Probability Antimatching

by

Peter DiBerardino

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Arts
in
Psychology

Waterloo, Ontario, Canada, 2023

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

This thesis contains data and discussion previously presented in my undergraduate thesis. This content, mainly in Study 1 and Appendix A, are included to provide important theoretical and empirical context for the novel work produced for this thesis.

Abstract

We present a conceptual inversion of probability matching called "probability antimatching." Where probability matching describes a decision strategy of stimulus pursuit, probability antimatching describes an analogous decision strategy of stimulus avoidance. We present three behavioural studies where participants played a computer game of hide-and-seek. Participants played hide-and-seek against a simulated computer opponent that selected rooms for hiding/seeking according to a given probability distribution. Seeking trials replicate traditional probability matching. Hiding trials demonstrate probability antimatching. In Study 1, we formally present our methodology of expressing participant seeking and hiding behaviour as a linear combination of Euclidean vectors. Participant seeking strategies, $\vec{s}$, are well-represented by a linear combination of the optimal maximizing strategy, $\vec{x}$, and the probability matching strategy, $\vec{m}$. Participant hiding strategies, $\vec{h}$, are equally well-represented by a linear combination of the optimal minimizing strategy, $\vec{n}$, and the probability antimatching strategy, $\vec{a}$. We define $\vec{a}$ as a vector reflection of $\vec{m}$ over the uniform distribution vector, $\vec{u}$. This operation is denoted $\vec{a} = refl_{\vec{u}}(\vec{m}) = 2\vec{u} - \vec{m}$. In Study 2, we replicate the findings of Study 1 using data collected online. In Study 3, we demonstrate that our conceptualization of probability antimatching extends to probability distributions that have non-unique optimal hiding/seeking strategies and distributions that have invalid reflections (that result in negative probability values). Across our three studies, we find that hiding/seeking strategies are influenced by the number of rooms presented during hide-and-seek, corresponding to the dimensionality of the underlying probability distributions. However, the direction of this effect fails to replicate across our studies.

Acknowledgements

I cannot overstate my gratitude to my supervisor and mentor Dr. Britt Anderson for his knowledge, wisdom, and support. My current and future successes can be attributed in large part to your influence.

Thank you to my past and present lab mates, collaborators, and readers, for your guidance and hours of great discussion.

To my wife, family, and friends, thank you for your love, support, and patience.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Probability matching is a phenomenon found in repeated choice experiments where participants predict events proportionately to the probability of each event occurring, rather than optimize by strictly selecting the most likely event. For example, consider two lights L and R. Each light appears on a given trial with probability P(L) and P(R), respectively. An individual predicting whether a light will appear on their left or right will tend to predict the most probable light, say L, P(L) of the time [36]. Under such conditions, the optimal strategy to maximize correct prediction is one where light L is predicted 100% of the time, P(guess L) = 1:

$$\text{P(correct)} = \text{P(guess L)} * \text{P(L)} + \text{P(guess R)} * \text{P(R)} = 1 * 0.7 + 0 * 0.3 = 0.7$$

rather than P(L) of the time, P(guess L) = P(L):

$$\text{P(correct)} = 0.7 * 0.7 + 0.3 * 0.3 = 0.58$$

If the underlying probabilities are known, why do people employ a matching strategy rather than a strategy the maximizes correct choices? The answer to this question is unclear. The leading perspective attributes probability matching to cognitive biases and the use heuristics [37, 28, 15, 10]. For example, [10] argues it may be that probability matching is a consequence of the well-documented human misperception of randomness, and our tendency to search for patterns. If a series of events occurred in a known deterministic pattern, the optimal strategy would be to follow that pattern, resulting in a 'matching' behaviour in the aggregate. If some underlying pattern could be found, it would no longer be optimal to maximize. The only way to have a perfect or near-perfect strategy, would be to use a strategy that looks like probability matching in aggregate. While associations between pattern searching and probability matching behaviours are documented [10, 35], it is not clear why pattern searching should *cause* probability matching or vice versa [21].

In predator/prey environmental contexts involving a group of agents, probability matching is consistent with the strategy for optimal access to food [36]. Consider a group of ducks on a pond. If a researcher throws bread into one side of the pond at double the frequency as the other side, one would expect 2/3 of the ducks to congregate on the higher frequency side, and 1/3 on the other. This behaviour is optimal in that it establishes a Nash equilibrium, as each singular duck will be unable to increase its food-to-competition ratio once such a frequency matched distribution is established [36]. Aside from a few aggressive ducks skewing the availability of food, this optimal matching behaviour in ducks has been observed [14]. To preserve this evolutionarily stable strategy [36], each new duck entering the pond should elect to feed at each side L or R, with probability P(L) and P(R), respectively. Of course, if a duck were the only duck in the pond, the optimal strategy would be to strictly maximize, a scenario closer to traditional probability matching experiments.

In real-world scenarios with richer information availability and social consequences, probability matching is less-obviously non-optimal. Consider the socially-embedded probability matching scenario presented in [35]: you are watching a football game with your friends, and join them in a guessing game: will the home team execute a running play or passing play the next time they have the ball? Despite never watching a football game before, you identify that the home team executes a passing play 70% of the time, and a running play 30% of the time. This scenario is presented in [35] as a example of where probability matching may occur, but its embedded social context provides an interesting thought-experiment for where and why probability matching may occur. This is a game with the same theoretical payoffs as the binary light experiments of traditional probability matching work. That is, the optimal strategy is to always guess "passing play" to ensure you are correct 70% of the time. But this is only optimal if you are unable to infer patterns in the home team's propensity to select a passing play. Perhaps, despite being a naive football observer, you believe the team never runs more than 3 consecutive passing plays. You may have also noticed contextual clues, such that running plays are much more likely when the home team is close to the opposing end-zone, or when players place themselves in particular arrangements.

The availability of contextual clues and the added likelihood the sequential events are not truly random in social contexts could motivate a search for patterns in order to implement a near-perfect strategy. Regarding the missing causal link between pattern searching and probability matching [21, 20], it may be that searching for patterns and optimizing your sequential choices in the meanwhile (by maximizing) is just too difficult.

The easier approach may be to "buy-in" with the current hypothesis you are testing and behave as though your current beliefs are true. Individuals who do not have access to sequential trial outcomes probability match similarly to those who do, which has been argued as evidence that there is no causal link between pattern searching and probability matching [20]. However, it has also been shown that people hold beliefs (priors) about a forthcoming sequence of events [5]. It may be that despite never having the opportunity to update their priors in a sequential learning condition, the priors participants hold about what patterns *may* exist still inform their estimates. As such, a search for patterns may still *cause* probability matching, since the case where probability matching occurs even when estimations are made aggregate [20] can be interpreted as a prior of the *potential* sequential patterns that has been initialized but not updated.

The ambiguous utility of probability matching aside, this behaviour indicates that the participant has learned the underlying probabilities of the task [36]. This internal representation of statistical information can be considered a mental model [34, 9, 16, 8, 5]. Probability matching is therefore a behavioural measure of a mental model. To our knowledge, probability matching has only been studied in the pursuit of a stimulus. How would probability matching appear in the inverted case, what we call *probability antimatching*, where the goal is to *avoid* presented stimuli? In many applied settings, people need to behave in a way that is adversarial to some environmental process. Similar to the pursuit case, the adversary to avoid may follow a certain probability distribution that an individual can represent as a mental model. Consequently, an individual should behave in a way that is maximally dissimilar to the established model of the adversary. The optimal strategy is unambiguous: always select the least likely outcome according to your model of the adversary. However, this optimal avoidance behaviour cannot be assumed, as evident by traditional probability matching.

Hypothesizing about probability antimatching is constrained by the lack of a formally defined "opposite" probability distribution. "Oppositeness" can be considered a special case of "similarity". If object Q is the opposite of P, then Q is minimally similar and maximally informative to P. The "maximally informative" feature is required to ensure the opposite of an opposite is the original. While there are many existing candidates to express similarity between probability distributions [23, 1, 31] it is unclear which most plausibly represents human cognition and reasoning patterns, and how a notion of "opposites" may be mathematically defined under these frameworks. In this work, we utilize a Euclidean vector framework to represent probability distributions. These representative Euclidean vectors can be reflected across the uniform vector/distribution to produce a candidate opposite distribution. This approach affords intuitive geometric interpretation and visualizations, while also providing a direct mapping between psychological constructs and neural models, where geometric representations of various behavioural, cognitive and neural activity models are already established [12, 11, 7, 22].

The children's game of hide-and-seek is an intuitive way to simultaneously explore matching and antimatching. In a simple hide-and-seek game with two players, the 'hider' hides in a set of possible hiding spots while the 'seeker' attempts to find the hider. The seeker wins by matching the hider's decision, while the hider wins by doing the opposite [3, 2]. In this set of studies, we had participants play against a simulated child, who always hides and searches a set of locations according to a given probability distribution. The participant, made aware of this distribution, searches (chases) the child (stimuli) as the seeker, emulating traditional probability matching, and hides from (avoids) the child (stimuli) as the hider, providing a basis to test probability antimatching. Since the child hides and seeks according to the same distribution, this allows investigation of the asymmetries between chasing and avoiding scenarios of probability matching with respect to the exact same distribution within each participant.

Here we present three hide-and-seek studies. Seeking data demonstrates the known phenomenon of probability matching. Hiding data affords us the first known exploration of probability antimatching. In Study 1, we outline an effective methodology to represent participant strategies as linear combination of Euclidean vectors. Using this methodology, we demonstrate participant seeking behaviour is a combination of the optimal maximizing strategy and probability matching, consistent with past work. We also demonstrate that hiding behaviour is a combination of the optimal minimizing strategy (only selecting the least likely outcome) and probability antimatching, defined by our proposed vector reflection. We also find that participant strategy mixes when hiding are more optimal than when seeking, but that this is only true when playing hide-and-seek in a five-room virtual house, not a two-room house. In Study 2, we replicate Study 1 for online data collection finding that seeking is still well defined by maximizing and matching, while hiding is defined by minimizing and antimatching. However, unlike Study 1, we find no differences in hide/seek strategy mix optimality across room-number conditions. In Study 3, we present new distributions that again demonstrate our expected hide/seek strategy mixes. We also generalize our vector reflection hypothesis to account for reflections that produce vectors with negative entries. By exploring two candidate projection methods, we discuss how our geometric model of probability antimatching can account for opposite probability distributions that cannot be expressed behaviorally. Finally, we again find that participants are more optimal when hiding than with seeking, but that this difference decreases from three to five to seven room conditions.

Chapter 2

Study 1: Original Hide-and-Seek

2.1 Methods

2.1.1 Task Environment

Stimuli

We developed a computerized version of the children's game "hide-and-seek". Each participant played a set of hide-and-seek games, against one simulated opponent each game. We introduced each opponent as a predictable agent, that selects hiding locations according to an assigned probability distribution. Each opponent hides according to the same probability distribution to which it seeks. This distribution is made known to each participant.

We presented stimuli on a computer monitor in our lab. Our user interface consisted of a cartoon scene of a house with either two or five rooms, presented on the computer monitor. The house was of a 'dollhouse style' where each room is visible to the participant. We presented a percentage within each room. Each percentage represented the probability that the opponent will hide/seek in that room. Room probabilities always summed to 100%. We also presented four room selection counters. Each tracker was a miniature version of the displayed house, where each room held a numerical value indicating the number of times each room was selected by the participant/opponent to hide/seek (Figure 2.1). Each opponent and their corresponding room selection distribution was drawn from Table 2.1. When a distribution was assigned to a given game, the values of said distribuition were assigned randomly to each of the rooms.

Procedure

We seated participants approximately 40cm away from the screen such that the experimental background encompassed 30 degrees of visual angle. The experiment began with a practice gameplay consisting of three seeks followed by three hides against Toby, one of four simulated opponents the participants faced. As in standard hide-and-seek games, each participant was instructed to find each opponent while seeking, and avoid each opponent while hiding.

As the seeker, the participant was notified that the child was hidden in one of the rooms via a dialogue box on the screen reading “Look for [child]!” where [child] is the name of the current opponent. Using the mouse, the participant selected a room to search with a click. Each room would turn light grey when clicked. If the child was in the selected room, a notification of success was presented as “You found [child]!” and the child was revealed within the room. Otherwise, a notification of failure was presented as “You didn’t find [child]!”. In such failed attempts the child was still revealed. As the hider, the dialogue box read “Hide from [child]!” The participant selected a room to hide using the mouse. If found by the child, the icon would be revealed in the selected room and the message “[child] found you!” was presented. If not found by the child, the child icon was still revealed, and the message “[child] didn’t find you!” was presented. The dynamics of seeking and hiding were otherwise identical. A single hide or seek by the participant is a ‘trial.’ The child’s hiding and seeking locations were drawn from the distribution displayed over the rooms, and were independent between trials. A series of 10 trials constitutes a ‘round.’ A round where the participant seeks the child is called a seeking round. A round where the participant hides from the child is called a hiding round.

The experiment alternated between seeking and hiding rounds, always beginning with seeking. A set of 10 seeking rounds and 10 hiding rounds against a particular child constitutes a ‘game.’ Each participant played a total of three games, one against each of Sally, Kala, and Bo, resulting in 100 seeking trials and 100 hiding trials against each of the three children. To prevent frustration and boredom from interfering with performance in any later games, the last game was always played against Bo, who always had one room with P(hide) = P(seek) = 100%. The first and seconds games were randomly selected as either Sally then Kala, or Kala then Sally. Following the three games, the participants answered demographic questions of gender, age, term of study, and academic program, along with their perceived competence in logical reasoning relative to other students in that program. The experiment was programmed in Python using the Psychopy module [30].

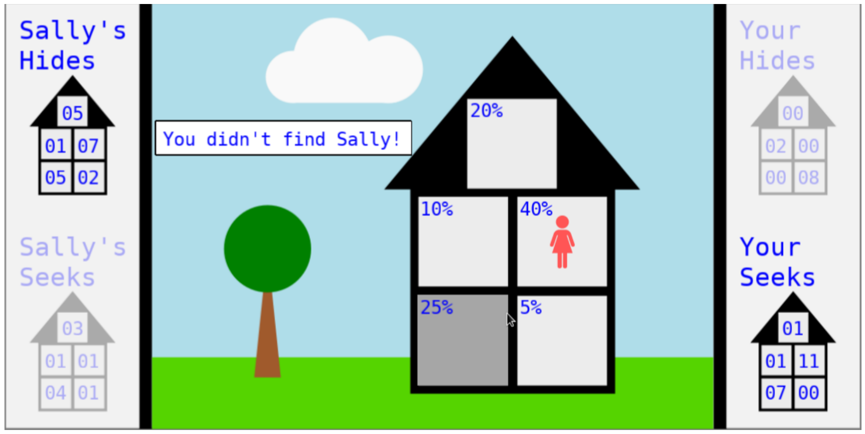

Figure 2.1: Participant view of screen in the five-room condition after selecting the bottom left room in a seek trial. Selected room turns grey, child is revealed, notification is presented, and counters are updated.

2.1.2 Participants

We recruited University of Waterloo students (N=50, 38 females) to participate in exchange for course credit. All participants gave informed consent and the study was cleared by a University of Waterloo Research Ethics Board (REB 41316). Participants were assigned to either the five-room or two-room condition. The first 17 participants were assigned to the five-room condition, as that protocol was created first. Every fourth participant thereafter was assigned to the five-room condition, and the rest to the two-room condition. The experiment took approximately 25 minutes to complete.

2.1.3 Data Analysis

The principal aim of this analysis is to quantitatively represent the strategies used by participants. Discussions of probability matching are typically framed as non-optimal, and contrasted with the optimal alternative strategy of maximizing. It is therefore reasonable to construct a model of stimulus pursuit behaviour (seeking) on the basis of some combination of probability matching and maximizing strategies. More generally, this can be expressed as:

$$\text{participant strategy} = \text{optimal strategy} + \text{non-optimal strategy}$$

Opponent name (order)	2 Room	5 Room
 Toby (Practice)	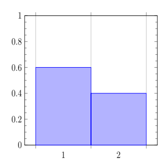	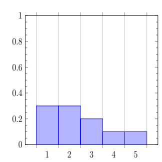
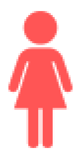 Sally (1st or 2nd)	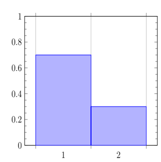	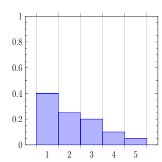
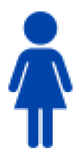 Kala (1st or 2nd)	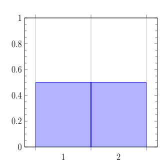	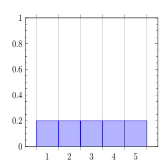
 Bo (3rd)	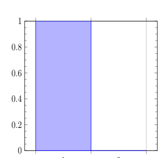	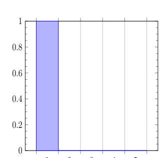

Table 2.1: Opponents faced by participants in each condition. Each distribution is presented here in decreasing room probabilities, though probabilities were randomly assigned to rooms during the experiment.

The purpose of this generalization is to also account for stimulus avoidance behaviour (hiding). When hiding, the optimal strategy is to *minimize* by always picking the least likely room. The non-optimal hiding strategy, however, is less clear. Unlike probability matching in the seeking case, there is no previous literature to describe what the analogous strategy may be for the hiding case. Nonetheless, we posit the combination of optimal and non-optimal hiding strategies should explain hiding behaviour just as well as the combination of optimal and non-optimal seeking strategies explain seeking strategies. The next section describes the mechanics of how we construct such a model, and the motivation for selecting a suitable non-optimal strategy, analogous to probability matching, to be used to describe hiding behaviour.

Histogram Vectors

Representing our data as Euclidean vectors allows for convenient geometric intuition and algebraic manipulations. Furthermore, we can take advantages of the known regularities within our data to simplify analysis. For example, any histogram representing a discrete probability distribution must contain values between 0 and 1 that sum to 1. Consider any histogram that could apply to the two-room hide-and-seek game (Figure 2.2). These histograms can be expressed as vectors in 2D space, and only exist in the positive quadrant on the diagonal line from (0,1) to (1,0). Any vector out of this region would have elements that do not sum to 1, or have negative values.

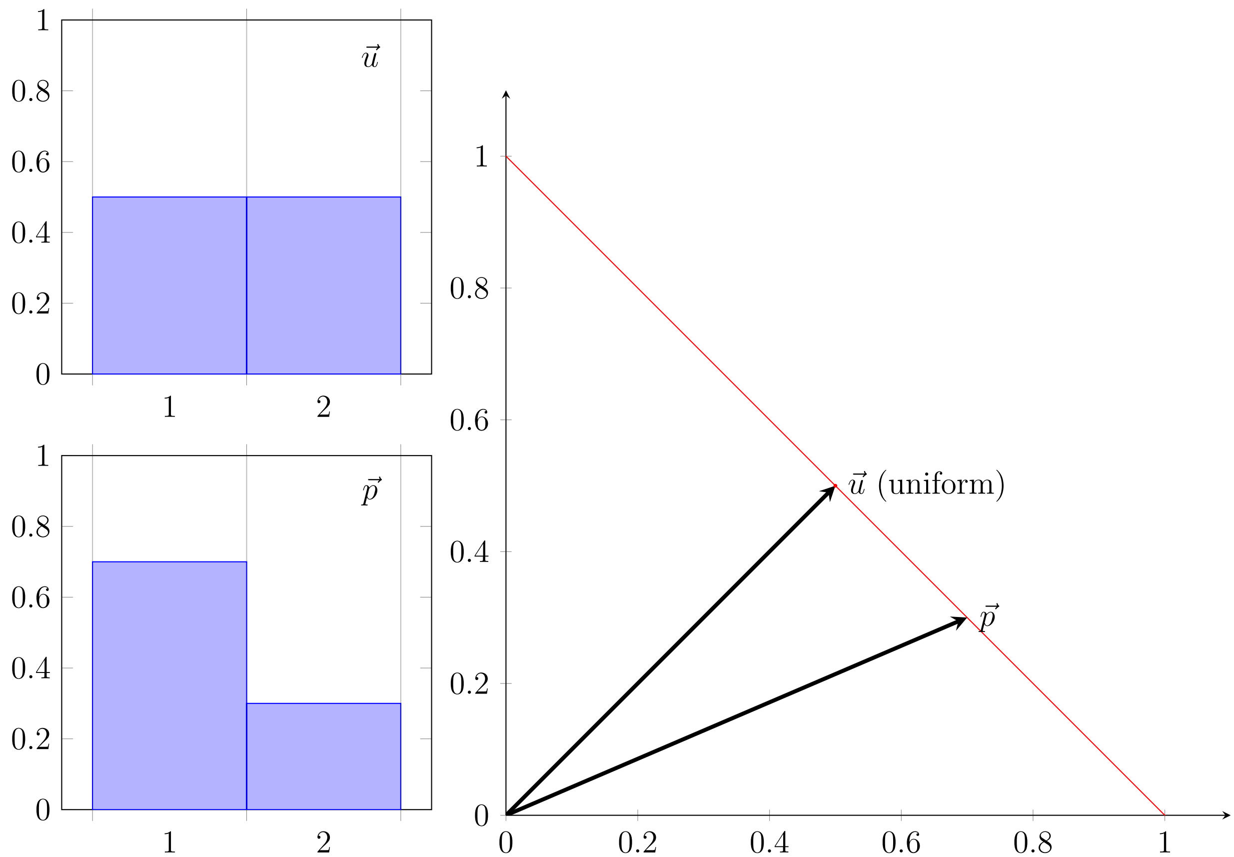

Figure 2.2: Any histogram with two bins has a unique representative 2D histogram vector. This vector must exist on the line from (0,1) to (1,0) since the sum of all histogram values must equal 1.

This representation holds in higher dimensions as well. That is, histograms from the five-room condition can be represented as vectors in 5D, similarly existing on the diagonal plane connecting the points (1,0,0,0,0), (0,1,0,0,0), ..., (0,0,0,0,1). For convenience, we will mainly refer to the 2D case when describing the analysis. All properties of the 2D case apply also in higher dimensional cases, unless otherwise noted. Expressing *participant seek strategy* = *maximizing strategy* + *matching strategy* algebraically, we have $\vec{s} = \alpha\vec{x} + \beta\vec{m}$, where $\vec{s}$ is the recorded participant room click frequencies when seeking, $\vec{x}$ is the strategy where only the most likely room is selected (maximizing), and $\vec{m}$ is the strategy where rooms are selected identically to the opponent strategy (probability matching). Therefore, the participant seek strategy can be expressed as a linear combination of a maximizing strategy and matching strategy. If the participant is purely maximizing, then $\vec{s} = \vec{x}$, implying $\alpha = 1$ and $\beta = 0$. If the participant is purely matching, then $\vec{p} = \vec{m}$, implying $\alpha = 0$ and $\beta = 1$. Otherwise, α and β are some values between 0 and 1. We constrain these values between 0 and 1 because one cannot conceivably use less than 0% or more than 100% of a particular strategy. For example, consider the case where $\vec{m} = (0.70, 0.30)$, and, necessarily, $\vec{x} = (1, 0)$. If we observe $\vec{s} = (0.85, 0.15)$, we can express $\vec{s}$ as $\frac{1}{2}\vec{m} + \frac{1}{2}\vec{x} = (0.85, 0.15)$. Therefore, we can express $\vec{p}$ as the pair of values $(\alpha, \beta) = (\frac{1}{2}, \frac{1}{2})$. This relationship with respect to α and β can be plotted for visual interpretation (Figure 2.3).

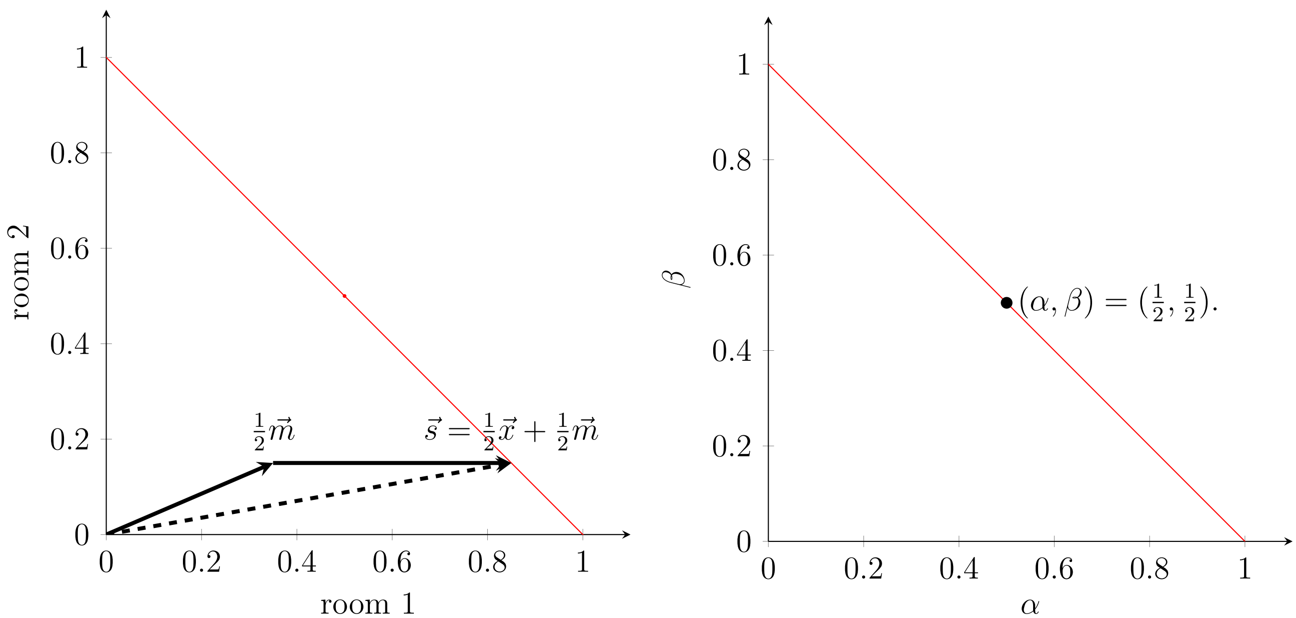

Figure 2.3: $\vec{s} = (0.85, 0.15)$ of the form $\vec{s} = \alpha\vec{x} + \beta\vec{m}$ where $\vec{m} = (0.7, 0.3)$ and $\vec{x} = (1, 0)$, can be represented as a point in the strategy space as $(\alpha, \beta) = (\frac{1}{2}, \frac{1}{2})$. Each dimension of the strategy space represents the amount of a particular strategy used by a participant (α for maximizing, β for matching. Each can hold a value between 0 and 1, inclusive.

Effectively, this is a change of variables from a 'room space' where dimensions are the room probabilities, to a 'strategy space' where dimensions are α and β, indicating the mix of strategy used. In this study, the room space is 2D in the two-room condition, and 5D in the five-room condition. Since the same mechanisms apply to the 5D interpretation, all participant data, regardless of condition, can be reduced to a 2D strategy space. However, it is possible that no combination of two strategy vectors alone will be able to create a particular $\vec{s}$. It may be that $\vec{s}$ lies outside of the region formed by every possible combination of $\vec{x}$ and $\vec{m}$, such that α and β are between 0 and 1. In such cases of imperfect fits, a third vector, representing error, denoted $\vec{\epsilon}$ is required to fully describe participant seeking behaviour (Figure 2.4).

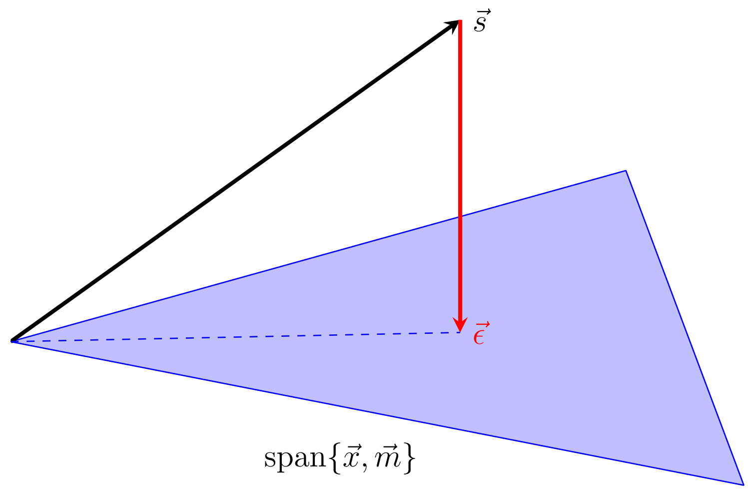

Figure 2.4: In higher dimensional spaces $\vec{s}$ (black) may not exist on the plane created by $\vec{x}$ and $\vec{m}$ (blue region). An error vector ϵ (red) can be added to return to the plane. The length of this error vector can be used to express the goodness of the model.

We thus revise our description of participant seeking behaviour as follows: *participant seek strategy = maximizing strategy + matching strategy - error*. Algebraically, $\vec{s} = \alpha\vec{x} + \beta\vec{m} - \vec{\epsilon}$. This error is visible in strategy space plots as well. Error is minimized when $\alpha + \beta = 1$, which implies participant strategies are explained exclusively by $\vec{x}$ and $\vec{m}$. This occurs on the diagonal line from (1,0) to (0,1) in the strategy space (Figure 2.5). As $\alpha + \beta$ approaches 1, error reduces, implying a better model fit. Error will be reported as residual sums of squares in the results section. This is equivalent to the sum of the squared elements of each error vector.

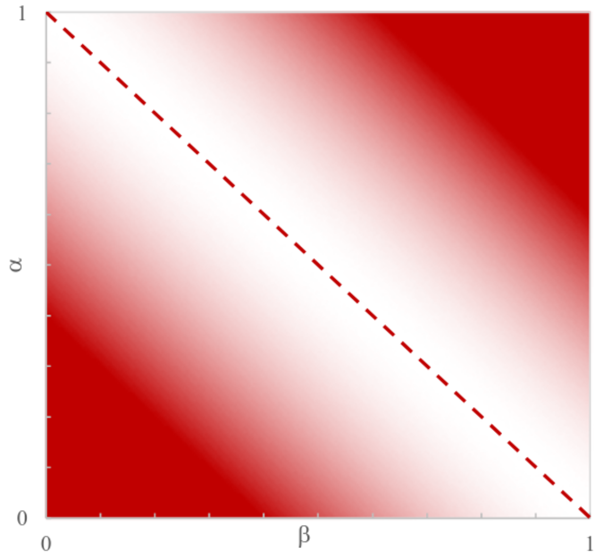

Figure 2.5: Error (red regions) is minimized for strategies on the line $\alpha + \beta = 1$, when $\vec{epsilon} = \vec{0}$. Many points on or near this line imply a good model fit. Many points in high error regions imply a poor fit.

We aim to express participant hiding behaviour in the same form as seeking. As such, *optimal* and *non-optimal* hiding strategies must be defined. The former is self-evident; when hiding, the optimal strategy is to always select the room with lowest probability of being found. The latter is not as clear. Utilizing a geometric argument, we propose a possible non-optimal hiding strategy that is symmetric to the non-optimal seeking strategy of probability matching already established by the literature. We propose a non-optimal hiding strategy that is literally symmetric to probability matching across the uniform distribution. Any 2D histogram vector forms an angle between itself and the uniform histogram vector. There is a unique vector across the uniform vector, opposite from the original, that forms this same angle (Figure 2.6).

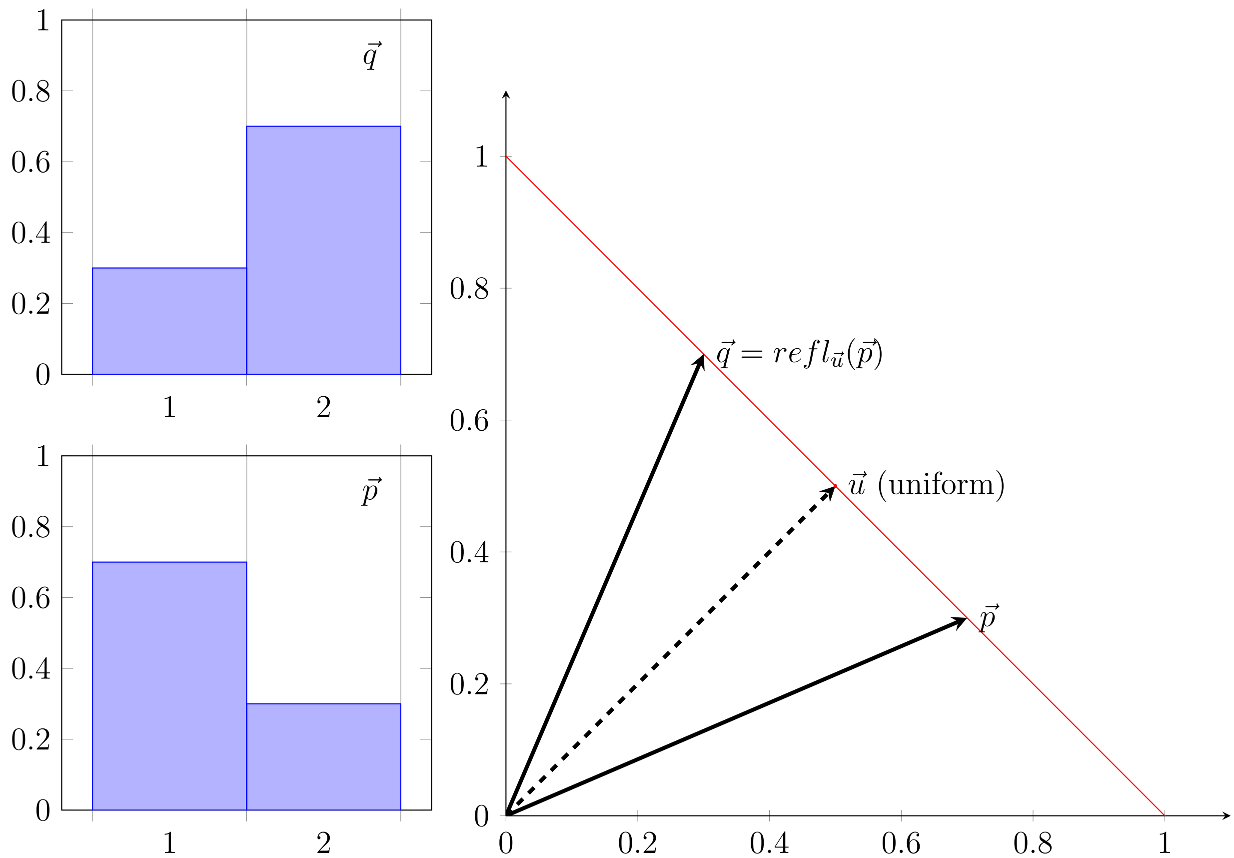

Figure 2.6: All histogram vectors form an angle with the uniform vector. The reflected vector forms that same angle with the uniform. Above, (0.7, 0.3) reflects to (0.3, 0.7) which appears visually opposite

This *histogram reflection* over the uniform, denoted $refl_{\vec{u}}(\vec{p})$ for any histogram $\vec{p}$, is computationally inexpensive, provides a geometric interpretation of 'opposite', and in our opinion, results in a visually opposite histogram. However, this method does not always produce a viable reflection. In higher dimensional spaces, the reflected vector lies within the plane formed by the original vector and the uniform, but may require negative elements in order for the reflected histogram to have the same angle to the uniform as the original to the uniform. No such degenerate cases were used in this Study, with the exception of the distribution with 100% probability mass assigned to one room. This distribution was not used to calculate opposite fits.

Based on the arguments above, we choose the histogram vector reflection of the opponent distribution to act as the non-optimal strategy component for participant hiding behaviour. We denote this strategy *probability antimatching*. Therefore, *participant hide strategy = minimizing strategy + antimatching strategy – error*. Algebraically, $\vec{h} = \gamma\vec{n} + \delta\vec{a} - \vec{\zeta}$, where $\vec{h}$ is the recorded participant room click requencies when hiding, $\vec{n}$ is the strategy where only the least likely room is selected (minimzing), $\vec{a}$ is the strategy where rooms are selecting exactly according to the vector reflection of the opponent strategy (probability antimatching), and $\vec{\zeta}$ is the error in representation. As with seeking, we can convert data from the room space to the strategy space (in terms of γ and δ) for analysis.

Error in each proposed model, $\vec{\epsilon}$ for seeking and $\vec{\zeta}$ for hiding, not only indicate the goodness of each model, but also assist in identifying whether the vector reflection is the appropriate antimatching strategy in hiding to complement the established matching strategy in seeking. That is, if in aggregate, $\vec{\epsilon}$ and $\vec{\zeta}$ are sufficiently close in magnitude, we gain confidence that a combination of minimizing and antimatching is just as good at explaining hiding behaviour as a combination of maximizing and matching is at describing seeking behaviour. The above models were constructed using the Stark-Parker algorithm for bounded-variable least squares via the 'bvls' package in R [27].

To ensure that our selected methods are not 'unreasonably forgiving' in fitting data into any given model, we perform sensitivity analyses. This constitutes attempting to describe participant hiding behaviour using two random vectors instead of using the optimal minimizing and non-optimal antimatching vectors. Algebraically, we test whether $\vec{h} = \gamma\vec{n} + \delta\vec{a} - \vec{\zeta}$ explains the data any better than $\vec{h} = \gamma_r\vec{r_1} + \delta_r\vec{r_2} - \vec{\zeta_r}$, where $\vec{r_1}$ and $\vec{r_2}$ are 'random histograms'. Since we require $\vec{r_1}$ and $\vec{r_2}$ to be valid histogram vectors, their elements must each sum to 1, where each element is between 0 and 1, inclusive. This added constraint means we cannot use a vector of random numbers, as each element cannot be completely independent of the others. Deriving random vector for the two-room data is fairly straightforward. For a random vector $\vec{r} = (r_1, r_2)$, let r_1 be a random number between 0 and 1, and $r_2 = 1–r_1$. For the five-room data, we generalize further. For random vector $\vec{r} = (r_1, r_2, r_3, r_4, r_5)$, let r_1 be a random number between 0 and 1, r_2 be a random number between 0 and $(1–r_1)$, r_3 be a random number between 0 and $(1–r_1–_2)$, r_4 be a random number between 0 and $(1–r_1–r_2–r_3)$, and $r_5 = 1–r_1–r_2–r_3–r_5$. In both the 2D and 5D case, the order of the elements is then shuffled. The stochastic nature of this analysis leaves results vulnerable to noise, where the random histograms may be abnormally similar or dissimilar to the $\vec{n}$ or $\vec{a}$ of the original model. To account for this, we repeat this analysis 1000 times and report a summary of this data.

