Social Influences in Behavioral Rigidity:

Shared Attention in a Three-Alternative Choice Task

A Dissertation Presented for the

Doctor of Philosophy

Degree

The University of Tennessee, Knoxville

Andrés Hernando García-Penagos

December 2018

Dedication

I want to dedicate this dissertation to the many people who have been part of the last decade in my life and who have been in some way or the other a source of encouragement and support emotionally and intellectually for me. First and foremost, my wonderful wife, Carolina, and my parents, Maria Eugenia and Hernando, as well as my brother Diego and his beautiful family, have nurtured me with unconditional love and endured patiently all the absences and the many sacrifices required to finally achieve this final point in my education. They have always believed in me, especially in the most difficult of times, and I cannot thank them enough.

I also want to dedicate this dissertation to my two mentors in the University of Tennessee. I cannot put into words how much I have learned from John Malone. His conceptual clarity, his outstanding knowledge of the history of psychology, his unrelenting criticism of our discipline, his passion for knowledge and for teaching, his commitment to truth no matter how uncomfortable or inconvenient, these characteristics are all unlike any I have found in anyone before. I can only hope I get to be even half of the scholar he is. Moreover, John and his wife Carlene have always been as close to a family as we have had in this country, and they will always be in our hearts and thoughts.

I also have nothing but gratitude and admiration for my second mentor, Garriy Shteynberg, who adopted me into his lab despite our very different theoretical, conceptual, and methodological backgrounds. His genuine passion for research and his thirst for understanding and knowledge convinced me of seeking for his collaboration, and I couldn’t be more fortunate for having done so. I will always be thankful for all his patience and encouragement with my ideas, and his willingness to engage in genuine discussion.

This dissertation is also dedicated to the many wonderful teachers I have had throughout my life, and whose many teachings I hope are somehow

reflected in whatever is good in this document, in particular, Andy Lattal, Michael Perone, Telmo Peña-Correal, Luis Flórez Alarcón, Michael Nash, Todd Freeberg, Gordon Burghardt, Erin Hardin, the late Dan Hursh, and many, many others. I also have had the fortune of having incredible colleagues and friends, who have also influenced my thinking in more ways than they know: Arturo Clavijo, Toshi Kuroda, August Holtyn, Carlos Cançado, David Jarmolowicz, Allison Tetreault, Alicia Roca, Rogelio Escobar, Mohammed Al-Hammouri and Tariq Al-Dwaikat, Naoki Yamagishi, Elise Fles, Chris Silver, Michael Finn, David Book, to mention but a few. Also, to the many friends we have had throughout these years, from Bogotá to Morgantown (esp. Isa, Alejo, Hever and Paola) to Knoxville (in particular our wonderful friends in Knoxville, Alex, Ana, Anya, Daniela, Diana, Eli, Jackie, Mafe, Marisol, Nelly, Virginia, Zuania, Andrés, Camilo, Gabriel, Horus, Jan, Julian, Ludo, Melqui, Nelson, and so many others who have been there always for us, loving us and helping us in more ways than I can express. ¡Les estaremos agradecidos por siempre!

Finally, I want to dedicate this dissertation to Juliana, who has had to endure my lack of time for her. I truly don't think I could've done this without Carolina and Juliana's happiness as my goal, and their endless and unconditional love as my support. ¡Las amo!

Acknowledgements

The research herein presented couldn't have been conducted without the help of many people. I want to thank, in particular, the many undergraduate research assistants that participated as confederates and ran almost all of the experimental sessions: Hashem Alawami, Kelly Carlson, Brittney Clevenger, Shelby Cooper, Delaney Eldridge, Hannah Crawford, Nathaly Hunter, Tanner Jones, Taylor Kidd, Douglas Kievit, Siyu Lin, Gabrielle Schmittou, Nicole Matis, Zoe Mackenzie, Justine Reat, Nichole Schafer, Heather Smith, Samantha Smith, Keyara Taylor, Delaney Thurston, Madlein Atia, and Ian Warren. My wonderful wife, Carolina, also helped me with some of the data tabulation.

Dr. Garriy Shteynberg allowed me to be part of his lab and explore and modify his experimental protocol for the present research. He also provided incredible methodological guidance and theoretical discussion. I also received technical and intellectual support from my wonderful lab mates, James Bramlett, Elise Fles, and Sydney Michelson.

Dr. Lowell Gaertner was also incredibly helpful with some of our methodological and data analysis issues. Finally, all members of my dissertation committee, Drs. Freeberg, Malone, Olson, and Skinner, provided incredibly useful advice on the conduct of this research, and I want to thank them especially for this.

Abstract

The understanding of the factors and conditions under which humans and other animals fail to adjust to changing circumstances is still incomplete. One possibility has to do with the role of social influences that might favor or hinder adequate control by contingencies. Three experiments are presented in which one such factor was manipulated, the presence of a coattending individual. Shared attention has been found to enhance cognitive resources devoted to a task in a variety of experimental settings. Participants played a computerized version of the Rock-Paper-Scissors game which was programmed so that a particular choice would be more likely to win than the remaining two within a block. The allocation of behavior towards these choices was used as a measure of the participants' ability to track unannounced changes in the probabilities of the program's choices and considered a measure of sensitivity. Although, shared attention had little effect on the overall outcomes, it had unsystematic effects on the participants' strategy (Exp. 1). A significant effect of the sequence of block presentation was found when a progressively more discriminable sequence of blocks was compared to a progressively less discriminable one (Exp. 2). Early failure appeared to hinder effective control by the programmed probabilities. This effect is attributed to verbal behavior concurrent to performance. In support of this, providing participants with accurate instructions increased the overall accuracy, and these instructions were still followed to an extent when they became inaccurate (Exp. 3). Instructions had an immediate effect in the participant's strategy, reflected in a reduced tendency to shift strategy following losses and, particularly, draws. Further manipulations of shared attention are required to ascribe a particular role of this social phenomenon to the rigidity of behavior and other instances of suboptimal choice.

Table of Contents

1. Introduction 1
1.1. The Experimental Analysis of Rigidity and Flexibility 3
1.1.1 The experimental study of behavioral rigidity in humans 4
1.2. The Study of Rigidity in the Experimental Analysis of Behavior 5
1.3 Exploring Social Influences in Behavioral Rigidity: The Case of Shared Attention 10
2. Experiment 1 13
2.1 Method 14
2.1.1 Design 14
2.1.2 Participants 14
2.1.3 Apparatus and Setting 15
2.1.4 Procedure 15
2.2 Results 19
2.3 Discussion 25
3. Experiment 2: Elucidating Shared Attention Effects 27
3.1 Method 28
3.1.1 Design 28
3.1.2 Participants 29
3.1.3 Apparatus and Setting 29
3.1.4 Procedure 30
3.2 Results 31
3.2.1 Self-Report Data 35
3.3 Discussion 36
4. Experiment 3: Exploring the Effects of Verbal Instructions 39
4.1 Method 41
4.1.1. Design 41
4.1.2 Participants 42
4.1.3 Apparatus and Setting 42
4.1.4 Procedure 42
4.2 Results 45
4.2.1 Self-Report Data 47
4.3 Discussion 48
5. General Discussion 50
References 58
Appendix 67
Vita 107

List of Tables

Table 2.1: Shared Attention Conditions for the Pilot Experiment..... 68

Table 3.1: Experimental Conditions in Experiment 2..... 69

Table 3.2: Sensitivity and Bias Parameters of the GML Analysis for Pairs of Choices across Conditions..... 70

Table 3.3: Correlation Between Performance and Strategy Outcomes & Self-Report Measures..... 71

Table 4.1: Correlation Matrix for the TIPI Personality Measure and the Other Self-Report Measures in Experiment 3..... 72

Table 4.2: Correlation Matrix for the TIPI Personality Measure and the Performance and Strategy of Participants in Experiment 3..... 73

List of Figures

Figure 2.1: A prototypical trial configuration..... 74

Figure 2.2: Mean proportion of correct and incorrect attempts as a function of the type of block..... 75

Figure 2.3: Proportion of correct and incorrect attempts across blocks..... 76

Figure 2.4: Proportion of correct and incorrect attempts across block portions..... 77

Figure 2.5: Log response ratios as a function of log reinforcement ratios for participants in the Solo condition of Experiment 1..... 78

Figure 2.6: Log response ratios as a function of log reinforcement ratios for participants in the Shared-Outcome-Only condition of Experiment 1..... 79

Figure 2.7: Log response ratios as a function of log reinforcement ratios for participants in the Shared-Performance-&-Outcome condition of Experiment 1. 80

Figure 2.8: Estimated marginal means of the response times across blocks, across block types, and across block portion..... 81

Figure 2.9: Overall strategy at trial *n + 1* as a function of the selection at trial *n*..... 82

Figure 2.10: Strategy at trial *n + 1* as a function of the *outcome* of trial *n* 83

Figure 2.11: Strategy at trial *n + 1* as a function of the outcome of trial *n* across blocks..... 84

Figure 2.12: Strategy at trial *n + 1* as a function of the outcome of trial *n* across (non-random) block portion..... 85

Figure 2.13: Variations in bias and sensitivity to reinforcement according to the generalized matching equation..... 86

Figure 3.1: Mean proportion of overall correct and incorrect attempts as a function of the sequence of probability presentations..... 87

Figure 3.2: Proportion of correct and incorrect attempts across blocks..... 88

Figure 3.3: Proportion of correct and incorrect attempts across block portions of 15 trials each..... 89

Figure 3.4: Overall strategy at trial *n + 1* as a function of the selection at trial *n*..... 90

Figure 3.5: Strategy at trial *n + 1* as a function of the *selection* of trial *n* across blocks..... 91

Figure 3.6: Strategy at trial *n + 1* as a function of the *selection* of trial *n* across block portions..... 92

Figure 3.7: Overall strategy at trial *n + 1* as a function of the outcome of trial *n*..... 93

Figure 3.8: Strategy at trial *n + 1* as a function of the outcome of trial *n* for all participants..... 94

Figure 3.9: Strategy at trial *n + 1* as a function of the outcome of trial *n* across blocks..... 95

Figure 3.10: Log response ratios as a function of log reinforcement ratios for participants in the Solo condition..... 96

Figure 3.11: Log response ratios as a function of log reinforcement ratios for participants in the Shared Outcome Only condition..... 97

Figure 3.12: Log response ratios as a function of log reinforcement ratios for participants in the Shared Attention & Feedback condition..... 98

Figure 4.1: Mean proportion of overall correct and incorrect attempts..... 99

Figure 4.2: Mean proportion of overall correct and incorrect attempts across blocks..... 100

Figure 4.3: Overall strategy at trial *n + 1* as a function of the selection at trial *n*..... 101

Figure 4.4 Overall strategy at trial *n + 1* as a function of the selection at trial *n* across blocks..... 102

Figure 4.5: Overall strategy at trial *n + 1* as a function of the outcome of trial *n*..... 103

Figure 4.6: Strategy at trial *n + 1* as a function of the outcome of trial *n* across blocks..... 104

Figure 4.7: Proportions of correct and incorrect attempts across successive block portions of 10 trials..... 105

Figure 4.8: Proportions of strategies at trial *n + 1* as a function of trial *n* across successive block portions of 10 trials..... 106

Figure 4.9: Proportions of strategies at trial *n + 1* as a function of the outcome of trial *n* across successive block portions of 10 trials..... 107

1. Introduction

In the face of the ever-changing environmental conditions that animals have to continuously adapt to, it might seem a tautology that the ability to respond flexibly to such changes has adaptive value. It might seem less obvious that some behavioral rigidity is also of adaptive value, until one considers how inefficient our lives would be if we had to learn everything anew, if we had, for instance, to find our way from home to work every day (cf. James, 1890, chapter 4). These different aspects and gradations of adaptation are well represented in folk language in words like "tenacity," "stubbornness," "plasticity" or "stiffness" among many others, that focus the attention in the fuzzy borders at which persistence stops being useful and starts becoming a barrier.

Not surprisingly, then, psychology has been interested in the distinction for a very long time (see Schultz & Searleman, 2002, for a review of 100 years of the psychological study of rigidity and flexibility), and the distinction, fuzzy as it is, between positive and negative effects of rigidity in the face of changing environmental contingencies is arguably at the root of the old concept of "habit," which was central to the thought of the school of functionalism, inspired by William James and others. The related constructs of behavioral rigidity and behavioral flexibility have been studied under a number of different names almost from the start of psychology as an independent discipline: persistence, adherence, perseveration, stereotypy, compulsion, plasticity, and many others (cf. Schultz & Searleman, 2002).

Underlying these differing approaches is a concern with *optimal* behavior, the ability (or inability) to adapt to changing circumstances. It is perhaps due to the variety of guises and approaches under which the study of flexible-rigid behavior has been addressed that despite a long history, this area of inquiry remains largely unexplored, regardless of its importance

in the understanding of adaptive behavior in general, and its relevance to a number of applied concerns, including topics as diverse and apart as the nature of compulsive behavior in OCD-diagnosed patients, attitude change, compliance with clinical recommendations, child food neophobia, or obedient aggression, to name but a few.

With such a large variety of somewhat related phenomena it's not surprising that nothing close to a consensus has been arrived at the simple definition of both rigidity and flexibility. Ionescu (2011), for example, collected thirteen definitions of cognitive flexibility in the experimental literature between 2003 and 2010. Schultz and Searleman (2006) favor a conception that is akin to that advanced by Gestalt psychologists in the area of problem solving, with the critical element being that of a "mental or behavioral set, [...] a learned mental or behavioral pattern that forms through repeated experience in a given situation" (Schultz & Searleman, 2006, p. 170). Such mental sets, according to Schultz and Searleman, are further characterized by their "perseverance." In recent years, the emphasis has been on the related notion of "cognitive flexibility" (e.g., Ionescu, 2011), that covers aspects as different as switching among tasks and overcoming functional fixedness.

Furthermore, the study of the behavioral rigidity-flexibility continuum can be roughly divided in two broad traditions. First is the "trait" or correlational approach, exemplified by the early work of both Charles Spearman and R. B. Cattell, both of whom viewed "rigidity" as a personality trait interfering with effective behavior in novel circumstances (see Chown, 1959, for a review). As a "trait" approach, the construct is addressed by way of personality inventories and the like, and then contrasted with behavioral measures to determine whether, and to what extent, rigidity predicts clinically or otherwise relevant patterns of behavior (e.g., eating disorders, academic achievement, or response to treatment).

Although the Rigidity-as-trait approach is maybe fundamental for the comprehension of the phenomenon, a second topic of interest regarding the flexibility/rigidity continuum is what we could call the "Rigidity-as-state" approach that emphasizes the occurrence of rigid behavior as a function of contextual factors. Since the research herein presented can be categorized in this latter experimental tradition, it will be briefly reviewed next.

1.1. The Experimental Analysis of Rigidity and Flexibility

Insofar as it emphasizes the understanding of environmental factors promoting or hindering rigid behavior in particular circumstances, the "Rigidity-as-state" approach might be described more appropriately as "experimental," as contrasted with the more correlational emphasis of the trait approach. Methodologically, this experimental approach often involves an organism facing a task or problem that requires adjustment to changing circumstances, and is typical of the studies in behavioral flexibility in species other than humans, generally by using the "discrimination-reversal learning task" (Izquierdo & Jentsch, 2012; Klanker, Feenstra & Denys, 2013).

This task involves the simultaneous presence of two stimuli, a response to one of which (formally called an S+ or S^D) is typically followed by food, while a response to the other (the S— or S^Δ) has no associated consequence. Once the positive and negative discriminations are well established (generally after many sessions) the functions of the stimuli are reversed, such that that a response to the previously negative stimulus is now the one followed by food, and vice versa. The rigidity/flexibility of the discrimination behavior is operationalized as the time or number of trials it takes for the organism to adjust to the reversed stimulus functions (see Izquierdo & Jentsch, 2012, for a brief review).

Renewed interest in this task has appeared as a result of the relatively recent emphasis on individual consistency in patterns of behavior in the guise of "behavioral syndromes" (e.g., Sih, Bell & Johnson, 2004; Sih,

Bell, Johnson & Ziemba, 2004). The flexibility of behavior of numerous species has been tested with this procedure, and in a variety of sensory modalities and behaviors (see e.g., Bond, Kamil & Balda, 2007 for its use with corvids, and see McMillan, Hahn, Congdon, et al., 2017, for a recent study with chickadees).

In the case of humans, while the task has been used as a measure of rigidity, the interest has been typically in underpinning the neural mechanisms underlying the reversal rather than in the study of the environmental conditions that favor or hinder flexible behavior, and accordingly most of this research has focused on lesion studies conducted with clinical populations (e.g., Hornak, O'Doherty, Bramham, et al., 2004; Xue et al., 2013). Human participants with no neurological lesions typically show fast discrimination learning and fast reversal learning even during a single session (e.g., Ghahremani et al., 2010), unlike studies in other species which often require many daily sessions.

1.1.1 The experimental study of behavioral rigidity in humans. Due to the limitations of the discrimination reversal task to measure adaptation to changing contingencies in normally developing, non-clinical populations (and particularly the fact these individuals tend to excel readily in such tasks) alternative methods of studying rigid/flexible behavior in humans have been used. One set of approaches comes from studies on problem-solving by Gestalt psychologists. What is characteristic of these studies is that, after successful problem solving, a novel problem is introduced.

In Luchins and Luchins' 'water jugs problem' (Luchins, 1942; Luchins & Luchins, 1959), for example, participants are required to solve a number of easy arithmetical problems (calculating how to measure a particular amount of water by using three hypothetical jars of different volumes) all of which can be solved in the same way. Novel problems are then introduced, some of which can be solved in the usual way, but can also be solved in an easier manner,

and some of which can only be solved by using the easier method. Luchins and Luchins noticed that typically most participants who were exposed to the initial problems failed to use the simpler solution when the other one was available, and were mostly unable to find out the simpler solution in extinction problems. According to the Luchins, the first problem produced an *Einstellung effect*, a mental or problem-solving set that prevented participants from discovering a novel situation.

Similar concepts related to rigid behavior developed from Gestalt psychology include those of *functional fixedness* and *negative transfer.* More recently, and inspired by cognitive models another set of methods for the study of these patterns of behaviors has become increasingly common. One of them is the "attentional set shifting" paradigm, which involves participants learning a discrimination depending on some physical variable (e.g., color) and then being tested in novel discriminations that depend on a different stimulus dimension (e.g., size) (see Ionescu, 2011, for an overview of these methods; see also, Koch, Poljac, Müller & Kiesel, A., 2018, for a review of recent methodological advances in the context of so-called "multitasking"; Vandierendonck Liefooghe & Verbruggen, 2010)).

1.2. The Study of Rigidity in the Experimental Analysis of Behavior

A final field of unrelated research that can be included in this broad categorization of experimental approaches to rigid/flexible behavior corresponds to an experimental tradition that developed mostly in the 1980s and 1990s in the area of human operant research, within the larger field of the experimental analysis of behavior developed by B. F. Skinner and his students. These studies originated in applied, clinical contexts where verbal instructions were used to facilitate the acquisition of socially-appropriate behaviors (see Vaughan, 1989, for a historical review). Verbal instructions were generally found to be quite helpful in reducing the time required to learn these discriminations, but it was later found that they might also

prevent participants from detecting unannounced changes in the experimental contingencies (Hayes, Zettle & Rosenfarb, 1989).

These early results gave way to a number of experiments where participants received verbal instructions about the most effective manner to respond in operant learning tasks that involved pressing a lever or a button which produced points or some other consequence according to particular schedules of reinforcement, and then changing the schedule of reinforcement without announcement (see Hayes, Zettle & Rosenfarb, 1989, for a review). If participants' behavior remained unchanged despite the change in contingencies, as was often the case, their behavior was said to be *insensitive to contingencies* (e.g., Shimoff, Catania, and Matthews, 1981). For instance, in the experiments by Matthews et al. (1977), even minimal instructions were enough to generate patterns of responding that were incongruent with the scheduled contingencies.

A number of factors have been shown to alter sensitivity to changing contingencies in human performance under schedules of reinforcement. Galizio (1979) found that only when the schedule incorporated "response cost" (money loss) as a consequence for not responding, did subjects respond more sensitively, which led him to conclude that insensitivity occurs as a result of failing to come into contact with the contingencies: instructions might generate behavior that prevents participants' from experiencing the discrepancy between the instructions and current contingencies. For example, an instruction to respond rapidly (e.g., according to a Variable Ratio schedule) might engender such high rates of behavior that participants might not realize that, unannounced to them the schedule of reinforcement was changed to one of interval in which it is the timing of responding rather than its frequency that determines the delivery of a reinforcer.

In support of this view, when instructions are vague and imprecise but do not directly contradict the contingencies, responding is often insensitive

(Buskist, Bennett & Miller, 1981; Ninness & Ninness, 1998), and the degree of discrepancy between instructions and contingencies correlates highly with sensitivity (DeGrandpre & Buskist, 1991; Hackenberg & Joker, 1994; cf. Fox & Kyonka, 2017). Newman, Buffington and Hemmes (1995) found that regardless of the accuracy of the instructions, responding was only sensitive when every response was reinforced. A history of variable responding (e.g., having been exposed to a number of different schedules of reinforcement rather than to a single one) also increases sensitivity (Baumann, Abreu-Rodrigues & Souza, 2009; Joyce & Chase, 1990; LeFrancois, Chase & Joyce, 1988).

A second, complementary approach to this phenomenon emphasizes the role of social influences (Hayes, Brownstein, Haas & Greenway, 1986; Hayes, Brownstein, Zettle, Rosenfarb & Corn, 1986), arguing that instructions bring into play additional (extra-experimental) contingencies related to a social history of complying with instructions and norms. These extra-experimental contingencies "compete" with experimental contingencies for the control of behavior. In this way, it's not that the behavior is as a rule less sensitive to contingencies, but rather that it's sensitive to extra-experimental (often implicit) contingencies (e.g., that you're expected to be compliant in an experiment). Hence, when both verbal instructions and direct contingencies are present, as in the typical experiments, responding might be sensitive to both types of contingencies or just to one, depending on which proves to be most advantageous to a participant.

Further evidence for the role of social factors in the rigidity of behavior was found in a study by Cerutti (1994), where some participants were told that they were being videotaped while they responded in a typical instructed operant task, resulting in higher levels of insensitivity. Likewise, Kroger-Costa & Abreu-Rodrigues (2012) found that the presence of the experimenter throughout the experimental session produced insensitive responding, as compared to the responding of subjects who were alone.

Although this paradigm within the broader field of the experimental analysis of human behavior provided some insights and pointed to interesting aspects of the complex continuum of behavioral rigidity and flexibility, experimental attention has dwindled in the last two decades, and few of these studies have been conducted in this decade. There are many reasons for this decline, among which is the increased interest in the research areas of symbolic behavior and derived stimulus relations. Another possibility is that too much emphasis was placed on arguing that humans were "less sensitive" to contingencies (see e.g., Kollins, Newland, & Critchfield, 1997; Madden, Chase, & Joyce, 1998), which led to a number of debates on the nature of sensitivity and the procedural differences in the study of human and nonhuman behavior. Arguably, many researchers were happy with showing that human behavior appeared "sensitive" to contingencies under particular conditions; thus, instead of pursuing some of the issues raised with human research, the whole paradigm was almost forgotten, and the baby had been thrown with the bathwater.

Whatever the case, with the reduced interest in the paradigm a number of potentially interesting issues that could be considered highly relevant in the understanding of human behavior in laboratory conditions and that could be procedurally addressed were simply relegated. One instance of this, which is the focus of the present research, is the analysis of social influences in the relative "rigidization" of behavior that is observed as a result of presenting instructions. With the exception of the study of Kroger-Costa and Abreu-Rodrigues (2012), almost no research has attempted to identify the role of social variables in the sensitivity to changing experimental contingencies.

Moreover, since the emphasis in the original studies was in the typical patterns of behavior under particular schedules of reinforcement (for example, whether humans exhibit the scallop patterns observed in other

species when responding under Fixed Interval schedules), the apparatus were essentially computerized versions of a Skinnner box, with participants receiving points in exchange for pressing a button in the keyboard. There is nothing particularly problematic in this, but as a result of this technicality the rate of responses became the most common dependent variable across these experiments, with "sensitivity" being a synonym for changes (or lack thereof) in these rates as a function of unannounced changes in the schedules of reinforcement.

The rate of responding, although certainly important to understand the behavioral control of particular natural or artificial contingencies of reinforcement and punishment, is limited in terms of ecological validity in that the reinforcement (and punishment) of many human behaviors doesn't depend on the frequency of the behavior but rather on its timing, pertinence, or duration. As a consequence of this emphasis on traditional research with schedules of reinforcement, I'd argue, no study on instructional control has explored sensitivity in the context of choice behavior, a very unfortunate situation insofar as the traditional study of choice has proved successful in extending interdisciplinary research problems in areas as varied as behavioral ecology, consumer behavior and behavioral economics, or substance use disorders.

The research herein presented was conducted with the aim of addressing some of the gaps in the understanding of social influences on rigidity and flexibility of behavior in a context of choice. Although there are many social factors that might be relevant to these issues, we opted to start this inquiry by focusing on the role of the coattending presence of another person as one is engaged in a particular task.

1.3 Exploring Social Influences in Behavioral Rigidity: The Case of Shared Attention

Although, procedurally, there is much in common between the methodological approaches to behavioral rigidity/flexibility described earlier, no unified theoretical framework has been proposed to integrate these dissimilar bodies of research (Ionescu, 2011; Schultz & Searleman, 2002). More importantly, research about the variables controlling rigid/flexible behavior is still inconclusive, and has generated little interest. Renewed interest in the Einstellung effect, however, has demonstrated that one variable, relatively unexplored, has to do with the way expert problem-solvers allocate their attention. Bilalić, McLeod and Gobet (2008) evaluated whether chess expert and master class players would find an optimal solution to a number of chess problems which could also be solved by a familiar but suboptimal solution that they had used before. Results showed that most players opted for the suboptimal solution. Further, when their eye movements were tracked they found out that players kept directing their gaze to those squares related to the familiar solution, even though they reported to be looking for alternatives (see also Sheridan & Reingold, 2013). According to the authors, this suggests that the Einstellung effect occurs as a result of the participants' attention being biased towards those features associated with the familiar solution, hence preventing the discovery of new solutions (cf. Galizio, 1979).

A relatively recent body of research has emphasized the social nature of human attention and its important role in intra-group coordination (see Shteynberg, 2015, for a review). According to this perspective, the perception that one is co-attending with others to some aspect of the world ("shared attention") leads to directing greater cognitive resources to such aspects of the world, hence affecting other basic processes such as memory, motivation, and social learning, among others (Shteynberg, 2015). Shteynberg

(2010), for example, found evidence that performance in a number of tasks (e.g., a time pressure requirement, a word recognition task, or a painting recognition task) was higher for participants who thought that other individuals (that they perceived as being more similar to themselves) were co-attending to those tasks.[1]

If rigid behavior is related to differential patterns of attention as a result of previous success, it is conceivable that shared attention might differentially affect the sensitivity of behavior to changing consequences, an issue that has yet to be explored experimentally. The research herein described addresses the question of whether experimental manipulation of shared attention leads to either increased, decreased, or no changes in the sensitivity of behavior to changes in programmed contingencies.

In order to address this question we have opted for a task that requires constant adaptation to changing contingencies, and which provides a *quantitative* measure of sensitivity, as described by the generalized matching law (GML; see Appendix I for a brief review and a description of the quantitative model, and Figure 19 for a schematic representation of how data are analyzed in this model). The task involves three simultaneous alternatives, each associated with a different probability of reinforcement, such that participants must allocate their behavior according to such probabilities to respond optimally.

The task, developed by Kangas, Berry, Cassidy et al (2009) requires participants to play a computerized version of the Rock-Paper-Scissors game which has been programmed such that the probability of each alternative being the correct one changes across blocks within the experimental session. Kangas et al. (2009) found that the participants' behavior matched the programmed

[1] Shteynberg has argued that these effects are not mere social facilitation effects insofar as they occur only when participants consider themselves and co-attending others as belonging to the same group (i.e., as a "we").

probabilities across a number of ten 100-trial blocks, but that near-exclusive responses for the most probable alternative (i.e., overmatching) was found when *accurate* verbal instructions were provided. Lie, Baxter and Alsop (2013) used a brief four-block version of the same task finding again sensitivity to the changing probabilities, but reduced sensitivity when participants were told that they were playing against another human participant versus when they were told that they would be playing against the computer. It needs to be noted, though, that the literature on choice with human participants is rather contradictory, with many studies finding that the matching law described very accurately the choices of participants (see Pierce & Epling, 1983, for a review), and many others showing significant and consistent departures from it (e.g., Horne & Lowe, 1993; see also Takahashi & Shimakura, 1998).

In this research, three experiments were conducted in which participants took part in a modified version of the choice allocation task developed by Kangas et al. (2009). The important modification was that participants completed the task either in the presence (shared attention) or absence of another individual (a confederate in Experiments 1 and 2, and a second participant in Experiment 3). While Experiment 1 served as a pilot study and consisted in a systematic replication of the study by Lie, Baxter and Alsop (2013), including shared attention conditions, Experiment 2 analyzed social effects as a function of the sequence of presentation of experimental blocks, and Experiment 3 introduced accurate and inaccurate verbal instructions.

2. Experiment 1

As described in the previous section, a computerized version of the Rock, Paper, Scissors game was used to determine the allocation of behavior to each particular choice by the computer, whose strategy was programmed so that it played one option preferentially across most blocks. Using the shared attention paradigm developed by Shteynberg (2015), participants first completed a creativity task designed to increase the extent to which participants feel similar to the confederates. Participants were then randomly assigned to one of three experimental groups (see Table 2.1; all Tables can be found in the Appendix section). Participants in the first group (*Solo condition*) were presented with the task while they were alone, while participants in the other two groups did the task in the presence of confederates who could aurally co-attend to the performance of the participants.

The first experiment was conducted to determine the feasibility of the experimental procedure, and its procedure and results are presented next. Plausibly, the co-attending presence of another individual might enhance the salience of different aspects of the experimental task. Among such different factors, we decided to determine first whether shared attention to the consequences delivered to the participant's choices would have any effect on matching. Moreover, we wanted to determine whether enhanced attention to the participant's choice pattern of allocation (her/his "strategy") would have any subsequent effect in addition to or against that of shared attention on the strategy's outcome. Accordingly, participants in the shared-attention condition were divided in two further groups, differing in the specificity of the auditory feedback presented after each participant's choice: in the *Shared-outcome-only* condition confederates were only able to hear whether the participants' choice was correct or incorrect, while in the *Shared-performance-&-outcome* condition the auditory feedback also contained

information about the participants' selection (see Table 2.1 for a summary of the experimental conditions).

2.1 Method

2.1.1 Design

To determine the role of shared attention in the participants' sensitivity to the changing probabilities of one response being correct, participants were randomly assigned to one of three groups. Participants in the first group (*Alone condition*), completed the task by themselves. Participants in both the second group (*Shared-performance-only* condition) and third group (*Shared-performance-&-feedback* condition) completed the task while a confederate was present in the room, with both groups differing on the type of auditory feedback that was delivered by the computer during each trial, as described in the procedure section (see Table 2.1).

The probability that one of the three options (i.e., rock, paper, or scissors) was correct was modified across blocks. Just like in the Lie, Baxter and Alsop (2013) study, the present experiment involved four blocks. The probabilities used in the present study were similar to those used in the Lie et al., study as well, that is, one block where each alternative was as likely as the other two, and three blocks where the probabilities for each alternative were 55%, 30%, and 15%, respectively, with either alternative being the most likely on one block.

2.1.2 Participants

Forty-two undergraduate students participated in the study. They were randomly assigned to one of the three experimental conditions. Data from three of the participants were excluded from the analysis because of errors in the experimental session (e.g., a participant turned off the volume of the computer). The final data analysis, then, included thirty-nine participants, with 14 in the Alone condition (female N = 7, mean age = 18.79 years, sd =1.12), 12 in the Shared-Outcome-Only condition (female N = 6, mean age =

18.67 years, sd =0.98), and 13 in the Shared-Outcome-&-Performance condition (female N = 8, mean age = 18.08 years, sd =0.27). Participants received academic credit for participating in the experiment.

2.1.3 Apparatus and Setting

All experimental sessions were conducted in the Group Attention Laboratory. Participants sat in individual cubicles where a computer screen was located above a table. Undergraduate confederates assisting in the task were sat down in a close cubicle during the experimental session, so that they had no visual access to the participant's responding, but were able to hear the auditory feedback provided by the computer. The experimental task was programmed in the FlashDevelop (IDE) by using the open-source HaxeFlixel library, and it was presented to the participants in the standalone Adobe Flash Player 25. Auditory feedback was delivered by the computer speakers, located below the table in the cubicle where participants sat down, and were set up at the highest volume so that anyone could easily hear it from any part of the experimental room. The volume control options of the computer were blocked so that participants couldn't change the sound or mute the auditory feedback. Sessions lasted about 45 minutes.

