An Analysis of Penalty Kick Behaviour

Yazdy Mistry

A research component submitted to Auckland University of Technology in partial fulfilment of the requirements for the degree of Master of Business

2026

Department of Economics and Finance

Abstract

This thesis examines the penalty kick game within a behavioural game theory context to better understand the decision making rationale behind players' actions. There is currently debate to the existence of goalkeeper 'Action Bias' where goalkeepers are shown to dive to the sides rather than staying in the centre, even when staying in the centre more often may be strategically optimal. This research examines the dataset collected by (Bar-Eli et al., 2007) where the researchers find evidence for action bias. This thesis formally parametrises and models action bias, which has not been done before in previous research. This thesis also explores logit best responding through the use of the Quantal Response Equilibrium framework to test if bounded rationality better corresponds to real world play as opposed to perfect rationality assumed under the commonly used Mixed Strategy Nash Equilibrium framework. This research finds evidence of action bias within the primary dataset and that the inclusion of action bias within specified payoffs along with bounded rationality provide the most explanatory power out of all the models tested. Supplementary testing finds no evidence for behavioural heterogeneity. The penalty kick game cannot be treated as a singular homogeneous game across all observations, and all results must be treated as dataset dependent. This research finds there is evidence of goalkeeper action bias within the primary dataset as well as evidence that players are shown to not be acting as perfectly rational utility maximisers.

Contents

Abstract 1
List of figures 5
List of tables 6
Attestation of Authorship 8
Acknowledgements 10
(1) Introduction 11
(2) Literature review 12
(2.1) Testing for mixed strategy equilibria in professional sports 12
(2.1.1) Research from tennis 12
(2.2) Penalty kicks 13
(2.3) MSNE as a benchmark for penalty kick research 13
(2.3.1) Bar-Eli et al. (2007) 13
(2.3.2) Azar and Bar-Eli (2011) 14
(2.3.3) Replication study counterclaim: Singh and Mukherjee (2026) 16
(2.3.4) Heterogeneity and aggregation issue, Chiappori et al. (2002) 16
(2.4) Synthesis on existing penalty kick literature 18
(2.5) Restrictions of MSNE as a descriptive model 18
(2.6) Logit QRE as an alternative behavioural model 19
(2.6.1) Evidence comparing MSNE and QRE in related games 19
(2.7) Research gap and contribution 22
(3) Data setup and previous methodologies 24
(3.1) Non-independence within the data 26
(3.2) Sequential play theory 27
(3.3) Kicker footedness 29
(4) MSNE analysis 33
(4.1) Action Bias implementation 37
(4.2) Parameter estimations 43
(4.3) Heterogeneous models 46

(4.3.1) Heterogeneity in ε 46
(4.3.2) Other ε distributions 50
(4.3.3) Interpretation of fitted MLE parameters 51
(4.3.4) Incomplete information assumption 55
(4.3.5) MSNE model summary 55
(5) Quantal response equilibrium 59
(5.1) QRE background and comparative statics 59
(5.1.1) General logit QRE behaviour 59
(5.2) Matching pennies example 62
(6) QRE application to Azar et al. 67
(6.1) Base Logit QRE 67
(6.2) QRE with action bias 74
(6.3) Parameter estimation 76
(6.4) Heterogeneity in ε 79
(6.4.1) Complete information QRE 79
(6.4.2) Incomplete information QRE 80
(6.5) Common λ restriction 80
(6.6) Heterogeneity in ε redundancy 83
(7) Azar et al. model comparison 85
(7.1) Action bias specification 85
(7.1.1) Conclusion on ε specification 87
(7.2) Lambda specification 87
(7.3) Comparison between preferred MSNE and preferred QRE 88
(7.4) Overall model comparison and narrative 89
(8) Singh and Mukherjee comparison 91
(8.1) Model estimations 91
(8.2) Conflicting results between the Azar et al. and Singh and Mukherjee estimations. 94
(9) Behavioural discussion 96

(9.1) Perfect versus bounded rationality 96
(9.2) Goalkeeper action bias 96
(9.3) Limitations of behavioural conclusions 97
(10) Conclusion 98
(11) References 99
(12) Appendix 101
(12.1) Derivation of ε critical value 101
(12.2) Unique equilibria testing 102
(12.3) Player vs joint likelihood 109
(12.4) QRE model results table 111
(12.5) Incomplete information result 113
(12.5.1) Proof of proposition 2 113

List of figures

Figure 1 - (Azar & Bar-Eli, 2011) Observed data 26
Figure 2 – MSNE vs observed data 35
Figure 3 – Player Best response regions 36
Figure 4 – Animated Best response regions 41
Figure 5 – Animated MSNE predictions given ε 42
Figure 6 – Equilibrium mixes at $\varepsilon*$ 43
Figure 7 – MSNE action bias fitted mixes 46
Figure 8 – Heterogeneous MSNE kicker mixes 52
Figure 9 – MSNE model comparison 56
Figure 10 – Two strategy logit best response curve 60
Figure 11 – Two strategy logit best response curve animated 62
Figure 12 – Matching pennies MSNE response curves 63
Figure 13 – Logit QRE matching pennies equilibria 64
Figure 14 – Traced logit equilibria for matching pennies 65
Figure 15 – Comparative statics as x increases 66
Figure 16 – Base game QRE equilibrium paths 68
Figure 17 – Base game goalie equilibrium path against fixed kicker 69
Figure 18 – Base game goalkeeper feasibility region 70
Figure 19 – Base game kicker feasibility region 71
Figure 20 – Base game QRE likelihood heatmap 72
Figure 21- Goalkeeper feasibility region as ε varies 75
Figure 22 – Fitted QRE plot 76
Figure 23 – Goalkeeper feasibility region at MLE ε 77
Figure 24 – Kicker feasibility region at MLE ε 78
Figure 25 – QRE fitted model comparison 82
Figure 26 – MLE equilibria as ε varies 83
Figure 27 – Model comparison flowchart 89
Figure 28 - (Singh & Mukherjee, 2026) Best response regions 92
Figure 29 - (Singh & Mukherjee, 2026) QRE model comparison 93
Figure 30 – Best response plot under full support 102
Figure 31 – Best response plot for knifes edge case 104
Figure 32 – Best response plot for reduced support case 106

List of tables

Table 1 – Tested strategic conjectures by (Azar & Bar-Eli, 2011) 15
Table 2 – (Azar & Bar-Eli, 2011) Joint distribution of jumps and kicks 24
Table 3 - (Azar & Bar-Eli, 2011) Scoring rates 24
Table 4 – Goalie diving vs staying centre scoring rates 25
Table 5 – Assumed independence joint action table 27
Table 6 – Percentage over/underrepresentation relative to independence 27
Table 7 - (Chiappori et al., 2002) over/underrepresentation 28
Table 8 - (Singh & Mukherjee, 2026) Over/underrepresentation 28
Table 9 – Left footed kicker table 30
Table 10 – Right footed kicker table 30
Table 11 – Predicted joint frequency under footedness assumption 31
Table 12 – Payoff matrix 33
Table 13 – Expected payoffs 33
Table 14 – Predicted MSNE distribution 34
Table 15 – Marginal scoring probabilities given observed frequencies 37
Table 16 – Action bias across whole action 38
Table 17 – Action bias only when goal is scored 38
Table 18 – Fitted MSNE results 44
Table 19 – Fitted MSNE results including MLE 45
Table 20 – Predicted mixes at MLE 45
Table 21 – Heterogeneous MSNE fitted parameters 48
Table 22 – Low type predicted MSNE mixes 48
Table 23 – High type predicted MSNE mixes 48
Table 24 – Pooled heterogeneous MSNE mixes 48
Table 25 – Player vs joint MSNE fit 49
Table 26 – Conditional goalie frequency rates 53
Table 27 – Conditional goalie predicted rates 53
Table 28 – Incomplete information MSNE mixes 55
Table 29 – MSNE model marginal predictions 56
Table 30 – Base MSNE cell likelihood 57
Table 31 – MSNE homogeneous action bias cell likelihood 57
Table 32 – MSNE heterogeneous CI action bias cell likelihood 57
Table 33 –Fitted QRE parameters 76

Table 34 – Predicted QRE marginal mixes 76
Table 35 – Heterogeneous QRE fitted parameters 79
Table 36 – Common λ fitted parameters 80
Table 37 – Common vs separate λ fitted parameters 82
Table 38 – MSNE model comparison 85
Table 39 – AIC and BIC testing for MSNE models 86
Table 40 – QRE fitted model comparison 87
Table 41 – Common vs separate λ BIC 88
Table 42 – Preferred QRE vs MSNE 88
Table 43 – QRE vs MSNE bootstrap test 88
Table 44 - (Singh & Mukherjee, 2026) Joint frequency table 91
Table 45 - (Singh & Mukherjee, 2026) Stopped kicks 91
Table 46 – BIC results (Azar & Bar-Eli, 2011) and (Singh & Mukherjee, 2026) datasets 94
Table 47 – Preferred BIC model (Azar & Bar-Eli, 2011) and (Singh & Mukherjee, 2026) 94
Table 48- Possible support testing for $\varepsilon < \varepsilon *$ 103
Table 49 – Possible support testing for $\varepsilon = \varepsilon *$ 105
Table 50 - Possible support testing for $\varepsilon > \varepsilon *$ 107
Table 51 - QRE MLE table 112

Attestation of Authorship

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person (except where explicitly defined in the acknowledgements), nor used artificial intelligence tools or generative artificial intelligence tools (unless it is clearly stated, and referenced, along with the purpose of use), nor material which to a substantial extent has been submitted for the award of any other degree or diploma of a university or other institution of higher learning.

Yazdy Mistry 

30th September 2026

ARTIFICIAL INTELLIGENCE DECLARATION

Specify the chapter number(s) in your research proposal/research component where AI was used.	
Chapter number(s):	**Chapters 1-9 and the appendices**
Briefly explain why the AI tool was used (e.g., idea generation, summarization, data analysis, image creation, etc.).	
Purpose of AI Use:	AI supported statistical programming and verification, plot coding, language editing, clearer explanations, and consistency checks. It was not used to select the research question, provide source evidence, or make final analytical judgements.
Clearly state the AI tools used.	
AI Tool(s) Used:	**OpenAI codex**
Prompts or queries entered into the AI tool (only the instructive portion of the prompt is required). Include all relevant variations.	
Prompts Used:	• Prompts and variations included: • Run an MLE over the following search parameters • Create a simplex plot that shows this resulting data • Check the thesis against the AUT Postgraduate Handbook and identify required formatting changes. • Calculate or verify likelihood-ratio comparisons between QRE models with and without action bias. • Review statistical code, equations, tables, outputs, terminology, and explanations for consistency.
A summary or description of the AI-generated output. (Optional: attach output in an appendix.)	
Output Received:	Outputs included fitted model parameters, recalculated statistics, figures and tables, alternative wording and structure, explanations of feedback, code and diagnostic checks, and a submission review.
Describe how the AI output was used, edited, or transformed.	
Post-AI Processing Methods:	I reviewed every output, rejected or made suggestions as necessary, checked factual claims against cited sources, and checked calculations against the original data and code. I retained responsibility for the research design, model choices, interpretations, and conclusions. All AI assisted plots were checked, and all equilibrium search code outputs were checked for robustness and numerical errors.

You may use this free text box to explain the intent for the use of the AI generated content. *For example, rather than stating what prompt(s) were used, you may articulate the purpose for which the prompt was used (supported by examples) here:*

AI was used for editing and computational support: calculating equilibria, plot making, clearer wording, responses to feedback, alternative explanations. The data, models, literature interpretation, and conclusions were all completed by me. AI was used primarily as a helper, for tasks like suggesting relevant literature (each paper was then read by me and relevance manually checked), helping with succinctness, and coding tasks. All outputs were manually checked and often changed before any plots or numerical results were included.

STUDENT DECLARATION

By signing you are confirming that the AI use stated in the table(s) above are accurate and follow AUT's recommended guidelines detailed in the AI Hub and Postgraduate Handbook.

Student Name: **Yazdy Mistry** **Signature:** **Date:** **28/09/2026**

This AI Declaration has been reviewed and endorsed by my supervisor, Matthew Ryan.

Acknowledgements

I would like to express my sincere gratitude to my supervisor, Professor Matthew Ryan, for his guidance, patience, and support throughout this research. His careful feedback, thoughtful questions, and willingness to deepen my reasoning greatly improved both this thesis and my understanding of the subject. I am especially grateful for the time and expertise he devoted to helping me refine the theoretical analysis and communicate my findings more clearly. This thesis would not have been possible without his encouragement and supervision.

(1) Introduction

This research sets out to analyse the strategic decision making of professional footballers when taking penalty kicks. This type of situation is naturally suited to game theory. The kicker and goalkeeper both have to make a decision at almost the same time, and neither player can know for certain what the other will do. If the kicker always chooses the same direction, the goalkeeper could exploit this. Similarly, if the goalkeeper always dived the same way every time, the kicker could respond by kicking somewhere else. Because of this, both players have an incentive to randomise their choices. This is why penalty kicks are often modelled using a Mixed Strategy Nash Equilibrium.

However, the observed data collected in previous research does not fully match what we might expect from the standard MSNE model. Previous research on penalty kicks finds goalkeepers heavily prefer diving left or right and stay in the centre only a small proportion of the time. Previous penalty kick research posits this phenomenon is caused by *action bias,* where the goalie feels pressure to dive to either side of the goal even when staying in the centre may be optimal. Survey evidence suggests that this is a true bias felt by goalkeepers, yet no current research formally models or estimates this bias.

This implies the existing models used in penalty kick research may not be adequately capturing players' true incentives. This research will model penalty kicks using behavioural models which have not yet been applied to this research space.

The central purpose of this research is to explore whether observed penalty kick behaviour is consistent with goalkeeper action bias, and to examine how the effect of this bias depends on equilibrium concept and parameter specification.

The analysis separates three potential explanations for departures from the base MSNE predictions: altered goalkeeper utility through action bias, stochastic payoff responsiveness through QRE, and unobserved player heterogeneity. The models are evaluated using likelihood comparisons and are then applied to an independent dataset to assess whether the conclusions are robust across samples.

(2) Literature review

(2.1) Testing for mixed strategy equilibria in professional sports.

Professional sports provide a useful setting for examining strategic decision making because they are naturally occurring, high stakes, competitive situations, with quantifiable strategies. This creates a situation suitable for game theory analysis, where each player's best action depends on what they believe their opponent will do.

Penalty kicks is a game where one player is trying to match the direction of their opponent, while the other player is trying to choose a direction away from their opponent. Existing research also shows evidence that penalty kicks are a simultaneous move game. Therefore, we treat the penalty kick game as a variant of the matching pennies game with the goalie being the matching player and the kicker being the non-matching player.

Tennis serves are a comparable game to penalty kicks as professional players choose directions strategically and benefit from being unpredictable. Thus, before analysing penalty kicks we can look at studies conducted on tennis serves to gain supplementary insight into the penalty kick game.

(2.1.1) Research from tennis

(Walker & Wooders, 2001) examine the question 'do players follow Nash Equilibrium strategies in real world play'. By using professional tennis serve data from Wimbledon they test a key implication of mixed strategy equilibrium. If a player is mixing optimally then each serve direction used in equilibrium should produce the same expected probability of winning the point. Their results show that professional tennis players' win rates were broadly consistent with the Nash Equilibrium, suggesting that elite athletes may approximate equilibrium play in real competitive settings.

However, (Anderson et al., 2025) reach a different conclusion using a more detailed dynamic model of tennis serves. Their model differs from Walker and Wooders because it includes three serve directions, the score at the time of the serve, and muscle-memory effects from previous serves. They find that players' serve choices do not always equalise win probabilities across directions. They also show that players could improve their chance of winning service games by using dynamic-programming best responses. This was done by using previous serves within the match and current score to estimate "point outcome probability" for each

direction that changes across serves within the same match. This suggests that athletes do randomise, but not always in the optimal way the Mixed Strategy Nash Equilibrium predicts.

Walker and Wooders show that elite athletes can follow behaviour that broadly satisfies the MSNE conditions. Anderson et al., however, show that this conclusion can change when the game is modelled dynamically and looked at on the match level rather than the aggregate level. These two papers show that explanatory conclusions can vary based on the specified game, dataset, heterogeneity across players, and level of aggregation.

(2.2) Penalty kicks

There is extensive research on equilibria of matching pennies games and while this research is highly relevant to analysing penalty kicks, it is not directly transferable without adjustment.

The penalty kick game differs from the standard matching pennies framework because players have three available pure strategies (Left, Centre, and Right) rather than two. Additionally, each kicker and goalie matchup may have different scoring probabilities, meaning a lack of observations on specific kicker/goalie pairs may lead to a heterogeneity issue when working with aggregate data.

Therefore, matching pennies provides a useful theoretical foundation for modelling penalty kicks, but the model must be adapted to reflect the left, centre, and right strategy structure used in this research.

(2.3) MSNE as a benchmark for penalty kick research

The penalty kick game has no pure strategy equilibrium, meaning the Mixed Strategy Nash Equilibrium model is treated as the baseline theoretical benchmark for the three action penalty kick game. A central implication of the MSNE model is that within equilibrium, a player should be indifferent across any action played with positive probability. Within the context of the penalty kick game, expected utilities are calculated based on scoring probability for each direction combination, this means equal scoring probability across directions for a given player.

Some research conducted on penalty kicks compute and fit test the MSNE directly, whilst others test for equal scoring probability and present that as evidence of equilibrium play.

(2.3.1) Bar-Eli et al. (2007)

(Bar-Eli et al., 2007) collected data from observing 286 penalty kicks from various leagues and championships worldwide, the original study does not report the individual leagues or competitions represented. The researchers recoded the kick direction, jump direction, and

whether a goal was scored or not. The kick and jump directions were categorised as either left, centre, or right from the goalkeeper's perspective.

(Bar-Eli et al., 2007) find that goalkeepers only stay in the centre 6 percent of the time, despite the fact that staying in the centre has the highest marginal save probability among the three available goalkeeper directions within their dataset. The researchers posit the reason for this behaviour is "Action Bias".