Variable	Definition
Seeking variables	
$\vec{s}$	Recorded participant seek strategy
$\vec{x}$	Optimal maximizing strategy
$\vec{m}$	Probability matching strategy
α	Amount of $\vec{x}$ needed to best describe $\vec{s}$
β	Amount of $\vec{m}$ needed to best describe $\vec{s}$
$\vec{\epsilon}$	Error in representing $\vec{s}$ with $\vec{x}$ and $\vec{m}$
Seeking variables	
$\vec{h}$	Recorded participant hide strategy
$\vec{n}$	Optimal minimizing strategy
$\vec{a}$	Probability anti-matching strategy $:= refl_{\vec{u}}(\vec{m})$
γ	Amount of $\vec{n}$ needed to best describe $\vec{h}$
δ	Amount of $\vec{a}$ needed to best describe $\vec{h}$
$\vec{\zeta}$	Error in representing $\vec{s}$ with $\vec{x}$ and $\vec{m}$

Note: Vectors and parameters with the subscript r are randomly generated for sensitivity analyses. $\vec{u}$ denotes the uniform distribution.

Table 2.2: Definition of variables

2.2 Results and Discussion

2.2.1 Seeking Sally (Two-Room)

The seeking rounds against Sally in the two-room condition most closely resemble traditional studies of binary choice probability matching. A plot of all participant seek frequencies overlaid with Sally's hide frequencies reveals a spectrum of strict maximizers, strict matchers, and some mixing of the two (Figure 2.7). Converting and plotting each participant into the strategy space of α and β, refines our interpretation (Figure 2.8). Participants employing a mixed strategy appear to be using a mix specifically of maximizing and matching, as all participants lie on or essentially on the line $\alpha + \beta = 1$ implying there is negligible error in this fit (Mean RSS = 0.0019034).

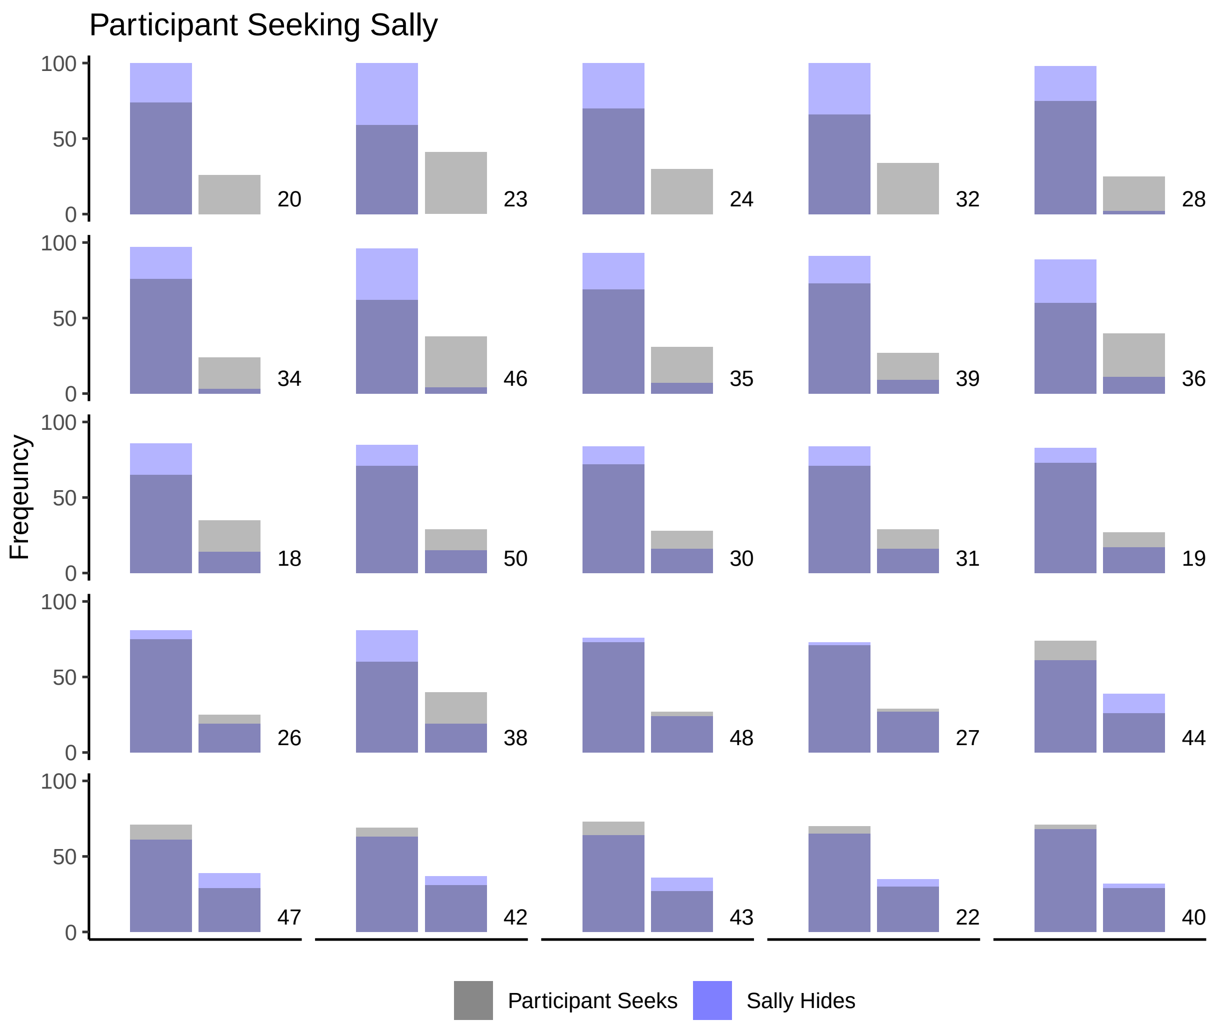

Figure 2.7: Participant seeks (blue) and Sally's hides (grey) sorted by increasing values of β. Histogram bins sorted by decreasing values of Sally's distribution. Blue participant seek bars at 100% indicate perfect maximizing (see top left). Participant seek bars perfectly overlapping Sally's hides indicate perfect matching (see bottom right)

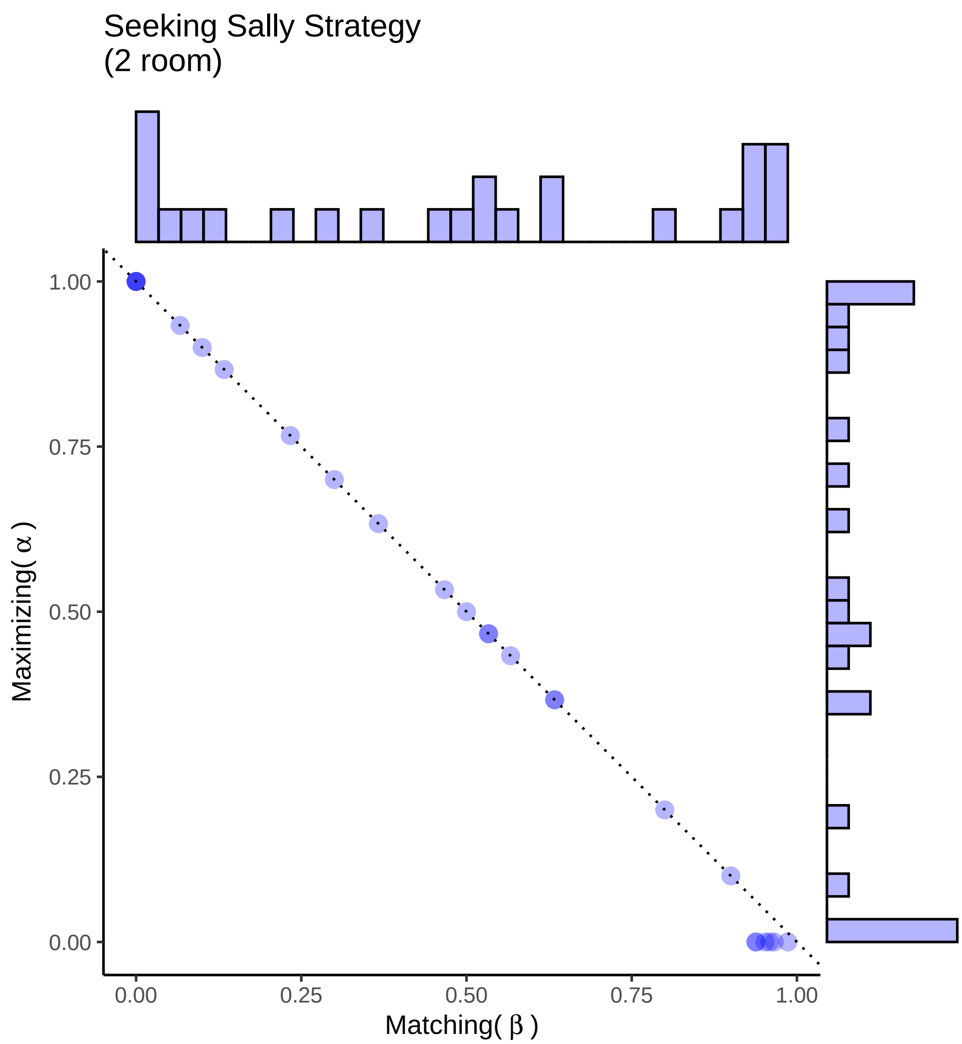

Figure 2.8: Strategy space of participant seeks against Sally. Bars along margin count participants using particular value of α or β. Perfect minimizers use $\alpha = 1$ and $\beta = 0$ (top left). Perfect matchers use $\alpha = 0$ and $\beta = 1$ (bottom right). Participant strategies existing on or near line $\alpha + \beta = 1$ implies near perfect model fit.

2.2.2 Hiding from Sally (Two-Room)

Hiding rounds against Sally provide stimulus avoidance data most analogous to stimulus pursuit in traditional binary choice probability matching. Plots of participant hide frequencies over Sally's seek frequencies reveal a pattern similar to that of the seeking data (Figure 2.9). Some participants behave optimally, minimizing by only hiding in Sally's least searched rooms. Others appear to employ some 'opposite' strategy to Sally's, hiding in Sally's most searched room as frequently as Sally searches the least searched room. Upon viewing this data in the strategy space, we can see the participants used a mix specifically of an optimal minimizing strategy and a non-optimal antimatching strategy, defined here as a 2D histogram reflection (Figure 2.10). All participants were on or near the line $\gamma + \delta = 1$, implying negligible error (Mean RSS = 0.0009379).

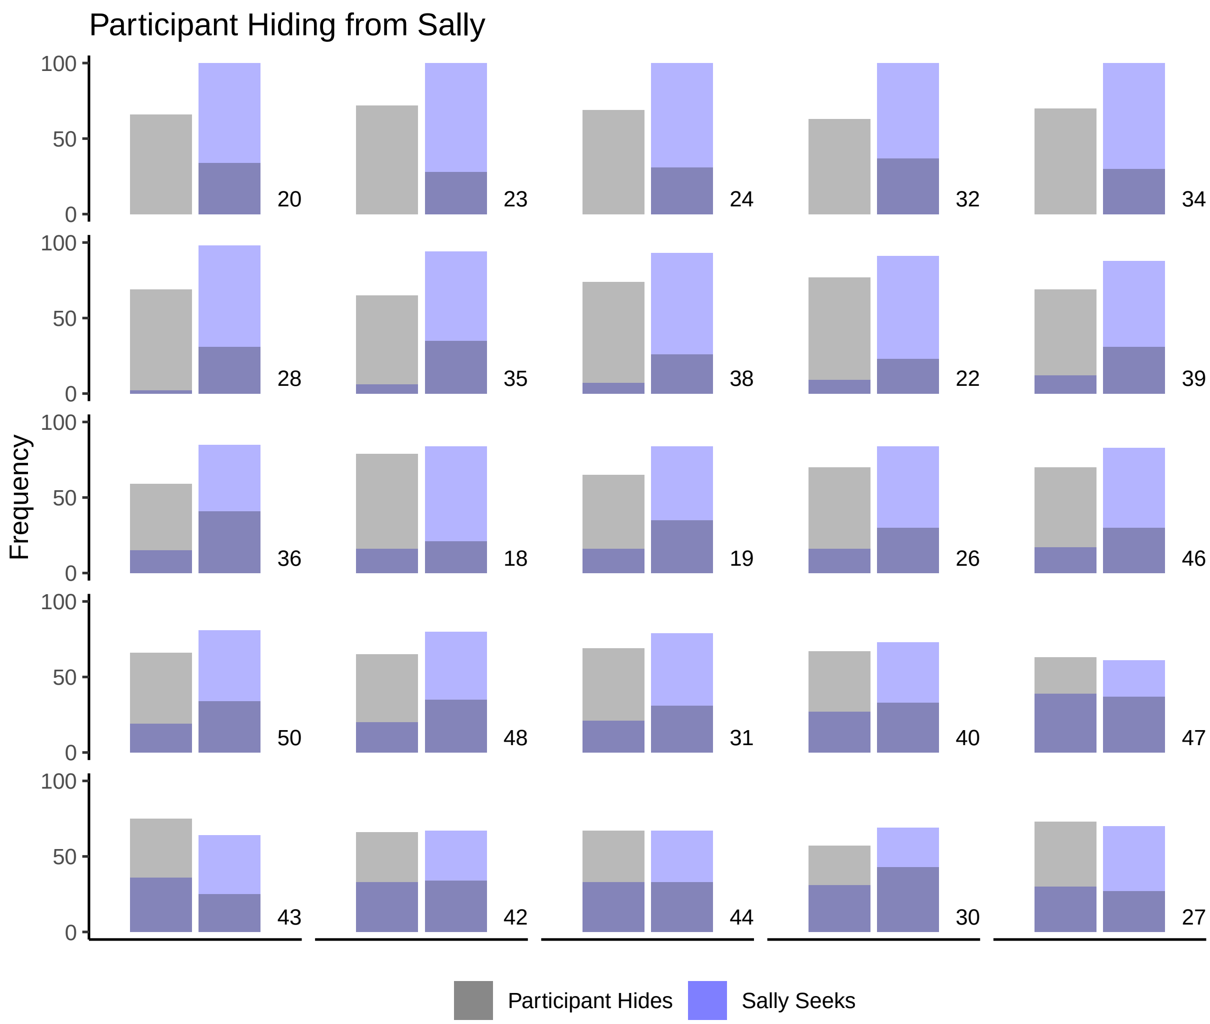

Figure 2.9: Participant hides (blue) and Sally's hides (grey) sorted by increasing values of δ (use of antimatching strategy). Histogram bins sorted by decreasing values of Sally's distribution. Blue participant hide bars at 100% indicate perfect minimizing (top left). Participant hide bars at $1 -$ Sally's seek bars indicate antimatching (bottom right)

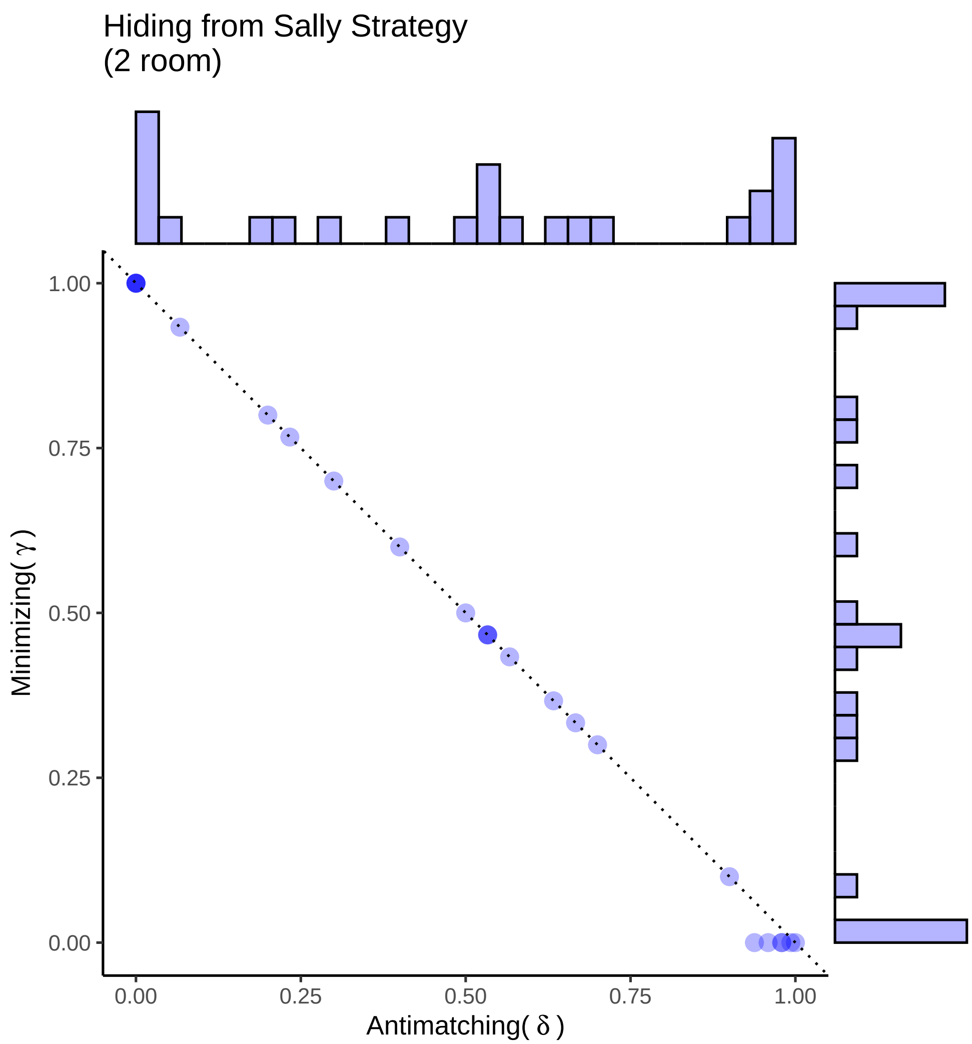

Figure 2.10: Strategy space of participant hides against Sally. Bars along margin count participants using particular value of γ or δ. Perfect minimizers use $\gamma = 1$ and $\delta = 0$ (top left). Perfect antimatchers use $\gamma = 0$ and $\delta = 1$ (bottom right). Participant strategies existing on or near line $\gamma + \delta = 1$ implies near perfect model fit.

2.2.3 Comparing Hiding and Seeking (Two-Room)

Comparing strategy spaces of hiding and seeking behaviours reveal similar mixes of optimal and non-optimal strategies (Figures 2.8 and 2.10). In aggregate, *maximizing* + *matching* describes seeking similarly as *minimizing* + *antimatching* describes hiding.

In general, participants did not become more or less optimal in their strategy choice when hiding instead of seeking against Sally. A large portion had an identical strategy mix between hiding and seeking, and of those who didn't, no discernible pattern was found to indicate that people generally become more or less optimal when hiding (Figure 2.11). A boxplot of the same data is consistent with this conclusion (Figure 2.12).

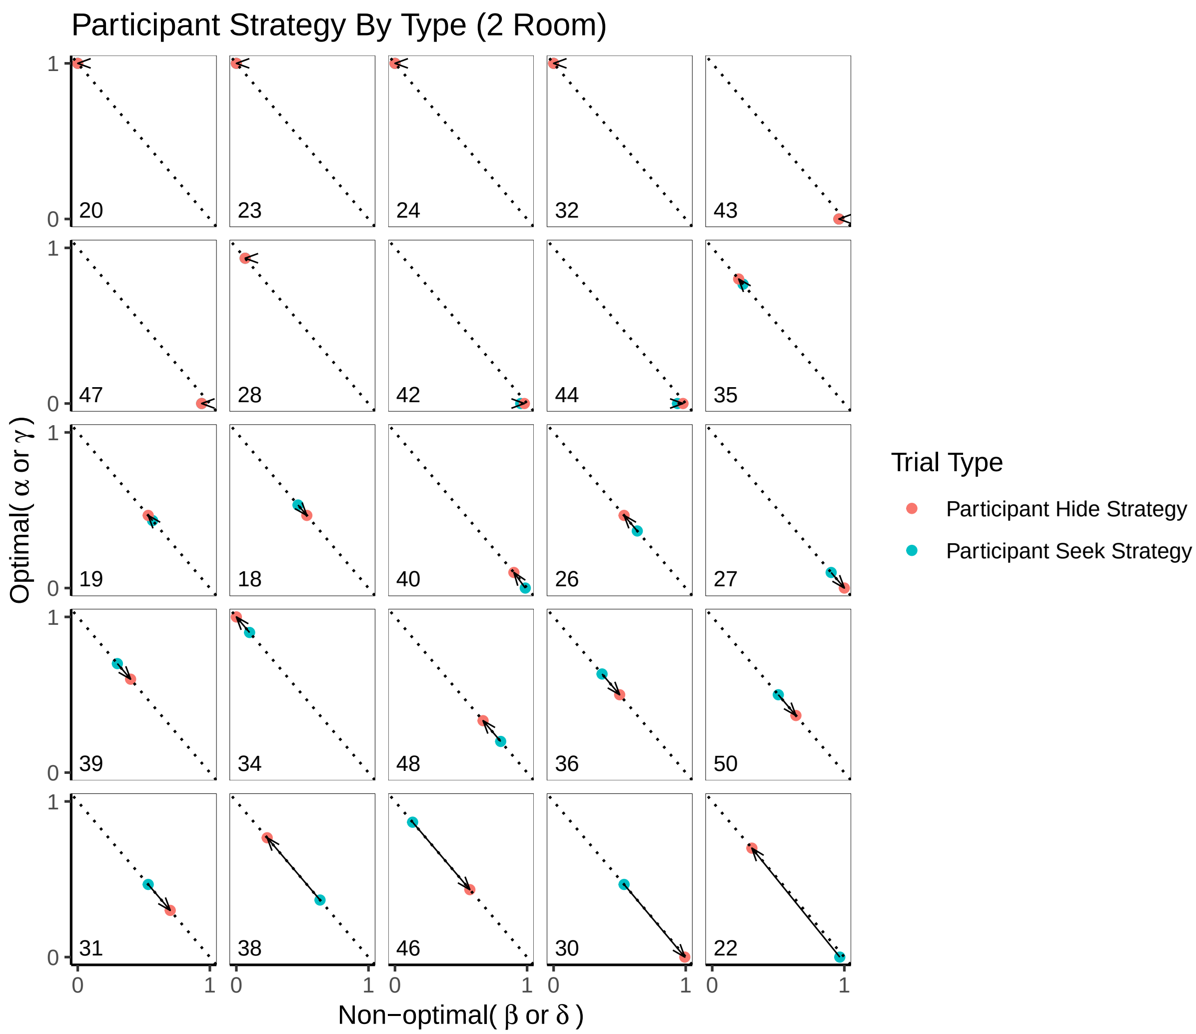

Figure 2.11: Each participant's strategy mix for seeking (α, β) hiding (γ, δ) are plotted together, ordered by increasing values of strategy change. Arrow points from seek strategy to hide strategy. Most participants have little change. Those who do, change in a non-predictable direction

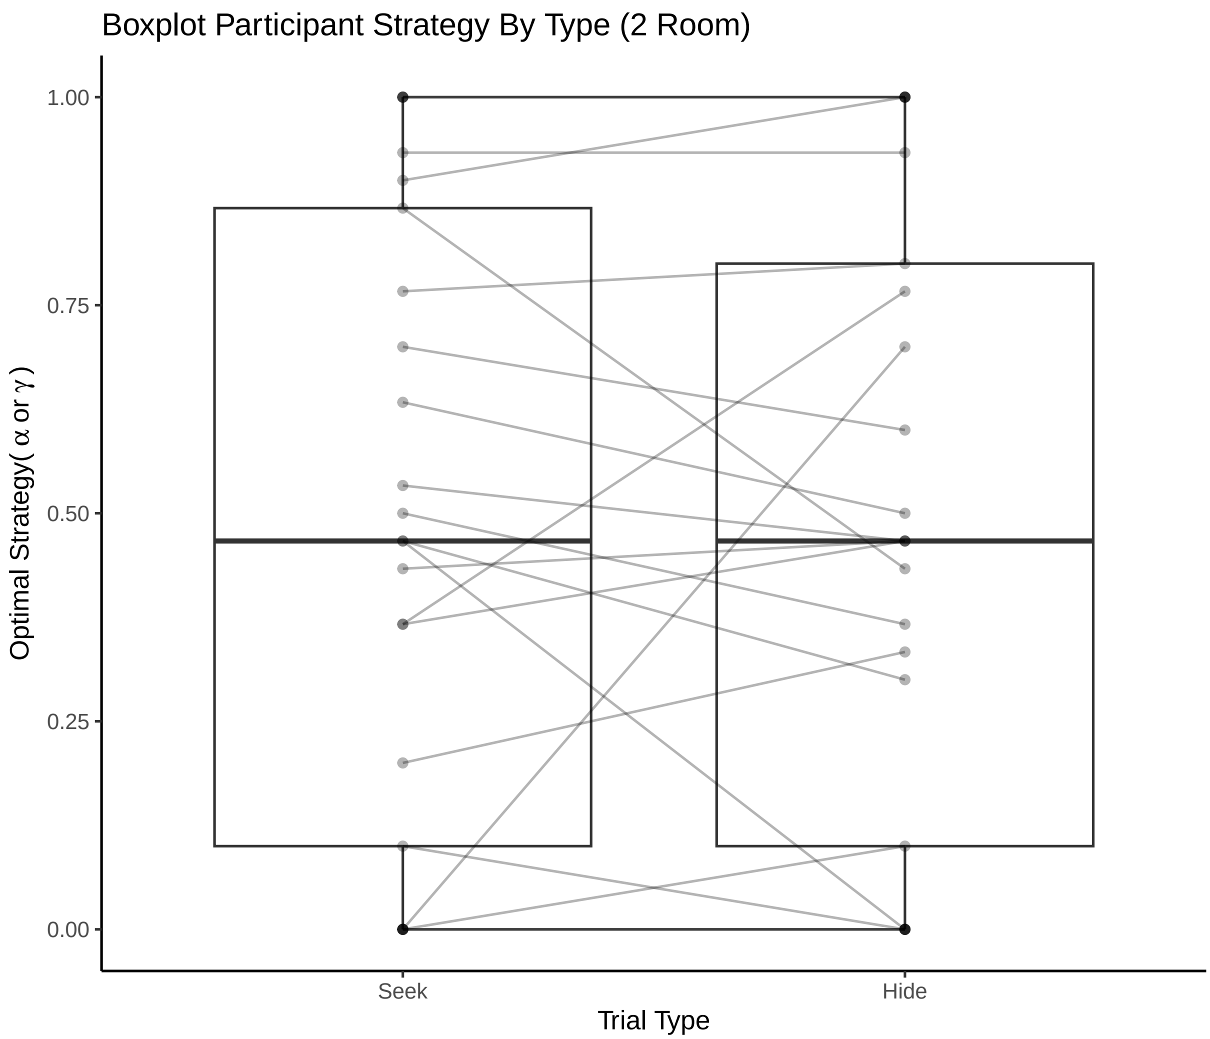

Figure 2.12: Most individual participants use a similar strategy mix between hiding and seeking. Some change strategies between hiding and seeking exist, though in an unpredictable direction and average out in aggregate. The result is a roughly equivalent use of the optimal strategy in hiding and seeking.

2.2.4 Sensitivity Analysis (Two-Room)

Performing a sensitivity analysis on our model of participant hiding by replacing the minimizing and antimatching strategy vectors with two random histogram vectors indicates (1) that $\vec{h} = \gamma\vec{n} + \delta\vec{a} – \vec{\zeta}$ is a suitable model, and (2) that using the bounded-variable least squares method to present and analyze data is not unreasonably forgiving to any arbitrary model it is given. The plot presenting the strategy space for participant hides against Sally is repeated four times using random strategy vectors instead of the minimizing ($\vec{n}$) and antimatching ($\vec{a}$) strategies of our model (Figure 2.13). All four random strategy spaces presented fail to explain participant behaviour effectively, resulting in few strategies lying on or near the line $\gamma_r + \delta_r = 1$ and inflated Mean RSS values ranging between 0.2345 and 0.5594. Repeating this procedure 1000 times and plotting Mean RSS draws the same conclusion: $\vec{h} = \gamma\vec{n} + \delta\vec{a} – \vec{\zeta}$ is a more reliable model than $\vec{h} = \gamma_r r\vec{1} + \delta_r r\vec{2} – \vec{\zeta_r}$ (Figure 2.14).

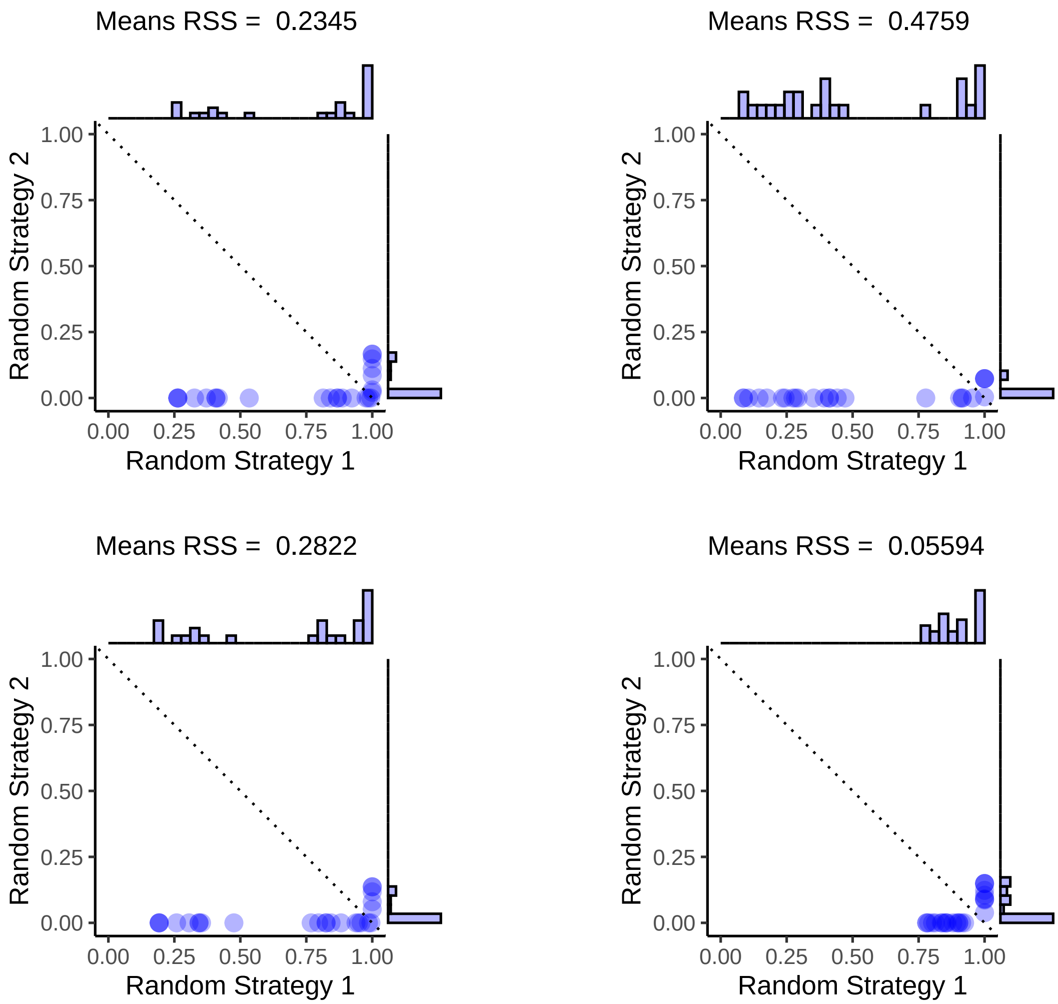

Figure 2.13: Plotting participant hides against sally in a strategy space using two random strategy vectors instead of minimizing and antimatching strategies results in a poor model fit. Many points are in region of high error, reflected in high Mean RSS values. Figure contains four instances of this plot, each with two new randomly generated strategy vectors.

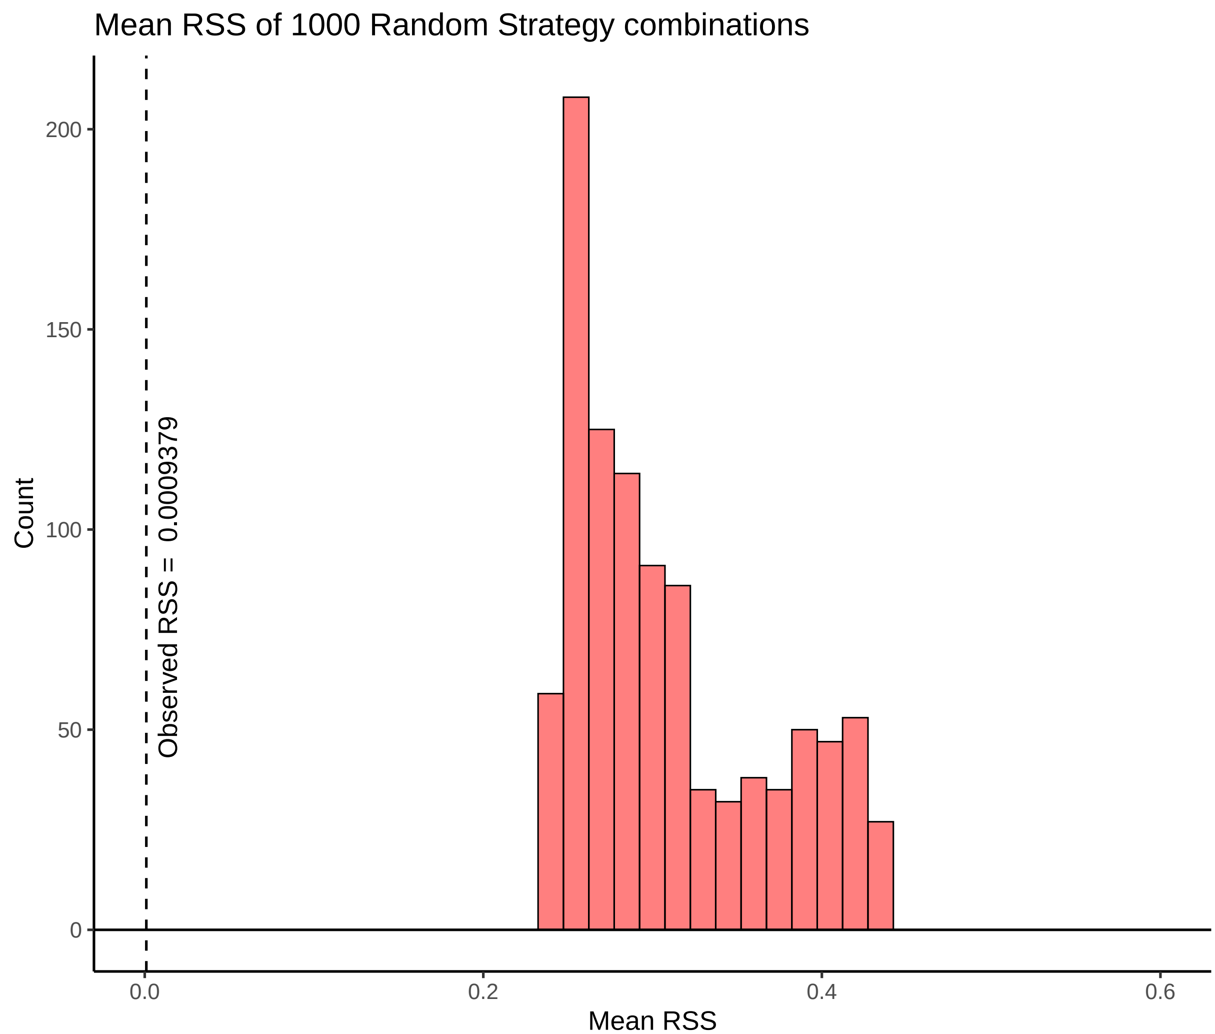

Figure 2.14: Fitting participant data to a strategy space of two random strategy vectors 1000 times. Figure shows histogram of resulting error. Our proposed model dramatically outperforms the 1000 tested here. This implies that our model is of meaningful accuracy, and our methods used are not biased to result in low error for any arbitrary model it is given.

2.2.5 Seeking Sally (Five-Room)

The seeking rounds against Sally in five-rooms are an attempt to generalize from the traditional binary choice stimulus pursuit form of probability matching. A plot of all participant seek frequencies overlaid with Sally's hide frequencies reveals a spectrum of nearly strict maximizers, nearly strict matchers, though many are using a strategy mix of predominantly matching (Figure 2.15). Converting and plotting each participant into the strategy space of α and β, refines our interpretation. Participants appear to be using a mix of maximizing and matching but strongly favour matching (Figure 2.16). All participants lie in close proximity to the line $\alpha + \beta = 1$ resulting in minimal error (Mean RSS = 0.0101012). This error in the five-room case is larger than in the two-room, though some increase in error can be expected as data increases in dimension and it becomes more difficult to express in a two-dimensional strategy space since participants have more degrees of freedom for how they can construct strategies.

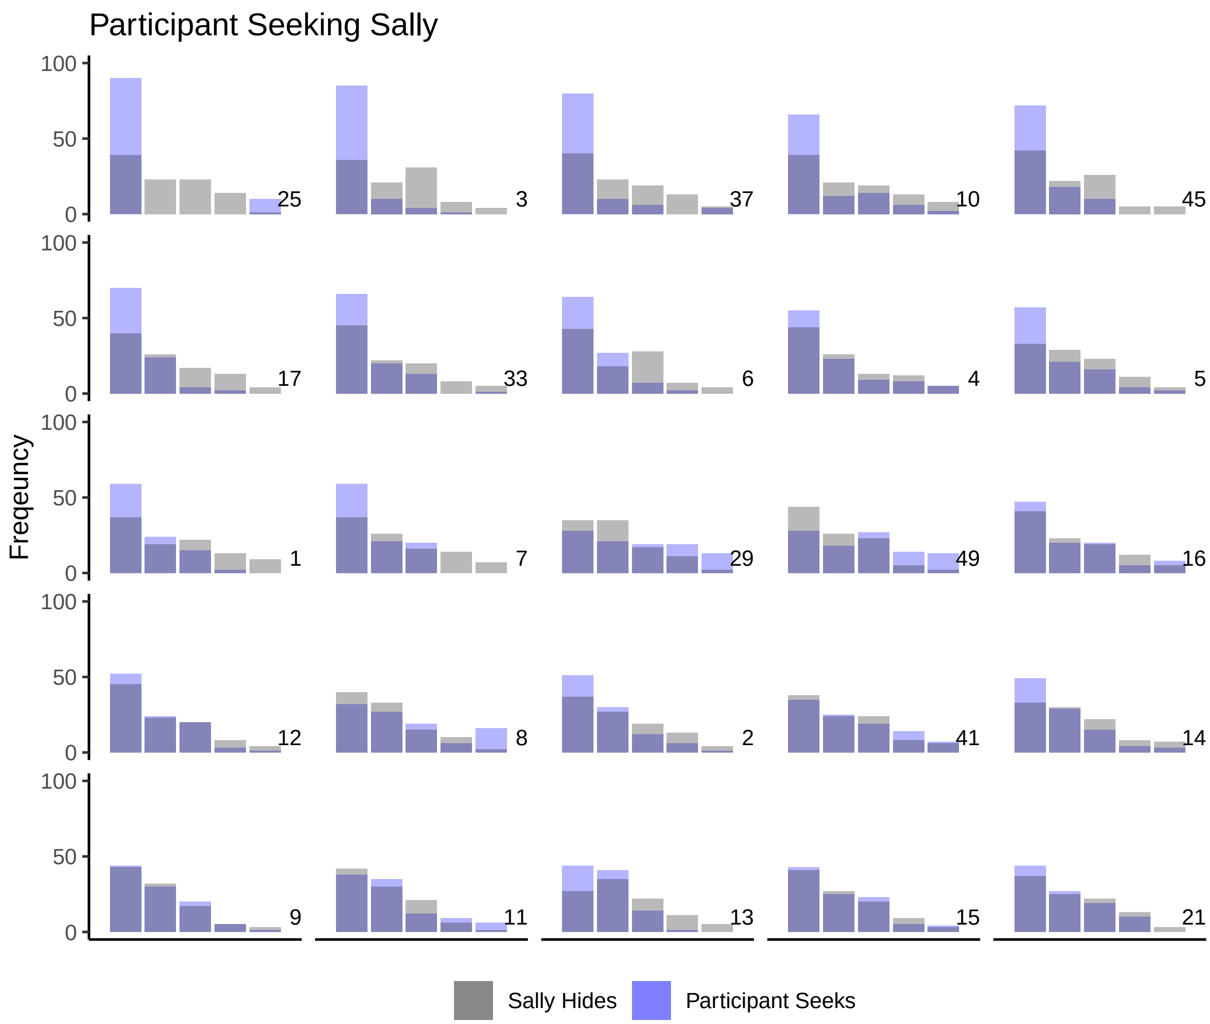

Figure 2.15: Participant seeks (blue) and Sally's hides (grey) sorted by increasing values of β. Histogram bins sorted by decreasing values of Sally's distribution. All participants use a mix of almost only maximizing and matching, though most tend toward a mix containing more matching than maximizing.

