blue curve) and IAT slope scores (green curve) with the state variable being group affiliation where y = 1 represents participants who identified as pro-choice and y = 0 represents participants who identified as pro-life. The diagonal reference line at (0, 0) is represented in grey.

3.4.3.1 Area Under the Receiver Operating Characteristic Curve

The AUC for FAST slope scores was 0.765, and for IAT slope scores was 0.816, as compared to that of the diagonal reference line which is always 0.5 (half the graph). That is, through conducting an AUC analysis it was found that FAST slope scores correctly classified randomly drawn pairs of pro-choice/pro-life participants 76.5% of the time, and that IAT slope scores correctly classified randomly drawn pairs of pro-choice/pro-life participants 81.6% of the time, all of which signify an excellent diagnostic power. Therefore, the IAT slope score had a higher predicted probability for distinguishing between voting behaviour (with 81.6% of the pairs) compared to that of the FAST slope score. Table 11 displays significance tests for the AUC analysis, showing results for FAST slope scores and IAT slope scores.

Table 11: Area Under the Curve significance tests for the FAST slope scores and IAT slope scores in relation to group affiliation.

Test	Area	Std. Error	Asymptotic Sig	Asymptotic 95% Confidence Interval	
				Lower Bound	Upper Bound
FAST Slope	.765	.068	.001	.630	.899
IAT Slope	.816	.056	.000	.706	.926

3.4.3.2 Cut-off points for FAST Slope Scores

An optimum cut-off point for FAST slope scores was calculated by conducting a ROC curve (see graph 10) and analysing the sensitivity and specificity values, where a cut off scale value of –0.0572 was determined (see Appendix X for the FAST slope scores sensitivity and specificity coordinate points for the ROC curve in relation to group affiliation). For this particular data set, participants who had an RFD score above -0.0572 were 88.6% likely to be pro-choice. Participants who had an RFD score below -0.0572 were 56% likely to be pro-life.

3.3.4 Converting IAT D Scores to RFD Scores

IAT D scores were converted to the newly developed RFD score metric (see section 2.4.2.1 for how RFD scores were calculated), and compared to that of FAST RFD scores using an ROC analysis as presented in Figure 12. The graph is constructed with sensitivity on the vertical axis and 1-specificity on the horizontal. Here, the true-positive rate (y = 1) represents a those who identify as and the false-positive rate (y = 0) represents those who identify as pro-life, in relation to FAST slope scores and IAT slope scores.

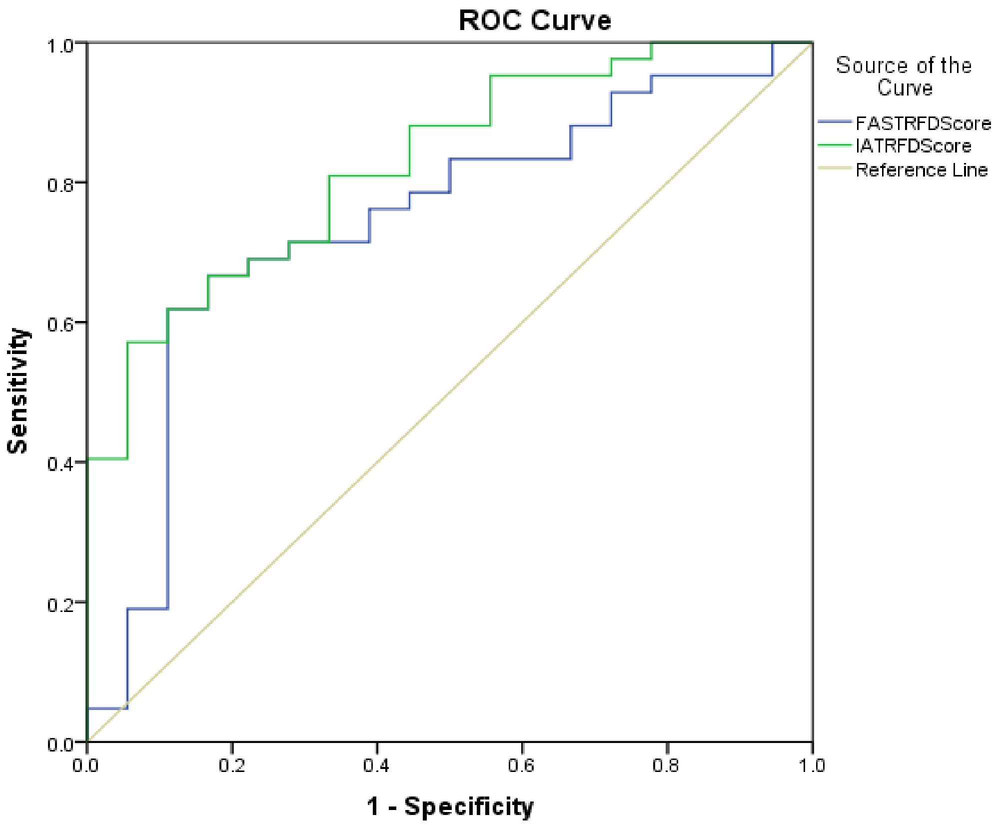

Figure 12: ROC curve with the test variables being FAST RFD scores (blue curve) and IAT RFD scores (green curve) with the state variable being group affiliation where y = 1

represents participants who identified as pro-choice and y = 0 represents participants who identified as pro-life. The diagonal reference line at (0, 0) is represented in grey.

3.4.4.1 Area Under the Receiver Operating Characteristic Curve

The AUC for FAST RFD scores was 0.742, and for IAT RFD scores was 0.827, as compared to that of the diagonal reference line which is always 0.5 (half the graph). That is, through conducting an AUC analysis it was found that FAST RFD scores correctly classified randomly drawn pairs of pro-choice/pro-life participants 74.2% of the time, and that IAT RFD scores correctly classified randomly drawn pairs of pro-choice/pro-life participants 82.7% of the time, all of which signify an excellent diagnostic power. Therefore, the IAT RFD score had a higher predicted probability for distinguishing between group affiliations (with 82.7% of the pairs) compared to that of the FAST RFD scores. Table 12 displays significance tests for the AUC analysis, showing results for FAST RFD scores and IAT RFD scores.

Table 12: Area Under the Curve significance tests for the FAST RFD scores and IAT RFD scores in relation to voting behaviour.

Test	Area	Std. Error	Asymptotic Sig	Asymptotic 95% Confidence Interval	
				Lower Bound	Upper Bound
FAST RFD	.742	.072	.003	.602	.882
IAT RFD	.827	.054	.000	.721	.933

3.3.5 Summary of Results

3.3.5.1 Area Under the Receiver Operating Characteristic Curve

The area under the (ROC) curve has allowed for the comparison of different implicit test scoring metrics for two implicit tests, the FAST and the IAT, in relation to participants group affiliation. A summary of the results can be seen in Table 13. Here, the FAST has the best predictive validity in terms of group affiliation when using the slope scoring metric, whereas the IAT has the best predictive validity in terms of group affiliation when using the D scoring metric. When comparing the FAST and IAT side by side, where both tests used the same scoring metric, the IAT outperforms the FAST in all scoring metrics (RFD scoring, Slope scoring and D scoring). However, neither tests outperform the results found for the explicit measure from the Attitudes to Abortion questionnaire at an AUC of .917. The AUC for the IAT D score of .889 is the closest result to that of the explicit measure.

Table 13: Summary Table of AUC analysis for all three scoring metrics for both the FAST and IAT in relation to group affiliation, as well as an AUC analysis of the explicit measure.

Test Measure	Method of Scoring	Area under the Curve
FAST	RFD Score	.742
	Slope Score	.765
	D Score	.759
IAT	RFD Score	.827
	Slope Score	.816
	D Score	.889
Attitudes to Abortion Questionnaire	Total Score	.917

3.4 Correlations

3.4.1 Relationship between FAST D Scores and IAT D Scores

The relationship between the FAST D scores and the IAT D scores were investigated using a Pearson product-moment correlation. There was a significantly moderate, negative correlation between the two variables, r = -.486, p < .001, with higher D scores in the FAST (higher scores indicative of pro-choice attitudes) associated with lower D scores in the IAT (also indicative of pro-choice attitudes). This indicates that scores obtained in the FAST reflected those obtained in the IAT. Thus, in effect, correlations were found to be as expected.

3.4.2 Relationship between FAST Slope Scores and IAT Slope Scores

The relationship between the FAST slope scores and the IAT slope scores were investigated using a Pearson product-moment correlation. There was a significantly moderate, negative correlation between the two variables, r = -.259, p < .05, with higher slope scores in the FAST (higher scores indicative of pro-choice attitudes) associated with lower slope scores in the IAT (also indicative of pro-choice attitudes). This indicates that scores obtained in the FAST reflected those obtained in the IAT. Thus, in effect, correlations were found to be as expected.

3.4.3 Relationship between FAST RFD Scores and IAT RFD Scores

The relationship between the FAST RFD scores and the IAT RFD scores were investigated using a Pearson product-moment correlation. There was a significantly moderate, negative correlation between the two variables, r = -.401, p < .01, with higher RFD scores in the FAST (higher scores indicative of pro-choice attitudes) associated with lower RFD scores in the IAT (also indicative of pro-choice attitudes). This indicates that scores obtained in the FAST reflected those obtained in the IAT. Thus, in effect, correlations were found to be as expected.

3.5 Discussion

This chapter focused on exploring different scoring methodologies for the FAST and IAT in relation to predictive validity of voting behaviour and group affiliation. Empirically different scoring metrics were explored, including the RFD scoring metric, the D score scoring metric and the slope score scoring method. The FAST and IAT original scoring methods were converted, so that scoring methods were constant across both tests, allowing for their predictive validity to be more accurately assessed and for the two implicit tests to be compared using varying scoring metrics. All scoring metrics used for the FAST and IAT were analysed using the ROC and AUC analysis. Here it was found that the FAST had the best predictive validity in terms of voting behaviour when using the RFD scoring metric, whereas the IAT has the best predictive validity in terms of voting behaviour when using the D scoring metric. In terms of group affiliation, it was found that the FAST had the best predictive validity when using the slope scoring metric, whereas the IAT had the best predictive validity when using the D scoring metric. When comparing the FAST and IAT side by side, where both tests use the same scoring metric, the IAT outperforms the FAST in all scoring metrics (RFD scoring, Slope scoring and D scoring) in relation to both voting behaviour and group affiliation. However, neither tests outperform the results found for the explicit measure from the Attitudes to Abortion questionnaire in relation to voting behaviour and group affiliation (see General Discussion).

As mentioned above, the AUC analysis shows that the RFD scoring metric for the FAST outperforms the slope scoring metric for the FAST in terms of predictive validity of voting behaviour. However, the opposite was true in the AUC analysis of the two scoring metrics for the FAST in terms of predictive validity of group affiliation, where the slope scoring metric for the FAST outperforms the RFD scoring metric for the FAST. This suggests that both the RFD scoring metric and slope scoring metric are potentially interchangeable, with no clear benefit of using one over the other, at least for purposes of predictive validity. However, the slope score is somewhat more conceptually-opaque than that of the RFD score. The use of a cumulative record (which the slope score employs) is well-established in behaviour analysis as a robust tool for the analysis of rates of responding (Skinner, 1959). However, fitting a regression line to this record (the process by which slope scores are calculated) adds a deal of conceptual opacity, in the sense that it is difficult to precisely define what is being measured (given that the score is in some way mathematically-abstract). In contrast, the RFD score is the result of simply subtracting the response rate differential (RRD) score of the inconsistent block from that of the consistent block. Conceptually, the RRD score for each block is relatively transparent: it simply reflects the difference in the rate of correct and incorrect responding. If random responding constitutes responding correctly and incorrectly at equal rates, then an RRD score for a given block will be zero. This method of scoring therefore minimises the potential effects of random responding in FAST scores. In contrast, the block-slope approach may suffer in some cases from inflation by random responding that is high in speed but low in accuracy (e.g., chance). In this case, random responding constitutes responding correctly and incorrectly at equal rates, but there is no counter measure built into the scoring method to reduce the impact of the 50% of correct responses made in what may be a very short period of time. In effect, FAST scores may be inflated disproportionately by a very rapid responding rate, even despite mere 50% accurate responding, should a participant be responding randomly (i.e., not under stimulus control of the test apparatus or feedback contingencies).

As mentioned above, it was found that the IAT outperforms the FAST in all scoring metrics (RFD scoring, Slope scoring and D scoring) in relation to both voting behaviour and group affiliation. Out of the three scoring metrics, the IAT had that highest AUC score when using the D score for both voting behaviour and group affiliation. Although the D score metric outperforms the two other scoring metrics, there are many questions outstanding relating to how the scoring metric interfaces with poorly understood procedural variables (see General Discussion

Correlating implicit measures, such as that of the IAT and the FAST, should reflect these test's convergent validity (e.g., Banaji, 2001). As the IAT and FAST measure very different constructs according to their conceptual paradigms, correlations were only carried out here when the scoring methods were constant across the two tests. In past studies, correlations between IATs and other implicit measures have typically been found to be weak, where the test scores are compared in the original formulation (e.g., Fazio & Olson, 2003; Sherman, Presson, Chassin, Rose, & Koch, 2003; Rudolph, Schröder-Abé, Schütz, Gregg, & Sedikides, 2008; Teige, Schnabel, Banse, & Asendorpf, 2004). For example, Fazio and Olson (2003) reported low correlations between IAT measures and priming measures of racial attitudes (r's ranging from −.13 to .05 across four studies), with Sherman and colleagues finding similar weak correlations between the priming and IAT measures (r's ranging from -.11 to .11) of smoking attitudes. For this study, however, correlations were strongest between the FAST and IAT when both tests used the D score (r = -.486). This was closely followed by the correlation found between the FAST and IAT when both tests used the RFD score (r = -.401). Correlations were weakest between the FAST and IAT when both tests used the slope score (r = -.259). Previous studies have found that low implicit– implicit consistency is often influenced by the structural fit of the measures (Payne, Burkley, & Stokes, 2008; see also De Houwer, 2008) and the structural differences among the measures (Rudolph et al., 2008).

Interestingly, however, there is quite a degree of *prima facia* structural fit between the FAST and IAT, but where differences do exist (e.g., the use of response correction in the IAT), these differences may have a significant effect on score magnitudes and stability across trials. Such nuanced artefacts will be discussed in more detail in the following chapter.