Action bias refers to a systematic preference for taking action even when inaction may produce an equal or better outcome. It is rooted in norm theory, which suggests that negative outcomes are experienced more intensely when they result from deviating from the perceived norm as opposed to following it (Kahneman & Miller, 1986). Thus, whether a bias favours action or inaction depends on what behaviour is treated as normal in a given context. When action is the norm, individuals may prefer to act because failing after inaction carries additional psychological or social disutility.

In the context of penalty kicks, diving to either side is widely perceived as the normative behaviour for goalkeepers, while remaining in the centre is considered to be deviating from the norm (Bar-Eli et al., 2007). As a result, a goalkeeper who stays in the centre and concedes a goal may experience greater disutility than one who dives and still fails to save the shot. The rationale for this extra disutility for the goalie is that staying in the centre and conceding a goal may look lazy, slow, or indecisive, whilst if they dived and still conceded the goal they would at least look like they had been trying.

Bar-Eli et al. then surveyed 32 professional goalkeepers and asked how bad they felt when they let in a penalty kick dependent on their jump direction. Goalkeepers mostly responded by saying that they felt worse when staying in the centre and conceding a goal compared to diving and conceding a goal. The observed underutilisation of centre and the survey responses led the researchers to conclude that action bias might be present within goalkeepers.

(2.3.2) Azar and Bar-Eli (2011)

The dataset collected by (Bar-Eli et al., 2007) was used in a second study conducted by (Azar & Bar-Eli, 2011) where a Mixed Strategy Nash Equilibrium was calculated under the assumption that all observations come from one common game. Their model generates predicted frequencies of each of the 9 possible kick/dive combinations.

The fit of this MSNE prediction to the observed data was then calculated by taking the absolute differences between the model predicted counts and observed counts for each of the 9 direction combinations. These directional differences were then summed together to get a final "absolute difference" value (AD). A second distance criterion using the sum of squared

differences (SD) was also used to more harshly punish larger deviations from the observed data. Lower AD and SD values indicate better estimated fit with the observed cell counts.

They compared the MSNE predictions against other strategic conjectures including probability matching. Probability matching means that players choose each action with a probability that is equal to the probability that the chosen action is the best response. For example, if the goalie knows the kicker kicks to the left 35% of the time, under probability matching the goalie will choose to dive to the left 35% of the time. The strategic conjectures, and their associated AD and SD values, are given in Table 1.

Method	**AD value**	**SD value**
Sequential-play methods		
Kicker plays first	326.0	15 614.0
Goalkeeper plays first	390.0	26 332.0
Simultaneous-play methods		
MSNE	75.2	817.0
Both play each strategy with equal probability	156.2	3585.6
Kickers use probability matching (based on the sample marginal distribution of jumps)	164.0	4361.7
Goalkeepers use probability matching (based on the sample marginal distribution of kicks)	130.2	2597.4
Both goalkeepers and kickers use probability matching	185.0	5856.6
Goalkeepers use probability matching (based on the perceived marginal distribution of kicks gathered from the survey conducted in the 2009 study)	116.2	1961.0

Table 1 – Tested strategic conjectures by (Azar & Bar-Eli, 2011)

All simultaneous-play conjectures fit better than the sequential-play conjectures, consistent with the assumption that penalty kicks are a simultaneous move game. The MSNE predictions had the lowest AD and SD values amongst all conjectures measured. This result is in favour of their aggregate MSNE model against the other conjectures tested, not evidence that it is a uniquely correct behavioural model.

The findings of this study give evidence of MSNE play, yet the MSNE predictions still overpredict goalie centre play. These results are consistent with the action bias hypothesis established by (Bar-Eli et al., 2007).

(2.3.3) Replication study counterclaim: Singh and Mukherjee (2026)

Recently, a paper replicating the Azar et al (2007) conducted by (Singh & Mukherjee, 2026) tested for the existence of action bias amongst goalkeepers. The researchers collected a new dataset of 292 on-target penalty kicks in the Men's Football World Cup from 1982 to 2022. They found goalies did stay in the centre significantly less often than diving. However, the scoring probabilities from this new dataset do not indicate that staying in the centre is statistically superior to diving for goalies. Quoting directly from the conclusion of this paper on page 5: "goalkeepers had statistically similar chances of saving shots while staying in the centre in comparison to jumping to either side. Thus, in our case, goalkeepers do jump to the sides more during penalty kicks, but not more than they should."

These two studies present different conclusions regarding the low goalkeeper centre rate. (Bar-Eli et al., 2007) posit the low goalkeeper centre rate is suboptimal as their dataset shows staying in the centre to have the highest marginal stop rate for all available goalie directions. (Singh & Mukherjee, 2026) also observe a low goalie centre rate but they do not find that staying in the centre has a significantly higher stop rate compared to when the goalie dives. Therefore, their result challenges that claim that goalkeepers dive more than they should.

Action Bias remains a plausible behavioural mechanism to explain the low centre rate seen in the Azar et al. dataset, yet Singh and Mukherjee find no statistically significant difference in the stop probability for the goalie when they dive compared to staying in the centre.

(2.3.4) Heterogeneity and aggregation issue, Chiappori et al. (2002)

The papers discussed above test for MSNE play on the aggregate level. With (Azar & Bar-Eli, 2011) computing the MSNE off all collected observations, and (Singh & Mukherjee, 2026) use aggregate data to test for equal scoring probabilities and presenting that as evidence of MSNE play.

However, (Chiappori et al., 2002) show that indifference conditions apply across all directions within a player's strategy support for each individual matchup, but do not necessarily apply to aggregated observations. Different penalty kicks may involve different payoff matrices because kickers and goalkeepers vary in ability, dominant foot, experience and familiarity with their opponent, while match circumstances may also affect scoring probabilities.

(Chiappori et al., 2002) establish this observation using 459 penalty kicks from the French and Italian first leagues. Chiappori et al. assume a different matrix for each kicker/goalie matchup, meaning each penalty kick is treated as its own game with unique payoffs. If equilibrium mixes

vary across heterogeneous players, pooling observations produces a mixture of different equilibria rather than the equilibrium of one common game.

The matrices used in the paper allow for heterogeneity across kickers while imposing and testing an Identical Goalkeepers assumption. Under this assumption, the payoff matrix can vary across kickers but does not vary across the goalkeepers facing a given kicker.

The paper also presents evidence that penalty kicks are a simultaneous-move game. This assumption was tested by examining whether, conditional on each player's aggregated past choices, the opponent's direction on the current kick predicts the player's current direction. The current direction was not a significant predictor, which is consistent with players not observing their opponent's current choice before acting.

Chiappori et al. tested and failed to reject the assumption that goalkeepers are homogeneous by testing whether the goalkeeper fixed effects were equal.

Under this maintained assumption, they tested whether, for a given kicker, scoring probabilities were equal across kicks directed left, middle, and right, and whether goalkeepers facing that kicker obtained the same expected payoff from jumping left or right. The joint equality restriction was not rejected at the 5% level for kickers or at the 10% level for goalkeepers.

In separate tests for the 27 kickers with at least five observations, equality was not rejected at the 10% level for 22 kickers and was rejected for five, compared with 2.7 rejections expected under the null.

These results are broadly consistent with equilibrium play but should be interpreted cautiously. The limited number of observations for many kicker/goalkeeper direction combinations may have left the test with low power to detect meaningful payoff differences. Additionally, equal scoring probabilities are only a necessary implication of MSNE and are not sufficient by themselves to establish MSNE play, and the 5 individual rejections indicate that some kickers may depart from the equilibrium restriction.

The aggregation issue presented by (Chiappori et al., 2002) emphasises that individual players may show signs of optimal play, yet when aggregating multiple observations, the pooled observations need not satisfy the restrictions of a common MSNE game. Specifically, individual players are shown to equalise scoring probability across directions, but due to heterogeneity across kickers, equality of scoring probability may not hold at the aggregate level. Kicker footedness is the primary cause of this heterogeneity. Other match specific effects are shown to be less impactful on equilibrium predictions. This will be discussed in more detail and analysed section (2.2) of this research.

(2.4) Synthesis on existing penalty kick literature

Penalty kicks can be treated as a simultaneous move game. Matchup specific evidence is broadly consistent with MSNE indifference conditions. Observations of individual penalty kick matchups are limited. However, literature on tennis serves have more plentiful individual observations, and the results from tennis serves support equal scoring probabilities more strongly.

This leaves gaps in current penalty kick research, with mixed evidence on whether observed penalty kick data is consistent with MSNE play based on scoring probability.

The research conducted by (Azar & Bar-Eli, 2011) found MSNE predictions fit the aggregate data they collected better than probability matching, but do not test any other equilibrium concepts other than MSNE. (Singh & Mukherjee, 2026) test for equal scoring probability across player directions, but do not compute the MSNE or test that their data indicates full support play.

(Bar-Eli et al., 2007) find evidence of action bias being present within goalkeepers, and hypothesise it explains deviations from equilibrium play, yet the researchers do not formally model action bias. Singh et al. conversely present evidence against action bias, but these researchers also do not model action bias.

(Chiappori et al., 2002) present heterogeneity issues that may impact aggregate analysis. Primarily, Heterogeneity across kickers. This means treating aggregate observations as coming from one common game may not be sound.

Even if Action Bias is accepted as a plausible explanation for goalie behaviour, it is not clear that including it into the goalie's payoffs will move the MSNE probabilities in the expected direction, this comparative statics issue will be discussed in the next section.

(2.5) Restrictions of MSNE as a descriptive model

A natural intuition is that the model's prediction could be corrected by including action bias and reducing the utility assigned to the goalkeeper for staying centre. However, within a Nash equilibrium framework, changing a player's outcome payoffs does not necessarily shift that player's equilibrium strategy in a simple or intuitive direction.

This concern is supported by (Crawford & Smallwood, 1984) analysis of comparative statics of equilibria which are strictly mixed. They show that payoff changes in Nash equilibrium can generate counterintuitive strategy adjustments. This occurs because equilibrium mixing is governed by indifference conditions across both players' strategies, rather than by the direct

attractiveness of a single action (Crawford & Smallwood, 1984). Altering the goalkeeper's payoff changes the kicker's equilibrium mix, not the goalkeeper's own mix. This means we must explore other equilibrium models is we want to use action bias as an explanatory parameter within penalty kick modelling.

(2.6) Logit QRE as an alternative behavioural model

The Quantal Response Equilibrium solution concept provides an alternative to Nash Equilibrium by allowing players to make imperfect strategic choices (McKelvey & Palfrey, Quantal response equilibria for normal form games, 1995). Rather than assuming that players always choose the best response with certainty, Logit QRE assumes that players choose higher utility strategies with higher probability, but suboptimal strategies may also be selected. This means that players are still responsive to incentives, but their behaviour is probabilistic rather than perfectly rational. Expected payoffs depend on the opponent's strategy probabilities, and each player's probabilities must therefore be mutually consistent with those of the opponent. Like MSNE, logit QRE is a fixed-point solution, but it replaces exact best responses with probabilistic responses.

A rationality parameter controls how strongly players respond to payoff differences. As this parameter increases, QRE moves closer to the Nash Equilibrium prediction, while lower values allow for more noise or error in decision making.

QRE has been shown to describe real world play better than MSNE in games with strictly mixed equilibria. Observed play can be fit better under QRE compared to MSNE due to the preferable comparative statics effects QRE has over MSNE when payoffs are adjusted.

(2.6.1) Evidence comparing MSNE and QRE in related games

The most relevant existing literature comes from games with unique, strictly mixed equilibria. These can be distinguished into two main categories: QRE applied to a fixed payoff specification, and comparative statics experiments that alter payoffs.

(Ochs, 1995) studies three non-constant sum matching pennies variants and finds that QRE generally fits the data better than MSNE, although the improvement is not uniform across games. (Selten & Chmura, 2008) test 12 games, 6 constant sum, and 6 non-constant sum. In their experiments the Quantal response equilibrium performs significantly better than the Nash on the aggregate level for all 12 games. (Friedman & Zhao, 2021) test two asymmetric matching pennies games. They vary the matching protocol, timing of decisions, and whether subjects choose mixtures directly or realize them through pure actions. Across these

treatments, logit QRE generally provides only a slight improvement, or no improvement, over Nash equilibrium and uniform randomisation.

(McKelvey et al., 2000) vary payoff magnitudes for 2x2 matching pennies games in a way that keeps the MSNE mixes constant. Homogeneous logit QRE predicts that larger payoff differences will systematically change behaviour. The observed choices showed little evidence of these predicted payoff-magnitude effects, and the homogeneous model generally overpredicted the frequency with which row players chose U. The researchers therefore allowed individuals to differ in their rationality parameters, capturing variation in responsiveness to payoff differences. This heterogeneous QRE specification fit the data significantly better than either homogeneous QRE or MSNE.

When testing matching pennies games (Goeree & Holt, 2001) first run a symmetric payoff game and find players mix close to evenly across both available options. They then increase the row player's 'left' payoffs relative to symmetry.

Symmetric game	Column: Left	Column: Right
Row: Top	80,40	40,80
Row: Bottom	40,80	80,40

MSNE for this game: Row $(0.50,0.50)$, Column $(0.50,0.50)$

Observed frequencies: Row $(0.48,0.52)$, Column $(0.48,0.52)$

Asymmetric game	Column: Left	Column: Right
Row: Top	320,40	40,80
Row: Bottom	40,80	80,40

MSNE for this game: Row $(0.50,0.50)$, Column $(0.125,0.875)$

Observed frequencies: Row $(0.96,0.04)$, Column $(0.16,0.84)$

The MSNE would predict that the row player would not adjust their mix and instead the column player would adjust their mix to keep the row player indifferent. Yet when this game was played, the row players chose left 96 percent of the time. Though interestingly, the column players did still mix closely to the new MSNE predictions.

This experiment does not compute any QRE predictions, yet it is useful for this research as it shows real world players tend to play strategies that seem intuitive rather than calculating and playing the Nash optimal strategy. This experiment highlights that the unintuitive comparative statics of the MSNE model are not replicated in real world play. This shows why a model that

allows players to respond to their own payoff changes may be beneficial when modelling real world behaviour.

The QRE framework allows for payoff and preference alterations which can intuitively move the predicted equilibrium in an intended direction. This addresses the own payoff limitation of the MSNE model discussed above.

(Goeree et al., 2003) discuss risk aversion within their matching pennies example. They argue that players in the lab experiments are shown to not be risk neutral. When running a 'safe vs risky' game:

	Column: Left (safe)	Column: Right (risky)
Row: Up (safe)	200,160	160,10
Row: Down (risky)	370,200	10,370

The participants played the safe strategy more often than the risky strategy despite the QRE predictions estimating otherwise. When solved using the classic Logit QRE model the equilibrium always predicted the column player playing right with probability ≥ 0.5 for all rationality levels. When this experiment was run in the laboratory the aggregate choice frequency of right was 0.35. This implied that laboratory participants were showing signs of risk aversion.

This led them to develop the risk aversion model where expected utilities were not just based on the monetary outcomes, instead the utility functions were of the CRRA class. The new risk-averse QRE model moved predictions much closer to the aggregate observed frequencies. The main conclusion of the (Goeree et al., 2003) paper is that "Neither risk aversion nor quantal response equilibrium alone can account for the data across the range of games we study" and that it is only when both are implemented into the explanatory model that we see that equilibrium predictions align with the data.

Other types of behavioural adjustments have also shown to move the QRE predictions in line with observed data.

A separate matching pennies experiment conducted by (Leng et al., 2023) replicate the (Goeree & Holt, 2001) asymmetric experiment and find participants in their experiment play probabilities closer to the MSNE compared to the original Goeree et al. study. When using the logit QRE along with a spite/altruism parameter the model was able to improve fit to both collected datasets. The researcher found that the Goeree et al. frequencies showed behavioural deviations towards altruism, whilst their collected dataset corresponded with a positive spite parameter. This study then shows how the QRE can be used as an explanatory

model that can be augmented using behavioural parameters to estimate behavioural differences between datasets for the same game.

Overall, the logit QRE model is shown to not universally outperform the MSNE. However, there is established evidence showing that the QRE model can produce more intuitive comparative statics when payoffs change compared to the MSNE model. Additionally, the QRE often has strongest empirical performance when paired with heterogeneity or a richer utility specification such as behavioural parameters or CRRA utility specification.

(2.7) Research gap and contribution

The existing literature has not arrived at a consensus for modelling penalty kicks. (Chiappori et al., 2002) present evidence that players play in accordance with the MSNE. Though, this is only on the individual match up level and these predictions do not necessarily hold after aggregation. (Azar & Bar-Eli, 2011) found MSNE predictions fit the aggregate data they collected better than probability matching, but do not test any other equilibrium concepts other than MSNE. (Singh & Mukherjee, 2026) and (Bar-Eli et al., 2007), provide competing conclusions with regards to the presence of goalkeeper action bias. The former provides no evidence of this bias, while the latter presents evidence for the bias. Yet neither of these papers estimate an equilibrium model using an explicit action bias parameter, and Singh and Mukherjee assume full support MSNE in their research, although their estimated payoffs imply that the MSNE for their dataset may instead have reduced support.

Related matching pennies experiments do not establish that homogeneous logit QRE consistently outperforms MSNE. Instead, their results show that explanatory performance depends on the game and model specification: QRE can capture responses to payoff changes that MSNE excludes, while allowing heterogeneity in payoff responsiveness can substantially improve QRE's fit.

This thesis addresses these gaps by incorporating an explicit goalkeeper action bias parameter into both MSNE and logit QRE models. Estimating the models within a common dataset makes it possible to distinguish the contribution of altered goalkeeper utility from that of stochastic payoff responsiveness. Homogeneous and heterogeneous specifications are also compared to determine whether the aggregate data identify meaningful variation in action bias across goalkeepers.

The primary analysis uses the dataset of Bar-Eli et al. (2007). The models are then applied to the independent dataset of Singh and Mukherjee (2026) to assess whether the findings are

robust and generalisable beyond the original sample. This comparison is particularly informative because Singh and Mukherjee reach a different conclusion about the presence of goalkeeper action bias.