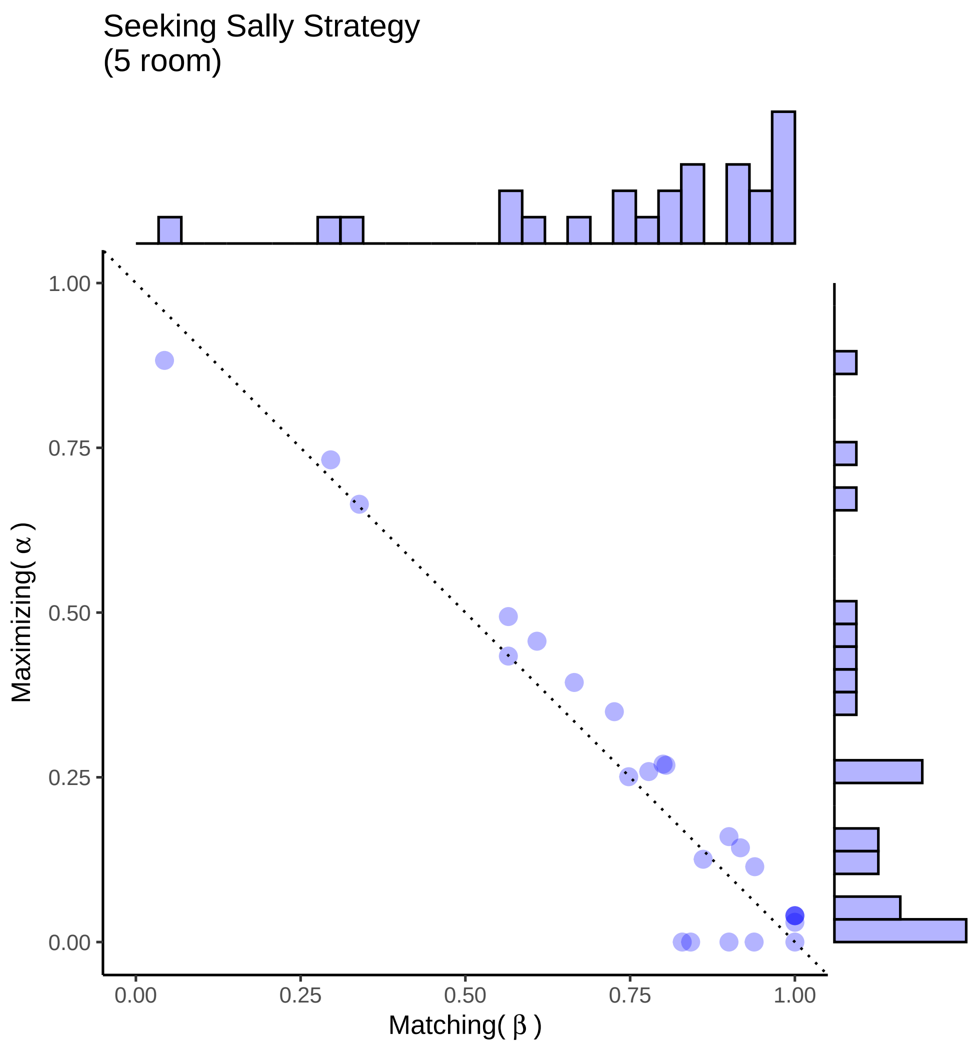

Figure 2.16: Strategy space of participant seeks against Sally. Bars along margin count participants using particular value of α or β, and indicate that most participants use a strategy with a low maximizing component (α) and high matching component (β).

2.2.6 Hiding from Sally (Five-Room)

The hiding rounds against Sally in the five-room condition combines the novel concept of probability antimatching in a more complicated environment. This added complexity provides the means to generalize the conclusions drawn from the two-room antimatching data. Plots of participant hide frequencies over Sally's seek frequencies reveal a pattern similar to hides of the two-room condition, where participants range from perfect minimizing to perfect antimatching (Figure 2.17). Viewing this data in strategy space reveals that participants are indeed using strategies made mostly of minimizing and antimatching (Figure 2.18). As with five-room seeking data, participant strategies did not lie exactly on the line $\gamma + \delta = 1$, though this is expected with higher dimensional data. Error was still relatively low (Mean RSS = 0.0171683).

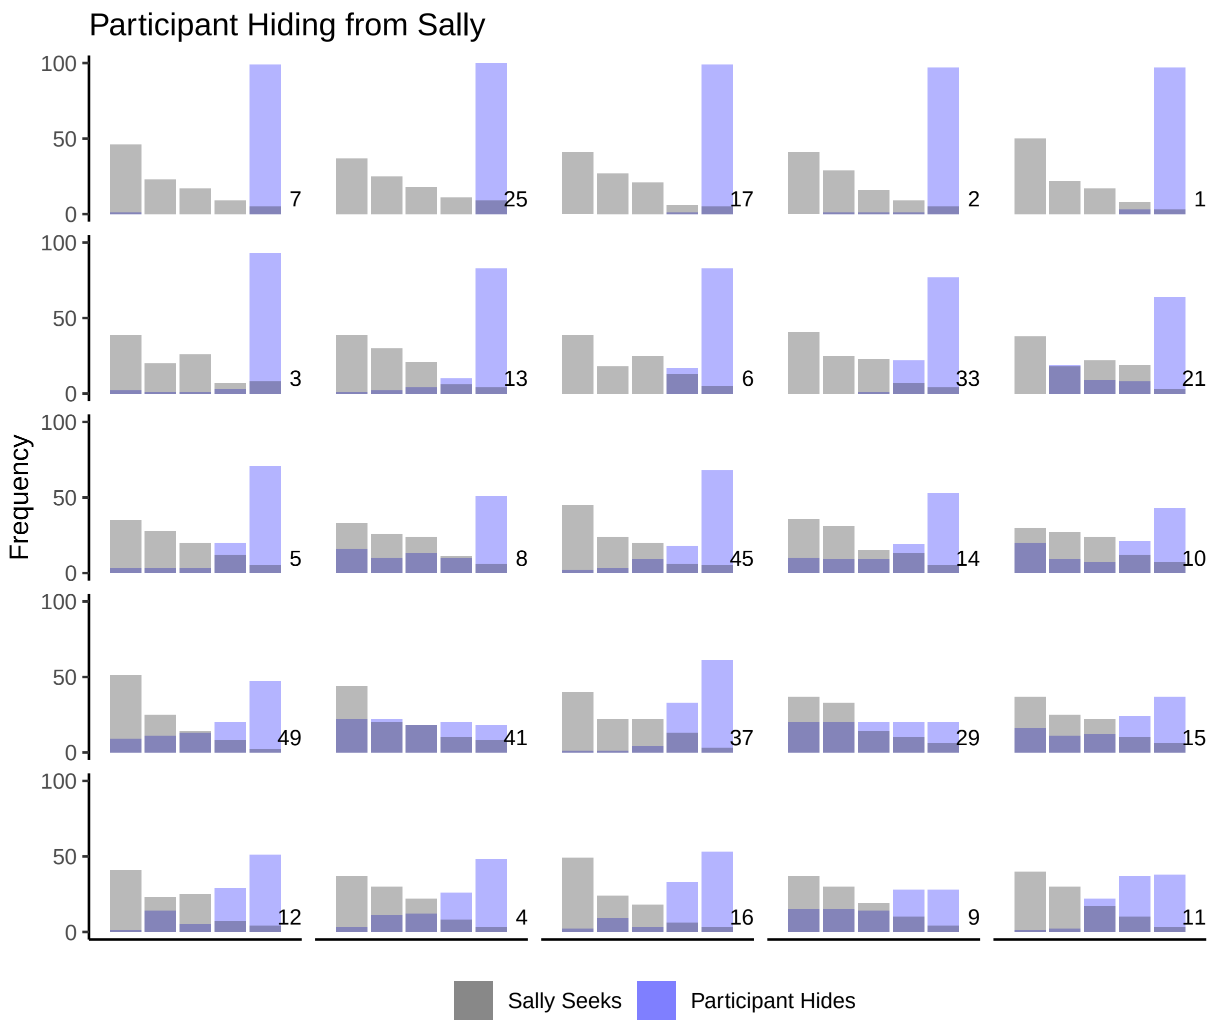

Figure 2.17: Participant hides (blue) and Sally's hides (grey) sorted by increasing values of δ (use of antimatching strategy). Histogram bins sorted by decreasing values of Sally's distribution. All participants use a mix of almost only minimizing and antimatching.

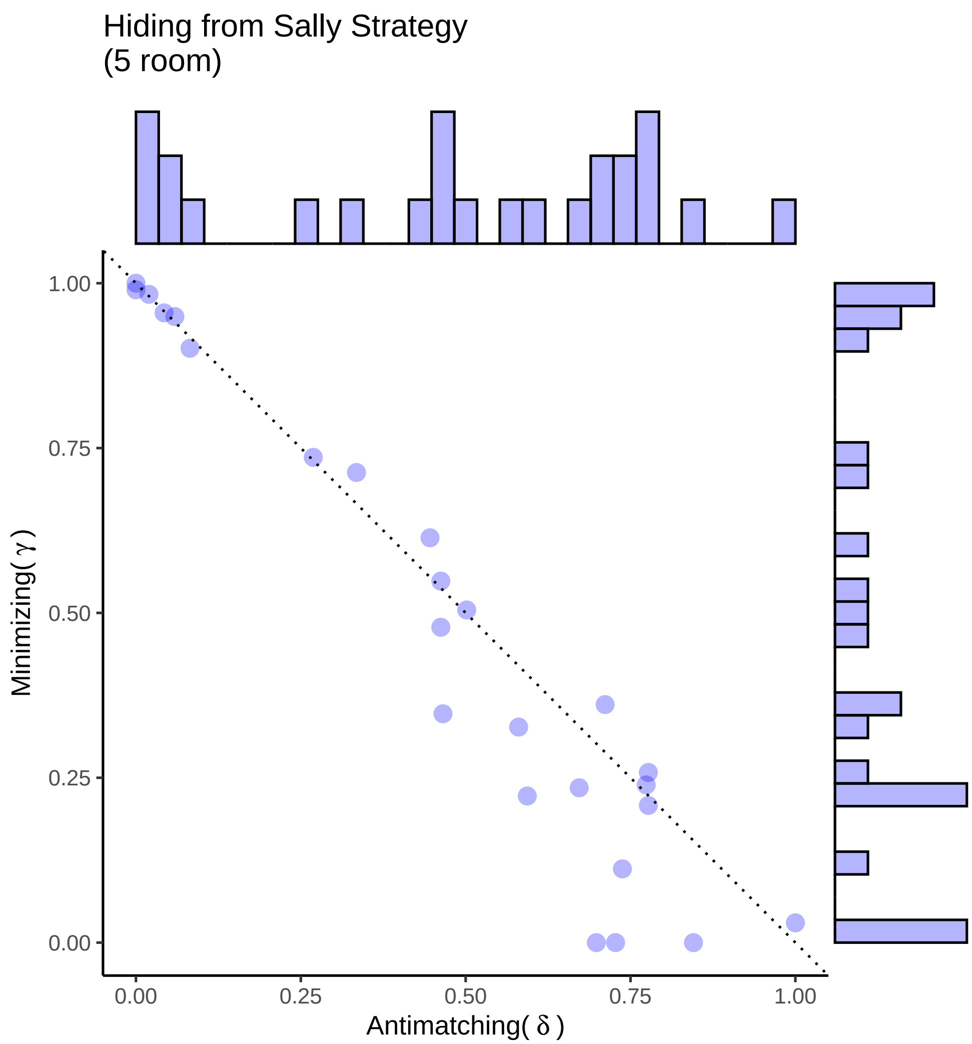

Figure 2.18: Strategy space of participant hides against Sally. Bars along margin count participants using particular value of γ or δ. Participants used a varied mix of minimizing and antimatching strategies. Model still shows relatively low error as strategies are clustered around the line $\gamma + \delta = 1$.

2.2.7 Comparing Hiding and Seeking (Five-Room)

The two-room data showed little difference in strategy mix between hiding and seeking. That is, strategy space plots expressing seeking behaviour as *maximizing* + *matching* looked nearly identical to plots expressing hiding behaviour as *minimizing* + *antimatching*. This is not the case for the five-room data (Figures 2.16 and 2.18). Here, we see a trend towards the non-optimal strategy during seeking, and a trend to the optimal strategy during hiding.

In general, participants became more optimal in their strategy choice when hiding, rather than seeking, against Sally. Like the two-room condition, a portion of the participants had a nearly identical strategy mix between hiding and seeking. Unlike the two-room condition, those who did not use identical strategies between hiding and seeking were almost always dramatically more optimal when hiding (Figure 2.19). A boxplot of the same data reveals the same trends, with average use of the optimal strategy notably higher in the hiding compared to the seeking condition (Figure 2.20).

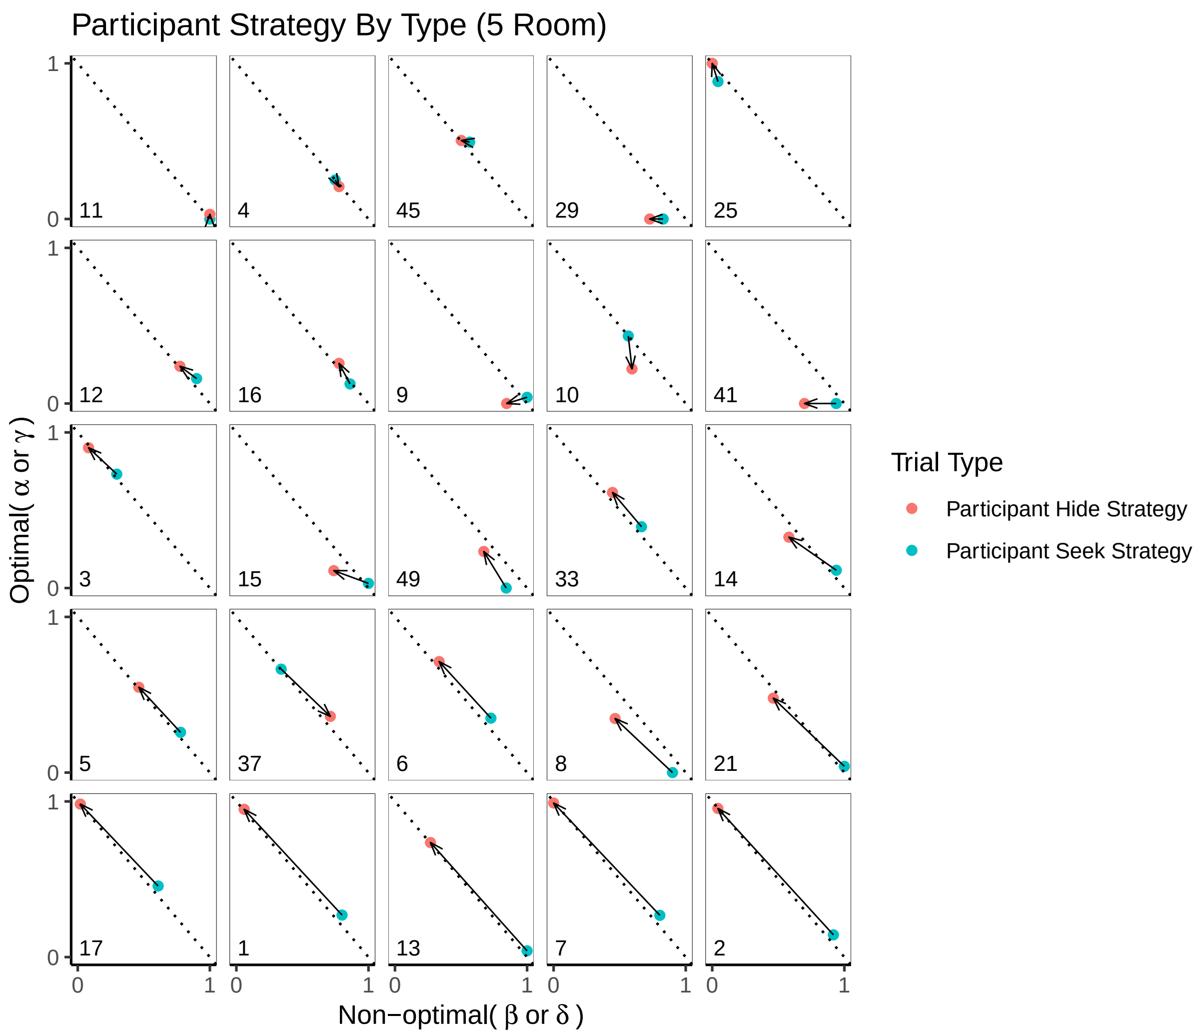

Figure 2.19: Each participant's strategy mix for seeking (α, β) hiding (γ, δ) are plotted together, ordered by increasing values of strategy change. Arrow points from seek strategy to hide strategy. Of those participants who change strategy, they change towards a more optimal mix.

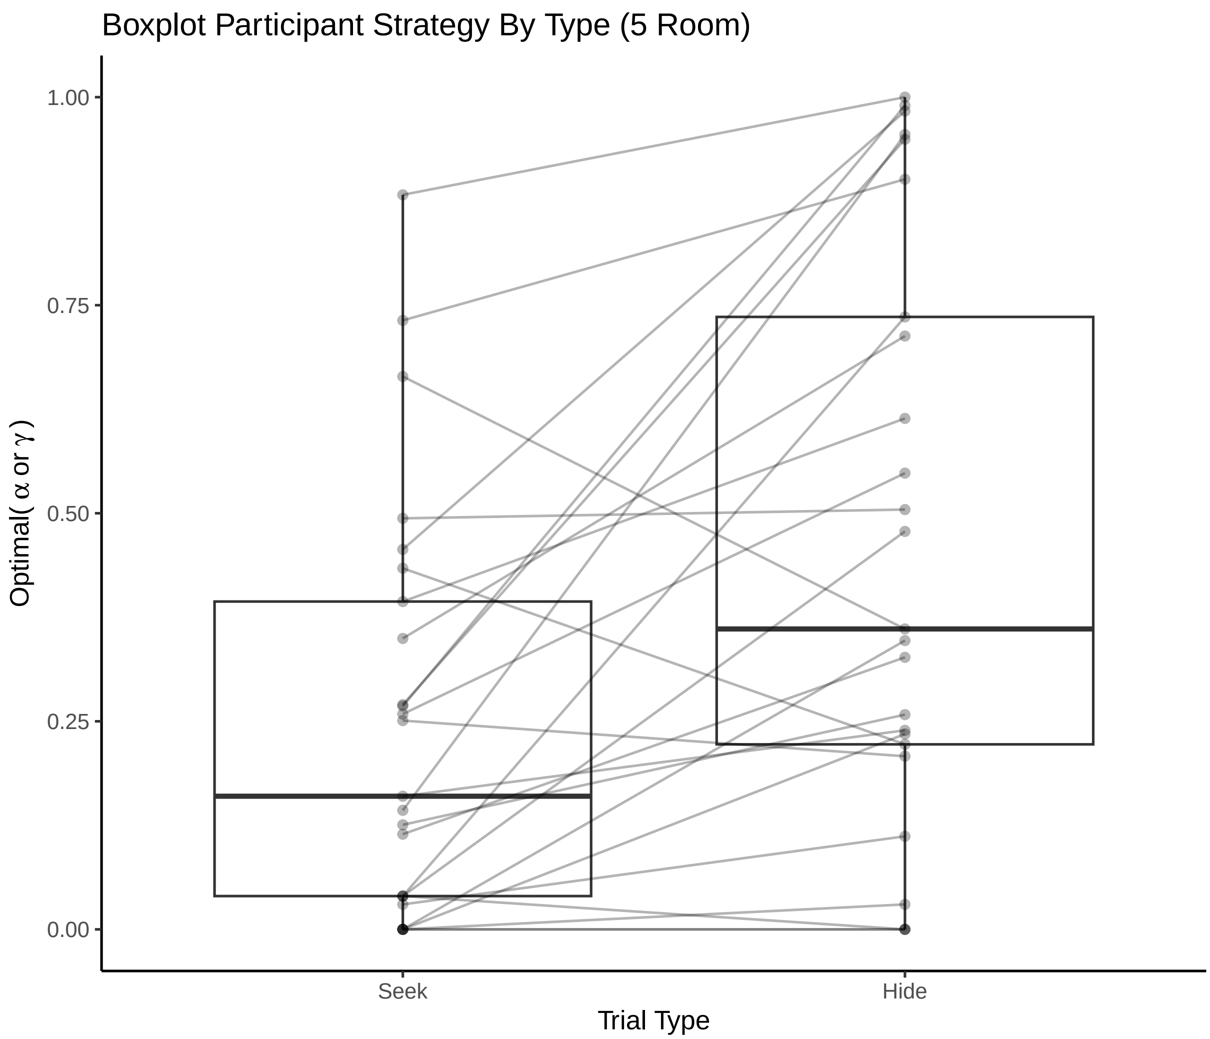

Figure 2.20: Most individual participants change strategies between hiding and seeking, any many change to a more optimal strategy mix when hiding.

2.2.8 Sensitivity Analysis (Five-Room)

While there are differences between hiding and seeking in the five-room condition that were not found in the two-room, we do not believe this to be evidence against the proposed model. We claim only that *maximizing* + *matching* should account for participant seeking strategy, and *minimizing* + *antimatching* should account for participant hiding strategy. This is generalized to *optimal* + *non-optimal*. If, in aggregate, participants use a different mix of optimal and non-optimal strategies depending on the context (pursuit vs avoidance), it does not necessarily mean that the model is accounting for more or less of the participant strategy. This should be decided by analysing RSS error relative to other models. Thus, we perform a second sensitivity analysis, this time on the five-room hiding data. The plot presenting the strategy space for participant hides against Sally is repeated four times using random strategy vectors instead of the minimizing ($\vec{n}$) and antimatching ($\vec{a}$) strategies of our model (Figure 2.21). All four random strategy spaces presented fail to explain participant behaviour effectively, resulting in few strategies lying on or near the line $\gamma_r + \delta_r = 1$ and inflated Mean RSS values ranging between *0.292 and 0.419*. Repeating this procedure 1000 times and plotting Mean RSS draws the same conclusion: $\vec{h} = \gamma\vec{n} + \delta\vec{a}–\vec{\zeta}$ is a more reliable model than $h = \gamma_r r\vec{1} + \delta_r r\vec{2}–\vec{\zeta_r}$ (Figure 2.22).

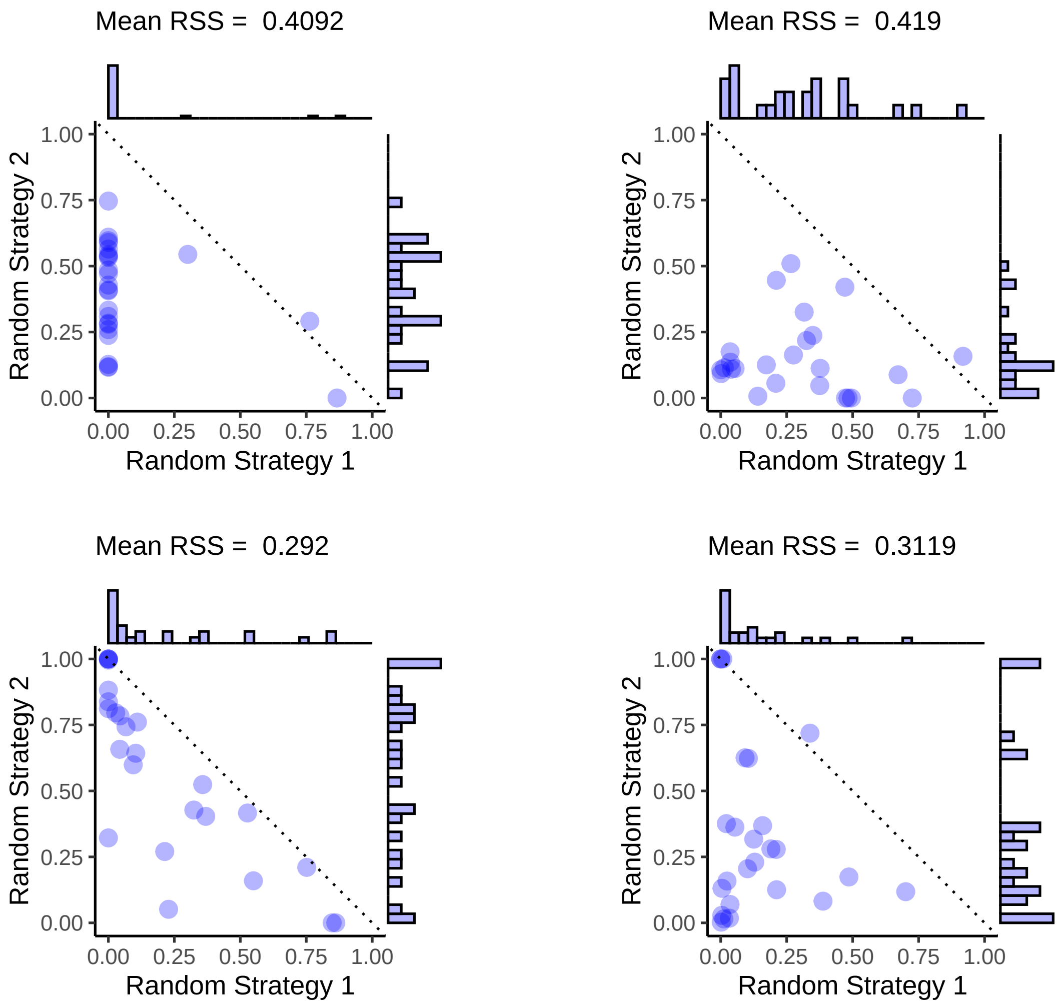

Figure 2.21: Plotting participant hides against sally in a strategy space using two random strategy vectors instead of minimizing and antimatching strategies results in a poor model fit. Many points are in region of high error, reflected in high Mean RSS values. Figure contains four instances of this plot, each with two new randomly generated strategy vectors.

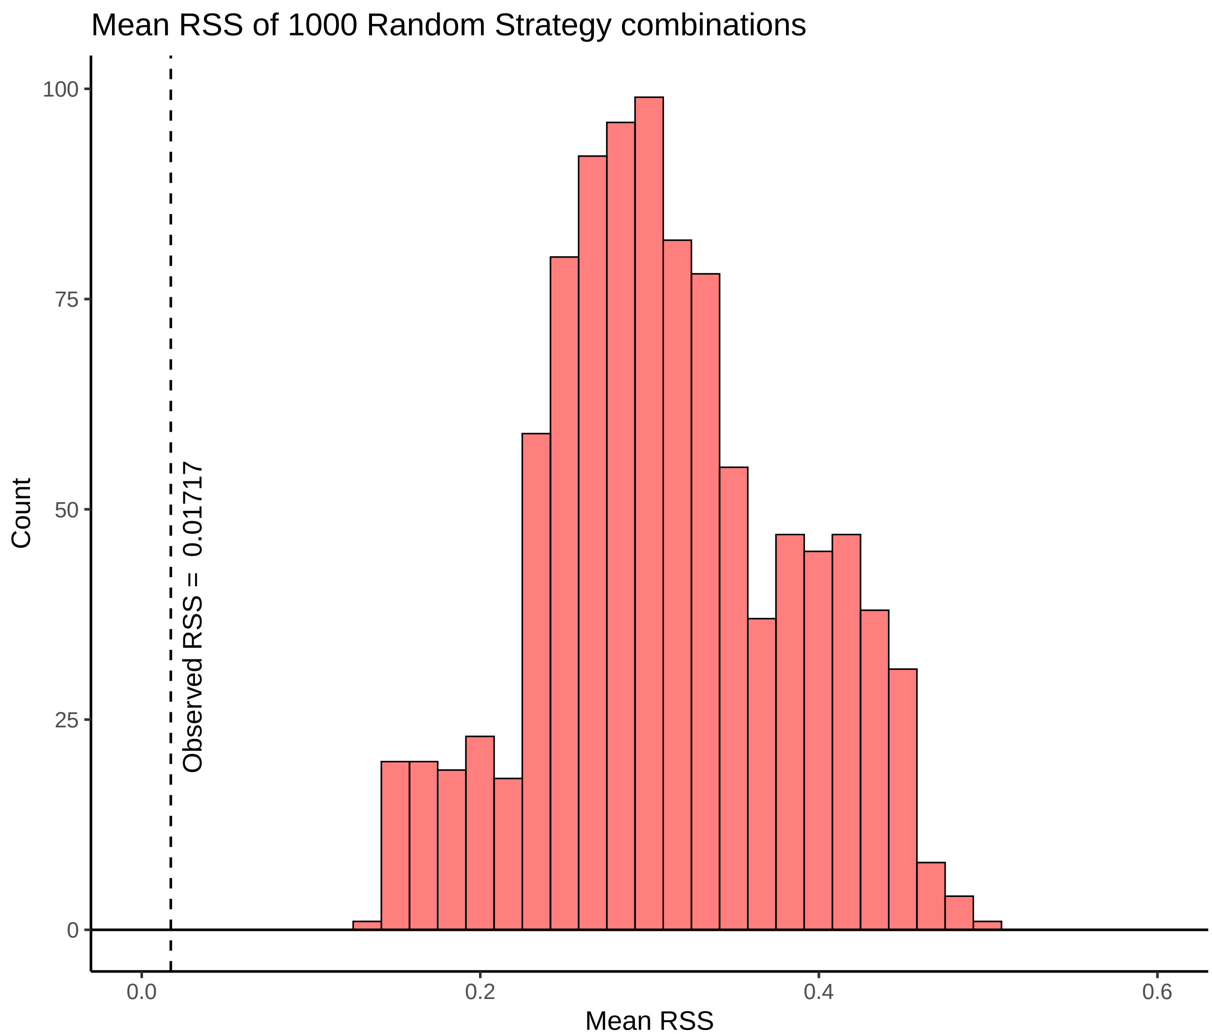

Figure 2.22: Fitting participant data to a strategy space of two random strategy vectors 1000 times. Figure shows histogram of resulting error. Our proposed model dramatically outperforms the 1000 tested here. This implies that our model is of meaningful accuracy, and our methods used are not biased to result in low any for any arbitrary model it is given.

2.2.9 Comparing Two-Room to Five-Room

In our analysis, we are essentially mapping n-dimensional data (where n is the number of rooms) to a two-dimensional strategy space. Therefore, we would expect to see participant strategy mixes well represented by the two-dimensional linear combination of maximizing/matching or minimizing/antimatching strategies when we are already working with two-dimensional data. This is not guaranteed to be true, however, since we constrain our strategy coefficients $(\alpha, \beta, \gamma, \delta)$ to be between 0 and 1. This is demonstrated by the poor representation of participant strategies when using a random distributions in our two-room sensitivity analysis.

We observe in both the two-room and five-room data that representational accuracy differs between maximizing/minimizing vs matching/antimatching strategies. In both seeking and hiding, the optimal strategy is easy to implement behaviorally – select the most/least likely room exclusively. In contrast, matching/antimatching strategies require a more nuanced behavioural pattern where participants need to ensure their rooms selections follow in close proportion to the distribution they wish to produce in aggregate. We should expect this to result in greater error (distance from $\alpha + \beta = 1$ or $\gamma + \delta = 1$) when a strategy mix is predominately of the more complicated matching/antimatching strategy. This explains why, in general, participant strategies are less tightly clustered in the bottom right corners of our plots compared the top left corners.

We observe differences in the relationship between hiding and seeking strategy optimality depending on the room condition. Specifically, participants used approximately a 50/50 optimal/non-optimal strategy mix in the two-room condition for both hiding and seeking. In the five-room condition, participants used approximately a 25/75 optimal/non-optimal strategy mix when seeking, but a 50/50 optimal/non-optimal strategy mix when hiding. In addition, all strategy means lie close to the line $\alpha + \beta = \gamma + \delta = 1$ implying our model appropriately accounts for participant hiding and seeking behaviour in both two-room and five-room conditions (Figure 2.23).

This can also be interpreted in terms of a correlation. Participant use of an optimal strategy is simultaneously negatively correlated to participant use of a non-optimal strategy. Thus, results of all participant hiding and seeking behaviour, across both conditions can be accurately represented with boxplots of participant use of optimal strategy (Figures 2.12 and 2.20).

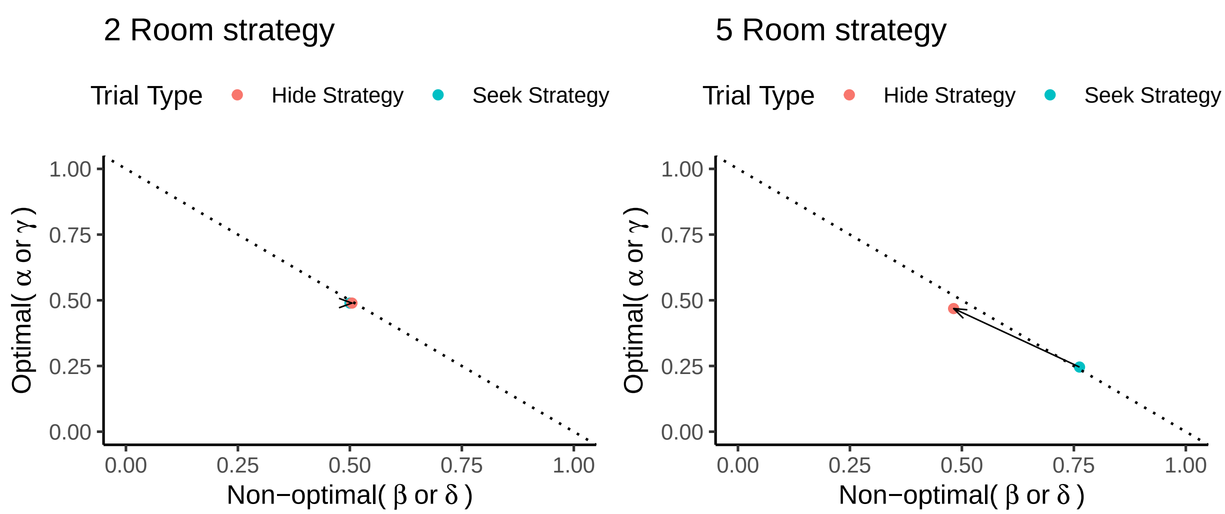

Figure 2.23: Mean strategy change from seek to hide split by room count (2-room on left panel, 5-room on right panel). In five-room condition, participants used a less optimal strategy when seeking than they used when hiding. This difference was not seen in the two-room condition, where participants used the same mix 50/50 mix of optimal/non-optimal strategy when hiding as they did with seeking. In addition, all strategy means lie close to the line $\alpha + \beta = \gamma + \delta = 1$ implying our model appropriately accounts for participant hiding and seeking behaviour in both two-room and five-room conditions.

Representational Complexity Hypothesis

What accounts for the difference in strategy mix found in the five-room condition, but not the two-room condition? Consider the representational complexity of a two-room vs a five-room hide-and-seek task. In general, visualization becomes harder in higher dimensions: a cube is easier to imagine than a tesseract (4D-cube), and a tesseract easier than a penteract (5D-cube). Perhaps individuals have a bias to probability antimatch when hiding, analogous to matching when seeking. But when this strategy becomes too difficult to compute in higher dimensions, perhaps people resort to minimizing; a strategy that, while easier, also happens to be optimal. We denote this as the *representational complexity hypothesis*. This can be tested by measuring how improvement from seeking to hiding is influenced by the dimensionality of the given problem. We explore this in Study 3. Two-dimensional space is the simplest space to compute a histogram vector reflection, where any histogram has a valid reflection (by swapping bin values), so the cost of computing an opposite should be essentially none, and would only increase with added dimensions. Except, perhaps, until the problem becomes so complicated that people would internally simplify it to a lower dimension. Consider a house of 1000 rooms. Participants would likely consider groups of adjacent rooms rather than each individual room on each trial.

2.2.10 Special case opponent strategies

While Sally was of primary theoretical interest in this study, we also recorded participant play against the uniform distribution, (Kala) and a maximally predictable distribution, where one room has 100% weight (Bo). As expected with our model, the uniform distribution acts an equilibrium, where there is no reason to deviate from this strategy when hiding or seeking against an opponent. Most participants appear to be following this pattern (Figure 2.24). Recall, any other strategy is equally as viable given the opponent is static in strategy use. Any participants who use anything other than a uniform strategy are at no disadvantage.

The distribution represented by the character Bo, where one room has probability 100% and the rest have 0%, is the one distribution used in this experiment that did not have valid histogram vector reflection in the five-room condition. Similar to the uniform case, any strategy exclusively over the equally 0% rooms are theoretically equivalent. Regardless, Bo acted as a useful assurance that participant understood the experimental protocol, by always selecting the 100% room when seeking, and never selecting it when hiding. Participants appear to have followed this pattern (Figure 2.25). Aside from a few potentially errant clicks, participants never selected a room that is guaranteed to be strictly worse than picking an optimal room, implying all participants understood and followed the experimental protocol.