Chapter 4

General Discussion

4.1 Research Summary and Main Findings

The current research sought to validate a novel behavioural analytic implicit test measure, the Function Acquisition Speed Test (FAST), using a known-groups paradigm approach, as well as systematically analysing the FAST in relation to the Implicit Association Test (IAT). Specifically, the FAST was employed to measure abortion attitudes of two polarising groups with openly expressed pro-life and pro-choice views. The first aim of this research was to assess the FAST"'s utility in providing converging test results to predict both group affiliation and voting behaviour in the 2018 referendum on the 8th amendment concerning the legalisation of abortion. In seeking to validate the FAST, the gold standard IAT, as well as explicit test measures (i.e., the Attitudes to Abortion Questionnaire and Demographic survey), were administrated to assess convergence and predictive validity. The second aim of this research was to meaningfully compare the two implicit test measures, by employing common data scoring methods across the tests. In this chapter, the experimental findings from Chapters 2 and 3 will be discussed, as well as the implications of these findings. In addition, research challenges that must be confronted in future research will also be examined, and possible applications of the FAST in both basic research and applied settings will be discussed.

4.2 Summary of Results

In Chapter 2, a known group's paradigm approach was employed in order to validate the FAST through comparing FAST test scores with explicit and implicit test scores. That is, 60 participants with pre-existing pro-choice (42 participants) and pro-life (18 participants) attitudes completed an Attitudes to Abortion Questionnaire, a demographic survey where they were also asked how they voted in the abortion referendum, a FAST, and an IAT. The results from the experiment in Chapter 2 firstly found that the Attitudes to Abortion

Questionnaire was valid. That is, the total score correlated significantly with the participants' self-reported group membership (i.e., pro-choice and pro-life). This outcome paralleled that of Hess and Rueb's (2005), suggesting that the questionnaire accurately depicted group affiliation. In addition, significant differences were found in the questionnaires total scores between pro-life and pro-choice groups, as well as between yes and no referendum voters. These results were as predicted, considering that those recruited for the study were members of pro-life and pro-choice societies in which extreme pro-life and pro-choice views would be expected. Nevertheless, validating the questionnaire using this participant sample was important because it ensured that it was in fact a valid test measure with which the FAST could be validated.

The FAST used a relatively novel scoring metric (the rate-fluency differential score or RFD score) suggested by Cummins in 2017, which involved subtracting the response rate accuracy differential for the inconsistent block from the response rate accuracy differential of the consistent block. Significant differences were found between pro-life and pro-choice groups and yes and no voters using this measure, even though both pro-choice and pro-life groups did show an overall FAST effect in the same direction. That is, abortion related terms were more easily established as common members of a functional response class with negative terms than with positive terms for both groups. Therefore, pro-choice groups in fact responded on the FAST in a way that can be taken as indicative of a mildly negative view of abortion, while that of the pro-life group was indicative of significantly stronger negative views of abortion.

For the IAT, the standard D score metric developed by Greenwald and colleagues (2003) was used. Similarly to the FAST, significant differences were found between D scores of pro-life and pro-choice groups, as well as yes and no voters. However, unlike the FAST, pro-choice and pro-life groups showed an overall IAT effect in different directions. As would be predicted, abortion related terms were more easily established as common members of a functional response class with negative terms than with positive terms for pro-life members on average, where the opposite was true for pro-choice members. What is of interest here is finding out what exactly is driving these different effects between the IAT and the FAST. In the FAST, Abortion/Negative functional response classes are easier to establish for all subjects, although to different degrees. However, in the IAT, subjects are finding it easier to establish Abortion/Positive functional response classes. Potential reasons as to why this is happening in the IAT and not in the FAST will be discussed in section 4.3.

A significant correlation between FAST RFD scores and explicit scores (r = .405) was found in the predicted direction. However, a stronger significant correlation was found between IAT D scores and explicit scores (r = -.647). Both of these correlations are stronger than those typcially found in previous studies using the IAT. A meta-analysis covering various content domains (including attitudes, stereotypes, and self- concept) revealed a somewhat low correlation of .24 between IATs and explicit measures (Hofmann, Gawronski, Gschwendner, Le, & Schmitt, 2005), whereas a large-scale analysis of Internet data from the IAT website (projectimplicit.com) yielded a higher implicit–explicit correlation of .37 (Nosek, 2005). The results found in the current study may have yielded higher correlations because in the context of abortion attitudes, higher implicit-explicit consistency might be expected due to the obvious strength of polarised opinions. Also, Likert scale responses, as employed here, may better correspond with implicit measures, in that Likert scales are a relative measure that may more directly tap into an affective attitude component (Gawronski & Payne, 2010).

In order to assess the predictive validity of the FAST and IAT, a binary logistic regression was performed for each in relation to voting behaviour. It was found that both the FAST and the IAT predicted voting behaviour, adding significantly to the regression model.

The FAST correctly classified 83.3% of the voters, and the IAT correctly classified 90% of the voters, compared to that of the Attitudes to Abortion Questionnaire that correctly classified 93.3% of the voters. However, when controlling for the explicit measure, neither the FAST nor IAT significantly improved the model fit for predicting voting behaviour. This was not as predicted, as previous research using the IAT found that implicit measures did add significantly to the explanation of voting behaviour through conducting a binary logistic regression, even when the explicit measure was a strong predictor of voting intention and actual voting behaviour (Roccato & Zogmaister, 2010).

Evidence for the predictive validity of implicit measures across various behavioural domains was provided in a meta-analysis conducted by Greenwald, Poehlman, Uhlmann, and Banaji in 2009. This study found that the IAT showed better predictive validity than explicit measures in socially sensitive domains such as stereotyping and prejudice. This would have been expected given that socially desirable responding may lead to social desirability bias in explicit measures, particularly in these domains. However, in this study, participants appear to not have been in any way hiding their views on abortion. They were recruited from pro-choice and pro-life societies, often wearing jumpers or badges expressing such views to the public. It was therefore expected that the explicit measure would in fact outperform implicit measures in terms of predictive validity. Here, implicit and explicit measures tended to overlap as a result of well elaborated attitudes toward abortion and the abortion referendum, leaving little room for implicit measures to predict beyond explicit measures. The meta-analysis conducted by Greenwald and colleagues (2009) revealed just this, where the predictive validity for the IATs were lower than that found for the explicit measures in studies that explored brand preferences or political attitudes. The results found in the current study were also in accordance with results found in previous studies using the IAT to predict voting behaviour, where explicit, but not implicit, attitudes more accurately predicted future voting (Friese et al., 2012).

In terms of administrating a binary logistic regression there is a possibility that there were not enough cases to produce accurate findings. A logistic regression would usually require much more data to achieve stable, meaningful results. The sample size requirement for a logistic regression has been discussed a great deal in literature (Bujang, Sa'at, Sidik, & Joo, 2018). In general, studies with small to moderate sample sizes such as less than 100 have been found to frequently overestimate the effect measure. A simulation study conducted by Nemes and colleagues (2009) showed that large sample sizes, preferably consisting of 500 participants or more, would increase the accuracy of the estimates. This rule of thumb, requiring a minimum of 500 participants has also been supported by other studies (Long, 1997; Peduzzi, Concato, Kemper, Holford, & Feinstein, 1996). In these studies, a sample size of 500 or more, yielded statistics which represented the parameters in the targeted population. As mentioned before, a logistic regression requires at least 50 data points per predictor to achieve stable results, where this current study only had 60 participants in total. Therefore, the findings from this study using a binary logistic regression should be interpreted cautiously.

Chapter 3 focused on a more direct and meaningful comparison of the IAT and FAST using the current data set. That is, a more focused analysis of their predictive utility was carried out using more advanced statistical methods. This analysis was conducted using the original scoring methods used for the FAST (RFD scores) and the IAT (D scores), which was complimented by analyses conducted on these test outcomes using D score, RFD scores, and also learning slope differential scores, where scores were held constant across both the IAT and FAST. This was carried out in relation to both group affiliation and voting behaviour.

Analysing that data in this way allowed for meaningful comparisons to be made between these two test measures.

Firstly, a Receiver Operating Characteristics (ROC) curve was conducted using participants FAST RFD scores, IAT D scores, and total scores from the Attitudes to Abortion Questionnaire. From this, the Area Under the ROC curve (AUC) could be calculated. In relation to voting behaviour, FAST RFD scores had an area of .809, IAT D scores had an area of .939, and the explicit scores had an area of .983. In relation to group affiliation, FAST RFD scores had an area of .742, IAT D scores had an area of .889, and explicit scores had an area of .917. An area of 1 would signify a perfect test and 0.5 would represent a somewhat useless test, where the AUC represents correctly classified randomly drawn pairs of yes/no voters, or pro-choice/pro-life members. Therefore, an AUC analysis identified patterns that correspond well to those reported in Chapter 2, in which IAT D scores predicted cases slightly better than FAST RFD scores, while neither outperformed the almost perfectly predictive explicit test score.

FAST RFD scores were then converted to D scores and the previous analysis was repeated. An area of .769 was found for FAST D scores in relation to voting behaviour, and an area of .742 was found for FAST D scores in relation to group affiliation. Therefore, converting FAST RFD scores to D scores, lowered the FASTs overall predictive validity in relation to both voting behaviour and group affiliation. So in this case, RFD scores are better at measuring the predictive validity of the FAST than that of D scores.

Next, FAST RFD scores and IAT D scores were converted to learning slope differential scores, a more appropriate measure of learning rate differentials used as the standard measure for several published FAST studies. Again, an ROC curve and AUC analysis were conducted. In relation to voting behaviour, an area of .794 was found for FAST slope scores and an area of .858 was found for IAT slope scores. In relation to group affiliation, an area of .765 was found for FAST slope scores, and an area of .816 was found for IAT slope scores. For both voting behaviour and group affiliation, the IAT still outperformed the FAST when using slope scores, but less so than when using D scores. Therefore, the slope score was intrinsically not as sensitive as the D score method for the IAT. However, for the FAST, the slope score outperformed the D score in relation to both voting behaviour and group affiliation. It can be inferred from this that the FAST and IAT are both conceptually coherent with respect to their scoring methods. In other words, they are both clearly using methodologies conceived in certain conceptual terms to measure different phenomena, and their scoring methods are appropriate to that, maximising the sensitivity in each case. Moreover, the FAST slope scores outperformed the FAST RFD scores in relation to group affiliation. Therefore, the slope score and RFD score for the FAST differ in sensitivity depending on whether we are trying to predict voting behaviour or group affiliation.

Lastly, IAT D scores were converted to RFD scores for a final AUC analysis. An area of .895 was found for IAT RFD scores in relation to voting behaviour, and an area of .827 was found for IAT RFD scores in relation to group affiliation. Here, the IAT still outperformed the FAST even when converted to RFD scores, similar to what was found when converting the FAST RFD scores and IAT D scores to slope scores. This could be due to a procedural artefact of the IAT such as the use of practice blocks. The IAT exposes participants to practice blocks using the experimental stimuli, producing much more stable data and eliminating much error variation. These effects caused by participants being exposed to practice blocks will be discussed in more detail in section 4.3.

All in all, it was found that the FAST had the best predictive validity in terms of voting behaviour when using the RFD scoring metric, whereas the IAT had the best predictive validity in terms of voting behaviour when using the D scoring metric. Therefore, each test naturally arrived at a scoring method that was sensitive to the very phenomenon the test method was devised to measure. However, in terms of group affiliation the FAST had the best predictive validity when using the slope scoring metric, whereas the IAT had the best predictive validity in terms of group affiliation when again using the D scoring metric. This suggests that for the IAT the D score algorithm is certainly the optimal of the three examined here, but for the FAST, the two scoring systems currently suggested as appropriate for it in the literature are of more or less equal utility.

Cut-off points were also identified for the FAST RFD scores, D scores and slope scores in relation to both voting behaviour and group affiliation using the ROC curve, taking into account both sensitivity and 1-specificity. In relation to voting behaviour, the optimum cut-off point for FAST RFD scores was -4.85, for FAST D scores was -.655, and for FAST slope scores was -.0513. In relation to group affiliation, the optimum cut-off point for FAST RFD scores was -1.942, for FAST D scores was -0.64, and for slope scores was -0.057. These results were consistent with what was found in Chapter 2 (see Figure 3 in section 2.4.2.2). That is, the optimum cut-off scores indicated that the majority of participants leaned towards negative scores, which represents a pro-life stance, although those who do identify as pro-life had a much lower negative score. This may not represent an ideal scenario for the FAST metric, insofar as, at least in the domain of abortion, it does not yield an equal split of scores around zero. For instance, recall that in Chapter 2 both pro-life and pro-choice groups showed an overall IAT effect in different directions (see Figure 4 in section 2.4.3.2). Here the IAT scores were more intuitively correct, where those who identified as pro-life had strong positive scores on average, and those who identified as pro-choice had negative scores on average, although weak. The FAST on the other hand, showed that both pro-life and pro-choice participants had negative scores on average, but to significantly different degrees. This is one of the biggest differences between the FAST and IAT found in this study, where the IAT picked up on the leaning of pro-choice views ever so slightly, whereas the FAST identified pro-choice participants as mildly pro-life. Potential reasons behind why we are seeing these differences will be discussed in section 4.3.

Correlations between the FAST and IAT using scoring metrics that were constant across both tests were also performed. A significantly moderate, negative correlation was found between FAST D scores and IAT D scores (r = -.486). A significantly moderate, negative correlation was also found between FAST slope scores and IAT slope scores (r = -.259). Lastly, a significantly moderate, negative correlation was found between FAST RFD scores and IAT RFD scores (r = -.401). All correlations were found to be as expected, where higher FAST D scores, FAST slope scores, and FAST RFD scores (where higher scores in the FAST were indicative of pro-choice attitudes) were associated with lower IAT D scores, IAT slope scores, and IAT RFD scores (where lower scores in the IAT were indicative of pro-choice attitudes). Therefore, scores obtained in the FAST reflected those obtained in the IAT for all scoring metrics, where the correlation between the FAST and IAT when both using the D score was the strongest. This latter outcome suggests that a response time focused scoring metric may allow scores on the IAT (which is conceived as a response time based measure) to converge most strongly with those of the FAST, when scored based on response times only. In contrast, when the both the FAST and IAT are scored in such a way that response accuracy weighs heavily in the resulting score, the FAST and IAT deviate. In effect, the FAST can adapt relatively well to a time based scoring metric, but the IAT suffers in terms of predictive validity and correlation with the FAST when it is scored in terms of accuracy. In other words, the FAST is sensitive to both response speed and accuracy, as one might expect for a test format conceived as a measure of both response speed and accuracy.