(3) Data setup and previous methodologies

The data collected by (Bar-Eli et al., 2007) and also used in (Azar & Bar-Eli, 2011) contains 286 observations from various leagues and championships. All kicks that struck the posts or crossbar or missed the goal, as well as kicks whose direction could not be classified reliably are not included in the sample. Thus, all 286 reported kicks are either scored or saved (stopped).

Table 2 shows the resulting data. Left, centre, and right are defined from the goalkeeper's perspective. A kick recorded as left therefore travels towards the goalkeeper's left, which is the kicker's right. Each cell reports the number of observations, with its percentage of the full 286 kick sample in parentheses.

Kicker vs goalie direction frequency table:

	Goalkeeper			
Kicker	Left	Centre	Right	Total
Left	54 (18.9%)	1 (0.3%)	37 (12.9%)	92 (32.2%)
Centre	41 (14.3%)	10 (3.5%)	31 (10.8%)	82 (28.7%)
Right	46 (16.1%)	7 (2.4%)	59 (20.6%)	112 (39.2%)
Total	141 (49.3%)	18 (6.3%)	127 (44.4%)	286 (100%)

Table 2 – (Azar & Bar-Eli, 2011) Joint distribution of jumps and kicks

Table 3. Goals scored and conditional scoring rates, by kicker and goalkeeper direction

Number of scored kicks based on player direction:

	Goalkeeper			
Kicker	Left	Centre	Right	Total
Left	38 (70.4%)	1 (100.0%)	37 (100.0%)	76 (82.6%)
Centre	37 (90.2%)	4 (40.0%)	30 (96.8%)	71 (86.6%)
Right	46 (100.0%)	7 (100.0%)	44 (74.6%)	97 (86.6%)
Total	**121 (85.8%)**	**12 (66.7%)**	**111 (87.4%)**	**244 (85.3%)**

Table 3 - (Azar & Bar-Eli, 2011) Scoring rates

Here the directions Left, Centre and Right are from the goalie's perspective, so Right, Right means both players go in the same direction.

Running a chi-squared test for equal scoring probabilities gives,

Goalkeeper [$\chi^2(2) = 5.466, p = 0.0650$]

Kicker [$\chi^2(2) = 0.793, p = 0.673$]

At the 5% significance level, we fail to reject the null hypothesis of equal scoring probabilities across the three actions for both the goalkeeper and the kicker. Thus, the tests do not provide sufficient evidence that scoring probability differs by action.

However, when testing for equal scoring probabilities for the goalie when diving vs staying centre we get:

Goalkeeper action	Not scored	Scored	Total	Scoring rate
Jumped (L or R)	36	232	268	86.6%
Stayed Centre	6	12	18	66.7%

Table 4 – Goalie diving vs staying centre scoring rates

$$\chi^2(1) = 5.332, \qquad p = 0.0209$$

In the (Bar-Eli et al., 2007) data, equality of scoring probabilities between jumping and staying centre is rejected at the 5% level. The scoring rate is 86.6% when the goalkeeper jumps and 66.7% when the goalkeeper remains in the centre. If goalkeepers only care about preventing goals, this result suggests that jumping is sub-optimal. Their strong preference for jumping in the observed data may therefore represent a departure from MSNE or indicate that goalkeeper payoffs include considerations beyond the match outcome, such as the psychological cost of appearing inactive.

(Singh & Mukherjee, 2026) conduct the same comparison using their World Cup dataset, though when using their dataset, they obtain $\chi^2(1) = 1.16, p = 0.281$. They therefore fail to reject equality between the two actions, but the researchers do not check if their observed scoring rates imply full support equilibrium. This is a limitation that will be addressed in section (8) of this paper.

These tests are not sufficient on their own to establish whether action bias is present within goalkeepers. Currently the literature presents conflicting suggestive results. This leaves open the opportunity to formally model action bias and test its explanatory power within the penalty kick game.

The observed Azar et al. strategy mixes are better visualised on a simplex plot.

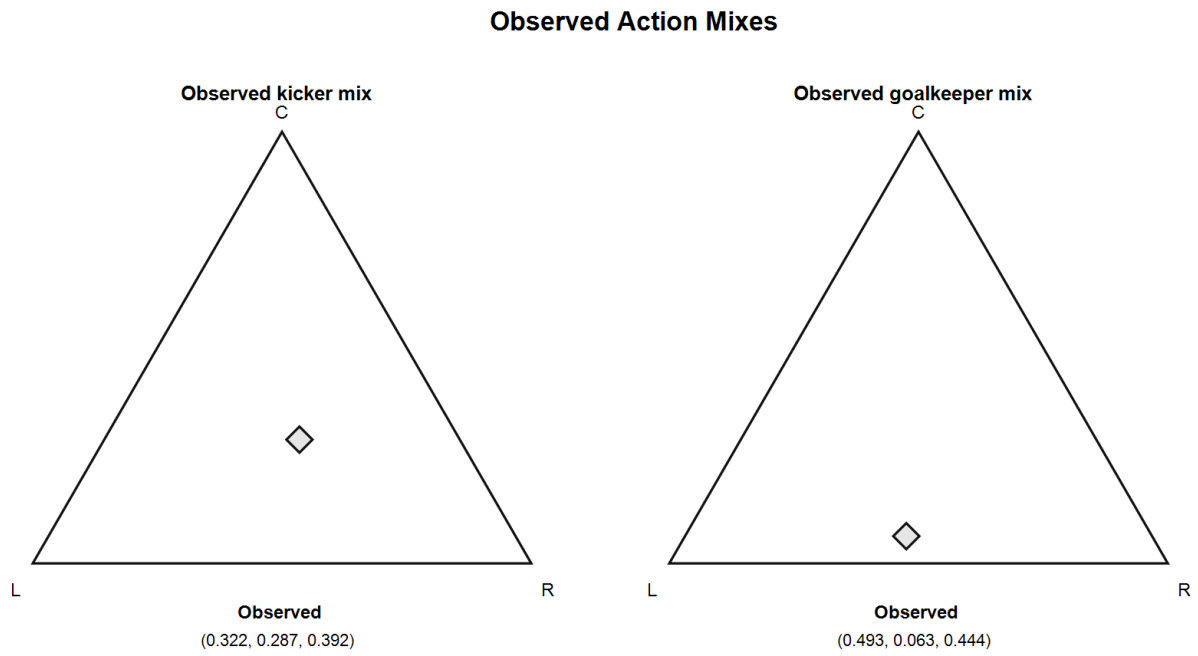


Figure 1 - (Azar & Bar-Eli, 2011) Observed data

Figure 1 shows the probability mix of each player. Each point on the plot represents a mix of probabilities that sum to 1; each corner represents putting 100% probability on one direction (left, right, or centre). The edges represent a mixture across two pure strategies, and any point inside the triangle represents a strictly mixed strategy across all three actions.

Therefore, the plots can be used to easily visualise the proportions of Left, Centre, and Right each player takes, essentially plotting the marginal frequencies of each strategy.

(3.1) Non-independence within the data

The plot represents only the marginal distributions of kicker and goalkeeper strategies. It does not display the joint distribution over the nine direction combinations or the conditional distribution of one player's action given the other's. Under independence, each joint probability would equal the product of the corresponding marginal probabilities. The observed joint frequencies in Table 2 do not coincide exactly with these products.

For example (Left, Left) has an observed percentage frequency of 18.9%, yet if we were to multiply the probability of both players going left we would get $.493 * .322 = .158 \; or \; 15.8\%$.

Assuming complete independence between player actions the observed frequency table would be:

	Goalkeeper			
Kicker	Left	Centre	Right	Total
Left	45.357 (15.9%)	5.790 (2.0%)	40.853 (14.3%)	92 (32.2%)
Centre	40.427 (14.1%)	5.161 (1.8%)	36.413 (12.7%)	82 (28.7%)
Right	55.217 (19.3%)	7.049 (2.5%)	49.734 (17.4%)	112 (39.2%)
Total	141 (49.3%)	18 (6.3%)	127 (44.4%)	286 (100%)

Table 5 – Assumed independence joint action table

The percentage to which each cell is over or underrepresented in the observed data when compared to independence is shown below

	goalkeeper		
Kicker	Left	Centre	Right
Left	+19.1%	-82.7%	-9.4%
Centre	+1.4%	+93.8%	-14.9%
Right	-16.7%	-0.7%	+18.6%

Table 6 – Percentage over/underrepresentation relative to independence

The non-independence within the data may indicate heterogeneity. Reiterating the (Chiappori et al., 2002) findings that aggregation across heterogeneous players may generate dependence even when actions are independent. Additionally, the goalkeeper may obtain information about the kicker's intended action before committing, creating sequential or partially sequential play. This explanation is presented by (Bar-Eli et al., 2007) and will be discussed in more detail in the following section.

Observing the simplex plots allows us to visualise the direction probabilities of each player but does not intuitively show how well the joint distribution is replicated. Therefore, model performance will be evaluated against the entire joint distribution using Maximum Likelihood Estimations to determine model parameters.

(3.2) Sequential play theory

From looking at Table 6 we see that the cells which represent both players choosing the same direction are heavily overrepresented. Giving suggestive evidence that alongside heterogeneity, the data may also be skewed by some proportion of the goalkeeper accurately predicting the kicker's direction and matching.

Recall, (Azar & Bar-Eli, 2011) fit test sequential play and simultaneous play methods and find all simultaneous play methods fit the observed data better than the sequential play methods (refer to Table 1). However, these conjectures assume 100% sequential play or 100%

simultaneous play. Real world play observations may be made up of some proportion of kicks which are simultaneous and others which are sequential.

(Bar-Eli et al., 2007) tested for this by assuming that the goalie is observing and matching the kickers direction M% of the time, and that all other kicks are independent/non-sequential. The independent kicks are assumed to have the same directional rates as the marginal strategies and are not estimated under any equilibrium. Matching is assumed to be equally likely for each direction.

Under hypothesis testing for different values of M, the highest p-value is obtained when M=9.1%. This implies that within the data a small percentage of observations are situations where the goalie is accurately predicting the kicker's direction and reacting accordingly when choosing his own action. All other observations are situations where the kick and jump are independent.

We see similar non-independence trends in the other two datasets.

The dataset collected by (Chiappori et al., 2002) has over/underrepresentation relative to independence of

	goalkeeper		
Kicker	Left	Centre	Right
Left	+0.3%	-19.0%	+0.7%
Centre	+7.3%	+58.5%	-13.5%
Right	-3.6%	-4.1%	+5.2%

Table 7 - (Chiappori et al., 2002) over/underrepresentation

Whilst the dataset collected by (Singh & Mukherjee, 2026) has over/underrepresentation of

	goalkeeper		
Kicker	Left	Centre	Right
Left	+24.6%	+34.1%	-27.8%
Centre	-0.2%	+6.8%	-1.4%
Right	-18.0%	-27.9%	+21.0%

Table 8 - (Singh & Mukherjee, 2026) Over/underrepresentation

We see that for all three datasets the same-direction cells are overrepresented. However, these are not the only overrepresented cells. For example, in Table 8 the cell (L,C) is heavily overrepresented relative to what independence would assume. The goalkeeper matching the kicker is one form of possible heterogeneity. Other forms of possible unobserved heterogeneity must also be considered.

(3.3) Kicker footedness

Skill heterogeneity implies that scoring probabilities for the same combination of kick and dive directions may differ across player matchups. It may also cause players with different abilities to choose different strategy mixtures. Aggregating across these matchups can therefore distort estimated scoring probabilities and produce correlation in the joint distribution, even when choices are independent within each matchup. For example, if more skilled kickers shoot towards one side more frequently, kicks to that side will disproportionately come from stronger players and will have a higher aggregate scoring rate.

(Chiappori et al., 2002) find evidence of kicker heterogeneity but reject goalkeeper heterogeneity under their 'Identical Goalkeepers' assumption. Most of the observed kicker heterogeneity is associated with footedness, with right and left footed kickers having opposite natural sides because kicking across the body is easier. Chiappori et al. therefore reverse the recorded directions for left footed kickers so that kicking left (from the kicker's perspective) denotes the natural side for every kicker. After this adjustment, the evidence for kicker heterogeneity becomes considerably weaker. The dominant form of heterogeneity can be interpreted as aggregation over two mirrored games: one for left footed kickers, and one for right footed kickers.

This presents a limitation for Azar and Bar-Eli's aggregate analysis, where the MSNE was calculated under the assumption that all observations came from one common game. If the collected data combines two footedness games with different payoffs or equilibria then, observed deviations from the aggregate MSNE may reflect unmodeled heterogeneity instead of genuine departure from equilibrium play.

This implies that if the Azar et al. data can be manipulated to include player footedness, the heterogeneity/selection bias issue may be partially mitigated.

Chiappori et al. established that conditional on the goalie choosing the correct direction, kicks to the kicker's natural side have a higher scoring probability than kicks to the nonnatural side.

We can assume that the aggregate Azar et al. stop probabilities are made up of some proportion of players kicking to their natural and nonnatural side. Note that Chiappori et al. assume that footedness does not affect scoring probability when kicking to the centre.

Therefore, we can separate the aggregate stop probability table into two separate tables, one for left footed kickers, and one for right footed kickers.

Because footedness affects kickers' left vs right scoring probabilities, we replace the recorded probabilities in cells (L,L) and (R,R) in Table 3 with variables. Directions are labelled from the goalkeeper's perspective, so Right denotes the natural side for right footed kickers.

There is no need to adjust (L,C) (L,R) (R,L) and (R,C) as the scoring probabilities of these cells is 1 meaning there is no heterogeneity within these cells.

Left footed kicker	Goalkeeper		
Kicker	Left	Centre	Right
Left	$X + a$	100%	100%
Centre	90.2%	40.0%	96.8%
Right	100%	100%	x

Table 9 – Left footed kicker table

right footed kicker	Goalkeeper		
Kicker	Left	Centre	Right
Left	x	100%	100%
Centre	90.2%	40.0%	96.8%
Right	100%	100%	$X + a$

Table 10 – Right footed kicker table

Here x is the scoring probability of a kicker kicking to their non-natural side and scoring given the goalie goes in the same direction

And a is the advantage in score probability a kicker receives from kicking to their natural side over their non-natural side.

Given that the aggregate score probability for (L,L) is .704, we assume this is made up of some left footed players kicking to their natural side (recall goalie's left is kicker's right) and some share of right footed kickers kicking to their nonnatural side.

This leads to the equation:

$$.704 = q(x + a) + (1 - q)\,(x) = x + aq$$

where q is the proportion of left footed kickers

Similarly, for (R,R)

$$.746 = q(x) + (1 - q)(x + a) = x + a - aq$$

Given there are three variables to be estimated and two equations, there is not enough information to solve these equations.

Chiappori et al. report that, when the goalkeeper chooses the same side as the kick, the scoring probability is 63.6% on the kicker's natural side and 43.7% on the opposite side. This is a difference of 19.9 percentage points, rounded here to an additive difference of $a = 0.2$

The following calculation assumes that this additive natural-side advantage can be transferred from the Chiappori et al. sample to the Azar and Bar-Eli payoff matrix. This is an assumption rather than a parameter identified from the Azar and Bar-Eli aggregate data.

We can now solve these equations, giving

$$q = .395,\ \ x = .625$$

These results imply that 39.5% of kickers in the data are left footed. Chiappori et al. report that left footed kickers make up 15% of the population. Other studies find left footed kickers make up 20%-25% of penalty kicks taken (Hart & Esteva, 2014), (Almeida & Volossovitch, 2023), (Almeida et al., 2016). The q estimation of 39.5% falls far outside this range.

Assuming this estimation is correct, we compute and solve for the aggregate equilibrium joint distribution.

Table 11 presents the aggregate predicted frequencies

	Goalkeeper			
Kicker	Left	Centre	Right	Total
Left	17.6%	2.8%	15.2%	35.6%
Centre	7.8%	1.6%	8.8%	18.2%
Right	17.2%	4.2%	24.9%	46.3%
Total	42.5%	8.6%	48.9%	100.0%

Table 11 – Predicted joint frequency under footedness assumption

The model does not reproduce the observed joint distribution closely. In particular, it predicts that kickers choose the centre 18.2% of the time, compared with 28.7% in the data, and substantially understates the observed (C,C) frequency.

Alternately, the (Chiappori et al., 2002) data indicates that kickers choose their natural side 1.185 times more frequently than their nonnatural side.

This implies that for the Azar et al. data,

$$Kicker\ aggregate\ left\ frequency = .322 = q(\alpha_l) + (1-q)\alpha_r$$

$$Kicker\ aggregate\ right\ frequency = .392 = q(\gamma_l) + (1-q)\gamma_r$$

subject to

$$\alpha_l = 1.185\gamma_l$$

$$1.185\alpha_r = \gamma_r$$

where q is the proportion of left footed kickers, and α_l and γ_l denote the probabilities a left footed kicker kicks left and right, respectively. Here, α_r and γ_r are defined similarly for right footed kickers.

Solving these equations give

$$\alpha_l = 1.185\gamma_l = \frac{-0.030576}{q}$$

meaning, for every $0 < q < 1$. $\alpha_l < 0,\ \gamma_l < 0$.

Therefore, these equations have no valid probability solution.

We then conclude that the aggregate frequency and scoring probability data alone is not enough to estimate the share of left footed kickers in the aggregate data. This means that going forward the footedness heterogeneity cannot be accounted for reliably in the Azar et al. data. Therefore, we must accept this limitation and interpret all results accordingly.

Additionally, there may also be heterogeneity in players' behavioural biases. This paper analyses players' payoffs through the inclusion of the behavioural bias "action bias". Just as players are shown to be heterogeneous in their abilities, it is also possible that they are heterogeneous in the extent they feel action bias. This possible behavioural heterogeneity will be explored throughout the analysis conducted.

(4) MSNE analysis

From Table 3 we see that when the kicker and the goalkeeper go in the same direction (highlighted in the table), the ball is much more likely to be stopped. While this may seem obvious, what is interesting is that when the goalie stays in the centre and the ball is also kicked to the centre, the goalie stops the ball at over double the rate of (L,L) and (R,R).