2.1.4 Procedure

At the beginning of the session, participants were allowed into the lab and asked to read an informed consent form. Participants of Groups 2 and 3 sat next to an undergraduate confederate who acted as if she/he were an unknowing participant. Once they provided consent and any questions by them were answered, participants (and confederates) were asked to enter the experimental room and sit at a table where they were presented with a creativity task consisting of writing up as many uses for a brick as they could think of for a period of three minutes. Participants from Groups 2 and 3 had to complete this creativity task by collaborating with the

confederates, who had been previously trained to try to collaborate actively in the task.

After the creativity task was completed, the experimenter took the participant (and confederates in the shared conditions) to the cubicle, and the participant was asked to sit down. Confederates were asked to stand behind the participant so that the participants were aware that the confederate could see what occurred in the screen, and could hear the auditory feedback from the computer. As the participant sat in front of the computer, the experimenter read the following general instructions, which were also presented in the screen:

"You will be playing a game of 'Rock-Paper-Scissors' against the computer. The rules are the usual ones: Rock beats Scissors, but loses against Paper; Paper beats Rock, but loses against Scissors; and Scissors beats Paper, but loses against Rock. If you don't know how to play this game, please ask the experimenter to explain. You will make your selection by clicking the mouse cursor on the option you choose. The computer will make its selection at the exact same time you do. The winner of each round will receive one point, and the loser will receive 0 points. If there is a tie, the score will remain unchanged. Your task is to obtain more points than the computer, so try to respond in the way that would make you get the most points. Do you have any questions?"

If no questions were asked, the participants were then told: "Let's do some practice, please click continue on the screen. (…) Alright, now you'll play a few times until you find a message." This started a block of five practice trials where the computer selection of an alternative was completely random. The trials were identical to those in the actual blocks, as described below, with the sole difference that they also provided written feedback in the screen in the form of "Correct. Paper beats Rock" or "Wrong, Scissors beats Paper." At the end of the five trials a message appeared in the screen

asking if they had any questions. Then, participants were told that they could start whenever they wanted by clicking on "Continue," and confederates were asked to sit down in an adjacent cubicle, such that they could clearly hear the auditory feedback from the computer, but had no visual access to the participants responding.

At the start of each trial the participants were presented with three images corresponding to Rock, Paper, and Scissors at the left side of the screen, while a circle of a similar side with an interrogation mark was presented on the right side of the screen (see Figure 1). The specific type of auditory feedback that participants received after making a choice depended on the experimental condition.

For participants from Group 2 (shared-outcome only), once they selected one of the options from the screen the two remaining alternatives disappeared and after three seconds the computer selection appeared, replacing the previous circle. If the participant had made a correct response (e.g., if the participant chose Rock and the computer chose Scissors), a green rectangle appeared behind the participant's choice, and auditory feedback was played consisting of the sound of applause for a correct answer, and one point was added to the participant score counter. If there was a tie (e.g., both participant and computer choosing "Scissors") the word "tie" was played, and no change occurred in the scores. Finally, if they made a wrong choice (e.g. the participant choosing Paper and the computer choosing Scissors), a red rectangle appeared behind the choice, a "thunk" tone was played, and the computer's score would increase by one. After a 5-s inter-trial interval (ITI) the rectangle disappeared, the computer choice would revert to the circle with the interrogation mark, and a new trial started.

The trials for participants from Groups 1 (alone) and 3 (shared-performance and feedback) were identical, and differed from those for Group 2 only in the type of auditory feedback they received. Upon making a choice

participants of these conditions would also hear their choice (e.g., the word "Rock"), and after three seconds they would hear the computer choice (e.g., "Paper") followed by the sounds of applause, tie, or the thud tone as in the other condition. In this way, while confederates in Group 2 would only hear the tones of applause, tie, or the thud tone, those in Group 3 heard also what the player's and computer's choices were.

The session consisted of four blocks of 60 trials each, for a total of 240 trials.[2] At the end of each block the screen displayed the following message: "Take a moment to rest and click on this message when you are ready to continue." At the end of the last trial of the last block they were shown the message: "Thank you for your participation. Your score has been recorded automatically. Please let the research assistant know that you are finished." Participants then were asked to complete a brief questionnaire asking for some demographic information and whether they used any particular strategy or game plan, and then were debriefed on the nature of the study and were allowed to leave.

The computer's responses were preprogrammed so that within each block of 60 trials the computer chose each particular alternative with a different probability. The order in which each block was presented was randomized by the computer. For one block, the probability of the computer choosing each alternative was the same (i.e., 33%). For the remaining three blocks the probabilities for each alternative were 55%, 30%, and 15%, with each of the three options being the most likely one in each of the blocks. We opted for these rather moderate probabilities as we wanted to probe the effect of shared attention in isolation from salient differences across blocks and also because they were similar to those used in the Lie et al. (2013) study.

[2] Lie, Baxter & Alsop (2013) presented their participants with significantly longer blocks consisting of 126 trials. We first had 100 trials per block, but the sessions lasted about one hour and a half, which we considered excessive.

Probabilities were automatically calculated by the computer for each trial, which meant that the actual number of trials only approximated the preprogrammed probabilities (i.e., the actual number of trials of one type was almost never exactly the same across all participants).

2.2 Results

The performance of participants will be analyzed separately into two types of blocks. One type of blocks comprises those where the computer chose any of the three alternatives with the same probability (i.e., .33; hereafter, random blocks), while the other type comprises those blocks where the computer chose one alternative with a higher probability (hereafter, non-random blocks). The task was programmed so that the order at which the random block would appear would be random. Across blocks participants would be expected to systematically select one alternative more than the others in non-random blocks.

Importantly, insofar as the selection of one particular alternative across a session might be reasonably inferred to depend not only on the current block probabilities but also on choice during previous blocks, we have considered that a multilevel modeling analysis might be the one most suitable to detect significant differences, if any, across experimental conditions. The participant sample size used in this first experiment (N = 42) is unfortunately too small and violates a number of assumptions of the model, and hence a finer-graded statistical analysis was not possible. Simpler analyses, like a chi-square test, are also affected by the fact that responses across blocks are not uncorrelated, and hence no statistical data about the differences across conditions will be provided in the analysis of these data.

Figure 2.2 depicts the overall proportions of correct and incorrect (i.e., losses and draws) responses for participants in the three experimental conditions. As can be seen, and as expected, the proportion of correct

responses in random blocks was around .33 across all three conditions. As for non-random blocks, and also as expected, the number of correct attempts was slightly higher across all three conditions, although in no case did the proportion go above .4. Shared attention conditions seemed to obtain slightly more correct responses than the Alone condition. Differences across groups in these overall proportions are analyzed in more detail in Figures 2.3 and 2.4.

Figure 2.3 depicts the proportion of wins, losses, and draws across blocks as the session progresses. As shown in the upper panel, although large variability is observed in the proportions of each outcome depending on the order at which the random block was presented, no systematic effect of block order can be inferred. More interestingly, it might have been expected that as participants gained experience with non-random blocks they might obtain higher numbers of wins towards the end of the session. As the bottom panel of Figure 3 shows, however, that doesn't seem to be the case. While differences seem to occur reliably across blocks, no linear trend can be inferred from this performance.

Rather than looking at the outcomes across blocks, Figure 2.4 depicts the proportion of each outcome across block portions. Since each block had 60 trials, each bar in Figure 2.4 represents an average of the first, second, third, and fourth fifteen trials of all blocks. As observed, no systematic trend is evident during random blocks (as should be expected), but in non-random blocks a slight trend is observed whereby wins were slightly more likely to occur in the last two portions of each block. Although moderate, these results suggest some adjustment to contingencies. Also importantly, no large systematic differences were observed across shared attention conditions.

Since this study closely followed the Lie et al. (2013) study, in regard to the number of blocks and the probabilities tested, we decided to analyze the present data in the same way as they did (see Figure 13 and

Appendix I for a summary of matching law analysis). Namely, we obtained log response ratios and log obtained reinforcer ratios for each possible pair of alternatives (paper vs. scissors, paper vs. rock, and rock vs. scissors) rather than by comparing the ratios of a single alternative to those of the other two alternatives together. For example, the response ratio of paper vs. scissors in the first block corresponded to the logarithm of the number of "rock" responses in the first block divided by the number of "rock" *and* "scissors" responses. Correspondingly, the obtained reinforcer ratio of paper vs. scissors in the first block corresponds to the log number of reinforcers (i.e., correct attempts) obtained when choosing "rock" in block one divided by the number of reinforcers obtained by choosing both "rock" and scissors."

Just like in the Lie et al. study, in those instances where participants never selected one alternative in a whole block, the value of 0 (zero) was replaced by 0.25, so that the logarithm was calculable. Furthermore, in their study they only calculated ratios for the last half of trials of each block, so as to allow participants enough contact with the programmed contingencies. We calculated the ratios for both the whole block and the last half. Figures 2.5, 2.6 and 2.7, accordingly, depict the log response ratios of each alternative as a function of the log programmed reinforcement ratios for all trials (left panels) and for just the latter half of trials (right panels) for each of the conditions. Figure 2.5 shows the results for participants in the Alone condition, while Figures 2.6 and 2.7 show the results for participants in the shared attention conditions, that is, the Shared-outcome-only and the Shared-performance-and-outcome conditions, respectively.

As expected from the proportion data above described, matching was not consistently observed across groups. The obtained R^2 coefficients were considerably smaller than those obtained in the Kangas et al. (2009) study (where they were higher than .9 in all cases), and those in the most similar

Lie et al. (2003) study (>.55 in almost all cases). In the current analysis, these coefficients tended to increase when only the latter half of trials was considered, particularly for the Shared-Outcome-only condition (Figure 2.6). When analyzed this way, and considering only the last half of trials of each block, it appears that participants from this condition outperformed those of the other two conditions.

The sensitivity and bias parameters of the three conditions were compared by doing an analysis of the covariance of the log reinforcer ratio for each of the conditions. As could be inferred by visual inspection, no significant differences were found for the Paper/Rock pair ($F[2, 160] = .229$, $p = .796$), the Paper/Scissor pair ($F[2, 160] = 1.67$, $p = .191$), or the Scissors/Rock pair ($F[2, 160] = .920$, $p = .4$), and no significant differences were found across conditions in posthoc tests. The obtained sensitivity parameters across groups are all inferior to 1, which indicates undermatching, although they seem to be consistently, if lightly, higher for participants from the Shared-Outcome-only condition, at least in the latter half of each block (Figure 2.6). The bias parameters, instead, and across all conditions are close to 0, which suggests little systematic bias for one particular alternative across conditions.

The time it took participants to make a selection after the start of a trial (hereafter called "response time") is depicted in Figure 8. The reaction time data was analyzed as a function of block, block type, and finally block portion (the first, second, third, and fourth 16 trials of each block). In all three analyses, participants from the Shared-Outcome-Only condition had slightly but significantly longer response times. The experimental condition affected reaction times, $F(2, 9192) = 15.85$, $p < .01$, and a significant interaction was found between condition and block, $F(6, 9192) = 6.312$, $p < .01$. Post-hoc tests indicated that while the Alone and Shared-Outcome-&-Performance conditions didn't significantly differ ($p =$

.56), the Shared-Outcome-Only differed significantly from both of the others ($p<.01$). While reaction times tended to decrease for the Shared-Outcome-&-Performance condition, and particularly for the Alone condition as the session progressed, they tended to increase for those in the Shared-Outcome-Only. When reaction times were analyzed within block portions (rightmost panel in Figure 8) it is evident that while response times tended to decrease as participants reached the end of a block of 60 trials, $F(2, 9192) = 15.66$, $p < .01$, they actually increased for the Shared-Outcome-Only ($p < .01$).

Analysis of the participants' strategies was conducted by following the procedure developed by Dyson, Wilbiks, Sandhu et al. (2016). According to these authors, a number of strategies can be observed when participants play the game of rock-paper-scissors, broadly encompassed by the win-stay/lose-shift strategy. Moreover, these authors have proposed that the shifting strategy can be divided into two further types, which leads to a categorization in terms of three cyclic strategies: (a) *stay*, consisting in playing the same alternative as in the previous trial, (b) *upgrading*, consisting in selecting the option that would have won in the previous trial, and (c) *downgrading*, consisting in selecting the option that would have lost in the previous trial.

Dyson and colleagues (2016) analyzed the responses of their participants in two ways. First, they compared the use of each strategy (i.e., stay, upgrade, and downgrade) as a function of the choice in the previous trial (i.e., rock, paper, or scissors). Second, they compared the use of each strategy as a function of the *outcome* of the previous trial (i.e., win, lose, or draw). The same analysis was conducted here. Figure 9 depicts the strategy of participants as a function of the choice of the previous trial for all participants. As can be seen, and unlike the participants in Dyson et al. (2016), "staying" was overall the preferred strategy, although important differences seem to emerge across groups.

Participants from the Alone condition tended to have less preference for a single strategy, although they tended to stay after choosing paper. In the Shared-Outcome-only condition, participants tended to stay following every choice. Finally, participants from the Shared-Outcome-&-Performance condition tended to keep playing Rock if that was the previous choice, and to downgrade and stay following paper. Although lacking a statistical analysis, visual inspection suggests that the shared attention conditions tended to produce the more variability in overall strategy.

Figure 2.10 depicts the participants' strategy as a function of the *outcome* (win, draw, or loss) of their last choice for random and non-random blocks. Unlike the results of Dyson and colleagues, and unexpectedly, "stay" remained a common strategy following losses and draws across all conditions. As expected, and just like in the study by Dyson, a win/stay strategy is quite common, but appears to be more marked in the current study. Interestingly, the tendency to stay following a win seemed to have increased in the shared attention conditions in non-random blocks. Losses seemed to affect differentially the strategy of participants compared to draws. In both random and non-random blocks, losses tended to lead to "downgrading" across all conditions. Participants didn't seem to prefer a particular strategy following a draw in the random blocks, but unexpectedly tended to "stay" in non-random blocks only for the shared attention conditions.

Figure 2.11 depicts the changes (or lack thereof) in strategy as a function of the outcome of the previous response across blocks only for non-random blocks. As can be seen, for all conditions, a win/stay strategy becomes more prevalent as the session progresses. It's also noticeable that staying becomes progressively more common following losses and particularly draws in both of the shared attention conditions. While this progressive preference for staying might be attributed to fatigue, it seems unlikely in the face of the fact that it seemed to affect mostly those participants in

the shared attention conditions, and in that it seemed to be more common following draws than losses.

Changes in strategy across block portions are depicted in Figure 12. As can be observed, the win-stay strategy also seems to become more favored as blocks are coming to an end. Otherwise, strategies remain relatively stable, although an increase in staying is also evident following draws only in the Shared-Outcome-only condition. As can be seen, the unexpected preference for staying following draws in both shared attention conditions appears to occur early within blocks.

2.3 Discussion

Undermatching was quite evident across all conditions in this pilot study. Participants failed to "track" the programmed reinforcer rates, regardless of the experimental manipulations, although matching seemed to increase in the Shared-Outcome-Only condition, albeit only when the latter half of trials across blocks are taken into account. In contrast, the shared attention manipulations had important although relatively unsystematic effects in response times and particularly in the strategy used by participants, but these effects were not generally reflected in matching the reinforcer rates. One important consideration derived from the analysis of strategy, and one that might be problematic for a matching analysis, is that draws and losses seemed to affect differentially the strategy of participants, and that this effect appears to be somewhat stronger in shared attention conditions.

One possible reason for these negative results is of a technical nature. The way the experiment was programmed, the probabilities that the computer would select a particular alternative were determined trial by trial. That is, at the start of each trial, and depending on the block, the computer selected an alternative with the programmed probability (e.g., rock: 55%, paper: 30%, scissors: 15%). Since the computer made such determination

with each trial, this means that the actual proportion that an alternative was chosen by the computer across a block differed across participants, and only approached the programmed probabilities. As a consequence the actual differences in probability for each alternative might have been artificially reduced for some participants, which might have blurred the differences in such proportions across blocks.

In addition, and as compared to the previous studies, the current study comprised a significantly lower number of trials (Four blocks of 60 trials each, versus ten blocks of 100 trials in the Kangas et al, 2009 study, and four blocks of 126 trials in the Lie et al., 2013 study). As mentioned above, we decided to reduce the number of trials because session times were exceeding one hour and a half, which we considered excessive. It seems reasonable that the technical issues in determining the probabilities (which might have made the differences across blocks less salient), combined with the reduced number of trials, might have prevented participants from making enough contact with the programmed contingencies. These considerations were addressed in the design of Experiment 2.

3. Experiment 2

Elucidating Shared Attention Effects

The previous experiment involved a shorter version of the task used by Lie et al. (2013) to attempt to determine whether social influences like shared attention play a role in the adjustment to changing contingencies. As discussed, although the social manipulation played some role in the strategy used by participants, regardless of the social manipulation participants in the previous experiment undermatched the programmed contingencies to an even larger extent than participants in the Lie et al. study, with sensitivity parameters ranging between 0.2 and 0.47 (versus 0.54-0.66 in the equivalent condition in the Lie et al. experiment).

One of the issues in the analysis of these results is related to the programmed probabilities across blocks that were selected. As described above, we chose block probabilities that were similar to those used by Lie et al., where one option was only slightly more probable than the others. Moreover, because of the way the computer program determined said probability the actual likelihood of one option being correct might have been lower than anticipated, which might suggest that the differences in the probability of each choice were so small that participants might have been unable to detect them, and hence responded randomly regardless of the changes in probabilities. In other words, since the differences in probabilities were relatively small, choosing at random on each trial was a good enough strategy.

As also pointed out by Lie et al. (2013), the original study by Kangas et al. (2009) involved more extreme ranges of programmed probabilities, including blocks where a choice would be correct up to 85% of the time. With increased experience of these larger differences in probability it might be anticipated that participants would respond more sensitively to the changes in contingencies across blocks. If so, it would be of interest to determine

whether early contact with largely discrepant proportions affects sensitivity more than later contact, and whether the presence of a co-attending confederate would play any positive or negative role in the adjustment to changing contingencies in this task. Accordingly, we modified our experimental task so that instead of having three blocks with the same programmed probabilities, participants experienced blocks with progressively larger discrepancies.

3.1 Method

3.1.1 Design

To further tease out these differences we decided to present these probabilities in two sequences: one sequence where block probabilities become more discrepant across blocks, and another one where probabilities start largely discrepant and become progressively less discriminable. As the discrepancy in probability is increased, participants should be expected to obtain correspondingly higher proportions of correct attempts. In light of the effects of the social conditions in the previous experiment, we decided to explore the effect of these changing probabilities using the same manipulation of shared attention as in the previous experiment. Hence, the second experiment involved a 3 x 2 design, wherein participants from the Solo, Shared-Outcome-only, and Shared-Outcome-&-Feedback conditions were presented with a sequence of either progressively more or progressively less discriminable response probabilities (see Table 3.1).

In the first sequence condition, the participants started with a block where the probabilities of each response being correct were the same, with each successive block having one alternative being progressively more likely. Participants of this "Hard-to-Easy" condition, hence, had a first block with a (0.33 – 0.33 – 0.33) probability, a second block with a (0.4 – 0.3 – 0.3) probability, a third block with a (0.6 – 0.2 – 0.2) probability, and a final block with a (0.8 – 0.1 – 0.1) probability. Participants from the second

condition, were exposed to the opposite sequence ("Easy-to-Hard"), that is, their first block was the one with the most extreme probabilities (0.8 – 0.1 – 0.1) and their final block was the one with the least extreme probabilities (0.33 – 0.33 – 0.33). As in the previous experiment, the computer determined pseudo-randomly which alternative is the most likely for each block, such that each option was the most likely for just one of the blocks.

Also, because of the way the computer choices were preprogrammed in the previous experiment, the actual number of trials of each type only approached the preselected probabilities. In consequence, in this second experiment we decided to control this by making sure that during each block the number of trials of each type will be kept constant across participants. In this manner, for instance, for the block with (0.6 – 0.2 – 0.2) probabilities, 36 out of the 60 trials were of one choice, while the computer chose each of the other alternatives for 12 trials each.

Other than these changes in the sequence and value of the probabilities and in the distribution of trials within each block, the experimental task was identical to that of the first experiment.

3.1.2 Participants

168 undergraduate students participated in the study. They were randomly assigned to one of the three experimental conditions. Data from two of the participants were excluded from the analysis because of errors in the experimental session. The final data analysis included 166 participants, with 58 in the Solo condition (female N = 29, mean age = 18.74 years, sd =2.59), 59 in the Shared-Outcome-Only condition (female N = 32, mean age = 18.78 years, sd =0.97), and 49 in the Shared-Outcome-&-Performance condition (female N = 25, mean age = 18.84 years, sd =0.87). 84 participants were in the "Easy-to-Hard" condition, while 82 were in the "Hard-to-Easy" condition. Participants received academic credit for participating in the experiment.

3.1.3 Apparatus and Setting

The experimental sessions were conducted in the same setting and with the same apparatus as in Experiment 1. Sessions lasted about 45 minutes.

3.1.4 Procedure

The procedure was identical to that of Experiment 1. Participants of Groups 2 (Shared-Outcome-only) and 3 (Shared-Outcome-&-Feedback) sat next to an undergraduate confederate who acted as an unknowing participant. All participants were presented with the same creativity task as in the previous experiment, with participants from Groups 2 and 3 completing it by collaborating with the confederates. Then, the participant (and confederates) was led to the cubicle, with confederates standing behind the participant. The same general instructions as in the previous experiment were read to the participant and presented in the screen, and following the five practice trials they were told that they could start whenever they wanted, while confederates sat down in an adjacent cubicle.

The trial configuration was identical to that of the previous experiment, and like that one, the type of auditory feedback that participants received after making a choice depended on the experimental condition: participants from Group 2 (Shared-Outcome-only) received auditory feedback for the outcome of their choice (applauses, "tie", or a "thunk" tone corresponding to wins, draws, and losses, respectively), while participants from Groups 1 (Solo) and 3 (Shared-Performance and Feedback) received auditory feedback for both their and the computer's choice (the words "Rock," "Paper," or "Scissors" according to their choice), as well as about the outcome of their choice (applauses, "tie", or a "thunk" tone). Thus, as in the previous experiment, while confederates in Group 2 would only hear the tones of applause, tie, or the thud tone, those in Group 3 heard also what the player's and computer's choices were.

Like the previous experiment the session consisted of four blocks of 60 trials each, for a total of 240 trials. The intertrial and interblock

intervals were the same as in the previous experiments, as well as the messages participants received between blocks and at the end of the session, and were asked to complete the brief survey about demographic information and whether they used any particular strategy or game plan, and then were debriefed on the nature of the study and were allowed to leave.

Participants also completed a brief self-report survey (see Appendix) that inquired about their own perception of different aspects of the task and their performance, and particularly about their interaction with the confederate in the case of the shared attention conditions. The survey consisted of 11 Likert-scale items that could be rated from 1 (Not all) to 7 (Very much). To rate this survey, items 4 through 6 were averaged under the label of "Partner Liking," while items 7 through 11 were averaged under the label of "Perceived closeness to partner."

3.2 Results

Insofar as the main dependent variables for the collected data on both the performance and strategy of participants were the proportions of outcomes (i.e., wins, draws, and losses) or the proportion of trials using a particular strategy (i.e., staying, upgrading, or downgrading), and these are categorical variables, and violate a number of the assumptions of the general linear model (GLM), the statistical analysis of these data demands a different approach. Moreover, since the proportions of each of these variables across blocks are likely correlated (as they correspond to responses by the same individual in time), the most appropriate model for analysis is the Generalized Estimating Equations or GEE model with an autoregressive working correlation structure, which allows for weaker correlations among data points farther away in time. All of the subsequent statistical analyses were obtained by using this model, and the significance of differences tested by conducting Wald Chi-squared tests.

Figure 3.1 shows the overall mean proportion of wins, draws, and losses for participants of each social and sequence condition. As can be observed, in general participants didn't differ significantly in the proportion of wins as a function of social condition, Wald $\chi^2(2, N = 39{,}840) = 1.129$, p=.569, with the overall proportion ranging between 0.38 and 0.40 across all conditions. The sequence of probabilities, instead, did affect the mean proportion of wins, with participants in the "Easy-to-Hard" sequence obtaining slightly more wins (39.6%) across all social conditions compared to those in the "Hard-to-Easy" condition (37.8%), Wald χ^2 (1, N = 39,840) = 6.963, p<.05. No interaction was found between the social and sequence conditions, χ^2 (2, N = 39,840) = 0.981, p= .612.

The two sequences affected differentially the proportion of wins across blocks for all social conditions, χ^2 (3, N = 39,840) = 381.1, p< 0.001. In particular, for participants in the "Hard-to-Easy" condition, the proportion of wins didn't increase progressively across blocks (Figure 3.2) as was expected, but instead remained relatively stable for the three first blocks and then increased abruptly in the final block. It should also be noticed that out of the three conditions, participants in the Shared-Outcome-only one showed the more progressive increase of wins. In contrast, in the "Easy-to-Hard" conditions the proportion of wins tended to decrease progressively as blocks became harder to discriminate, as would be expected of behavior that is sensitive to the changing contingencies. Also, participants in the Solo condition obtained the highest proportion of wins, but only in the Easy-to-Hard condition.

Within blocks the proportion of wins changed across block portions, with more wins overall in the last half of the trials (see Figure 3.3), Wald χ^2 (3, N = 39,840) = 419.997, p<.001. No significant interaction was found between block portion and either social condition, χ^2 (6, N = 39,840) = 4.157, p=.655, or sequence condition, χ^2 (3, N = 39,840) = 3.873, p=.276.

Figures 3.10 through 3.11 depict the log response ratios as a function of the log reinforcement ratios for pairs of choices (i.e., rock vs. paper, rock vs. scissors, and paper vs. scissors) for all participants for the three social conditions. As was the case in the previous experiment, across all social and sequence conditions the participants' pattern of choice was one of undermatching (sensitivity, $s < 1$). Bias didn't occur systematically across conditions (see Table 3.2). The sensitivity parameters tended to be larger for participants in the "Easy to Hard" sequence conditions, with the exception of participants in the Shared Outcome & Feedback condition, whose parameters were very similar for both sequence conditions. This increased sensitivity in "Easy to Hard" sequences appeared to be enhanced in the Solo condition. Since the GLM delivered such low accounted variance estimates (ranging between 1.8% and 42.3%), however, these differences across conditions need to be taken with care, particularly when considered together with the lack of overall differences in the proportion of correct attempts across conditions that we reviewed above.

Since the probabilities of the program choosing each option were systematically manipulated across blocks it could be argued that certain strategies would be most optimal depending on the particular block probabilities. Namely, it could be argued that in those blocks in which the computer chose each option with equal probability the optimal strategy for participants would be to switch their choice often. As the differences in probability become larger, the optimal strategy would be sticking with one option progressively more. Accordingly, "Staying" was the most common strategy across all social and sequence conditions (Figure 3.4). Although participants from the Solo condition tended to stay more overall (41.5%) than participants from the shared attention conditions (38.0% and 37.8%, for conditions 2 and 3, respectively), this effect was not significant, Wald χ^2 $(3, N = 39{,}176) = 2.548$, $p=.28$. In contrast, and as occurred in the case of

the proportion of total wins, the sequence of presentation did have a significant effect, (1, N = 39,176) = 12.224, p<.001, with participants in the "Easy-to-Hard" conditions staying for 43% of the trials versus 35% in the "Hard-to-Easy" conditions. No interaction was found between sequence and social condition.

The proportion of Staying as the preferred strategy varied as a function of the block, as expected, Wald χ^2 (3, N = 39,176) = 47.337, p<.001. Furthermore, this effect was moderated by the sequence of presentation, Wald χ^2 (3, N = 39,176) = 112.736, p<.001 (see Figure 3.5). Across social conditions, participants in the "Hard-to-Easy" conditions didn't vary their relatively inefficient strategy across the first three blocks and only switched to Staying abruptly by the fourth block. In contrast, participants in the "Easy-to-Hard" conditions tended to opt for Staying gradually less and less, as would be expected of a better contact with the programmed probabilities. As could be inferred from Figure 3.5, the presence or absence of a coattending confederate didn't appear to play a role, Wald χ^2 (2, N = 39,176) = 2.514, p=.285. No interaction between social and sequence conditions was found, although participants in the Solo condition tended to opt for Staying (56% of trials) more than participants in the Shared Attention conditions (46% and 42%, respectively) in the Easy-to-Hard sequence, and tended to do so across blocks. Overall, and regardless of social and sequence conditions, Staying tended to become progressively more frequent throughout block portions, Wald χ^2 (3, N = 39,176) = 48.043, p=<.001 (see Figure 3.6).

The participants' strategy was affected by the outcome of the previous trial, χ^2 (2, N = 39,176) = 162.279, p=<.001 (Figure 3.7). Overall, participants tended to stay more following wins (45.6% of trials), and to switch (upgrading or downgrading) following draws and losses. Furthermore, across social and sequence conditions, stays were slightly more frequent

following draws than when they followed losses. Overall, the strategy of participants as a function of the outcome of the previous trial was the same regardless of the social manipulation of shared attention (see Figure 3.8). In contrast, an effect of sequence was found, with stays after wins being much more frequent in the Easy-to-Hard conditions (about 62% of trials) than in Hard-to-Easy conditions (about 42% of trials, see Figure 3.8).

These results are much more evident when analyzed through blocks (Figure 3.9). Since the optimal strategy for those blocks with the larger discrepancies among probabilities (i.e., the 0.8 – 0.1 – 0.1 blocks) is staying, it's not surprising that participants chose staying preferentially in those blocks following wins. In the case of participants in the Easy-to-Hard conditions, who experienced first this type of blocks, staying remained the most common strategy across blocks following wins and decreased only slightly even in the more "ambiguous" blocks. In contrast, when participants first experienced the progressively easier conditions the rate of wins following stays remained low and increased only in the final, easiest, block. This was true regardless of the social condition. Following draws and loses, across blocks participants preferred switching (i.e., upgrading or downgrading) than staying. Interestingly, the rate of staying was slightly but reliably higher following draws than losses, which suggests that, indeed, they affect participants differentially, regardless of the social condition.

3.2.1 Self-Report Data

As can be observed (Table 3.3), overall and across conditions participants described the task as not difficult (ranging between 1.73 and 1.86), and relatively enjoyable (ranging between 3.24 and 4.05), and participants were relatively satisfied with their performance. Interestingly, this self-reported satisfaction with performance correlated with the total of wins only for the Solo conditions, $r(46)=.301$, $p<.05$, two-tailed, while in the Shared Attention conditions the satisfaction with their performance

correlated instead with the total losses, $r(56) = -.498$, $p < .001$, for the Shared Outcome Only condition, and $r(46) = -.341$, $p < .05$, for the Shared Outcome & Feedback condition. In regard to closeness to the liking and feelings of closeness with partner, they didn't appear to correlate significantly with either the strategy or overall outcomes of participants, although a small negative correlation appeared between closeness and total wins, only for the Shared Outcome Only condition, $r(56) = -.273$, $p < .05$.

3.3 Discussion

Compared to the previous experiment, participants obtained only slightly more correct attempts and then mostly on those blocks where the programmed probabilities of the computer's choices were the most discrepant. Overall, the social manipulation of shared attention by a coattending confederate didn't have any significant effect in the outcome, although participants in the Solo condition did slightly better in the most discrepant block when it was presented first. Participants across all social and attention conditions undermatched the programmed reinforcer probabilities, although undermatching was slightly less frequent in the "Easy to hard" sequence conditions.

The social manipulation didn't appear to systematically affect the participant's strategy as a function of either the previous choice or its outcome, although again, those in the Solo condition tended to opt for staying as a preferred strategy, only when the sequence of block presentation was Easy-to-Hard. This experiment is unique in that the different probabilities were deliberately presented in progressively decreasing or increasing order. Although this introduces the complicated issue of order effects, we did find that the sequence of presentation was a major predictor of both performance and, particularly, strategy in this brief three-choice task. Experiencing "hard" blocks, that is, those where each option was as probable as the others had deleterious effects on the strategy of

participants in the subsequent blocks. The fact that these participants only changed their strategy in the final, most discriminable block suggests that their responding was less sensitive than that of participants who experienced the easiest block first and whose strategy changed gradually in correspondence with the decreased probability of one response being the most frequently correct one for a particular block.