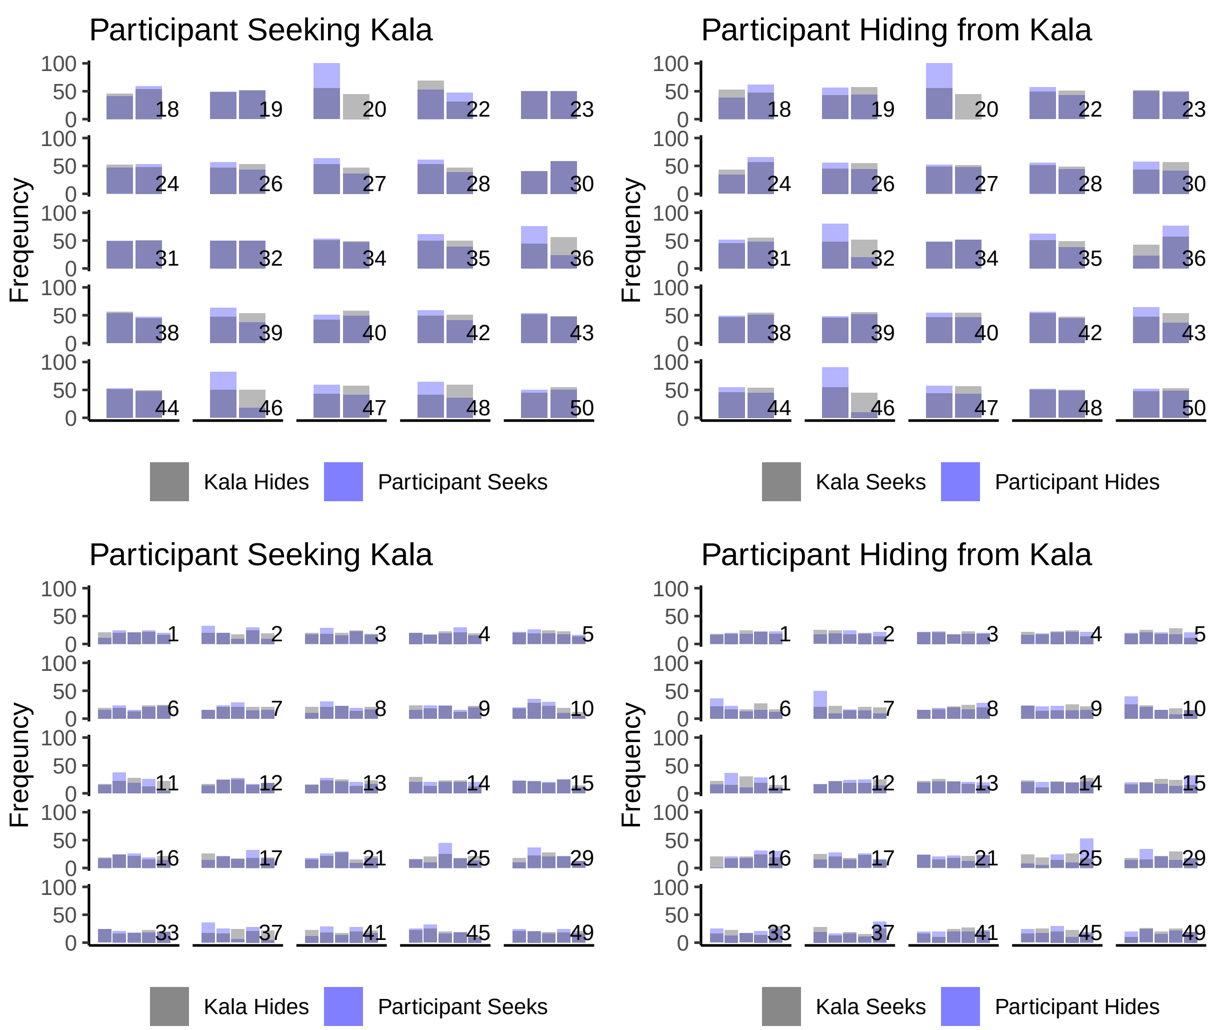

Figure 2.24: All participant rounds against Kala. Seek data in left column, hide in right column. Two-room on top row, five-room on bottom row. Most participants use an approximately uniform strategy as expected, though any other strategy has a theoretically equivalent payoff for the participant

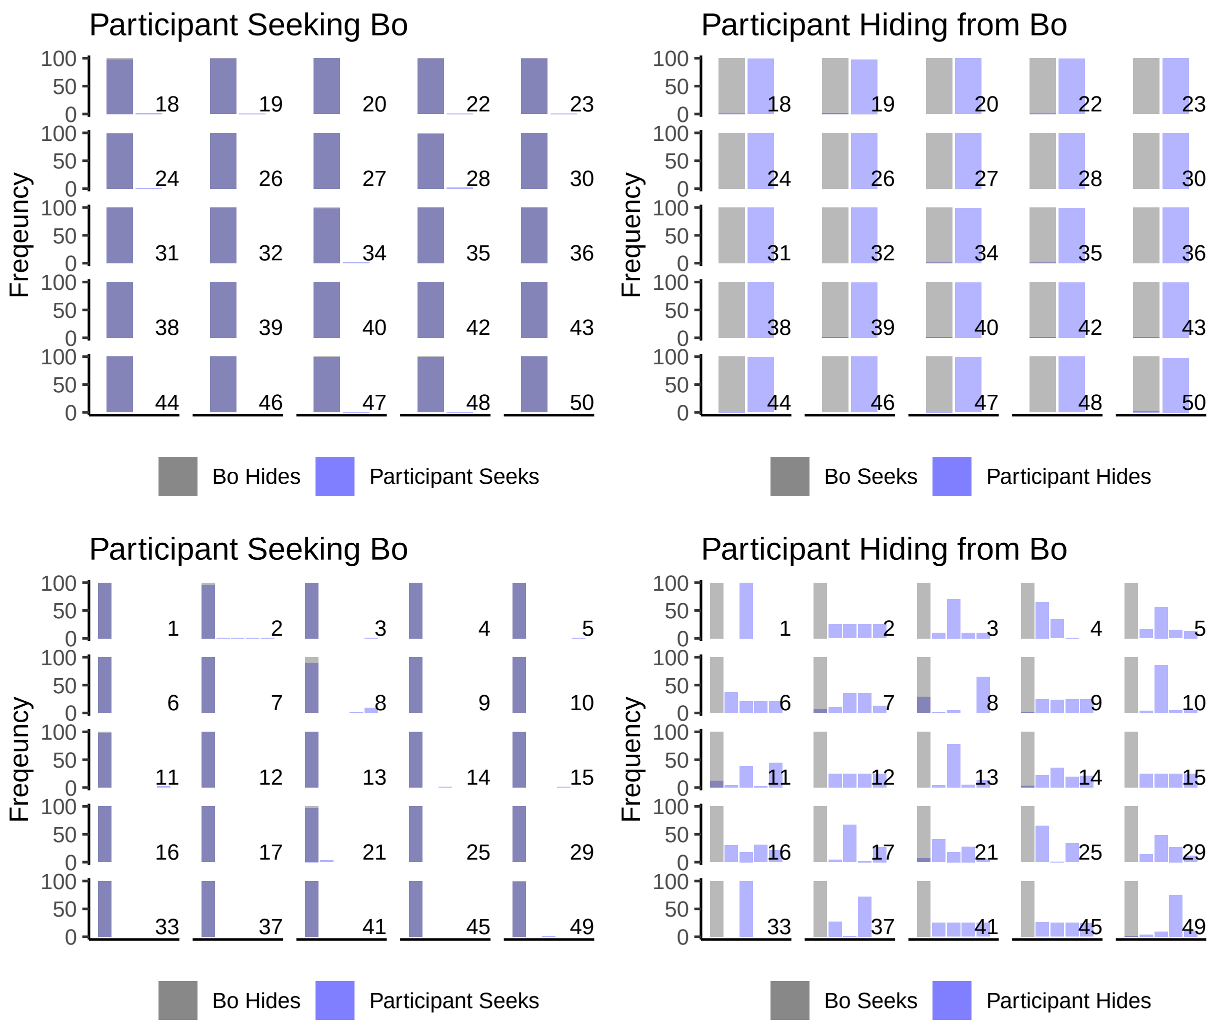

Figure 2.25: All participant rounds against Bo. Seek data in left column, hide in right column. Two-room on top row, five-room on bottom row. Participants behaved optimally aside from a few apparently erroneous clicks, ensuring the experiment was understood. Hiding behaviour in the five-room condition was least consistent, and is also the only dataset incompatible with our vector reflection model.

2.2.11 Demographics

If it were the case that individuals have a natural propensity to probability antimatch until it becomes too difficult, how would an increase in mathematical ability impact this tendency? Would someone with more mathematical reasoning experience be better able to compute opposite probability distributions, and thus be more likely to continue antimatching in higher dimensional problems? Or would they instead be better able to recognize the optimal strategy is to purely maximize or minimize? Though our sample size is likely too small to provide more than speculation, it may be that students who study math are better at recognizing the optimal choice of maximizing when seeking, although this benefit is mitigated when hiding, since everyone else who does not recognize the optimal strategy defaults to it (Figure 2.26). This is perhaps because conforming to the antimatching bias becomes too difficult.

Following the hide-and-seek gameplay, participants answered a short survey that asked which faculty and program they were enrolled in. Students from the faculty of Mathematics potentially utilized a more optimal seeking strategy against Sally. Though across all hiding rounds against Sally, mathematics students no longer stand out. It is most likely that there are too few participants for this result to be robust with n = 9, 21, 2, 8, 10 for Applied Health Sciences, Arts, Engineering, Mathematics, and Science, respectively.

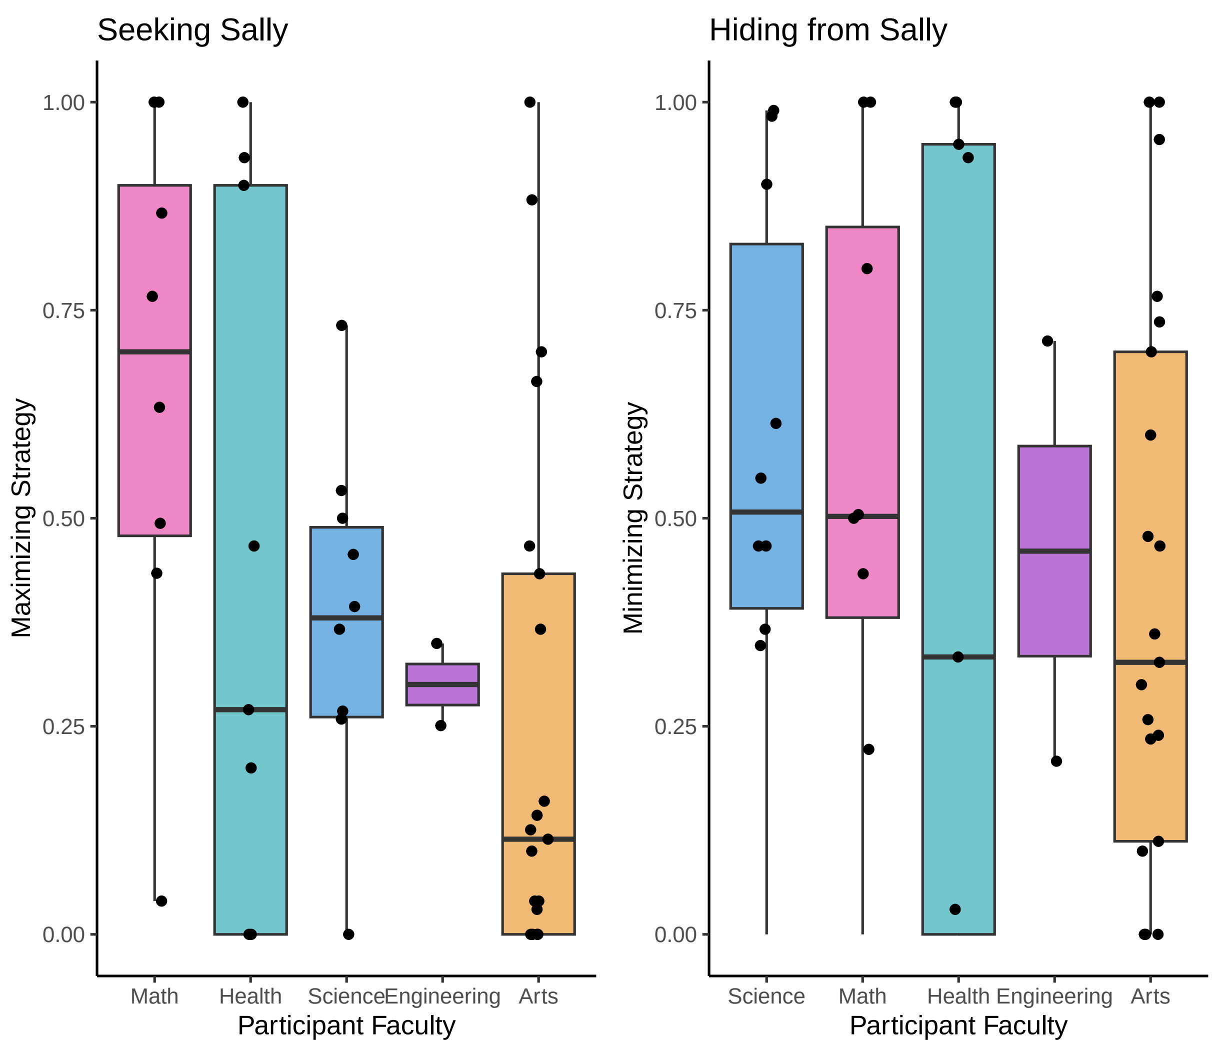

Figure 2.26: Mathematics students appear to use a more optimal strategy when seeking, though this difference is not apparent in hiding.

Chapter 3

Study 2: Online Replication

3.1 Methods

3.1.1 Task Environment

The COVID-19 pandemic prevented further data collection using the in-lab PsychoPy implementation from Study 1. We instead approximated Study 1 using JsPsych [4]. This JavaSpript implementation allowed for online data collection. This study had identical stimuli and procedures as Study 1, with the exception of slight visual differences in the dollhouse and background artwork (Figure 3.1). Since we had no control over the device participants used to complete the task, we had participants click a button with their cursor to advance to the next trial rather than hit the space bar as in Study 1. Both methologies achieve the the same goal of preventing participants from keeping the cursor over a single room and automatically repeating the same room selection. Finally, we also did not enforce the participants' device screens to encompass a specific visual angle. Stimuli and procedures are otherwise identical to Study 1.

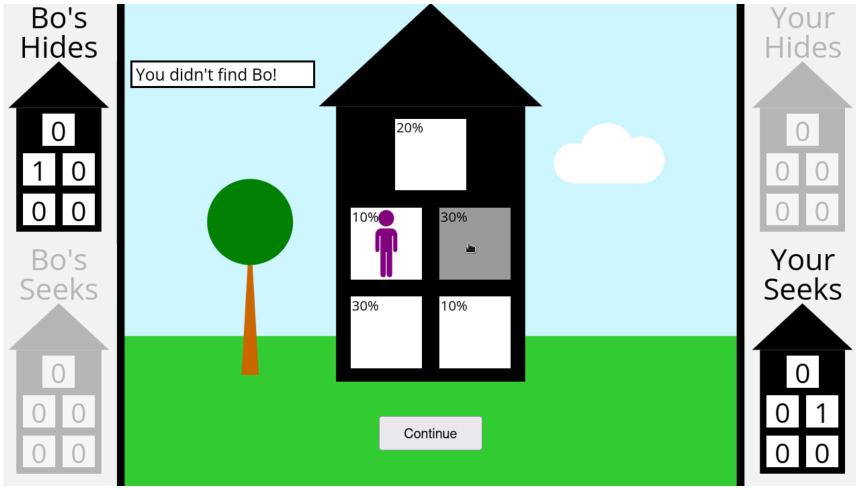

Figure 3.1: Participant view of screen in the five-room condition. Materials differ slightly from Study 1. There are slight differences in background artwork, and participants are required to click the 'continue' button instead of the space bar to move to the next trial. Procedure is otherwise identical to Study 1.

3.1.2 Participants

We recruited University of Waterloo students (N=54, 44 females) to participate in exchange for course credit. Participants were assigned randomly to the five-room or two-room condition. The experiment took approximately 21 minutes to complete.

3.1.3 Data Analysis

Data analysis of this Study is identical to that of Study 1. This includes investigation of participant strategy mix against Sally for both seeking and hiding in both two- and five-room conditions, a sensitivity analysis of the vector reflection model, and the relationship between academic major and strategy mix.

3.2 Results and Discussion

3.2.1 Seeking Sally (Two-Room)

As in Study 1, a plot of all participant seek frequencies overlaid with Sally's hide frequencies reveals a spectrum of strict maximizers, strict matchers, and some mixing of the two (Figure 3.2). Converting and plotting each participant into the strategy space of α and β, refines our interpretation (Figure 3.3). Participants employing a mixed strategy appear to be using a mix specifically of maximizing and matching, as all participants lie on or essentially on the line $\alpha + \beta = 1$ implying there is negligible error in this fit (Mean RSS = 0.0009453).

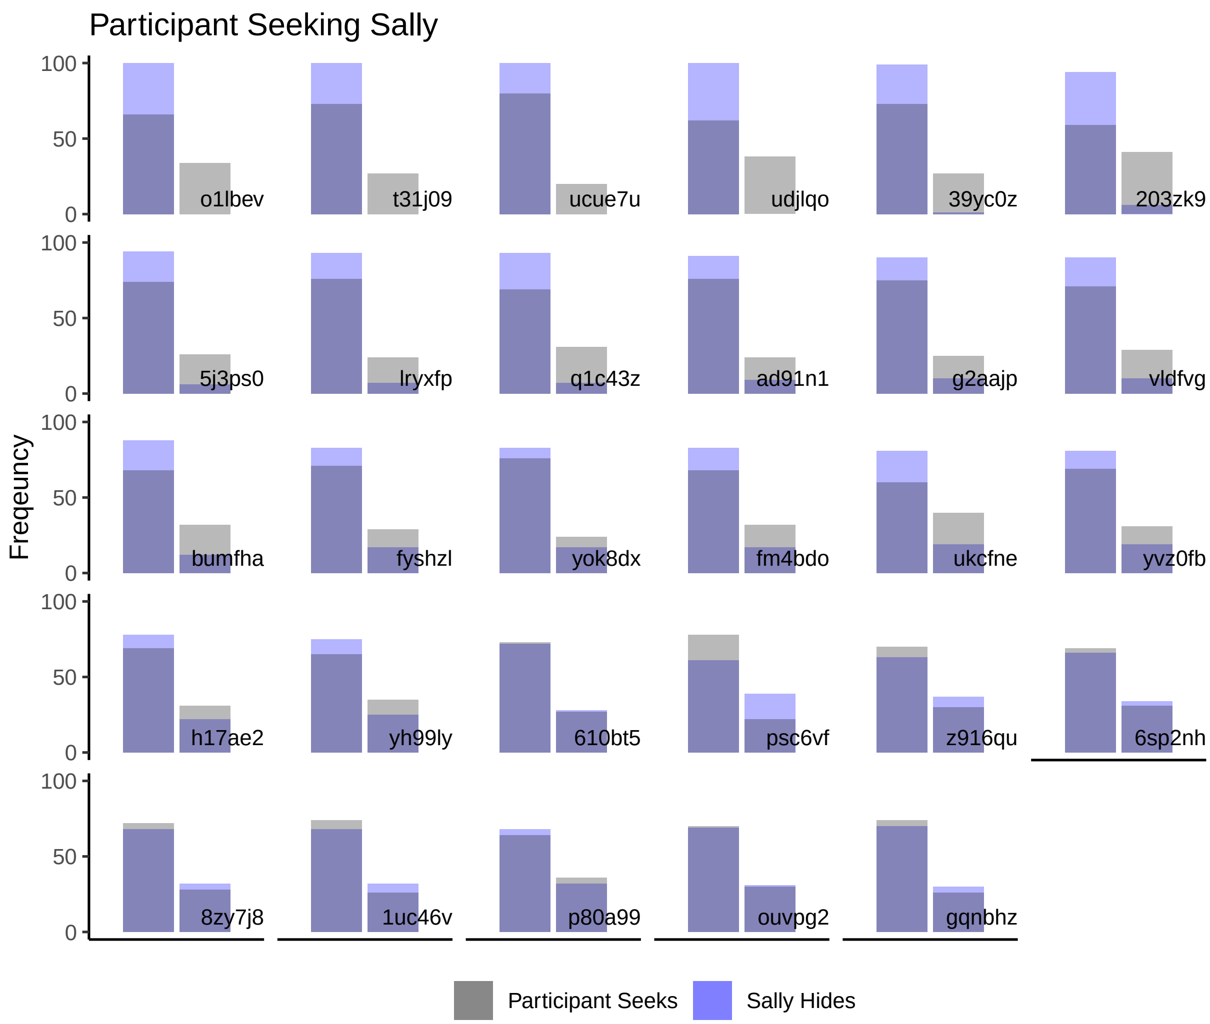

Figure 3.2: Participant seeks (blue) and Sally's hides (grey) sorted by increasing values of β. Histogram bins sorted by decreasing values of Sally's distribution. Blue participant seek bars at 100% indicate perfect maximizing (see top left). Participant seek bars perfectly overlapping Sally's hides indicate perfect matching (see bottom right)

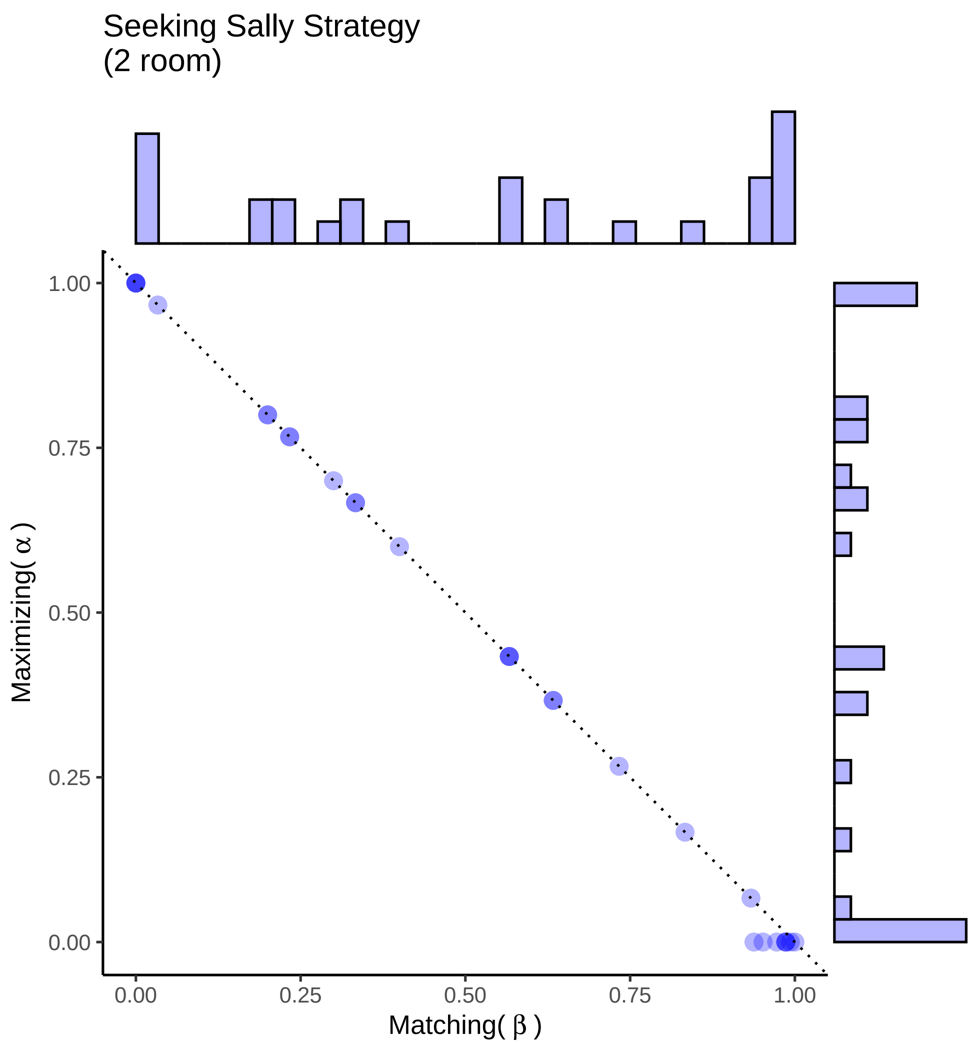

Figure 3.3: Strategy space of participant seeks against Sally. Bars along margin count participants using particular value of α or β. Perfect maximizers use $\alpha = 1$ and $\beta = 0$ (top left). Perfect matchers use $\alpha = 0$ and $\beta = 1$ (bottom right). Participant strategies existing on or near line $\alpha + \beta = 1$ implies near perfect model fit.

3.2.2 Hiding Sally (Two-Room)

Plots of participant hide frequencies over Sally's seek frequencies reveal a pattern similar to that of Study 1, and the seeking data from this Study (Figure 3.4). Some participants behave optimally, minimizing by only hiding in Sally's least searched rooms. Others appear to employ some 'opposite' strategy to Sally's, hiding in Sally's most searched room as frequently as Sally searches the least searched room. Upon viewing this data in the strategy space, we can see the participants used a mix specifically of an optimal minimizing strategy and a non-optimal antimatching strategy, defined here as a 2D histogram reflection (Figure 3.5). All participants were on or near the line $\gamma + \delta = 1$, implying negligible error (Mean RSS = 0.0068193).

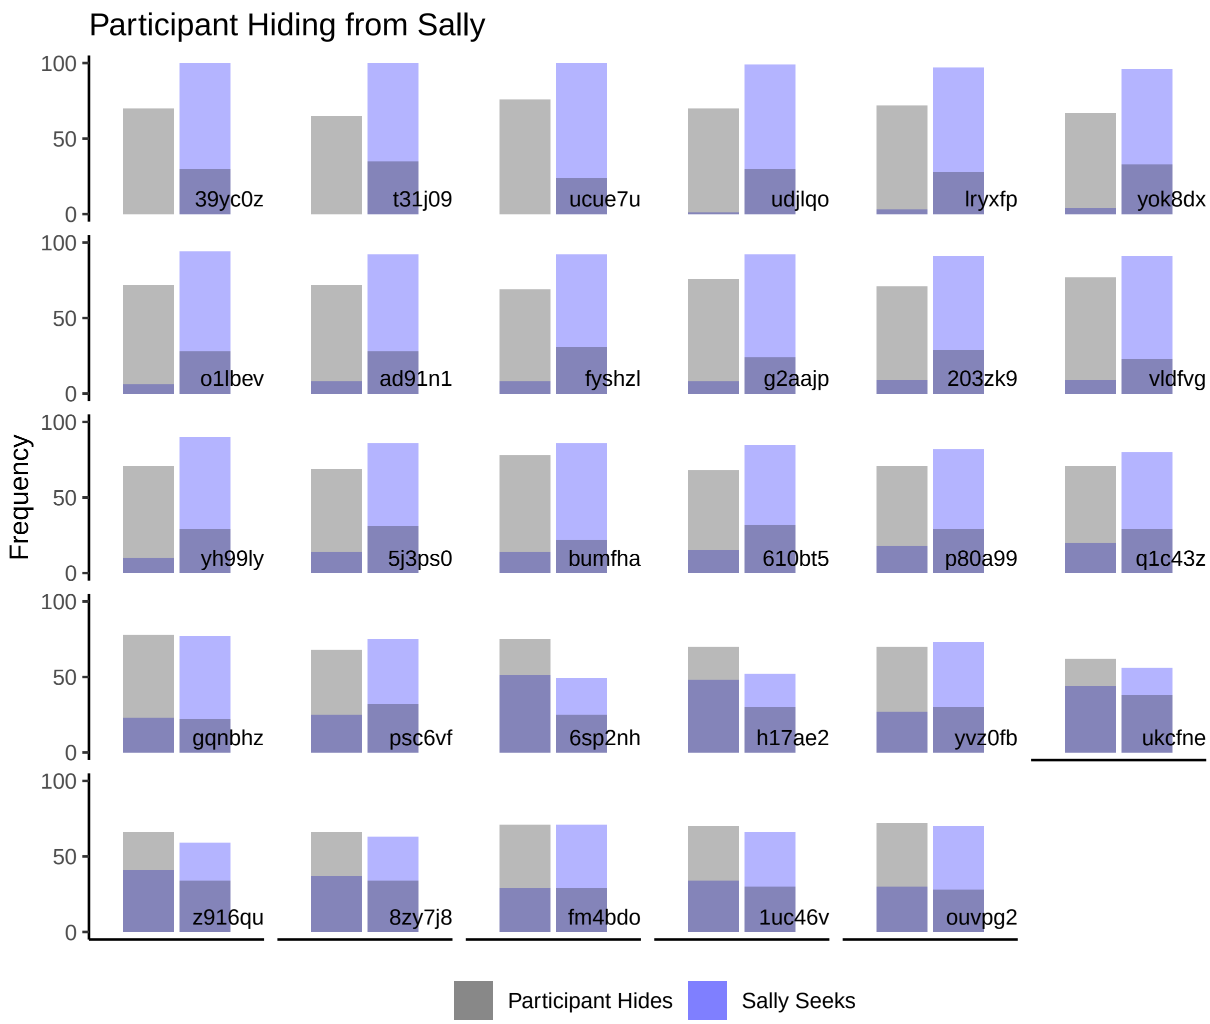

Figure 3.4: Participant hides (blue) and Sally's hides (grey) sorted by increasing values of d (use of antimatching strategy). Histogram bins sorted by decreasing values of Sally's distribution. Blue participant hide bars at 100% indicate perfect minimizing (top left). Participant hide bars at 1 – Sally's seek bars indicate antimatching (bottom right)

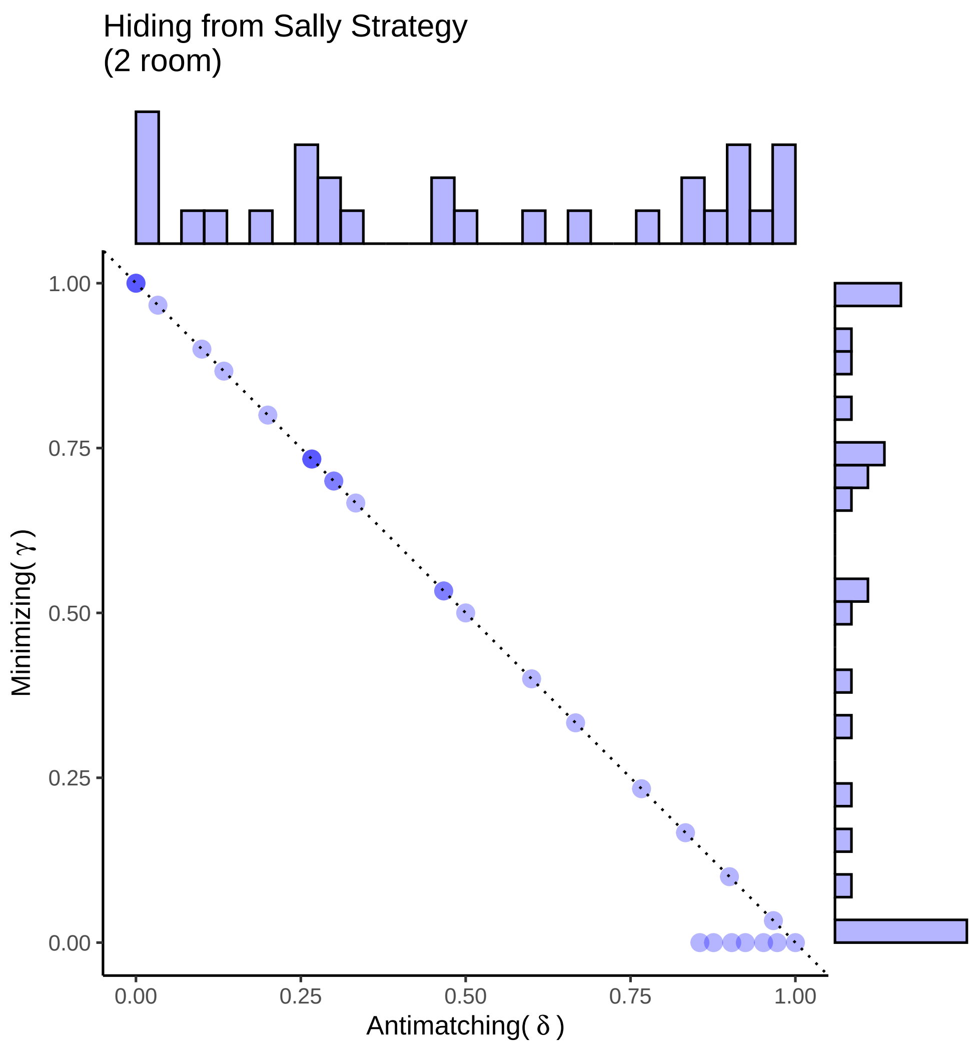

Figure 3.5: Strategy space of participant hides against Sally. Bars along margin count participants using particular value of γ or δ. Perfect minimizers use $\gamma = 1$ and $\delta = 0$ (top left). Perfect antimatchers use $\gamma = 0$ and $\delta = 1$ (bottom right). Participant strategies existing on or near line $\gamma + \delta = 1$ implies near perfect model fit.

3.2.3 Comparing Hiding and Seeking (Two-Room)

Comparing strategy spaces of hiding and seeking behaviours reveal similarities to the seeking strategy space (Figures 3.3 and 3.5). This suggests that, in aggregate, *maximizing* + *matching* describes seeking similarly as *minimizing* + *antimatching* describes hiding. This pattern is consistent with those found in Study 1.

As in Study 1, participants did not become more or less optimal in their strategy choice when hiding instead of seeking against Sally in the two-room condition. A large portion had an identical strategy mix between hiding and seeking, and of those who didn't, no discernible pattern was found to indicate that people generally become more or less optimal when hiding (Figure 3.6). A boxplot of the same data is consistent with this conclusion (Figure 3.7).

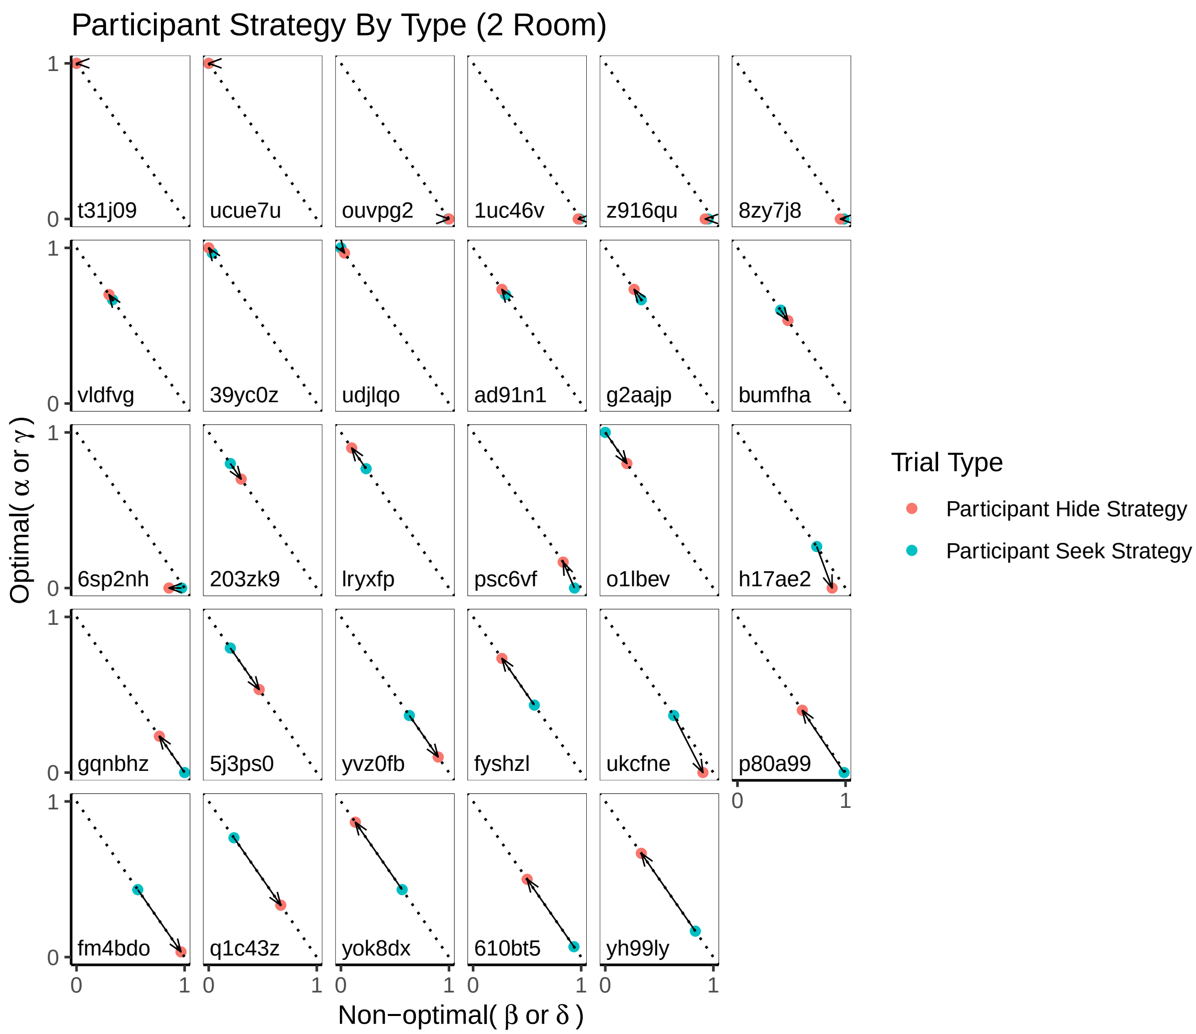

Figure 3.6: Each participant's strategy mix for seeking (α, β) hiding (γ, δ) are plotted together, ordered by increasing values of strategy change. Arrow points from seek strategy to hide strategy. Most participants have little change. Those who do, change in a non-predictable direction

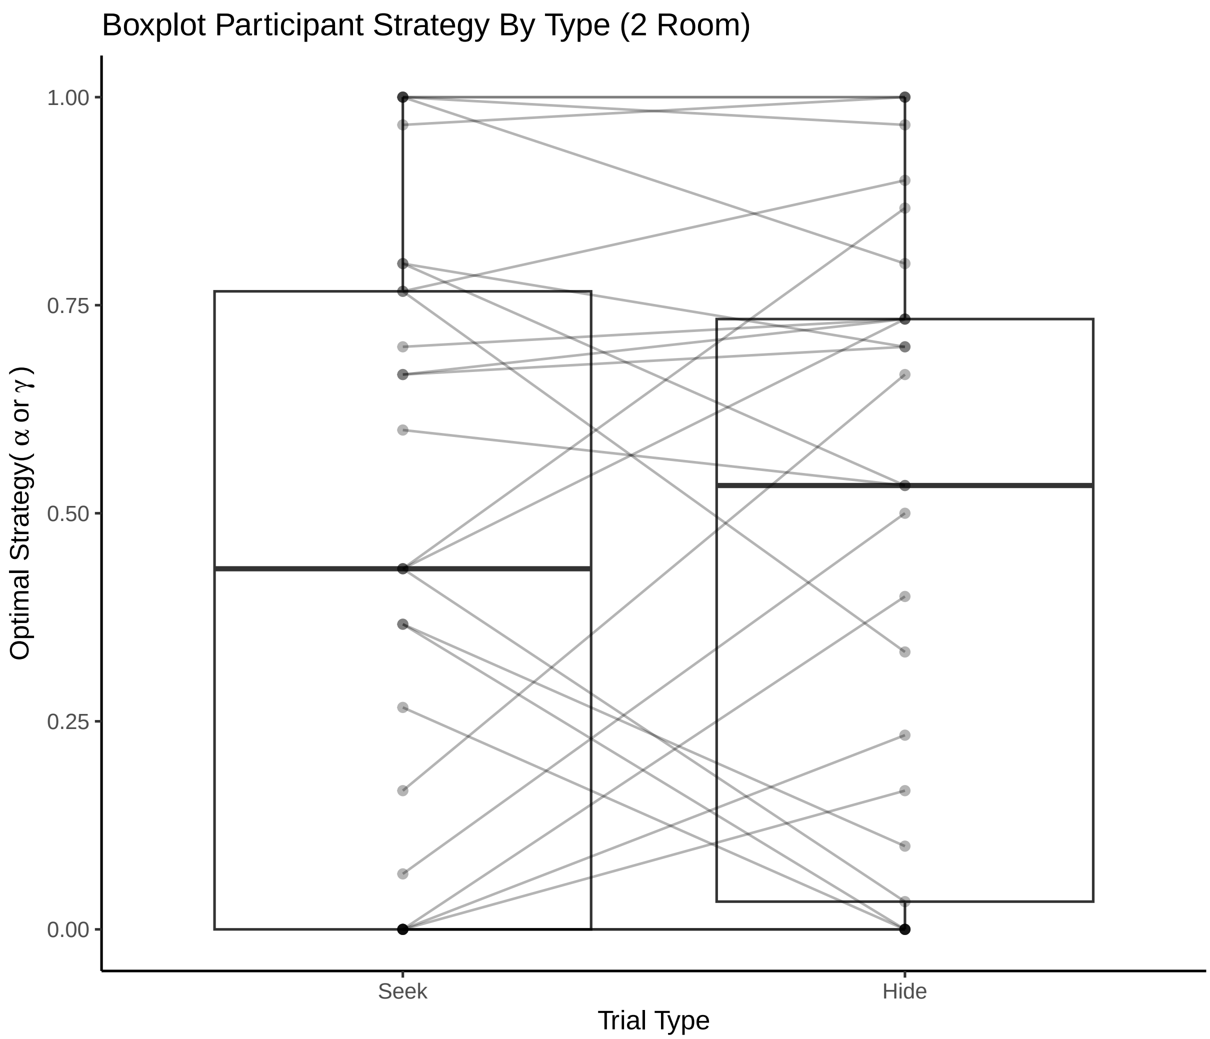

Figure 3.7: Most individual participants use a similar strategy mix between hiding and seeking. Some change strategies between hiding and seeking exist, though in an unpredictable direction and average out in aggregate. The result is a roughly equivalent use of the optimal strategy in hiding and seeking.