Under MSNE logic this leads to both players putting less probability on centre, as the goalie would need to stay in the centre less, making centre a more attractive option to the kicker. The kicker would then need to kick to the centre less, making it less attractive for the goalie to stay in the centre. However, in the recorded data goalies seem to stay in the centre even less than expected under MSNE conditions. To understand this, we first look at the original MSNE calculations done by (Azar & Bar-Eli, 2011).

To quantify the expected utility of each directional strategy, Azar et al. first assigned outcome utilities based on whether the goal is scored or not.

	Goal is scored	**Goal is not scored**
Kicker payoff	1	0
Goalkeeper payoff	-1	0

Table 12 – Payoff matrix

These payoffs can be used in conjunction with the scoring probabilities to calculate to expected utility for each player given their chosen action and their opponents chosen action.

For example, when both players choose left (L,L):

$$U_K(L) = (1 - s) * 1 + s * 0 = 0.704 * 1 + 0 = 0.704$$

$$U_G(L) = (1 - s) * -1 + s * 0 = 0.704 * -1 + 0 = -0.704$$

where s is the chance of the ball being stopped given the players' actions

The outcome utilities for each situation can be calculated in the same way, giving the following table:

Given that this is a zero sum game, only the kicker's payoffs are shown.

For simplicity we can shorten Left, Centre, Right to L,C,R respectively

	Goalie		
Kicker	L	C	R
L	0.704	1	1
C	0.902	0.4	0.968
R	1	1	0.746

Table 13 – Expected payoffs

Expected utility functions can then be made using this data.

We denote the probability that the goalkeeper jumps to the left by p, and the probability that he stays in the centre by q; denoting the probability that the kicker kicks towards the left by α and the probability he kicks to the centre by β. Hence, the keeper jumps right with probability $1 - p - q$. For simplicity this can be referred to as r. The kicker kicks right with probability $1 - \alpha - \beta$. This can be referred to as γ

Kicker expected utility functions:

$$U_k(\text{Left}) = 0.704p + q + (1 - p - q)$$

$$U_k(\text{Centre}) = 0.902p + 0.4q + 0.968(1 - p - q)$$

$$U_k(\text{Right}) = p + q + 0.746(1 - p - q)$$

Goalie expected utility functions:

$$U_g(\text{Left}) = -0.704\alpha - 0.902\beta - (1 - \alpha - \beta)$$

$$U_g(\text{Centre}) = -\alpha - 0.4\beta - (1 - \alpha - \beta)$$

$$U_g(\text{Right}) = -\alpha - 0.968\beta - 0.746(1 - \alpha - \beta)$$

Setting these expected utility functions equal to each other identifies any strictly mixed Nash Equilibria:

$$U_k(L) = U_K(C) = U_K(R)$$

$$U_G(L) = U_G(C) = U_G(R)$$

Solving this set of equations gives us the following equilibrium probabilities:

$p = 0.411,\ q = 0.110,\ \alpha = 0.344,\ \beta =$0.203

It is worth noting for this game setup this equilibrium is unique. This is explained later by Proposition 1 and proven in the appendix.

Table 14 below shows the predicted MSNE outcome distribution (i.e., the expected number of observations of each type) for the dataset.

	Goalie			
Kicker	Left	Centre	Right	Total
Left	40.42	10.83	47.10	98.34
Centre	23.83	6.38	27.77	57.99
Right	53.29	14.27	62.10	129.67
Total	117.54	31.48	136.98	286.00

Table 14 – Predicted MSNE distribution

To visualise how the Nash Equilibrium compares to the observed data.

The Nash equilibrium strategies have been plotted on the simplex plot alongside the observed data from earlier.

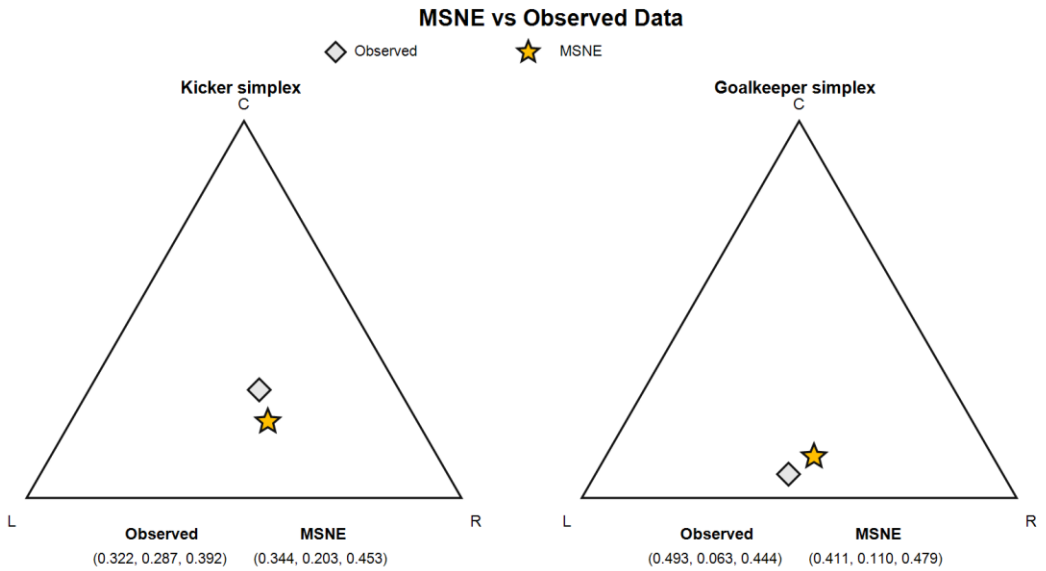


Figure 2 – MSNE vs observed data

It can be seen that the Nash Equilibrium model overpredicts the amount the goalie stands in the centre and does not dive, whilst the observed kicker overplays the centre relative to what

equilibrium play would suggest. This implies that the goalie will underplay C relative to the MSNE prediction.

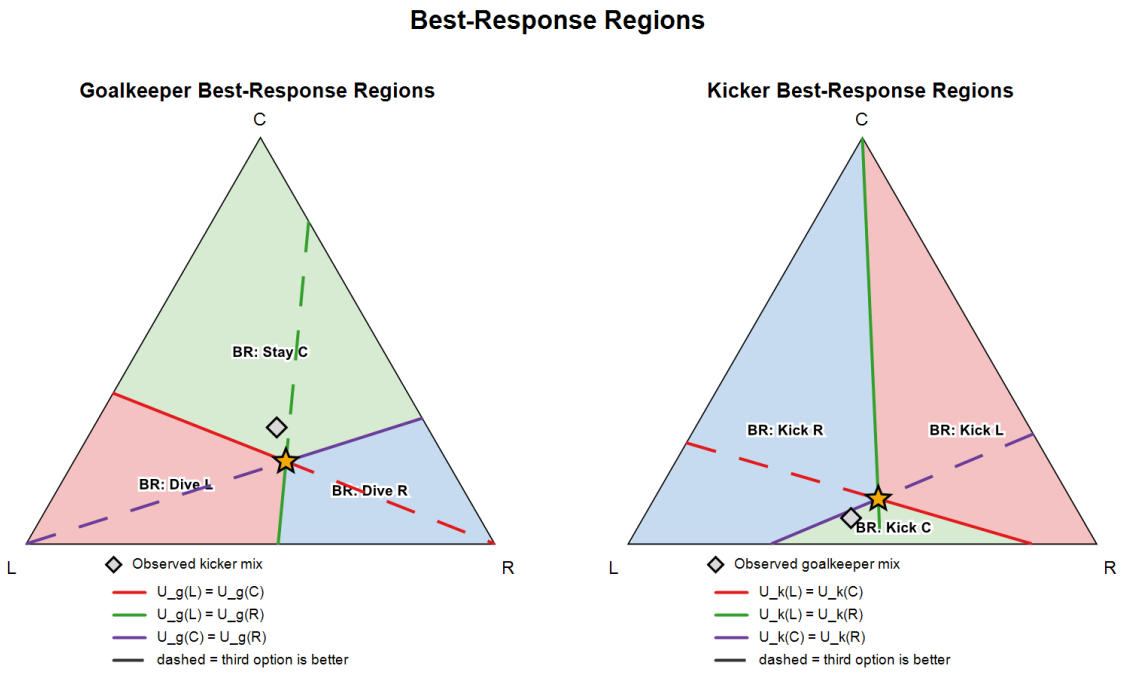


Figure 3 – Player Best response regions

Figure 3 shows what a player's best response would be given their opponent's strategy mix. The left plot shows what the goalkeeper's best response would be given the kicker's strategy mix and vice versa for the right plot. The coloured regions show opponent mixes that lead to the player responding with a pure strategy, and the lines show points where the responding player is indifferent across two actions, leading them to best respond with a two strategy mix. We see that each player's best response is to exclusively play centre when responding to their observed opponent mix.

Figure 3 can be used to assess how far players are from best responding to their opponent's observed behaviour.

From the data we get the scoring probabilities used to make Table 15 below. For the kicker row, each cell represents the scoring probability when the kicker selects that direction, given the goalie plays the observed frequencies. For the goalkeeper row, each cell represents the scoring probability when the goalie plays that direction given the kicker plays their observed frequencies. The mixed strategy column refers to the overall score probability given both players play the observed direction mixes.

Scoring probability	left	centre	right	Mixed strategy
Kicker	85.39%	**89.98%**	88.71%	88.01%
Goalkeeper	87.67%	**82.80%**	89.12%	88.01%

Table 15 – Marginal scoring probabilities given observed frequencies

Note that the overall score probability is different to that of Table 3, this is because these scoring probabilities were calculated from the observed mixes of each player rather than the total 3x3 stopped counts.

From Figure 3 we can see that the kicker's best response to the observed goalkeeper is to kick to the centre. The goalkeeper's best response to the observed kicker is to stay in the centre.

At the observed frequencies, the overall scoring probability is .8801

If the kicker were to best respond to the observed goalkeeper the improvement in their scoring probability would be, $0.8998 - 0.8801 = 0.0197$. This means that the kicker could increase their expected utility by 0.0197 utils by best responding as opposed to playing the observed kicker mix

And the goalkeeper could improve their expected utility by $(-0.8280) - (-0.8801) = 0.0521$ utils if they were to best respond to the observed kicker.

At the MSNE the overall scoring probability is 0.8783. Since the scoring probability at the marginal mixes is 0.8801, the goalie is worse off when playing the observed marginal mix compared to the MSNE and the kicker is better off when playing their observed marginal mix instead of equilibrium play.

The fact that we do not see observed players playing the MSNE predicted frequencies tells us that either the MSNE is not the right model of strategic behaviour and/or the current utility equations being used are failing to accurately capture the actual utilities faced by real world players.

The possibility of players acting non-optimally will be explored using the QRE framework and stochastic best responding. Incorporating action bias tests whether the specified utility functions adequately capture the incentives experienced by real world players.

(4.1) Action Bias implementation

(Bar-Eli et al., 2007) conclude that goalkeepers exhibit action bias because they choose jumping (action) when staying in the centre (inaction) would lead to a higher probability of stopping the ball. They reinforce this claim with the goalkeeper survey results.

In this thesis, action bias is incorporated into the model as an additional cost attached to inaction. This allows the model to capture the possibility that goalkeepers are not only responding to the objective probability of saving the shot, but also to the psychological cost of violating the norm of action.

Modifying the outcome payoffs to differentiate between the goalie diving and staying in the centre allows action bias to be incorporated into the model.

Action bias can be implemented in the model uniformly when the goalie plays centre, or only implemented if the goalie concedes the goal whilst staying in the centre:

	Goal is scored	Goal is not scored
Kicker payoff	1	0
Goalkeeper payoff when diving (L/R)	-1	0
Goalkeeper payoff when staying centre (C)	-1-ε	-ε

Table 16 – Action bias across whole action

VS

	Goal is scored	Goal is not scored
Kicker payoff	1	0
Goalkeeper payoff when diving (L/R)	-1	0
Goalkeeper payoff when staying centre (C)	-1-ε	0

Table 17 – Action bias only when goal is scored

Here, ε is a positive scalar that represents the additional burden felt by the goalkeeper if they stay in the centre. This ε can therefore be considered the quantifiable burden of action bias.

Algebraically the two implementations affect the utility functions in the following way

Whole-action bias:

$$U_g(Centre) = -\alpha - 0.4\beta - \gamma - \varepsilon$$

Scored-only bias:

$$U_g(Centre) = -\alpha - 0.4\beta - \gamma - \varepsilon(1 - 0.6\beta)$$

Assuming the values for $\alpha, \beta,$ and γ are fixed, the only difference across the two utility equations is the way ε enters the functions. This means we can use either implementation to achieve the same expected utility value if we use the following conversion:

$$\varepsilon_{whole} = \varepsilon_{scored}(1 - 0.6\beta)$$

Because this conversion requires the rate of kicking centre (β) to be known, we cannot create a generic conversion that scales the two across all strategy mixes. We can only use the two interchangeably for a specific equilibrium mix.

For any equilibrium value of β, the expected utility of the goalie staying in the centre is the same for either action bias implementation after the conversion is used. Therefore, the payoff differences across actions, and therefore best response conditions are the same under both action bias specifications after conversion. This means that any Nash Equilibrium mix generated by one action bias specification can be replicated by the other. The same conversion holds true for the QRE model.

Therefore, the choice for which one to implement comes down to the intuitive logic of football, rather than working backwards from the model. Action bias exists because deviating from the norm in situations leads to a larger perceived utility loss when incorrect (Kahneman & Miller, 1986). This is because when deviating from the norm and failing, it becomes easier to dwell on the alternative outcome if the player had stuck to the norm. In the case of penalty kicks diving is considered the norm for the goalie, conversely, staying in the centre is deviating from the norm. If the goalie were to stay in the centre and concede a goal, they may be viewed as lazy or indecisive. Norm theory implies that the regret of not sticking with the norm sets in after the decision to deviate is proven to be wrong. When trying something different and failing people will dwell on “What if” and regret not sticking to norm behaviour. Applying this logic to penalty kick games, it follows that action bias should only be included in the model in the case that the goalie stays in the centre and a goal is scored. If the goalie stays in the centre and stops the goal, they were correct in their decision and therefore do not dwell on “what if”, meaning they do not face action bias in that situation. Therefore, action bias will be implemented in this model only when a goal is scored, as per Table 17.

This gives the following utility equations for the action bias game,

Kicker expected utility functions:

$$U_k(\text{Left}) = 0.704p + q + (1 - p - q)$$

$$U_k(\text{Centre}) = 0.902p + 0.4q + 0.968(1 - p - q)$$

$$U_k(\text{Right}) = p + q + 0.746(1 - p - q)$$

Goalie expected utility functions:

$$U_g(\text{Left}) = -0.704\alpha - 0.902\beta - (1 - \alpha - \beta)$$

$$U_g(Centre) = -\alpha - 0.4\beta - \gamma - \varepsilon(1 - 0.6\beta)$$

$$U_g(\text{Right}) = -\alpha - 0.968\beta - 0.746(1 - \alpha - \beta)$$

When best responding, a player must be indifferent between chosen action within their mix. Therefore, if the goalkeeper's best response includes diving and staying in the centre, then they must receive equal expected utility from both.

Consequently,

$$U(diving) = U(centre)$$

We can use this equality assumption to make a behavioural interpretation for the value of ε.

$$-Pr(score|dive) = -Pr(score|\, centre)(1 + \varepsilon)$$

We will refer to $Pr(score|dive)$ as D and $(score|\, centre)$ as C

$$-D = -C - C\varepsilon$$

$$\frac{D - C}{C} = \varepsilon$$

When action bias is present, $D > C$ for the goalie to remain indifferent. Thus, ε represents the proportional decrease in scoring probability the goalie is willing to forgo to indulge their mental bias.

Recall that under the MSNE model with a strictly mixed equilibrium, changing a player's payoffs will not change their equilibrium mix, but rather their opponent adjusts their mix to keep the altered player indifferent. This implies that including action bias into the goalie's payoffs will not change their equilibrium mix and instead the kicker's equilibrium mix will need to adjust to keep the goalie indifferent. However, this assumes that the MSNE will always be strictly mixed, with players putting positive probability on all available pure strategies. For this game, when action bias is large enough, we will see the equilibrium move to reduced support, meaning players no longer include all available pure strategies within their mix. To explore how action bias changes the best responses of players observe the animation below.

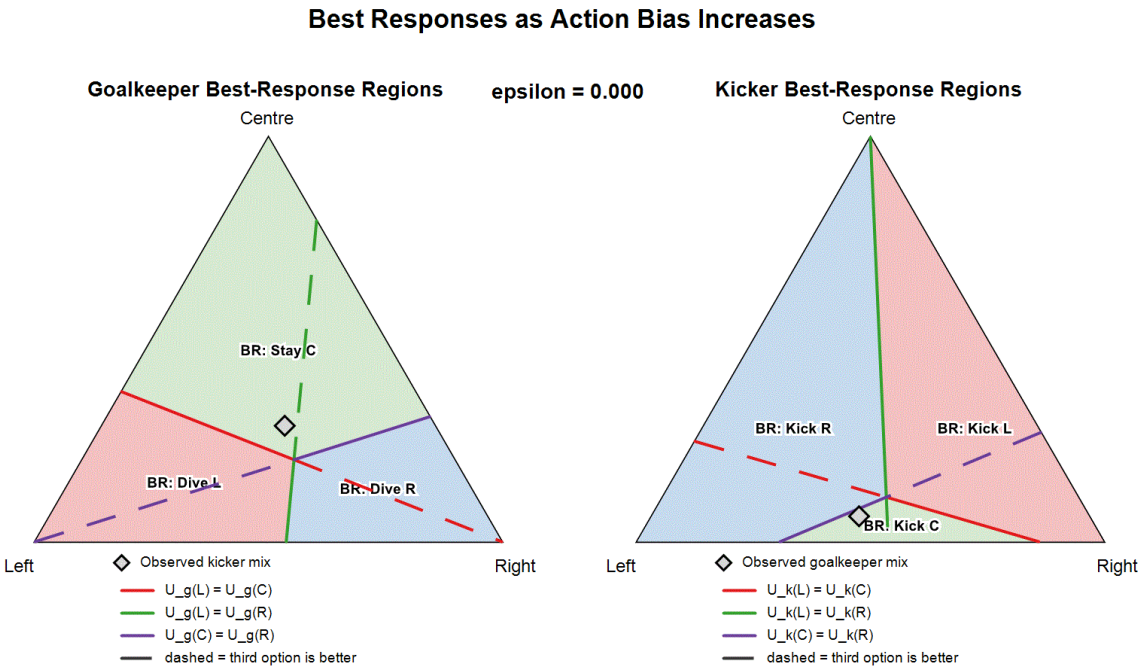


Figure 4 – Animated Best response regions

Looking at the animated plots above, we see that action bias changes the intercepts of the indifference lines but does not change the equilibrium support until action bias reaches a critical value. For low values of epsilon, the relevant best response intersection lies in the interior of both simplexes, implying a full support equilibrium in which both players mix over left, centre, and right. As epsilon rises the kicker chooses C more often so, staying centre becomes less attractive for the goalkeeper, causing the centre best response region to shrink. Once epsilon passes a certain value, the full support equilibrium is no longer sustained, and the equilibrium moves to a reduced support in which the goalkeeper no longer uses centre, and the kicker never chooses to kick left. We will call the epsilon value where the MSNE moves from full support to reduced support the 'critical value' denoted by ε^*.