Rodriguez and Logue (1986) reported relatively similar effects of order of presentation in a study with pigeons, showing bias towards alternatives that were initially preferred as they become progressively less beneficial (i.e., as their magnitude of reinforcement decreases and their delays to reinforcement increase), but to the best of our knowledge, no published study has presented the varying probabilities in this particular ascending or descending order. This is probably a result of the differences in the experimental study of choice when conducted in humans versus other species, where subjects have daily sessions extending for weeks. Human participants, instead, and for a variety of reasons, generally participate in very few or most commonly a single experimental session. In the traditional nonhuman animal preparation for choice behavior, the varying probabilities or rates of reinforcement associated with different choices are presented only once per experimental session. It might well be, of course, that with extended experimental sessions the behavior of human participants might resemble more that of other species (cf. Madden, Chase & Joyce, 1998), but insofar as ethical, practical or monetary issues make this possibility largely unrealistic this "insensitivity" effect occurring as a result of within-session presentation order needs to be considered in studies of choice behavior involving human participants.

There is at least another possibility. As will be discussed in more detail in the next chapter, appropriate levels of matching are obtained occasionally with human participants in concurrent choice procedures

relatively similar to this one (e.g., Kangas et al., 2003), but in many other circumstances deviations from optimal matching are frequent, unless precise verbal descriptions of the contingencies are provided to participants (e.g, Takahashi & Shimakura, 1998). Due to their presumably large history playing the Rock-Paper-Scissors game, participants are likely to infer that their opponent (the computer, in this case) is adapting its strategy throughout the session in response to the participant's strategy. In other words, participants might be inferring more complicated patterns of strategy than the computer is actually using (that is, choosing one option preferentially), and are hence being "too smart for their own good," as Arkes and Ayton (1999) suggest happens in other research on suboptimal choice (e.g., the sunk-cost effect). The participant's inferences about the contingencies would then work as self-provided instructions for their further responding, and might hinder control by experimental contingencies. If so, this might explain the discrepancies in strategy following wins when the more discriminable sequence is presented first versus when it's presented last. Participants might assume the computer is actually adapting to their own choices, particularly in the truly random blocks, and hence assume that no strategy will be optimal. When participants experience the easier block first, instead, they appear to become slightly more sensitive to the computer's strategy.

With the technical and procedural modifications made in this experiment as compared to the first one, the effects of the presence of a coattending confederate on participants' performance and strategy became even less obvious. If anything, shared attention seemed to only affect (and negatively) performance in the easiest, most discriminable of blocks, with participants in the Easy-to-Hard condition who did the task alone obtaining more correct attempts in the first block than those participants in the corresponding shared attention conditions. This resulted in a slightly larger, albeit non-significant, proportion of wins overall for this condition.

4. Experiment 3

Exploring the Effects of Verbal Instructions

As reviewed in the first chapter, in the tradition of the experimental analysis of human behavior, the study of rigidity in human responding to programmed experimental contingencies involving schedules of reinforcement has been centered on the role of verbal experimental instructions provided to participants. As mentioned there, verbal instructions can be seen as somewhat of a double-edge sword. On the one hand, they appear to bypass the need for extensive training in the establishing of behavioral repertoires, an aspect that has practical implications for example in the reduction of the number and duration of experimental sessions, or in the duration and efficacy of applied interventions. Verbal instructions, of course, are a very important aspect of teaching, a cooperative communication strategy that appears to be particularly characteristic of human cultures, and has been suggested as a critical factor in the so-called "ratchet effect," a term used to refer to the cumulative character of human culture (Tomasello, 2011).

By the other hand, and as reviewed in the first chapter, verbal instructions might on occasion hinder effective control by experimental contingencies, a phenomenon that is often called "insensitivity to contingencies." To the best of our knowledge, only a handful of studies involving instructional manipulations with human participants has been conducted using a choice paradigm. In one of these paradigms participants have to choose between two schedules of reinforcement with one of the schedules yielding progressively lower rates of reinforcement (a progressive-time or PT schedule of reinforcement), and participants are given relatively accurate or inaccurate description of these changes (Fox & Pietras, 2013; Fox & Kyonka, 2017; Hackenberg & Joker, 1994), all of them showing insensitivity as a function of the presentation of accurate instructions first. A second procedure involved a self-control task where participants could choose

between an option that provided reinforcers immediately but kept increasing the delays of future choices or one that provided reinforcers about a specific delay (Kudadjie-Gyamfi & Rachlin, 2002). Participants in this study were given an accurate instruction, which increased their choice of the "self-control" key.

Finally, in a more traditional concurrent choice paradigm, a few studies involved some direct manipulation of verbal instructions, consisting generally of a brief, vague description of the contingencies, (e.g, Horne & Lowe, 1993). Takahashi and Shimakura (1998) had participants press two levers according to two different concurrent variable interval (VI) schedules, and provided them with instructions that described either the time contingency or the reinforcement rate associated with each lever for each VI value, or both. They found that the matching law described very accurately the pattern of responding of participants, but only when they received instructions describing the frequency of behavior.

Kangas, Berry, Cassidy and colleagues (2009), using the rock-paper-scissors procedure herein used, conducted an experiment where participants were given accurate verbal instructions at the start of each of their ten 100-trials blocks. These accurate instructions resulted in overmatching, that is, "excessive" preference for the choice described as being the most probable. Considering the failure of our shared attention conditions in affecting systematically the participants' tracking of the programmed probabilities across blocks, it seems worth exploring whether providing instructions to participants might moderate the effect of the presence of a co-attending confederate on the performance and strategy of participants.

Moreover, to the best of our knowledge, no studies of choice in humans have involved the use of *inaccurate* instructions. In all reviewed literature in this area (cf. Takahashi & Shikamura, 1998), instructions have been at worst vague. As a first exploration of this issue in the context of our

paradigm, we decided to expose participants to a modified version of our original task, involving the presentation of accurate and inaccurate instructions, when completing the task by themselves or in the presence of a second, coattending participant.

4.1 Method

4.1.1. Design

Major modifications on our original paradigm were done in this experiment. First, since our interest in this experiment was mostly on the role of verbal instructions, we decided to have only two social conditions: a Solo condition and a Shared Instruction condition. Moreover, unlike our two previous experiments, this one didn't involve a confederate, and instead participants attempting the task either by themselves or with another participant.

More importantly, participants in the Shared Instruction condition were informed, accurately, that the computer program would be the same for both of them, and that in consequence it would be playing the same strategy against both of them. Because adding auditory performance feedback might instead prove a distracting factor with participants playing simultaneously, and to better isolate shared attention on instructions only, unlike the two previous experiments participants didn't receive any auditory feedback on their performance.

Two further major modifications on the original task were conducted. First, we wanted to increase the number of trials of each block to make them more similar to those used in the Kangas et al. (2009) and Lie et al. (2013) studies. Also, and in line with the second experiment of Kangas et al. (2009), we provided an instruction at the start of each of the blocks. As a consequence, and because we wanted to explore the utility of this paradigm to study instructional control, we decided to reduce the number of blocks. Participants were hence presented with three blocks. First, a control block

which had only 40 trials, and in which the computer played each choice with the same probability (i.e., 0.33, 033, and 0.33). Second, an Accurate Instructions block where participants received an accurate instruction at the start of a block of 100 trials. Finally, an Inaccurate Instructions block, with an inaccurate instruction at the start, consisting of 100 trials.

4.1.2 Participants

110 undergraduate students (46 females; mean age = 19.32 years, sd = 1.209) participated in this experiment, randomly assigned to each of the two experimental conditions. 43 participants were in the Solo condition, while 67 were in the Shared Attention condition. Participants received academic credit for participating in the experiment.

4.1.3 Apparatus and Setting

The experimental sessions were conducted in the same setting and with the same apparatus as in the two previous experiments. Sessions lasted about 45 minutes.

4.1.4 Procedure

Like in the two previous experiments, participants from both conditions were asked to complete a creativity task at the start of the experimental session, with participants of the Shared Instruction condition being urged to complete the task collaboratively. Following the completion of this task, participants were seated in an experimental cubicle, with those in the Shared Instruction condition sitting in adjacent cubicles. The cubicles prevented visual contact with the other participant's performance. All participants were asked to enter their subject number and click on a button reading "Condition 1." The goal of this was that participants in the Shared Instruction condition were aware that they were playing the same task as the other participant next to them.

General instructions were then read to all participants. These general instructions were almost identical to those used in the previous experiments:

"You will be playing a game of 'Rock-Paper-Scissors' against the computer. The rules are the usual ones: Rock beats Scissors, but loses against Paper; Paper beats Rock, but loses against Scissors; and Scissors beats Paper, but loses against Rock. If you don't know how to play this game, please ask the experimenter to explain. You will make your selection by clicking the mouse cursor on the option you choose. The computer will make its selection at the exact same time you do. The winner of each round will receive one point, and the loser will receive 0 points. If there is a tie, the score will remain unchanged. Your task is to obtain as many points as you can, so try and respond in the way that would make you get the most points."

Participants in the Solo condition were also presented with the following sentences to encourage them to perform better and to try to figure out the computer's strategy:

"To encourage you to do your best, we will donate some money to a charity depending on who obtains the most points. We'll provide you with a list of charities at the end of experiment, so, the higher your score, the higher the chances we'll donate to that particular charity. Also, notice that the computer tends to stick to a particular strategy across the game, generally choosing one option slightly more often than the others. We encourage you to try to figure out what that strategy is. Do you have any questions?"

For participants in the Shared Attention condition the instructions emphasized shared attention and were aimed to induce a more collaborative approach:

"To encourage you to do your best, we will donate some money to a charity depending on how well both of you perform. At the end of the experiment we will combine both of your scores. We'll provide both of you with a list of charities at the end of experiment so that the higher your combined score is, the higher the chances we'll donate to that particular

charity. Also, notice that the computer tends to stick to a particular strategy across the game, generally choosing one option slightly more often than the others. The computer will at all times play with exactly the same strategy against both of you. We encourage you to try to figure out what that strategy is. Do you have any questions?”

The trial configuration was identical to that of the previous experiments with the difference that no auditory feedback was provided for either choice or outcome for any of the participants.

The session consisted of three blocks, the first of 40 trials, and the remaining two of 100 trials each, for a total of 240 trials. Following completion of the first trial participants were presented with the following message in the screen:

"From now on, the computer will be programmed to play Rock ___% of the time, Paper ___% of the time, and Scissors ___% of the time. Click the 'Continue' button when you're ready to continue."

These were the same instructions provided in the experiment by Kangas and colleagues (2009). In the case of the instructions for the second block, the actual, corresponding probabilities of the computer choosing each option were provided. For the third block, the same message appeared but the depicted proportions described inaccurately the programmed contingencies.

To better tease out the effects of these instructions, we opted for intermediate probabilities. For both of these blocks, then, the computer chose one option slightly more often than the other two (.55, .15, .30, respectively). The computer determined randomly which option would be the most frequent, and it was always a different one of the second and third blocks.

After the last trial was completed, participants received a message informing them of the end of the session, and were asked to complete the brief survey about demographic information and whether they used any

particular strategy or game plan. In addition, and so as to determine whether any personality factors might be predictive of rigidity in the participants' strategy, they completed a brief personality inventory, the *Ten Item Personality Inventory* (TIPI; Gosling, Rentfrow & Swann, 2003). They also completed a form indicating to which charity they would prefer money was donated. They were then were debriefed on the nature of the study and were allowed to leave.

4.2 Results

Since this Experiment included only three blocks, and one of them with significantly less trials (40) than the other two (100 each), a matching analysis is precluded in the present experiment. Moreover, since the first block was briefer, data from the last two blocks was emphasized in the analysis.

The overall mean proportion of wins, draws, and losses for participants of both conditions is presented in Figure 4.1, which are slightly higher than those for participants in the previous experiment. Although participants in the Solo condition obtained slightly more wins, this difference was not significant, $\chi^2(1, N = 26{,}400) = .956$, p=.328. Figure 4.2 shows the proportions of each outcome for both conditions across blocks. As in the previous experiment, a GEE model with an autoregressive working correlation structure was used to analyze these differences in the performance of participants across blocks. The proportion of wins changed as a function of the instructional manipulations, Wald $\chi^2(2, N = 26{,}400) = 129.835$, p<.001, but the effects of introducing instructions didn't change as a function of the social condition, $\chi^2(2, N = 26{,}400) = 1.000$, p=.607.

As can be observed, for both conditions the introduction of accurate instructions increased the proportion of wins (see Figure 4.2). It must be remembered, though, that the first block consisted of merely 40 trials as compared to 100 trials for the other two blocks. Nonetheless, the proportion

of wins for both conditions was above 40%, which is slightly higher than those in the intermediate blocks of the previous experiment. A GEE analysis excluding the trials in the first block indicates, importantly, that performance deteriorated slightly but significantly with the introduction of inaccurate instructions (that is, in the third block), Wald $\chi^2(1, N = 22{,}000) = 22.062$, $p<.001$. Once again, no main effect or interaction was obtained as a function of the social condition.

The effect of the introduction of verbal descriptions of the experimental probabilities is clearer when the strategy of participants was observed. As can be seen, and unlike the previous experiments, the presentation of verbal instructions induced a marked preference for Staying versus the other two strategies, $\chi^2(2, N = 26{,}070) = 224.234$, $p<.001$. Moreover, when considering only the instruction blocks, participants tended to Stay slightly more in the block with inaccurate instructions, Wald $\chi^2(1, N = 21{,}780) = 3.967$, $p=0.046$. Once again, no significant main effect or interaction of the social manipulation was obtained.

As can be observed in the accompanying Figure 4.5, and as expected, following wins participants tended overall to stay. Moreover, the proportion of trials for which staying followed wins was comparable to that of the blocks with the most extreme differences in reinforcement rate in the previous experiment. Interestingly, this proportion was also larger for the block with inaccurate instructions, $\chi^2(1, N = 9{,}201) = 6.500$, $p<0.05$ (Figure 4.6). The data were further analyzed in successive block portions of ten trials. When analyzed this way, as can be seen, while both conditions had similar rates of correct attempts, the social attention condition tended to produce less variable rates within block (Figure 7). Moreover, as shown in Figure 4.8, the difference lied in that staying became progressively more common in the Accurate instructions block, while it started at a high level in the Inaccurate instructions block regardless of the social condition

manipulation. The difference in these sequences was significant, $\chi^2(18, N = 9{,}201) = 100.786$, $p<.001$.

Also, and unlike the previous experiment, staying remained the preferred strategy following both draws and losses for participants in both conditions (Figure 4.9). As can be seen, the effect was more pronounced with draws, suggesting again that they might affect participants differently than losses.

4.2.1 Self-Report Data

Participants rated the difficulty of the task as low ($M = 1.96$), with moderate enjoyment and moderate satisfaction with their performance. Those in the shared attention condition also reported moderate liking of the other participant and relatively low levels of closeness to them. The reported difficulty of the task didn't correlate significantly with the outcomes of performance (see Table 4.1), while some moderate correlations were found between difficulty and the participant's strategy as a function of the previous choice: the more participants "stayed" the lower the difficulty rating, $r = -2.6$, $p<.001$, two-tailed. Unlike the previous experiment, enjoyment of the task didn't correlate significantly with the outcomes of the task, which can be attributed to the fact that no auditory feedback was provided on performance in this experiment. A couple of small correlations were found between these self-reported items and the scores for the brief personality inventory, with enjoyment of the task and satisfaction with the performance correlating positively with "emotional stability," and enjoyment correlating negatively with openness to experience. Although Schultz and Searleman (2002) had suggested that Openness to Experience should be most related to behavioral rigidity, no significant correlations were found between scores in any of the personality items and the participants' performance or strategy.

4.3 Discussion

The introduction of verbal descriptions of the programmed contingencies appeared to produce a number of effects on the participants' performance and strategy in blocks where the computer chose a particular response slightly more than the others. This experiment included only two of these blocks, which in addition had a larger number of trials than those in the previous experiments. It also included a different shared attention condition: participants who completed the task simultaneously to a second participant were informed that the computer would be playing the same task and the computer would choose the same strategy against both of them. As described in a previous section, they were also told that their scores would be combined, so as to encourage engagement with the task. Overall, however, and just like in the previous experiments, the manipulation of shared attention didn't produce noticeable effects on either performance or strategy.

Verbal instructions generally had a beneficial effect on the participants' performance, with participants obtaining larger proportions of correct attempts than participants in blocks with comparable proportions in the previous two experiments. These results were probably a consequence of participants in this experiment opting for "staying" as the preferred strategy, even in the face of negative results (i.e., draws and losses), unlike those in the other experiments. The fact that participants did slightly worse in the third block in which *inaccurate* instructions were presented indicates that although the verbal instructions provided at the start of the block were merely brief descriptions of the probabilities they were at least to some extent followed in spite of their inaccuracy, as shown in that participants kept choosing to stay in spite of being unsuccessful. As has been suggested in the literature, one important factor in the likelihood of instructions promoting "insensitivity" is the extent at which they are incompatible or discrepant with the actual contingencies. Although our

instructions were of course inaccurate, participants could still go by them and win every now and then, a suboptimal strategy, certainly, but not completely bad. Following instructions is, to put it in lay terms, a "good enough" strategy, particularly as the negative consequences of draws and loses in these experiments are not particularly salient.

Interestingly, and related to this last point, our results replicate some of the findings regarding strategy in the rock-paper-scissors study by Dyson and colleagues (2016): our participants, like theirs, tended to stay more following draws than losses and, when switching, they tended to favor upgrading after draws and to favor downgrading following losses. It is not possible to ascertain the reason for such pattern of preferences with the present data, although an important possibility that needs to be explored more precisely might be related to the phenomenon of loss aversion. In support of this possibility, one of the factors that appear to enhance more reliably the sensitivity to contingencies in experiments studying instructional control under schedules of reinforcement is the introduction of punishment contingencies (like "response cost") for the failure to respond according to contingencies (e.g., losing monetary rewards; Fox & Kyonka, 2017; Fox & Pietras, 2013; Galizio, 1979; Hackenberg & Joker, 1994).

5. General Discussion

David Premack is an experimental psychologist that is better known for coining the interesting –and very problematic– notion of "theory of mind." Before that, however, he was better known by a less controversial and arguably more fruitful development, which carries his name: the Premack Principle (Premack, 1965). Although the basic idea behind this principle, the notion that high-probability activities can be used to reinforce low-probability ones, is very well known, the conceptual and potentially revolutionary implications of his ideas are unfortunately less understood (Killeen, 2014).

The fundamental point in Premack's ideas is that "reinforcers" are not things, or more precisely that "things" don't reinforce behavior. Instead, he proposed, activities reinforce other activities: Food is not necessarily the reinforcer, but eating; salaries are not reinforcers, but what you do with them is reinforcing. This is more than merely a linguistic turn: By emphasizing reinforcement in terms of activities not only is the traditionally hedonistic notion of reinforcement as pleasure discredited, but a more complex and dynamic notion of reinforcement is advanced. In particular, further refinements on Premack's ideas determined that for an activity to function as a reinforcer all that is required is that it is prevented (for whatever reason) from occurring at the frequency at which it normally does.

For instance, a rat in a traditional learning experiment isn't allowed to eat as much or as often as it does. Since eating is occurring below its usual frequency, it can now be used as a powerful consequence for another activity, say, pressing a lever. As can be inferred, this approach suggests a sort of behavioral homeostasis in that organisms would work to be able to reach a previous pattern of activity. Since for the rat this means, however, increased effort, the rat will settle for an *optimal* pattern that will allow

it to approach as much as possible the desired level of activity with as little effort as necessary. It is this emphasis in *optimality* that made possible an alliance of sorts with evolutionary biologists interested in adaptation (e.g., optimal foraging theory) and economists interested in, for instance, maximization of utility, in what is now called "behavioral economics."

Our understanding of optimal allocation of behavior is, to date, largely incomplete (see Grace & Hucks, 2013, and Staddon, 2001, for reviews) but one of the most surprising conclusions of this literature is that under some conditions humans behave suboptimally in tasks in which individuals of other species behave optimally. Herbranson and Schroeder (2010), for example, found that while pigeons used an optimal strategy in a version of the Monty Hall dilemma, undergraduate students failed to adopt such an optimal strategy. Another instance of suboptimal human choice versus optimal choice in other species is the famous sunk-cost effect (e.g., Arkes & Ayton, 1999; see also Fantino, 2012), that refers to tendency to persist in a decision after one has invested money or effort on it.

The approach to the phenomenon of behavioral rigidity/flexibility that has been advanced in the present document has emphasized, likewise, optimality. Participants in the three experiments played a computerized version of the old Rock, Paper, Scissors game in which the probabilities that the computer would play one option were systematically manipulated across blocks. The Rock, Paper, Scissors game resembles to a large extent a set of problems in probability theory that are often referred to as "bandit games" (see e.g., Berry, & Fristedt, 1985). The "bandits" in the name refer to the situation in which a gambler has a large number of slot machines ("one-armed bandits") to play, each with a different probability. The gambler has a limited amount of money or time, so a dilemma is created regarding which of the machines should be played and for how long. The key issue here is that

there is a tradeoff for the gambler: should he stick to a relatively good machine ("exploitation") or should he attempt others that might be more profitable ("exploration")?

One key distinction between the Rock-Paper-Scissors (RPS) game and the basic bandit game is that unlike the latter, the opponent in an RPS game also adapts his strategy with each repetition of the game (a so-called "adversarial bandit"). Because of this, as Dyson et al. (2016) point out, Rock-Paper-Scissors is interesting insofar as no particular strategy will ensure success. According to the analysis of Dyson and colleagues, under natural circumstances where players are attempting to beat each other it should be expected that players would choose each of the options "with a random distribution but equal probability (33.33%)" (p. 1).

By manipulating the probabilities across blocks, as originally proposed by Kangas et al. (2009), the game becomes "beatable," and as their results show, it would be expected for players to adapt to these changed probabilities by increasing the frequency of the appropriate counter-choice. The fact that in general our participants, as well as those of the experiment by Lie et al. (2013) undermatched the programmed probabilities suggests that this paradigm is overall useful in determining the occurrence of suboptimal choice. Undermatching isn't uncommon in studies of choice with human participants, even in experiments that attempt to resemble more those conducted in traditional operant boxes with animals of other species by, for example, using a lever as an apparatus, or particularly by increasing the number and duration of experimental sessions (see Pierce & Epling, 1983) for a review.

The present task differs in two important respects from most traditional literature on the allocation of behavior under concurrent schedules of reinforcement. First, instead of a two-choice alternative, this task involves allocation among three choices. In their literature review

Kangas et al. (2009) noticed that at least with nonhumans the generalized matching equation appeared to describe allocation to more than just two alternatives just as well, and the fact that their participants did demonstrate matching seems to support this result. The second critical difference is that instead of the typical concurrent variable interval (VI) schedules of reinforcement often used in matching research, this task involves instead probability matching. The main difference here lies in that in the typical procedure, while the organism is responding to one choice, the other alternative remains available. In the present design, instead, choosing one option eliminates (for that trial) the other two alternatives. As pointed out by Kangas et al., probabilistic schedules tend to produce exclusive allocation to the most probable alternative, and might be better described by a model different to that of matching.

That being said, it's important to remember that the original interest in this task isn't so much in how well does matching describe the data, but in whether this is a useful paradigm to understand rigid decision making. One critical factor that appeared to be related to the differences between the results of our experiments and those by Kangas et al. (2003), and to some extent those by Lie et al. (2013) was the length of exposure to the experimental contingencies, both of which involved a much larger amount of trials per block. In the traditional paradigm used with other species, animals are often exposed to the experimental contingencies in daily sessions that can extend for weeks. In humans, even in those cases where participants were exposed to daily sessions this didn't mean warranty of matching (e.g., Horne & Lowe, 1993; Takahashi & Shikamura, 1998). Furthermore, since the matching model emphasizes behavior that has reached a steady state, it has been suggested that it might be an inappropriate model to understand dynamical changes in transitional states (e.g., Grace & Hucks, 2013), which appears to be particularly relevant to the problem of rigidity herein

explored. It is for this reason precisely that a consensus is being created that, at least as regards choice in transition, matching will be soon discarded as model in favor of more complete ones (see McDowell, 2013 for a review of some of the alternatives). In summary, although our results suggest some limitations of the Rock, Paper, Scissors game as a paradigm for the study of rigid patterns of behavior, the fact that it has allowed as to identify order effects and the fact that it allowed an important measure of instructional control suggest that the task has potential to further study dynamic effects in the allocation of behavior.

An unanticipated result across our three experiments was the absence of a systematic effect of our social manipulation, at least in regards to performance. Shteynberg (2015) reviewed a variety of studies using a variety of manipulations demonstrating reliable and systematic effects of believing that someone was coattending to the same set of circumstances. This shared attention meant generally a facilitating or strengthening effect, indicative of greater devotion of cognitive resources (e.g, attention, memory) to the targets of such shared attention. In the present experiments, in contrast, a coattending presence didn't appear to affect positively or negatively the participants' performances in our task. Shared attention appeared to decrease the frequency of the "lose-shift" strategy following draws versus losses in our first experiment. In the second experiment, which involved the sequential manipulation of the discrepancies among choice probabilities, shared attention appeared to actually work against optimality when in an increasing difficulty sequence by slightly reducing the prevalence of staying as a strategy. Shared attention didn't modulate either the effect of accurate or inaccurate instructions in Experiment 3, even if unlike the previous experiments participants were explicitly encouraged to work cooperatively, which putatively should have increased the devotion of cognitive resources to the detection of discrepancies in the contingencies. If anything,

participants in the Shared Attention condition appeared to adopt a strategy (e.g., staying) more coherently throughout blocks, which might point out to a certain rigidization of behavior, even though it didn't actually affect performance.

As is the case with all negative results, the exact reason for which no reliable, systematic effect of shared attention was found is obscure and will require further testing. There are of course many possible explanations. For one, the use of a creativity task to enhance feelings of "sameness" or "closeness" in participants might have been not strong enough. Furthermore, although participants in the 3rd experiment were encouraged to cooperate, future manipulations could make the task genuinely collaborative in that optimal performance from both participants would be required to achieve a particular end or to get access to a commonly preferred consequence. It is also possible, of course, that due to the characteristics of the task, like its repetitive nature or the relatively low level of strategizing required to win (with children being outstanding rivals to adults, as pointed out by Dyson et al., 2016), or fatigue or boredom, shared attention might not play any important role. Only systematic replications of the paradigm in other learning tasks will determine whether this is the case.

To the best of our knowledge, this is the first study to provide evidence for strong order effects depending on the sequence of presentation of the experimental contingencies in studies of human behavior allocation. Although more research is required to determine exactly what is it that tends to alter the participants' performance when exposed first to either saliently different or indifferent discrepancies in the probabilities for choice, a strong candidate for these results is the participants' verbal behavior as they are exposed to the contingencies. As Schoenfeld and Cumming (1963) suggested, it is an unfortunate problem in research with human participants that instructions of all sorts are provided to them with little or no

consideration of the effect they might have or even if they have been understood (cf. Tversky & Kahneman, 1981).

Schoenfeld and Cumming advanced a model to understand the function of instructions that seems worthy to be described. According to their view, instructions tend to originate a "perseverative rather than transient" response (they call it "instruction response" or R_i) which persists throughout an experimental session continuously modifying the stimulus variables that participants are being exposed to. Schoenfeld and Cumming speculate that such response probably consists mainly of patterns of subvocal speech which under some circumstances might effectively override effective control by other independent variables. Our study is the first to our knowledge to present participants with *inaccurate* descriptions of contingencies in a behavior allocation task. Although participants continued to win at higher levels, a small but significant reduction in wins occurred in this block with provides some measure of increased "insensitivity" to contingencies, and which was also reflected in altered patterns of strategy, compared with conditions where no instructions were delivered.

A similar case for the role the participants' self-delivered instructions has been suggested in the case of supobtimal choice in the so-called sunk-cost dilemma (e.g., Arkes & Ayton, 1999), who noticed that children are less affected by it and might actually behave more optimally than adults. Adults, according to these authors, tend to behave suboptimally often in such tasks because they don't want to "appear to be wasteful," among other heuristics. Although more direct evidence is required, the fact that providing verbal instructions had such dramatic effects on participants' performance (i.e, more wins) and strategy (an increased frequency of "staying" even in the face of losses) appear to support this approach. These verbal instructions were both beneficial and detrimental for participants in

these experiments. On the one hand, they produced the most success of all manipulations across experiments. On the other hand, the participants' strategy rigidized as a function of their success, leading to slightly lower proportions of winning. A number of further instructional manipulations needs to be conducted in conjunction with our probability manipulations to more precisely ascertain the roles these self-provided and experimenter-provided instructions might play in the allocation of behavior in tasks like the current one.

A final aspect that will require further experimental exploration, and that might be related to this last point is related to the differential effect that draws and losses appeared to have in the participants' strategy, as reflected in a tendency to stay more following draws, which our data suggested might have been enhanced by shared attention. This effect might be related to the long-known phenomenon of "loss aversion" (Kahneman & Tversky, 1979), and might be related to the shared attention effects insofar as social coattention might enhance the salience of losses. Although anecdotal, participants in our shared attention conditions reported reduced enjoyment of the task compared to participants in the solo condition. If this is to some extent an accurate representation, future studies involving differential consequences for losses versus draws might provide some insight into the possibility of the loss aversion effect being moderated by social contingencies like shared attention.

References

Arkes, H.R., & Ayton, P. (1999). The sunk cost and Concorde effects: Are humans less rational than lower animals? *Psychological Bulletin, 125, 591-600*. doi:10.1037/0033-2909.125.5.591

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior, 22*, 231–242.

Baum , W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior, 32*, 269–281.

Beeby, E. & Alsop. B. (2017). Choice among two and three alternatives. *Journal of the Experimental Analysis of Behavior, 107* (3), 369–387.

Berry, D. A., & Fristedt, B. (1985). *Bandit problems: Sequential allocation of experiments*. New York: Chapman and Hall.

Bilalić, M., McLeod, P., & Gobet, F. (2008). Why good thoughts block better ones: The mechanism of the pernicious Einstellung (set) effect. *Cognition, 108* (3), 652–661. doi:10.1016/j.cognition.2008.05.005

Bond A. B., Kamil A. C., & Balda, R. P. (2007). Serial reversal learning and the evolution of behavioral flexibility in three species of North American corvids (*Gymnorhinus cyanocephalus, Nucifraga columbiana, Aphelocoma californica*). *Journal of Comparative Psychology, 121*, 372-379.

Cretenet, J. & Dru, V. (2009). Influence of peripheral and motivational cues on rigid–flexible functioning: Perceptual, behavioral, and cognitive aspects. *Journal of Experimental Psychology: General, 138 (2)*, 201–217. doi:10.1037/a0015379

Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Hillsdale, NJ: Erlbaum.

de Villiers, P. A. (1977). Choice in concurrent schedules and a quantitative formulation of the law of effect. In W. K. Honig & J. E. R. Staddon

(eds.), *Handbook of operant behavior* (pp. 233–287). Englewood Cliffs, NJ: Prentice-Hall.

Deak, G. O. & Narasimham, G. (2003). Is perseveration caused by inhibition failure? Evidence from preschool children's inferences about word meanings. Journal of Experimental Child Psychology, 86, 194–222.

Dyson, B. J., Wilbiks, J. M. P., Sandhu, R., Papanicolaou, G., & Lintag, J. (2016). Negative outcomes evoke cyclic irrational decisions in Rock, Paper, Scissors. *Scientific Reports, 6*, 20479. doi: 10.1038/srep20479

Fantino, E. (2012). Optimal and non-optimal behavior across species. *Comparative Cognition & Behavior Reviews,* 7, 44-54.

Fox, A. E. & Kyonka, E. G. E. (2017).Searching for the variables that control human rule-governed "insensitivity." *Journal of the Experimental Analysis of Behavior, 108* (2), 236–254. doi: 10.1002/jeab.270

Fox, A. E., & Pietras, C. J. (2013). The effects of response cost punishment on instructional control during a choice task. *Journal of the Experimental Analysis of Behavior, 99*, 346-361. doi:10.1002/jeab.20

Frank, M. J., & Claus, E.D. (2006). Anatomy of a decision: Striato-orbitofrontal interactions in reinforcement learning, decision making, and reversal. *Psychological Review, 113*, 300-326.

Ghahremani, D.G., Monterosso, J., Jentsch, J.D., Bilder, R.M., & Poldrack, R.A. (2010). Neural components underlying behavioral flexibility in human reversal learning. *Cerebral Cortex, 20*, 1843–1852. doi:10.1093/cercor/bhp247

Gosling, S. D., Rentfrow, P. J., & Swann, W. B., Jr. (2003). A very brief measure of the big five personality domains. *Journal of Research in Personality, 37*, 504-528.