3.2.4 Sensitivity Analysis (Two-Room)

Performing a sensitivity analysis on our model of participant hiding by replacing the minimizing and antimatching strategy vectors with two random histogram vectors indicates (1) that $\vec{h} = \gamma\vec{n} + \delta\vec{a} – \vec{\zeta}$ is a suitable model, and (2) that using the bvls methods to present and analyze data is not unreasonably forgiving to any arbitrary model it is given. The plot presenting the strategy space for participant hides against Sally is repeated four times using random strategy vectors instead of the minimizing ($\vec{n}$) and antimatching ($\vec{a}$) strategies of our model (Figure 3.8). All four random strategy spaces presented fail to explain participant behaviour effectively, resulting in few strategies lying on or near the line $\gamma_r + \delta_r = 1$ and inflated Mean RSS values ranging between 0.00134 and 0.04726. Repeating this procedure 1000 times and plotting Mean RSS draws the same conclusion: $\vec{h} = \gamma\vec{n} + \delta\vec{a} – \vec{\zeta}$ is a more reliable model than $\vec{h} = \gamma_r r\vec{1} + \delta_r r\vec{2} – \vec{\zeta_r}$ (Figure 3.9). This test may be redundant given the high degree of similarity between the findings of Studies 1 and 2, and the large sample size used in the sensitivity analyses. However, it's inclusion serves as a sanity check demonstrating the stability of our findings across repeated analyses on new data.

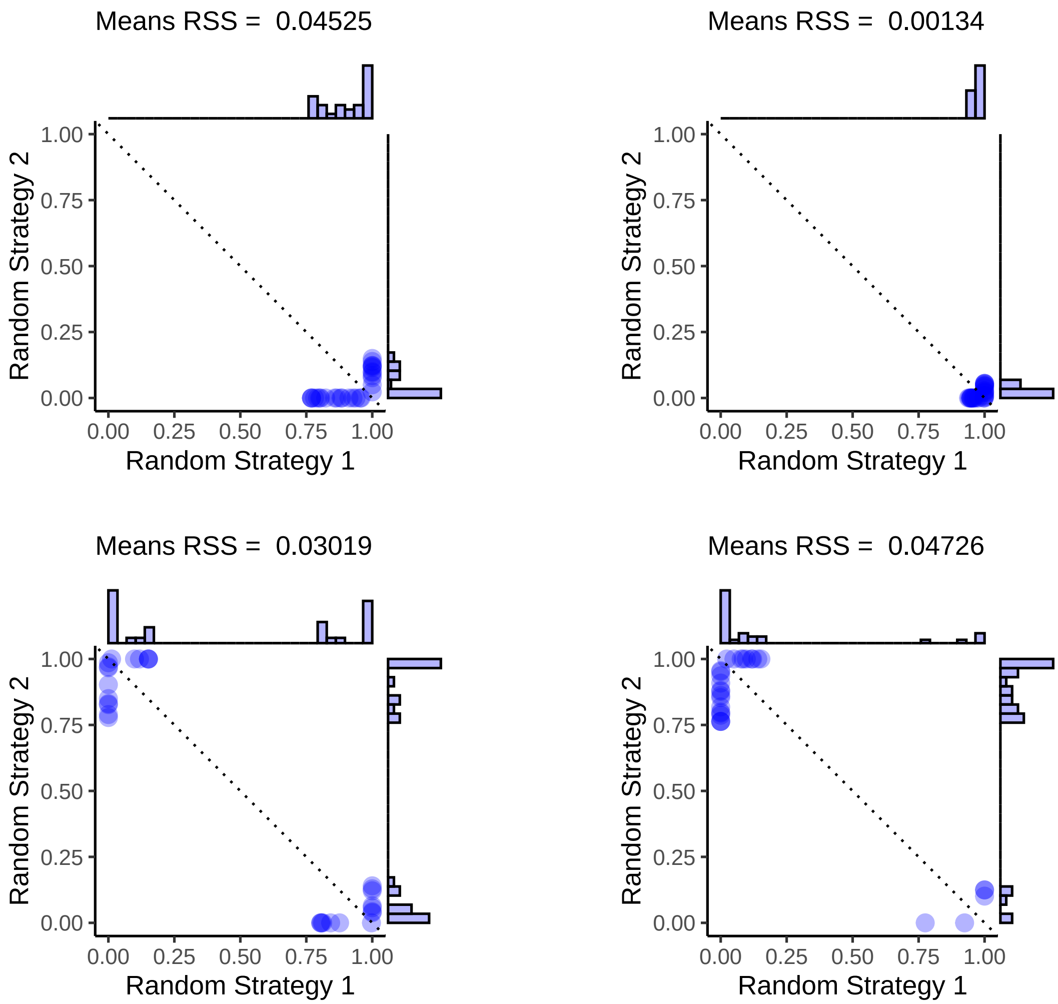

Figure 3.8: Plotting participant hides against sally in a strategy space using two random strategy vectors instead of minimizing and antimatching strategies results in a poor model fit. Many points are in region of high error, reflected in high Mean RSS values. Figure contains four instances of this plot, each with two new randomly generated strategy vectors.

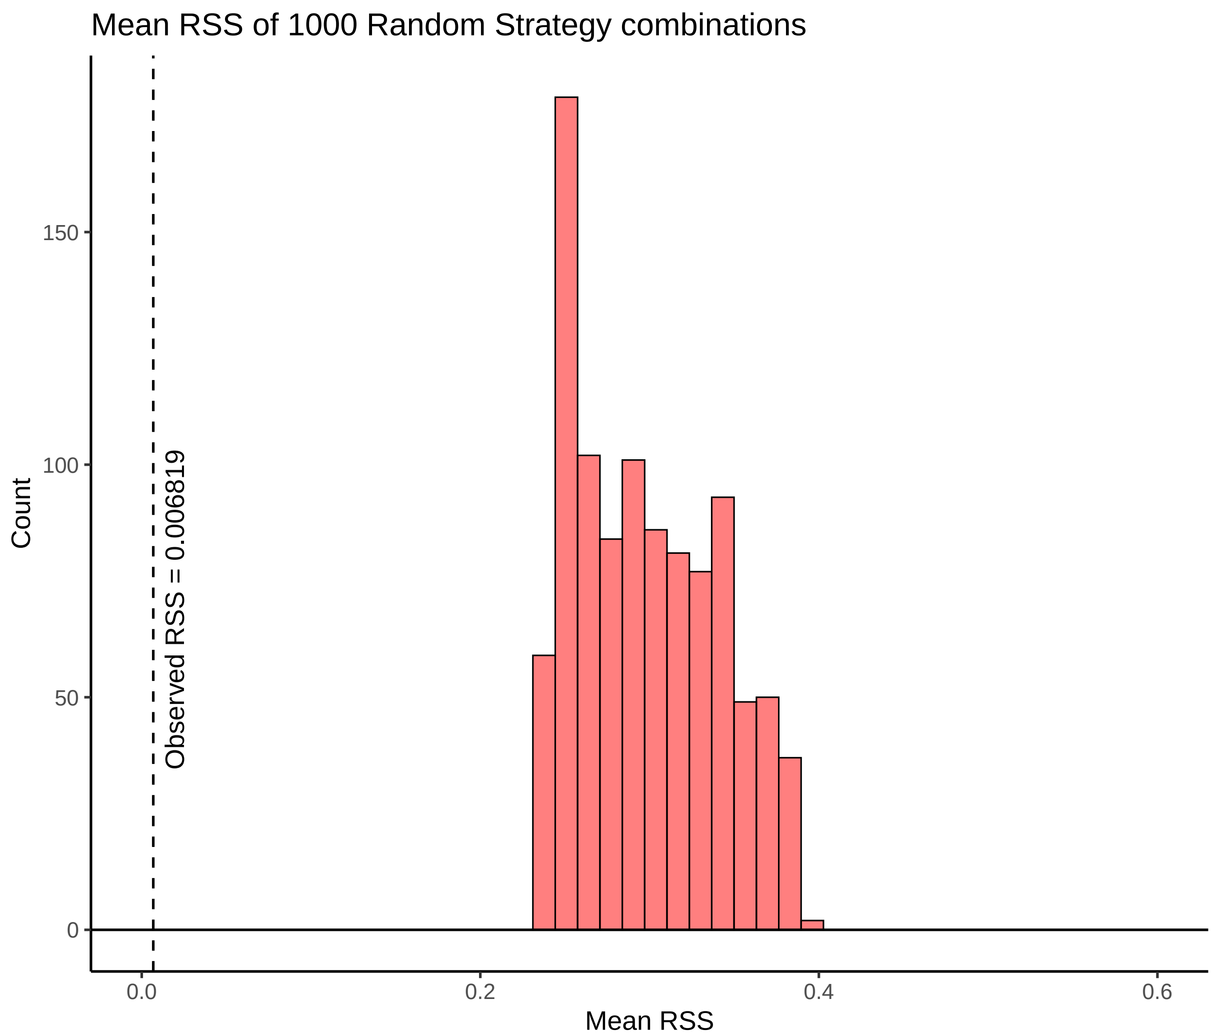

Figure 3.9: Fitting participant data to a strategy space of two random strategy vectors 1000 times. Figure shows histogram of resulting error. Our proposed model dramatically outperforms the 1000 tested here. This implies that our model is of meaningful accuracy, and our methods used are not biased to result in low any for any arbitrary model it is given.

3.2.5 Seeking Sally (Five-Room)

A plot of all participant seek frequencies overlaid with Sally's hide frequencies reveals a spectrum of nearly strict maximizers, nearly strict matchers, though many are using a strategy mix of predominantly matching (Figure 3.10). This is consistent with Study 1. Converting and plotting each participant into the strategy space of α and β, refines our interpretation (Figure 3.11). Participants appear to be using a mix of maximizing and matching but strongly favour matching. All participants lie in close proximity to the line $\alpha + \beta = 1$ resulting in minimal error (Mean RSS = 0.0089585). This error in the five-room case is larger than in the two-room, though some increase in error can be expected as data increases in dimension.

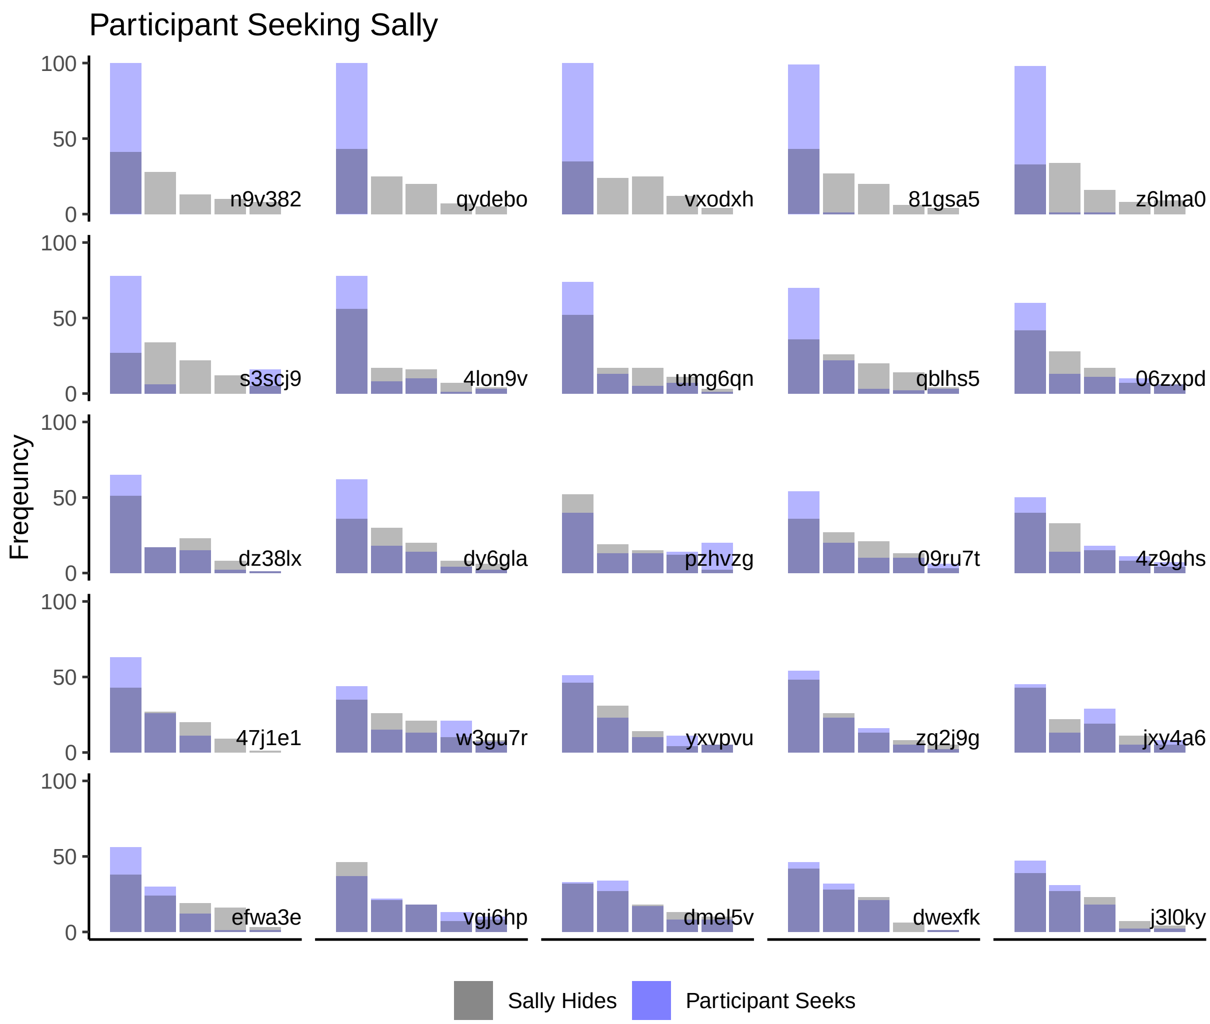

Figure 3.10: Participant seeks (blue) and Sally's hides (grey) sorted by increasing values of β. Histogram bins sorted by decreasing values of Sally's distribution. All participants use a mix of almost only maximizing and matching, though most tend toward a mix containing more matching than maximizing.

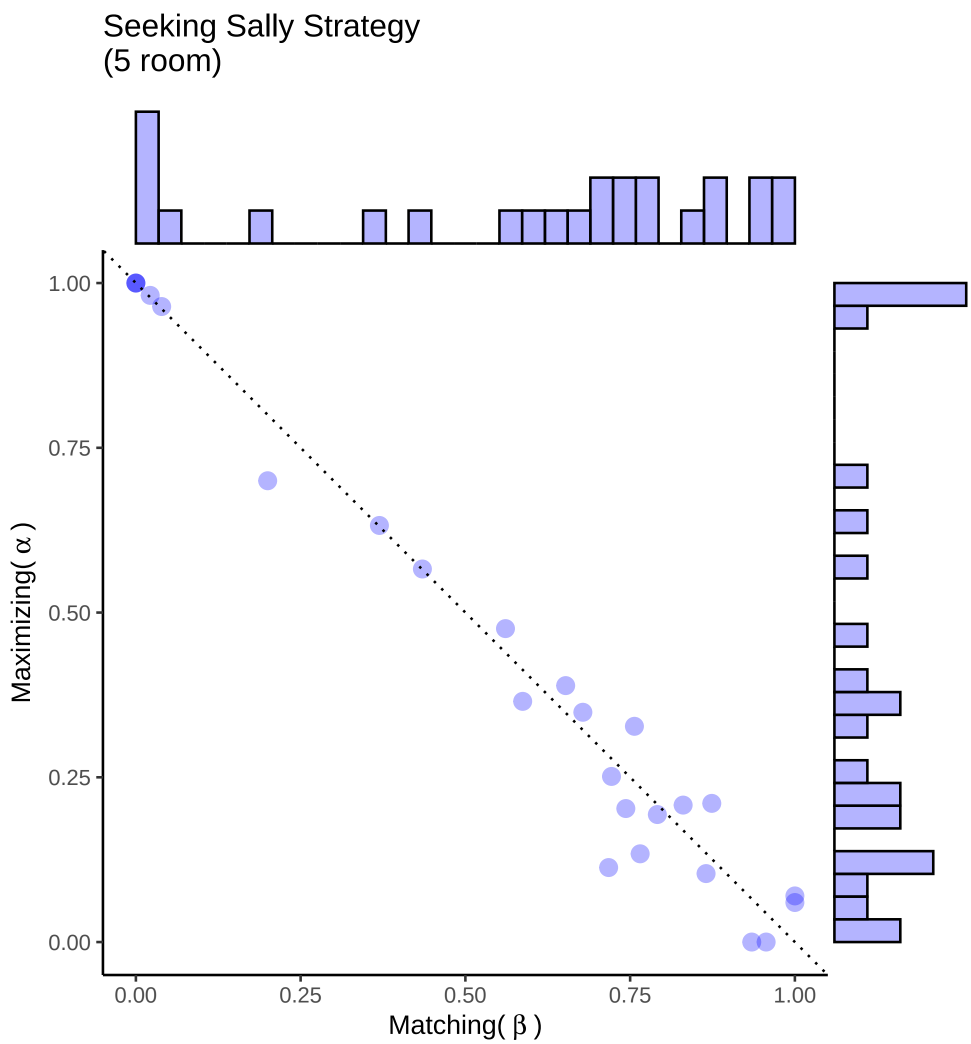

Figure 3.11: Strategy space of participant seeks against Sally. Bars along margin count participants using particular value of α or β, and indicate that most participants use a strategy with a low maximizing component (α) and high matching component (β).

3.2.6 Hiding from Sally (Five-Room)

Plots of participant hide frequencies over Sally's seek frequencies reveal a pattern similar to hides of the two-room condition, where participants range from perfect minimizing to perfect antimatching (Figure 3.12). Viewing this data in strategy space reveals that participants are indeed using strategies made mostly of minimizing and antimatching (Figure 3.13). As with five-room seeking data, participant strategies did not lie exactly on the line $\gamma + \delta = 1$, though this is expected with higher dimensional data. Error was still relatively low (Mean RSS = 0.0399392). In contrast to Study 1, this dataset includes at least five participants who use a strategy not near $\gamma + \delta = 1$. These participants appear to be using a strategy suitable for seeking – a combination of maximizing and matching. It could be that in the online environment, some participants were less focused on the task than would be if in the lab. This may result in some participants missing the indicators that they were to hide, rather than seek, on some trials.

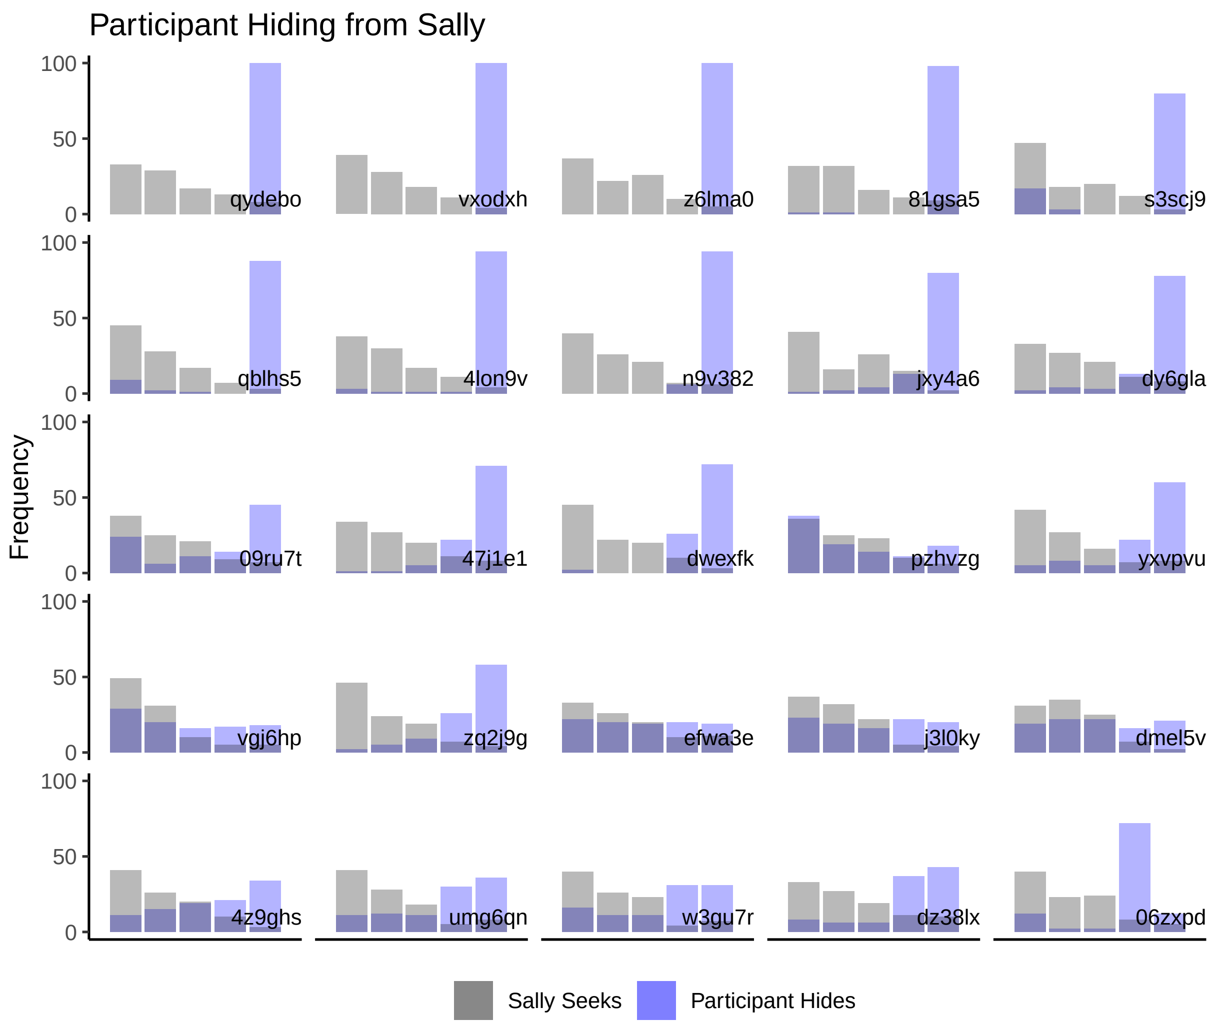

Figure 3.12: Participant hides (blue) and Sally's hides (grey) sorted by increasing values of δ (use of antimatching strategy). Histogram bins sorted by decreasing values of Sally's distribution. All participants use a mix of almost only minimizing and antimatching.

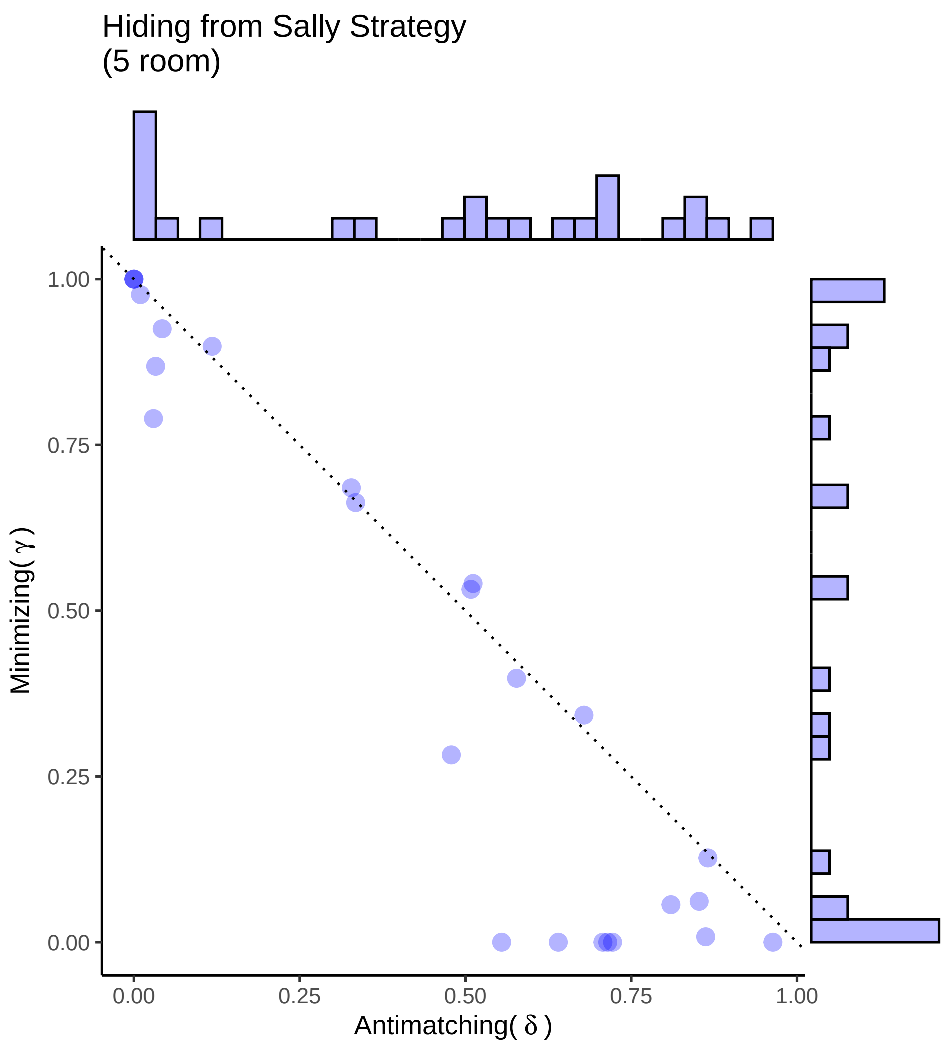

Figure 3.13: Strategy space of participant hides against Sally. Bars along margin count participants using particular value of γ or δ. Participants used a varied mix of minimizing and mismatching strategies. Model still shows relatively low error as strategies are clustered around the line $\gamma + \delta = 1$.

3.2.7 Comparing Hiding and Seeking (Five-Room)

Both seeking and hiding data cluster around the theoretical strategy mix line (Figures 3.11 and 3.13). However, this Study shows more noise in the hiding data than seen in Study 1. Also, while the five-room condition of Study 1 showed a marked shift to a more optimal when hiding (compared to seeking), this difference is less pronounced in Study 2. This diminished effect appears to be driven by the use of a more optimal *seeking* strategy in Study 2 compared to Study 1 (Figures 3.14 and 3.15). It could be that a distracted online participant, free from the social conformity pressures of a research lab, puts less effort into the task. This would result in a optimal heuristic effect, where instead of carefully probability matching when seeking, the simpler (but optimal) strategy of maximizing is used instead. The strategy mix of hiding may already be at a ceiling, which is why the online hiding data is similar to in-person, but the online seeking data approaches that of hiding. Finally, an online environment also introduces a potential compliance issue, where some participants appeared to have used a seeking strategy in the hiding trials.

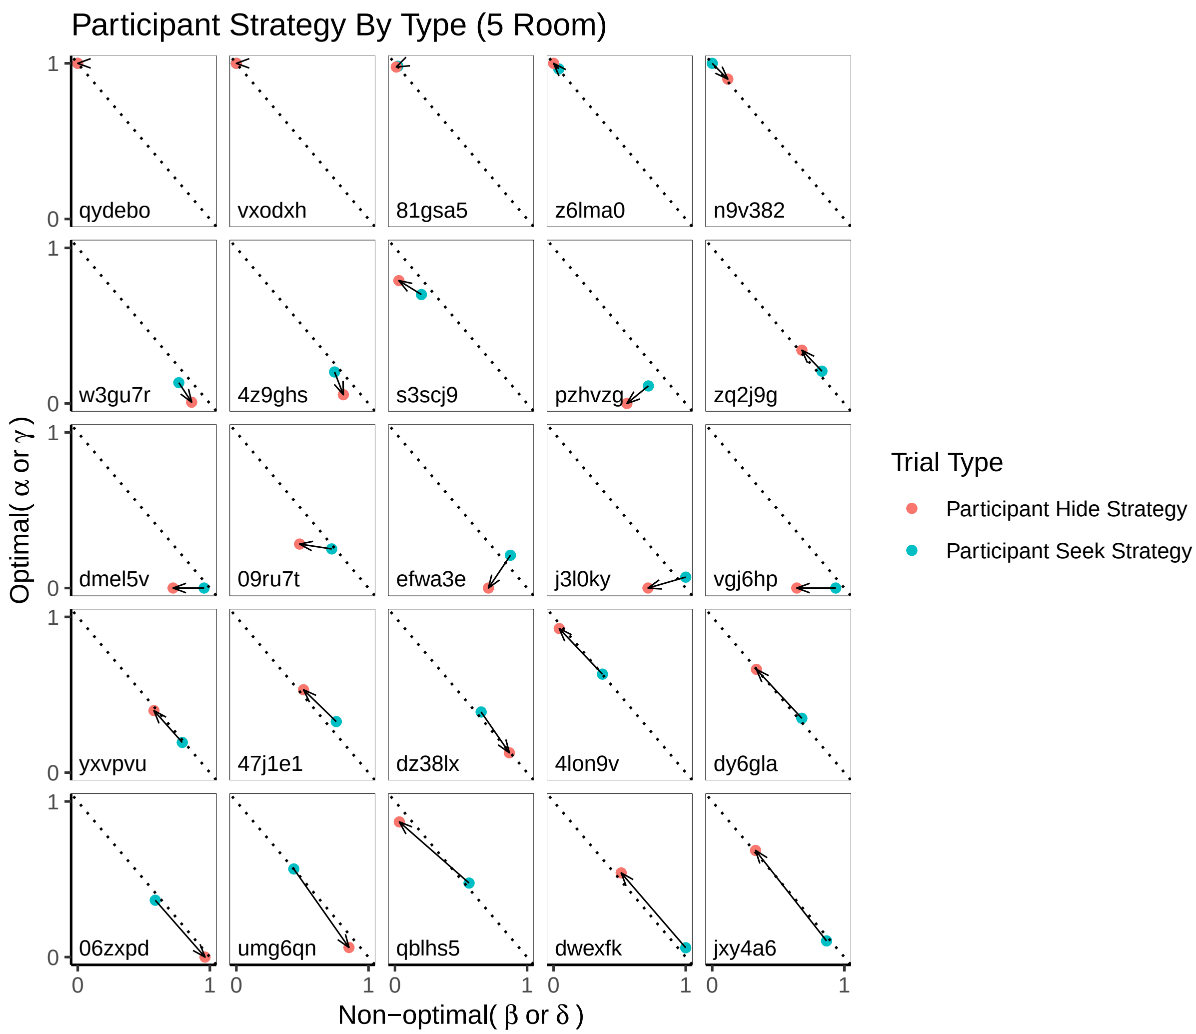

Figure 3.14: Each participant's strategy mix for seeking (α, β) hiding (γ, δ) are plotted together, ordered by increasing values of strategy change. Arrow points from seek strategy to hide strategy.

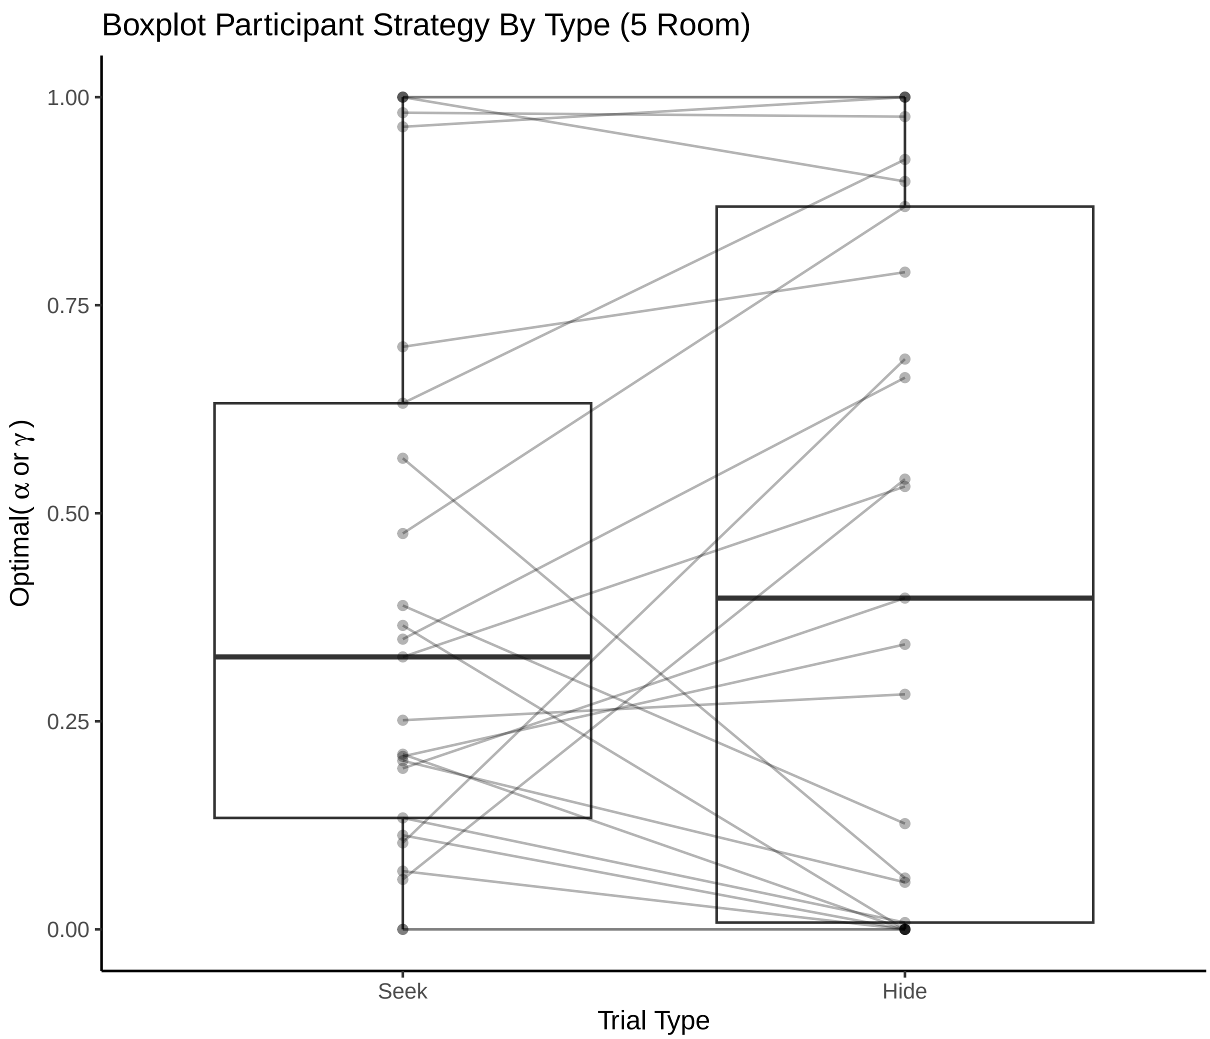

Figure 3.15: Most individual participants change strategies between hiding and seeking, any many change to a more optimal strategy mix when hiding, though this difference is less pronounced compared to Study 1.

3.2.8 Sensitivity Analysis (Five-Room)

We again repeat the sensitivity analysis, this time for the five-room data. The plot presenting the strategy space for participant hides against Sally is repeated four times using random strategy vectors instead of the minimizing ($\vec{n}$) and antimatching ($\vec{a}$) strategies of our model (Figure 3.16). All four random strategy spaces presented fail to explain participant behaviour effectively, resulting in few strategies lying on or near the line $\gamma_r + \delta_r = 1$ and inflated Mean RSS values ranging between 0.1756 and 0.3868. Repeating this procedure 1000 times and plotting Mean RSS draws the same conclusion: $\vec{q} = \gamma\vec{n} + \delta\vec{s} – \vec{\zeta}$ is a more reliable model than $q = \gamma_r r\vec{1} + \delta_r r\vec{2} – \vec{\zeta}_r$ (Figure 3.17).