Proposition 1: Set of equilibria for each ε

> Let $G(\varepsilon)$ denote the action bias game with action bias parameter ε, where the permitted range of ε is $\varepsilon \geq 0$. For this game the critical value of ε is $\varepsilon^* = \frac{31877}{41900} \approx 0.7608$.

The Nash Equilibria of $G(\varepsilon)$ are as follows:

1. If $0 \leq \varepsilon < \varepsilon^*$, then $G(\varepsilon)$ has a unique equilibrium which is strictly mixed (full support).
2. If $\varepsilon = \varepsilon^*$, then $G(\varepsilon)$ has multiple Nash Equilibria. The equilibrium set contains both full support and reduced support mixes. In this case, the kicker mixes over C and R, whilst the goalie has a continuum of equilibrium strategies that maintain the kicker's indifference between C and R.

3. If $\varepsilon > \varepsilon^*$, $G(\varepsilon)$ has a unique reduced support equilibrium that is independent of ε. In this equilibrium each player mixes over two actions. The kicker mixes over CR, and the goalie mixes over LR.

The proof of proposition 1 can be found in the appendix below.

The change in equilibrium as ε varies is shown on Figure 5 below.

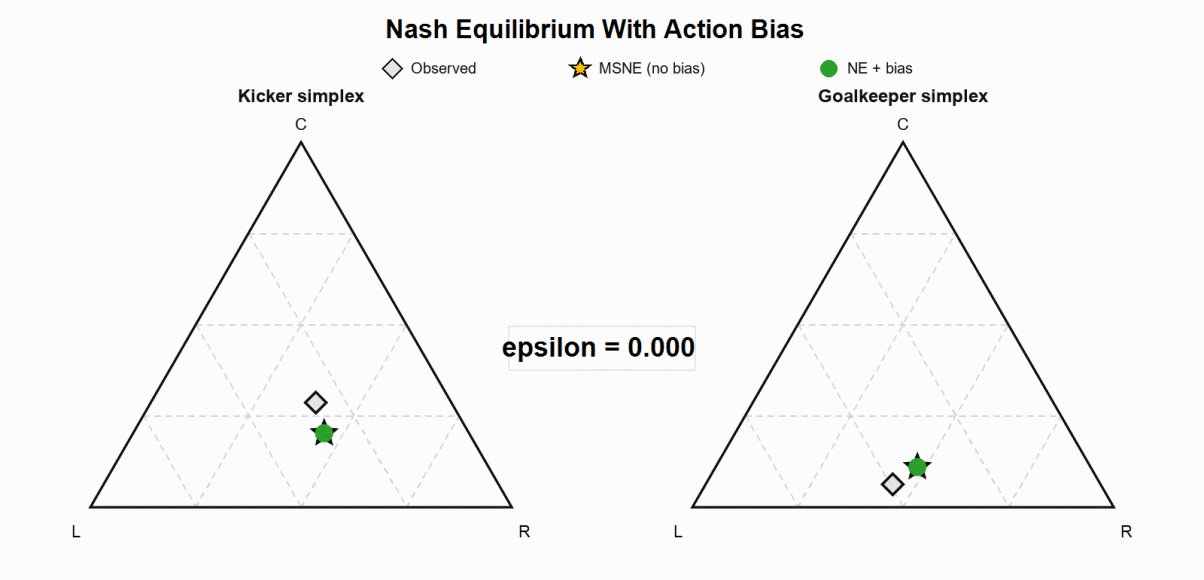


Figure 5 – Animated MSNE predictions given ε

From the plot we see that even though action bias affects the goalie's payoffs, it is the kicker that initially changes their equilibrium strategy. This is due to the indifference condition required by the Nash Equilibrium model. As the payoff for staying in the centre reduces for the goalie, they will be incentivised to stay in the centre less. Therefore, the kicker must kick to the centre more to avoid the goalie having a profitable deviation. It is only when $\varepsilon \geq \varepsilon^*$ that we see the goalie deviate their mix from the no bias Nash Equilibrium.

It is not possible to see the set of goalie strategy mixes at the knife edge case ($\varepsilon = \varepsilon^*$) on the above plot, so a static plot where $\varepsilon = \varepsilon^*$has been generated to show the range of equilibrium mixes the goalie may take at that point.

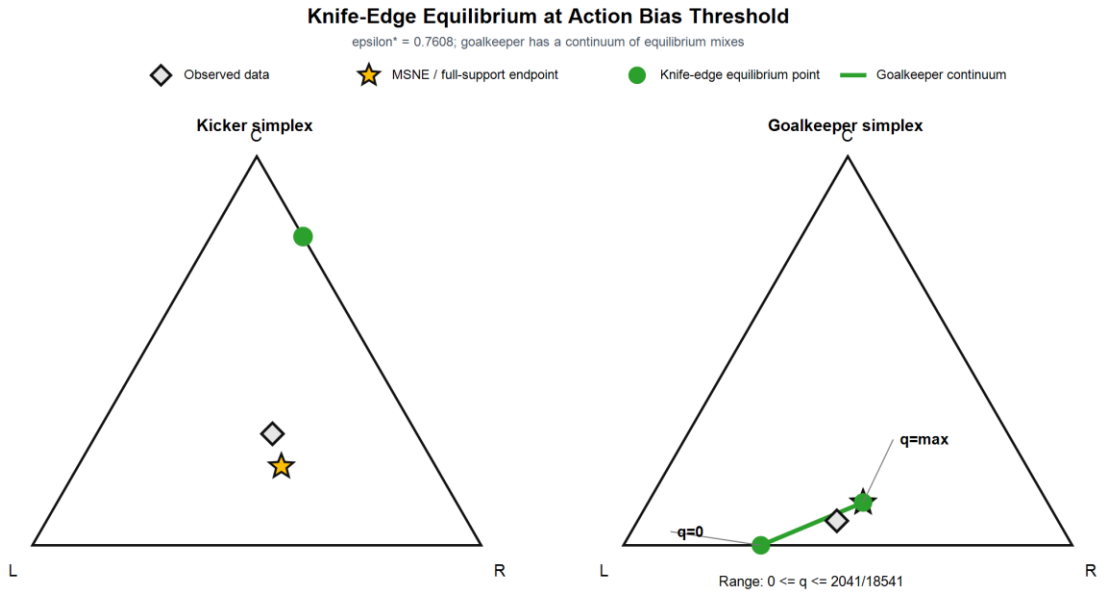


Figure 6 – Equilibrium mixes at ε^*

Once action bias is included in the model, there is still a unique equilibrium for each value of epsilon other than $\varepsilon = \varepsilon^*$.

(4.2) Parameter estimations

For this research maximum likelihood estimations are used to assess the fit of each model. The likelihood function is

$$\mathrm{LL}(\theta) = \sum_{i \in \{L,C,R\}} \sum_{j \in \{L,C,R\}} n_{ij} \log\left(P_{ij}(\theta)\right)$$

where, n_{ij}, is the number of observed penalties in which the kicker chooses direction i and the goalie chooses direction j; $P_{ij}(\theta)$ is the model's predicted probability of this kick/dive combination, given the parameter vector (θ).

This likelihood function will be referred to as the joint likelihood model, as it is matching the data to the model implied joint distribution. Because the likelihood calculations are matching to the frequency of each individual cell rather than the marginal direction frequences for each player, any non-independence in the data will be reflected in the model estimates. This also means that a particular model that can match the observed marginal strategies perfectly, but cannot replicate the dependence observed across individual cells, will be penalised by the likelihood estimations.

However, any reduced support play will result in a log likelihood value of $-\infty$. This is because if a player does not play an action then for certain direction combinations $P_{ij}(\theta)$=0. The likelihood function cannot compute an estimate for this as log(0) is undefined.

Therefore, any equilibria at or above ε^* cannot be likelihood tested, including the critical value case itself.

For models where this is a possibility the (Azar & Bar-Eli, 2011) metrics will be used instead. These are, the Absolute Differences (AD) and Squared Differences (SD). The Maximum Likelihood metric is still the preferred metric as it allows for statistical testing and is standard across existing literature. Therefore, we will only use the AD and SD metrics for models that could possibly return reduced support equilibria.

The AD and SD use the sum of the difference between the observed counts and the predicted counts for each direction combination, meaning the smaller the summed differences, the better the model replicates the observed data. The AD and SD are defined as:

$$AD = \sum_{i\in\{L,C,R\}} \sum_{j\in\{L,C,R\}} \left| n_{ij} - NP_{ij}(\theta) \right|$$

$$SD = \sum_{i\in\{L,C,R\}} \sum_{j\in\{L,C,R\}} (n_{ij} - NP_{ij}(\theta))^2$$

Here, N refers to the total number of observed counts (286), meaning $NP_{ij}(\theta)$ is the number of predicted counts for action ij, given the parameter vector θ.

The AD can range from 0 (best possible fit) to 570, and the SD can range from 0 to 93,898 (Azar & Bar-Eli, 2011).

To assess what value of ε best represents the observed data, two estimations were run.

A minimum absolute difference estimation defined as:

$$\hat{\varepsilon}_{AD} = \arg\min_{\varepsilon\geq 0} \mathrm{AD}(\theta)$$

And a minimum squared difference estimation defined as:

$$\hat{\varepsilon}_{SD} = \arg\min_{\varepsilon\geq 0} SD(\theta)$$

The results are:

Fit criterion	$\hat{\varepsilon}$	Distance measure result
Minimum AD	0.095692	57.1620
Minimum SD	0.051547	607.5921

Table 18 – Fitted MSNE results

The two estimations return similar $\hat{\varepsilon}$ values all less than 0.1 meaning all lie within the full support region[1].

This is because when action bias is included into the MSNE model, the kicker's equilibrium strategy changes, but the goalie's strategy only changes once the critical epsilon value is reached. At that point the kicker's mix is so far from the observed data that any improvement in goalie fit is offset by the reduction in kicker fit, leading to the search models fitting the kicker as well as possible.

This suggests that it is reasonable to restrict the search to ε values below ε^*, as the best fitting ε lies within the full support region. This means we can use the MLE search criterion as all equilibria searched will not result in any log(0) errors.

We run the MLE and compare to the AD and SD search results:

Fit criterion	$\hat{\varepsilon}$	Distance measure result
MLE	0.064059	-570.9426
Minimum AD	0.095692	57.1620
Minimum SD	0.051547	607.5921

Table 19 – Fitted MSNE results including MLE

The MLE $\hat{\varepsilon}$ lies between the AD and SD estimates. Across all three estimations, the goalie is willing to forgo an approximate 5%-10% proportional decrease in scoring probability, as per the behavioural interpretation established earlier.

The predicted mixes at equilibrium implied by the MLE are:

	L	C	R
Kicker	0.2979	0.2817	0.4204
Goalie	0.4110	0.1101	0.4789

Table 20 – Predicted mixes at MLE

A likelihood-ratio test comparing the homogeneous action-bias MSNE with the nested (Azar & Bar-Eli, 2011) MSNE produced $LR = 10.614$. Because the no bias restriction lies on the boundary of the parameter space, statistical significance was assessed using a parametric bootstrap. None of the 499 bootstrap samples generated an LR statistic as large as the observed value, giving $p = 0.002$. The no bias restriction is therefore rejected at the 5% significance level, indicating that allowing for action bias significantly improves model fit.

We can compare the model estimations against the observed data on the plot below:

[1] Recall that at ε^* there are multiple equilibria, all equilibria at ε^*were computed fit tested, none led to better AD or SD values than the reported $\hat{\varepsilon}$ results.

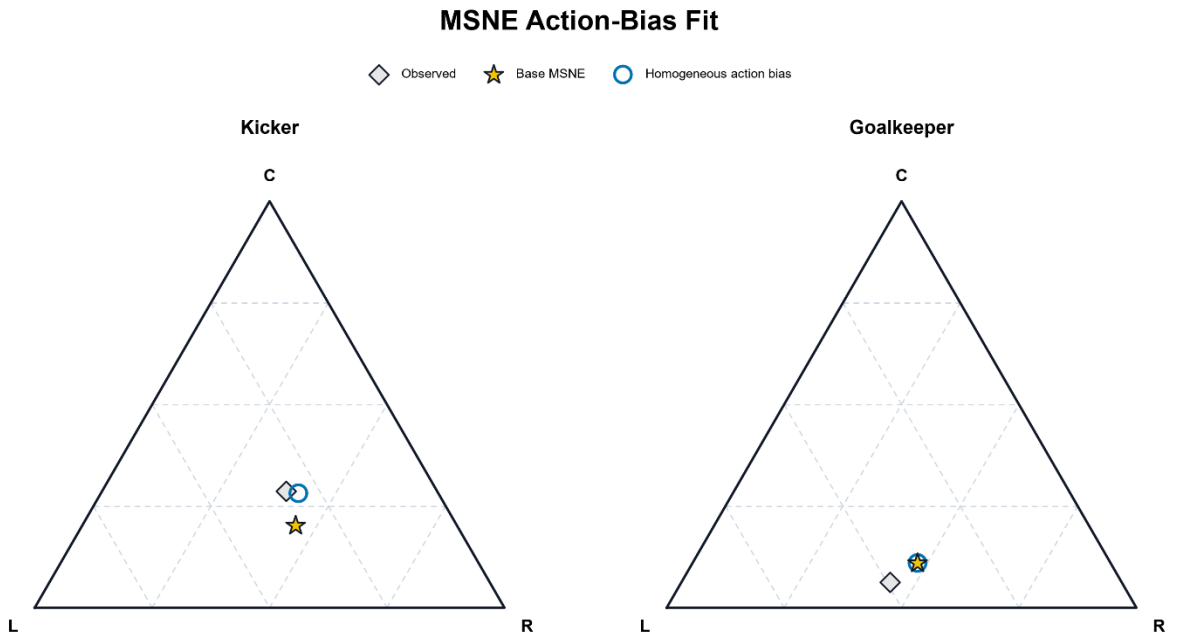


Figure 7 – MSNE action bias fitted mixes

Azar et al. created an MSNE structure to try and model the data they collected. They concluded that the MSNE fit their data best compared with other decision making frameworks, however the model they created still over predicted the goalie's rate of centre compared to real world play. Whilst they posit that the low centre rate of the goalie may be caused by action bias, they did not model action bias. By adjusting the goalie's payoffs to include action bias we saw a statistically significant improvement in fit over the base MSNE. Due to the indifference condition required for the MSNE model, the goalie's equilibrium mix cannot be changed without moving to a reduced support equilibrium. However overall fit still improved when action bias was included since the kicker's mix was brought closer to the observed data.

(4.3) Heterogeneous models

(4.3.1) Heterogeneity in ε

Chiappori et al. discuss the issue that there is not just one penalty kick game, but rather each individual matchup is different due to variation in kickers' footedness. This research has already attempted to reconcile this through footedness estimations and found that the aggregate scoring probabilities and direction frequencies alone are not enough to deconstruct heterogeneity. However, there may be another form of heterogeneity across players. The previous model assumes that all goalkeepers share the same value of epsilon, but in real

world play, goalkeepers may differ. Some goalkeepers may experience only a small psychological cost from staying centre and conceding, while others may experience a much larger cost. To see if this sort of behavioural heterogeneity is present within the data, heterogeneous action bias models will be tested.

To capture this possible heterogeneity, the model has been extended to include two goalkeeper types: a low action bias type and a high action bias type. This two type restriction will be justified below

The low type has action bias parameter:

$$\varepsilon_{low}$$

and the high type has action bias parameter:

$$\varepsilon_{high}$$

where:

$$\varepsilon_{high} > \varepsilon^* > \varepsilon_{low}$$

The proportion of high action bias goalkeepers in the population is denoted by:

$$\pi$$

Therefore, the proportion of low action bias goalkeepers is:

$$1 - \pi$$

This new model will be assessed under the complete information assumption[2]. The kicker knows which type of goalkeeper they are facing. This means that if the goalkeeper is a low action bias type, the kicker plays the equilibrium strategy for the low type. If the goalkeeper is a high action bias type, the kicker plays the equilibrium strategy for the high type. Therefore, the model solves two separate equilibria, one for each type.

The observed (aggregate) distribution is then a weighted average of the two type specific equilibria:

$$Pr(i,j) = (1-\pi)Pr(i,j|\varepsilon_{low}) + \pi Pr\big(i,j|\varepsilon_{high}\big)$$

[2] The incomplete information case will be covered after in section (4.3.4)**Error! Reference source not found.**

Because the likelihood estimation is run over the aggregate high/low type mixes, the heterogeneous model allows some share of reduced support play without causing the zero probability log likelihood issue discussed earlier. This means running AD and SD estimations are no longer necessary.

As ε is unidentified beyond the critical value, we fix ε_{high} to 0.7609 and search over ε_{low} and π. The assumption of fixing ε_{high} to 0.7609 will be justified below.

Using maximum likelihood, the model estimates:

ε_{low}	π	log likelihood
0.028702	0.077042	-570.107401

Table 21 – Heterogeneous MSNE fitted parameters

This means that the model estimates around 7.7 percent of goalkeepers to be high action bias types, while the remaining 92.3 percent are low action bias types.