Grace, R. C. & Hucks, A. D. (2013). The allocation of operant behavior. In G. J. Madden (Ed.), *APA handbook of behavior analysis* (Vol. 1, pp. 307-337). Washington, DC: American Psychological Association.

Hackenberg, T. D., & Joker, V. R. (1994). Instructional versus schedule control of humans' choices in situations of diminishing returns. *Journal of the Experimental Analysis of Behavior, 62*, 367-383. doi:10.1901/jeab.1994.62-367

Hamilton, D.A., & Brigman, J. L. (2015). Behavioral flexibility in rats and mice: Contributions of distinct frontocortical regions. *Genes, Brain and Behavior, 14*, 4-21.

Hayes, S. C., Zettle, R. D., & Rosenfarb, I. (1989). Rule-following. In S. C. Hayes (Ed.), *Rule-governed behavior: Cognition, contingencies, and instructional control* (pp. 191-220). NY: Plenum Press.

Herbranson, W. T., & Schroeder, J. (2010). Are birds smarter than mathematicians? Pigeons (*Columba livia*) perform optimally on a version of the Monty Hall dilemma. *Journal of Comparative Psychology, 124*, 1-13. doi:10.1037/a0017703

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior, 4*, 267–272.

Hornak, J., O'Doherty, J., Bramham, J., Rolls, E.T., Morris, R.G., et al. (2004). Reward-related reversal learning after surgical excisions in orbito-frontal or dorsolateral prefrontal cortex in humans. *Journal of Cognitive Neuroscience, 16*, 463–478.

Horne, P. J., & Lowe, C. F. (1993). Determinants of human performances on concurrent schedules. *Journal of the Experimental Analysis of Behavior, 59*, 29-60.

Ionescu, T. (2012). Exploring the nature of cognitive flexibility. *New Ideas in Psychology, 30* (2), 190–200. doi: 10.1016/j.newideapsych.2011.11.001

Izquierdo, A., & Jentsch, J. D. (2012). Reversal learning as a measure of impulsive and compulsive behavior in addictions. *Psychopharmacology, 219*(2), 607–620. doi: 10.1007/s00213-011-2579-7

James, W. (1890). *The principles of psychology*. Retrieved January 2018 from http://psychclassics.yorku.ca/James/Principles/

Kahneman D., & Tversky A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263–291. doi:10.1007/s11336-014-9425-x

Kangas, B.D., & Bergman, J. (2014). Repeated acquisition and discrimination reversal in the squirrel monkey (*Saimiri sciureus*). *Animal Cognition, 17*, 221-228.

Kangas, B.D., Berry, M.S., Cassidy, R.N., Dallery, J., Vaidya, M., & Hackenberg, T.D. (2009). Concurrent performance in a three-alternative choice situation: Response allocation in a rock/paper/scissors game. *Behavioural Processes, 82*, 164–172. doi:10.1016/j.beproc.2009.06.004

Killen, P. (2014). Pavlov + Skinner = Premack. *International Journal of Comparative Psychology, 27*(4), 544-568.

Klanker, M., Feenstra, M., & Denys, D. (2013). Dopaminergic control of cognitive flexibility in humans and animals. *Frontiers in Neuroscience*, 7, Art. 201. doi: 10.3389/fnins.2013.00201

Koch, I., Poljac, E., Müller, H., & Kiesel, A. (2018). Cognitive structure, flexibility, and plasticity in human multitasking: An integrative review of dual-task and task-switching research. *Psychological Bulletin, 144* (6), 557–583. doi:10.1037/bul0000144

Kollins, S. H., Newland, M. C., & Critchfield, T. S. (1997). Human sensitivity to reinforcement in operant choice: How much do consequences matter? *Psychonomic Bulletin & Review, 4*(2), 208-220.

Kroger-Costa, A., & Abreu-Rodrigues, J. (2012). Effects of historical and social variables on instruction following. *The Psychological Record, 62*, 691–706

Kubanek, J. (2017). Optimal decision making and matching are tied through diminishing returns. *PNAS, 114* (32), 8499–8504.

Lie, C., Baxter, J., & Alsop, B. (2013). The effect of opponent type on human performance in a three-alternative choice task. *Behavioural Processes, 99*, 87–94. doi:10.1016/j.beproc.2013.07.005

Luchins, A. S. (1942). Mechanization in problem solving: The effect of Einstellung. *Psychological Monographs, 54*(6), 95.

Luchins, A. S., & Luchins, E. H. (1959). *Rigidity of behavior: A variational approach to the effect of Einstellung*. Eugene, OR: University of Oregon Books.

McMillan, N., Hahn, A. H., Congdon, J. V., Campbell, K. A., Hoang, J., et al. (2017). Chickadees discriminate contingency reversals presented consistently, but not frequently. *Animal Cognition, 2*, 655–663. doi: 10.1007/s10071-017-1088-4

Madden, G. J., Chase, P. N., & Joyce, J. H. (1998). Making sense of sensitivity in the human operant literature. *The Behavior Analyst, 21*, 1-12. doi:1007/BF03392775

Masicampo, E. J., & Baumeister, R. F. (2012). Committed but closed-minded: When making a specific plan for a goal hinders success. *Social Cognition, 30* (1), 37–55.

McDowell, J. J. (1988). Matching theory in natural human environments. *The Behavior Analyst, 11*, 95–109.

McDowell, J. J. (2005). On the classic and modern theories of matching. Journal of the Experimental Analysis of Behavior, 84, 111–127.

McDowell, J. J. (2013). On the theoretical and empirical status of the matching law and matching theory. *Psychological Bulletin, 139*, 1000–1028.

Molet, M., Miller, H. C., Laude, J. R., Kirk, C., Manning, B., & Zentall, T. R. (2012). Decision making by humans in a behavioral task: Do humans, like pigeons, show suboptimal choice? *Learning & Behavior, 40*(4), 439-447. doi: 10.3758/s13420-012-0065-7.

Page, R. A., & Ryan, M. J. (2005). Flexibility in assessment of prey cues: Frog-eating bats and frog calls. *Proceedings of the Royal Society B-Biological Sciences, 272*, 841-847.

Pierce, W. D., & Epling, W. F. (1983). Choice, matching, and human behavior: A review of the literature *The Behavior Analyst 1983, 6* (1), 57-76.

Poling, A., Edwards, T. L., Weeden, M., & Foste, T. M. (2011). The matching law. *The Psychological Record, 61*, 313–322.

Reed, D. D., & Kaplan, B. A. (2011). The matching law: A tutorial for practitioners. *Behavior Analysis in Practice, 4*(2), 15–24

Roberts, A. C. (2006). Primate orbitofrontal cortex and adaptive behaviour. *Trends in Cognitive Sciences, 10*, 83-90.

Roberts, W. A., MacDonald, H., & Lo, K. H. (2018). Pigeons play the percentages: Computation of probability in a bird. *Animal Cognition, 21*, 575–581. doi:10.1007/s10071-018-1192-0

Rodríguez, M. L., & Logue, A. W. (1986). Independence of the amount and delay ratios in the generalized matching law. *Animal Learning & Behavior, 14* (1), 29-37.

Schoenfeld, W. N., & Cumming, W. W. (1963). Behavior and perception. In S. Koch (Ed.), *Psychology: A study of a science* (vol. 5, pp. 213-252). New York: McGraw-Hill.

Schultz, P. W., & Searleman, A. (2002). Rigidity of thought and behavior: 100 years of research. *Genetic, Social, and General Psychology Monographs, 128*(2), 165-207.

Sheridan, H., & Reingold, E. M. (2013). The mechanisms and boundary conditions of the Einstellung effect in chess: Evidence from eye movements. *PLoS ONE, 8*(10), e75796. doi:10.1371/journal.pone.0075796

Sheridan, H., & Reingold, E. M. (2014). Expert vs. novice differences in the detection of relevant information during a chess game: evidence from

eye movements. *Frontiers in Psychology, 5*, 941. doi: 10.3389/fpsyg.2014.00941

Sheridan, H., & Reingold, E. M. (2017). Chess player's eye movements reveal rapid recognition of complex visual patterns: Evidence from a chess-related visual search task. *Journal of Vision, 17*(3), 1-12. doi: 10.1167/17.3.4

Shimoff, E., Catania, A. C., & Matthews, B. A. (1981). Uninstructed human responding: Sensitivity of low-rate performance to schedule contingencies. *Journal of the Experimental Analysis of Behavior, 36*, 207-220.

Shteynberg, G. (2010). A silent emergence of culture: The social tuning effect. *Journal of Personality and Social Psychology, 99*, 683–689.

Shteynberg, G. (2015). Shared attention. *Perspectives on Psychological Science, 10*(5) 579–590. doi:10.1177/1745691615589104

Sih, A., Bell, A., & Johnson, J. C. (2004). Behavioral syndromes: An ecological and evolutionary overview. *Trends in Ecology and Evolution, 19* (7), 372-378. doi:10.1016/j.tree.2004.04.009

Sih, A., Bell, A. M., Johnson, J. C. & Ziemba, R. E. (2004). Behavioral syndromes: An integrative overview. *The Quarterly Review of Biology, 79*, 241–277.

Spetch, M. L., Belke, T. W., Barnet, R. C, Dunn, R., & Pierce, W. D. (1990). Suboptimal choice in a percentage-reinforcement procedure: Effects of signal condition and terminal-link length. *Journal of the Experimental Analysis of Behavior, 53*, 219-234. doi:10.1901/jeab.1990.53-219

Staddon, J. E. R. (2001). *Adaptive dynamics: The theoretical analysis of behavior.* Cambridge, MA: Bradfod/MIT Press.

Takahashi, M. & Shimakura, T. (1998). The effects of instructions on human matching. *Psychological Record, 48* (1), 171-181.

Tomasello, M. (2011). Human culture in evolutionary perspective. In M. J. Gelfand, C. Chiu, & Y. Hong (Eds.), *Advances in culture and psychology: Vol. 1. Advances in culture and psychology* (pp. 5-51). New York, NY, US: Oxford University Press.

Tversky, A. & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211* (4481), 453–458. doi:10.1126/science.7455683

Vallée-Tourangeau, F., Euden, G., & Hearn, V. (2011). Einstellung defused: Interactivity and mental set. *The Quarterly Journal of Experimental Psychology, 64* (10), 1889-1895. doi:10.1080/17470218.2011.605151

Vandierendonck, A., Liefooghe, B., & Verbruggen, F. (2010). Task switching: Interplay of reconfiguration and interference control. *Psychological Bulletin, 136*, 601–626.

Vaughan, M. (1989). Rule-governed behavior in behavior analysis: A theoretical and experimental history. In S. C. Hayes (Ed.), *Rule-governed behavior: Cognition, contingencies, and instructional control* (pp. 97-118). NY: Plenum Press.

Wang, Z., Xu, B., Zhou, H.-J. (2014). Social cycling and conditional responses in the in Rock-Paper-Scissors game. *Scientific Reports, 4*, 5830. doi: 10.1038/srep20479

Weir, M. W. (1964). Developmental changes in problem-solving strategies. *Psychological Review, 71*, 473-490. doi: 10.1037/h0041785

Williams, B. A. (1988). Reinforcement, choice, and response strength. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Stevens' handbook of experimental psychology: Vol. 2: Learning and cognition* (pp. 167–244). New York: Wiley.

Xue, G., Xue, F., Droutman, V., Lu, Z-L., Bechara, A., & Read, S. (2013). Common neural mechanisms underlying reversal learning by reward and punishment. *PLoS ONE, 8*, e82169. doi:10.1371/journal.pone.0082169

Appendix

Table 2.1

Shared Attention Conditions for the Pilot Experiment

Conditions	Co-attending Observer	Response Feedback	Outcome Feedback
Alone (Group 1)	No	Yes	Yes
Shared-Outcome-Only (Group 2)	Yes	No	Yes
Shared Performance & Outcome (Group 3)	Yes	Yes	Yes

Note: Response feedback consisted in auditory description of the participant response (e.g., "Paper" when the corresponding choice was selected). Outcome feedback consisted in auditory feedback corresponding to the outcome of the trial (e.g., applause if the trial was a win).

Table 3. 1
Experimental Conditions in Experiment 2

Conditions	Co-attending Observer	Response Feedback	Outcome Feedback	Discriminability Sequence
Alone HtE (Group 1)	No	Yes	Yes	Hard-to-Easy
Shared-Outcome- Only HtE (Group 2)	Yes	No	Yes	Hard-to-Easy
Shared Performance & Outcome HtE (Group 3)	Yes	Yes	Yes	Hard-to-Easy
Alone EtH (Group 4)	No	Yes	Yes	Easy-to-Hard
Shared-Outcome- Only EtH (Group 5)	Yes	No	Yes	Easy-to-Hard
Shared Performance & Outcome EtH(Group 6)	Yes	Yes	Yes	Easy-to-Hard

Note: Discriminability sequence refers to the varying degrees of discrepancy in the probability that the computer will choose a particular option, ranging from "Hard," where the computer chooses each option with the same probability, to "Easy," where the computer chooses one option 80% of the time. In the Hard-to-Easy sequence, the computer starts with equal probability and starts to increase the probability of choosing one option across blocks. In the Easy-to-Hard sequence, the computer starts by choosing one option with the highest probability (80% of the time) and that probability starts decreasing across blocks until in the final block each option is played 33% of the time.

Table 3.2

Sensitivity and Bias Parameters of the GML Analysis for Pairs of Choices across Conditions

		Hard to Easy			Easy to Hard		
		S	*b*	R^2	*s*	*b*	R^2
Solo	R/P	0.46	0.06	0.51	0.547	0.052	0.707
	R/S	0.45	0.03	0.58	0.470	0.068	0.739
	P/S	0.36	0.01	0.54	0.474	0.015	0.677
Shared Outcome Only	R/P	0.343	0.028	0.538	0.392	0.048	0.704
	R/S	0.418	0.054	0.652	0.492	0.032	0.698
	P/S	0.335	0.031	0.566	0.502	0.028	0.759
Shared Outcome & Feedback	R/P	0.374	0.003	0.648	0.462	0.062	0.687
	R/S	0.433	0.063	0.574	0.393	0.087	0.603
	P/S	0.362	0.062	0.590	0.376	0.022	0.674

Table 3.3

Correlation Between Performance and Strategy Outcomes & Self-Report Measures

	Diffic	Enjoy	Satisf	Liking	Close	Wins	Draw	Lose	Stay	Upgr	Downgr
Perceived Difficulty		.21**	-.02	-.185*	-.06	-.08	-.04	.15	-.01	.05	-.04
Enjoyment	.21**		.37**	.10	.05	-.08	.08	.03	-.13	.13	.10
Satisfaction with Performance	-.02	.37**		.06	.04	.26**	.00	-.36**	.04	-.02	-.06
Partner Liking	-.19*	.10	.06		.46**	-.12	.02	.13	-.10	.04	.13
Closeness	-.06	.05	.04	.46**		-.08	.06	.04	-.05	.04	.04
Wins	-.08	-.08	.26**	-.12	-.08		-.65**	-.68**	.59**	-.48**	-.53**
Draws	-.04	.08	.00	.02	.06	-.65**		-.11	-.44**	.39**	.37**
Loses	.15	.03	-.36**	.13	.04	-.68**	-.11		-.35**	.25**	.34**
Stay	-.01	-.13	.04	-.10	-.05	.58**	-.44**	-.35**		-.87**	-.85**
Upgrade	.05	.13	-.02	.04	.04	-.47**	.39**	.25**	-.87**		.48**
Downgrade	-.04	.10	-.06	.13	.04	-.53**	.37**	.34**	-.85**	.48**	

* p<.05
** p<.001

Table 4.1

Correlation Matrix for the TIPI Personality Measure and the Other Self-Report Measures in Experiment 3

	Extrav	Agree	Consc	Emot	Open	Diffic	Enjoy	Satisf	Liking	Close
Extraversion	-	-.18	-.01	.08	.17	.03	-.10	-.04	.07	-.08
Agreeableness	-.18	-	.00	-.02	.16	.00	.02	.05	-.04	.24
Conscientiousness	-.01	.00	-	.17	.02	.03	.08	.06	.23	.05
Emotional Stability	.08	-.02	.17	-	.01	-.14	.190*	.291**	.08	-.06
Openness to Experience	.17	.16	.02	.01	-	-.10	-.217*	.13	-.05	-.06
Difficulty	.03	.00	.03	-.14	-.10	-	.11	-.13	.03	.04
Task Enjoyment	-.10	.02	.08	.190*	-.217*	.11	-	.16	.268*	.23
Satisfaction with Performance	-.04	.05	.06	.291**	.13	-.13	.16	-	.21	.21
Partner Liking	.07	-.04	.23	.08	-.05	.03	.268*	.21	-	.486**
Partner Closeness	-.08	.24	.05	-.06	-.06	.04	.23	.21	.486**	-

Table 4.2

Correlation Matrix for the TIPI Personality Measure and the Performance and Strategy of Participants in Experiment 3

	Wins	Draws	Loses	Stay	Upgr	Downg	Win-Stay	Win-Upgr	Win-Downg	Loss-Stay	Loss-Upgr	Loss-Downg	Draw-Stay	Draw-Upgr	Draw-Downgr
Extraversion	-.065	.036	.043	-.010	.013	.006	-.009	-.028	-.024	-.032	-.103	.162	.009	.125	-.151
Agreeableness	-.004	.097	-.083	-.057	.042	.069	-.060	.075	.074	-.057	-.014	-.013	-.026	.035	.123
Conscien-tiousness	-.088	-.024	.123	-.108	.146	.062	-.102	.133	.042	-.071	.159	.076	-.098	.086	.035
Emotional Stability	.092	-.109	-.008	.059	-.055	-.059	.082	-.042	-.063	.024	.055	-.079	.020	-.121	.004
Openness to Experience	.069	.054	-.128	.025	.013	-.062	.064	-.017	-.066	-.067	-.101	.023	.016	.115	-.126
Difficulty	-.18	.059	.154	-.260**	.268**	.234*	-.229*	.215*	.175	-.208*	.237*	.208*	-.232*	.223*	.205*
Task Enjoyment	-.040	-.092	.131	-.087	.081	.087	-.081	.078	.083	-.039	.207*	.016	-.093	-.039	.129
Satisfaction with Performance	.30**	-.181	-.186	.203*	-.210*	-.182	.190*	-.098	-.047	.158	-.154	-.246**	.166	-.26**	-.166
Partner Liking	-.14	.313**	-.111	-.121	.125	.112	-.126	.073	.100	-.280*	.068	.133	.038	.16	.044
Partner Closeness	.11	-.012	-.113	.076	-.080	-.069	.045	.029	-.005	.048	-.113	-.076	.112	-.111	-.105

* p<.05

**p.<001

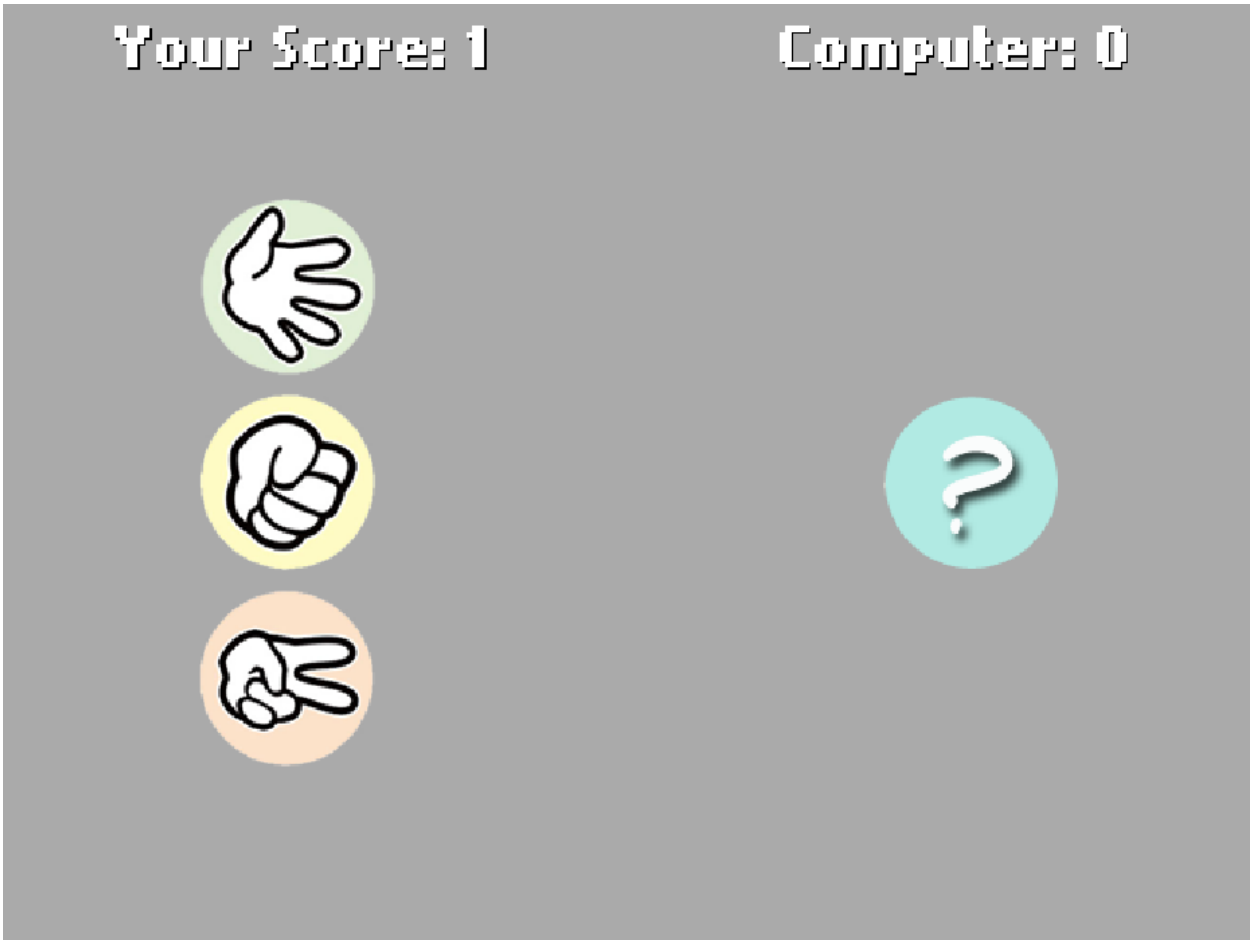

Figure 2.1. A prototypical trial configuration.

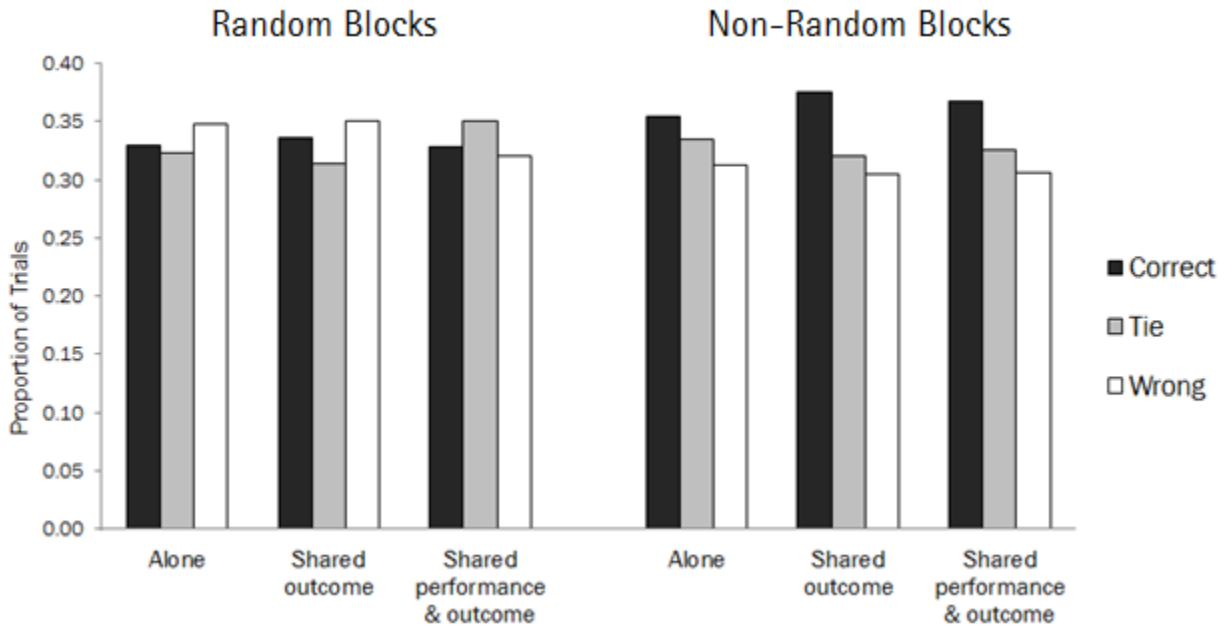

Figure 2.2. Mean proportion of correct and incorrect (i.e., ties and losses) attempts for participants in the tree experimental conditions as a function of the type of block, random (i.e., the block where the computer chose each alternative with equal probability), or non-random (comprising the three blocks at which the computer chose one alternative with a higher probability).

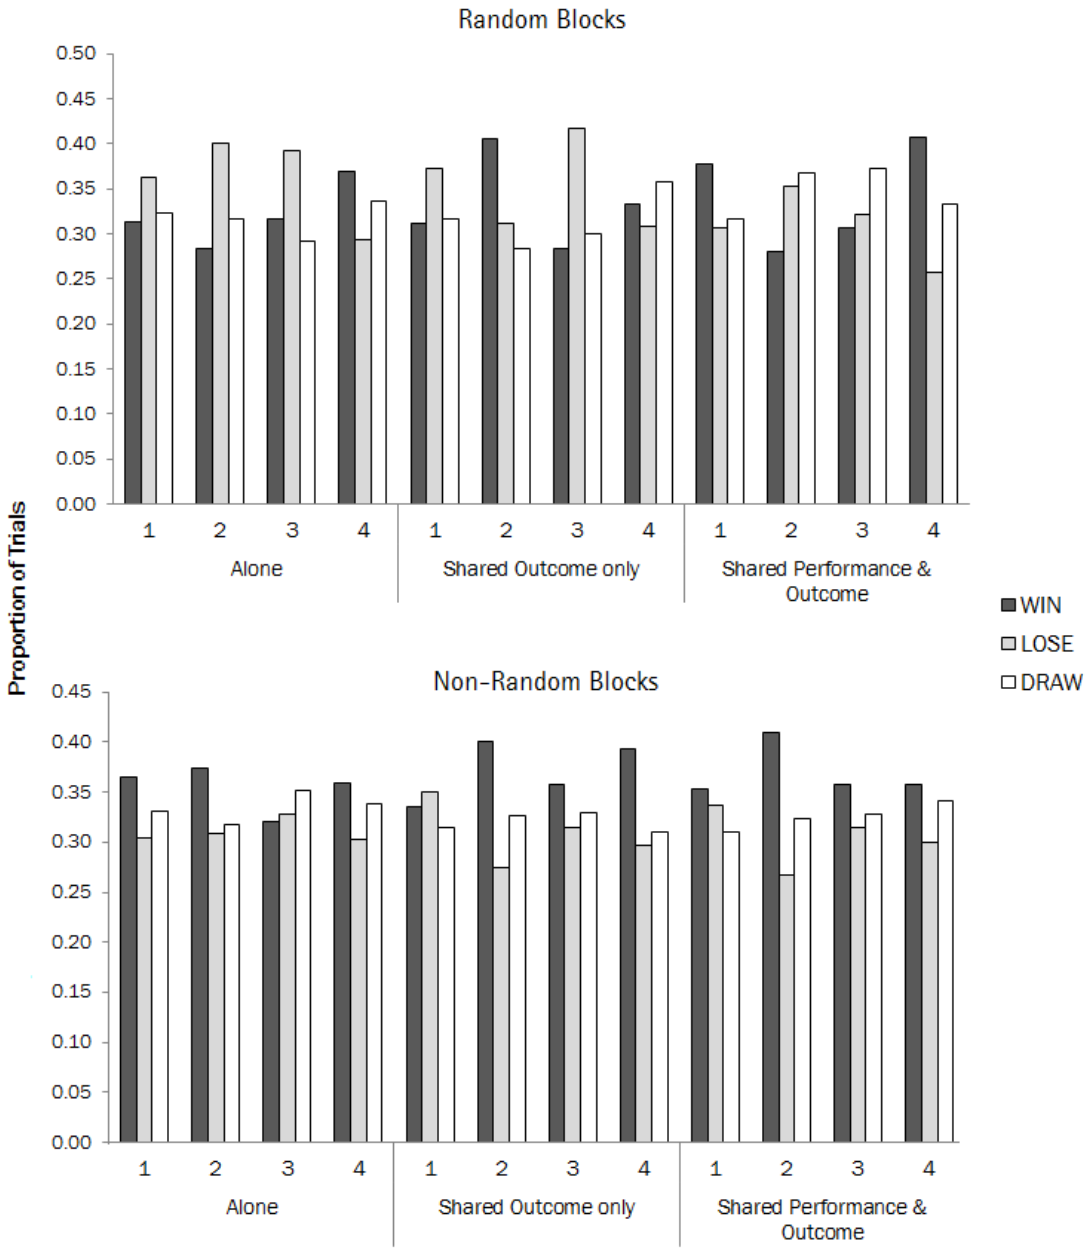

Figure 2.3. Proportion of correct and incorrect (i.e., draws and losses) attempts for participants in the three experimental conditions across blocks. The upper panel corresponds to random blocks (i.e., blocks where the computer chose all alternatives with the same probability), while the bottom panel corresponds to the strategy in non-random blocks (i.e., those where the computer chose one alternative with a higher probability).

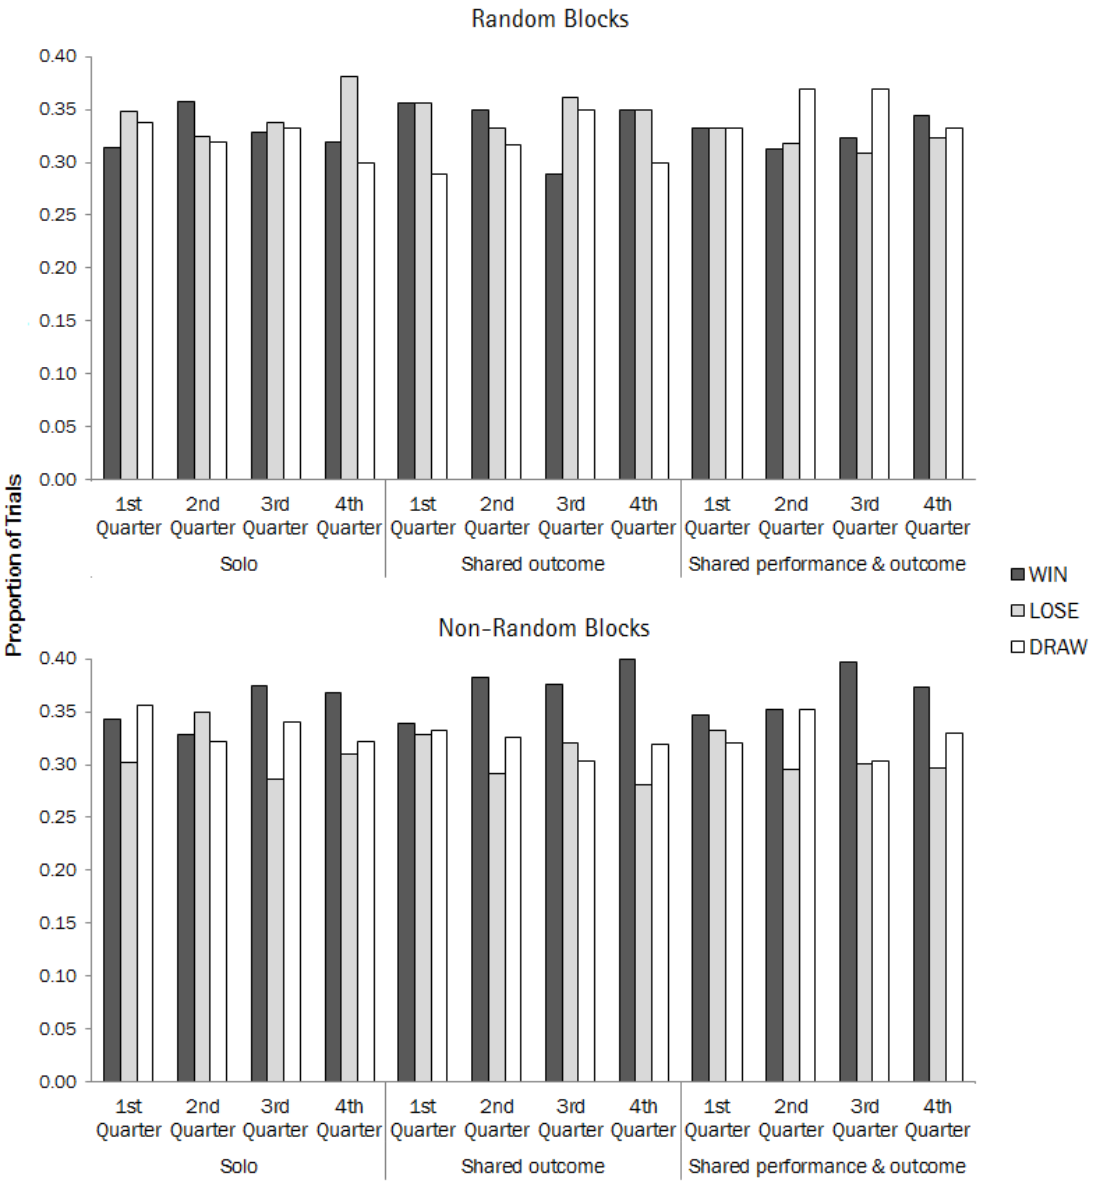

Figure 2.4. Proportion of correct and incorrect (i.e., draws and losses) attempts for participants in the three experimental conditions across block portion. The upper panel corresponds to random blocks (i.e., blocks where the computer chose all alternatives with the same probability), while the bottom panel corresponds to the strategy in non-random blocks (i.e., those where the computer chose one alternative with a higher probability).