Similarly, the IAT is only particularly sensitive to response speed and not accuracy, as one might expect for a test format conceived as a measure of response speed.

4.3 Interpretations and Implications of the Current Findings

The IAT was ostensibly very like the FAST from the participant's perspective, but with some key and well noted differences (see O'Reilly et al., 2015). The FAST consisted of two 50-trial learning task blocks, whereas the IAT consisted of three 20-trial sorting task blocks, two 20-trial practice task blocks, and two 40-trial critical task blocks. The sorting task blocks present in the IAT are designed to familiarize the participants with the stimuli and to check that they can indeed distinguish the attributes from each other as well as the target categories from each other as intended by the experimenters. The FAST does not contain any sorting task blocks, and because of this we cannot be sure that participants were indeed able to discriminate the different attribute and target words as required for a meaningful performance on the test.

The IAT also exposes participants to practice blocks using the experimental stimuli, whereas the FAST does not contain any practice blocks. As mentioned, the practice blocks exposed to participants during the IAT produces much more stable data and eliminates much error variation. That is, the IAT developers discovered that retaining practice trials and error trials increased overall IAT effects, where correlations with explicit measures were actually higher for the measure based on the practice blocks than the measure based on the test blocks (Greenwald et al., 2003). Subsequently, retaining practice blocks meant that the IAT was more sensitive to a "true effect" (i.e. a stronger IAT effect). Throughout the current study comparing the IAT with the FAST, the IAT outperforms the FAST throughout in terms of predictive validity. The practice blocks present in the IAT and not present in the FAST could be driving these effects, where the inclusion of practice blocks might have helped in reducing randomness and stabilizing response times so that they were more reflective of current contingencies coming into contact with historical contingencies (i.e., the main effect was being captured more directly with less noise in the data).

Results showed that when enhancing the role of accuracy, and relatively reducing the role of speed in a score for the IAT (i.e., using slope scores), the predictive utility is lowered to a point where it matches the FAST, at least in relation to predicting group affiliation. This strongly suggests that the IAT is superior in its current format and with its current scoring system (the D score) because it is response-time based. However, this may be a misplaced conclusion. Specifically, the IAT it is not truly time based, insofar as part of the response time measure consists of procedurally arranged response time penalties as part of the response correction procedure as mentioned briefly above. In brief, the time penalties were formerly added post-hoc to all response times recorded for error responses, prior to the advent of the 2003 D-score algorithm. Since that time, they are usually buried in the procedure, making it no less arbitrary but masking the arbitrariness of the penalty as a way of converting errors into elongated response times (i.e., enhancing its sensitivity to possibly weak time differences vis-à-vis the transcoding of real accuracy differences). It is important to understand that the time penalties added to error trial response times, were arrived through a process of reverse engineering. That is, being of arbitrary magnitude, they were of use only if they were sufficient to create a significant response time differential across trial blocks where one was absent or was statistically insignificant. In the Greenwald (2003) paper there is a rather surprisingly open discussion regarding the fact that in the new IAT procedure to include response correction, there was a several hundred millisecond penalty needed to be introduced for error response times in order for statistical effects to be significant. In that paper, Greenwald and colleagues assessed the effects on response time recordings (i.e., including the arbitrary penalty delivered by the response correction requirement) on the statistical significance of the resulting IAT score. They concluded that the arbitrary time extension was similar to that formerly imposed by the manual post-hoc time penalty procedure and therefore did not threaten IAT effects. Such a suggestion is inconceivable in the experimental analysis of behaviour insofar as control over the behavioural phenomenon of interest is sought strictly through improved stimulus control, and not by the backward engineering of response metric amplification methods that are entirely numerical, post-hoc and arbitrary. Incidentally, it is worth remembering at this point that the IRAP procedure inherited this method in whole cloth without modification, and to this day no study has unpacked the real response fluency difference across critical blocks in the IRAP under different response quantification systems. In this regard, the IRAP has not advanced our understanding of the core behavioural process involved in contingency juxtaposition test methods such as the IAT, FAST and IRAP, except insofar as offering elegant and rather coherent, albeit largely untestable, interpretive accounts in terms of Relational Frame Theory.

Another issue that relates to potential methodological concerns is that optimum cut-off scores identified for the FAST RFD scores, FAST D scores and FAST slope scores were all negative. This indicated, similar to the results found in Chapter 2, that the scores were in the same direction for both pro-life and pro-choice groups. However, as mentioned previously, the IAT picked up on the pro-choice views of the pro-choice participants, where scores were in different directions for the pro-life and pro-choice groups. This may be because an implicit test effect is more valid when conceived in terms of response times rather than accuracy as mentioned above. This may work as follows. The IAT allows a limited number of very large response times to be included, up to a maximum of 10 seconds. It also raised response times under 300ms up to 300ms and truncated response times over 10s to 3000ms. While intended to stabilise the response time measure and reduce error variance, it may also have the effect of allowing sufficient variation to render the metric discriminative of performances across participants. In contrast, by the FAST cutting off response times at a maximum of 3s in order to increase time pressure and thereby yield high error rate differences across test blocks (i.e., an emphasis on accuracy is integral to the method and metric), participants' reaction time variances across blocks is diminished. While most responses are comfortably under 3s in any case, if allowed, several responses would exceed the 3s limited hold. This may partly explain how the FAST and IAT can position the response proficiency difference (conceived in terms of speed or fluency) around a non-zero point while at the same time still managing to detect a difference, even of a similar magnitude.

In the FAST, participants may be slow to form Abortion/Positive functional response classes, because even though they do not think abortion is bad, they cannot bring themselves to say that it is good. However, in the IAT there appears to be an ease with which participants respond to Abortion/Positive functional response classes. The IAT could be showing these effects because of the way it arranges responding and scoring as mentioned above. However, when the scoring is equalised across the tests, the difference in direction of effect across groups persists, with the optimum cut-off score still being negative for FAST D scores. Even when converting the scores, the IAT still benefited from longer artificial response times every time the subject made an error. So this may be the reason that the IAT seems to capture the effect in the "correct" direction as expected. In effect, response time could be a more sensitive measure for the IAT, particularly when it is amplified artificially whenever there is an error response.

Cut-off points were calculated for FAST RFD scores, D scores and slope scores in relation to both voting behaviour and group affiliation based on the ROCs and the AUCs sensitivity and 1-specificity pairs. Finding these cut-off scores makes a considerable contribution to the FAST literature insofar as it provides a scoring scale according to which real life voting behaviour and group affiliation can be predicted with up to 90% accuracy in some cases. However, issues with how cut-off scores are calculated have been identified. In this particular study, the point on the curve closest to (0, 1) is identified and the corresponding cut-off point is labelled "optimal" (Coffin & Sukhatme, 1997). The rationale behind this approach is that the point on the curve closet to perfection should be the optimal cut-off point, thus intuitively minimizing misclassification. However, it is argued that calculating a cut-off score in this way is merely a guess based on the data using probability, and may not necessarily be accurate. Another measure suggested for evaluating cut-off points more accurately is the Youden index (J; Perkins & Schisterman, 2006). The intuitive interpretation of the Youden index is that J is the point on the curve farthest from chance. J reflects the intention of maximizing overall correct classification rates and thus minimizing misclassification rates, while choosing the point closest to (0,1). It is for this reason that some researchers advocate for the use of J to find the optimal cut-off point.

As mentioned in the summary, correlations were strongest between the FAST and IAT when both tests used the D score (r = -.486). Again, this was most likely because the D score stabilised responses in the FAST by cutting off very rapid responses under 300ms, and ignoring accuracy. Therefore, the FAST and IAT correlated best when both tests used the D score because they indeed were then more closely measuring the same underlying "construct". However, this does not mean that the FAST should use a D score. It could just as well be argued that the IAT had compromised convergent validity because it correlated less with the FAST when using the slope score, which is preferable for conceptual reasons over the D score. It is important not to take the results from the IAT as a given, although it can be employed here as a moving target with which to compare the FAST, insofar as the IAT may one day revise its methodology and scoring method based precisely on these types of interrogative analyses. It would behove experimental analysts of behaviour interested in behavioural processes not to take as a given the psychometric validity of the IAT in developing their tests, and thereby compounding rather than expunging the methodological, epistemological and statistical issues that compromise the IAT as an approach to measuring the strength of association between stimuli.

The IAT has committed to a very precise scoring method and is consciously psychometric in its approach, basing itself on the concepts of traditional psychological test theory (i.e., aspiring to good internal consistency, reliability, and validity, despite being agnostic on what exactly the construct of interest is). However, the FAST has committed to avoiding a psychometric approach and reifying as constructs ("unconscious bias") the learning phenomena to which it is sensitive (i.e., resistance to change in functional response class formation). Of course, different scoring methods may suit different purposes and there is a tension between perfecting the scoring method employed by any test and losing sight of the importance of how scoring methods interact intimately with even the smallest methodological variation (e.g., in the case of the FAST the slope score may be very sensitive to inter-trial intervals, as these partly determine the slope of the continuously recorded learning curve). In effect, any ideal scoring system would be immediately threatened by procedural modifications, which are welcomed, encouraged and necessary for a vital experimental analysis of behavioural processes to evolve and improve over time. The IAT, and arguably the IRAP, in contrast, are now relatively immune to such evolutions in process analysis due to the reification and therefore potential petrification of their current methodologies.

It may not be acceptable to experimental analysts of behaviour to introduce social-cognitive style scoring methods for the FAST whereby for example, reaction times are simply elongated on an error trial. For instance, if 400ms were to be added to every error response reaction time (as is more or less achieved by the response correction procedure in the IAT and IRAP), the FAST effect size would increase significantly. However, this would not be achieved by increased behavioural control, as is the aim of the experimental analysis of behaviour, but by *post hoc* data manipulation. To put it simply, data points are being increased in IAT analyses, in the direction of hypotheses, on the basis that errors have the latent value of around 400ms of processing time. Presumably, had this not worked in early incarnations of the IAT, the penalty would have been increased to 600ms, or more. This strategy would always inevitably lead to significant and large test effects but on the back of underlying response *accuracy differentials* across critical test blocks, mistakenly reported as *response speed differences*. In effect, the research tradition from which the FAST has emerged, confronts it with serious epistemological objections to converting behaviour accuracy measures into behavioural response time measures in a conceptually opaque way, in the service of chasing statistical power rather than behavioural control and the illumination of behavioural process.

In order to address the issue of the FAST showing lower predictive validity than the IAT, it may be beneficial to adjust the relative difficulty of the FAST tasks. For example, instantiating shorter response windows may increase the difficulty of both FAST blocks, leading to a greater degree of variance in scores across blocks, and therefore more fine grained discrimination of verbal histories.

Overall, the FAST offers a functional model of implicit attitudes that is wanting for the IAT and so offers the advantage of continual and ongoing development and refinement, which has not happened with the IAT. Indeed, basic research surrounding implicit attitudes using the IAT has stagnated in recent years. Specifically, publications involving the use of the IAT has started to decrease in the last three years, with 403 papers published in 2017, 384 in 2018, and 307 in 2019 (trends found through the Web of Science database; see Figure 13 for trend line from 1998 to 2019). Publications and research using the IAT may have declined due the developer's refusal to change or improve it and therefore narrowing the opportunity to explore core process. As mentioned in section 1.4.2, there are many issues associated with the IAT and how it is measured that the developers do recognise. However, they are unable to as of yet satisfactorily explain the underlying causes of various artefacts, and have therefore avoided making any changes to the measure, arguing instead that it's apparent utility is grounds enough to leave the procedure untouched. That is not an acceptable approach, however, within the experimental analysis of behaviour tradition, in which correctly identifying behavioural process is a necessary part of a complete behavioural explanation of any given effect, which should also be supported by how it allows behavioural control and prediction.

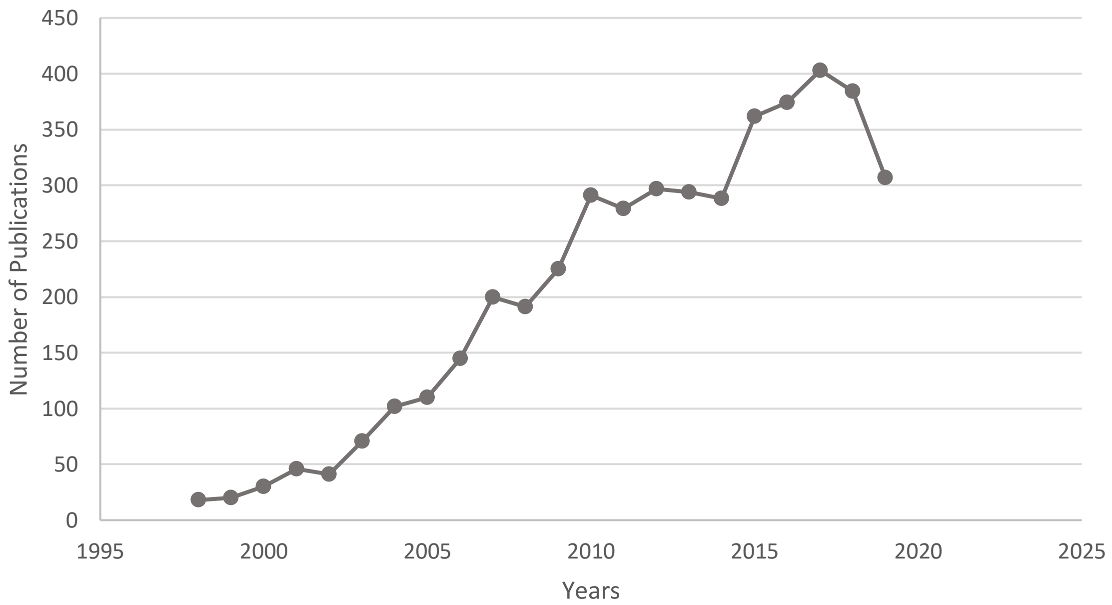

Figure 13: Trend line of the number of publications using the IAT from 1998 to 2019.

4.4 Future Research

The results of the current study show that the IAT reverses the effect direction for many subjects than what was found in the FAST. Future research should investigate what could be driving this effect (e.g. adding sorting task blocks and practice blocks in the FAST before comparing it to the IAT). Specifically, investigating the FAST to see whether the block that participants find easier (i.e. Abortion/Negative) really actually is easy or is it simply less hard than the Abortion/Positive blocks for those who are pro-choice. These types of questions need examining and answering but are complex because they are, after all, different procedures. What is recorded as “easy” is relative to the stimuli, the other block, and the scoring method. A more basic analysis is required using transparent procedures that systematically examine the effect of each procedural parameter on test outcomes.