Looking at the predicted marginal strategies of each type will give insight into the structure of the aggregated mix.

For the low type, the marginal strategies are:

	L	C	R
Kicker	0.3226	0.2392	0.4381
Goalie	0.4110	0.1101	0.4789

Table 22 – Low type predicted MSNE mixes

For the high type, the marginals are:

	L	C	R
Kicker	0.0000	0.7938	0.2063
Goalie	0.6938	0.0000	0.3063

Table 23 – High type predicted MSNE mixes

The pooled model prediction is:

	L	C	R
Kicker	0.2978	0.2819	0.4203
Goalie	0.4328	0.1016	0.4656

Table 24 – Pooled heterogeneous MSNE mixes

We see that for the high type equilibrium, each player only plays two of their available pure strategies. This mix is consistent with the reduced support calculations found in the appendix.

Under the homogeneous action bias setup, we can only fit the kicker, but under the heterogeneous setup some reduced support play can be used to increase goalie fit without damaging kicker fit.

The player specific likelihoods evaluate how closely the model reproduces each player's marginal frequencies of choosing left, centre and right. As an oppose to joint likelihood used previously which evaluates the predicted frequencies of all nine kicker/goalkeeper direction combinations.

Model	**Kicker log likelihood**	**Goalkeeper log likelihood**	**Joint log likelihood**
Homogeneous MSNE + action bias	−312.351196	−258.591382	−570.942578
Heterogeneous MSNE + action bias	−312.351235	−256.329583	−570.107401
Improvement from heterogeneity	−0.000040	+2.261799	**+0.835177**

Table 25 – Player vs joint MSNE fit

We see that all improvement in the joint likelihood is driven by improvement in goalie fit. Allowing for heterogeneity within ε can only replicate the correlation within the data through variation in the player's marginal strategies. From tables 6 and 8 we know that fitting each player's marginal frequencies perfectly will not fully replicate the joint distribution of the 9 direction combinations.

Note that for row two the sum of the player likelihoods is not the same as the joint likelihood like we see for row 1. This is due to the way the joint likelihood is calculated compared to the separate player likelihoods. Further explanation can be found in the appendix. This means that once heterogeneity within ε is introduced the joint distribution over direction combinations is no longer equal to the product of the marginals. It is now:

$$P_{ij}(\theta) = (1-\pi)p_i(\varepsilon_{Low})q_j(\varepsilon_{Low}) + \pi p_i(\varepsilon_{High})q_j(\varepsilon_{High})$$

Whereas the pooled (or aggregate) player strategies are:

$$p_i^{pooled} = (1-\pi)p_i(\varepsilon_{Low}) + \pi p_i(\varepsilon_{High})$$

$$q_j^{pooled} = (1-\pi)q_j(\varepsilon_{Low}) + \pi q_j(\varepsilon_{High})$$

Therefore, $P_{ij}(\theta) \neq p_i^{pooled} q_j^{pooled}$

This explains why we see larger fit improvements to the player marginals than to the joint likelihood. Improving marginal fit can improve the overall model fit through directional frequency improvements, yet this is separate from replicating the correlation within the observed data.

(4.3.2) Other ε distributions

Because equilibrium strategies above ε^* are independent of ε (as shown by Proposition 1), any ε distribution can be represented by a distribution conditional on $(\varepsilon\,|\varepsilon < \varepsilon^*)$ and a separate mass of reduced support play defined to be π.

The generic heterogeneous equation then becomes.

$$Pr(ij) = (1-\pi)Pr(ij|\varepsilon < \varepsilon^*) + \pi Pr(ij|\varepsilon > \varepsilon^*)$$

The probability of observing each direction combination is the weighted average of observing the combination in the full and reduced support types.

Given that the equilibrium mix is constant across all reduced support equilibria, $Pr(ij|\varepsilon > \varepsilon^*)$ is constant. Going forward we will denote this as$(\Pr ij|(\varepsilon^*)$

$$given\ \varepsilon < \varepsilon *$$

$$Pr(ij|\varepsilon) = p_i^K(\varepsilon) * p_j^G$$

Since p_j^G is constant for every $\varepsilon < \varepsilon^*$, we do not need to express it as a function of ε

Therefore, for all distributions of ε,

$$Pr(ij) = (1-\pi)\,\mathrm{E}\left(p_i^K(\varepsilon) * p_j^G|\varepsilon < \varepsilon^*\right) + \pi\,\Pr(ij|\varepsilon^*)$$

$$= (1-\pi)\,p_j^G\,\mathrm{E}\left(p_i^K(\varepsilon)|\varepsilon < \varepsilon^*\right) + \pi\,\Pr(ij|\varepsilon^*)$$

Showing, the only things the data can identify within the model are $E\left(p_i^K(\varepsilon)|\varepsilon < \varepsilon^*\right)$ and the share of reduced support play (π).

Define:

$$p^K(\varepsilon) = \begin{pmatrix} \alpha(\varepsilon) \\ \beta(\varepsilon) \\ 1-\alpha(\varepsilon)-\beta(\varepsilon) \end{pmatrix}, \qquad 0 \le \varepsilon < \varepsilon^*$$

The equilibrium probabilities for the kicker's mix within the full support region are:

$$\alpha(\varepsilon) = \frac{127}{275} - \frac{32}{55}\beta(\varepsilon)$$

$$\beta(\varepsilon) = \frac{\left(\varepsilon + \frac{4699}{34375}\right)}{\frac{18541}{27500} + \frac{3}{5}\varepsilon}$$

$$1 - \alpha(\varepsilon) - \beta(\varepsilon) = \frac{148}{275} - \frac{23}{55}\beta(\varepsilon)$$

Therefore,

$$p^K(\varepsilon) = \begin{pmatrix} \frac{127}{275} - \frac{32}{55}\beta(\varepsilon) \\ \beta(\varepsilon) \\ \frac{148}{275} - \frac{23}{55}\beta(\varepsilon) \end{pmatrix}$$

$$p^K(\varepsilon) = \begin{pmatrix} \frac{127}{275} \\ 0 \\ \frac{148}{275} \end{pmatrix} + \beta(\varepsilon) \begin{pmatrix} -\frac{32}{55} \\ 1 \\ -\frac{23}{55} \end{pmatrix}$$

Because $\beta(\varepsilon)$ is a strictly increasing continuous function, $\{\beta(\varepsilon): 0 \leq \varepsilon < \varepsilon^*\}$ is an interval and therefore convex.

Because $\alpha(\varepsilon)$ and $1 - \alpha(\varepsilon) - \beta(\varepsilon)$ are both linear functions of $\beta(\varepsilon)$, the set $\{p^K(\varepsilon) | \varepsilon < \varepsilon^*\}$ is convex.

From this it follows that $E[p^K(\varepsilon) | \varepsilon < \varepsilon^*] = p^K(\hat{\varepsilon}) for\ some\ \hat{\varepsilon} \in [0, \varepsilon^*]$

Therefore, the average strategy produced by any distribution below ε^* can be represented by a single value of $\varepsilon_{low} < \varepsilon^*$. Because strategies above $\varepsilon^*$are identical, it follows that any distribution of ε can be replicated by a two point high/low distribution where $\varepsilon_{high} > \varepsilon^* > \varepsilon_{low}$.

(4.3.3) Interpretation of fitted MLE parameters

Given that the kicker's strategy is convex, their aggregated strategy mix will only ever lie on the green line below:

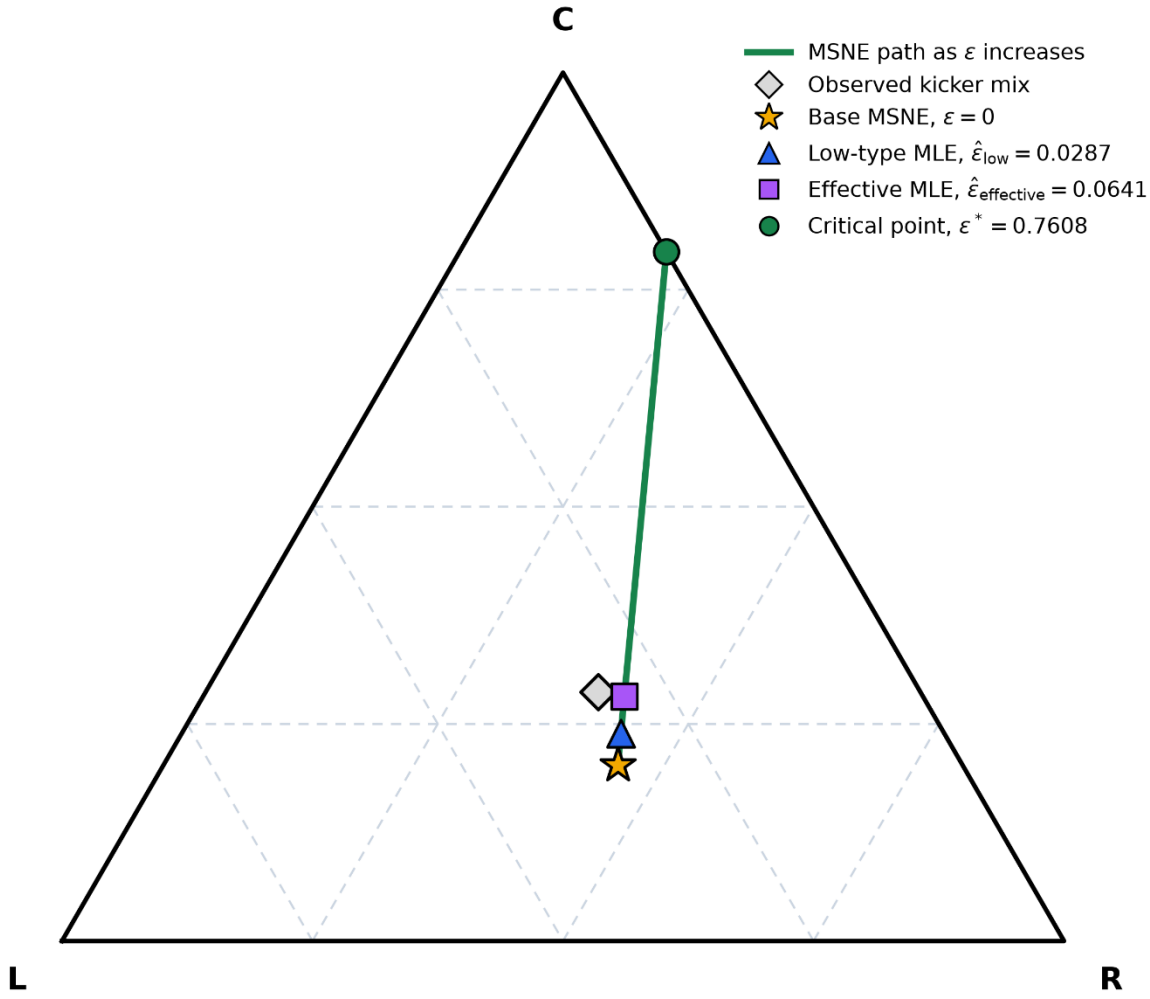


Figure 8 – Heterogeneous MSNE kicker mixes

Figure 8 represents the pooled kicker's strategy.

$$p_i^{pooled} = p_i\ (\varepsilon_{effective}) = (1-\pi)p_i(\varepsilon_{Low}) + \pi p_i(\varepsilon_{High})$$

The blue triangle represents the full support strategy associated with ε_{Low}. The green circle represents the reduced support strategy associated with ε_{High}. The purple square represents the pooled strategy. The distance from the blue triangle to the purple square is determined by π.

The goalie's pooled strategy is only determined by π, and from looking at Figure 6 we can see that the MLE goalie's marginal fit could be improved by increasing the amount of reduced support play (π). To fit the goalkeeper's centre rate π would need to be 0.43. However, doing so would worsen kicker fit which is why we see the MLE value of π being 0.077. Essentially too much reduced support play would drag the kicker's pooled mix (purple square) too far

along the green line, thus worsening kicker fit. Therefore, we would expect to see ε_{low} predicted to be 0. That way we could maximise reduced support play without damaging kicker fit by dragging the mix too far along the green line. However, the MLE imputes ε_{low} to be 0.0287.

This seemingly unintuitive MLE result is caused by the correlation in the data discussed earlier.

The heterogeneous action bias model is limited in the extent to which it can replicate the correlations within the observed data. Comparing the observed versus the MLE predicted centre marginal rates for each player:

Centre rate	Kicker	Goalie
Heterogeneous MLE	0.2819	0.1016
Observed data	.287	.063

We see the model slightly underpredicts the kicker's centre rate and overpredicts the goalie's centre rate. A larger π and lower ε_{low} would lead to a reduction in the goalie's marginal centre rate whilst keeping the kicker's pooled centre rate the same.

Recall (C,C) is overrepresented in the data relative to the overall goalie centre rate.

Given the kicker kicks to the centre the conditional goalie rates are:

Goalkeeper action (j)	**Observed conditional rate** $P(j \mid i = C)$	**Observed marginal goalie rate** $P(j)$
Left	50.0%	49.3%
Centre	12.2%	6.3%
Right	37.8%	44.4%
Total	**100.0%**	**100.0%**

Table 26 – Conditional goalie frequency rates

The MLE predicted goalie rates given the kicker kicks to the centre are:

Goalkeeper action	**predicted conditional rate** $P(j \mid i = C)$
Left	47.2%
Centre	8.6%
Right	44.1%
Total	**100.0%**

Table 27 – Conditional goalie predicted rates

We see that the model imputes a conditional goalie centre rate higher than the observed goalie marginal frequency, but still less than the observed conditional frequency. $6.3\% < 8.6\% < 12.2\%$

Given the kicker kicks to the centre, the predicted goalie centre rate is:

$$P(G = C \mid K = C) =$$

$$(1-\pi)\frac{p_C(\varepsilon_{\text{low}})q_C(\varepsilon_{\text{low}})}{(1-\pi)p_C(\varepsilon_{\text{low}}) + \pi p_C(\varepsilon_{\text{high}})} + \pi\frac{p_C(\varepsilon_{\text{high}})q_C(\varepsilon_{\text{high}})}{(1-\pi)p_C(\varepsilon_{\text{low}}) + \pi p_C(\varepsilon_{\text{high}})}.$$

Because $q_C(\varepsilon_{\text{high}}) = 0$, the conditional probability is:

$$P(G = C \mid K = C) = (1-\pi)\frac{p_C(\varepsilon_{\text{low}})q_C(\varepsilon_{\text{low}})}{(1-\pi)p_C(\varepsilon_{\text{low}}) + \pi p_C(\varepsilon_{\text{high}})} + \pi(0)$$

Therefore, we see that an increase in π would reduce $P(G = C \mid K = C)$. Thus, although a smaller ε_{low} and higher π would fit the goalie's marginal strategy better, it would lead to a worse fit to the correlation within the data. Therefore, even though the imputed result of ε_{low} = 0.028702 may seem unintuitive, it is an accurate fitted parameter caused by the overrepresentation of (C,C) in the observed data.

Comparing the log likelihood values of heterogeneous vs homogeneous action bias, we see in Table 25 that the heterogeneous version of MSNE improves log likelihood fit over the homogeneous model. This improvement of fit is caused by the small proportion of reduced support play. Goalie strategy cannot be altered without moving to a reduced support equilibrium. Because doing so reduces kicker fit drastically, the homogeneous model cannot fit the data by adjusting the goalie's mix. The heterogeneous model can adjust the pooled goalie's strategy by adjusting the proportion of reduced support play. The pooled prediction fits the observed data better and allows reduced goalie centre play compared to (Azar & Bar-Eli, 2011) initial MSNE prediction.

To see if this improvement in fit is statistically significant over the homogeneous model a Likelihood ratio test was done.

LR statistic $= 2 \times (-570.10741 + 570.942578) = 1.670354$

Using a standard chi-squared approximation with one degree of freedom, this gives an approximate p-value of 0.196, showing the improvement is too small to reject the homogeneous model at the 10% level using the LR comparison. However, under the homogeneous null, the second goalkeeper type has zero weight and is not identified, so the

standard chi-square assumptions do not hold. Therefore, we use a direct bootstrap comparison, which returns $p = 0.098$. The heterogeneous model is not significant at 5% but provides weak evidence at 10% under the bootstrap comparison.

(4.3.4) Incomplete information assumption

We now also test for heterogeneity under the incomplete information assumption.

An MLE was first run for the MSNE model using a continuous ε distribution under the incomplete information assumption. This results in unidentified ε distribution. The kicker mix and pooled goalkeeper mix are well identified, but there are many parameter vectors that can replicate the same pooled mixes. The log likelihood value at these mixes is −570.942578. The fitted mixes under incomplete information are

	L	C	R
Kicker	0.2979	0.2817	0.4204
Goalie	0.4110	0.1101	0.4789

Table 28 – Incomplete information MSNE mixes

These are the same mixes as the homogeneous action bias MLE result. Therefore, the incomplete information assumption does not improve fit over the simpler homogeneous model.

Proposition 2. For any incomplete information MSNE with full support there exists an equivalent homogeneous MSNE. That is, it has the same joint distribution of kicks and dives.

The proof of proposition 2 can be found in the appendix.

(4.3.5) MSNE model summary

Looking at Figure 9 below we see that allowing for action bias improves kicker fit for both the homogeneous and heterogeneous action bias models. However, allowing for some reduced support play in the complete information heterogeneous action bias model also slightly improves goalie fit, yet this fit improvement is not statistically significant. The incomplete information model reproduces the same equilibrium mixes and therefore the same likelihood as the homogeneous ε model. Therefore, heterogeneity within action bias has weak explanatory power within the MSNE framework.