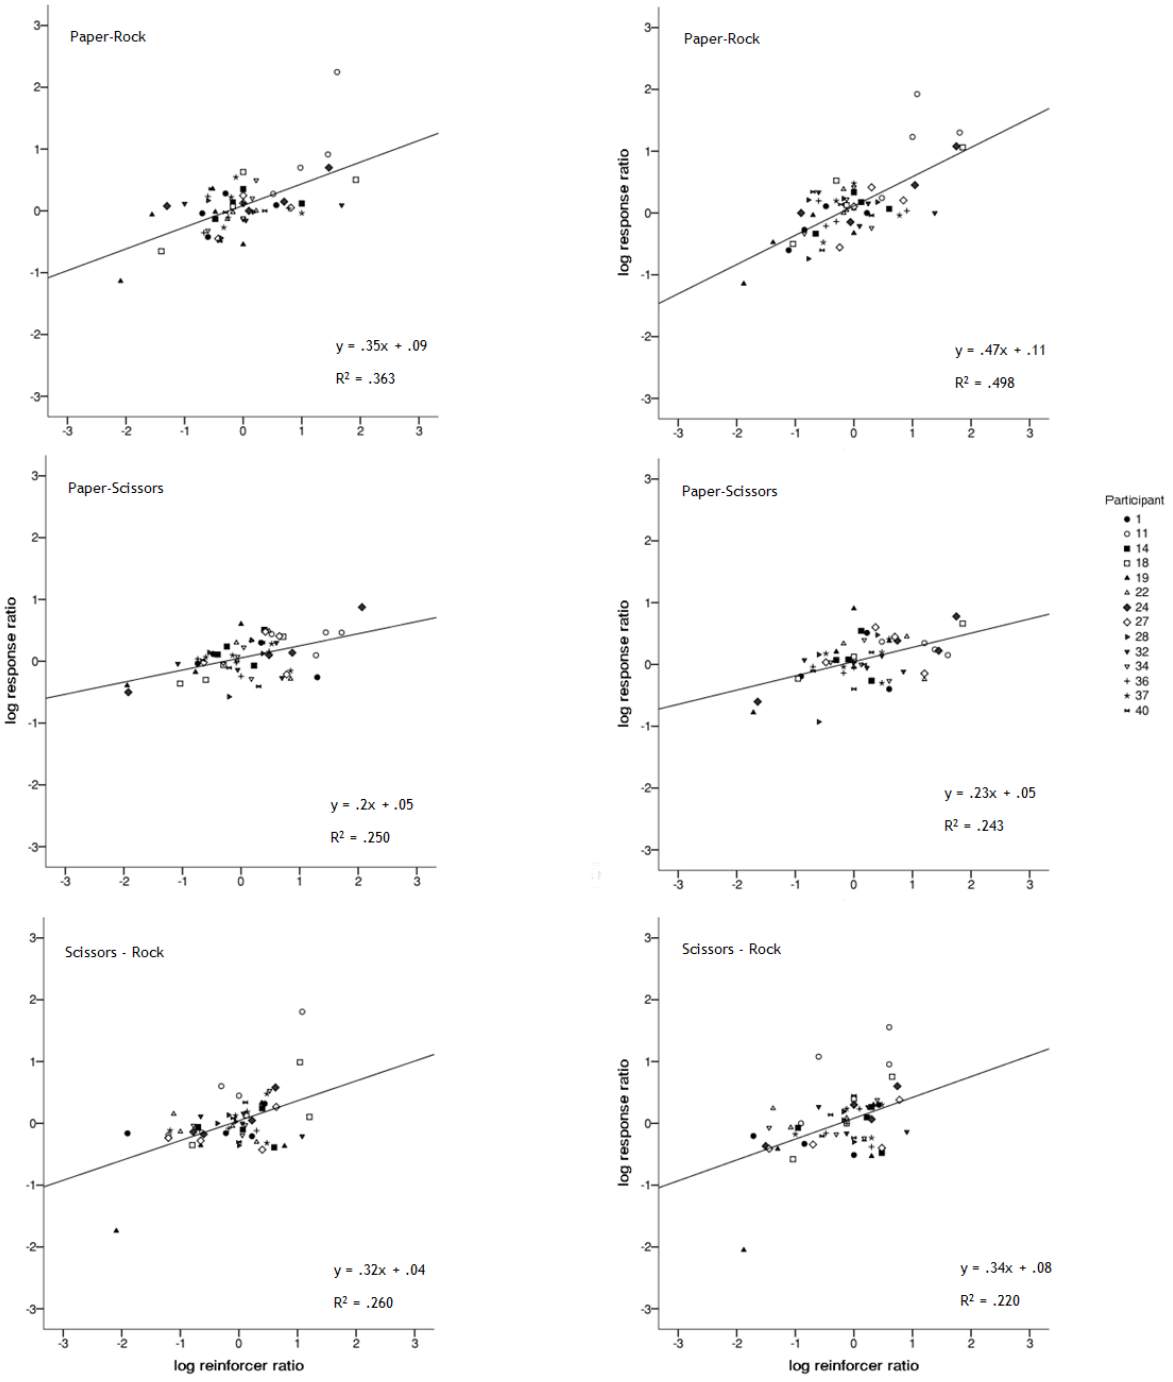

Figure 2.5. Log response ratios as a function of log reinforcement ratios for participants in the Solo condition of Experiment 1. The left panel depicts results for the entire blocks, while the right panel depicts results from trials from the latter half of each block.

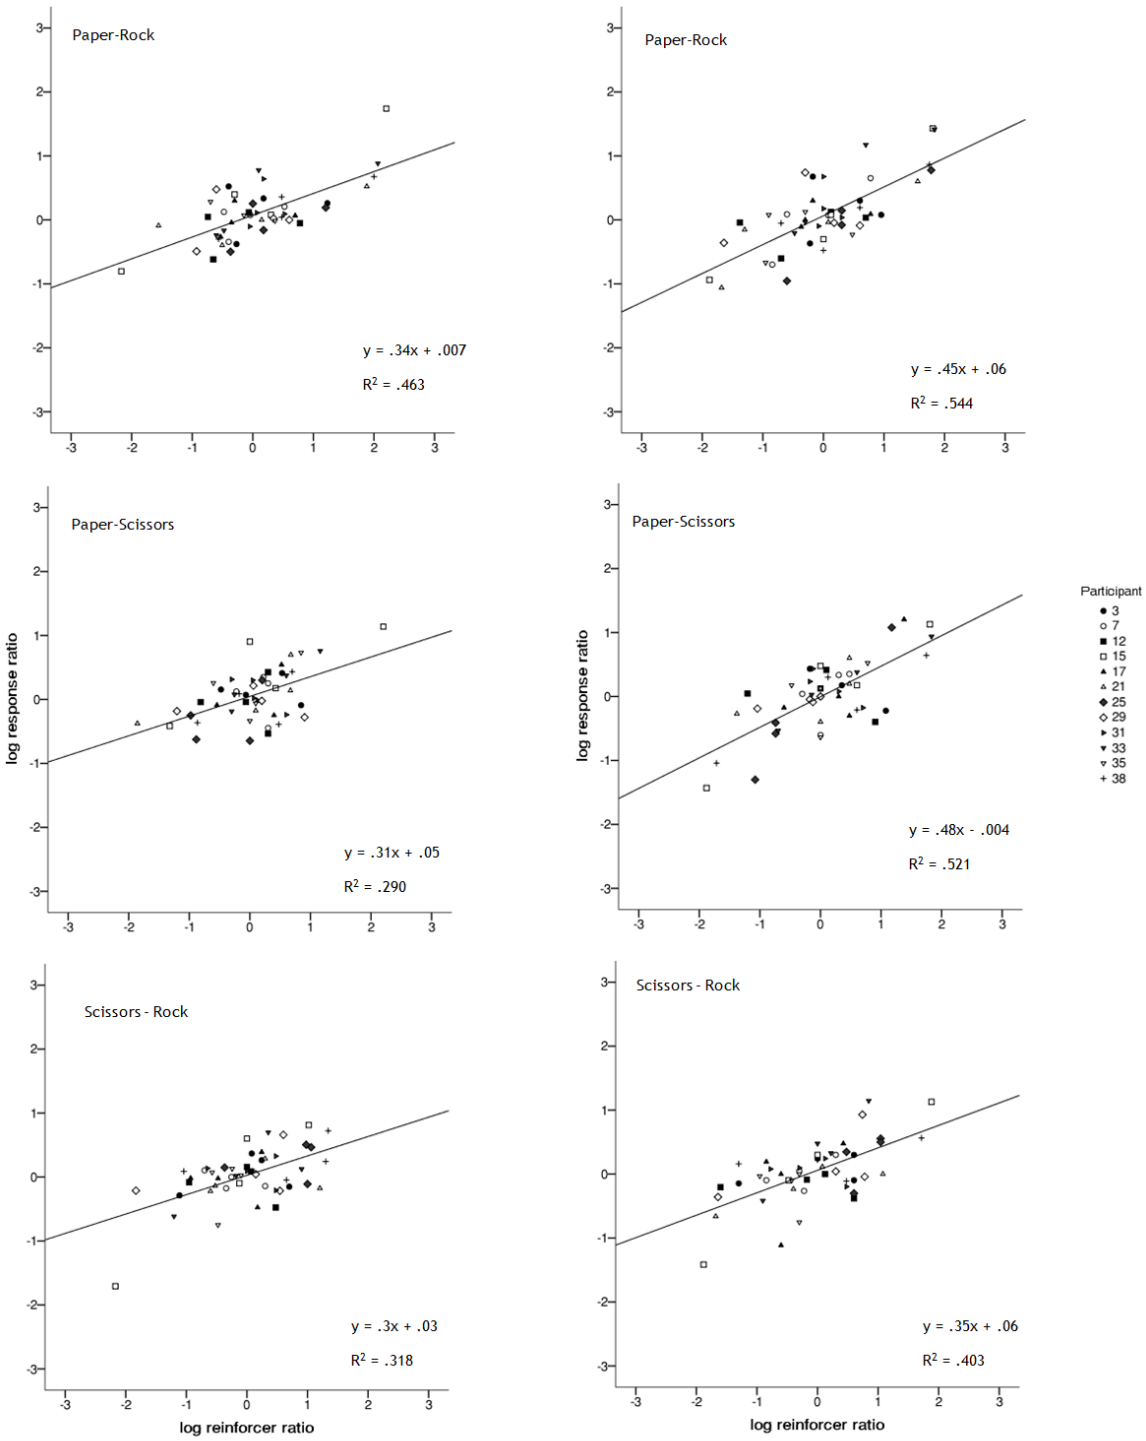

Figure 2.6. Log response ratios as a function of log reinforcement ratios for participants in the Shared-Outcome-Only condition of Experiment 1. The left panel depicts results for the entire blocks, while the right panel depicts results from trials from the latter half of each block.

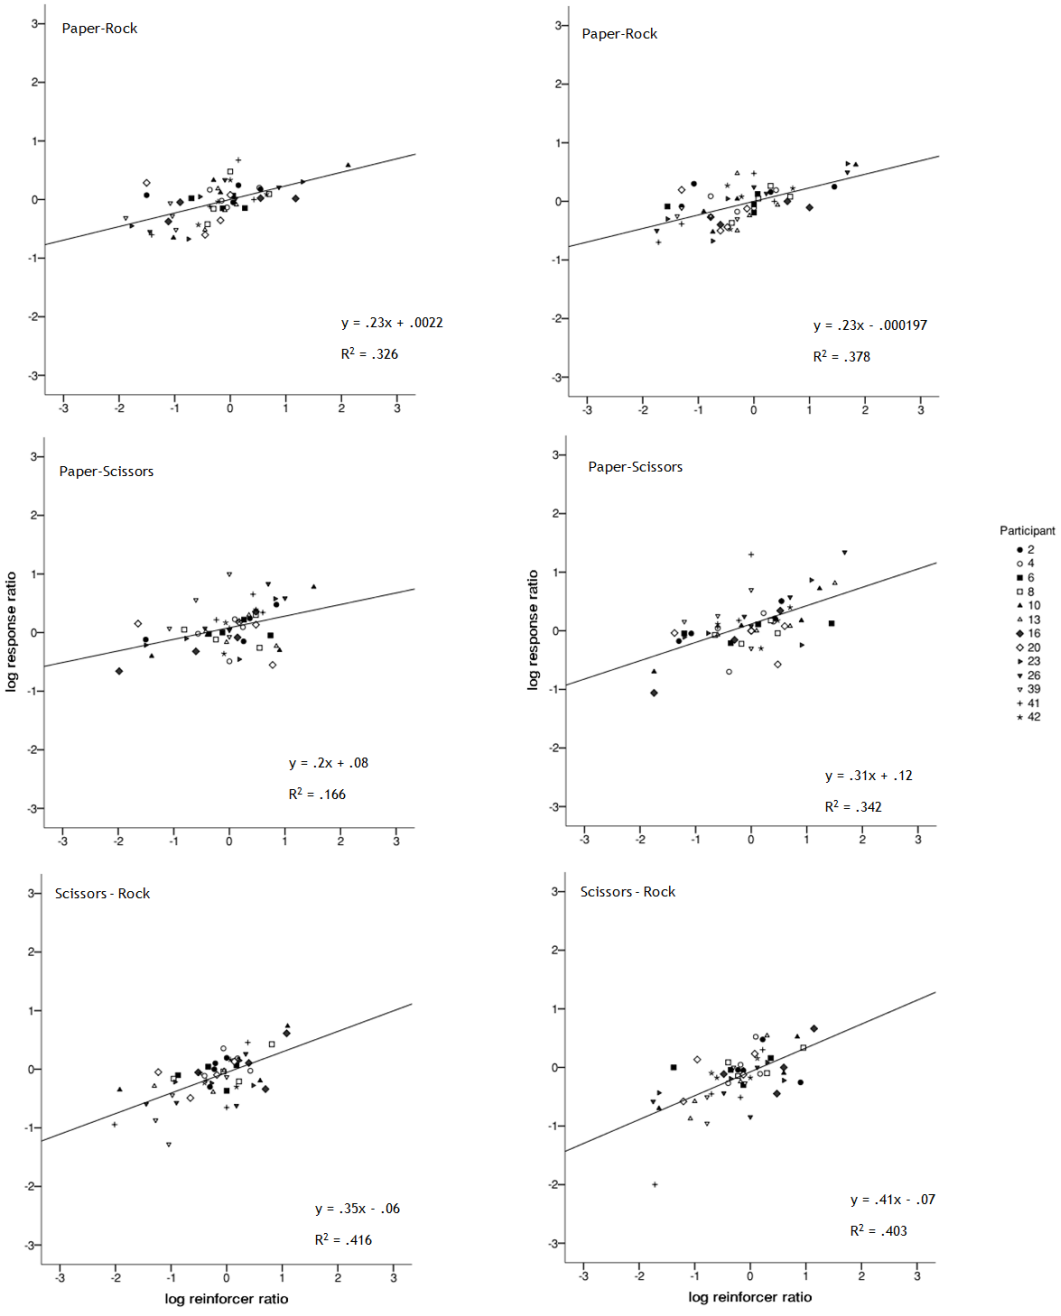

Figure 2.7. Log response ratios as a function of log reinforcement ratios for participants in the Shared-Performance-&-Outcome condition of Experiment 1. The left panel depicts results for the entire blocks, while the right panel depicts results from trials from the latter half of each block.

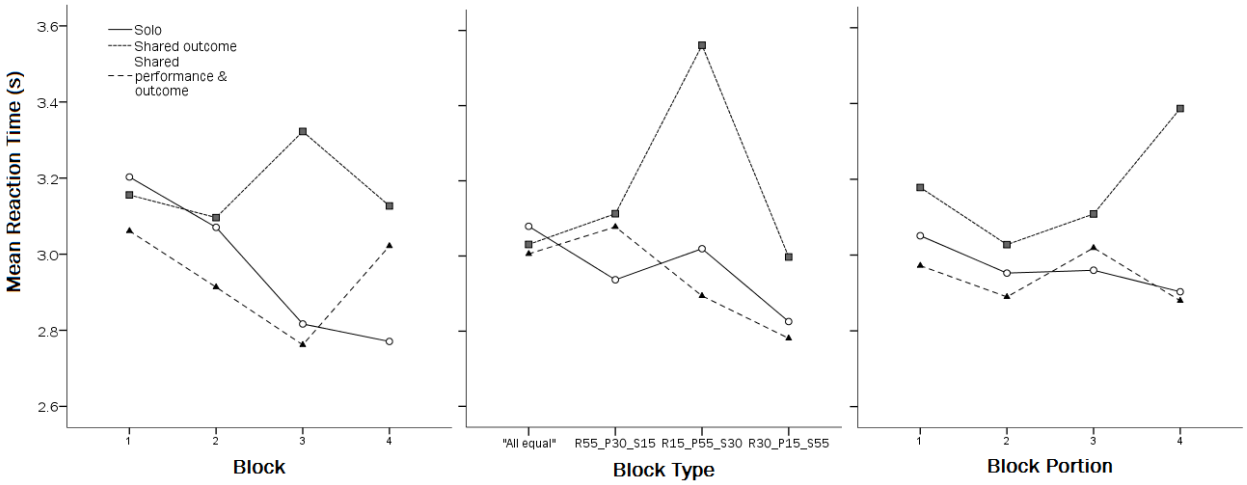

Figure 2.8. Estimated marginal means of the response times across blocks, across block types, and across block portion for the participants in the three experimental groups.

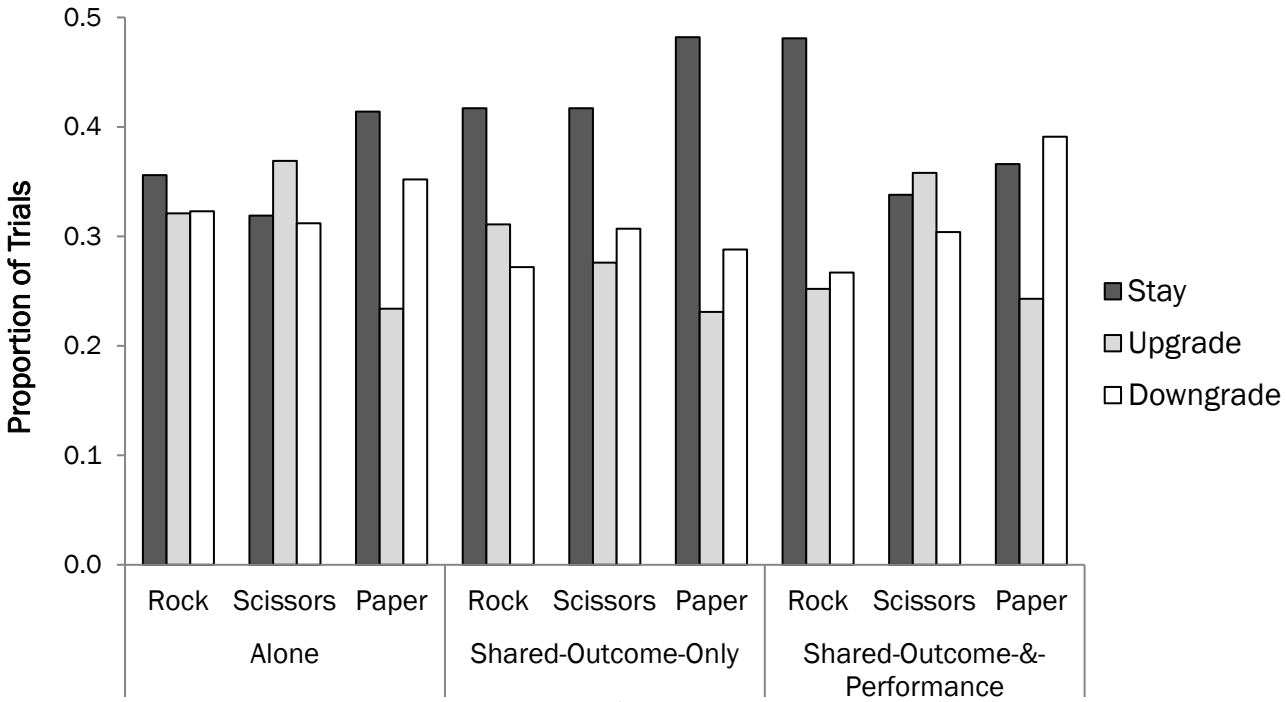

Figure 2.9. Overall strategy at trial *n + 1* as a function of the selection at trial *n* for participants in the three experimental conditions.

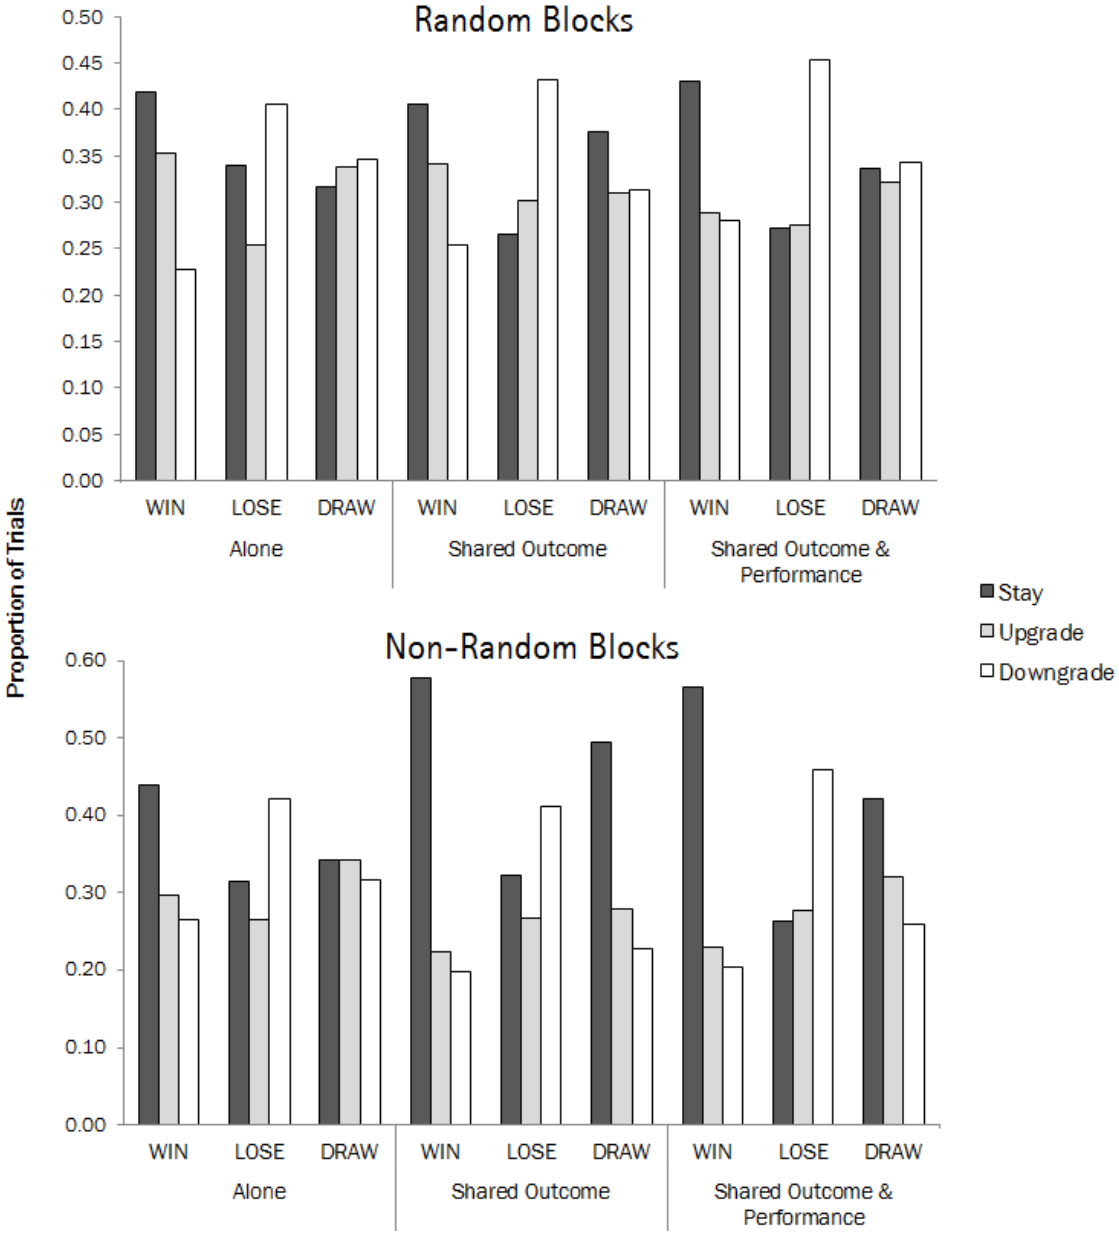

Figure 2.10. Strategy at trial *n + 1* as a function of the *outcome* of trial *n* for participants in the three experimental conditions. The upper panel corresponds to random blocks (i.e., blocks where the computer chose all alternatives with the same probability), while the bottom panel corresponds to the strategy in non-random blocks (i.e., those where the computer chose one alternative with a higher probability).

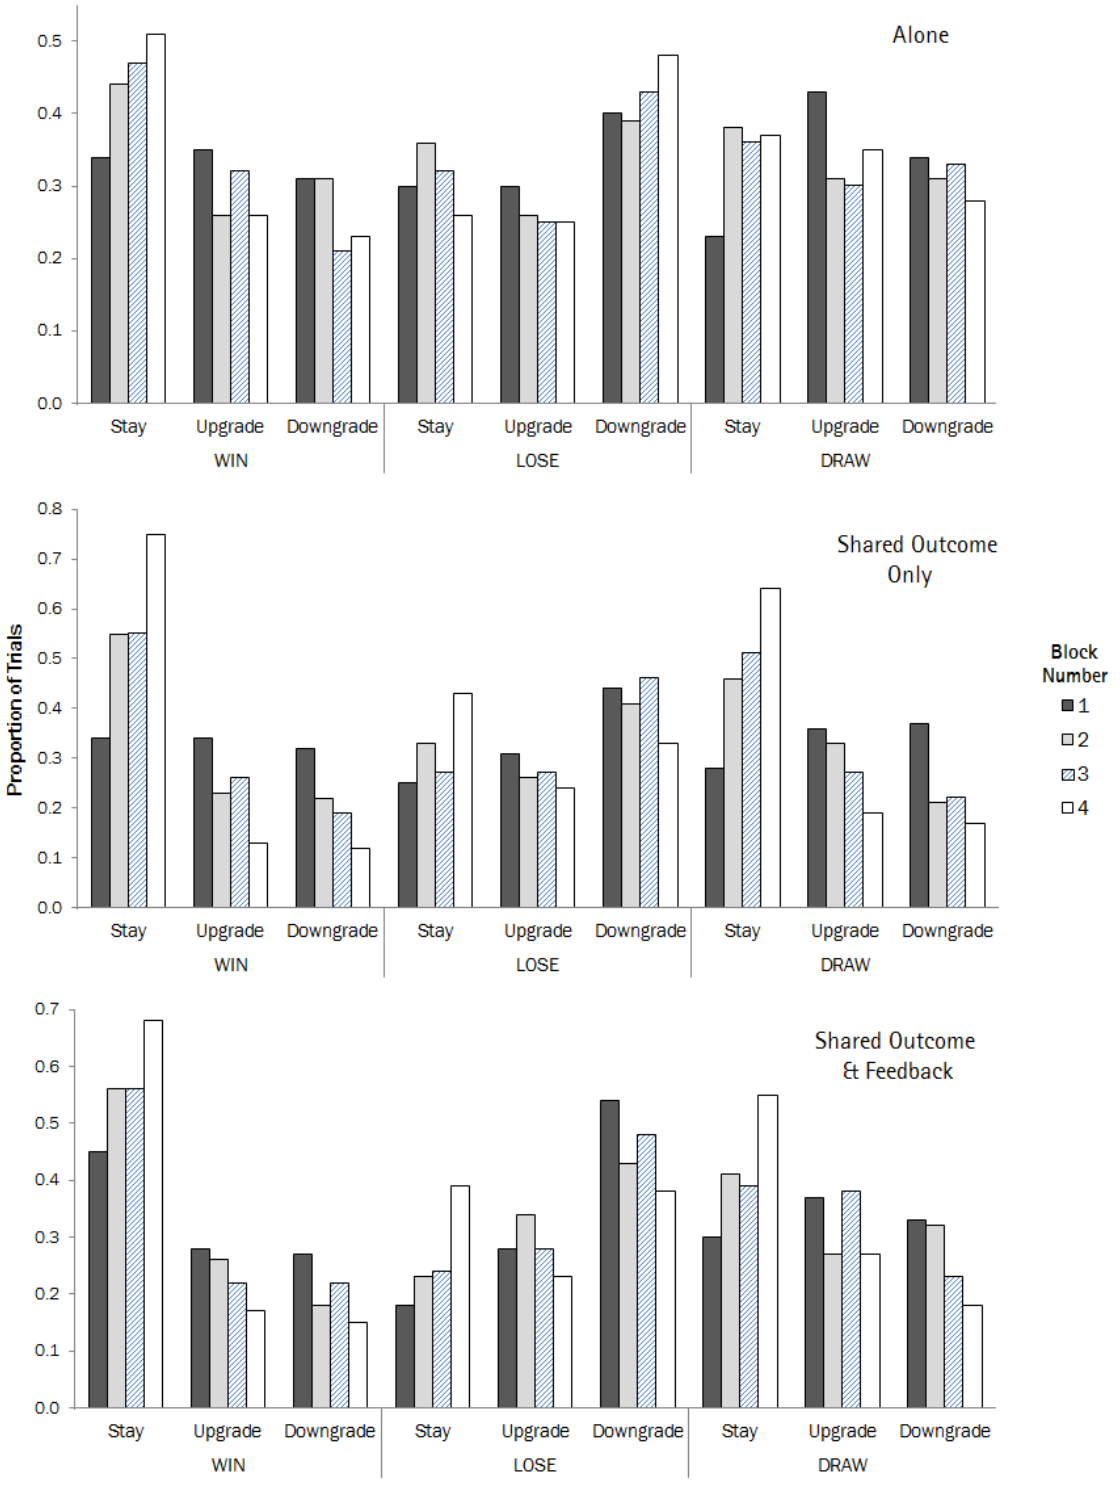

Figure 2.11. Strategy at trial *n + 1* as a function of the outcome of trial *n* across (non-random) blocks for participants in the three experimental conditions. Random blocks have been excluded.