Up to this point, the FAST has mostly only been discussed in terms of core process and the predictive validity of its scores. Similarly, the developers of the IAT and the IRAP have focused somewhat on predicting real world behaviour, such as voting behaviour, consumption, racially prejudicial decision making, and so on (but see Carpenter et al., 2012; Dawson et al., 2009). However, most of the research focuses on their scores based on standardized mean response latency differences across test blocks in terms of how it can distinguish group membership or correlate with explicit measures. What is sorely needed, however, is the beginnings of an analysis of how specific scores indexing stimulus relatedness relate to specific known strengths of relation between words in the vernacular (e.g., the words “love” and “good” as indexed on word association indices) or behaviours of known probabilities. At present, implicit test scores do not sit on a meaningful continuum which might allow us to make some face-valid assessment of the “strength” of a bias in relation to other biases. Therefore, providing a meaningful scoring system would represent a major transformation in how psychologists assess attitudes and represent a major advance over self-report tests that psychologists are calling for. This study and its findings are the first step in that direction. In assessing the FASTs predictive validity using a known group’s paradigm approach against explicit tests and other implicit tests, and using scoring methods constant across test measures, and providing cut-off scores for predicting voting behaviour with near certainty, new and important avenues of research are opened up regarding the real world utility of these test formats. For instance, future studies can now reasonably attempt to predict voting behaviour of individuals from unknown groups or undecided voters. This was the original hope for implicit tests, and a sound experimental analysis of the core process involved in one such test, not only makes this realising this hope more likely, but also puts the use of the test in real world settings on a far more solid footing.

As mentioned previously, a strong known group's paradigm approach was used for this study, which made it likely that neither of the implicit test would provide increased predictive validity over the self-reported measures. While this was not the aim of the current study, it may be a goal for future studies testing attitudes in more socially sensitive domains using groups such as terrorists or sex offenders. In effect, while the questionnaire used here performed impressively as a predictor of group membership and voting behaviour, this does not diminish the utility of the FAST in more ambiguous or sensitive attitude contexts (e.g., in cases where access to self-report data is difficult or in which social desirability is an challenge). Of course, it remains an empirical question, therefore, whether or not the IATs predictive utility would fall in the context of a different topic or greater social desirability but it is at least possible that the FAST would hold up and outperform a self-report even in such contexts.

In line with classical test theory, more stable data collected over a larger number of trials (i.e., including practice blocks), may provide a more accurate representation of any latent construct, such as the strength of an association between two stimuli. Results found that the IAT had better predictive validity than the FAST for all scoring metrics. Therefore, the FAST may benefit from considering the possibility that either practice trials may be required to stabilise responding or perhaps longer trial blocks are required. However, it may be that the error correction procedure in the IAT, which elongates response times on error trials, is responsible for the IATs superiority across the board. This error correction would have made the RFD score and slope score lower on the inconsistent block when converting scores, therefore attenuating inter-block differences. The use of negative feedback only on the IAT leads to further errors, disrupting the acquisition of smooth learning curves, for example. However, there is only so much that be gleaned about the two different methods from the current in depth data analysis. Therefore, an experimental analysis is also needed in future research, in which features such as the effect of negative feedback only, and response time elongation on test effect sizes are systematically analysed. This is an important insight and it could only have been arrived at with confidence by the current analysis in which scoring metrics were held constant across tests for purposes of direct comparison.

4.5 Conclusion

It can be concluded that the FAST can indeed predict group affiliation and voting behaviour satisfactorily. FAST RFD scores were found to be the optimal scoring metric when predicting voting behaviour, whereas FAST slope scores were found to be the optimal scoring metric when predicting group affiliation. Therefore, the slope score and RFD score for the FAST differ in sensitivity depending on whether we are trying to predict voting behaviour or group affiliation. Cut-off points were also provided for both scoring metrics, contributing to the production of a meaningful scoring system. Although the IAT D scores had the best overall predictive validity of the IAT scoring systems examined, the IAT is plagued with conceptual issues. The FAST offers a functional model of implicit attitudes that is wanting for the IAT and so offers the advantage of continual and ongoing development and refinement, hopefully sufficient to increase our prediction and influence over the well explicated behavioural processes involved.

References

Abramowitz, A. (1995). It's Abortion, Stupid: Policy Voting in the 1992 Presidential Election. *Journal of Politics, 57,* 176-186.

Aiken, A. R. A., Digol, I., Trussell, J., & Gomperts, R. (2017). Self-reported outcomes and adverse events after medical abortion through online telemedicine: Population based study in the republic of Ireland and Northern Ireland. *Bmj, 357*.

Aiken, A., Gomperts, R., & Trussell, J. (2017). Experiences and characteristics of women seeking and completing at-home medical termination of pregnancy through online telemedicine in Ireland and Northern Ireland: A population-based analysis. *BJOG: An International Journal of Obstetrics & Gynaecology, 124*(8), 1208-1215.

Aitken, K., Patek, P., & Murphy, M. E. (2017). The opinions and experiences of Irish obstetrics and gynaecology trainee doctors in relation to abortion services in Ireland. *Journal of Medical Ethics, 43*(11), 778-783.

Ajzen, I. (1991). The theory of planned behaviour. *Organizational Behaviour and Human Decision Processes, 50*(2), 179-211.

Ajzen, I., & Fishbein, M. (1980). *Understanding attitudes and predicting social behaviour*. Englewood Cliffs, NJ: Prentice-Hall.

Allport, G. W. (1935). Attitudes. In *A Handbook of Social Psychology* (pp. 798-844). Worcester, MA, US: Clark University Press.

Alvarez, R. M., & Brehm, J. (1995). American Ambivalence toward Abortion Policy: Development of a Heteroskedastic Probit Model of Competing Values. *American Journal of Political Science, 39,* 1055-1082.

Arcuri, L., Castelli, L., Galdi, S., Zogmaister, C., & Amadori, A. (2008). Predicting the vote: Implicit attitudes as predictors of the future behaviour of decided and undecided voters. *Political Psychology, 29*(3), 369-387.

Arkes, H. R., & Tetlock, P. E. (2004). Attributions of Implicit Prejudice, or "Would Jesse Jackson 'Fail' the Implicit Association Test?" *Psychological Inquiry, 15*(4), 257-278.

Ashburn-Nardo, L., Voils, C. I., & Monteith, M. J. (2001). Implicit associations as the seeds of intergroup bias: How easily do they take root? *Journal of Personality and Social Psychology, 81*(5), 789-799.

Banaji, M. R. (2001). Implicit attitudes can be measured. In H. L. Roediger, III, J. S. Nairne, I. Neath, & A. Surprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 117–150). Washington, DC: American Psychological Association.

Barnes-Holmes, D., Barnes-Holmes, Y., Hayden, E., Milne, R., Power, P., & Stewart, I. (2006). Do you really know what you believe? Developing the Implicit Relational Assessment Procedure as a direct measure of implicit beliefs. *The Irish Psychologist, 32*, 169-177.

Barnes-Holmes, D., Barnes-Holmes, Y., Hussey, I., & Luciano, C. (2015). Relational Frame Theory: Finding its historical and intellectual roots and reflecting upon its future development. In R. D. Zettle, S. C. Hayes, D. Barnes-Holmes, & A. Biglan (Eds.), *The Wiley Handbook of Contextual Behavioural Science* (pp. 115–128). UK: Wiley Blackwell.

Barnes-Holmes, D., Barnes-Holmes, Y., Luciano, C., & McEnteggart, C. (2017). From the IRAP and REC model to a multi-dimensional multi-level framework for analyzing the dynamics of arbitrarily applicable relational responding. *Journal of Contextual Behavioural Science, 6*(4), 434-445.

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*(3), 527-542.

Barnes-Holmes, D., Hayden, E., Barnes-Holmes, Y., & Stewart, I. (2008). The implicit relational assessment procedure (IRAP) as a response-time and event-related-potentials methodology for testing natural verbal relations: A preliminary study. *The Psychological Record, 58*(4), 497-516.

Barnes-Holmes, D., Murphy, A., Barnes-Holmes, Y., & Stewart, I., (2010). The Implicit Relational Assessment Procedure: 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,* 57-80.

Barnes-Holmes, D., Regan, D., Barnes-Holmes, Y., Commins, S., Walsh, D., Stewart, I., Smeets, P. M., Whelan, R., & Dymond, S. (2005). Relating derived relations as a model of analogical reasoning reaction times and event-related potentials. *Journal of the Experimental Analysis of Behaviour, 84*(3), 435-451.

Beaumont, C. (1997). Women, citizenship and Catholicism in the Irish free state, 1922-1948. *Women's History Review, 6*(4), 563-585.

Bem, D. J. (1972). Self-Perception Theory. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 6, pp.1-62). New York: Academic Press.

Binder, C. (1996). Behavioural fluency: Evolution of a new paradigm. *The Behaviour Analyst, 19*(2), 163-197.

Blanton, H., & Jaccard, J. (2006). Arbitrary metrics in psychology. *American Psychologist, 61*(1), 27-41.

Bluemke, M., & Friese, M. (2007). Reliability and validity of the Single-Target IAT (ST-IAT): Assessing automatic affect towards multiple attitude objects. *European Journal of Social Psychology, 38,* 977-997.

Blumer, H. (1955). Attitudes and the social act. *Social Problems, 3*(2), 59-65.

Bohner, G., & Wänke, M. (2002). *Attitudes and attitude change*. New York: Psychology Press.

Breen, R. (2000). Why is support for extreme parties underestimated by surveys? A latent class analysis. *British Journal of Political Science, 30,* 375–382.

Brown, A. S., Gray, N. S., & Snowden, R. J. (2009). Implicit measurement of sexual associations in child sex abusers: Role of victim type and denial. *Sexual Abuse: A Journal of Research and Treatment, 21*(2), 166-180.

Bujang, M. A., Sa'at, N., Sidik, T. M. I. T. A. B., & Joo, L. C. (2018). Sample size guidelines for logistic regression from observational studies with large population: Emphasis on the accuracy between statistics and parameters based on real life clinical data. *The Malaysian Journal of Medical Sciences, 25*(4), 122-130.

Cairns, E., & Mercer, G. W. (1984). Social identity in Northern Ireland. *Human Relations, 37*(12), 1095-1102.

Calkin, A. B. (2005). Precision teaching: The standard celeration charts. *The Behaviour Analyst Today, 6*(4), 207-215.

Cameron, C. D., Brown-Iannuzzi, J. L., & Payne, B. K. (2012). Sequential priming measures of implicit social cognition: A meta-analysis of associations with behaviour and explicit attitudes. *Personality and Social Psychology Review, 16*(4), 330-350.

Campbell, D. T. (1963). Social Attitudes and Other Acquired Behavioural Dispositions. In S. Koch, *Psychology: A study of a science. Study II. Empirical substructure and relations with other sciences. Vol. 6. Investigations of man as socius: Their place in psychology and the social sciences* (pp. 94-172). New York, NY, US: McGraw-Hill.

Carlsson, R., & Agerstrom, J. (2016). A closer look at the discrimination outcomes in the IAT literature. *Scandinavian Journal of Psychology, 57*(4), 278-287.

Carpenter, K. M., Martinez, D., Vadhan, N. P., Barnes-Holmes, D., & Nunes, E. V. (2012). 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*(2), 146-154.

Cartwright, A., Roche, B., Gogarty, M., O'Reilly, A., & Stewart, I. (2016). Using a modified function acquisition speed test (FAST) for assessing implicit gender stereotypes. *The Psychological Record, 66*(2), 223-233.

Coakley, J., & Gallagher, M. (2018). *Politics in the republic of Ireland*. Abingdon, Oxon: Routledge.

Coffin, M., & Sukhatme, S. (1997). Receiver operating characteristic studies and measurement errors. *Biometrics, 53*(3), 823-837.

Cook, E. A., Jelen, T. G. & Wilcox, C. (1992). *Between Two Absolutes: Public Opinion and the Politics of Abortion Boulder*. CO: Westview.

Cook, E. A., Jelen, T. G. & Wilcox, C. (1994). Issue Voting in U.S. Senate Elections: The Abortion Issue in 1990. *Congress and the Presidency, 21,* 99-112.

Cronbach, L. J., & Meehl, P. E. (1955). Construct validity in psychological tests. *Psychological Bulletin, 52*(4), 281-302.

Cullen, C., & Barnes-Holmes, D. (2008). Implicit pride and prejudice: A heterosexual phenomenon? In T. G. Morrison & M. A. Morrison (Eds.), *Modern prejudice* (pp. 195-223). New York: Nova Science.

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.

Cummins, J. (2017). *Quantifying differential stimulus relatedness using the function acquisition speed test* (Unpublished doctoral dissertation). Maynooth University, Kildare, Ireland.

Cummins, J., Roche, B., Tyndall, I., & Cartwright, A. (2018). The relationship between differential stimulus relatedness and implicit measure effect sizes. *Journal of the Experimental Analysis of Behaviour, 110*(1), 24-38.

Cummins, J., Tyndall, I., Curtis, A., & Roche, B. (2019). The function acquisition speed test (FAST) as a measure of verbal stimulus relations in the context of condom use. *The Psychological Record, 69*(1), 107-115.

Cvencek, D., Greenwald, A. G., Brown, A. S., Gray, N. S., & Snowden, R. J. (2010). Faking of the Implicit Association Test is statistically detectable and partly correctable. *Basic and Applied Social Psychology, 32*(4), 302-314.

Dambrun, M., & Guimond, S. (2004). Implicit and explicit measures of prejudice and stereotyping: Do they assess the same underlying knowledge structure? *European Journal of Social Psychology, 34,* 663-676.

Dankers F. J. W. M., Traverso A., Wee L., & van Kuijk S. M. J. (2019). Prediction Modelling Methodology. In P. Kubben, M. Dumontier, & A. Dekker (Eds), *Fundamentals of Clinical Data Science* (pp. 101-120)**.** Cham: Springer.

Dawson, D. L., Barnes-Holmes, D., Gresswell, D. M., Hart, A. J., & Gore, N. J. (2009). Assessing the implicit beliefs of sexual offenders using the implicit relational assessment procedure: A first study. *Sexual Abuse: A Journal of Research and Treatment, 21*(1), 57-75.