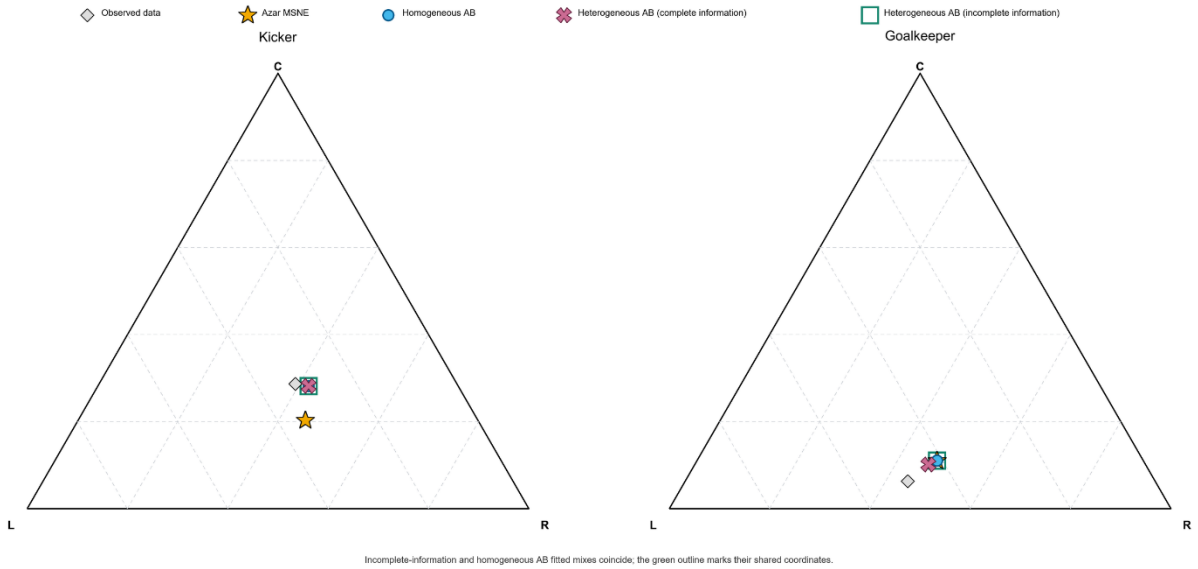


Figure 9 – MSNE model comparison

Model	Log likelihood	Kicker L	Kicker C	Kicker R	Goalie L	Goalie C	Goalie R
Observed data	—	0.3217	0.2867	0.3916	0.4930	0.0629	0.4441
Azar MSNE, no action bias	−576.2497	0.3439	0.2028	0.4534	0.4110	0.1101	0.4789
MSNE + Homogeneous AB	−570.9426	0.2979	0.2817	0.4204	0.4110	0.1101	0.4789
MSNE + Heterogeneous AB (complete information)	−570.1074	0.2978	0.2819	0.4203	0.4328	0.1016	0.4656
MSNE + Heterogeneous AB (incomplete information)	−570.9426	0.2979	0.2817	0.4204	0.4110	0.1101	0.4789

Table 29 – MSNE model marginal predictions

We also compare the cell likelihood calculated as $n_{ij} \ln(p_{ij})$. This allows us to see where the fit differences from each model is coming from.

In the tables below, the column and row totals represent the sum of the individual cell likelihoods not the marginal likelihood of the overall strategy. This is why we see an improvement in the goalie's totals in the homogeneous AB table over the no AB table. The improvement in this row isn't caused by a change in goalkeeper strategy, but rather a joint improvement driven by the kicker's change in strategy against a fixed goalie. The incomplete

information case is not presented as the joint distribution is the same as the homogeneous action bias case and is therefore redundant.

Azar et al. MSNE (no action bias)

	Goalkeeper			
Kicker	**Left**	**Centre**	**Right**	**Total**
Left	−105.664	−3.274	−66.738	−175.676
Centre	−101.884	−38.023	−72.291	−212.198
Right	−77.289	−20.983	−90.103	−188.375
Total	−284.838	−62.280	−229.132	**−576.250**

Table 30 – Base MSNE cell likelihood

MSNE + homogeneous AB

	Goalkeeper			
Kicker	**Left**	**Centre**	**Right**	**Total**
Left	−113.408	−3.417	−72.043	−188.869
Centre	−88.400	−34.734	−62.096	−185.230
Right	−80.767	−21.512	−94.564	−196.843
Total	−282.575	−59.664	−228.703	**−570.943**

Table 31 – MSNE homogeneous action bias cell likelihood

MSNE + Heterogeneous AB (complete information)

	Goalkeeper			
Kicker	**Left**	**Centre**	**Right**	**Total**
Left	−113.432	−3.418	−72.060	−188.910
Centre	−82.664	−37.171	−64.595	−184.430
Right	−79.596	−21.783	−95.388	−196.767
Total	−275.692	−62.373	−232.042	**−570.107**

Table 32 – MSNE heterogeneous CI action bias cell likelihood

Table 31 is shaded to make it easier to compare the no action bias model to the homogeneous action bias model. The red shading shows the cell fits worse relative to the Table 30 equivalent, whilst green shading shows likelihood improvement over the base estimation. Table 32 is

shaded relative to Table 31 in the same fashion. We see that adding action bias improves the kicker's centre row, yet all other joint cells fit worse. This gives interesting insight into the mechanics of the MSNE model, the overall joint likelihood still improves to a statistically significant degree when action bias is included, yet all this improvement is being driven by the kicker increasing C against a fixed goalkeeper.

For Table 32, only two of the 9 direction combinations fit better under complete information heterogeneity compared to homogeneity, however these improvements are large enough to compensate for the worse fitting cells and still improve the overall joint likelihood, though not to a significant level. Reinforcing the conclusion that heterogeneity within action bias has little explanatory power.

When action bias is large enough it can rationalise the goalkeeper's underuse of centre within a Nash equilibrium framework, but only at the cost of distorting the kicker's mix far beyond what is observed in the data. Once the cost of staying centre and conceding becomes large enough, the goalkeeper completely removes centre from their equilibrium strategy. This is consistent with the idea that a goalkeeper may avoid staying in the centre because conceding from that position creates an additional psychological cost.

However, the way this happens in the Nash equilibrium model is quite sharp. The goalkeeper does not gradually reduce their probability of staying centre as ε increases. Instead, the goalkeeper's centre probability remains constant in the full support equilibrium, then falls to zero once the model moves to the reduced support equilibrium. This is important because it shows both the strength and the limitation of the Nash framework. The model can explain why the goalkeeper may avoid staying centre, but it does so through a sudden support change rather than through a smooth behavioural response to increasing action bias.

This adds additional support on top of the reasons given in the literature review for using the Quantal Response Equilibrium model in the next section. In the QRE framework, all available actions are always played with positive probability, meaning the likelihood function can be used universally without any log(0) errors occurring. Additionally, changes in expected utility can affect the probability of choosing an action directly and smoothly in a more intuitive direction. Therefore, if action bias makes staying centre less attractive for the goalkeeper, the QRE model should allow the goalkeeper's probability of staying centre to fall gradually as the size of the bias increases.

(5) Quantal response equilibrium

(5.1) QRE background and comparative statics

(5.1.1) General logit QRE behaviour

In a quantal response equilibrium, the probability of picking an action is an increasing function of the expected utility of that action relative to all other choices. The logit quantal response model is by far the most used specification of the quantal response model.

$$p(_k) = \frac{e^{\lambda EU(k)}}{\sum_{l \in A} e^{\lambda Eu(l)}}$$

Where $p(k)$ is the probability of the player choosing strategy (k)

$EU(k)$ is the expected utility of choosing strategy (k)

A represents the set of all available actions

λ represents how close to optimally the player will act. If λ is 0 the player will fully randomise their actions, as $\lambda \to \infty$, the QRE converges to the mixed strategy Nash equilibrium.

The denominator represents the summed exponentiated utility of all available actions. Therefore, this formula shows how players use bounded rationality to make probabilistic best response choices. Players choose to play strategy j based on that action's expected utility relative to the summed utility of all possible actions. Actions with higher expected payoffs are chosen at higher rates. The QRE model represents stochastic choice, unlike ordinary best responding which assumes deterministic choices.

Therefore, the formula represents a player logit best responding by playing action j with a probability that is proportional to its expected utility relative to the total expected utility when summing all possible actions

The rationality parameter λ is a model parameter varied in comparative statics and subsequently estimated, so it is important to understand how varying it can impact the model's equilibrium predictions.

Let's assume a player has two pure strategy choices in a game.

Here $A = \{a, b\}$ therefore, $l \in \{a, b\}$

The Logit QRE probabilities for this player then become

$$P(a) = \frac{e^{\lambda EU(a)}}{e^{\lambda EU(a)} + e^{\lambda EU(b)}} = \frac{1}{1 + e^{\lambda[EU(a) - EU(b)]}}$$

$$P(b) = 1 - P(a) = \frac{1}{1 + e^{\lambda[EU(b) - EU(a)]}}$$

Note: if $\lambda = 0, P(a) = \frac{1}{2}$ if $\lambda \to \infty$ $then$ $P(a) \to \begin{cases} 1 \; if \; EU(a) > EU(b) \\ \frac{1}{2} \; if \; EU(a) = EU(b) \\ 0 \; if \; EU(a) < EU(b) \end{cases}$

When the expected payoff between strategies is high, the strategy with the larger payoff is chosen at a higher rate relative to if the utility difference was smaller.

The rationality parameter (λ) then multiplies the effect of these payoff differences. When λ = 0 the probability ratio is always 1 regardless of utility differences. This gives the uniformly mixed strategy shown above. As rationality increases, the gap between chosen probabilities grows for any game where outcome utilities for each strategy are not exactly equal. Eventually at high enough levels of λ, the gap becomes so large that any lower expected utility strategy is played a probability rate near zero. This mimics the sharp best responding of the MSNE model.

This choice behaviour can be visualised on a best response plot.

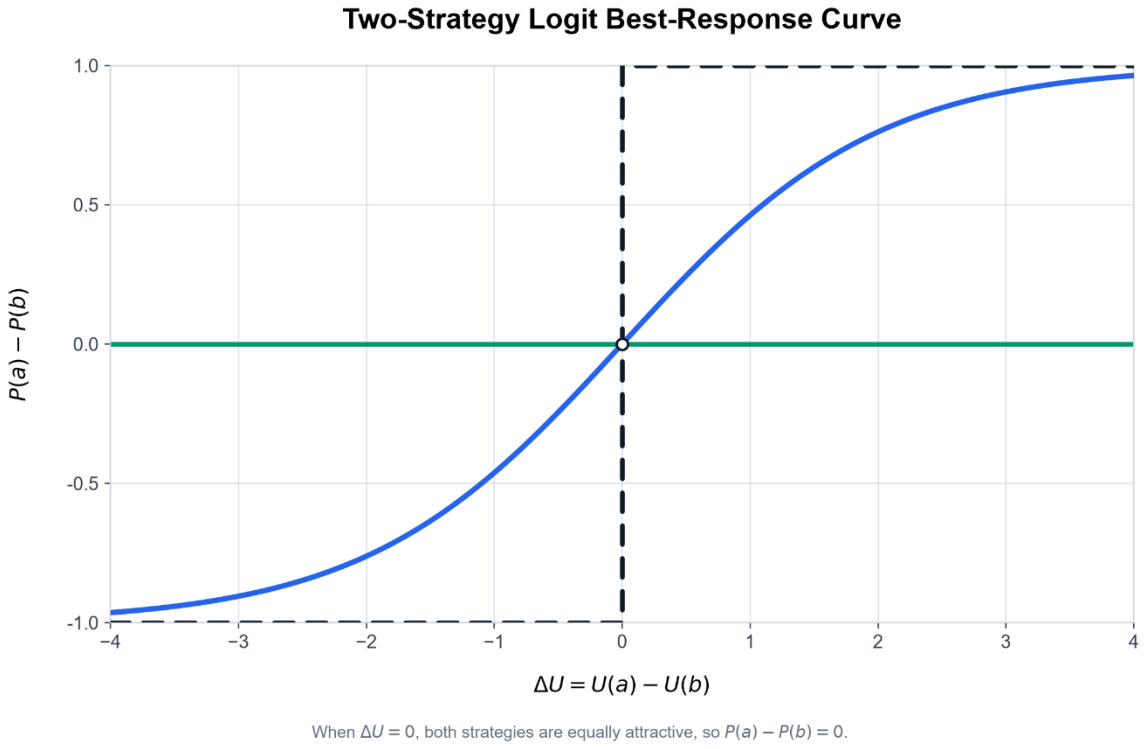


Figure 10 – Two strategy logit best response curve



This plot shows the chosen mixing frequencies of a two strategy game as the expected utility of each strategy varies. Note the y axis on the plot represents the probability difference not the probability ratio. This was done to make the plot more intuitive to read and still represents the same choice behaviour.

y=1 on the plot represents strategy a being chosen with probability 1.

y=-1 on the plot represents strategy b being chosen with probability 1.

y=0 on the plot represents uniform randomisation.

We see, for a fixed λ, increasing the utility difference increases the gap between chosen probabilities.

Here, λ changes the bending of the logit best response function, lower λ leads to a flatter curve, while higher λ leads to sharper best responding.

This is shown in the animation below.

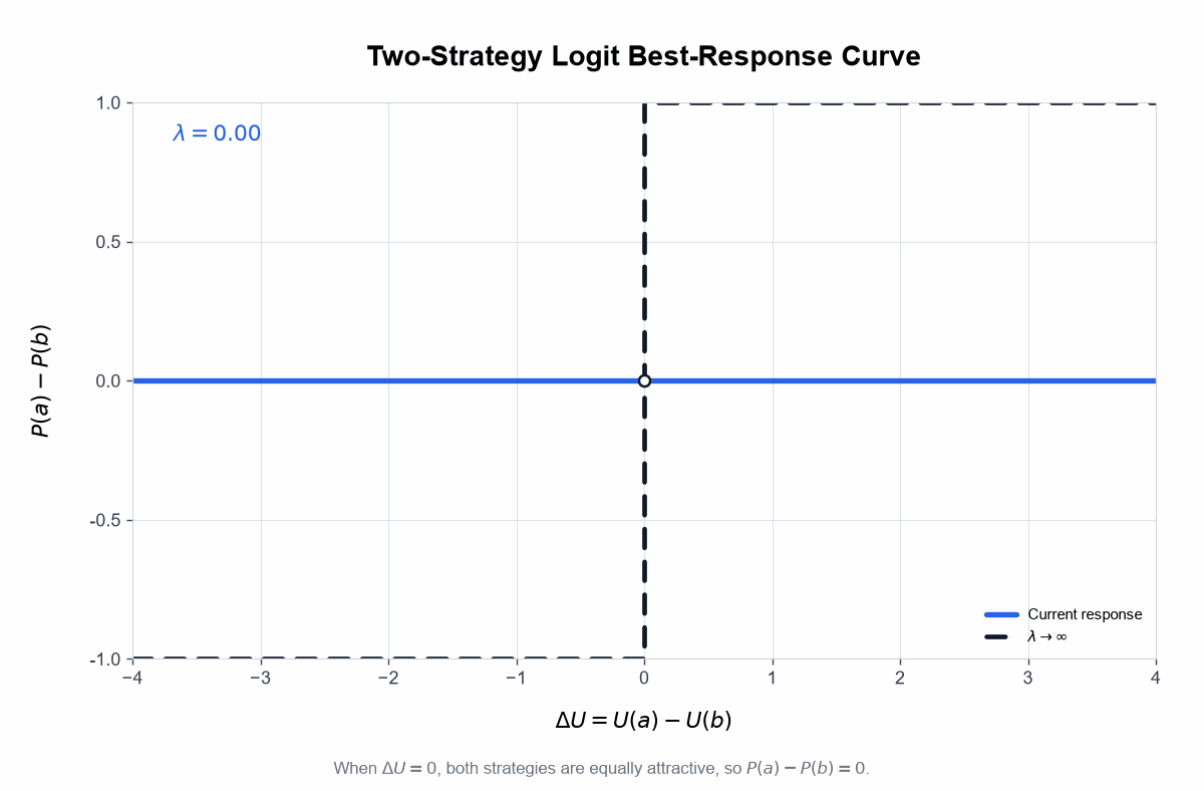


Figure 11 – Two strategy logit best response curve animated

At low levels of rationality, large utility differences (+4,-4) still leave both strategies with meaningful choice probabilities. As λ increases the curve asymptotically approaches the upper and lower boundaries (y=1, y=-1). These boundaries represent behaviour that is almost entirely concentrated on the higher utility strategy. One technical caveat is that the logit QRE model always places positive probability on all pure strategies, given a finite lambda. A larger λ makes choices more sensitive to utility differences, even a relatively small utility advantage can then produce a high probability of choosing the better performing strategy.

(5.2) Matching pennies example

An analysis of an asymmetric matching pennies game can be used to compare the equilibrium properties of the MSNE model and the QRE model.

(Goeree et al., 2003) present the following generalised matching pennies game.