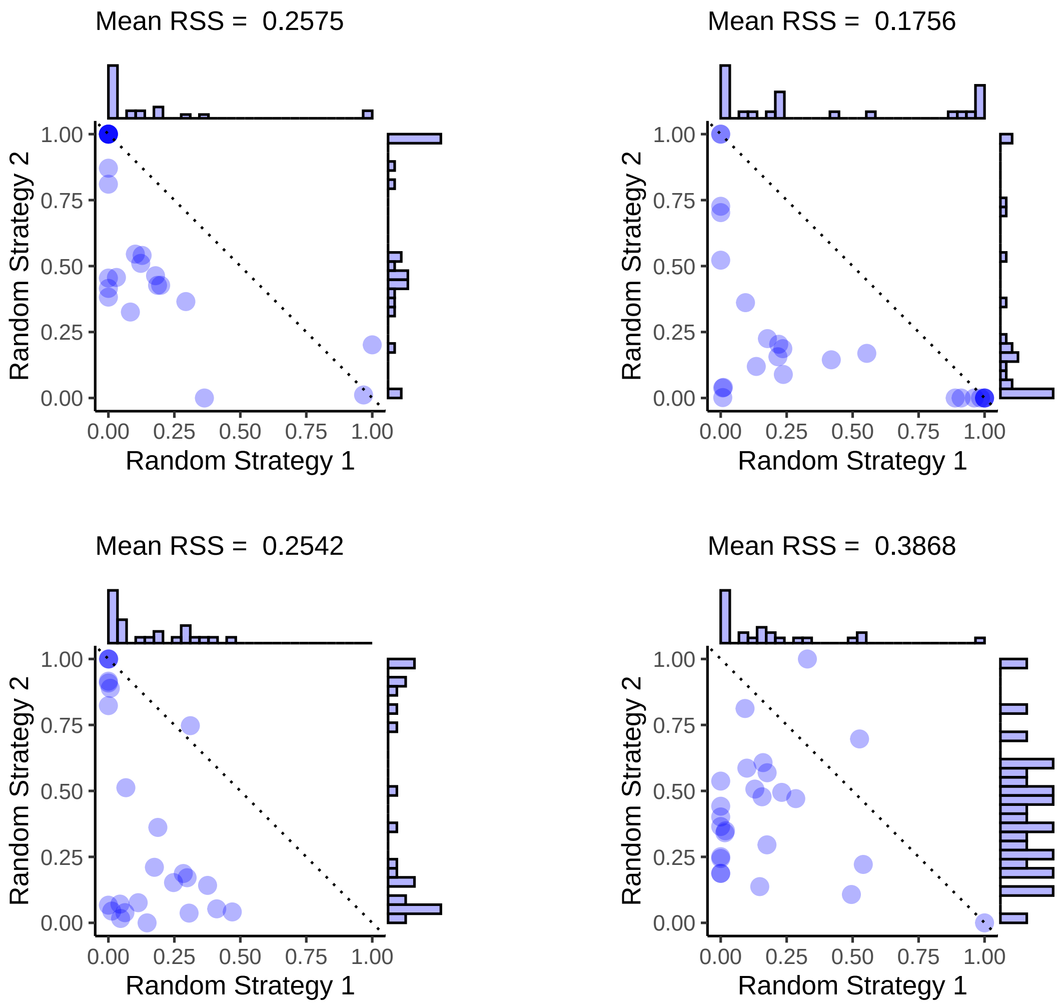

Figure 3.16: Plotting participant hides against sally in a strategy space using two random strategy vectors instead of minimizing and antimatching strategies results in a poor model fit. Many points are in region of high error, reflected in high Mean RSS values. Figure contains four instances of this plot, each with two new randomly generated strategy vectors.

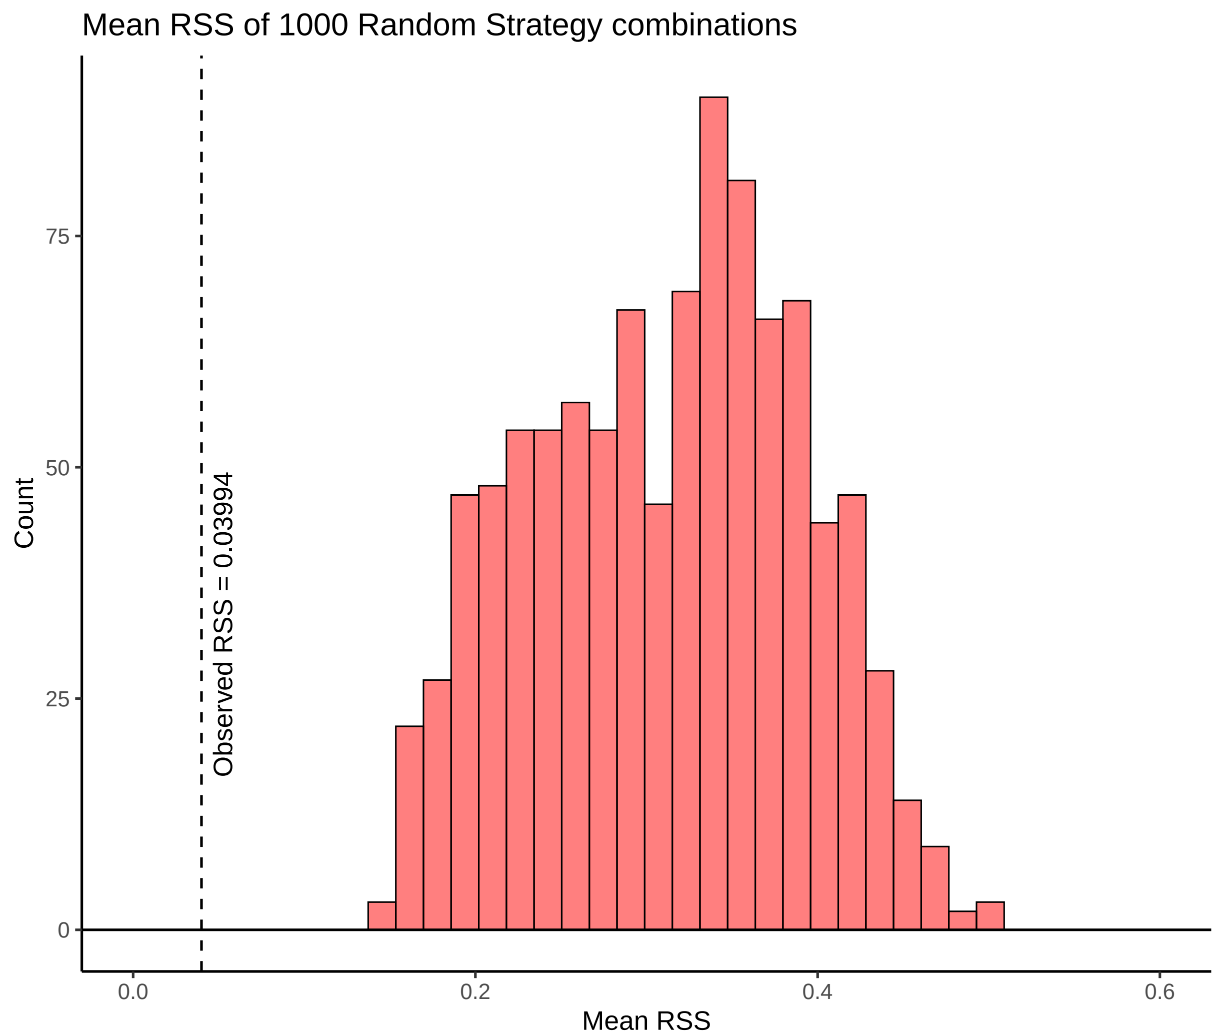

Figure 3.17: Fitting participant data to a strategy space of two random strategy vectors 1000 times. Figure shows histogram of resulting error. Our proposed model dramatically outperforms the 1000 tested here. This implies that our model is of meaningful accuracy, and our methods used are not biased to result in low any for any arbitrary model it is given.

3.2.9 Comparing Two-Room to Five-Room

Unlike in Study 1, we do not see a notable difference in the strategy mix between the two-room and the five-room conditions. This appears to be largely driven by the fact that participants in Study 2 did not utilize the non-optimal seeking strategy (probability matching) as much as participants in Study 1. However, the lack of probability matching was replaced with maximizing, so we still observe that participant strategies lie close to the line $\alpha + \beta = \gamma + \delta = 1$ implying our model appropriately accounts for participant hiding and seeking behaviour in both two-room and five-room conditions (Figures 3.7, 3.15, and 3.18). Again, the line $\alpha + \beta = \gamma + \delta = 1$ represents a strategy mix of exclusively matching and maximizing for seeking, and exclusively antimatching and minimizing for hiding.

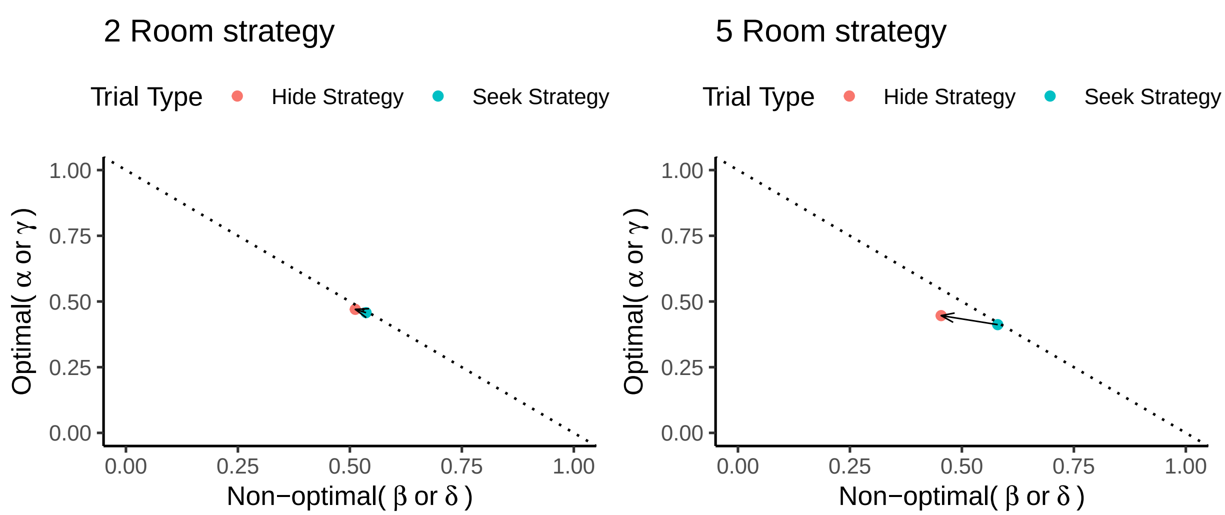

Figure 3.18: All strategy means lie close to the line $\alpha + \beta = \gamma + \delta = 1$ implying our model appropriately accounts for participant hiding and seeking behaviour in both two-room and five-room conditions. But we do not see a strategy shift between hiding and seeking in the five-room condition as extreme as seen in Study 1.

3.2.10 Special case opponent strategies

Participant behaviour against Kala (uniform) and Bo (100%) follow closely to that of Study 1. While all strategies against a uniform distribution are equivalent, participants mostly produce a nearly uniform distribution, while hiding and seeking (Figure 3.19). Note that the reflection of the uniform is the uniform, thus the vector reflection hypothesis appears to still function in this ambiguous case.

Participant seeking against the 100% distribution utilizes the unique reasonable strategy – only picking the 100% room. With respect to hiding, participants again show a mix between exploiting a particular 0% room, and distributing hides across all of the 0% rooms (Figure 3.20). Note that no unambiguous reflection of the 100% distribution exists, making this heterogeneous strategy choice when hiding against this distribution particularly notable.

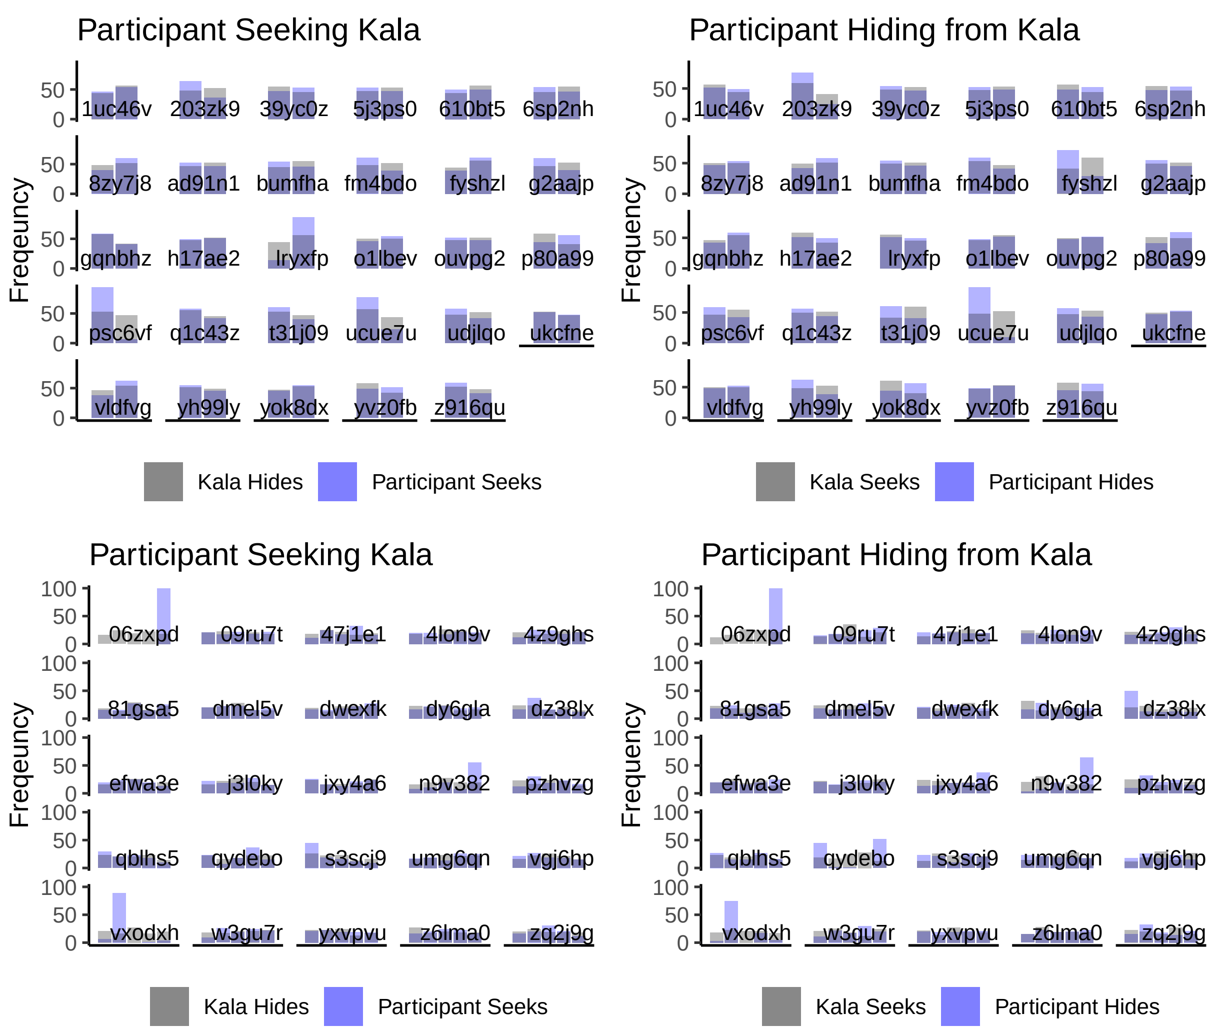

Figure 3.19: All participant rounds against Kala. Seek data in left column, hide in right column. Two-room on top row, five-room on bottomrow. Most participants use an approximately uniform strategy as expected, though any other strategy has a theoretically equivalent payoff for the participant

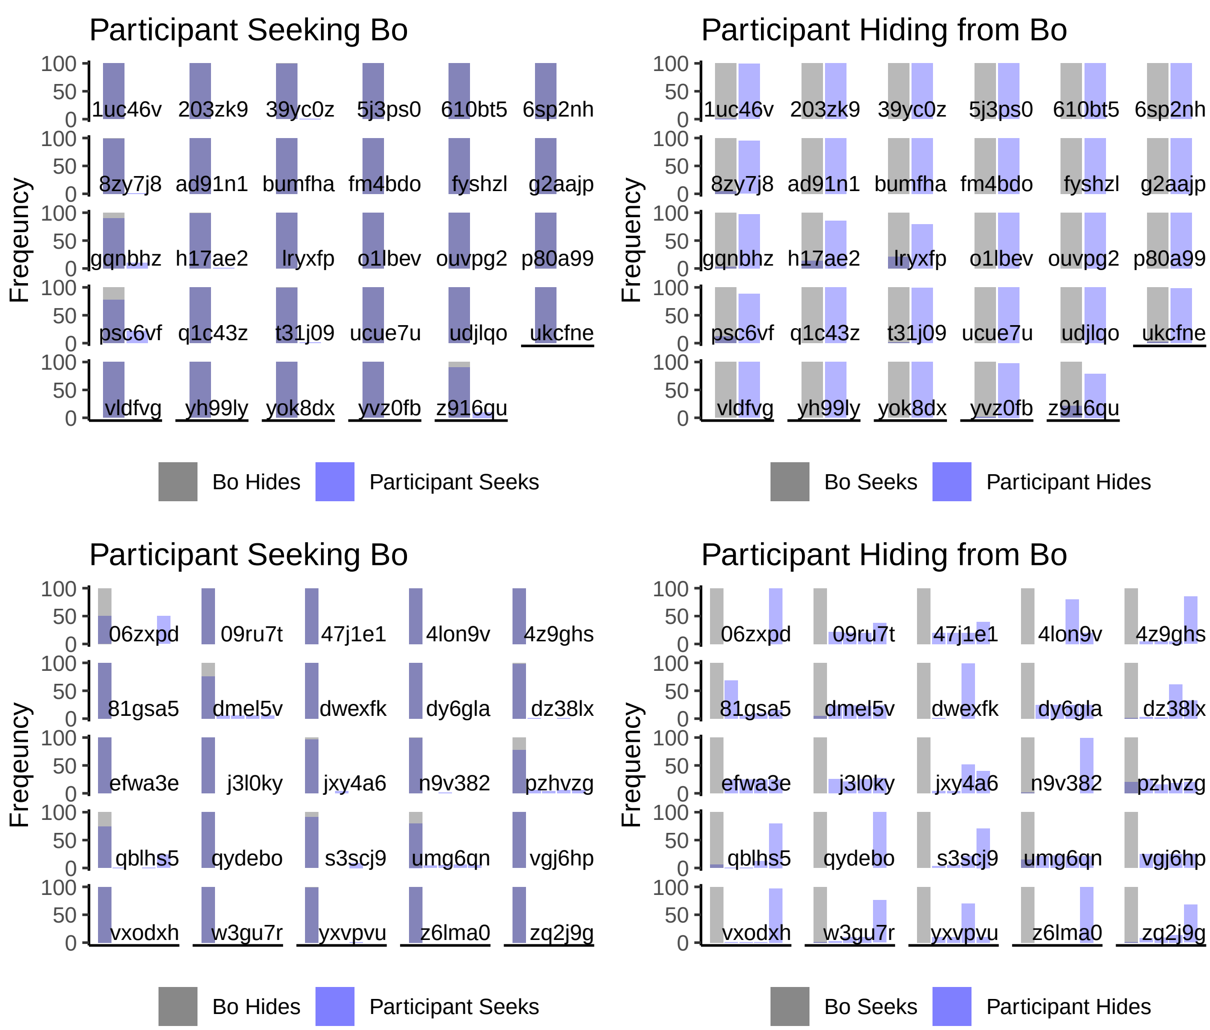

Figure 3.20: All participant rounds against Bo. Seek data in left column, hide in right column. Two-room on top row, five-room on bottom row. Participants behaved optimally aside from a few apparently erroneous clicks, ensuring the experiment was understood. Hiding behaviour in the five-room condition was least consistent, and is also the only dataset incompatible with our vector reflection model.

3.2.11 Demographics

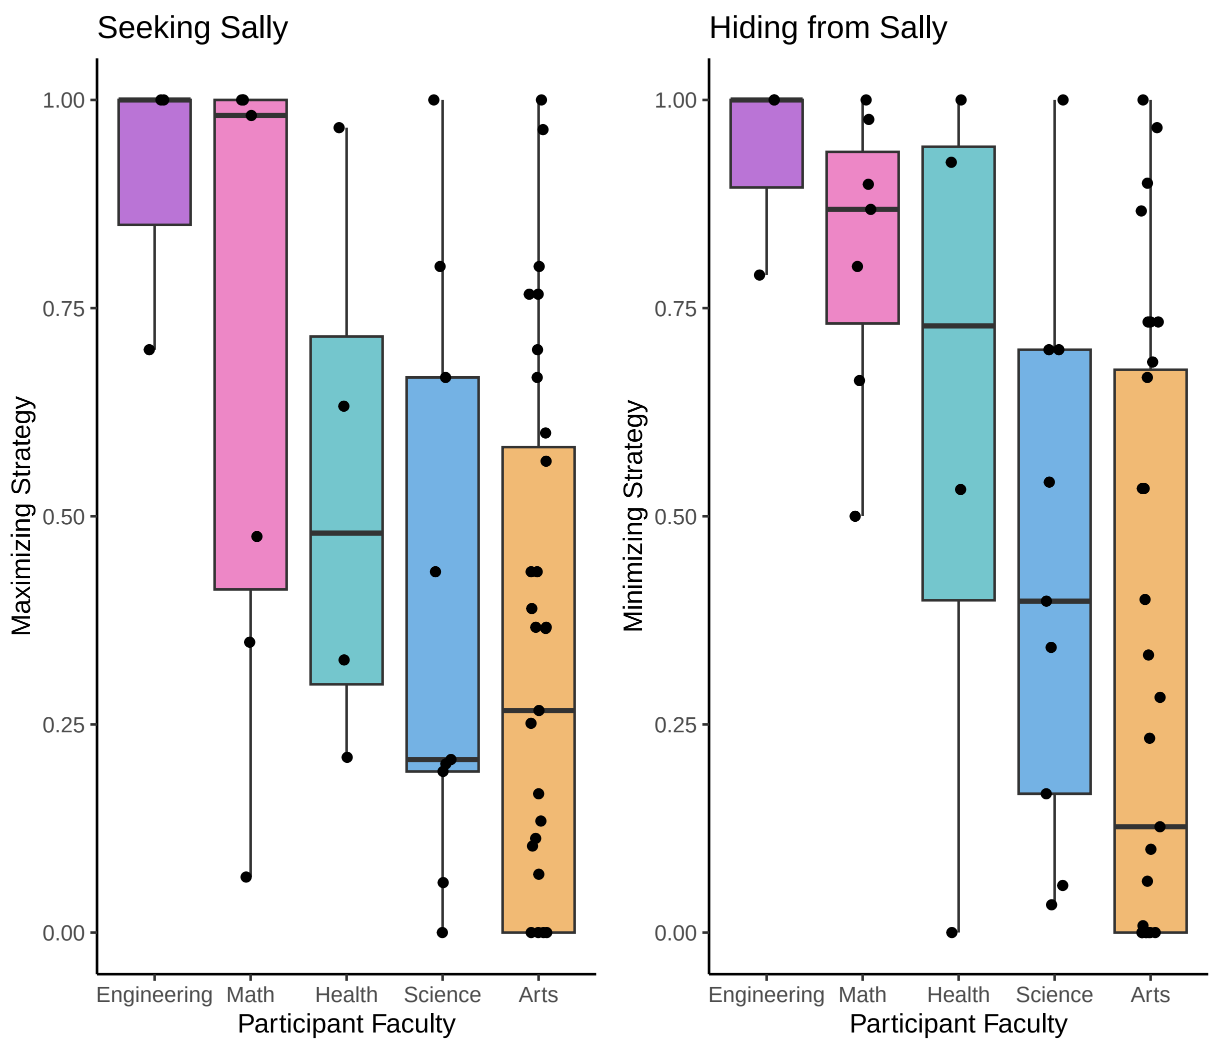

Figure 3.21: Optimal strategy use by academic faculty of student. Rankings do not follow that of Study 1.

Unlike in Study 1, we observe a less surprising trend where participants from more math-oriented faculties utilize the theoretically optimal strategy of maximizing when seeking and minimizing for hiding (Figure 3.21). This is evidence that math experience increases a participant's ability to recognize the optimal minimizing strategy, more than it facilitates the use of the cognitively challenging non-optimal antimatching distribution.

Chapter 4

Study 3: More Hide-and-Seek Distributions

In Studies 1 and 2, we demonstrate the efficacy of vector reflections in defining probability antimatching. In this study, we further test our notion of probability antimatching by testing participants with more probability distributions. These include distributions that produce negative probability entries when reflected as a vector across the uniform, and distributions that have infinitly many optimal maximizing or minimizing strategies. In addition, we also set dimensionality (the number of rooms) as a within-participant variable for this study. By having each participant face opponent hide-and-seek distributions from multiple dimensions, we can more easily test the representational complexity hypothesis.

4.1 Methods

4.1.1 Task Environment

Stimuli

Visual stimuli are idential to that of Study 2, except with a new set of probability distributions and the addition of another opponent character (Table 4.1).

Procedure

Participants played a short practice round consisting of three seeking trials followed by three hiding trials, identical to previous studies. Following the practice round, participants completed a total of four hide-and-seek games, where each game is structured identically to games in Studies 1 and 2. That is, each game consists of 100 seeking trials and 100 hiding trials.

The distribution faced in practice rounds was always (30, 30, 20, 10, 10). The distribution each participant faced in the first game was randomly selected from one of three sets: 3-, 5-, or 7-room distributions. The second game distribution was randomly selected from the remaining two sets that did not contain the distribution selected for the first game. The third distribution was randomly selected from the remaining set. The fourth game was randomly selected from a set that only contains distributions with probability mass all in one room (i.e. an entry of 100% and n-1 entries of 0%'s). For all games (including practice), one of the five characters were assigned to be identity of the opponent. Each opponent identity was used for exactly one game.

The nature of this procedure is efficiently described with mathematical notation. The symbol $\in_R$ denotes 'randomly selected from'. A backslash in set notation denotes exclusion. For example, $\{A, B, C, D\} \backslash \{A, C\} = \{B, D\}$.

Let

$$D^3 = \left\{ \begin{bmatrix} 50 \\ 25 \\ 25 \end{bmatrix}, \begin{bmatrix} 42 \\ 42 \\ 16 \end{bmatrix}, \begin{bmatrix} 80 \\ 10 \\ 10 \end{bmatrix} \right\},$$

$$D^5 = \left\{ \begin{bmatrix} 35 \\ 30 \\ 15 \\ 15 \\ 5 \end{bmatrix}, \begin{bmatrix} 35 \\ 25 \\ 25 \\ 10 \\ 5 \end{bmatrix}, \begin{bmatrix} 45 \\ 35 \\ 10 \\ 5 \\ 5 \end{bmatrix} \right\},$$

$$D^7 = \left\{ \begin{bmatrix} 25 \\ 25 \\ 20 \\ 14 \\ 10 \\ 4 \\ 2 \end{bmatrix}, \begin{bmatrix} 26 \\ 24 \\ 18 \\ 15 \\ 9 \\ 4 \\ 4 \end{bmatrix}, \begin{bmatrix} 50 \\ 18 \\ 12 \\ 8 \\ 5 \\ 5 \\ 2 \end{bmatrix} \right\}, \text{and}$$

$$D' = \left\{ \begin{bmatrix} 100 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 100 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 100 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right\}.$$

- Each participant plays 4 games: G_1, G_2, G_3, G_4
- Each game G_i has one opponent that follows a particular hide-and-seek distribution, $\vec{g_i}$
- For any game G_i with distribution $\vec{g_i}$, each element of $\vec{g_i}$ is randomly assigned to each room of the house presented on the screen
- $\vec{g_1} \in_R D^n$, where $n \in_R \{3, 5, 7\}$
- $\vec{g_2} \in_R D^m$, where $m \in_R \{3, 5, 7\} \backslash \{n\}$
- $\vec{g_3} \in_R D^k$, where $k \in_R \{3, 5, 7\} \backslash \{n, m\}$
- $\vec{g_4} \in_R D'$

4.1.2 Participants

177 (134 female) participants completed the online hide and seek task. Each participants had hide and seek distributions assigned to them identically and independently, as described above. The Study took approximately 32 minutes to complete.

Bo	Jill	Kala	Sally	Toby

Table 4.1: Study 3 opponent identities are randomly assigned to each game without replacement.

4.1.3 Data Analysis

As in the previous Studies, all participants who play a game with a given distribution are collapsed together, ignoring the context in which that distribution is presented. This includes room probability assignments, order of distribution presentation, and opponent identity.

Accounting for infinite optimal strategies

Recall, *participant strategy* = *optimal strategy* + *non-optimal strategy*. For seeking, that is *participant seeking strategy* = *maximizing strategy* + *matching strategy*. For hiding, that is *participant hiding strategy* = *minimizing strategy* + *antimatching strategy*. In previous Studies, the optimal strategies were unique (i.e. picking exclusively the maximum probability room when seeking, and exclusively the minimum probability room when hiding). In this experiment, a subset of the distributions have more than one optimal strategy. This occurs for seeking when the maximum probability room is not unique (a tie for the maximum probability room), and for hiding where the minimum probability room is not unique (a tie for the minimum probability room).

Interestingly, if there is no unique optimal strategy, then there are *infinite* optimal strategies. With more than one room tied for maximum (minimum) value, these rooms can be considered equivalent. Therefore *any* distribution of seeks (hides) across the maximum (minimum) probability rooms will result in an equivalently optimal strategy. Note, there are only a theoretically infinite number of optimal strategies. Since participants only complete 100 trials for seeking/hiding for each distribution, there are actually finitely many for any discrete dataset. Our analysis will also account for the infinite theoretical case, making this analysis generalizable for future work.

For any distribution that has a non-unique maximum or minimum room, the maximizing and minimizing strategies are generalized from the vectors $\vec{x}$ and $\vec{n}$, respectively, to the sets of vectors X and N, respectively. The number of vectors in X is determined by the number of the rooms that share the maximum probability value. Analogously, the number of vectors in N is determined by the number of rooms that share the minimum probability value. Both X and N contain only standard unit vectors, denoted $\vec{e_i}$, where the i'th element is 1 and all other elements are 0. The set X contains vectors $\vec{e_i}$, where i is selected from the room numbers that have the maximum probability value. For example, if $\vec{m} = (0.42, 0.42, 0.16)$, then $X = \{(1, 0, 0), (0, 1, 0)\}$. If $\vec{m} = (0.26, 0.24, 0.18, 0.15, 0.09, 0.04, 0.04)$, then $N = \{(0, 0, 0, 0, 0, 1, 0), (0, 0, 0, 0, 0, 0, 1)\}$.

As in previous experiments, we apply a the Stark-Parker algorithm for bounded-variable least squares via the 'bvls' package in R [27]. In this case, when analyzing seeking data, we include all vectors in X along with $\vec{m}$, to determine the strategy mix α and β. We define α as $\sum \alpha_i$, where α_i is the coefficient for each $\vec{e_i}$ in X that results from the bvls algorithm. We do this because the use of any mix of maximizing strategy, represented by vectors in X, is equivalent and can thus be combined into one value. We perform an analogous analysis for hiding data using N.

Testing invalid reflections

A distribution $\vec{x}$ is said to have an invalid distribution if and only if $refl(\vec{x}) = \vec{y}$, where $\vec{y}$ contains elements that are less than zero. Three of our tested distributions have invalid distributions:

- (80,10,10) reflects to (-13.3,56.6,56.6)
- (45,35,10,5,5) reflects to (-5,5,30,35,35)
- (50,18,12,8,5,5,2) reflects to (-21.4, 10.6, 16.6, 20.6, 23.6, 23.6, 26.6)

By the nature of our task, participants are required to present their hiding strategy behaviorally via room selection frequencies. Therefore, an antimatching strategy (a reflection of the seek distribution of their opponent) cannot be expressed behaviorally, because the antimatching component of our strategy mix contains negative values, and negative rooms selections are not possible.

All frequency-based representations of probability must be represented as a vector within the probability simplex. In addition to expressing participant hiding behaviour with the invalid reflections that contain negative entries, we also consider two projection methods that compute a vector that is close to the original vector reflection, but still exists within the probability simplex. We can then represent participant hiding behaviour as a linear combination of the optimal minimizing strategy, and this “closest possible” antimatching strategy. This allows us to explore possible consistencies across individuals and potentially generalize what it means to use an antimatching strategy against a strategy that has no physically observable opposite.

Shift back to the uniform Shifting invalid reflections back to the simplex in the direction of the uniform distribution is computationally simple. Consider a $\vec{q} = refl(\vec{p})$ where at least one element of $\vec{q}$ is strictly less than 0. Let q_{min} be the minimum valued element of $\vec{q}$. Then, let $\vec{v} = (q_{min}, ..., q_{min})$. We now know that $\vec{q} - \vec{v}$ will have a minium entry of exactly 0. After normalizing $\vec{q} - \vec{v}$ such that it sums to 1, we now have a valid probability histogram vector that is along the same trajectory between, $\vec{p}$, $\vec{q}$, and the uniform distribution (Figure 4.1). Note that this method of returning to the simplex may not result in $\vec{q} - \vec{v}$ being the closest possible point on the simplex to $\vec{q}$.

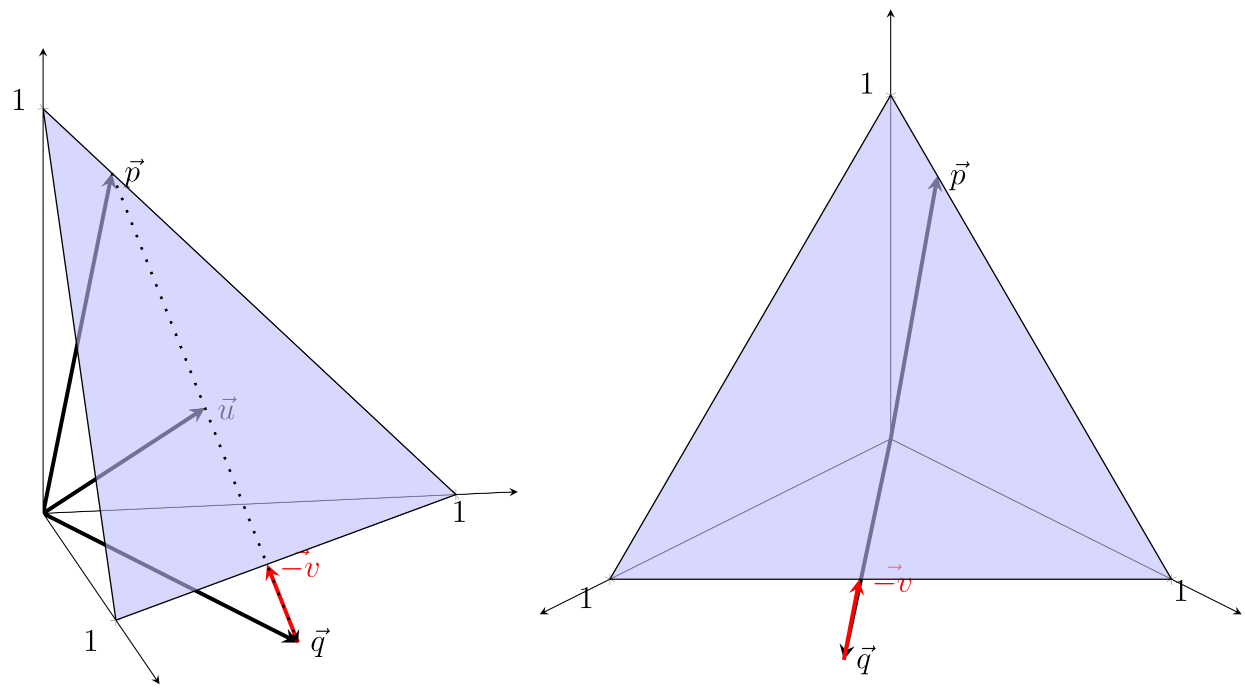

Figure 4.1: Two different viewpoints of the probability simplex in 3-dimensional space. Here, the reflection of $\vec{p}$, denoted $\vec{q}$, lies outside of the probability simplex. Subtracting a uniform vector containing most negative element of $\vec{q}$, denoted $\vec{v}$, returns $\vec{q}$ to the simplex along the original trajectory of the reflection of $\vec{p}$ over $\vec{u}$.

Project to the closest point on the simplex Our second projection methods considers a simple idea: if $\vec{q} = refl(\vec{p})$ is outside of the simplex, then find the point on the simplex that is closest to $\vec{q}$ (Figure 4.2). Mathematically, this is a Euclidean projection onto the simplex requiring a optimization problem to find some vector $\vec{w}$ that minimizes $\vec{w} - \vec{q}$ such that $\vec{w}$ is within the probability simplex. A solution to this problem can be found in [6]. By the nature of this proof, we considered participant strategy mixes using this simplex projection method for the cases where $\vec{q}$ exists outside of the probability simplex. We adopted an implementation of this projection method from [29].

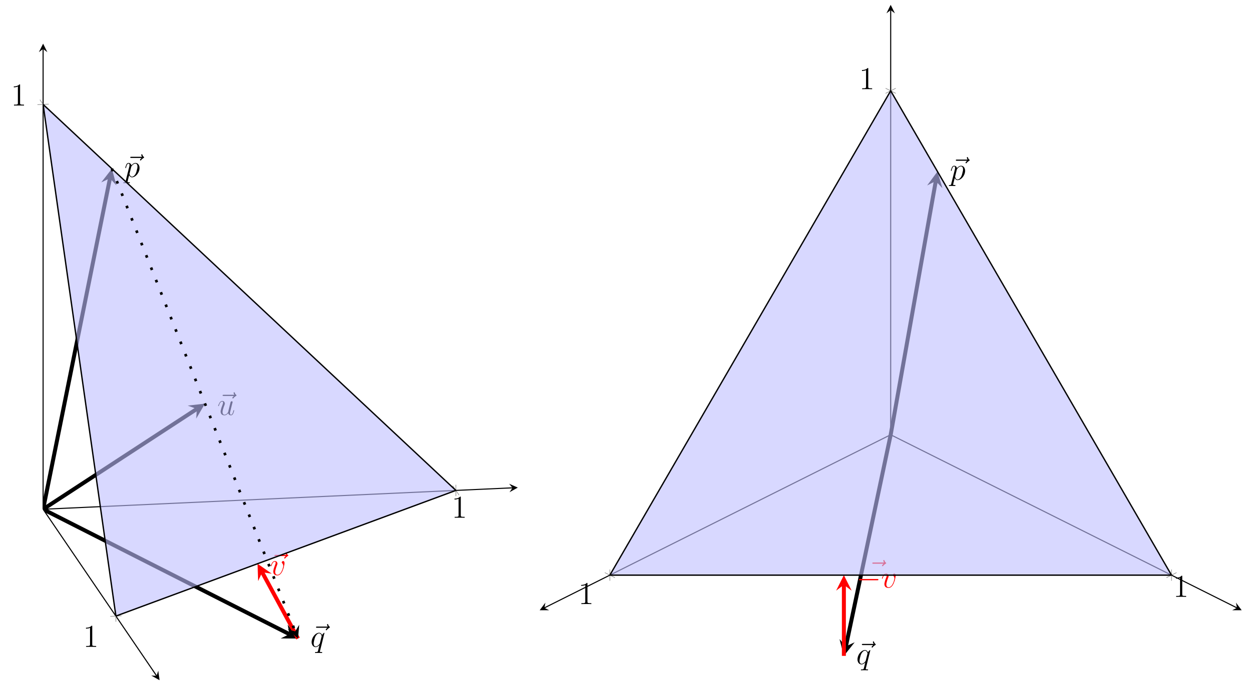

Figure 4.2: Two different viewpoints of the probability simplex in 3-dimensional space. Here, the reflection of $\vec{p}$, denoted $\vec{q}$, lies outside of the probability simplex. Performing the Euclidean projection onto the simplex produces the shortest possible vector $\vec{v}$ to return to closest vector in the simplex to $\vec{q}$, $\vec{w}$. Unlike the uniform shift method, this does not always shift $\vec{q}$ back along the original reflection trajectory from $\vec{p}$.