De Houwer, J. (2002). The implicit association test as a tool for studying dysfunctional associations in psychopathology: Strengths and limitations. *Journal of Behaviour Therapy and Experimental Psychiatry, 33*(2), 115-133.

De Houwer, J. (2003). The extrinsic affective Simon task. *Experimental Psychology, 50*(2), 77-85.

De Houwer, J. (2008). Implicit measures: Similarities and differences. In B. Gawronski & B. K. Payne (Eds.), *Handbook of implicit social cognition.* New York: Guilford.

De Houwer, J. (2011). Why the cognitive approach in psychology would profit from a functional approach and vice versa. *Perspectives on Psychological Science, 6*(2), 202-209.

De Houwer, J., Teige-Mocigemba, S., Spruyt, A., & Moors, A. (2009). Implicit measures: A normative analysis and review. *Psychological Bulletin, 135,* 347-368.

Demo, D. (1985). The measurement of self-esteem: Refining our methods. *Journal of Personality and Social Psychology, 48,* 1490–1502.

Deutscher, I. (1966). Words and deeds: Social science and social policy. *Social Problems*, *13*(3), 235-254.

Donders, F. C. (1969). Over de snelheid van psychische processen [On the speed of psychological processes]. In W. Koster (Ed.), *Attention and Performance: II* (Original work published 1868). Amsterdam: North-Holland.

Dovidio, J. F., & Fazio, R. H. (1992). New technologies for the direct and indirect assessment of attitudes. In J. M. Tanur (Ed.), *Questions about questions: Inquiries into the cognitive bases of surveys* (pp. 204–237). New York: Sage.

Dovidio, J. F., Kawakami, K., & Gaertner, S. L. (2002). Implicit and explicit prejudice and interracial interaction. *Journal of Personality and Social Psychology*, *82*(1), 62-68.

Dovidio, J. F., Kawakami, K., Johnson, B., Johnson, C., & Howard, A. (1997). On the nature of prejudice: Automatic and controlled processes. *Journal of Experimental Social Psychology*, *33*(5), 510-540.

Drake, C. E., Kellum, K. K., Wilson, K. G., Luoma, J. B., Weinstein, J. H., & Adams, C. H. (2010). Examining the implicit relational assessment procedure: Four preliminary studies. *The Psychological Record*, *60*(1), 81-100.

Drake, C. E., Primeaux, S., & Thomas, J. (2018). Comparing implicit gender stereotypes between women and men with the implicit relational assessment procedure. *Gender Issues*, *35*(1), 3-20.

Drake, C. E., Seymour, K. H., & Habib, R. (2016). Testing the IRAP: Exploring the reliability and fakability of an idiographic approach to interpersonal attitudes. *The Psychological Record, 66*(1), 153-163.

Dymond, S., & Barnes, D. (1994). A transfer of self-discrimination response functions through equivalence relations. *Journal of the Experimental Analysis of Behaviour, 62*(2), 251-267.

Dymond, S., & Barnes, D. (1995). A transformation of self-discrimination response functions in accordance with the arbitrarily applicable relations of sameness, more-than, and less-than. *Journal of the Experimental Analysis of Behaviour, 64*, 163-184.

Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes.* Orlando, FL, US: Harcourt Brace Jovanovich College Publishers.

Eagly, A. H., & Chaiken, S. (2005). Attitude Research in the 21st Century: The Current State of Knowledge. In D. Albarracín, B. T. Johnson, & M. P. Zanna (Eds.), *the handbook of attitudes* (pp. 743-767). Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers.

Earp, B. D., & Trafimow, D. (2015). Replication, falsification, and the crisis of confidence in social psychology. *Frontiers in Psychology, 6,* 621.

Edwards, A. L. (1957). *The social desirability variable in personality assessment and research.* Ft Worth, TX, US: Dryden Press.

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.

Eighth Amendment of the Constitution Act, Article 40.3.3, Act C8, Bill 35, Minister for Health, Dáil Éireann (1983). Retrieved from: http://www.irishstatutebook.ie/eli/1983/ca/8/enacted/en/print.html

Elias, N. (1978). *The history of manners*. Oxford: Blackwell.

European Court of Human Rights (1992). Case of open door and Dublin well woman v. Ireland. *Publications of the European Court of Human Rights. Series A, Judgments and Decisions. European Court of Human Rights, 246*, 1.

Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters, 27*(8), 861-874.

Fazio, R. H. (1986). How do attitudes guide behaviour? In R. M. Sorrentino & E. T. Higgins (Eds.), Handbook of motivation and cognition (pp. 204–243). New York: Guilford.

Fazio, R. H. (1990). A practical guide to the use of response latency in social psychological research. In C. Hendrick & M. S. Clark (Eds.), *Review of personality and social psychology, Vol. 11. Research methods in personality and social psychology* (pp. 74-97). Thousand Oaks, CA, US: Sage Publications, Inc.

Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology, 54,* 297–327

Fazio, R. H., & Roskos-Ewoldsen, D. R. (2005). Acting as We Feel: When and How Attitudes Guide Behaviour. In T. C. Brock & M. C. Green (Eds.), *Persuasion: Psychological insights and perspectives* (pp. 41-62). Thousand Oaks, CA, US: Sage Publications, Inc.

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.

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50,* 229–238.

Felzmann, H. (2014). Bringing abortion to Ireland? The protection of life during pregnancy act 2013. *IJFAB: International Journal of Feminist Approaches to Bioethics, 7*(1), 192-198.

Fenton, K. A., Johnson, A. M., McManus, S., & Erens, B. (2001). Measuring sexual behaviour: Methodological challenges in survey research. *Sexually Transmitted Infections, 77,* 84-92.

Ferriter, D. (2004). *The transformation of Ireland, 1900-2000*. London: Profile.

Ferriter, D. (2012). *Occasions of sin: Sex and society in modern Ireland*. London: Profile Books.

Festinger, L. (1957). *A theory of cognitive dissonance*. Evanston, IL: Row, Peterson.

Festinger, L. (1964). *Conflict, decision, and dissonance.* Palo Alto, CA: Stanford University Press.

Fiedler, K., & Bluemke, M. (2005). Faking the IAT: Aided and unaided response control on the implicit association tests. *Basic and Applied Social Psychology, 27*(4), 307-316.

Fishbein, M., & Ajzen, I. (1974). Attitudes towards objects as predictors of single and multiple behavioural criteria. *Psychological Review, 81*(1), 59-74.

Fishbein, M., & Ajzen, I. (1975). *Belief, Attitude, Intention, and Behaviour: An Introduction to Theory and Research.* Reading, MA: Addison-Wesley.

Fishbein, M., & Ajzen, I. (2010). *Predicting and changing behaviour: The reasoned action approach*. New York: Psychology Press.

Fletcher, R. (2000). National crisis, supranational opportunity: The Irish construction of abortion as a European service. *Reproductive Health Matters, 8*(16), 35-44.

Fourteenth Amendment of the Constitution Act, Article 40.3, Act C14, Bill 26, Minister for Justice, Dáil Éireann (1992). Retrieved from: http://www.irishstatutebook.ie/eli/1992/ca/14/enacted/en/html?q=The+Fourteenth+Amendment

Friese, M., Bluemke, M., & Wänke, M. (2007). Predicting voting behaviour with implicit attitude measures: The 2002 German parliamentary election. *Experimental Psychology, 54*(4), 247-255.

Friese, M., Smith, C. T., Plischke, T., Bluemke, M., & Nosek, B. A. (2012). Do implicit attitudes predict actual voting behaviour particularly for undecided voters? *PloS One, 7*(8), e44130-e44130.

Gaertner, S. L., & Dovidio, J. F. (1986). The aversive form of racism. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 61–89). Orlando, FL: Academic.

Galdi, S., Arcuri, L., & Gawronski, B. (2008). Automatic mental associations predict future choices of undecided decision-makers. *Science, 321*(5892), 1100-1102.

Gavin, A., Roche, B., & Ruiz, M. R. (2008). Competing contingencies over derived relational responding: A behavioural model of the implicit association test. The Psychological Record, 58, 427–441.

Gavin, A., Roche, B., Ruiz, M. R., Hogan, M., & O'Reilly, A. (2012). A behaviour analytically modified implicit association test for measuring sexual categorization of children. *The Psychological Record, 62*(1), 55-68.

Gawronski, B. (2009). Ten frequently asked questions about implicit measures and their frequently supposed, but not entirely correct answers. *Canadian Psychology, 50,* 141-150.

Gawronski, B., & Payne, B. K. (2010). *Handbook of implicit social cognition: Measurement, theory, and applications*. New York: Guilford Press.

Gawronski, B., Hofmann, W., & Wilbur, C. J. (2006). Are "implicit" attitudes unconscious? *Consciousness and Cognition, 15*(3), 485-499.

Green, B. F. (1954). Attitude Measurement. In G. Lindzey (Ed.), *Handbook of Social Psychology*. Reading, MA: Addison Wesley.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* Oxford, England: John Wiley.

Greenwald, A. G. (1992). Unconscious cognition reclaimed. *American Psychologist, 47,* 766–779.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102*(1), 4-27.

Greenwald, A. G., & Nosek, B. A. (2008). Attitudinal dissociation: What does it mean? In R. E. Petty, R. H. Fazio, & P. Briñol (Eds.), *Attitudes: Insights from the new implicit measures* (pp. 65-82). New York, NY, US: Psychology Press.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology, 74*(6), 1464-1480.

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 2008 U.S. presidential election. *Analyses of Social Issues and Public Policy, 9*(1), 241-253.

Grey, I., & Barnes, D. (1996). Stimulus equivalence and attitudes. *The Psychological Record, 46*, 243-270.

Griner, P. F., Mayewski, R. J., Mushlin, A. I., & Greenland, P. (1981). Selection and interpretation of diagnostic tests and procedures. *Annals of Internal Medicine, 94*, 555-600.

Hattie, J., & Cooksey, R. W. (1984). Procedures for Assessing the Validities of Tests Using the “Known-Groups” Method. *Applied Psychological Measurement*, *8*(3), 295–305.

Hawkins, D. I., & Coney, K. A. (1981). Uninformed response error in survey research. *Journal of Marketing Research, 18,* 370–374.

Hayes, S. C. (1994). Content, context, and the types of psychological acceptance. In S. C. Hayes, N. S. Jacobson, V. M. Follette, & M. J. Dougher (Eds.), *Acceptance and change: Content and context in psychotherapy* (pp. 13-32). Reno, NV: Context Press.

Hayes, S. C., Barnes-Holmes, D., & Roche, B. (2001). *Relational frame theory: A post-Skinnerian account of human language and cognition.* New York, NY, US: Kluwer Academic/Plenum Publishers.

Hayes, S. C., Barnes-Holmes, D., & Wilson, K. G. (2012). Contextual behavioural science: Creating a science more adequate to the challenge of the human condition. *Journal of Contextual Behavioural Science, 1*(1-2), 1-16.

Health (Regulation of Termination of Pregnancy) Act, Act 31, Bill 105, Minister for Health, Dáil Éireann (2018). Retrieved from http://www.irishstatutebook.ie/eli/2018/act/31/enacted/en/html?q=termination+of+pregnancy

Hess, J. A., & Rueb, J. D. (2005). Attitudes toward abortion, religion, and party affiliation among college students. *Current Psychology, 24*(1), 24-42.

Hofmann, W., Gawronski, B., Gschwendner, T., Le, H., & Schmitt, M. (2005). A meta-analysis on the correlation between the implicit association test and explicit self-report measures. *Personality and Social Psychology Bulletin, 31*(10), 1369-1385.

Holbrook, A. L., & Krosnick, J. A. (2010). Social desirability bias in voter turnout reports: Tests using the item count technique. *Public Opinion Quarterly, 74,* 37-67.

Holt, N., Bremner, A., Sutherland, E., Vliek, M., Passer, M., & Smith, R. (2015). *Psychology: The Science of Mind and Behaviour*. McGraw-Hill, Higher Education.

Holtgraves, T. (2004). Social desirability and self-reports: Testing models of socially desirable responding. *Personality and Social Psychology Bulletin, 30*(2), 161-172.

Hughes, S., & Barnes-Holmes, D. (2011). On the formation and persistence of implicit attitudes: New evidence from the implicit relational assessment procedure (IRAP). *The Psychological Record, 61*(3), 391-410.

Huijding, J., de Jong, P. J., Wiers, R. W., & Verkooijen, K. (2005). Implicit and explicit attitudes toward smoking in a smoking and a non-smoking setting. *Addictive Behaviours, 30*(5), 949-961.

Humphries, S. (1998). *Sex in a Cold Climate.* Testimony Films: Channel 4.

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.

Inglis, T. (1998). *Lessons in Irish sexuality*. Dublin: University College Dublin Press.

Inglis, T. (2005). Origins and legacies of Irish prudery: Sexuality and social control in modern Ireland. *Éire-Ireland, 40*(2), 9-37.

Inquisit 5 computer software (2018). Retrieved from https://www.millisecond.com.

Jacoby, L. L., Lindsay, D. S., & Toth, J. P. (1992). Unconscious influences revealed. *American Psychologist, 47,* 802–809.

Juarascio, A. S., Forman, E. M., Timko, C. A., Herbert, J. D., Butryn, M., & Lowe, M. (2011). Implicit internalization of the thin ideal as a predictor of increases in weight, body dissatisfaction, and disordered eating. *Eating Behaviours*, *12*(3), 207-213.

Karpinski, A., & Hilton, J. L. (2001). Attitudes and the implicit association test. *Journal of Personality and Social Psychology, 81*(5), 774-788.

Kawakami, K., & Dovidio, J. F. (2001). The reliability of implicit stereotyping. *Personality and Social Psychology Bulletin*, *27*(2), 212-225.

Keating, A. (2004). Church, state, and sexual crime against children in Ireland after 1922. *Radharc, 5*(7), 155-180.

Kennedy, F. (2000). The suppression of the Carrigan report: A historical perspective on child abuse. *Studies: An Irish Quarterly Review, 89*(356), 354-363.

Kihlstrom, J. F. (2004). Implicit methods in social psychology. In C. Sansone, C. C. Morf, & A. T. Panter (Eds.), *The Sage handbook of methods in social psychology* (pp. 195–212). Thousand Oaks, CA: Sage.

Killian, M., & Wilcox, C. (2008). Do abortion attitudes lead to party switching? *Political Research Quarterly, 61*(4), 561-573.

Kim, D. Y. (2003). Voluntary controllability of the Implicit Association Test (IAT). *Social Psychology Quarterly, 66*(1), 83-96.