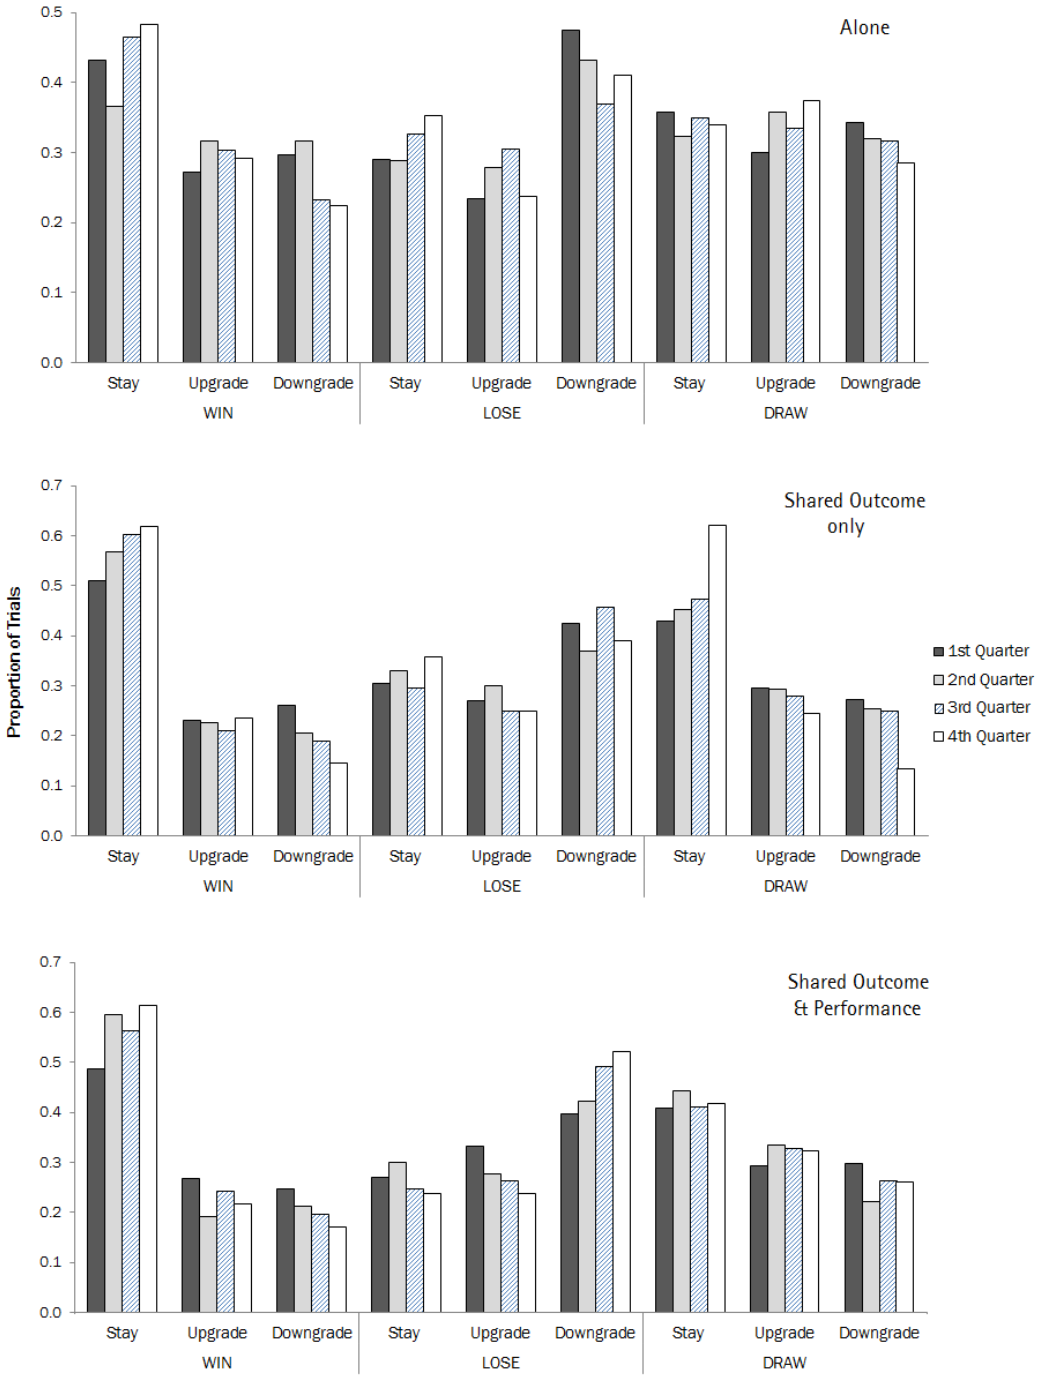

Figure 2.12. Strategy at trial *n + 1* as a function of the outcome of trial *n* across (non-random) block portion (each portion corresponding to one fourth of the block) for participants in the three experimental conditions. Random blocks have been excluded.

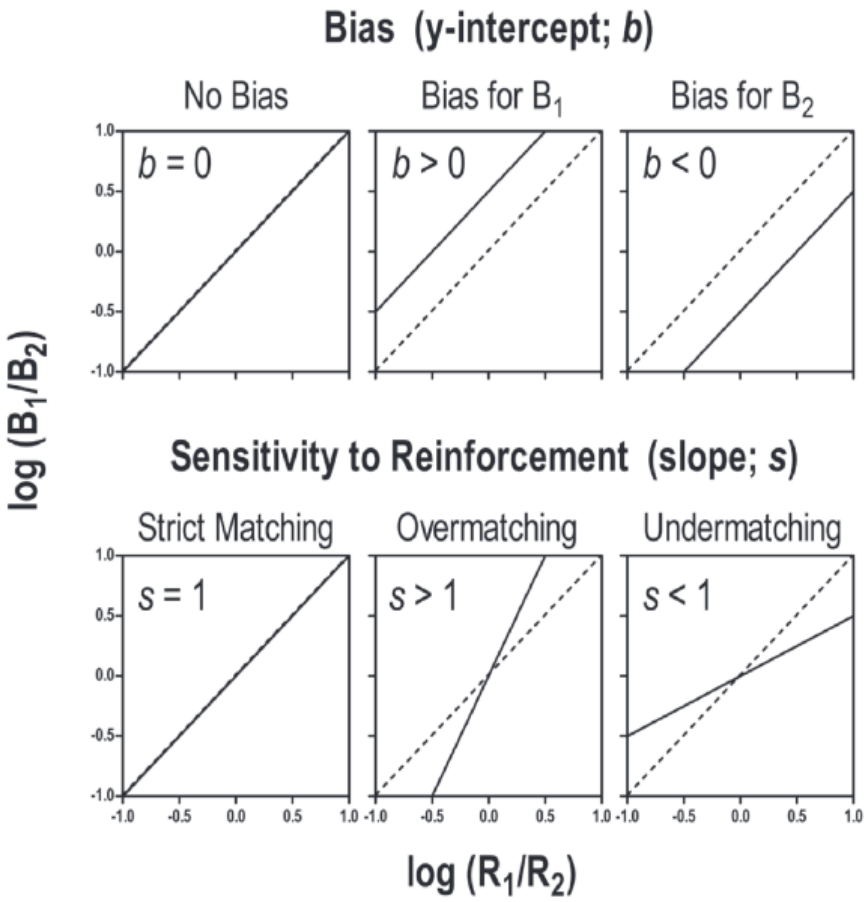

Figure 2.13. Variations in bias and sensitivity to reinforcement according to the generalized matching equation (taken from Reed & Kaplan, 2001, p. 17).

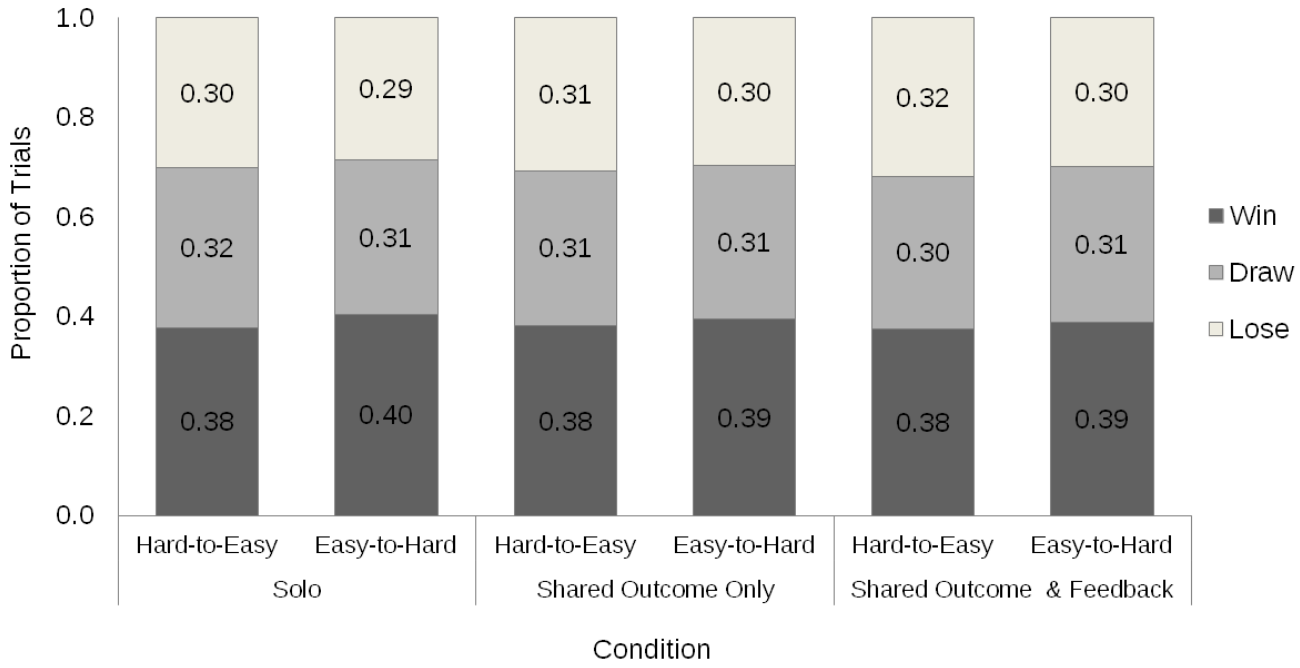

Figure 3.1. Mean proportion of overall correct and incorrect (i.e., draws and losses) attempts for participants in the tree experimental conditions as a function of the sequence of probability presentations.

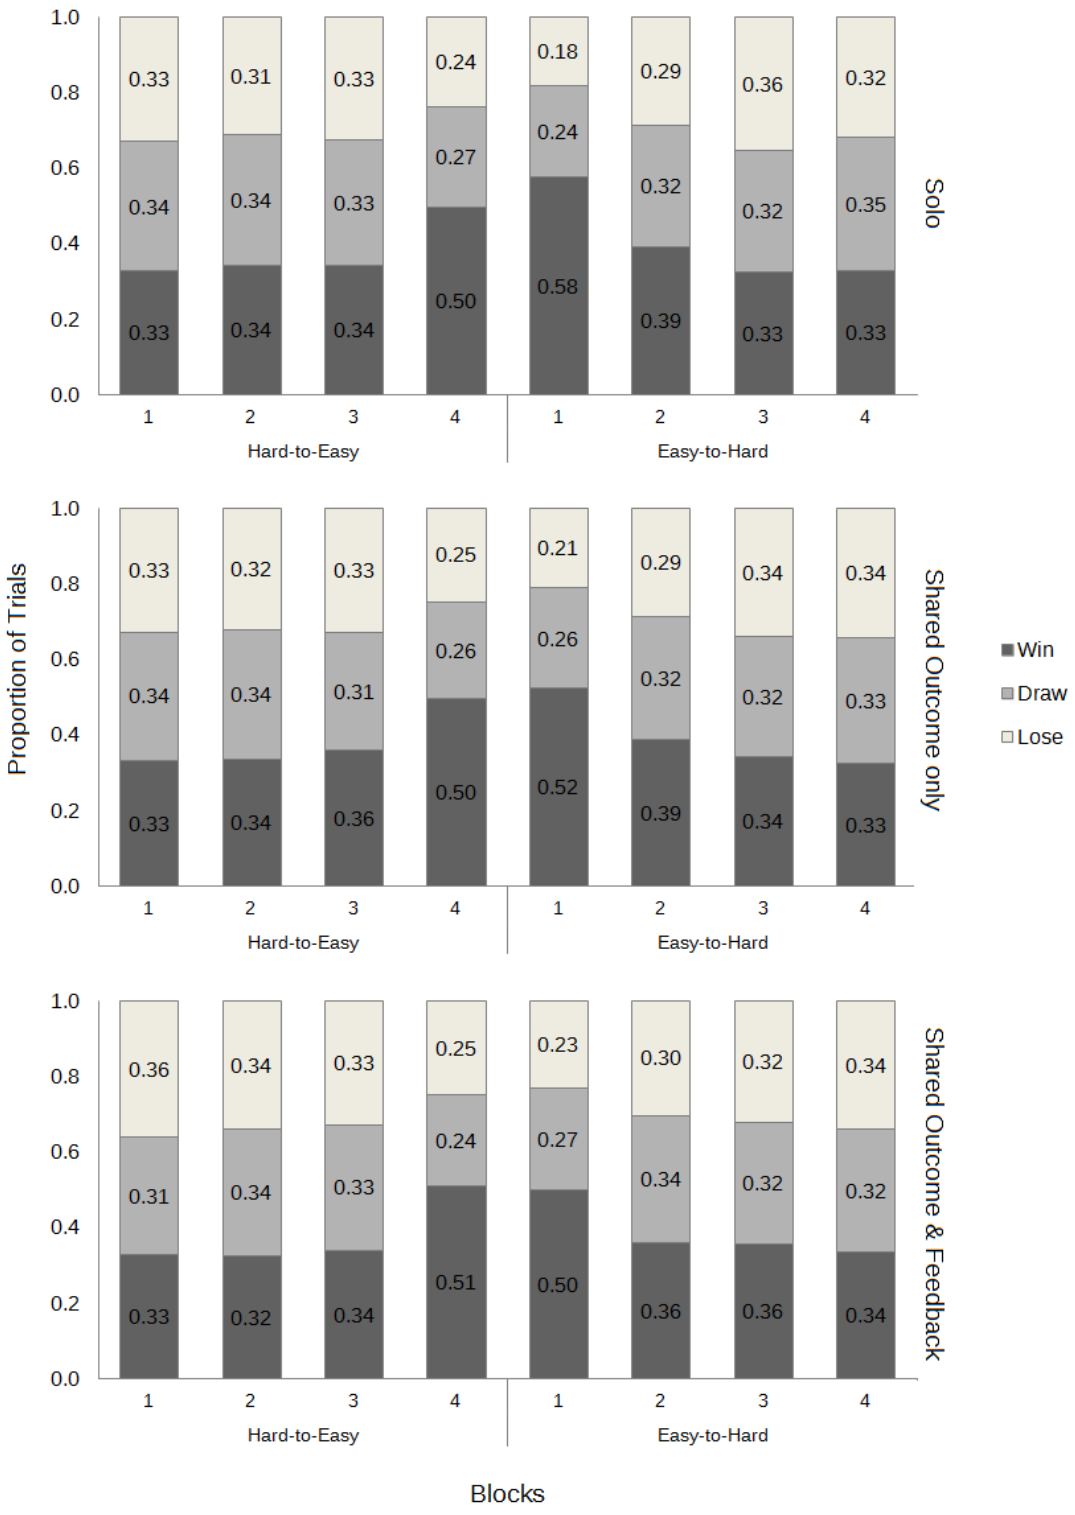

Figure 3.2. Proportion of correct and incorrect (i.e., draws and losses) attempts for participants in the three experimental conditions across blocks.

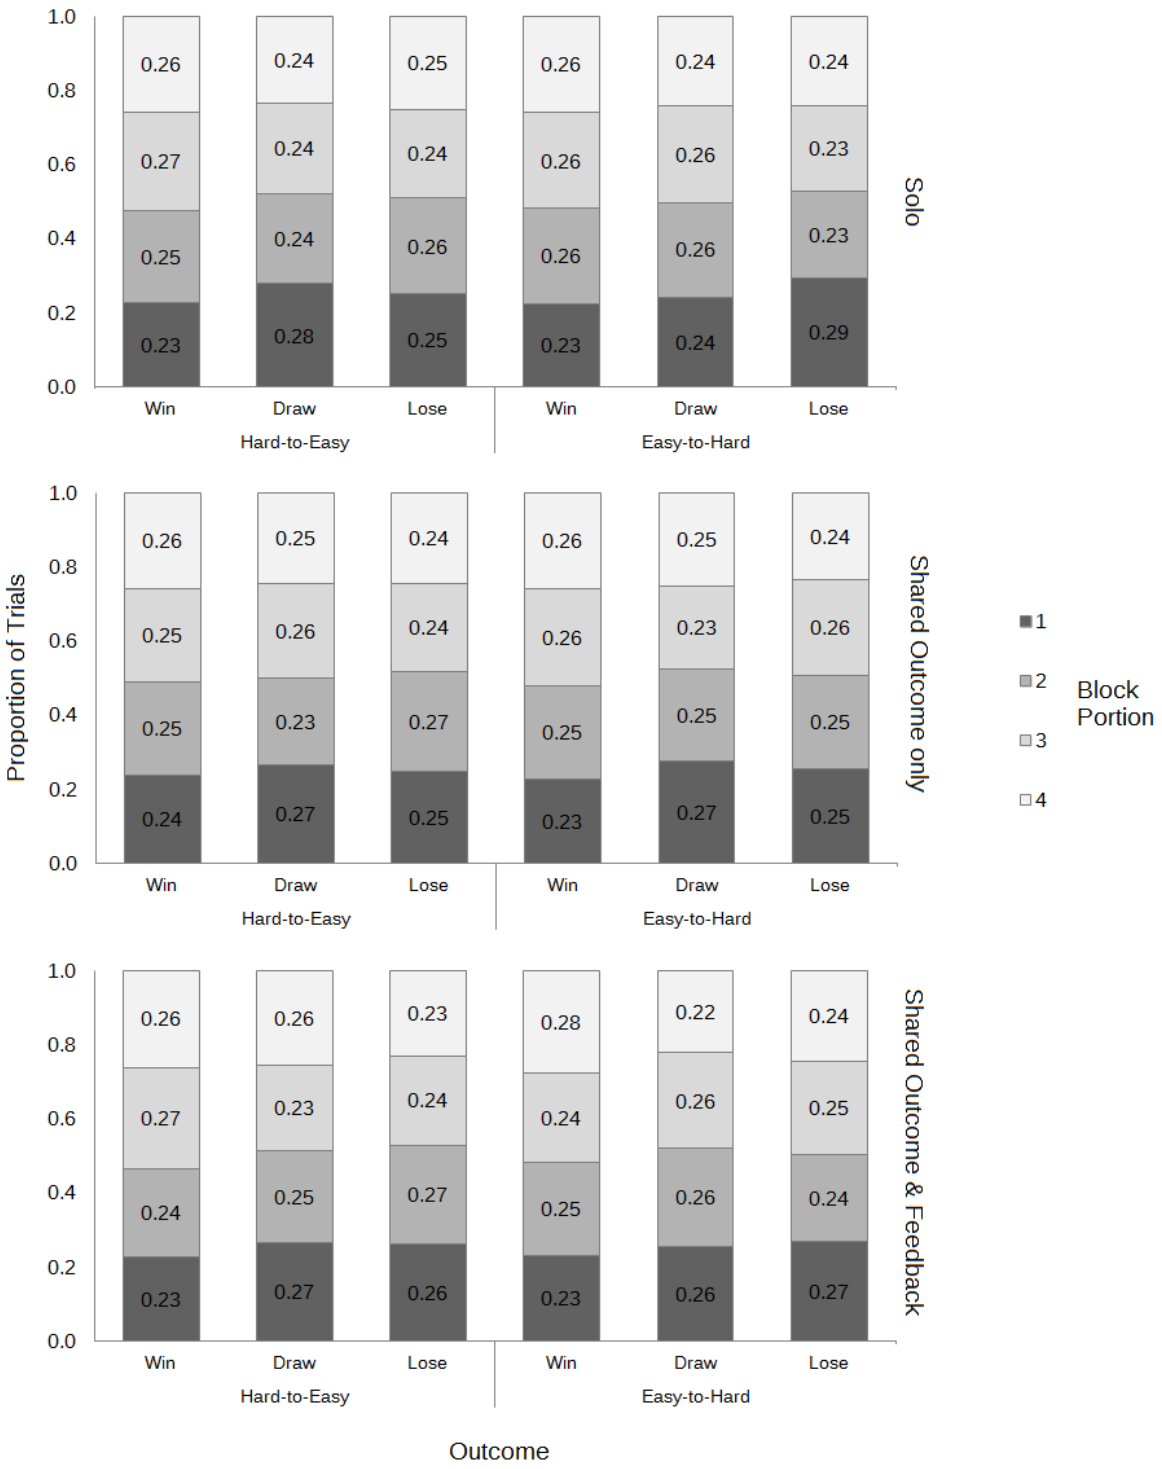

Figure 3.3. Proportion of correct and incorrect (i.e., draws and losses) attempts for participants in the three experimental conditions across block portions of 15 trials each.

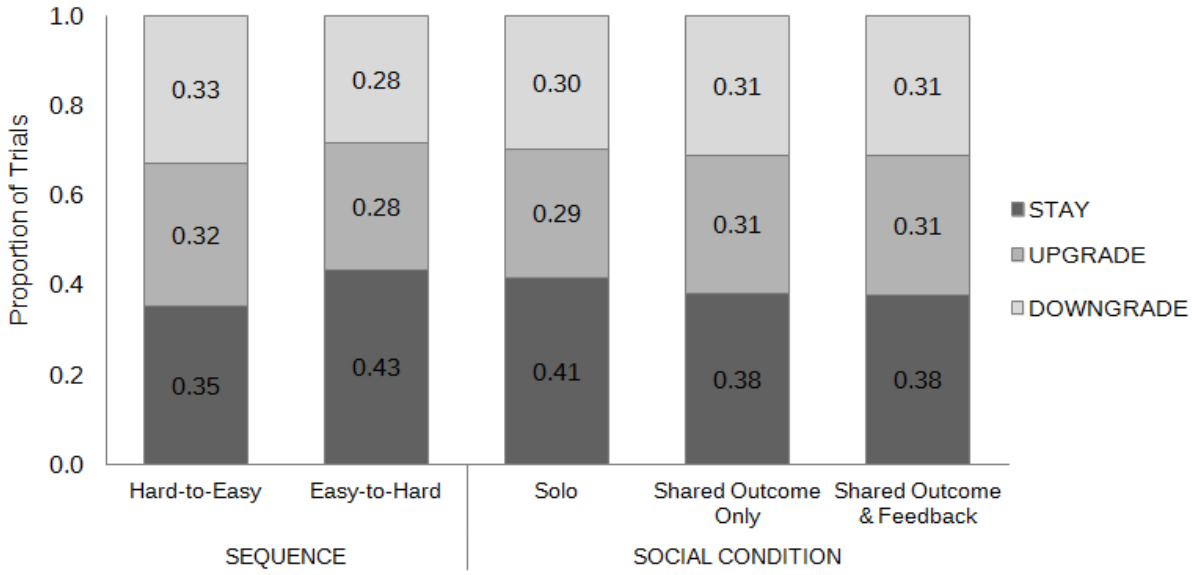

Figure 3.4. Overall strategy at trial *n + 1* as a function of the selection at trial *n* for participants in the sequence conditions and in the three experimental conditions.

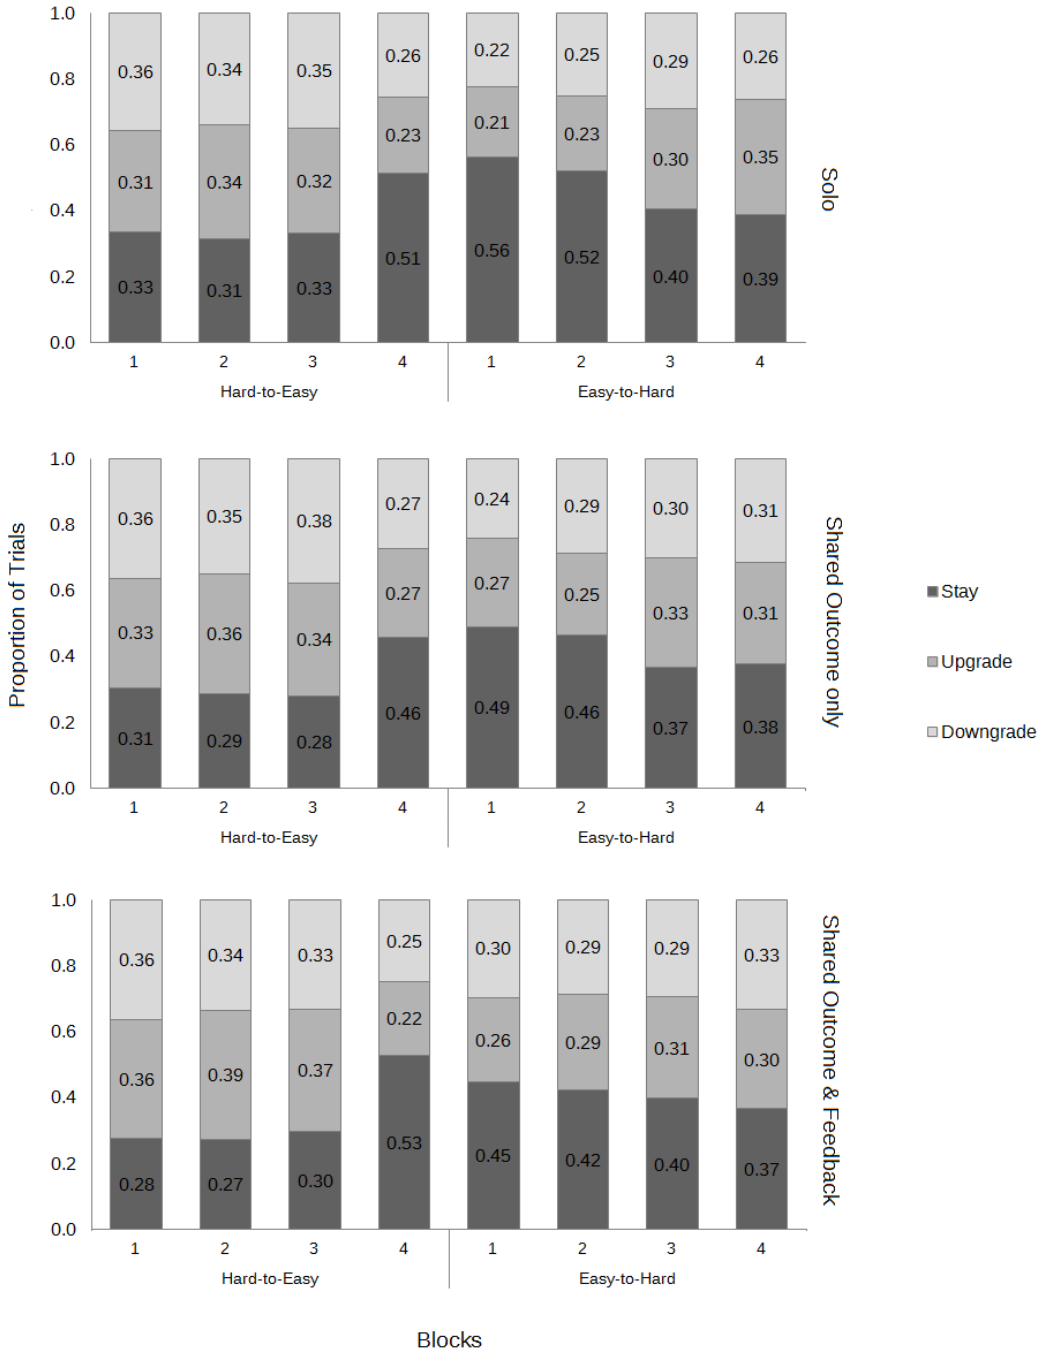

Figure 3.5. Strategy at trial *n + 1* as a function of the *selection* of trial *n* for participants in the three experimental conditions across blocks. The left panel corresponds to the decreasing difficulty (Hard-to-Easy) sequence, while the right panel corresponds to the strategy in the increasing difficulty (Easy-to-Hard) sequence.

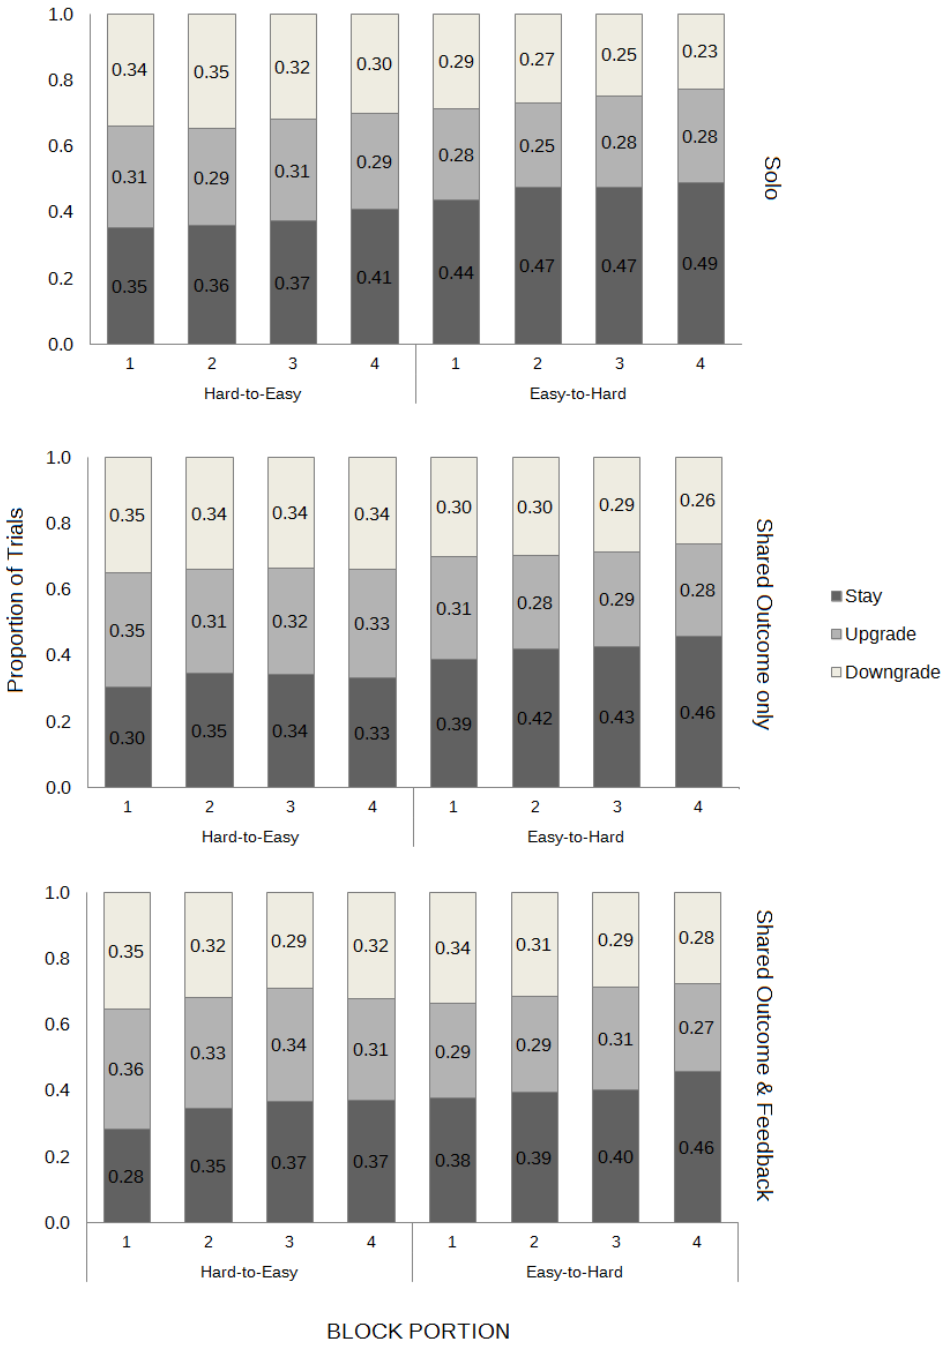

Figure 3.6. Strategy at trial *n + 1* as a function of the *selection* of trial *n* for participants in the three experimental conditions across block portions of 15 trials. The left panel corresponds to the Hard-to-Easy sequence, while the right panel corresponds to the strategy in the Easy-to-Hard sequence.