Testing representational complexity

To test the representational complexity hypothesis, we compare the optimality of hiding vs seeking strategy mixes across room-number conditions. If hiding strategies become more optimal than seeking strategies as dimensionality increases, this would support the idea that participants use a more optimal hiding strategy because it is easier to compute than antimatching. If hiding strategies become less optimal than seeking strategies as dimensionality increases, this would support the idea that participants hold priors that failed hides are worse than failed seeks, but the stakes of each room selection is lesser when there are more rooms (chance of being found is lower).

4.2 Results and Discussion

4.2.1 Matching and antimatching behaviours are observed and infinite optimal strategies are adequately represented

The strategy plots of this Study (Figures 4.3, 4.4, and 4.5) resemble those from Studies 1 and 2. That is, individual participant strategy points cluster around the diagonal line $\alpha + \beta = 1$ for seeking, and $\gamma + \delta = 1$ for hiding. Proximity to this theoretically expected line is nearly error-free in the 3-dimensional conditions, much like in the 2-dimensional conditions observed in previous studies. Recall, this is because fewer hide-and-seek rooms translates to fewer degrees of freedom, represented by a lower dimensional strategy space. Our 2-dimensional model of behaviour (*optimal* + *non-optimal*) more easily spans the entire strategy space when the strategy space itself is lower-dimensional, resulting in lower representational error. Representational error increases at a diminishing rate as participant degrees of freedom increase, since participant strategy mixes are similarly clustered around the diagonal in both the 5- and 7-dimensional cases. This implies that our model captures genuine features of participant behaviour, and is not artificially assisted by low dimensional data.

The possibility of infinitely many optimal strategies is accounted for by our methodology of using optimal strategy spans. This occurs in cases where the maximum probability room is not unique for seeking, and when the minimum probability room is not unique for hiding. These cases are denoted with a green cross (+) in Figures 4.3, 4.4, and 4.5, and are indistinguishable from strategy mixes where only one unique optimal strategy exists.

Some histogram vectors produce negative values when reflected. This renders these reflections "invalid" because there is no meaningful interpretation of a probability distribution with negative probability mass, let alone any way for participant behaviour to ever manifest *negative room selections*, at least in our protocol. Therefore, representing participant hiding behaviour as a linear combination of the optimal minimizing strategy, and, in this particular case, an antimatching strategy that has no behavioral manifestation, should result in a poor strategy mix representation. This is exactly what we observe in our data. These cases are denoted with a red asterisk in Figures 4.3, 4.4, and 4.5.

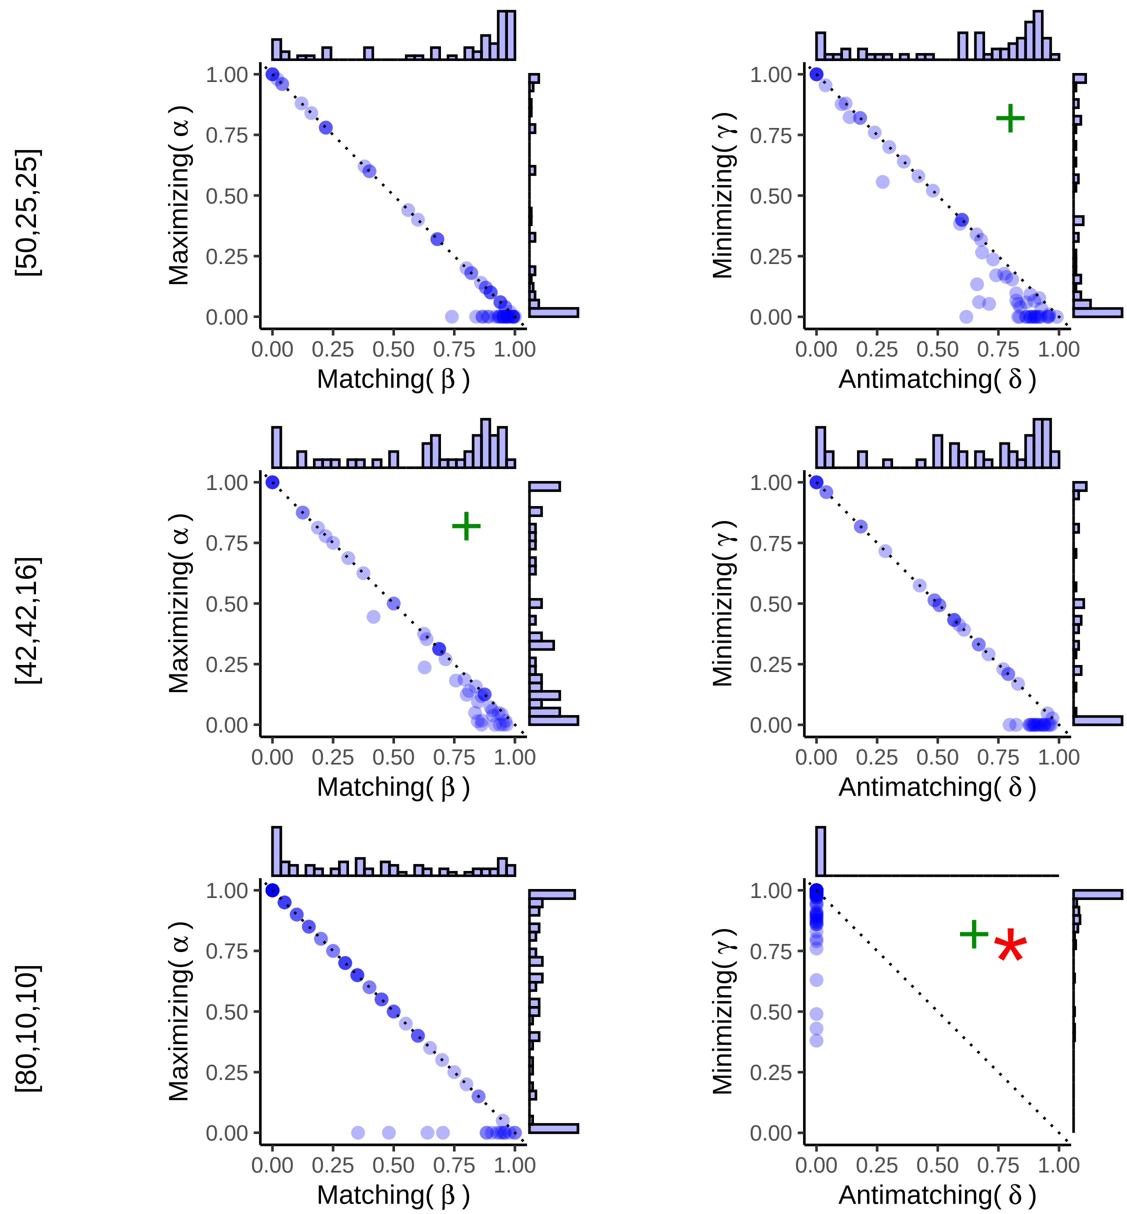

Figure 4.3: Strategy plots for seek (left) and hide (right) distributions from the set D^3. Strategy mixes calculated with infinite optimal strategies are denoted with a green cross (+). Strategy mixes that are computed with a distribution that has an invalid reflection are denoted with a red asterisk (*).

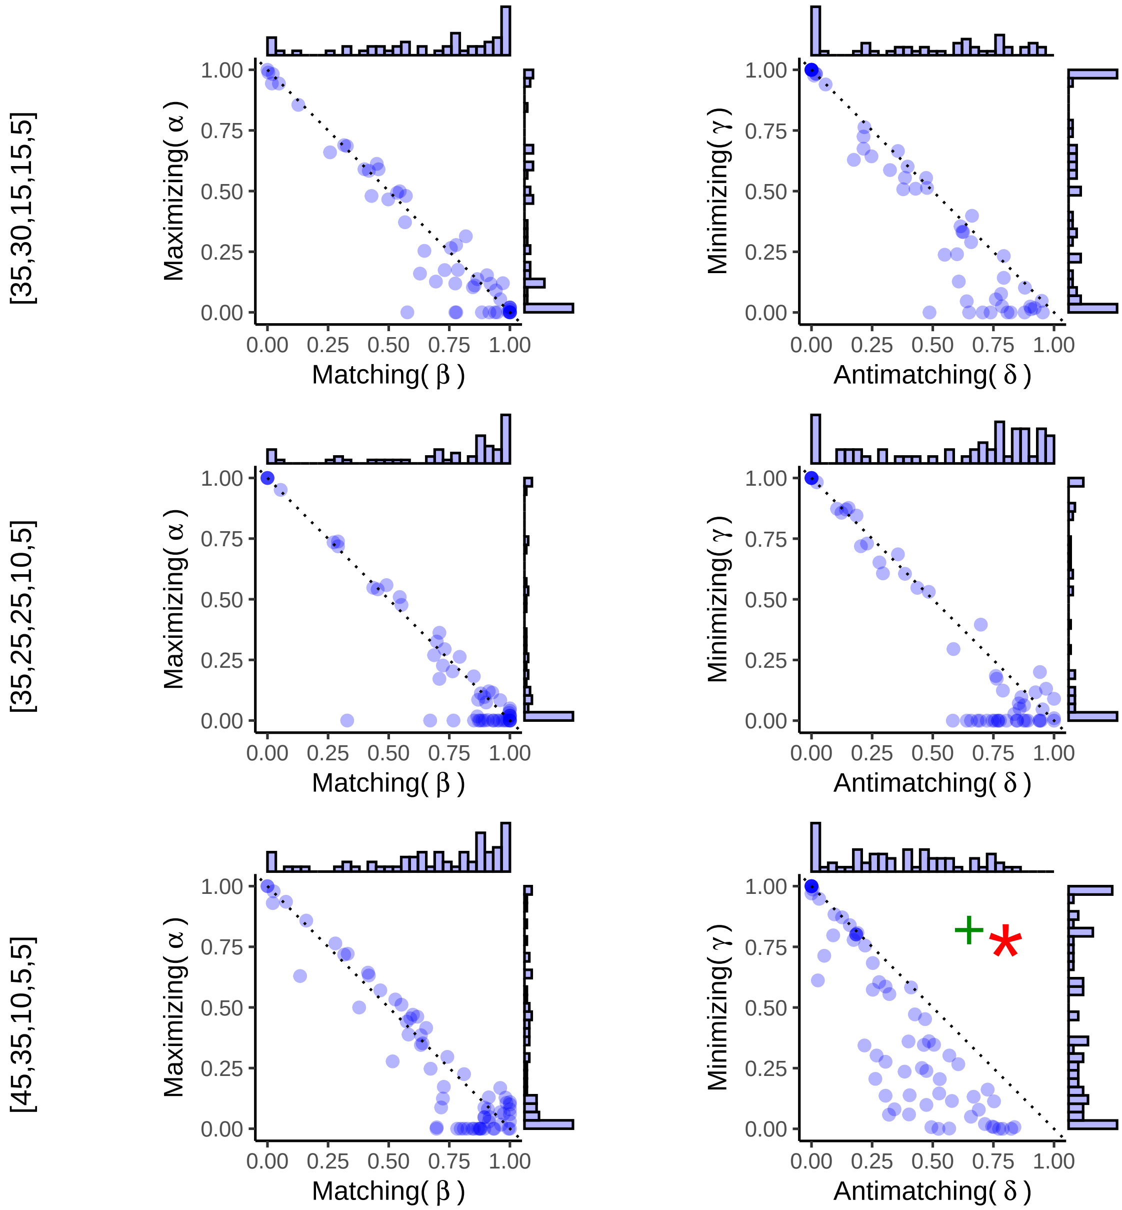

Figure 4.4: Strategy plots for seek (left) and hide (right) distributions from the set D^5. Strategy mixes calculated with infinite optimal strategies are denoted with a green cross (+). Strategy mixes that are computed with a distribution that has an invalid reflection are denoted with a red asterisk (*).

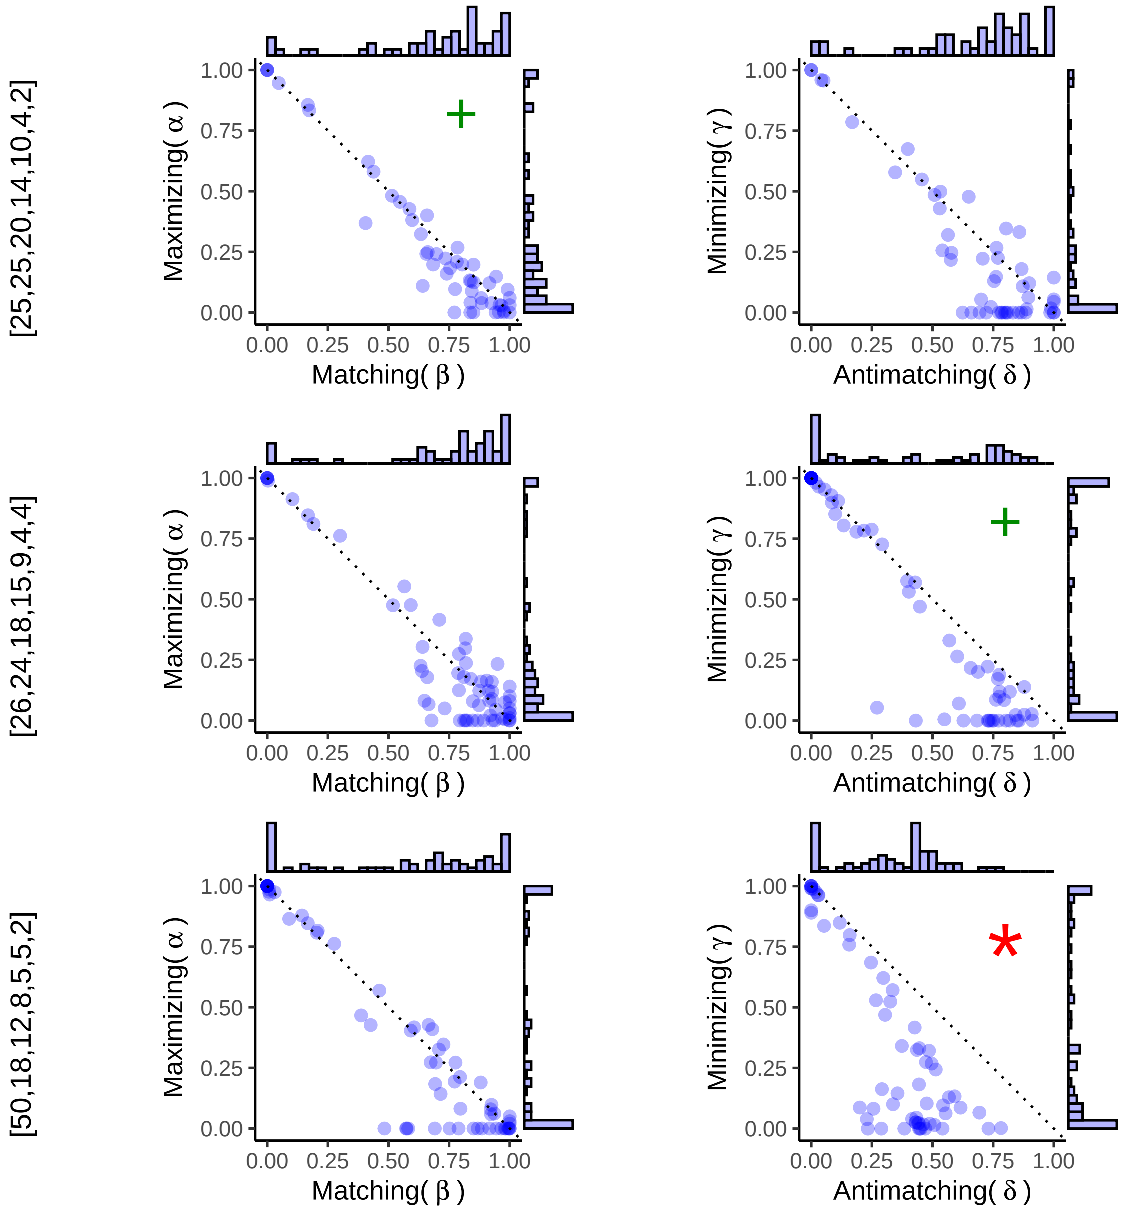

Figure 4.5: Strategy plots for seek (left) and hide (right) distributions from the set D^7. Strategy mixes calculated with infinite optimal strategies are denoted with a green cross (+). Strategy mixes that are computed with a distribution that has an invalid reflection are denoted with a red asterisk (*).

4.2.2 Participant hiding strategies follow a projection back to the simplex when reflections are invalid

Figures 4.6, 4.7, and 4.8 show the hiding strategy mixes for the invalid reflection case in 3-, 5-, and 7-room conditions, respectively. For each figure, the leftmost panel shows the strategy mix calculated with the unadjusted invalid reflection that contains negative values. Recall, a distribution with negative entries cannot produce a hiding strategy in terms of relative room selection frequencies, because participants cannot make negative room selections. As such, we expect, and observe here, that hiding strategy is not well represented as a linear combination of the optimal minimizing strategy, and our antimatching strategy that has no behavioural manifestation. The exception to this is the 3-room case, where most participants use a strictly maximizing strategy, and the antimatching component is not needed to represent hiding strategy.

Each middle panel shows hiding strategy mixes where the invalid reflections are projected back to the simplex via the uniform shift method. Each rightmost panel shows hiding strategy mixes where the invalid reflections are projected back to the simplex via the closest point method. In all three conditions, the results of each pair of shifts are very similar to each other. It is likely that any reasonable method to project an invalid reflection back to simplex will result in very similar distributions. Future mathematical work can explore which methods may differ the most, and which particular reflections may result in the largest discrepancy between projection methods.

In the 3-room condition, each projection method happens to produce the same adjusted reflection (0, 50, 50), resulting in an uninformative strategy plot 4.6. This is because the shortest path back to the simplex (via the Euclidean projection) lies along the original trajectory from the uniform (via the uniform shift). Moreover, the adjusted reflection (0, 50, 50) is also one of the infinity many optimal seeking strategy of hiding any arbitrary mix of rooms, so long as you only hide in the lowest probability rooms. Therefore, in this particular case, the antimatching strategy is also an optimal minimizing strategy. This explains why our adjusted hiding strategy plots only differ from the original by representing some exclusive minimizers as exclusive antimatchers. As such, hiding strategy is not better represented when using adjusted reflections, compared to without. From our analysis it is unclear how the bounded values least squares algorithm assigns either $\gamma = 1$ or $\delta = 1$ when $\vec{n}$ and $\vec{a}$ are identical.

In the 5-room condition, both projection methods induce small improvements to representational error compared to using the invalid reflection. This is visually apparent in Figure 4.7, as each projection method pulls participant strategy mixes closer to the line $\gamma + \delta = 1$.

In the 7-room condition, we observe considerable representational accuracy induced by both projection methods (Figure 4.8). Where participant strategy mixes drift from $\gamma + \delta = 1$ as δ approaches 1 when computing strategy mix using the invalid reflection, this does not occur if we define the antimatching strategy as the projection of the invalid reflection back to the simplex. This implies that when a reflection does not exist within the simplex, participants instead use a distribution that is near this invalid reflection. It is unclear why the projection back to the simplex induces better representational accuracy in the 7-room case than the 5-room case. It may be some idiosyncratic feature of the particular distributions tested, random variation in participant performance, or some theoretically interesting feature that differs between the probability simplex in 5-dimensions vs 7-dimensions.

While we may not be able to tell which method is better here, each projection method comes paired with distinct theoretical implications. For the uniform shift method, any internal representations of probability distributions need not require that the invalid vector reflection (of negative entries) ever exist. That is, the adjusted reflection may be computed geometrically by reflecting $\vec{p}$ over $\vec{u}$, but only following that trajectory until the boundary of the simplex is met. This would imply humans only represent what exists within the probability simplex in order to compute opposite probability distributions. Algorithmically: to find $\vec{q}$, the opposite of a given distribution $\vec{p}$, move $\vec{p}$ towards and beyond $\vec{u}$ until you double the distance from $\vec{p}$ to $\vec{u}$ or you hit the boundary of probability space, whichever occurs first.

Alternatively, the Euclidean simplex projection requires that the invalid $\vec{q}$ be explicitly represented, as $\vec{q}$ is required to find a point in the simplex $\vec{w}$ that minimizes the distance between $\vec{q}$ and $\vec{w}$. This would imply that humans can represent a broader interpretation of probabilities that are not constrained by non-negativity and a sum of 1. Under this hypothesis, probability could be represented as any other geometric/visuospatial problems are, and finding the closest point on the simplex is only a means to behaviorally express this representation in a task that is constrained by the formal definition of probability. Algorithmically: to find $\vec{q}$, the opposite of a given distribution $\vec{p}$, move $\vec{p}$ towards and beyond $\vec{u}$ until you double the distance from $\vec{p}$ to $\vec{u}$. Express the belief $\vec{w}$, which is the closest "true" probability distribution to $\vec{q}$.

The benefit of our analysis here is that future researchers can propose new methods to account for invalid reflections. The effectiveness of alternative models can be directly compared to each other by examining the quantitative differences between strategy vector coefficients, and model error.

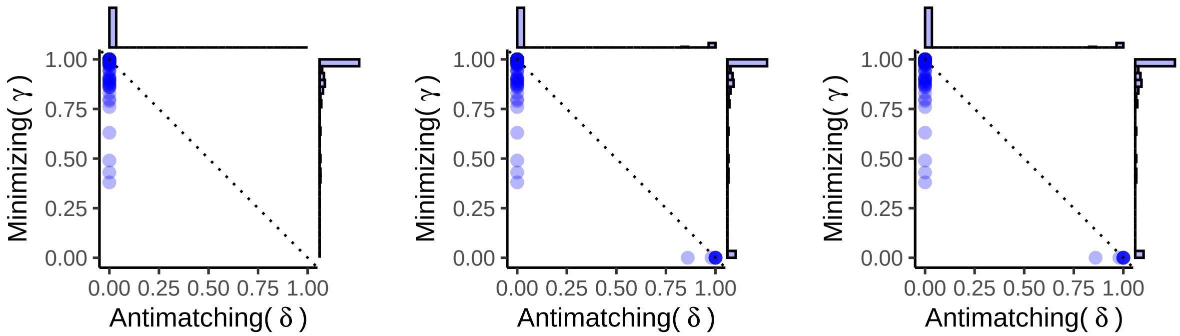

Figure 4.6: Applying original (left), unishift (centre), and closest simplex point (right) methods for accounting for invalid reflections in the 2D case.

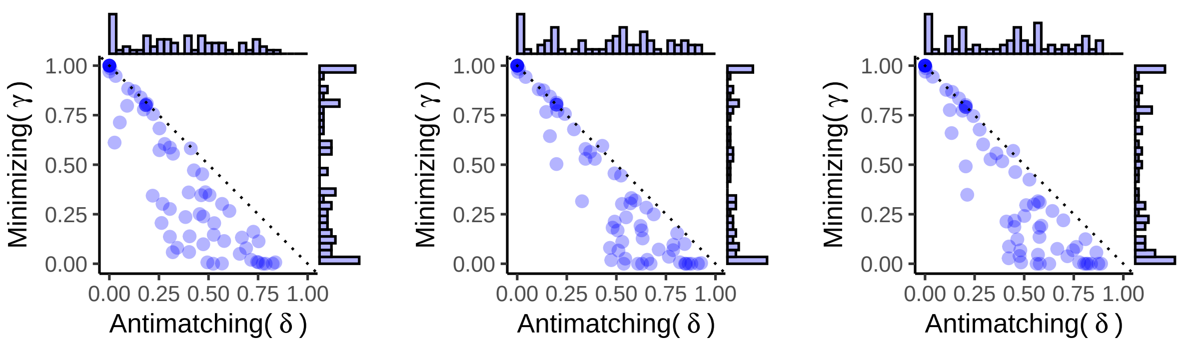

Figure 4.7: Applying original (left), unishift (center), and closest simplex point (right) methods for accounting for invalid reflections in the 5D case.

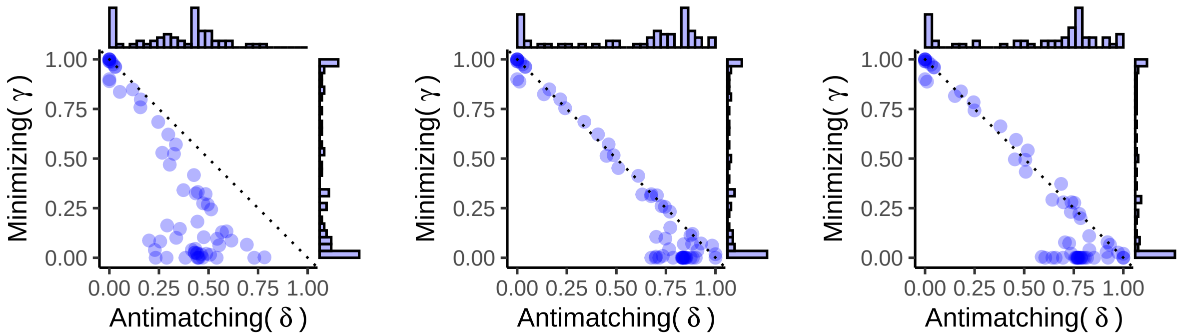

Figure 4.8: Applying original (left), unishift (center), and closest simplex point (right) methods for accounting for invalid reflections in the 7D case.

4.2.3 Hiding strategies are more optimal than seeking strategies, but less so in higher dimensions

Figure 4.9 compares participant use of the optimal strategy between hiding and seeking trials in the 3-, 5-, and 7-room conditions. In general, participants use an optimal strategy when hiding more than they do while seeking. However, this effect diminishes as room-count increases. This is mostly driven by a reduced use of the optimal hiding strategy as room-count increases. This finding is the opposite of what we expected following Study 1. Recall, initial findings found that participants were equally optimal in hiding and seeking in 2-room condition, but more optimal when hiding in the 5-room condition. We hypothesized that this was due to representational complexity, such that n-dimensional geometry is harder to represent than m-dimensional geometry where n > m. This would result in participants being disincentivized to compute the vector reflection required to utilize an antimatching strategy, and instead default to an 'easier' strategy that also turned out to be optimal in our task.

Under the representational complexity hypothesis, we would expect participants to bias more strongly to the optimal strategy in higher room (dimension) conditions. Moreover, this Study used a within-participant design for room condition. Under the representational complexity hypothesis, we would expect each individual participant to use strictly more of the optimal strategy when hiding in n-rooms, compared to m-rooms, where n > m. Given Figure 4.10, this is clearly not the case.

Future work is required to determine the relationship between optimal strategy use and dimensionality. It may be that representational complexity does not scale with dimensionality. While it is more difficult to imagine a tesseract than a cube, it is reasonable to concede that explicit geometric visualization is not the cognitive mechanism required to evade a stimulus that follows a particular pattern. It may be that with more room numbers, the perceived stakes of a failed hide are lessened, as the chance of being found in any particular room also decreases. Participants might therefore be less concerned with utilizing an optimal strategy. The finding that optimality asymmetry is negatively related to dimensionality is more consistant with an *assumed-payoff* hypothesis, where hiding strategies are generally more optimal than seeking strategies because participants hold prior beliefs that a failed hide is worse than a failed seek. Perhaps with more room options, the perceived importance of each individual trial is lessened because the overall probability of being found is lesser, thus reducing motivation to select an optimal strategy. This hypothesis can be tested in future studies by manipulating payoffs of successful/failed hide/seeks via gamified points, money, or additional course credit.

It's possible that each dimension has a strategy mix "signature". That is, can we identify the number of rooms from a given dataset only by inspecting the relationship between hide strategy mixes and seek strategy mixes? If geometric representations are in fact driving this cognition, then other geometric properties in addition to increasing dimensionality could be relevant. For example, we only tested odd-numbers of rooms in this study. The Euler characteristic (relating the number of vertices, edges, faces, etc.) for a simplex is 2 in odd dimensions, and 0 in even dimensions. It may be that participant strategy use will vary according the geometry that represents the task environment.

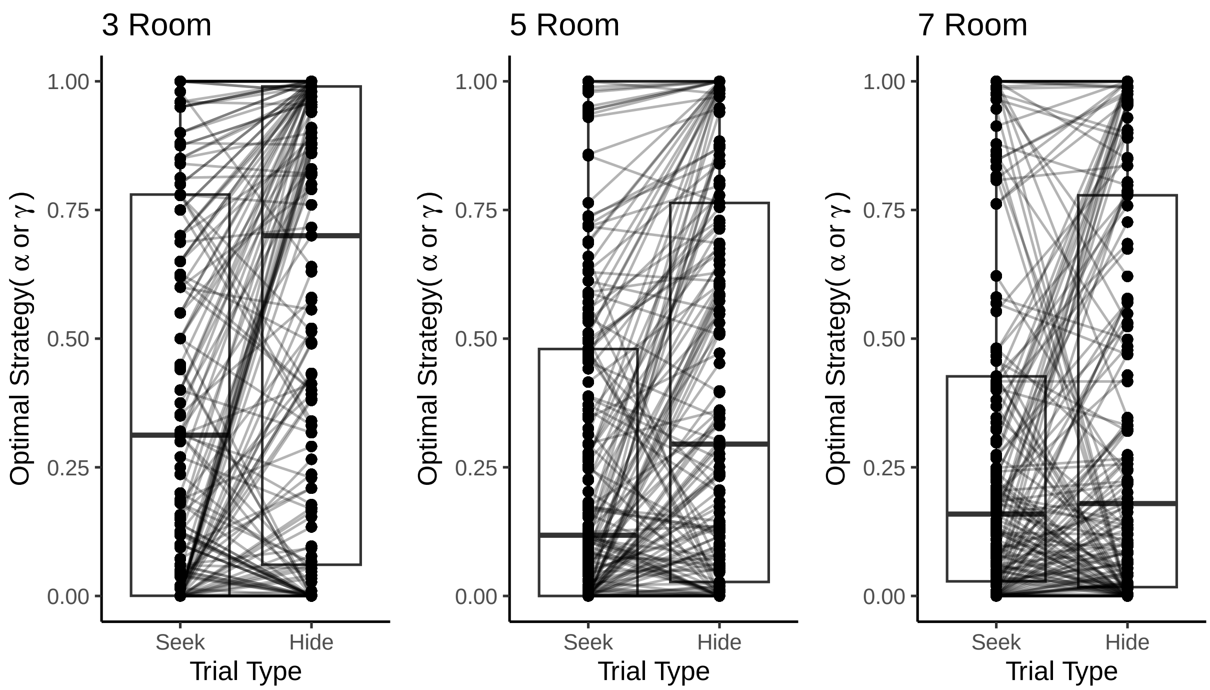

Figure 4.9: Participants may use a more optimal strategy mix when hiding than when seeking in the three room condition. The difference between hiding and seeking optimality is not observed as strongly in the five or seven room condition.

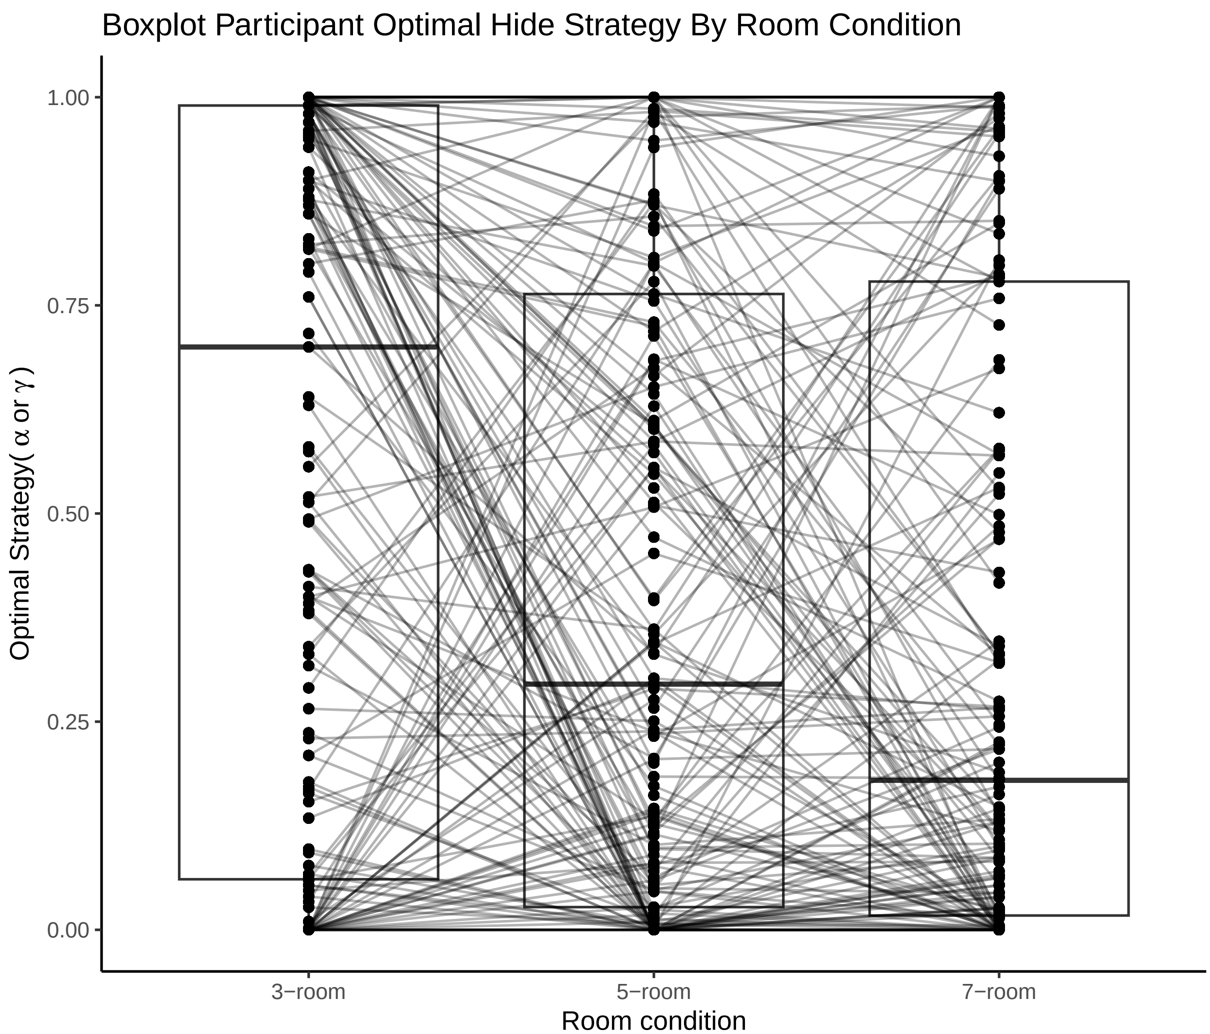

Figure 4.10: Participant use of the optimal hiding strategy does not increase as dimensionality increases. This evidence counters the representational complexity hypothesis, which states that participants will antimatch less (and optimize more) becasue computing the opposite probability distrubutions becomes difficult in higher dimensions.

Chapter 5

General Discussion

The aim of this work is to establish an analogous strategy space for stimiulus avoidance as exists for stimulus pursuit. Through our three hide-and-seek studies, we have found that participants use a mixture of the optimal maximizing strategy and the non-optimal matching strategy when seeking. When hiding, participants use a mixture of the optimial minimizing strategy, and the non-optimal antimatching strategy. The antimatching strategy when evading a stimulus generated from a probability distribution, expressed as a Euclidean vector $\vec{p}$, is defined as a reflection of $\vec{p}$ over the uniform distribution $\vec{u}$, such that $refl_{\vec{u}}(\vec{p}) = 2\vec{u} - \vec{p}$.

In Study 1, we first established the efficacy of the hide-and-seek task in measuring pursuit and avoidance strategy against a fixed probability distribution in an in-lab environment. Both between-subject conditions of 2-room and 5-room games found that participants use a nearly exclusive mix of maximizing and matching when seeking, and minimizing and antimatching when hiding. We also found that participants were more likely to use the optimal minimizing strategy when hiding, than the optimal maximizing strategy when seeking. However, this was only the case in the 5-room condition. The interaction between optimal-strategy-use-asymmetry and number of rooms motivated our hypothesis that participants have a natural propensity to use an antimatching matching strategy in the same sense as probability matching. However, with more rooms to consider, the difficulty of computing the antimatching strategy (via a vector reflecton) scales as dimensionality of the problem increases. This added cognitive demand may induce a shortcut – only hiding the single least likely room – a strategy that also happens to be optimal in the i.i.d. nature of our task.

Study 2 attempted to reproduce Study 1 using an online procedure. In this sample, we did not find differences between strategy use across hide/seek trials, or across room-number conditions. More importantly however, we found similar results to Study 1 with respect to strategy mixes. That is, we again found that participant hiding behaviour is just as well explained as mixture of minimizing and antimatching, as seeking behaviour is explained as a mixture of maximizing and matching.

In Study 3, we adopted a within-subjects design such that every participant faced at least one 3-room, 5-room, and 7-room condition. Some participants also faced distributions that had invalid reflections, allowing us to expand our mathematical definition of probability antimatching. Again, we found a near exclusive mix of maximizing and matching when seeking, and a near exclusive mix of minimizing and antimatching when hiding. We also found that when hiding against distributions with invalid reflections, participants' antimatching strategy may be selected as a valid distribution on the probability simplex that is "close" to the original invalid reflection. However, it is unclear how this "close" alternative is selected/produced as both of our candidate methodologies produced similar results. Like Study 1, and unlike Study 2, we found that participants utilize a more optimal strategy mix when hiding than when seeking. However, this difference reduces as room-number increases. Importantly, the results of these three studies indicate support for our primary interest – that probability antimatching defines hiding strategy analogously to how probability matching defines seeking strategy.