Klassen, D., Anderson, P. B., & Hornstra, R. K. (1975). Influence of social desirability on symptom and mood reporting in a community survey. *Journal of Consulting and Clinical Psychology, 43,* 448-452.

Klauer, K. C., Voss, A., Schmitz, F., & Teige-Mocigemba, S. (2007). Process components of the Implicit Association Test: A diffusion-model analysis. *Journal of Personality and Social Psychology, 93*(3), 353-368.

Krosnick, J. A., Judd, C. M., Wittenbrink, B. (2005). Attitude measurement. In D. Albarracin, B. T. Johnson, & M. P. Zanna (Eds.), *Handbook of attitudes and attitude change*. Mahwah, NJ: Erlbaum.

Kruschke, J. K. (2011). *Doing Bayesian data analysis: A tutorial with R and BUGS*. New York, NY: Academic Press; 2011.

Kühnen, U., Schiessl, M., Bauer, N., Paulig, N., Pöhlmann, C., & Schmidthals, K. (2001). How robust is the IAT? Measuring and manipulating implicit attitudes of east- and west-Germans. *Zeitschrift Fur Experimentelle Psychologie: Organ Der Deutschen Gesellschaft Fur Psychologie, 48*(2), 135.

Lane, K. A., Banaji, M. R., Nosek, B. A., & Greenwald, A. G. (2007). Understanding and Using the Implicit Association Test: IV: What We Know (So Far) about the Method. In B. Wittenbrink & N. Schwarz (Eds.), *Implicit measures of attitudes* (pp. 59-102). New York, NY, US: The Guilford Press.

LaPiere, R. T. (1934). Attitudes vs. actions. *Social Forces, 13*(2), 230-237.

Leigland, S. (1997). Is a new definition of verbal behaviour necessary in light of derived relational responding? *The Behaviour Analyst, 20*(1), 3-9.

Likert, R. (1932). A technique for the measurement of attitudes. *Archives of Psychology, 22*(140), 55.

Lipkens, G., Hayes, S. C., & Hayes, L. J. (1993). Longitudinal study of derived stimulus relations in an infant. *Journal of Experimental Child Psychology, 56*, 201-239.

Long, J. S. (1997). *Advanced quantitative techniques in the social sciences series, Vol. 7. Regression models for categorical and limited dependent variables.* Thousand Oaks, CA, US: Sage Publications, Inc.

Maguire, M. J. (2001). The changing face of catholic Ireland: Conservatism and liberalism in the Ann Lovett and Kerry babies' scandals. *Feminist Studies*, *27*(2), 335-358.

Mahon, E., Conlon, C., & Dillon, L. (1998). *Women and crisis pregnancy: A report presented to the department of health and children.* Dublin: Stationery Office.

Maxwell, J. C. (2002). *Pro-life activists in America: Meaning, motivation, and direct action.* New York: Cambridge University Press.

Mayock, P., & Kitching, K. (2007). *Relationships and sexuality education (RSE) in the context of social, personal and health education (SPHE): An assessment of the challenges to full implementation of the programme in post-primary schools*. Dublin: Crisis Pregnancy Agency.

McCarthy, R. L. (2010). *Origins of the Magdalene laundries: An analytical history*. Jefferson, N.C: McFarland & Company, Inc., Publishers.

McKenna, I. M., 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 & Psychological Therapy*, *7*(2), 253-268.

McNamee, K and Power, R. (2017) 'Varadkar Announces Plan to Hold a Referendum on the Eighth Amendment', The University Times, 14 June, available: http://www.universitytimes.ie/2017/06/varadkar-announces-plan-to-hold-areferendum-on-the-eighth-amendment/ [accessed October 2017].

Miller, D. W. (1975). Irish Catholicism and the great famine. *Journal of Social History, 9*(1), 81-98.

Mischel, W. (1968). *Personality and assessment.* New York: Wiley.

Monteith, M. J., Voils, C. I., & Ashburn-Nardo, L. (2001). Taking a look underground: Detecting, interpreting, and reacting to implicit racial biases. *Social Cognition, 19*(4), 395-417.

Murphy, C., MacCarthaigh, S., & Barnes-Holmes, D. (2014). Implicit relational assessment procedure and attractiveness bias: Directionality of bias and influence of gender of participants. *International Journal of Psychology and Psychological Therapy, 14*(3), 333.

Murray, C. (2016). The Protection of Life During Pregnancy Act 2013: suicide, dignity and the Irish discourse on abortion. *Social and Legal Studies, 25*(6), 667-698.

Nederhof, A. J. (1985). Methods of coping with social desirability bias: A review. *European Journal of Social Psychology, 15*(3), 263-280.

Nemes, S., Jonasson, J., Genell, A., & Steineck, G. (2009). Bias in odds ratios by logistic regression modelling and sample size. *BMC Medical Research Methodology, 9*, 56.

Nolan, J. (1989). *Ourselves alone: Women's emigration from Ireland, 1885-1920*. Lexington: University Press of Kentucky.

Norrander, B., & Wilcox, C. (1999). Public opinion and policymaking in the states: The case of post-roe abortion policy. *Policy Studies Journal, 27*(4), 707-722.

Nosek, B. A. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology: General, 134*(4), 565-584.

Nosek, B. A., & Banaji, M. R. (2001). The Go/No-go Association Task. *Social Cognition, 19*(6), 625-666.

Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics: Theory, Research, and Practice, 6*(1), 101-115.

Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2005). Understanding and using the Implicit Association Test: II. Method variables and construct validity. *Personality and Social Psychology Bulletin, 31(2),* 166–180.

Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2007). The Implicit Association Test at age 7: A methodological and conceptual review. In J. A. Bargh (Ed.), *Automatic processes in social thinking and behaviour* (pp. 265–292). New York: Psychology Press.

Nosek, B. A., Hawkins, C. B., & Frazier, R. S. (2011). Implicit social cognition: From measures to mechanisms. *Trends in Cognitive Sciences, 15*(4), 152-159.

O'Reilly, A., Roche, B., & Cartwright, A. (2015). Function over form: A behavioural approach to implicit attitudes. In Z. Jin, *Exploring Implicit Cognition: Learning, Memory, and Social Cognitive Processes* (pp. 162-182). Philadelphia: Hershey.

O'Reilly, A., Roche, B., Gavin, A., Ruiz, M. R., Ryan, A., & Campion, G. (2013). A function acquisition speed test for equivalence relations (faster). *The Psychological Record*, *63*(4), 707-724.

O'Reilly, A., Roche, B., Ruiz, M., Tyndall, I., & Gavin, A. (2012). The function acquisition speed test (FAST): A behaviour analytic implicit test for assessing stimulus relations. *The Psychological Record, 62*(3), 507-528.

O'Hora, D., Barnes-Holmes, D., & Stewart, I. (2014). Antecedent and consequential control of derived instruction-following. *Journal of the Experimental Analysis of Behaviour, 102*(1), 66-85.

Olson, J. M., & Zanna, M. P. (1993). Attitudes and attitude change. *Annual Review of Psychology, 44,* 117-154.

Olson, M. A., & Fazio, R. H. (2004). Reducing the Influence of Extrapersonal Associations on the Implicit Association Test: Personalizing the IAT. *Journal of Personality and Social Psychology, 86*(5), 653-667.

Oswald, F. L., Mitchell, G., Blanton, H., Jaccard, J., & Tetlock, P. E. (2013). Predicting ethnic and racial discrimination: A meta-analysis of IAT criterion studies. *Journal of Personality and Social Psychology, 105*(2), 171-192.

Paulhus, D. L. (1984). Two-component models of socially desirable responding. *Journal of Personality and Social Psychology, 46,* 598-609.

Payne, B. K., Burkley, M. A., & Stokes, M. B. (2008). Why do implicit and explicit attitude tests diverge? The role of structural fit. *Journal of Personality and Social Psychology, 94*(1), 16-31.

Payne, B. K., Krosnick, J. A., Pasek, J., Leikes, Y., Akhtar, O., & Tompson, T. (2010). Implicit and explicit prejudice in 2008 American presidential election. *Journal of Experimental Social Psychology, 46,* 367-374.

Peduzzi, P., Concato, J., Kemper, E., Holford, T. R., & Feinstein, A. R. (1996). A simulation study of the number of events per variable in logistic regression analysis. *Journal of Clinical Epidemiology, 49*(12), 1373-1379.

Perkins, N. J., & Schisterman, E. F. (2006). The inconsistency of "optimal" cutpoints obtained using two criteria based on the receiver operating characteristic curve. *American Journal of Epidemiology, 163*(7), 670.

Pinter, B. (2002). Medico-legal aspects of abortion in Europe. *European Journal of Contraception and Reproductive Healthcare, 7*(1), 15-19.

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*(4), 621-640.

Prislin, R., & Wood, W. (2005). Social Influence in Attitudes and Attitude Change. In D. Albarracín, B. T. Johnson, & M. P. Zanna (Eds.), *the handbook of attitudes* (pp. 671-705). Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers.

Protection of Life during Pregnancy Act, Act 35, Bill 66, Minister for Health, Dáil Éireann (2013). Retrieved from: http://www.irishstatutebook.ie/eli/2013/act/35/enacted/en/html?q=Protection+of+Life+During+Pregnancy+Act+

Raccuia, L. (2016). Single-target implicit association tests (ST-IAT) predict voting behaviour of decided and undecided voters in Swiss referendums. *PloS One, 11*(10), e0163872-e0163872.

Randall, J. R., Rowe, B. H., Dong, K. A., Nock, M. K., & Colman, I. (2013). Assessment of self-harm risk using implicit thoughts. *Psychological Assessment, 25*(3), 714-721.

Regulation of Information (Services outside the State for Termination of Pregnancies) Act, Act 5, Bill 8, Minister for Health, Dáil Éireann (1995). Retrieved from: https://www.oireachtas.ie/en/bills/bill/1995/8/

Rhinehart, M. A. (2013). Abortions in Ireland: Reconciling a history of restrictive abortion practices with the European court of human rights' ruling in A., B. & C. v. Ireland. *Penn State Law Review, 117*(3), 959.

Richards, M. (1981). *Two to tango*. Dublin: Ward River Press.

Richetin, J., Costantini, G., Perugini, M., & Schönbrodt, F. (2015). Should we stop looking for a better scoring algorithm for handling implicit association test data? Test of the role of errors, extreme latencies treatment, scoring formula, and practice trials on reliability and validity. *PloS One, 10*(6), e0129601.

Ridgeway, I., Roche, B., Gavin, A., & Ruiz, M. R. (2010). Establishing and eliminating implicit association test effects in the laboratory: Extending the behaviour-analytic model of the IAT. *European Journal of Behaviour Analysis, 11*(2), 133-150.

Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J., & Müller, M. (2011). pROC: An open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinformatics, 12*(1), 77-77.

Roccato, M., & Zogmaister, C. (2010). Predicting the vote through implicit and explicit attitudes: A field research. *Political Psychology, 31*(2), 249-274.

Roche, B., & Barnes, D. (1996). Arbitrarily applicable relational responding and human sexual categorization: A critical test of the derived difference relation. *The Psychological Record, 46*, 451-475.

Roche, B., O'Reilly, A., Gavin, A., Ruiz, M., & Arancibia, G. (2012). Using behaviour-analytic implicit tests to assess sexual interests among normal and sex-offender populations. *Socioaffective Neuroscience & Psychology, 2*(1), 17335-17335.

Roche, B., Ruiz, M., O'Riordan, M., & Hand, K. (2005). A relational frame approach to the psychological assessment of sex offenders. In e. Quayle & m. Taylor, Viewing Child Pornography on the Internet: Understanding the Offence, Managing the Offender, and Helping the Victims (pp. 109-125). Dorset: Russell House Publishing.

Rönspies, J., Schmidt, A. F., Melnikova, A., Krumova, R., Zolfagari, A., & Banse, R. (2015). Indirect measurement of sexual orientation: Comparison of the implicit relational assessment procedure, viewing time, and choice reaction time tasks. *Archives of Sexual Behaviour, 44*(5), 1483-1492.

Rossiter, A. (2009). *Ireland's hidden diaspora: The 'abortion trail' and the making of a London-Irish underground, 1980-2000*. London: Iasc.

Rothermund, K., & Wentura, D. (2004). Underlying processes in the implicit association test: Dissociating salience from associations. *Journal of Experimental Psychology: General, 133*(2), 139-165.

Rudman, L. A. (2011). *Implicit measures for social and personality psychology*. London: SAGE Publications Ltd.

Rudolph, A., Schröder-Abé, M., Schütz, A., Gregg, A. P., & Sedikides, C. (2008). Through a glass, less darkly? Reassessing convergent and discriminant validity in measures of implicit self-esteem. *European Journal of Psychological Assessment, 24*(4), 273-281.

Saletan, W. (2003). *Bearing right: How conservatives won the abortion war.* Berkeley, CA: University of California Press.

Schwarz, N., & Bohner, G. (2001). The construction of attitudes. In A. Tesser & N. Schwarz (Eds.), Blackwell handbook of social psychology: Intraindividual processes (pp. 413–435). Oxford, England: Blackwell.

Schweppe, J. (2008). *The unborn child, article 40.3.3◦ and abortion in Ireland: Twenty-five years of protection?* Dublin: Liffey Press.

Sherman, S. J., Rose, J. S., Koch, K., Presson, C. C., & Chassin, L. (2003). Implicit and explicit attitudes toward cigarette smoking: The effects of context and motivation. *Journal of Social and Clinical Psychology, 22*(1), 13-39.

Sidman, M. (1971). Reading and auditory-visual equivalences. *Journal of Speech & Hearing Research, 14*(1), 5-13.

Sidman, M., & Tailby, W. (1982). Conditional discrimination vs. matching to sample: An expansion of the testing paradigm. *Journal of the Experimental Analysis of Behaviour, 37*(1), 5-22.

Skinner, B. F. (1937). The behaviour of organisms. Acton, Massachusetts.: Copley Publishing Group.

Skinner, B. F. (1959). *Cumulative record.* Acton, Massachusetts: Copley Publishing Group.

Smith, J. (2007). *Ireland's Magdalen Laundries and the nation's architecture of containment.* Indiana: University of Notre Dame Press.

Smith, J. M. (2004). The politics of sexual knowledge: The origins of Ireland's containment culture and the Carrigan report (1931). *Journal of the History of Sexuality, 13*(2), 208-233.