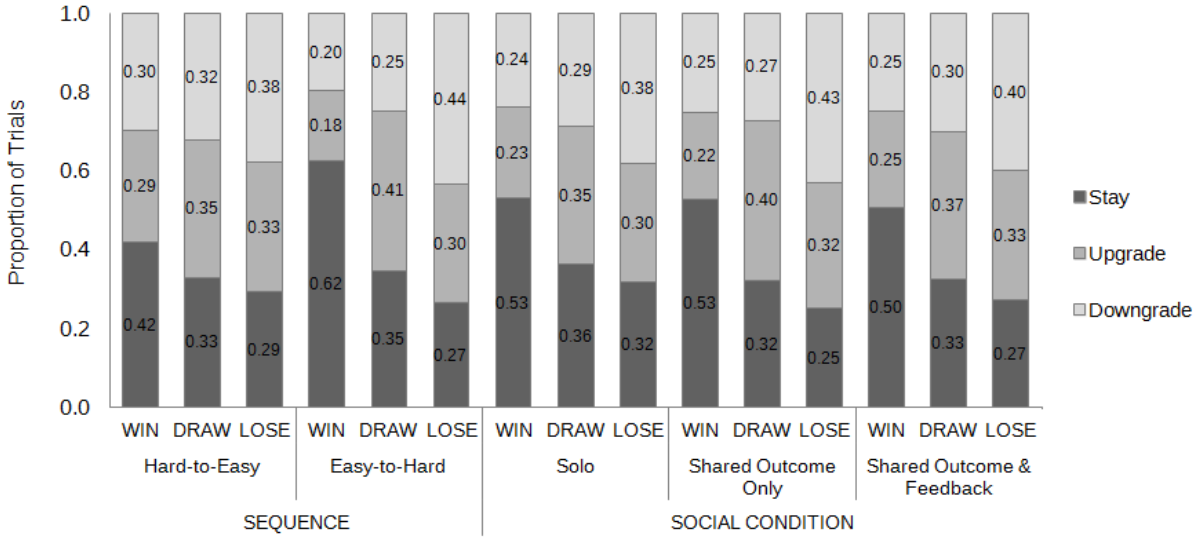

Figure 3.7. Overall strategy at trial *n + 1* as a function of the outcome of trial *n* for participants collapsed across the three social conditions and the two difficulty sequences.

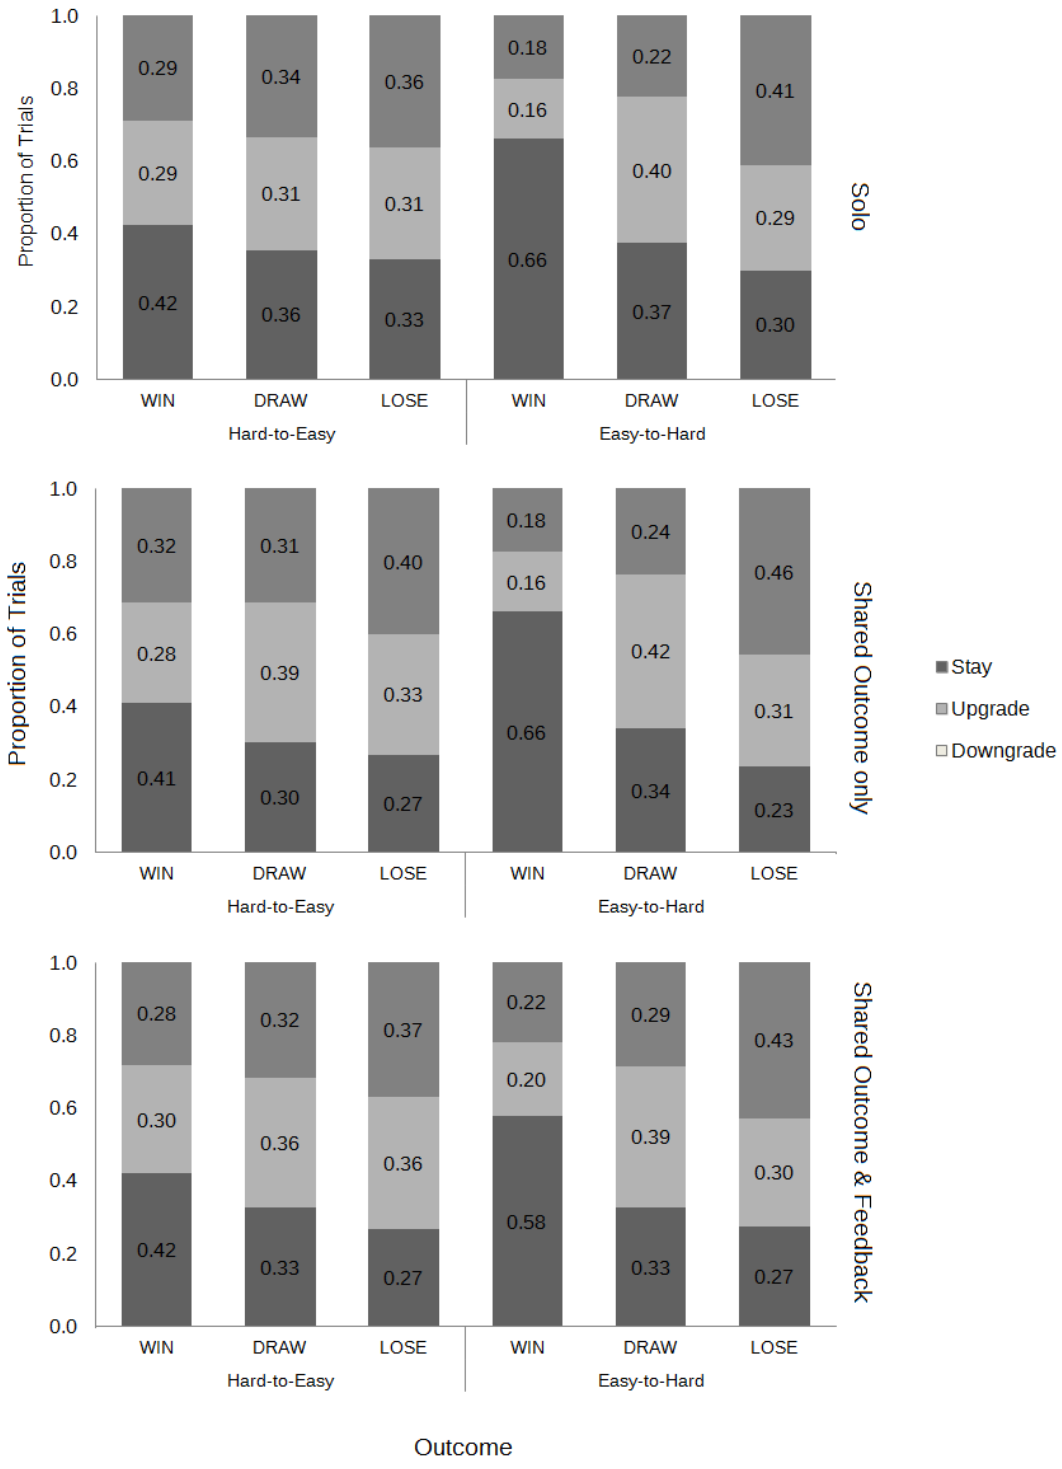

Figure 3.8. Strategy at trial *n + 1* as a function of the outcome of trial *n* for all participants.

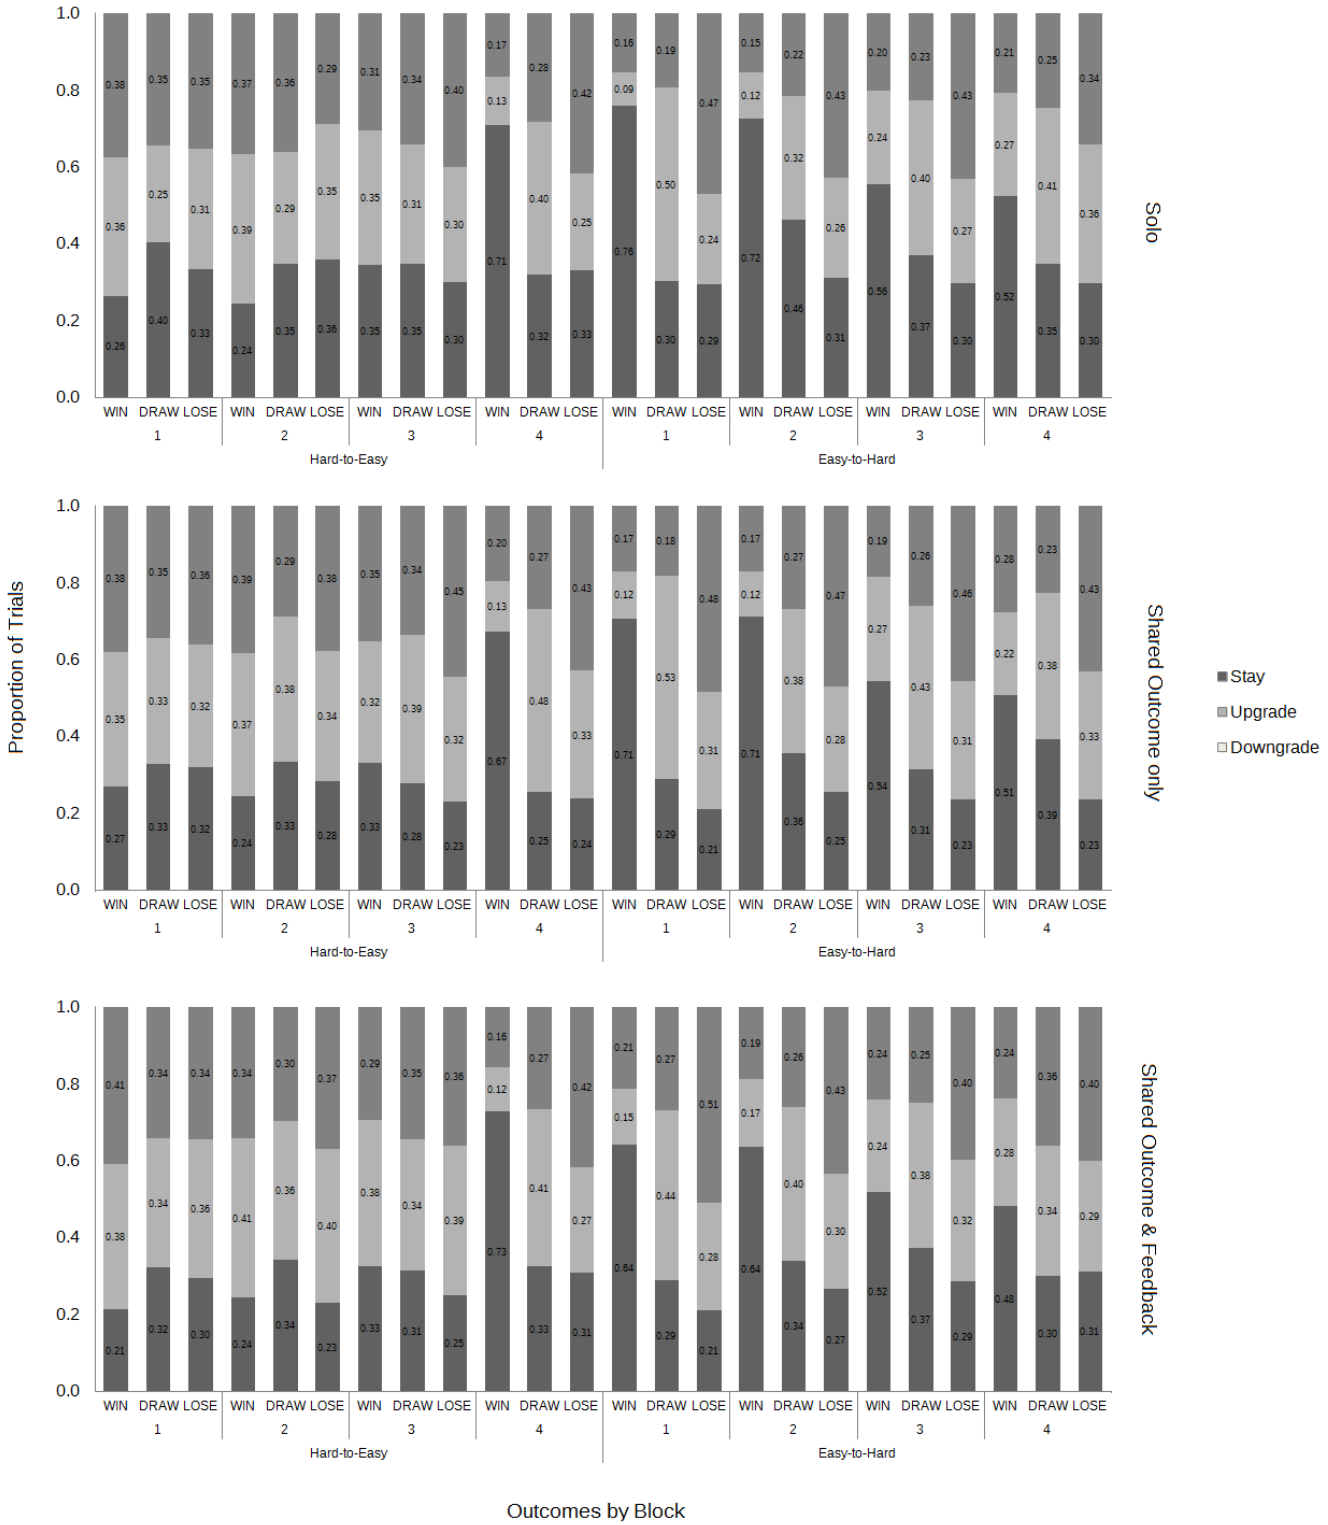

Figure 3.9. Strategy at trial *n + 1* as a function of the outcome of trial *n* across blocks for participants in all experimental conditions.

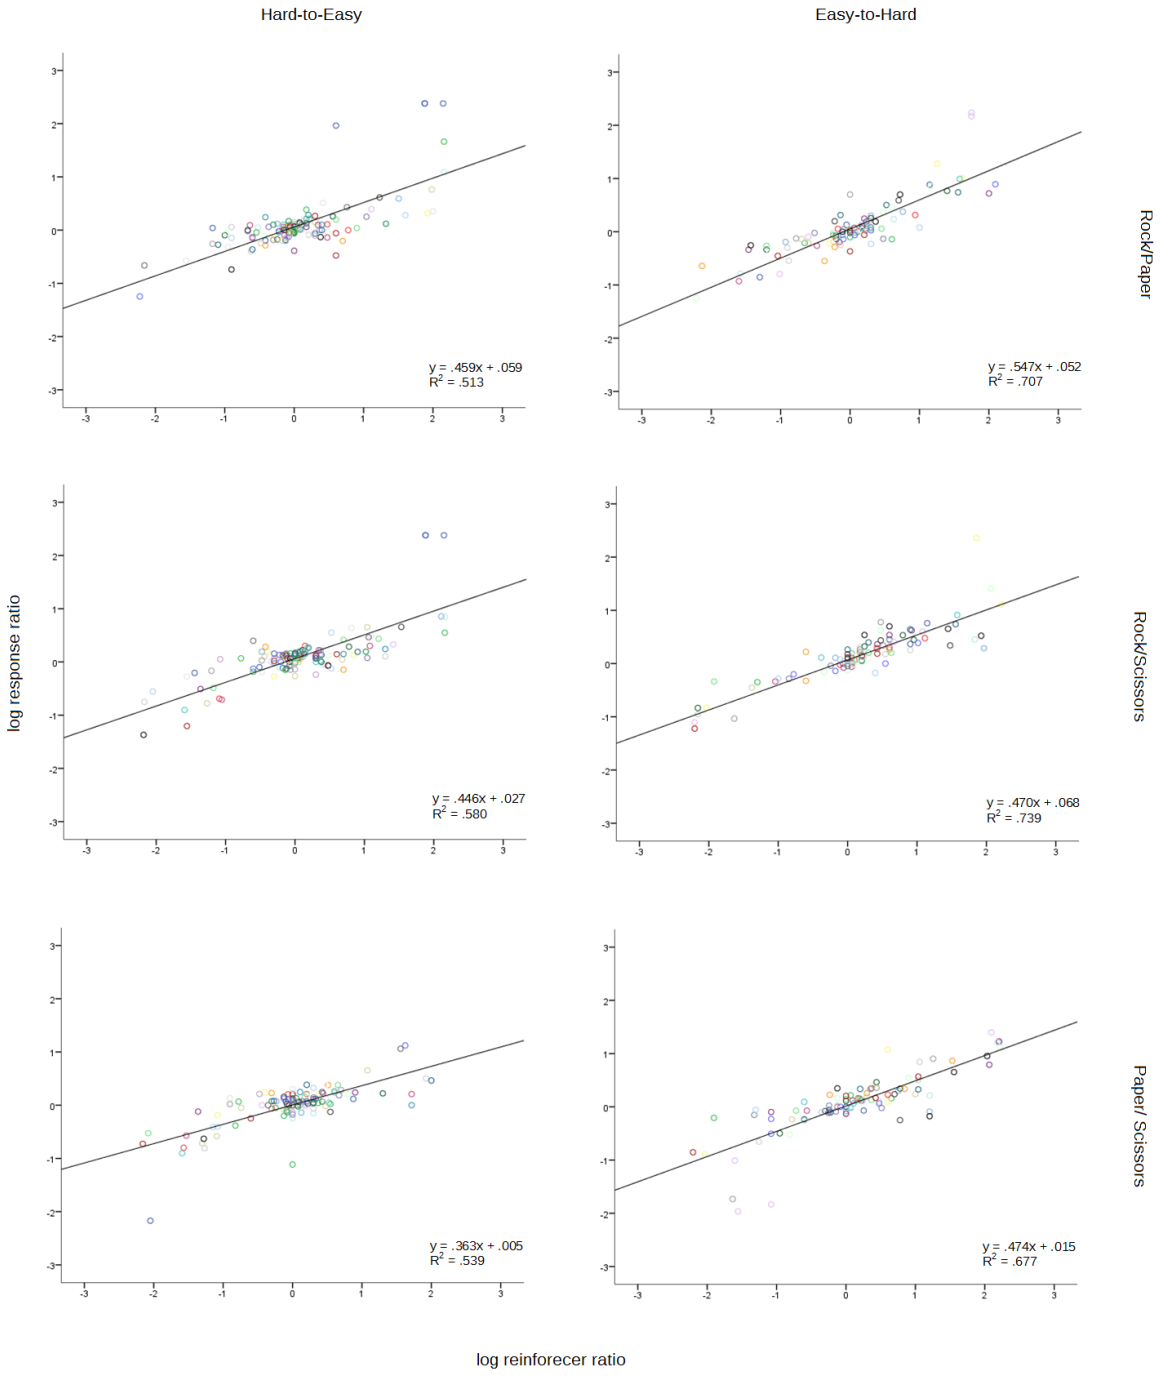

Figure 3.10. Log response ratios as a function of log reinforcement ratios for participants in the Solo condition. The left panel depicts results for the Hard-to-Easy sequence, while the right panel depicts results from the Easy-to-Hard sequence.

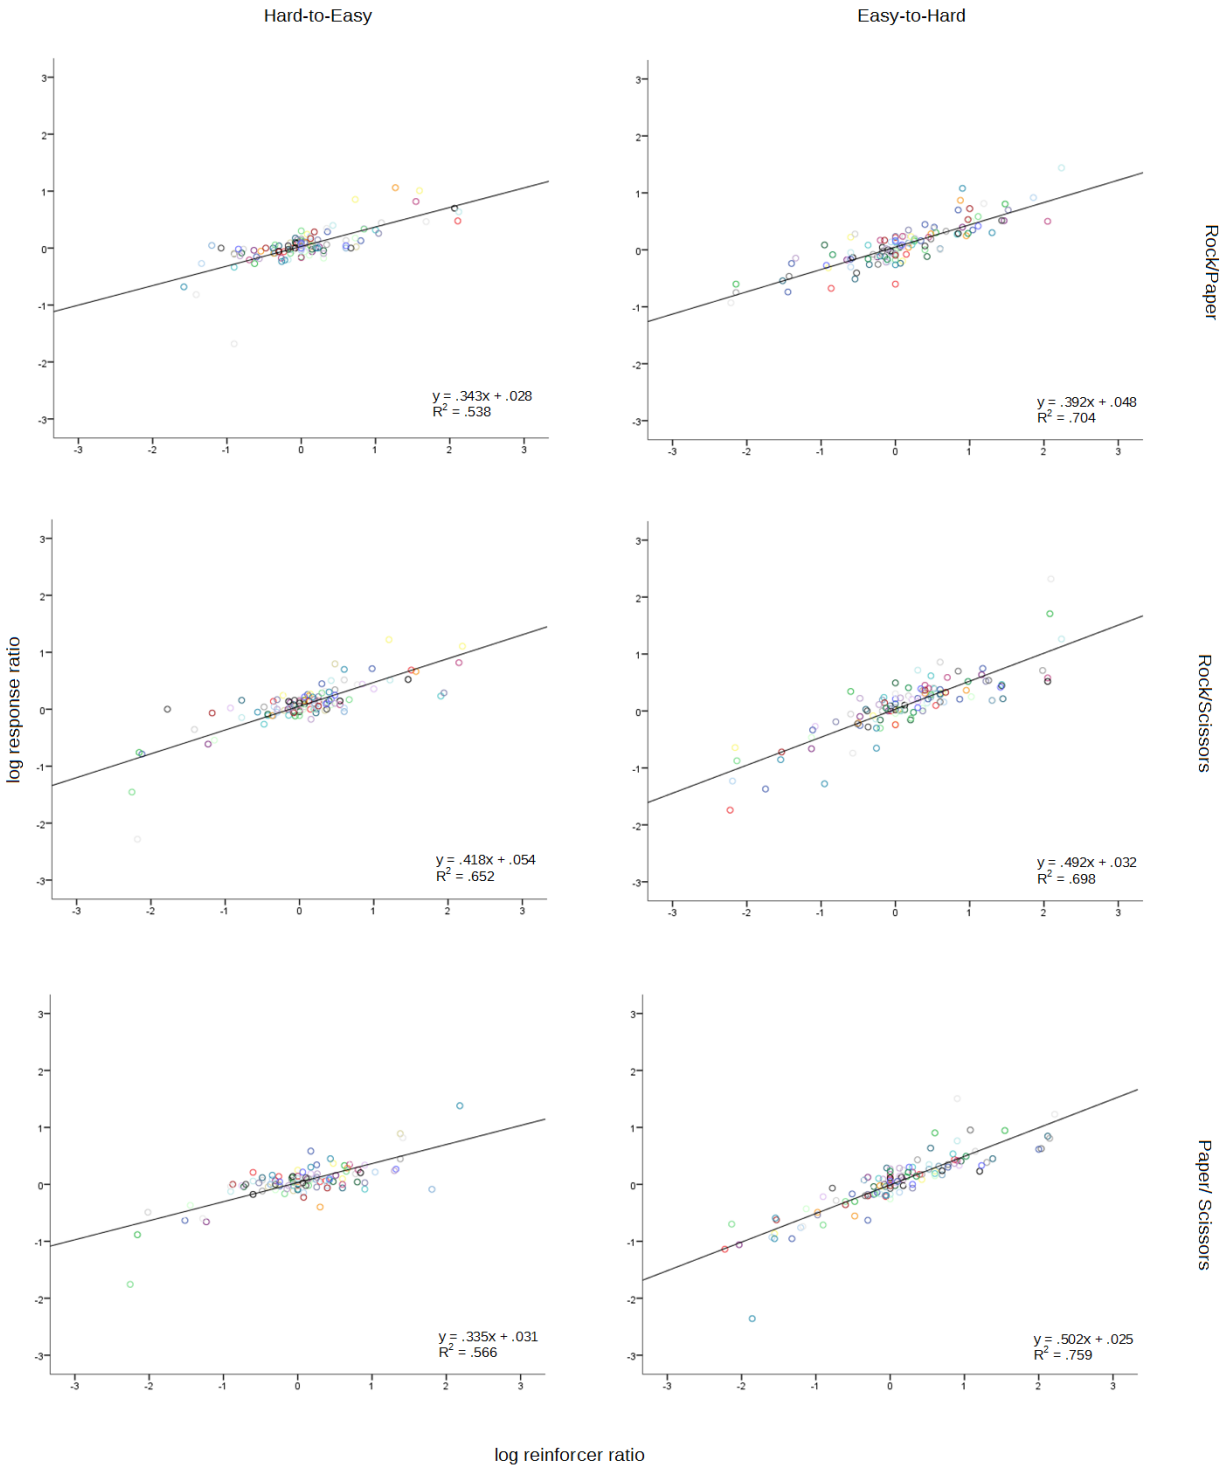

Figure 3.11. Log response ratios as a function of log reinforcement ratios for participants in the Shared Outcome Only condition. The left panel depicts results for the Hard-to-Easy sequence, while the right panel depicts results from the Easy-to-Hard sequence.

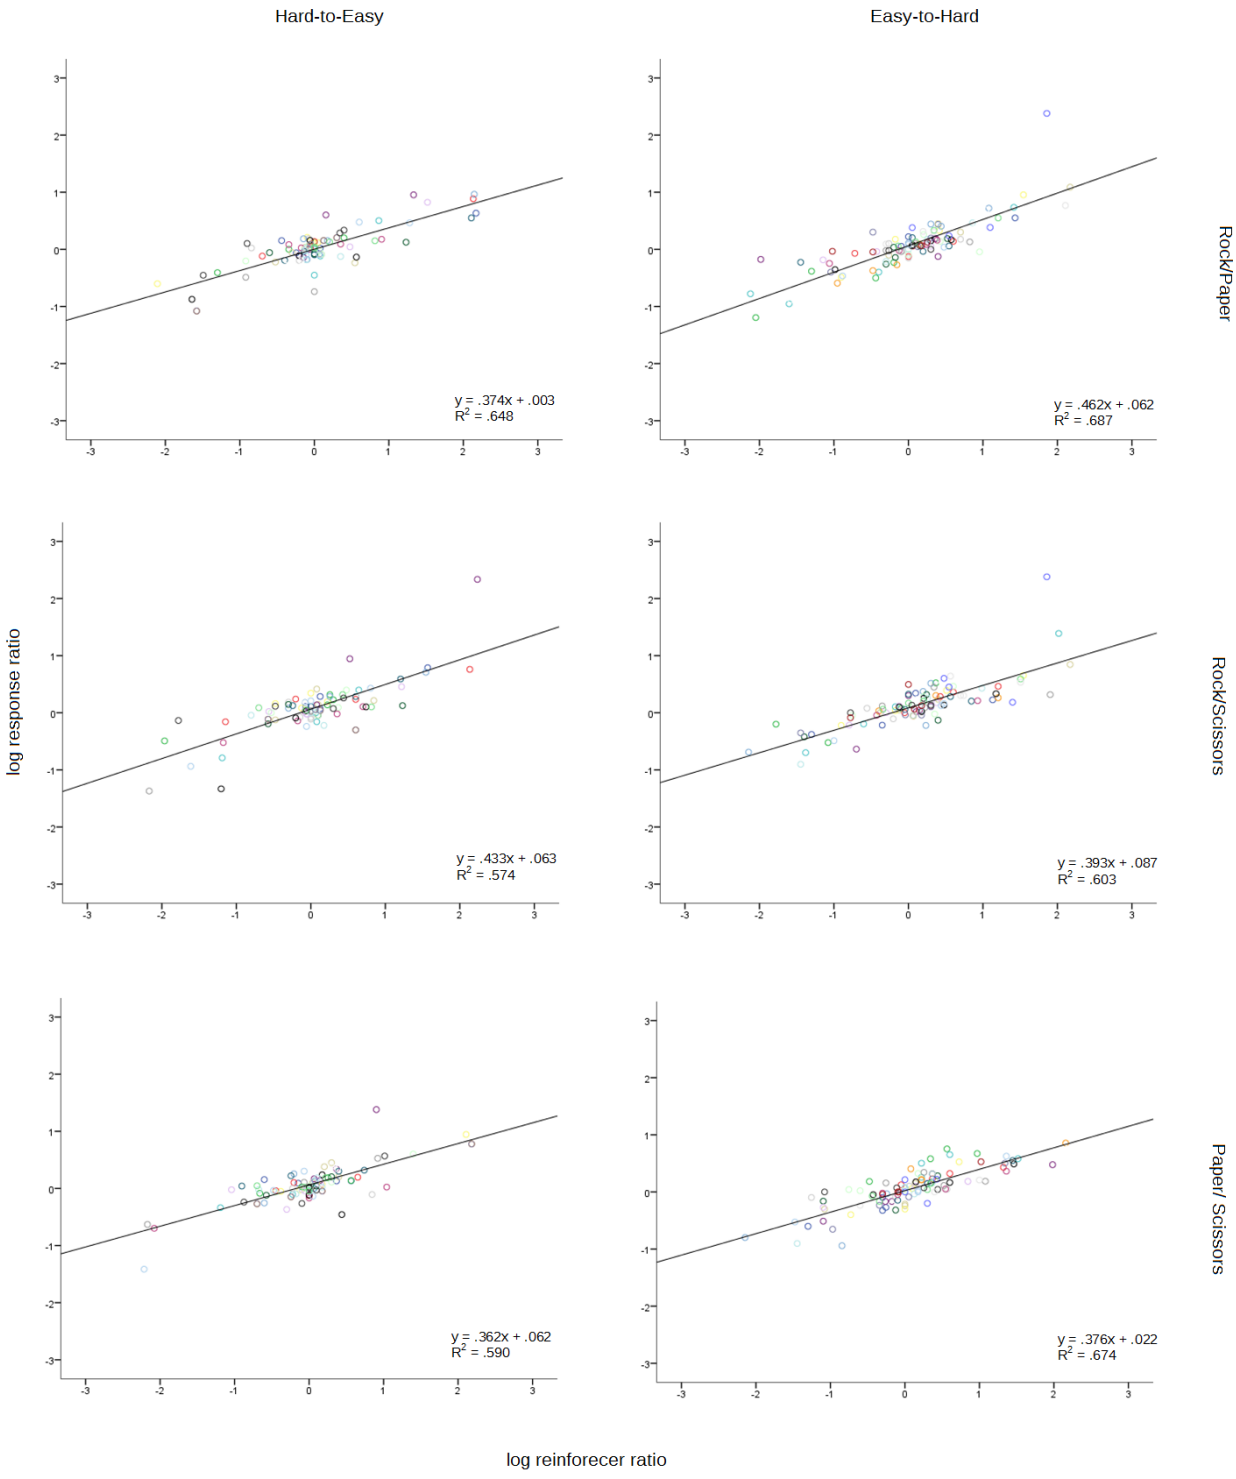

Figure 3.12. Log response ratios as a function of log reinforcement ratios for participants in the Shared Attention & Feedback condition. The left panel depicts results for the Hard-to-Easy sequence, while the right panel depicts results from the Easy-to-Hard sequence.

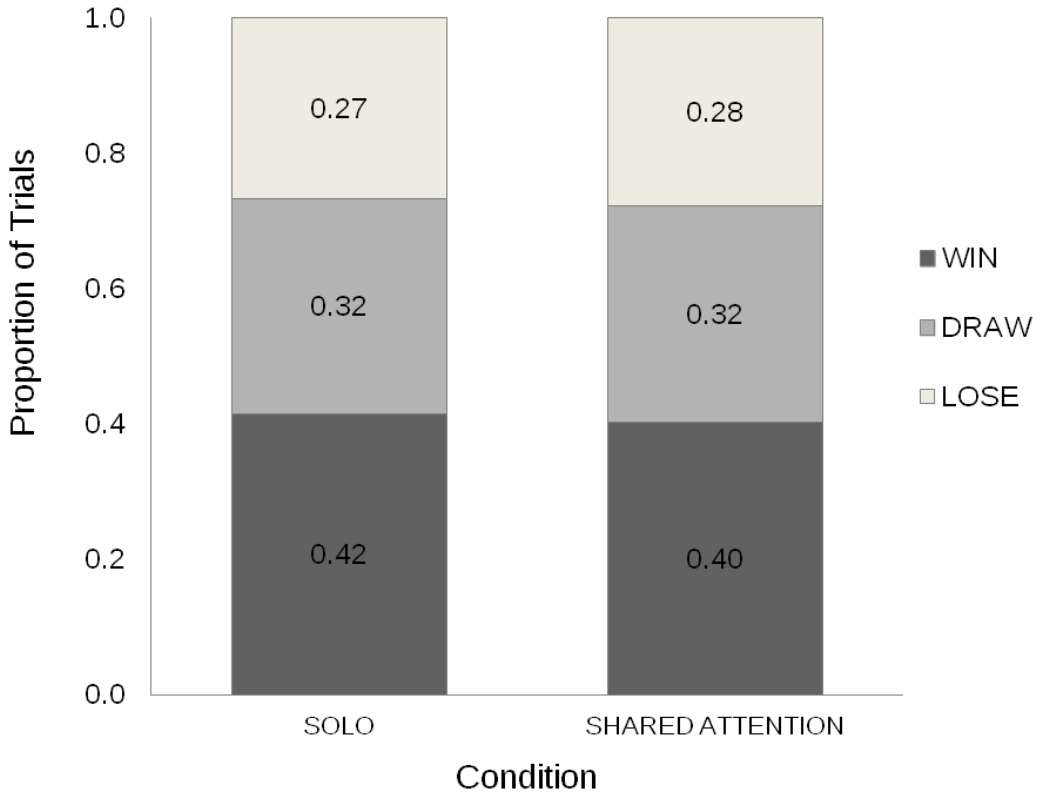

Figure 4.1. Mean proportion of overall correct (wins) and incorrect (i.e., draws and losses) attempts for participants in the two experimental conditions.