It is important to emphasize that our vector reflection model of probability antimatching is not a model of individual participant behaviour. That is, a histogram vector reflection does *not* act as a direct mapping from a participant's individual seeking strategy to their individual hiding strategy. Rather, our work demonstrates that participant seeking (stimulus pursuit) strategy is adequately defined as a mix of two strategies: traditional probability matching, and optimal maximizing. Moreover, we demonstrate that participant hiding behaviour is also adequately defined by a mix of two strategies: probability antimatching, which we define as probability histogram vector reflection, and the optimal minimizing strategy. We emphasize here that we have been able to articulate human stimulus avoidance strategies with just two dimensions (antimatching and minimizing) despite there existing no unambiguous mathematical definition an opposite probability distribution.

Across our three studies, we observe a pattern of findings where participants use a more optimal strategy when hiding than they do when seeking. The strength of this finding, and the manner in which it is affected by the number of room choices (dimensionality) is unclear. There exists an evolutionary basis to support this finding: a failed hide often carries greater consequence (eaten by a predator) than a failed seek (being hungry). This might induce a default payoff matrix that skews to value the outcomes of hide success over seek success. Future work can explore this idea further by modifying our hide-and-seek task to accommodate different payoffs, via arbitrary in-task points, monetary, or course credit. It may be that a certain distribution of payoffs would result in a persistent optimality asymmetry between hiding and seeking. Interesting implications may follow from this result if future work is able to better understand this relationship. Many problems of daily life can be presented in a pursuit/avoidance dichotomy. It may be that simply framing a problem as something to hide from, rather than something to pursue, will make people act more optimally. For example, would people make better financial decisions if the problem were framed as an avoidance of losses/costs (including opportunity cost) rather than a pursuit of gains [33, 25, 18]?

Our work presents a testable hypothesis for strategy utilization that accounts for both stimulus pursuit and stimulus evasion, therefore creating a meaningful extension to current probability matching literature. The formal geometric terms used to represent participant strategy mixes allows for richer analysis and quantitative theoretical development. For example, if another researcher proposes an alternative definition of antimatching, we can again represent participant strategy mix as a linear combination of Euclidean vectors and compare model fit via the magnitudes of strategy coefficients $(\alpha, \beta, \gamma, \delta)$, and error vectors $(\vec{\epsilon}, \vec{\zeta})$.

Theory expressed in mathematics promotes unifying theories between behaviour and the brain. Geometry has a long history in representing neural activity, cognition, and behaviour in general [12, 11, 7, 22]. The ability to express behaviour using our model, and brain activity in the same mathematical language allows for richer exploration of how the brain works. For example, consider a artificial neural network (ANN). This network can be trained to perform the same hide-and-seek task presented here. With both human and ANN behavioral output expressed mathematically, both can be compared directly to the task representations of the ANN's internal layers, which is already often analyzed geometrically as vectors. Various versions of this network could be built to emulate the structure of different candidate brain areas that drive phenonomenon of interest. For probability matching/antimatching, these may include the right dorsolateral prefrontal cortex, and left frontal and prefrontal areas, [26, 38]. The candidate model that best approximates human behavioural output, and perhaps, human brain activity recordings, may provide evidence to the neural underpinnings of the phenonomenon of interest.

With this work, we are not advocating that people are literally and explicitly computing vector reflections to determine their hiding strategies. Claiming this would imply that people are performing non-trivial algebra in their heads in order to complete our task. The benefit of a model with geometric interpretations is that it can be explained and conceptualized *without* algebra. It may be difficult to ask an individual to calculate square roots and evaluate trigonometric functions, but it is easier to imagine a point in space then imagine a new point that is double the distance away from you in the direction of the original point. This is similar to how catching a thrown ball is fundamentally different than solving a physics problem about a thrown ball. Our model with a built-in geometric interpretation, holds a "feasibility advantage" over other models such that the cognitive mechanisms required to emulate its output may already exist and be used for other visual-spatial capacities. Consider the analogy of Bayesian cognition. That is, the claim that people are "Bayesian" is infeasible as a literal claim since we cannot assume that people calculate the computationally difficult integrals in their head required by Bayes' rule, despite often presenting behavioral patterns that loosely emulate Bayesian models [17]. It may be that some other mechanism is the driving force of this emulating behaviour. Claiming people to be Bayesian is equivalent to claiming that people literally "use" vector reflections. While the geometric nature of vector representations provides some degree of feasibility over and above literal Bayesian cognition, we still stress that our model is effective in describing *behaviour*. Cognitive and neural extensions are left for future work.

References

[1] Anil Bhattacharyya. On a measure of divergence between two statistical populations defined by their probability distribution. *Bulletin of the Calcutta Mathematical Society*, 35:99–110, 1943.

[2] Martin Chapman, Gareth Tyson, Peter McBurney, Michael Luck, and Simon Parsons. Playing hide-and-seek: an abstract game for cyber security. In *Proceedings of the 1st International Workshop on Agents and CyberSecurity*, pages 1–8, 2014.

[3] Vincent P Crawford and Nagore Iriberri. Fatal attraction: Salience, naivete, and sophistication in experimental "hide-and-seek" games. *American Economic Review*, 97(5):1731–1750, 2007.

[4] Joshua R De Leeuw. jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior research methods*, 47:1–12, 2015.

[5] Peter A. V. DiBerardino, Alexandre L. S. Filipowicz, James Danckert, and Britt Anderson. Plinko: Eliciting beliefs to build better models of statistical learning and mental model updating, 2022.

[6] John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. Efficient projections onto the l 1-ball for learning in high dimensions. In *Proceedings of the 25th international conference on Machine learning*, pages 272–279, 2008.

[7] Chris Eliasmith. *How to build a brain: A neural architecture for biological cognition.* Oxford University Press, 2013.

[8] Alex Filipowicz, Britt Anderson, and James Danckert. Adapting to change: The role of the right hemisphere in mental model building and updating. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 70(3):201, 2016.

[9] József Fiser, Pietro Berkes, Gergő Orbán, and Máté Lengyel. Statistically optimal perception and learning: from behavior to neural representations. *Trends in cognitive sciences*, 14(3):119–130, 2010.

[10] Wolfgang Gaissmaier and Lael J Schooler. The smart potential behind probability matching. *Cognition*, 109(3):416–422, 2008.

[11] Peter Gardenfors. *Conceptual spaces: The geometry of thought*. MIT press, 2004.

[12] Apostolos P Georgopoulos, Joseph T Lurito, Michael Petrides, Andrew B Schwartz, and Joe T Massey. Mental rotation of the neuronal population vector. *Science*, 243(4888):234–236, 1989.

[13] Michel Goossens, Frank Mittelbach, and Alexander Samarin. *The LaTeX Companion*. Addison-Wesley, Reading, Massachusetts, 1994.

[14] DGC Harper. Competitive foraging in mallards:"ideal free'ducks. *Animal Behaviour*, 30(2):575–584, 1982.

[15] Derek J. Koehler and Greta James. Probability matching and strategy availability. *Memory & cognition*, 38(6):667–676, 2010.

[16] Philip N Johnson-Laird. Mental models and cognitive change. *Journal of Cognitive Psychology*, 25(2):131–138, 2013.

[17] Matt Jones and Bradley C Love. Bayesian fundamentalism or enlightenment? on the explanatory status and theoretical contributions of bayesian models of cognition. *Behavioral and brain sciences*, 34(4):169–188, 2011.

[18] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.

[19] Donald Knuth. *The TeXbook*. Addison-Wesley, Reading, Massachusetts, 1986.

[20] Derek J Koehler and Greta James. Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition*, 113(1):123–127, 2009.

[21] Derek J Koehler and Greta James. Probability matching, fast and slow. In *Psychology of learning and motivation*, volume 61, pages 103–131. Elsevier, 2014.

[22] Nikolaus Kriegeskorte and Rogier A Kievit. Representational geometry: integrating cognition, computation, and the brain. *Trends in cognitive sciences*, 17(8):401–412, 2013.

[23] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.

[24] Leslie Lamport. *LaTeX — A Document Preparation System*. Addison-Wesley, Reading, Massachusetts, second edition, 1994.

[25] Irwin P Levin, Sandra L Schneider, and Gary J Gaeth. All frames are not created equal: A typology and critical analysis of framing effects. *Organizational behavior and human decision processes*, 76(2):149–188, 1998.

[26] Michael B Miller, Monica Valsangkar-Smyth, Sarah Newman, Heather Dumont, and George Wolford. Brain activations associated with probability matching. *Neuropsychologia*, 43(11):1598–1608, 2005.

[27] Katharine M. Mullen. *bvls: The Stark-Parker algorithm for bounded-variable least squares*, 2013. R package version 1.4.

[28] Ben R Newell, Derek J Koehler, Greta James, Tim Rakow, and Don Van Ravenzwaaij. Probability matching in risky choice: The interplay of feedback and strategy availability. *Memory & Cognition*, 41:329–338, 2013.

[29] ngloe. olpsr - on-line portfolio selection with r - projection onto a simplex. https://github.com/ngloe/olpsR/blob/master/R/projsplx.R, 2015.

[30] Jonathan Peirce, Jeremy R Gray, Sol Simpson, Michael MacAskill, Richard Höchenberger, Hiroyuki Sogo, Erik Kastman, and Jonas Kristoffer Lindeløv. Psychopy2: Experiments in behavior made easy. *Behavior research methods*, 51:195–203, 2019.

[31] Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. The earth mover's distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99, 2000.

[32] David R Shanks, Richard J Tunney, and John D McCarthy. A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15(3):233–250, 2002.

[33] Dilip Soman. Framing, loss aversion, and mental accounting. *Blackwell handbook of judgment and decision making*, pages 379–398, 2004.

[34] Joshua B Tenenbaum, Charles Kemp, Thomas L Griffiths, and Noah D Goodman. How to grow a mind: Statistics, structure, and abstraction. *science*, 331(6022):1279–1285, 2011.

[35] Sharon L Thompson-Schill, Michael Ramscar, and Evangelia G Chrysikou. Cognition without control: When a little frontal lobe goes a long way. *Current directions in psychological science*, 18(5):259–263, 2009.

[36] Nir Vulkan. An economist's perspective on probability matching. *Journal of economic surveys*, 14(1):101–118, 2000.

[37] Richard F West and Keith E Stanovich. Is probability matching smart? associations between probabilistic choices and cognitive ability. *Memory & Cognition*, 31:243–251, 2003.

[38] George Wolford, Michael B Miller, and Michael Gazzaniga. The left hemisphere's role in hypothesis formation. *The Journal of Neuroscience*, 20(6):RC64, 2000.

APPENDICES

Appendix A

Mathematics of Vector Reflections

What is an opposite probability distribution? This question motivated the development of the mathematical theory this thesis is built upon. In this Appendix, we will formally define Histogram Vectors, Vector Reflections, and the relevant properties of both.

There is no mathematical definition for an "inverse" or "opposite" probability distribution. Given a histogram, we can derive a "reflected" histogram using vector projections.

Definition A.0.1. Histogram Vector
A vector $\vec{v} \in \mathbb{R}^n$ where each element v_i, for $1 \leq i \leq n$, represents the value of bin i of a histogram with n bins, is called a histogram vector.

Remark. Each element of a histogram vector must have a value greater than or equal to 0.

Example 1

The vector $\vec{v} = \begin{bmatrix} v_1 \\ v_2 \\ v_3 \\ v_4 \\ v_5 \end{bmatrix} = \begin{bmatrix} 10 \\ 10 \\ 20 \\ 35 \\ 25 \end{bmatrix} \in \mathbb{R}^5$ represents the following histogram:

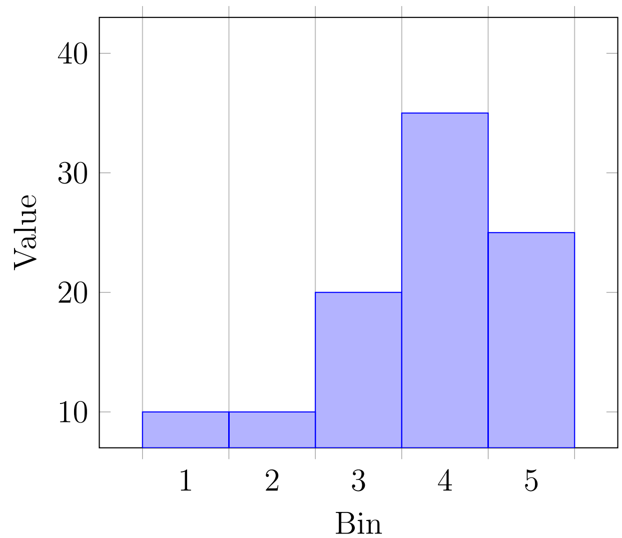

Remark. A histogram vector with all elements of equal value represents a uniform distribution.

Now that we can represent histograms as vectors with certain properties, we are able to use established tools and operations from linear algebra as fact.

Definition A.0.2. Vector Reflection
For vectors $\vec{u}$ and $\vec{v} \in \mathbb{R}^n$, the vector $\vec{w}$ where

$$\begin{aligned}\vec{w} &= \text{proj}_{\vec{u}}\vec{v} - \text{perp}_{\vec{u}}\vec{v} \\ &= \text{proj}_{\vec{u}}\vec{v} - (\vec{v} - \text{proj}_{\vec{u}}\vec{v}) \\ &= 2\text{proj}_{\vec{u}}\vec{v} - \vec{v} \\ &= 2\frac{\vec{v} \cdot \vec{u}}{\vec{u} \cdot \vec{u}}\vec{u} - \vec{v}\end{aligned}$$

is called the reflection of $\vec{v}$ over $\vec{u}$.

Remark. For vectors $\vec{u}$ and $\vec{v} \in \mathbb{R}^n$, $\text{proj}_{\vec{u}}\vec{v} + \text{perp}_{\vec{u}}\vec{v} = \vec{v}$. By subtracting $\text{perp}_{\vec{u}}\vec{v}$, from $\text{proj}_{\vec{u}}\vec{v}$, we get a vector $\vec{w}$ that has the "opposite" non-$\vec{u}$ component vector to $\vec{v}$. The equivalent expression of $2\text{proj}_{\vec{u}}\vec{v} - \vec{v}$ is derived from a simple algebraic manipulation of the definition of $\text{perp}_{\vec{u}}\vec{v}$ (see above).

Example 2

What is the reflection of $\vec{v} = \begin{bmatrix} 1 \\ 5 \end{bmatrix}$ over $\vec{u} = \begin{bmatrix} 4 \\ 4 \end{bmatrix}$?

Solution: By Definition 1.2, the reflection of $\vec{v} = \begin{bmatrix} 1 \\ 5 \end{bmatrix}$ over $\vec{u} = \begin{bmatrix} 4 \\ 4 \end{bmatrix}$ is

$$\begin{aligned}
\vec{w} &= 2\mathrm{proj}_{\vec{u}}\vec{v} - \vec{v} \\
&= 2\frac{\vec{v} \cdot \vec{u}}{\vec{u} \cdot \vec{u}}\vec{u} - \vec{v} \\
&= 2\frac{\begin{bmatrix} 1 \\ 5 \end{bmatrix} \cdot \begin{bmatrix} 4 \\ 4 \end{bmatrix}}{\begin{bmatrix} 4 \\ 4 \end{bmatrix} \cdot \begin{bmatrix} 4 \\ 4 \end{bmatrix}} \begin{bmatrix} 4 \\ 4 \end{bmatrix} - \begin{bmatrix} 1 \\ 5 \end{bmatrix} \\
&= 2\frac{24}{32} \begin{bmatrix} 4 \\ 4 \end{bmatrix} - \begin{bmatrix} 1 \\ 5 \end{bmatrix} \\
&= \begin{bmatrix} 6 \\ 6 \end{bmatrix} - \begin{bmatrix} 1 \\ 5 \end{bmatrix} \\
&= \begin{bmatrix} 5 \\ 1 \end{bmatrix}
\end{aligned}$$

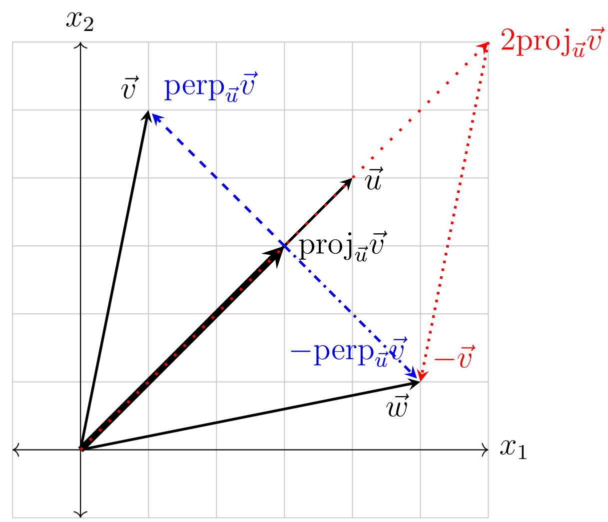

Comparing $\vec{v}$ to $\vec{w}$ from Example 2, we find that they appear to be "opposites" of sorts. In fact, we can show $\vec{v}$ and $\vec{w}$ are truly "reflections" of each other, with the angle from $\vec{v}$ to $\vec{u}$, equivalent to the angle from $\vec{u}$ to $\vec{w}$.

Theorem A.0.3. *If $\vec{v}$, $\vec{u}$ and $\vec{w} = 2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v} \in \mathbb{R}^n$, with θ_1 and θ_2 the angle between $\vec{v}$ and $\vec{u}$, and $\vec{u}$ and $\vec{w}$, respectively, then $\theta_1 = \theta_2$.*

Proof.

$$\|\vec{w}\|^2 = \vec{w}\cdot\vec{w} = (2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v}) \cdot (2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v}) \tag{A.1}$$

$$= 2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} \cdot (2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v}) - \vec{v} \cdot (2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v}) \tag{A.2}$$

$$= 4\frac{(\vec{v}\cdot\vec{u})^2}{\vec{u}\cdot\vec{u}} - 2\frac{(\vec{v}\cdot\vec{u})^2}{\vec{u}\cdot\vec{u}} - 2\frac{(\vec{v}\cdot\vec{u})^2}{\vec{u}\cdot\vec{u}} + \vec{v}\cdot\vec{v} \tag{A.3}$$

$$= \vec{v}\cdot\vec{v} \tag{A.4}$$

$$= \|\vec{v}\|^2 \tag{A.5}$$

$$\therefore \|\vec{w}\| = \|\vec{v}\| \qquad \text{(Since vector lengths are non-negative)} \tag{A.6}$$

Similarly:

$$\vec{w}\cdot\vec{u} = (2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v}) \cdot \vec{u} \tag{A.7}$$

$$= 2\vec{v}\cdot\vec{u} - \vec{v}\cdot\vec{u} \tag{A.8}$$

$$= \vec{v}\cdot\vec{u} \tag{A.9}$$

$$\therefore \vec{w}\cdot\vec{u} = \vec{v}\cdot\vec{u} \tag{A.10}$$

By the geometric definition of dot products:

$$\vec{v} \cdot \vec{u} = \|\vec{v}\|\|\vec{u}\| \cos(\theta_1) \qquad \text{and} \qquad \vec{w} \cdot \vec{u} = \|\vec{w}\|\|\vec{u}\| \cos(\theta_2) \tag{A.11}$$

$$\therefore \theta_1 = \arccos\left(\frac{\vec{v} \cdot \vec{u}}{\|\vec{v}\|\|\vec{u}\|}\right) \qquad \text{and} \qquad \theta_2 = \arccos\left(\frac{\vec{w} \cdot \vec{u}}{\|\vec{w}\|\|\vec{u}\|}\right) \tag{A.12}$$

$$\theta_1 = \arccos\left(\frac{\vec{v} \cdot \vec{u}}{\|\vec{v}\|\|\vec{u}\|}\right) \tag{A.13}$$

$$= \arccos\left(\frac{\vec{w} \cdot \vec{u}}{\|\vec{v}\|\|\vec{u}\|}\right) \qquad \text{(By 10)} \tag{A.14}$$

$$= \arccos\left(\frac{\vec{w} \cdot \vec{u}}{\|\vec{w}\|\|\vec{u}\|}\right) \qquad \text{(By 6)} \tag{A.15}$$

$$= \theta_2 \qquad \text{(By 12)} \tag{A.16}$$

$$\therefore \theta_1 = \theta_2 \tag{A.17}$$

$\square$

Given these established properties of vector reflections, we may be able to reflect a histogram vector over the uniform to find an "opposite" histogram.

A.1 Reflecting Histogram Vectors

Proposition A.1.1. *For any histogram vector* $\vec{v}$ *and the uniform histogram vector* $\vec{u} = \begin{bmatrix} a \\ \vdots \\ a \end{bmatrix} \in \mathbb{R}^n$, $\vec{w} = 2\dfrac{\vec{v} \cdot \vec{u}}{\vec{u} \cdot \vec{u}}\vec{u} - \vec{v}$ *is the histogram vector representing the "opposite" histogram represented by* $\vec{v}$.

Example 3

Consider the histogram vector $\vec{v} = \begin{bmatrix} 10 \\ 10 \\ 20 \\ 35 \\ 25 \end{bmatrix} \in \mathbb{R}^5$ from Example 1. By Proposition 2.1, the opposite histogram vector $\vec{w} = 2\dfrac{\vec{v} \cdot \vec{u}}{\vec{u} \cdot \vec{u}}\vec{u} - \vec{v} = \begin{bmatrix} 30 \\ 30 \\ 20 \\ 5 \\ 15 \end{bmatrix}$, which represents the following histogram:

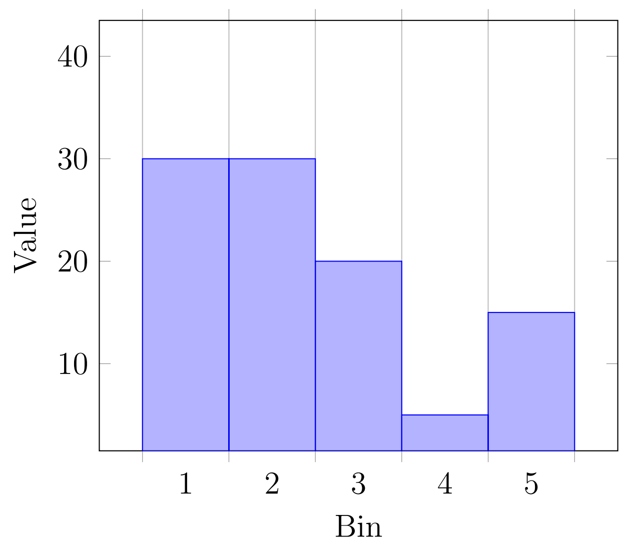

Remark. As long as the uniform vector $\vec{u}$ has all n elements of equal value, by the properties of vector projection, the particular value of the elements within $\vec{u}$ will not change the result of $\vec{w}$.

We also notice that the sum of all the bin values of the original histogram is equal to the sum of all the bin values of the opposite histogram. In fact, this is true in the general case.

Theorem A.1.2. *For any vector $\vec{v}$, uniform histogram vector $\vec{u} = \begin{bmatrix} a \\ \vdots \\ a \end{bmatrix}$, and reflected vector $\vec{w} = 2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v} \in \mathbb{R}^n$, $\sum_{i=1}^{n} v_i = \sum_{i=1}^{n} w_i$.*

Proof.

By the algebraic definition of dot product:

$$\begin{aligned}
\sum_{i=1}^{n} w_i &= \vec{w}\cdot\vec{1} && \text{(A.1)}\\
&= (2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v})\cdot\vec{1} && \text{(A.2)}\\
&= (2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}})\vec{u}\cdot\vec{1} - \vec{v}\cdot\vec{1} && \text{(A.3)}\\
&= \Big(2\frac{\sum_{i=1}^{n} a(v_i)}{\sum_{i=1}^{n} a^2}\Big)\sum_{i=1}^{n} a - \sum_{i=1}^{n} v_i && \text{(A.4)}\\
&= 2\frac{a\sum_{i=1}^{n} v_i}{na^2}na - \sum_{i=1}^{n} v_i && \text{(A.5)}\\
&= 2\sum_{i=1}^{n} v_i - \sum_{i=1}^{n} v_i && \text{(A.6)}\\
&= \sum_{i=1}^{n} v_i && \text{(A.7)}
\end{aligned}$$

□

This result is very important in the context of histogram vectors. We know that the sum of all bins of the original histogram is equal to the sum of all bins of the opposite histogram.

A.2 When Things Go Wrong

Unfortunately, not all histogram vector reflections work out so nicely. Consider the following example:

Example 4

Let $\vec{v} = \begin{bmatrix} 5 \\ 20 \\ 50 \\ 10 \\ 15 \end{bmatrix}$ and $\vec{u} = \begin{bmatrix} 20 \\ 20 \\ 20 \\ 20 \\ 20 \end{bmatrix}$. Therefore, $\vec{w} = 2\dfrac{\vec{v} \cdot \vec{u}}{\vec{u} \cdot \vec{u}}\vec{u} - \vec{v} = \begin{bmatrix} 35 \\ 20 \\ -10 \\ 30 \\ 25 \end{bmatrix}$.

Note that while the sum of all elements of $\vec{w}$ is still equal to the sum of all elements of $\vec{v}$, we have a negative element in $\vec{w}$ ($w_3 = -10$). This result disproves Proposition 2.1, since $\vec{w}$ contains a negative element and thus is not a valid histogram vector (by Definition 1.1).

As it turns out, there is a condition that each element v_i of a histogram vector $\vec{v}$ must satisfy so that the reflected vector $\vec{w} = 2\dfrac{\vec{v} \cdot \vec{u}}{\vec{u} \cdot \vec{u}}\vec{u} - \vec{v}$ is also a valid histogram vector. However, we must first uncover the structure of each element w_i of $\vec{w}$ in terms of the corresponding element v_i of $\vec{v}$.

$$\vec{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} = 2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v} \tag{A.1}$$

$$= 2\frac{\begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \cdot \begin{bmatrix} a \\ \vdots \\ a \end{bmatrix}}{\begin{bmatrix} a \\ \vdots \\ a \end{bmatrix} \cdot \begin{bmatrix} a \\ \vdots \\ a \end{bmatrix}} \begin{bmatrix} a \\ \vdots \\ a \end{bmatrix} - \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \tag{A.2}$$

$$= 2\frac{a(v_1 + ... + v_n)}{na^2} \begin{bmatrix} a \\ \vdots \\ a \end{bmatrix} - \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \tag{A.3}$$

$$= 2\frac{(v_1 + ... + v_n)}{na} \begin{bmatrix} a \\ \vdots \\ a \end{bmatrix} - \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \tag{A.4}$$

$$= \begin{bmatrix} \frac{2}{n}(v_1 + ... + v_n) \\ \vdots \\ \frac{2}{n}(v_1 + ... + v_n) \end{bmatrix} - \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \tag{A.5}$$

$$= \begin{bmatrix} \frac{2}{n}(v_1 + ... + v_n) - v_1 \\ \vdots \\ \frac{2}{n}(v_1 + ... + v_n) - v_n \end{bmatrix} \tag{A.6}$$

$$\therefore w_i = \frac{2}{n}(v_1 + ... + v_n) - v_i \text{ for } 1 \leq i \leq n \tag{A.7}$$

Theorem A.2.1. *If vectors* $\vec{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}$, $\vec{u} = \begin{bmatrix} a \\ \vdots \\ a \end{bmatrix}$, *and* $\vec{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} \in \mathbb{R}^n$, *such that* $\vec{w} = 2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v}$, *then* $w_i \geq 0$ *if and only if* $v_1+, \ldots, v_{i-1} + v_{i+1} + \ldots + v_n \geq (\frac{n}{2} - 1)v_i$, *for every* v_i *where* $1 \leq i \leq n$.

Proof.

Assume $w_i \geq 0 \qquad (1 \leq i \leq n)$

$$\tag{A.1}$$

$$\iff \qquad \frac{2}{n}(v_1 + ... + v_n) - v_i \geq 0 \tag{A.2}$$

$$\iff \qquad \frac{2}{n}v_1 + ... + \frac{2-n}{n}v_i + ... + \frac{2}{n}v_n \geq 0 \tag{A.3}$$

$$\iff \qquad \frac{2}{n}v_1 + ... + \frac{2}{n}v_{i-1} + \frac{2}{n}v_{i+1} + ... + \frac{2}{n}v_n \geq \frac{n-2}{n}v_i \tag{A.4}$$

$$\iff \qquad \frac{1}{n}v_1 + ... + \frac{1}{n}v_{i-1} + \frac{1}{n}v_{i+1} + ... + \frac{1}{n}v_n \geq \frac{\left(\frac{n}{2} - 1\right)}{n}v_i \tag{A.5}$$

$$\iff \qquad v_1 + ... + v_{i-1} + v_{i+1} + ... + v_n \geq \left(\frac{n}{2} - 1\right)v_i \text{ (Since } n \geq 0) \tag{A.6}$$

□

By replacing all $\geq$ with $=$, we immediately get the following result:

Corollary A.2.1.1. *If vectors* $\vec{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}$, $\vec{u} = \begin{bmatrix} a \\ \vdots \\ a \end{bmatrix}$, *and* $\vec{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} \in \mathbb{R}^n$, *such that* $\vec{w} = 2\frac{\vec{v} \cdot \vec{u}}{\vec{u} \cdot \vec{u}}\vec{u} - \vec{v}$, *then* $w_i = 0$ *if and only if* $v_1 +, \ldots, v_{i-1} + v_{i+1} + \ldots + v_n = (\frac{n}{2} - 1)v_i$, *for every* v_i *where* $1 \leq i \leq n$.

Example 5

Consider the vector $\vec{v}$ $\begin{bmatrix} 10 \\ 10 \\ 20 \\ 35 \\ 25 \end{bmatrix} \in \mathbb{R}^5$ from Example 1 and 3. We know $\vec{v}$ had a valid opposite histogram, therefore by Theorem 3.1, it must satisfy

$v_1 + ... + v_{i-1} + v_{i+1} + ... + v_5 \geq \left(\frac{5}{2} - 1\right) v_i$ for $1 \leq i \leq 5$.

$$
\begin{aligned}
i = 1 \qquad\qquad & v_2 + v_3 + v_4 + v_5 \geq \left(\frac{5}{2} - 1\right) v_1 \\
& 10 + 20 + 35 + 25 \geq \left(\frac{3}{2}\right) 10 \\
& 90 \geq 15 \quad \checkmark \\
i = 2 \qquad\qquad & v_1 + v_3 + v_4 + v_5 \geq \left(\frac{5}{2} - 1\right) v_2 \\
& 10 + 20 + 35 + 25 \geq \left(\frac{3}{2}\right) 10 \\
& 90 \geq 15 \quad \checkmark \\
i = 3 \qquad\qquad & v_1 + v_2 + v_4 + v_5 \geq \left(\frac{5}{2} - 1\right) v_3 \\
& 10 + 10 + 35 + 25 \geq \left(\frac{3}{2}\right) 20 \\
& 80 \geq 30 \quad \checkmark \\
i = 4 \qquad\qquad & v_1 + v_2 + v_3 + v_5 \geq \left(\frac{5}{2} - 1\right) v_4 \\
& 10 + 10 + 20 + 25 \geq \left(\frac{3}{2}\right) 35 \\
& 65 \geq 52.5 \quad \checkmark \\
i = 5 \qquad\qquad & v_1 + v_2 + v_3 + v_4 \geq \left(\frac{5}{2} - 1\right) v_5 \\
& 10 + 10 + 20 + 35 \geq \left(\frac{3}{2}\right) 25 \\
& 75 \geq 37.5 \quad \checkmark
\end{aligned}
$$

Example 6

Consider the vector $\vec{v} = \begin{bmatrix} 5 \\ 20 \\ 50 \\ 10 \\ 15 \end{bmatrix}$ from Example 4. We know $\vec{v}$ did not have a valid opposite histogram $\vec{w}$, since $w_3 = -10$. Therefore, by Theorem 3.1, v_3 did not satisfy the condition.

We can verify this is the case:

$$i = 3 \qquad\qquad \begin{aligned} v_1 + v_2 + v_4 + v_5 &\geq \left(\frac{5}{2} - 1\right) v_3 \\ 5 + 20 + 10 + 15 &\geq \left(\frac{3}{2}\right) 50 \\ 50 &\ngeq 75 \end{aligned}$$

A.3 Adjusting Reflections

In the previous section, we discovered that not all reflected histogram vectors create valid opposite histogram vectors because the resulting vector may contain a negative element. However, we also discovered the conditions that must be satisfied by the original vector in order to create a successful reflection. In this section, we will explore a possible adjustment algorithm to make any reflected histogram vector a valid opposite histogram vector.

Lemma A.3.1. *For any vector* $\vec{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix}$ *and* $\vec{c} = \begin{bmatrix} -min(w_1, ..., w_n) \\ \vdots \\ -min(w_1, ..., w_n) \end{bmatrix} \in \mathbb{R}^n$, $\vec{x} = \vec{w} + \vec{c}$ *is a histogram vector* ($x_i \geq 0$ *for* $1 \leq i \leq n$).

Proof.

$$\text{Consider } \vec{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_{m-1} \\ w_m \\ w_{m+1} \\ \vdots \\ w_n \end{bmatrix} \text{ where } w_m = min(w_1, ..., w_n). \tag{A.1}$$

$$\therefore \vec{x} = \begin{bmatrix} w_1 \\ \vdots \\ w_{m-1} \\ w_m \\ w_{m+1} \\ \vdots \\ w_n \end{bmatrix} + \begin{bmatrix} -min(w_1, ..., w_n) \\ \vdots \\ -min(w_1, ..., w_n) \\ -min(w_1, ..., w_n) \\ -min(w_1, ..., w_n) \\ \vdots \\ -min(w_1, ..., w_n) \end{bmatrix} \tag{A.2}$$

$$= \begin{bmatrix} w_1 \\ \vdots \\ w_{m-1} \\ w_m \\ w_{m+1} \\ \vdots \\ w_n \end{bmatrix} + \begin{bmatrix} -w_m \\ \vdots \\ -w_m \\ -w_m \\ -w_m \\ \vdots \\ -w_m \end{bmatrix} \tag{A.3}$$

$$= \begin{bmatrix} w_1 - w_m \\ \vdots \\ w_{m-1} - w_m \\ 0 \\ w_{m+1} - w_m \\ \vdots \\ w_n - w_m \end{bmatrix} \tag{A.4}$$

$\therefore \vec{x}$ is a valid histogram vector since w_m is the minimum value and thus $w_i - w_m \geq 0$ for $1 \leq i \leq n$. $\square$

Lemma 4.1 allows us to subtract the value of the "maximally negative" element from all elements in a vector, creating a new vector with no negative values. In the context of histogram vectors, if a reflected histogram vector is not valid (has negative element(s)), we can add the most negative value to all elements to create a valid histogram vector.

Definition A.3.2. For any histogram vector $\vec{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}$, uniform histogram vector $\vec{u}$, and $\vec{w} = 2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v} \in \mathbb{R}^n$, we define $\vec{c} = \begin{bmatrix} -min(w_1, ..., w_n) \\ \vdots \\ -min(w_1, ..., w_n) \end{bmatrix} \in \mathbb{R}^n$ as the "Correction Vector" of $\vec{w}$ and $\vec{x} = \vec{w} + \vec{c}$ as the "Corrected Histogram Vector" of $\vec{w}$.

However, by correcting an invalid histogram vector reflection by using $\vec{c}$ to create $\vec{x}$, we lose many of the nice properties we established earlier.

Example 7

Consider $\vec{v} = \begin{bmatrix} 5 \\ 20 \\ 50 \\ 10 \\ 15 \end{bmatrix}$, $\vec{u} = \begin{bmatrix} 20 \\ 20 \\ 20 \\ 20 \\ 20 \end{bmatrix}$, and $\vec{w} = 2\frac{\vec{v}\cdot\vec{u}}{\vec{u}\cdot\vec{u}}\vec{u} - \vec{v} = \begin{bmatrix} 35 \\ 20 \\ -10 \\ 30 \\ 25 \end{bmatrix}$ from Example 4. According to Lemma 4.1, $\vec{c} = \begin{bmatrix} -min(w_1, ..., w_n) \\ \vdots \\ -min(w_1, ..., w_n) \end{bmatrix} = \begin{bmatrix} 10 \\ 10 \\ 10 \\ 10 \\ 10 \end{bmatrix}$, and $\vec{x} = \vec{w} + \vec{c} = \begin{bmatrix} 35 \\ 20 \\ -10 \\ 30 \\ 25 \end{bmatrix} + \begin{bmatrix} 10 \\ 10 \\ 10 \\ 10 \\ 10 \end{bmatrix} = \begin{bmatrix} 45 \\ 30 \\ 0 \\ 40 \\ 35 \end{bmatrix}$.
Since we are now working with $\vec{v}$ and $\vec{x}$, rather than $\vec{v}$ and $\vec{w}$, Theorems 1.3 and 2.2 no longer apply. In this example, we find:

$$\sum_{i=1}^{n} x_i = 150 \neq 100 = \sum_{i=1}^{n} v_i$$

Remark. We can see that in the general case,
$\sum_{i=1}^{n} x_i = \left(\sum_{i=1}^{n} w_i\right) - n \times min(w_1, ..., w_n)$ since $\vec{c}$ subtracts $min(w_1, ..., w_n)$ from all n elements of $\vec{w}$ to create $\vec{x}$.

Similarly, we can also see the angle between $\vec{v}$ and $\vec{u}$ (θ_1) is not equal to the angle between $\vec{u}$ and $\vec{x}$ (θ_2):

$$\theta_1 = \arccos\left(\frac{\vec{v}\cdot\vec{u}}{\|\vec{v}\|\|\vec{u}\|}\right) \qquad \theta_2 = \arccos\left(\frac{\vec{u}\cdot\vec{x}}{\|\vec{u}\|\|\vec{x}\|}\right)$$
$$\theta_1 = \arccos\left(\frac{2000}{\sqrt{3250}\times\sqrt{2000}}\right) \qquad \theta_2 = \arccos\left(\frac{3000}{\sqrt{2000}\times\sqrt{5750}}\right)$$
$$\theta_1 \approx 0.67 \qquad \theta_2 \approx 0.49$$

$\therefore \theta_1 \neq \theta_2$

Remark. It is important to note that while our corrected vector $\vec{x}$ may no longer share the same propeties as $\vec{v}$, since our correction vector $\vec{c}$ contains all the same elements, it is a scalar multiple of the uniform vector $\vec{u}$. Therefore, like $\vec{w}$, $\vec{x}$ can be written as a linear combination of $\vec{v}$ and $\vec{u}$ and $\in \text{Span}\{\vec{v},\vec{u}\}$.

A.4 Conclusion

We have seen how histograms can be represented as vectors called histogram vectors. Reflecting a histogram vector across a uniform histogram vector gives us an "opposite" histogram. These vectors and their reflections have similar properties, such as equal sum of all elements (Theorem 1.3), and equal angle from the uniform (Theorem 2.2). However, in some cases (see Theorem 3.1), our reflected vector contains a negative value and thus is not a valid histogram. This can be corrected by subtracting the "maximally negative" component from each element of the reflected vector. However, this method no longer ensures the original vector will have the same properties as its corrected opposite.