Smyth, L. (1998). Narratives of Irishness and the problem of abortion: The X case 1992. *Feminist Review, 60*(1), 61-83.

Spalding, L. R., & Hardin, C. D. (1999). Unconscious unease and self-handicapping: Behavioural consequences of individual differences in implicit and explicit self-esteem. *Psychological Science, 10*(6), 535-539.

Steffens, M. C. (2004). Is the implicit association test immune to faking? *Experimental Psychology, 51*(3), 165-179.

Streiner, D. L., & Cairney, J. (2007). What's under the ROC? An introduction to receiver operating characteristics curves. *Canadian Journal of Psychiatry 52(2),* 121.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18*(6), 643-662.

Swanson, J., Rudman, L., & Greenwald, A. (2001). Using the Implicit Association Test to investigate attitude-behaviour consistency for stigmatized behaviour. *Cognition and Emotion, 15,* 207-230

Sweetman, R. (1979). *On our backs: sexual attitudes in a changing Ireland.* London: Pan Books.

Teachman, B. A., Gregg, A. P., & Woody, S. R. (2001). Implicit associations for fear-relevant stimuli among individuals with snake and spider fears. *Journal of Abnormal Psychology, 110*(2), 226-235.

Teige, S., Schnabel, K., Banse, R., & Asendorpf, J. B. (2004). Assessment of Multiple Implicit Self-Concept Dimensions Using the Extrinsic Affective Simon Task (EAST). *European Journal of Personality, 18*(6), 495-520.

Tesser, A., & Shaffer, D. R. (1990). Attitudes and attitude change. *Annual Review of Psychology, 41*, 479-523.

The Offences against the Person Act, Act 24 & 25 Vict. c. 100 (1861). Retrieved from: http://www.irishstatutebook.ie/eli/1861/act/100/enacted/en/print.html

Törneke, N. (2010). Learning RFT: An introduction to relational frame theory and its clinical application. Oakland, CA, US: Context Press/New Harbinger Publications.

Victoir, A., Eertmans, A., den Bergh, O. V., & den Broucke, S. V. (2005). Learning to drive safely: Social-cognitive responses are predictive of performance rated by novice drivers and their instructors. *Transportation Research Part F: Psychology and Behaviour, 8*(1), 59-74.

Watson, J. B. (1925). *Behaviourism.* New York: Norton.

Watt, A., Keenan, M., Barnes, D., & Cairns, E. (1991). Social categorization and stimulus equivalence. *The Psychological Record, 41*(1), 33-50.

Weinberger, J. (2000). William James and the unconscious: Redressing a century-old misunderstanding. *Psychological Science*, *11*, 439–445.

Whelan, R. (2008). Effective Analysis of Reaction Time Data. *The Psychological Record, 58(3)*, 475-482.

White, K. M., Hogg, M. A., & Terry, D. J. (2002). Improving attitude-behaviour correspondence through exposure to normative support from a salient ingroup. *Basic and Applied Social Psychology, 24*(2), 91-103.

Whitty, N. (1993). Law and the regulation of reproduction in Ireland: 1922-1992. *The University of Toronto Law Journal, 43*(4), 851-888.

Wilentz, G. (2016). The importance of European standards and a human rights-based approach in strengthening the implementation of sexuality education in Ireland. *Sex Education, 16*(4), 439-445.

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.

Wolff, N., von Hippel, C., Brener, L., & von Hippel, W. (2015). Implicit identification with drug and alcohol use predicts retention in residential rehabilitation programs. *Psychology of Addictive Behaviours, 29*(1), 136-141.

Youngstrom, E. A. (2014). A primer on receiver operating characteristic analysis and diagnostic efficiency statistics for pediatric psychology: We are ready to ROC. *Journal of Pediatric Psychology, 39*(2), 204-221.

Zweig, M. H., & Campbell, G. (1993). Receiver-operating characteristic (ROC) plots: a fundamental evaluation tool in clinical medicine. *Clinical Chemistry, 39*, 561-577.

Appendix I

Information Sheet

This research is being conducted by Isabella Lalor (contact *Isabella.lalor.2015@mumail.ie*), a postgraduate student at the Department of Psychology, Maynooth University, under the supervision of Dr. Bryan Roche (contact: *Bryan.T.Roche@mu.ie* / +353 (1) 708 6026). It is the responsibility of this student to adhere to professional ethical guidelines in their dealings with participants and the collection and handling of data. If you have any concerns about participation you may refuse to participate, or withdraw at any stage.

This study involves examining the effectiveness of psychological computer-based implicit tests, which are purported to measure the strength of word associations. This in turn can serve as a measure of someone's attitude towards any topic. In this study, the topic of interest is abortion. There are two implicit tests to be completed in this study. Both are very similar to each other. One is called the Function Acquisition Speed Test (FAST) and the other is called the Implicit Association Test (IAT). Both involve the presentation of individual words on a computer screen, one at a time, and each test requires you to press a left hand or right hand computer key when you see these words, depending on instructions provided by the software. The software records your speed and accuracy at responding as instructed on each task.

The words that will appear on the screen are either positive words (e.g., Good), negative words (e.g., bad), pro-life related terms (e.g., Carry to term) or abortion related terms (e.g., terminate pregnancy). You should be aware that several such words will be presented repeatedly as you learn to press the appropriate keyboard keys. If you would find this in any way unacceptable or overly upsetting for any reason please simply decline to participate in the study.

As part of the experiment you will also be asked to identity your gender, your age and religious affiliation if any. You will also be asked to rate on a scale how religious you are, how morally conservative you are, how you feel about abortion (positive or negative), and how you voted in the recent referendum on the 8th Amendment (or how you intended to vote if you did not). Finally, you will be asked to complete two questionnaires regarding attitudes to abortion and moral issues. These tests will take approximately 30 minutes in total to complete, where you will also be given the option to take short breaks in between these tests. Again, if completing these scales and questionnaires might be unacceptable to you or overly upsetting for any reason, please simply decline to participate in this study.

All data from the study will be treated confidentially and data will not be identified by name at any stage of the data analysis or in the final report. The data will be compiled, analysed at a group level, and submitted in a postgraduate thesis. This data may also be used in as part of analyses for a scientific publication. All data collected will be retained in a locked filing cabinet drawer and / or digitally or on a University computer in the Department of Psychology for the duration of 10 years as per University regulations. No personally identifying information will be stored in any form.

At the conclusion of your participation, any questions or concerns you have will be fully addressed. You will be provided with your scores on each test on request at the end of the experiment. However, we cannot provide this information after you leave the experimental setting because your name, or any other identifier will not be linked to your data. Furthermore, we will not be in a position to interpret these scores for you because (a) FAST and IAT scores are poorly understood, and making sense of these scores is the aim current study, and (b) the questionnaire scores are not linked to any meaningful scale so that it is not possible to say that a person is scoring high or low relative to their peers in an Irish context. Nevertheless, we will endeavour to explain what these tests are hoping to measure and how the research study was designed and what questions it is attempting to answer.

Participants should be over 18 years of age, and should not suffer from any known condition that will make concentrating or problem solving difficult in any way.

You are of course free to terminate your participation in this study at any stage and have the right to ask for any data provided to not be used as part of the data analysis. However, once you have left the research setting withdrawal of you data will not be possible due to the fact that your name or any other identifiers will not be linked to your data in any way.

It must be recognized that, in some circumstances, confidentiality of research data and records may be overridden by courts in the event of litigation or in the course of investigation by lawful authority. In such circumstances the University will take all reasonable steps within law to ensure that confidentiality is maintained to the greatest possible extent.

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.

If you are upset by any issues related to abortion or if it has affected you in a way that requires assistance from someone supportive, you can contact the Crisis Pregnancy Helpline at **1850 6733 33**. For help after an abortion you can contact the International Helpline 24/7 at **1 866 482 5433**.

Appendix II

Consent Form

This research is being conducted by Isabella Lalor, a postgraduate student at the Department of Psychology, Maynooth University. The method proposed for this research project has been approved in principle by the Departmental Ethics Committee, which means that the Committee does not have concerns about the procedure. It is the responsibility of the researcher to adhere to ethical guidelines in dealing with participants and the collection and handling of data. If you have any concerns about participation you may refuse to participate or withdraw at any stage. You should understand, however, that you cannot withdraw your data after you have left the research setting as it will be fully anonymised and will therefore be irretrievable.

In this study we will ask you to complete a short questionnaire, and two computer-based implicit tests. You will be given the option to take a short break in between these tasks if you wish. If you are under the age of 18, have photo epilepsy, or feel uncomfortable with the topic of this research, or for any other reason wish to not participate, please inform the experimenter and the study will end now. Please self-exclude also if English is not your first language, or if you have vision difficulties that cannot be corrected with spectacles.

Data will be anonymised and analysed only on a group level. This means that your name will not be recorded or appear in relation to any information gleaned from the research. When you start you will be assigned a code number and this will be used to identify all of your data. All of the data collected in this study will be aggregated and will be included in a final thesis report completed by the researcher. This thesis will be submitted to the University. There is a possibility that the research will be published in a research/science journal, but no participant will be named and data will be reported only at a group level.

Your participation in this study will require approximately 30 minutes. If you have any concerns or queries about the study during or after participation you can ask me directly or ask the Supervisor named on the information sheet which you can take with you upon completion of the study. While you are entitled to know your scores on the questionnaires that you will complete, these scores will not be interpreted for you.

Please tick each statement below:

The purpose and nature of the study has been explained to me verbally & in writing. I've been able to ask questions, which were answered satisfactorily. ☐

I am participating voluntarily. ☐

I am over 18 years of age and do not suffer from any medical condition, such as photo epilepsy. ☐

I understand that I can withdraw from the study, without repercussions, at any time, whether that is before it starts or while I am participating. ☐

I understand that I can withdraw permission to use the data immediately following the study, but not thereafter as the data will be anonymized and analyzed at a group level. ☐

I understand the limits of confidentiality as described in the information sheet. ☐

I understand that my data, in an anonymous format, may be used in further research projects and any subsequent publications if I give permission below: ☐

I am aware of the sensitive nature of the topic being studied in the questionnaire and in the computer task portion of the study. ☐

I consent and am happy to complete the: Computer based test components ☐
Attitudes to abortion questionnaire ☐
Short personal values and demographic information survey ☐

Signed……………………………………. Date………………

Participant Name in block capitals ……………………………………………...

I the undersigned have taken the time to fully explain to the above participant the nature and purpose of this study in a manner that they could understand. I have explained the risks involved as well as the possible benefits. I have invited them to ask questions on any aspect of the study that concerned them.

Signed……………………………………. Date………………

Researcher Name in block capitals ……………………………………………...

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 Maynooth University Ethics Committee at research.ethics@mu.ie or +353 (0)1 708 6019. Please be assured that your concerns will be dealt with in a sensitive manner.

For your information the Data Controller for this research project is Maynooth University, Maynooth, Co. Kildare. Maynooth University Data Protection officer is Ann McKeon in Humanity house, room 17, who can be contacted at ann.mckeon@mu.ie. Maynooth University Data Privacy policies can be found at https://www.maynoothuniversity.ie/data-protection.

Appendix III

Attitudes to Abortion Questionnaire

Participant Number _______

Section 1 – Abortion Opinions
This section deals with your attitudes on abortion. Choose the answer that best describes your opinions by circling the appropriate number on your answer sheet. Please leave a question blank if you prefer not to answer.

Strongly Disagree = 1	Disagree = 2	Neutral = 3	Agree = 4	Strongly Agree = 5

1. Life begins at conception.

   1 2 3 4 5

2. Life begins at birth.

   1 2 3 4 5

3. I want children in the future.

   1 2 3 4 5

4. I might abort a foetus I did not intend to create.

   1 2 3 4 5

5. I always use birth control.

   1 2 3 4 5

6. I could be a responsible parent if I had a child now.

   1 2 3 4 5

7. Abortion should be legal in all situations.

   1 2 3 4 5

8. Abortion should be illegal in all situations.

   1 2 3 4 5

9. Abortion should be legal in the cases of rape or incest.

   1 2 3 4 5

10. Abortion should be legal if the mother’s life or long-term health is at risk.

    1 2 3 4 5

11. Abortion should be legal if the foetus has a birth defect.

    1 2 3 4 5

12. Abortion should be legal if the parents cannot afford the baby.

    1 2 3 4 5

13. Abortion should be legal if the parents do not want the particular sex of the child.

    1 2 3 4 5

14. Abortion should be legal if the parents do not want the child.

    1 2 3 4 5

15. Abortion is morally wrong.

    1 2 3 4 5

16. Women should make the decision for or against abortion since it is their bodies.

    1 2 3 4 5

17. Men and women should have equal influence regarding abortion.

    1 2 3 4 5

18. Abortion is murder.

    1 2 3 4 5

19. I am Pro-Choice.

    1 2 3 4 5

20. I am Pro-Life.

    1 2 3 4 5

21. My religious beliefs influence my opinions regarding abortion.

    1 2 3 4 5

22. I am influenced by family members' opinions about abortion.

    1 2 3 4 5

23. I am influenced by friend's opinions about abortion.

    1 2 3 4 5

Section 2 –Demographic Survey

This section deals with other factors that may influence opinions on abortion. Where appropriate, tick the box next to your answer. Please leave a question blank if you prefer not to answer.