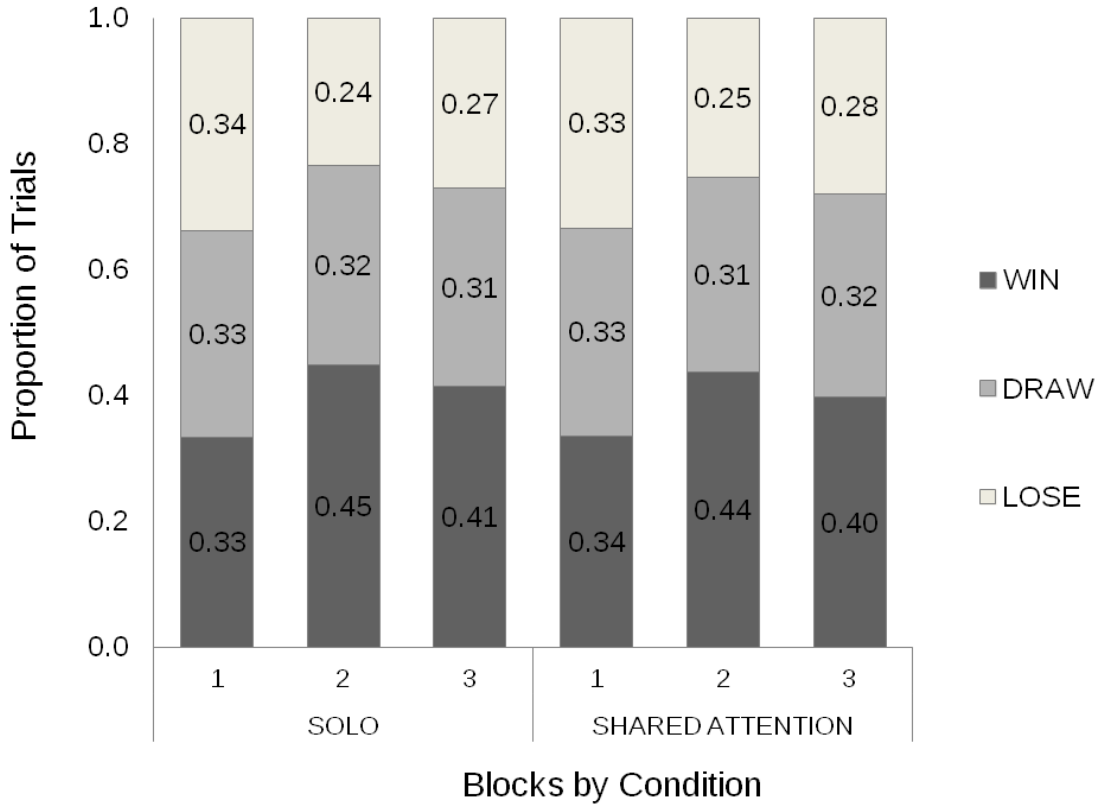

Figure 4.2. Mean proportion of overall correct and incorrect (i.e., draws and losses) attempts across blocks for participants in the experimental conditions, Solo (left panel) and Shared Attention (right panel).

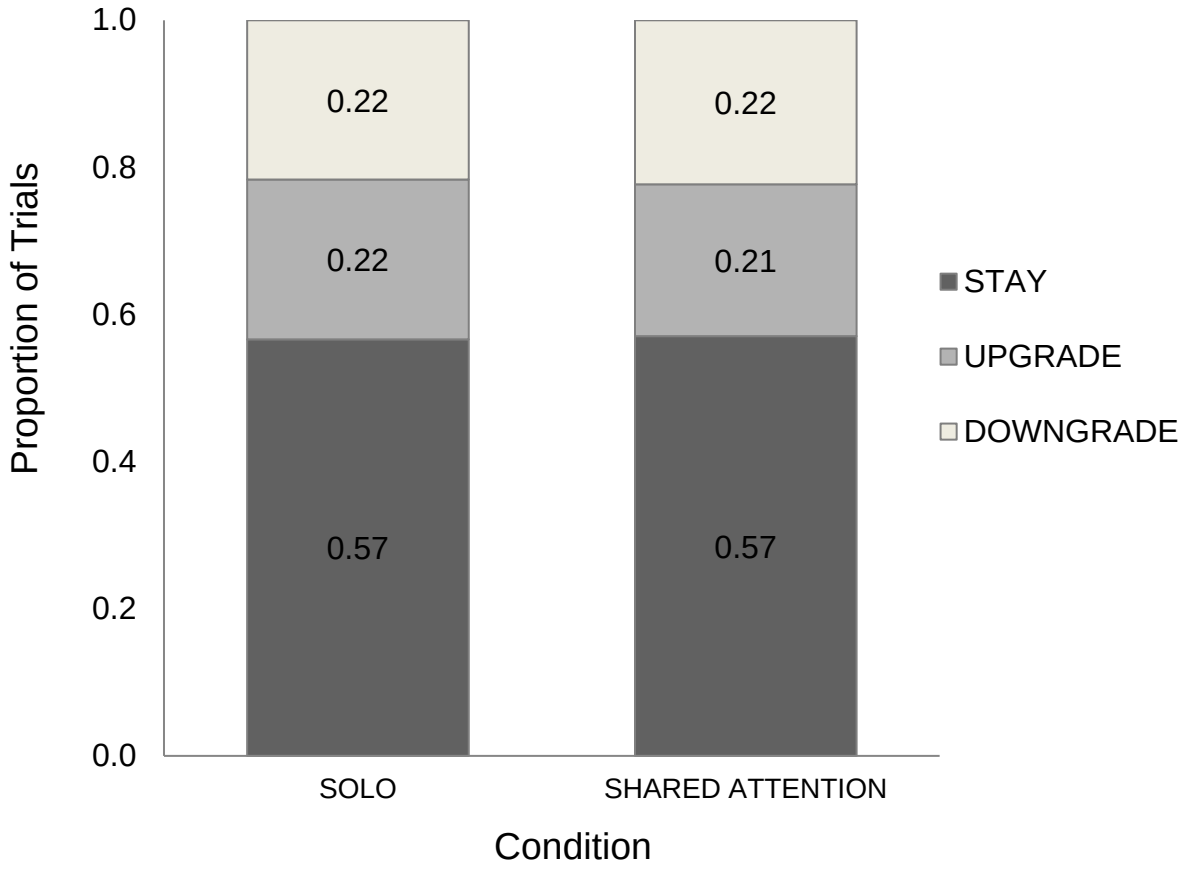

Figure 4.3. Overall strategy at trial *n + 1* as a function of the selection at trial *n* for participants in the Solo (left) and Shared Attention (right) conditions.

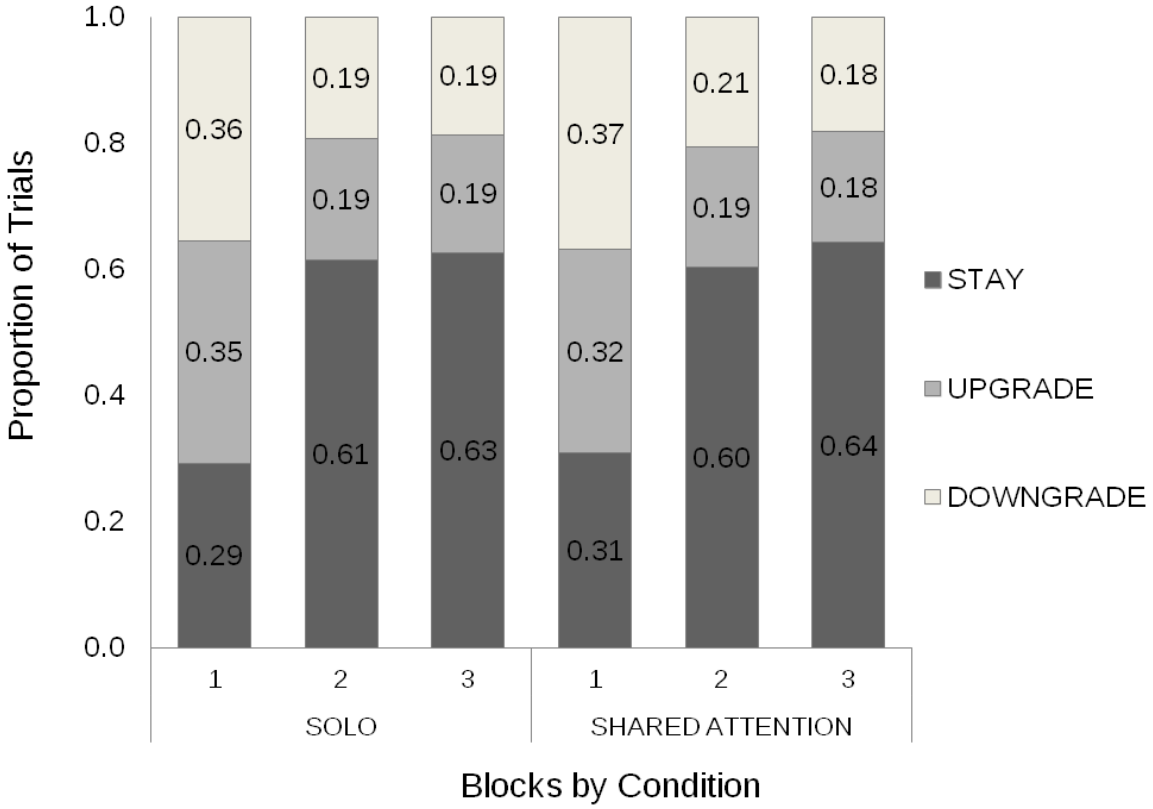

Figure 4.4. Overall strategy at trial *n + 1* as a function of the selection at trial *n* across blocks for participants in the Solo (left) and Shared Attention (right) conditions.

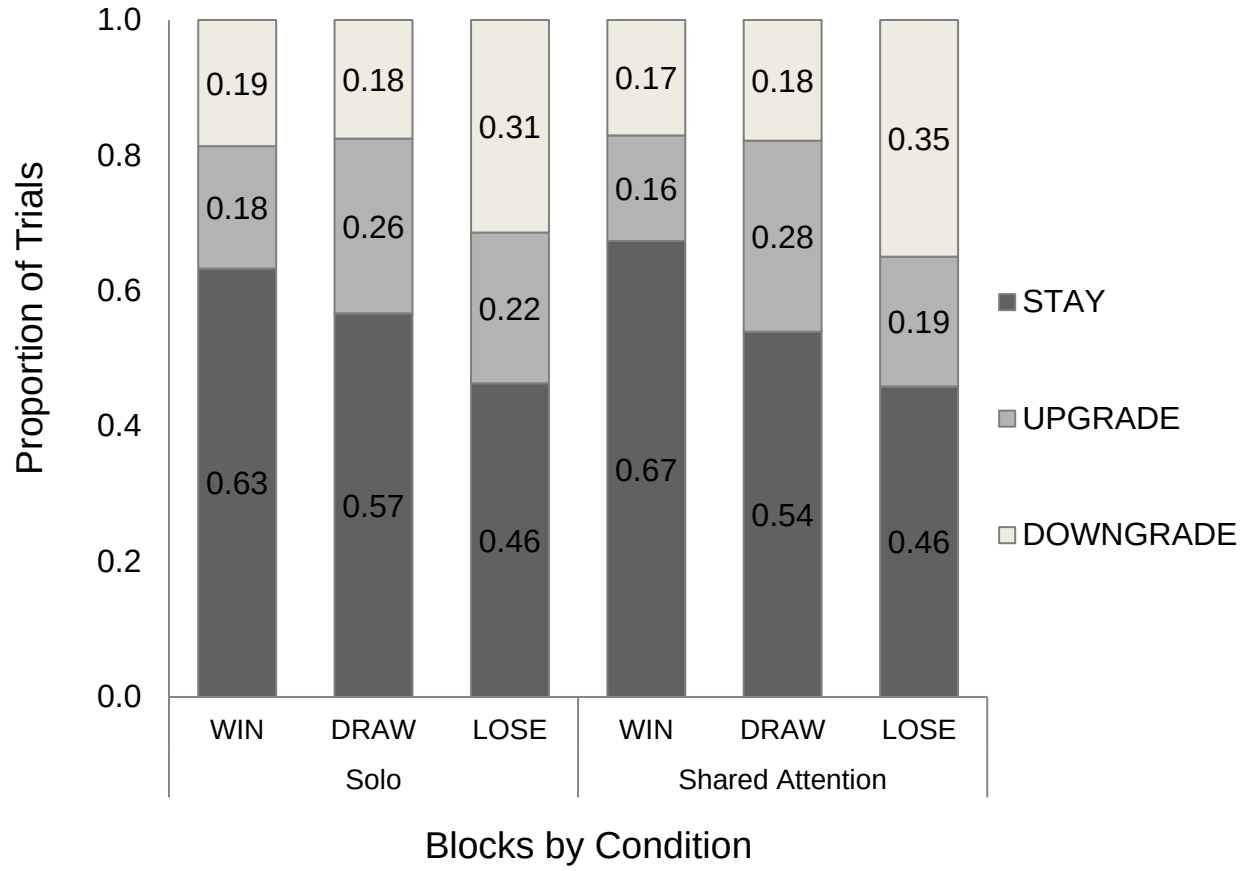

Figure 4.5. Overall strategy at trial *n + 1* as a function of the outcome of trial *n* for participants in the Solo (left) and Shared Attention (right) conditions.

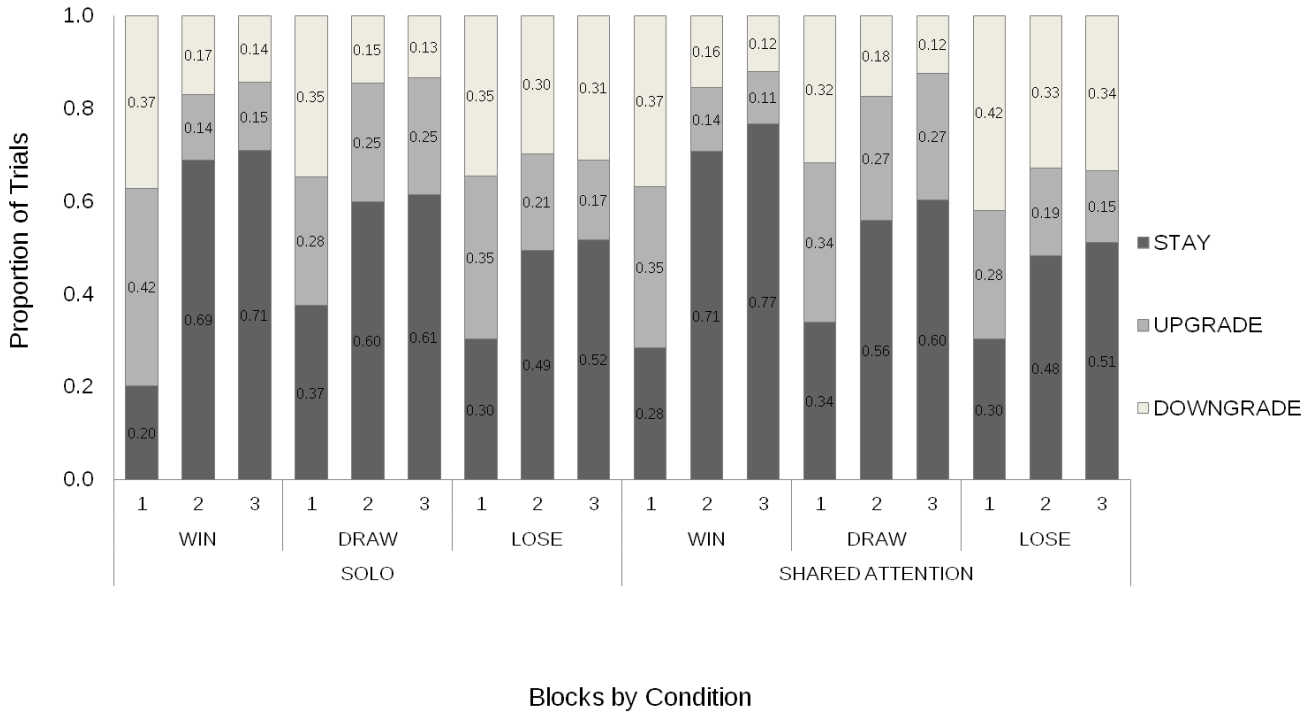

Figure 4.6. Strategy at trial *n + 1* as a function of the outcome of trial *n* for participants in the Solo (left) and Shared Attention (right) conditions across blocks.

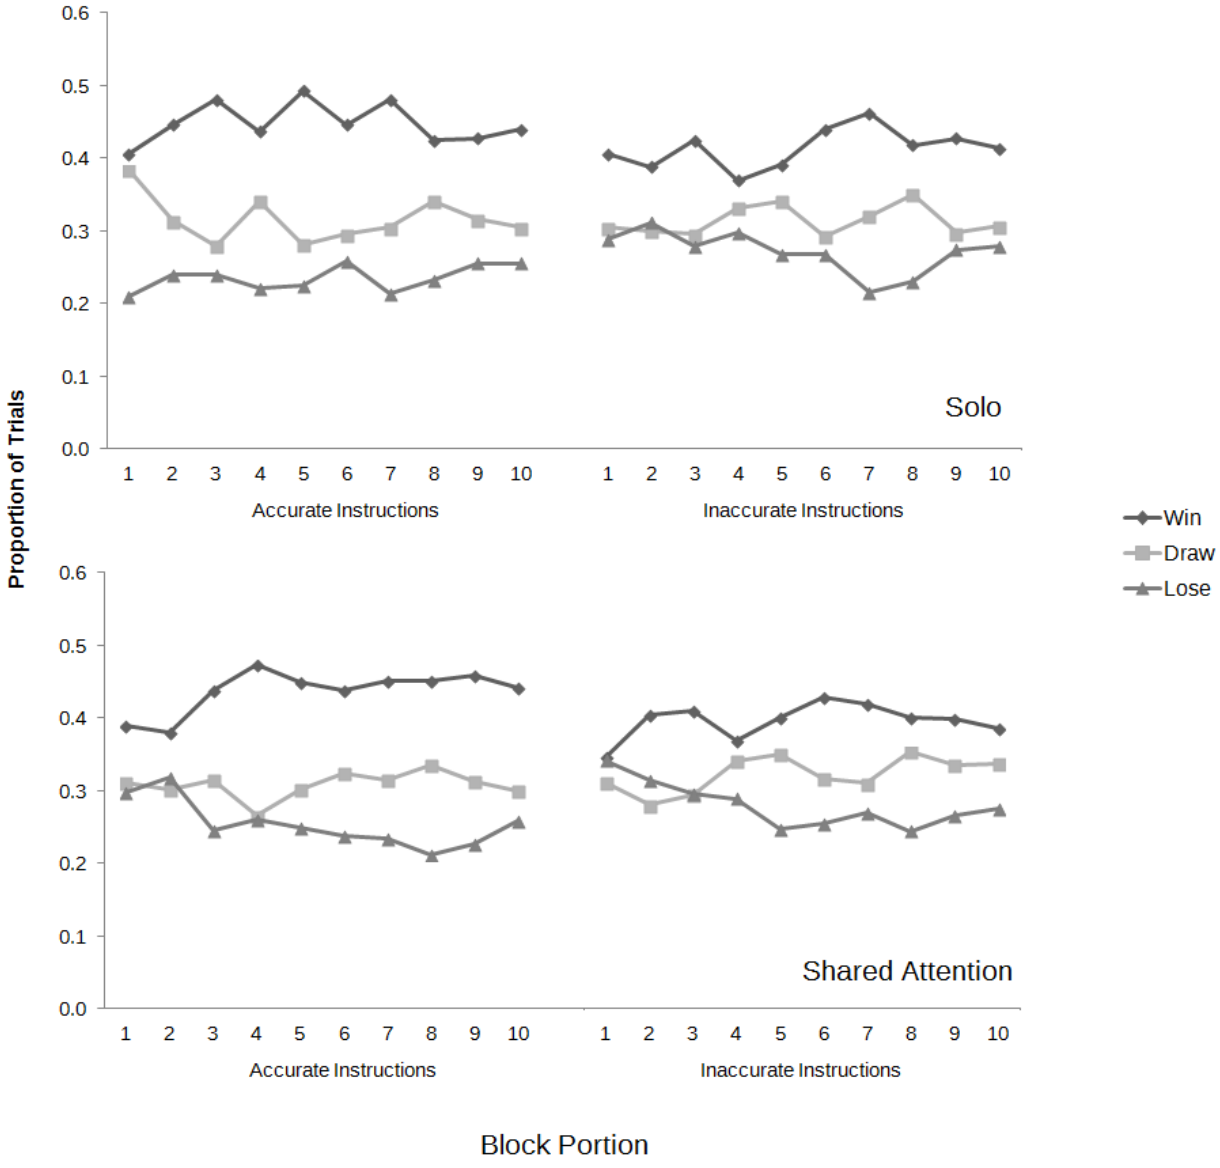

Figure 4.7. Proportions of wins, draws, and losses across successive block portions of 10 trials each in Block 2 (Accurate Instructions, left panel) and Block 3 (Inaccurate Instructions, right panel) for participants of the two conditions, Solo (upper panel), and Shared Attention (lower panel).

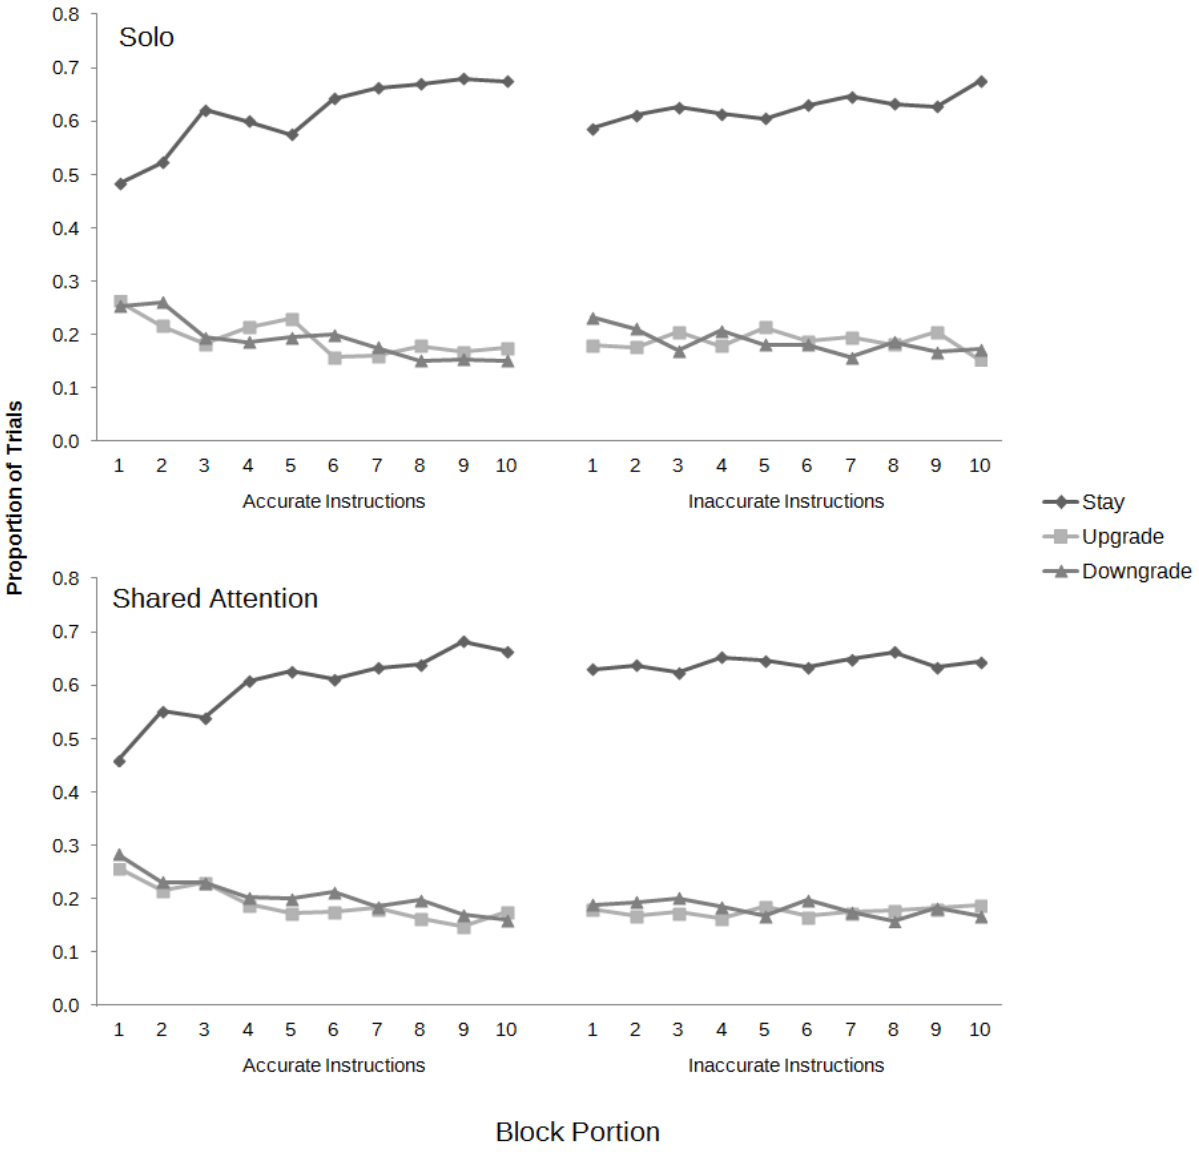

Figure 4.8. Proportions of strategies (i.e., staying, upgrading, downgrading) at trial *n + 1* as a function of trial *n* across successive block portions of 10 trials each in Block 2 (Accurate Instructions, left panel) and Block 3 (Inaccurate Instructions, right panel) for participants of the two conditions, Solo (upper panel), and Shared Attention (lower panel).

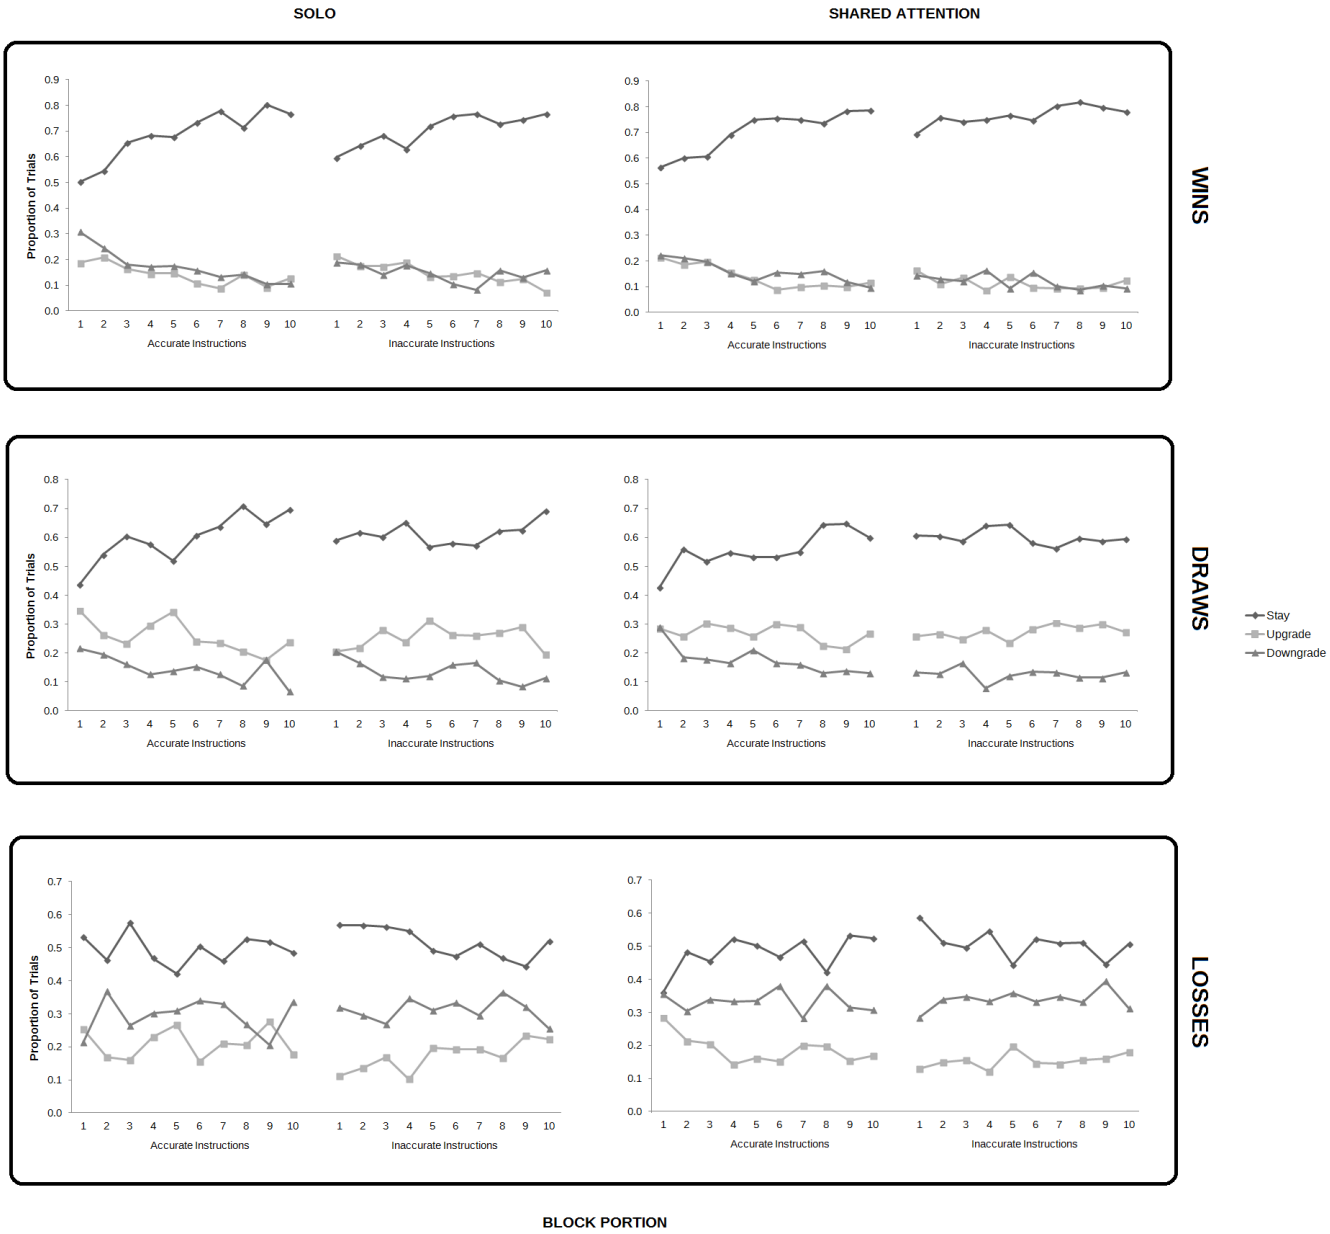

Figure 4.9. Proportions of strategies (i.e., staying, upgrading, downgrading) at trial *n + 1* as a function of the outcome of trial *n* (i.e., win, draw, or lose) across successive block portions of 10 trials each in Block 2 (Accurate Instructions) and Block 3 (Inaccurate Instructions) for participants of the two conditions, Solo (left panels) and Shared Attention (right panels).

Vita

Andres Garcia-Penagos was born in Bogota, Colombia. He attended the National University of Colombia, in Bogota, obtaining his undergraduate degree in Psychology. After a brief stint as a school psychologist in a couple of private schools in Bogota, he came to the United States to complete his graduate education. He obtained a M.Sc. degree in Psychology from West Virginia University, with an emphasis in Behavior Analysis. There he conducted basic operant research in the areas of learning and stimulus control with pigeons and rats. He went then to the University of Tennessee in 2012 where he continued doing research with pigeons and then joined the Group Attention and Behavior Lab, doing research at the intersection of learning theory, decision making, and social psychology, focusing on social influences in human choice behavior. In addition to his experimental interest in human and non-human social behavior, he also has a profound interest in the history, theory and philosophy of psychology and the philosophy of mind, particularly in anti-representationalism as exemplified in modern radical behaviorism and enactive cognition. He is currently at Delta State University as an Assistant Professor of Psychology.