24. My age to the nearest 6 months is: _____

25. My sex is:

    Female ☐ Male ☐ Other ☐

26. My vote (or my intended vote) for the 2018 referendum on abortion was:

    Yes ☐ No ☐

    Undecided ☐ I do not remember ☐

27. I voted according to my own moral beliefs:

    Yes ☐ No ☐

28. Choose the religious group you affiliate with:

Buddhist	☐	Muslim	☐
Catholic	☐	Protestant	☐
Hindi	☐	None	☐
Jewish	☐	Other	☐

29. How often do you attend activities at a place of worship each month:

    Everyday ☐

    Over once per week ☐

    Once a week ☐

    Once or twice per month ☐

    Only on holidays and occasions ☐

    I never go to a place of worship ☐

30. Rate how religious you are:

    Very religious ☐

    Moderately religious ☐

    Slightly religious ☐

    Not religious ☐

31. I identify most with:

    Liberals (left-Wing) ☐

    Conservatives (Right Wing) ☐

    Centre Parties ☐

    Nationalists/Socialists (far left) ☐

    Far Right ☐

    Green ☐

    Religious ☐

32. What region of Ireland are you originally from:

Leinster	☐	Munster	☐
Ulster	☐	Connaught	☐

    I am not originally from Ireland ☐

Appendix IV

RFD Scores from the FAST for Group One (Pro-Choice) and Group Two (Pro-Life)

Participant	Group	RFD Score	Participant	Group	RFD Score
1	2	-12.510644	**31**	1	-14.3287409
2	2	-5.1049418	**32**	1	.8901940
3	2	-10.289225	**33**	1	-1.8481229
4	1	-1.3727582	**34**	1	25.0221613
5	1	-4.5968151	**35**	1	2.0381024
6	2	-4.0451506	**36**	1	-15.4815094
7	1	-4.5890599	**37**	1	-1.8711017
8	1	-2.7611781	**38**	1	.6442831
9	2	-9.5503231	**39**	1	-1.4746896
10	1	3.2019056	**40**	1	-9.1208235
11	1	-2.4994849	**41**	2	-9.9554956
12	1	6.4599266	**42**	1	3.5225209
13	2	-4.3226119	**43**	2	-10.5023667
14	1	-1.7822512	**44**	1	-3.9252831
15	1	1.1767887	**45**	1	-1.8250821
16	2	-3.5580910	**46**	1	-10.5378324
17	2	-13.2179736	**47**	1	-1.7004405
18	2	-16.8204342	**48**	1	-10.3656670
19	1	7.8840239	**49**	1	-9.1681641
20	1	13.3762952	**50**	1	-3.4526018
21	1	-1.0779511	**51**	1	-.8610149
22	2	-3.3673240	**52**	1	2.1842680
23	1	-.1289819	**53**	1	-10.4619368
24	1	4.6596716	**54**	2	-2.0131200
25	1	-12.2948978	**55**	1	3.1989971
26	1	-1.1448007	**56**	1	.7635672
27	2	3.1960825	**57**	1	2.2046687
28	2	-14.0311550	**58**	2	-11.4808419
29	1	-7.1294717	**59**	1	-11.2917905
30	2	11.9220573	**60**	2	-7.4741367

Appendix V

Coordinates of the Curve

Test Result Variable(s): FAST RFD Score in relation to voting behaviour

Positive if Greater Than or Equal To[a]	Sensitivity	1 - Specificity
-17.820434222	1.000	1.000
-16.150971796	1.000	.923
-14.905125143	.979	.923
-14.179947963	.979	.846
-13.624564287	.979	.769
-12.864309008	.979	.692
-12.402771131	.979	.615
-11.887869844	.957	.615
-11.386316191	.957	.538
-10.914811452	.936	.538
-10.520099539	.915	.538
-10.482151761	.915	.462
-10.413801905	.894	.462
-10.327446281	.872	.462
-10.122360599	.851	.462
-9.752909351	.830	.462
-9.359243576	.830	.385
-9.144493765	.809	.385
-8.297480084	.787	.385
-7.301804199	.787	.308
-6.117206769	.766	.308
-4.850878468	.766	.231
-4.592937479	.745	.231
-4.455835878	.723	.231
-4.183881246	.702	.231
-3.985216861	.681	.231
-3.741687075	.660	.231
-3.505346384	.660	.154
-3.409962875	.638	.154
-3.064251042	.638	.077
-2.630331490	.617	.077
-2.256302434	.596	.077
-1.942110862	.574	.077
-1.859612349	.553	.077
-1.836602519	.532	.077
-1.803666668	.511	.077
-1.741345880	.489	.077
-1.587565039	.468	.077
-1.423723883	.447	.077
-1.258779430	.426	.077
-1.111375900	.404	.077
-.969483014	.383	.077
-.494998389	.362	.077
.257650617	.340	.077
.703925160	.319	.077

.826880601	.298	.077
1.033491341	.277	.077
1.607445522	.255	.077
2.111185177	.234	.077
2.194468332	.213	.077
2.700375597	.191	.077
3.197539831	.170	.077
3.200451367	.149	.077
3.362213226	.128	.077
4.091096250	.106	.077
5.559799094	.085	.077
7.171975206	.064	.077
9.903040605	.043	.077
12.649176300	.043	.000
19.199228285	.021	.000
26.022161319	.000	.000

a. The smallest cutoff value is the minimum observed test value minus 1, and the largest cutoff value is the maximum observed test value plus 1. All the other cutoff values are the averages of two consecutive ordered observed test values.

Appendix VI

Coordinates of the Curve

Test Result Variable(s): FAST D Score in relation to voting behaviour

Positive if Greater Than or Equal To[a]	Sensitivity	1 - Specificity
-2.250000000	1.000	1.000
-1.140000000	1.000	.923
-.980000000	.979	.923
-.905000000	.979	.769
-.865000000	.957	.769
-.835000000	.957	.692
-.810000000	.936	.615
-.765000000	.915	.615
-.720000000	.915	.538
-.695000000	.915	.462
-.665000000	.894	.462
-.640000000	.894	.385
-.620000000	.872	.385
-.590000000	.851	.385
-.555000000	.830	.385
-.535000000	.809	.385
-.525000000	.787	.385
-.515000000	.766	.385
-.470000000	.745	.385
-.425000000	.723	.385
-.415000000	.681	.385
-.395000000	.681	.308
-.370000000	.660	.308
-.355000000	.617	.308
-.345000000	.617	.231
-.325000000	.596	.231
-.305000000	.574	.231
-.295000000	.553	.231
-.265000000	.532	.154
-.230000000	.532	.077
-.190000000	.511	.077
-.150000000	.489	.077
-.130000000	.468	.077
-.110000000	.447	.077
-.080000000	.426	.077
-.030000000	.404	.077
.040000000	.383	.077
.090000000	.362	.077
.125000000	.340	.077
.155000000	.319	.077
.165000000	.298	.077
.185000000	.277	.077

.205000000	.255	.077
.215000000	.234	.077
.255000000	.213	.077
.335000000	.170	.077
.385000000	.149	.077
.395000000	.128	.077
.420000000	.106	.077
.470000000	.085	.077
.550000000	.064	.077
.610000000	.043	.077
.690000000	.043	.000
.790000000	.021	.000
1.820000000	.000	.000

The test result variable(s): FASTD has at least one tie between the positive actual state group and the negative actual state group.
a. The smallest cutoff value is the minimum observed test value minus 1, and the largest cutoff value is the maximum observed test value plus 1. All the other cutoff values are the averages of two consecutive ordered observed test values.

Appendix VII

Coordinates of the Curve

Test Result Variable(s): FAST Slope Score in relation to voting behaviour

Positive if Greater Than or Equal To[a]	Sensitivity	1 - Specificity
-1.1680	1.000	1.000
-.1606	1.000	.923
-.1517	1.000	.846
-.1471	1.000	.769
-.1392	.979	.769
-.1339	.957	.769
-.1324	.957	.692
-.1314	.936	.692
-.1304	.936	.615
-.1252	.915	.615
-.1204	.915	.538
-.1173	.894	.538
-.1122	.872	.538
-.1091	.851	.538
-.1081	.830	.538
-.1073	.830	.462
-.1062	.830	.385
-.0979	.809	.385
-.0850	.809	.308
-.0793	.787	.308
-.0783	.766	.308
-.0733	.745	.308
-.0649	.723	.308
-.0572	.723	.231
-.0513	.723	.154
-.0459	.702	.154
-.0417	.681	.154
-.0368	.660	.154
-.0332	.638	.154
-.0323	.617	.154
-.0295	.596	.154
-.0269	.574	.154
-.0264	.553	.154
-.0256	.553	.077
-.0238	.532	.077
-.0199	.511	.077
-.0166	.489	.077
-.0160	.468	.077
-.0157	.447	.077
-.0155	.426	.077
-.0123	.404	.077
-.0062	.383	.077

-.0026	.362	.077
-.0007	.340	.077
.0063	.319	.077
.0139	.298	.077
.0205	.277	.077
.0258	.255	.077
.0280	.234	.077
.0311	.213	.077
.0336	.191	.077
.0350	.170	.077
.0443	.149	.077
.0536	.128	.077
.0557	.106	.077
.0599	.085	.077
.0794	.064	.077
.1139	.043	.077
.1481	.043	.000
.2075	.021	.000
1.2501	.000	.000

a. The smallest cutoff value is the minimum observed test value minus 1, and the largest cutoff value is the maximum observed test value plus 1. All the other cutoff values are the averages of two consecutive ordered observed test values.

Appendix VIII

Coordinates of the Curve

Test Result Variable(s): FAST RFD Score in relation to group affiliation

Positive if Greater Than or Equal To[a]	Sensitivity	1 - Specificity
-17.820434222	1.000	1.000
-16.150971796	1.000	.944
-14.905125143	.976	.944
-14.179947963	.952	.944
-13.624564287	.952	.889
-12.864309008	.952	.833
-12.402771131	.952	.778
-11.887869844	.929	.778
-11.386316191	.929	.722
-10.914811452	.905	.722
-10.520099539	.881	.722
-10.482151761	.881	.667
-10.413801905	.857	.667
-10.327446281	.833	.667
-10.122360599	.833	.611
-9.752909351	.833	.556
-9.359243576	.833	.500
-9.144493765	.810	.500
-8.297480084	.786	.500
-7.301804199	.786	.444
-6.117206769	.762	.444
-4.850878468	.762	.389
-4.592937479	.738	.389
-4.455835878	.714	.389
-4.183881246	.714	.333
-3.985216861	.714	.278
-3.741687075	.690	.278
-3.505346384	.690	.222
-3.409962875	.667	.222
-3.064251042	.667	.167
-2.630331490	.643	.167
-2.256302434	.619	.167
-1.942110862	.619	.111
-1.859612349	.595	.111
-1.836602519	.571	.111
-1.803666668	.548	.111
-1.741345880	.524	.111
-1.587565039	.500	.111
-1.423723883	.476	.111
-1.258779430	.452	.111
-1.111375900	.429	.111
-.969483014	.405	.111

-.494998389	.381	.111
.257650617	.357	.111
.703925160	.333	.111
.826880601	.310	.111
1.033491341	.286	.111
1.607445522	.262	.111
2.111185177	.238	.111
2.194468332	.214	.111
2.700375597	.190	.111
3.197539831	.190	.056
3.200451367	.167	.056
3.362213226	.143	.056
4.091096250	.119	.056
5.559799094	.095	.056
7.171975206	.071	.056
9.903040605	.048	.056
12.649176300	.048	.000
19.199228285	.024	.000
26.022161319	.000	.000

a. The smallest cutoff value is the minimum observed test value minus 1, and the largest cutoff value is the maximum observed test value plus 1. All the other cutoff values are the averages of two consecutive ordered observed test values.

Appendix IX

Coordinates of the Curve

Test Result Variable(s): FAST D Score in relation to group affiliation

Positive if Greater Than or Equal To[a]	Sensitivity	1 - Specificity
-2.250000000	1.000	1.000
-1.140000000	1.000	.944
-.980000000	.976	.944
-.905000000	.976	.833
-.865000000	.976	.778
-.835000000	.952	.778
-.810000000	.952	.667
-.765000000	.929	.667
-.720000000	.929	.611
-.695000000	.929	.556
-.665000000	.929	.500
-.640000000	.929	.444
-.620000000	.905	.444
-.590000000	.881	.444
-.555000000	.857	.444
-.535000000	.833	.444
-.525000000	.833	.389
-.515000000	.810	.389
-.470000000	.786	.389
-.425000000	.762	.389
-.415000000	.714	.389
-.395000000	.714	.333
-.370000000	.690	.333
-.355000000	.643	.333
-.345000000	.643	.278
-.325000000	.619	.278
-.305000000	.595	.278
-.295000000	.571	.278
-.265000000	.571	.167
-.230000000	.571	.111
-.190000000	.548	.111
-.150000000	.524	.111
-.130000000	.500	.111
-.110000000	.476	.111
-.080000000	.452	.111
-.030000000	.429	.111
.040000000	.405	.111
.090000000	.381	.111
.125000000	.357	.111
.155000000	.333	.111
.165000000	.310	.111

.185000000	.286	.111
.205000000	.262	.111
.215000000	.238	.111
.255000000	.214	.111
.335000000	.167	.111
.385000000	.167	.056
.395000000	.143	.056
.420000000	.119	.056
.470000000	.095	.056
.550000000	.071	.056
.610000000	.048	.056
.690000000	.048	.000
.790000000	.024	.000
1.820000000	.000	.000

a. The smallest cutoff value is the minimum observed test value minus 1, and the largest cutoff value is the maximum observed test value plus 1. All the other cutoff values are the averages of two consecutive ordered observed test values.

Appendix X

Coordinates of the Curve

Test Result Variable(s): FAST Slope Score in relation to group affiliation

Positive if Greater Than or Equal To[a]	Sensitivity	1 - Specificity
-1.1680	1.000	1.000
-.1606	.976	1.000
-.1517	.976	.944
-.1471	.976	.889
-.1392	.952	.889
-.1339	.952	.833
-.1324	.952	.778
-.1314	.929	.778
-.1304	.929	.722
-.1252	.905	.722
-.1204	.905	.667
-.1173	.881	.667
-.1122	.881	.611
-.1091	.857	.611
-.1081	.833	.611
-.1073	.833	.556
-.1062	.833	.500
-.0979	.810	.500
-.0850	.810	.444
-.0793	.786	.444
-.0783	.786	.389
-.0733	.786	.333
-.0649	.762	.333
-.0572	.762	.278
-.0513	.762	.222
-.0459	.738	.222
-.0417	.714	.222
-.0368	.690	.222
-.0332	.667	.222
-.0323	.643	.222
-.0295	.643	.167
-.0269	.619	.167
-.0264	.595	.167
-.0256	.595	.111
-.0238	.571	.111
-.0199	.548	.111
-.0166	.524	.111
-.0160	.500	.111
-.0157	.476	.111
-.0155	.452	.111
-.0123	.429	.111

-.0062	.405	.111
-.0026	.381	.111
-.0007	.357	.111
.0063	.333	.111
.0139	.333	.056
.0205	.310	.056
.0258	.286	.056
.0280	.262	.056
.0311	.238	.056
.0336	.214	.056
.0350	.190	.056
.0443	.167	.056
.0536	.143	.056
.0557	.119	.056
.0599	.095	.056
.0794	.071	.056
.1139	.048	.056
.1481	.048	.000
.2075	.024	.000
1.2501	.000	.000

a. The smallest cutoff value is the minimum observed test value minus 1, and the largest cutoff value is the maximum observed test value plus 1. All the other cutoff values are the averages of two consecutive ordered observed test values.