	Player 2		
Player 1		Left	Right
	Up	x,0	0,1
	Down	0,1	1,0

We define U as the probability player 1 chooses Up and define L as the probability player 2 chooses Left

Substituting payoffs, x=1, x=4, and x=9, we get the following best response curves

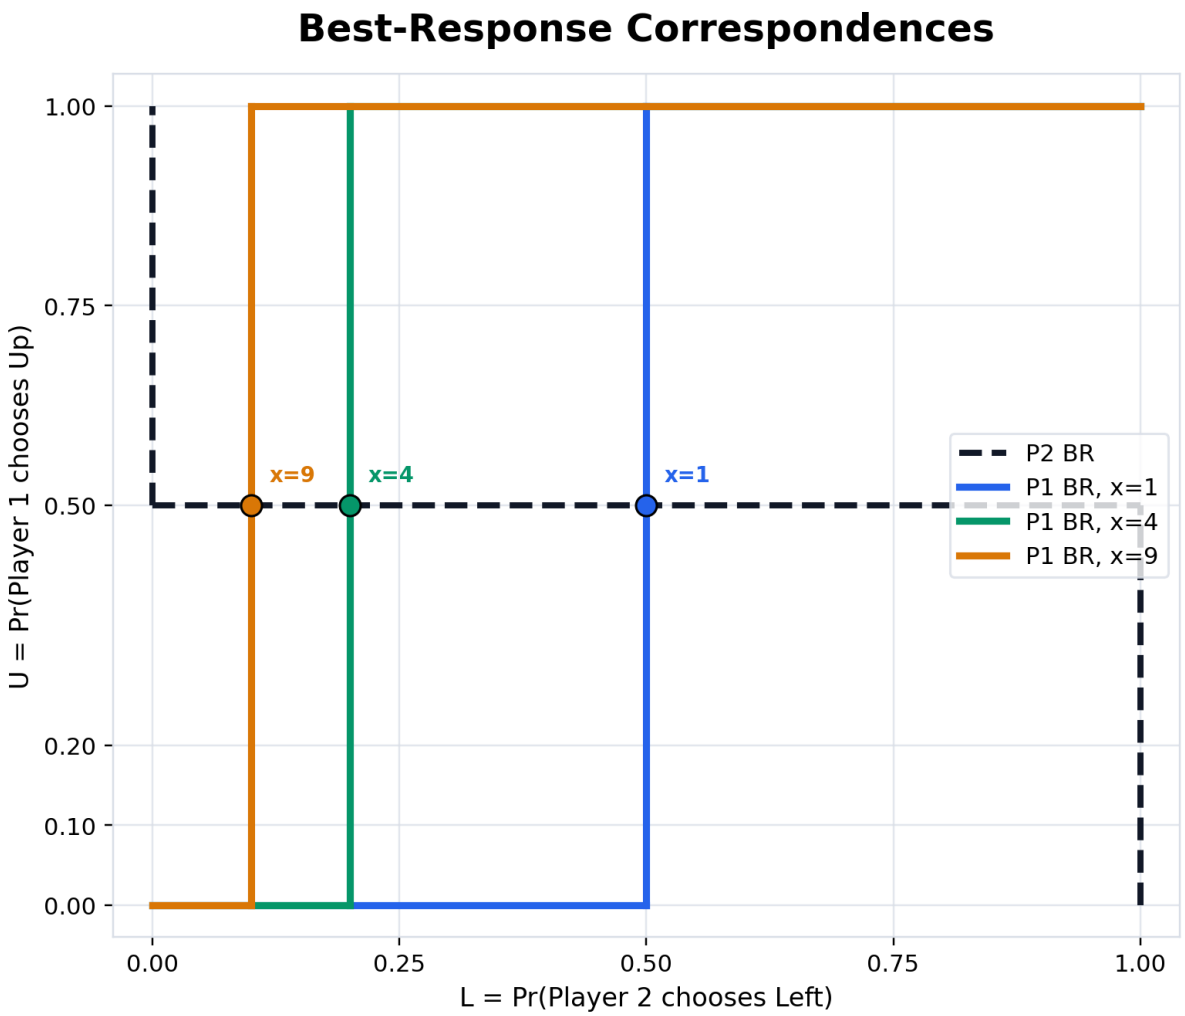


Figure 12 – Matching pennies MSNE response curves

The intersection of these curves gives the following MSNE points:

x	$Pr(Up)$	$Pr(Left)$
1	0.5	0.5
4	0.5	0.2
9	0.5	0.1

This illustrates the counterintuitive nature of a strictly mixed equilibrium under the MSNE model, we are adjusting player 1's payoffs, yet it is player 2's equilibrium mix that changes.

Looking now at the logit QRE for the same game we get the following equilibria:

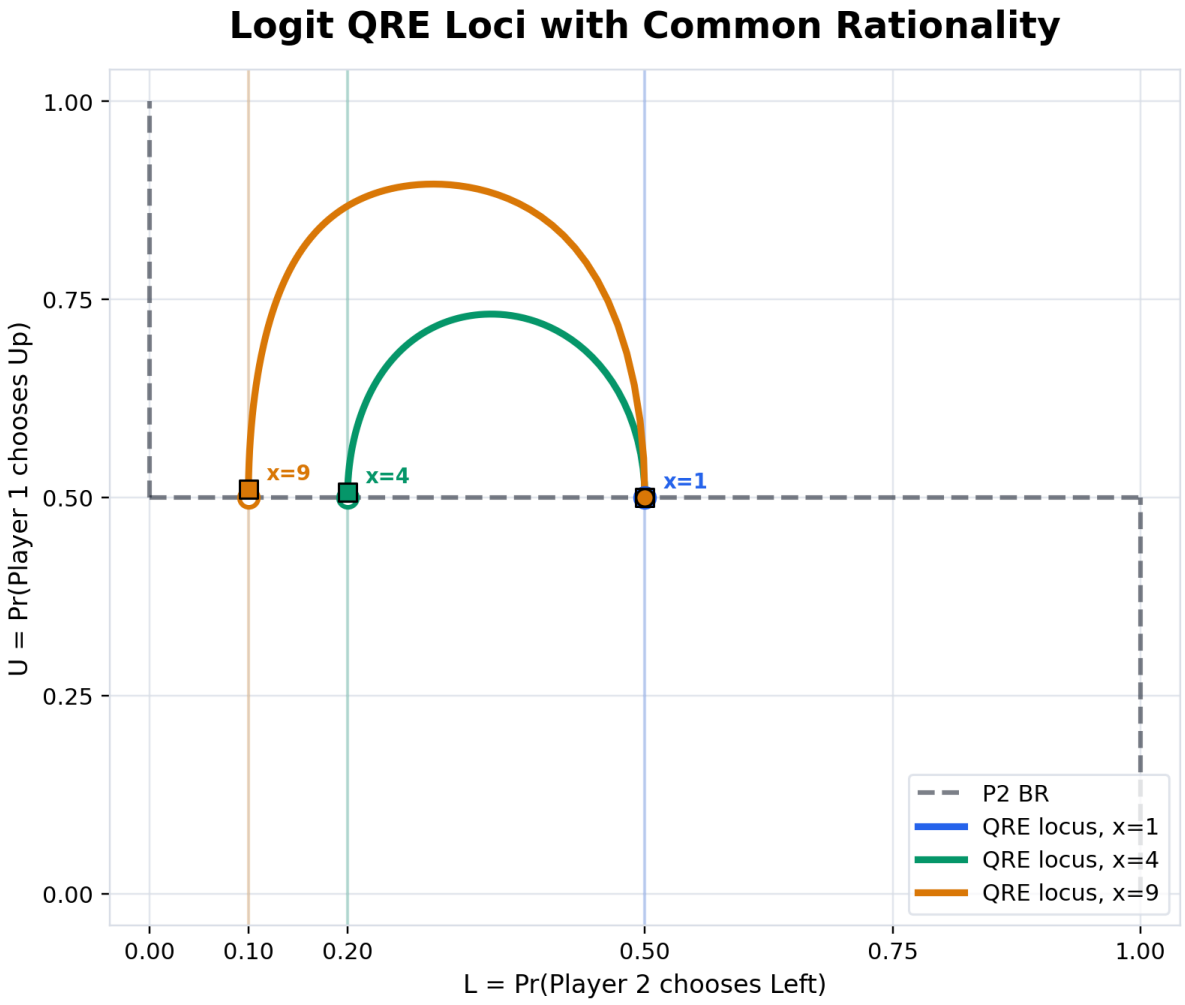


Figure 13 – Logit QRE matching pennies equilibria

The curves above show the locus of equilibria traced as the rationality parameter λ is increased from 0 to 100. At $\lambda = 0$ both players assign equal probability to all options, and at $\lambda = 100$ the Quantal Response Equilibrium approaches the Mixed Strategy Nash Equilibrium. When $x = 1$ uniform randomisation is the same as the MSNE mixes which is why we do not see a blue path being traced. However, when $x = 4$ and $x = 9$ we see that player 1 increases their probability of playing UP before reducing it to the MSNE probability of 0.5. Keeping expected utilities fixed and adjusting the rationality parameter (λ) we get different equilibria. Therefore, the Quantal response framework does not just give us one equilibrium, but a set of equilibria that depend on the chosen rationality parameter. For each value of λ, the QRE is found by solving both players' logit response equations simultaneously. The plotted paths therefore show how the equilibrium pair (U,L) changes as λ increases.

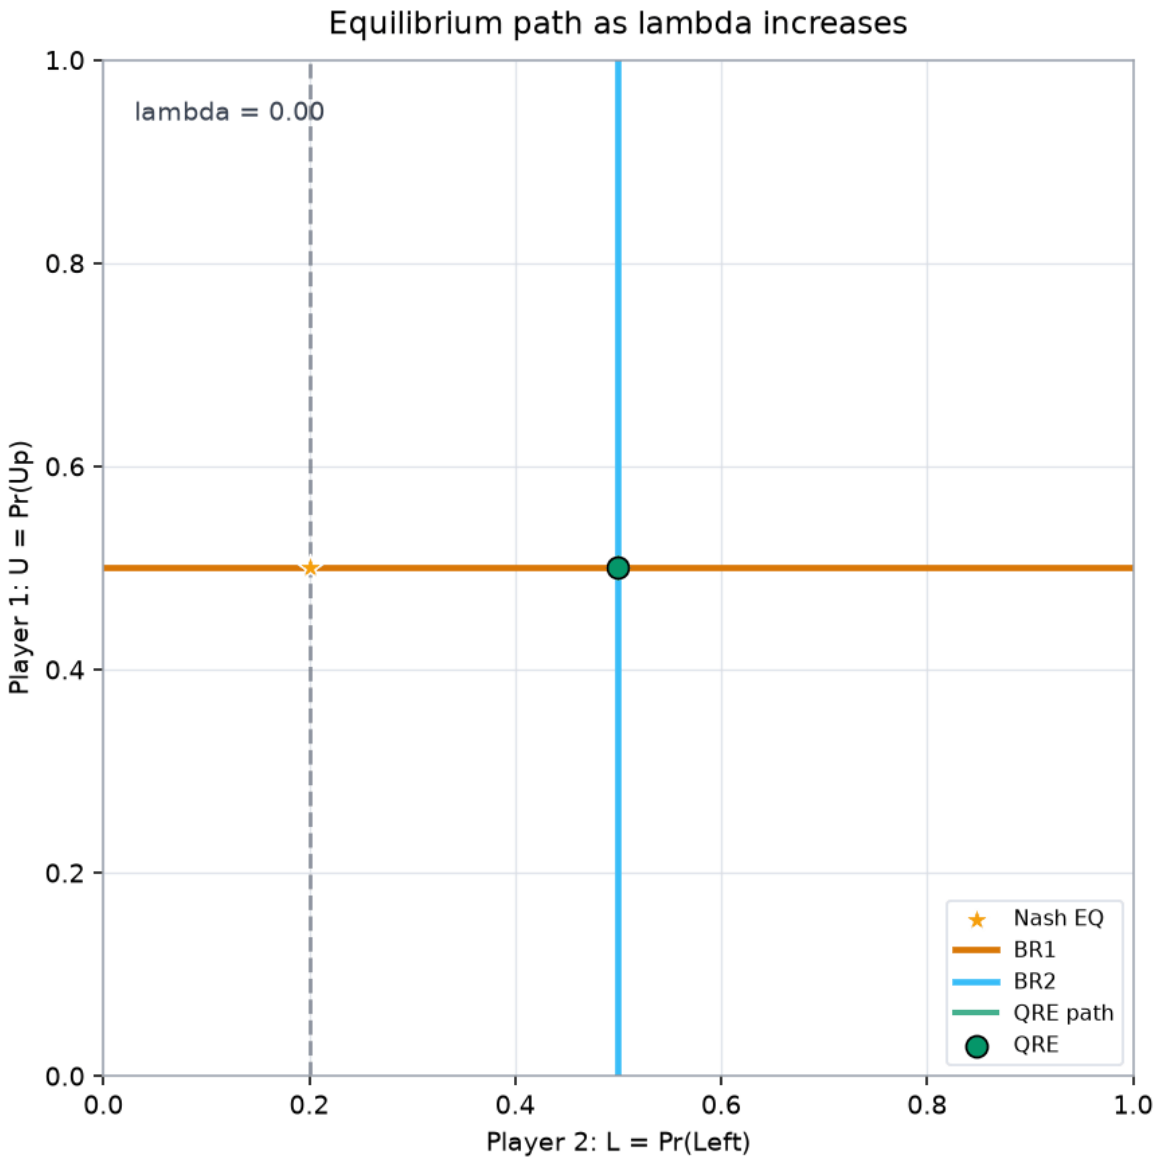


Figure 14 – Traced logit equilibria for matching pennies

Figure 14 gives an animation to show how the x=4 curve is traced out.

For this discussion we use x=4

Figure 13 illustrates an important difference between the comparative static predictions of MSNE and logit QRE. Under the strictly mixed Nash equilibrium, increasing x does not change player 1's equilibrium probability of Up, which remains $U = 0.5$. Instead, player 2 adjusts by reducing the probability of Left from $L = 0.5$ when $x = 1$ to $L = 0.2$ when $x = 4$ and $L = 0.1$ when $x = 9$. Thus, a change in player 1's payoff changes only the opponent's equilibrium strategy.

QRE produces a more direct comparative static response.

This is illustrated on Figure 15 below

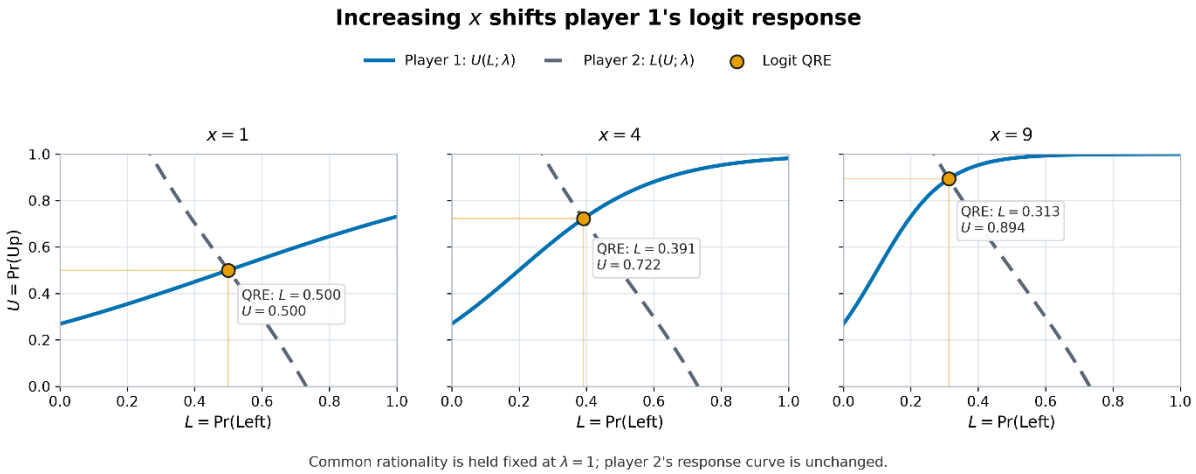


Figure 15 – Comparative statics as x increases

Here $\lambda = 1$ and x is varied. Unlike MSNE, QRE allows for an increase in x to directly increase player 1's probability of playing Up.

Figure 13 shows that a larger x value leads to a higher peak in the equilibrium path before bending back down. Figure 15 shows that for a given λ increasing x leads to a higher U in equilibrium. Togher these show the preferable comparative statics the QRE has when manipulating payoffs. However, even when payoffs are fixed the equilibrium can be moved in a desired direction by manipulation of the λ parameter.

This matching pennies exercise illustrates the main comparative statics advantage QRE has over MSNE. Adjusting a player's payoffs move their equilibrium probabilities in an intuitive direction. QRE does not require players to be indifferent between the actions they choose with positive probability. Instead, randomisation reflects stochastic responses to payoff differences, governed by the rationality parameter. This produces more intuitive comparative statics while ensuring that equilibrium play retains full support.

(6) QRE application to Azar et al.

From the previous MSNE section we observed that including action bias into the penalty kick model allowed for some improvement in fit. However, the MSNE indifference conditions prevent the goalie's equilibrium strategy from changing without introducing some reduced support play. Consequently, under full support, action bias and distorts the kicker's equilibrium mix rather than the goalie's.

(6.1) Base Logit QRE

For now, we will use the original (Azar & Bar-Eli, 2011) EU functions and scoring probabilities. We will reintroduce action bias later to show that QRE shifts probabilities smoothly.

To remind us of the original EU functions without action bias:

Kicker expected utility functions:

$$U_k(\mathrm{L}) = 0.704p + q + (1 - p - q)$$

$$U_k(\mathrm{C}) = 0.902p + 0.4q + 0.968(1 - p - q)$$

$$U_k(\mathrm{R}) = p + q + 0.746(1 - p - q)$$

Goalie expected utility functions:

$$U_g(\mathrm{L}) = -0.704\alpha - 0.902\beta - (1 - \alpha - \beta)$$

$$U_g(\mathrm{C}) = -\alpha - 0.4\beta - (1 - \alpha - \beta)$$

$$U_g(\mathrm{R}) = -\alpha - 0.968\beta - 0.746(1 - \alpha - \beta)$$

In equilibrium, each player is logit best responding to their opponent's mix.

Let $\theta_k = (\alpha, \beta, \gamma)$ denote the kicker's probabilities of choosing Left, Centre, and Right, and $\theta_G = (p, q, r)$ denote the corresponding goalie probabilities.

For the kicker we define:

$$Z_k = e^{\lambda U_K(Left;\theta_G)} + e^{\lambda U_K(Centre;\theta_G)} + e^{\lambda U_K(Right;\theta_G)}$$

We define Z_G for the goalie similarly.

Then the equilibrium probabilities of each action become:

$$\alpha = \frac{e^{\lambda U_K(Left;\theta_G)}}{Z_k}, \beta = \frac{e^{\lambda U_K(centre;\theta_G)}}{Z_k}, \gamma = \frac{e^{\lambda U_K(Right;\theta_G)}}{Z_k}$$

The same can be done for the goalie

$$p = \frac{e^{\lambda U_G(Left;\theta_k)}}{Z_G}, q = \frac{e^{\lambda U_G(Centre;\theta_k)}}{Z_G}, r = \frac{e^{\lambda U_g(Right;\theta_k)}}{Z_G}$$

Each value of λ results in its own unique equilibrium, varying λ from 0 to 100 allows us to trace the corresponding equilibria for the no action bias penalty kick game, similar to Figure 14. This is done below on Figure 16

Using computer modelling, the game was solved and plotted on the following simplexes:

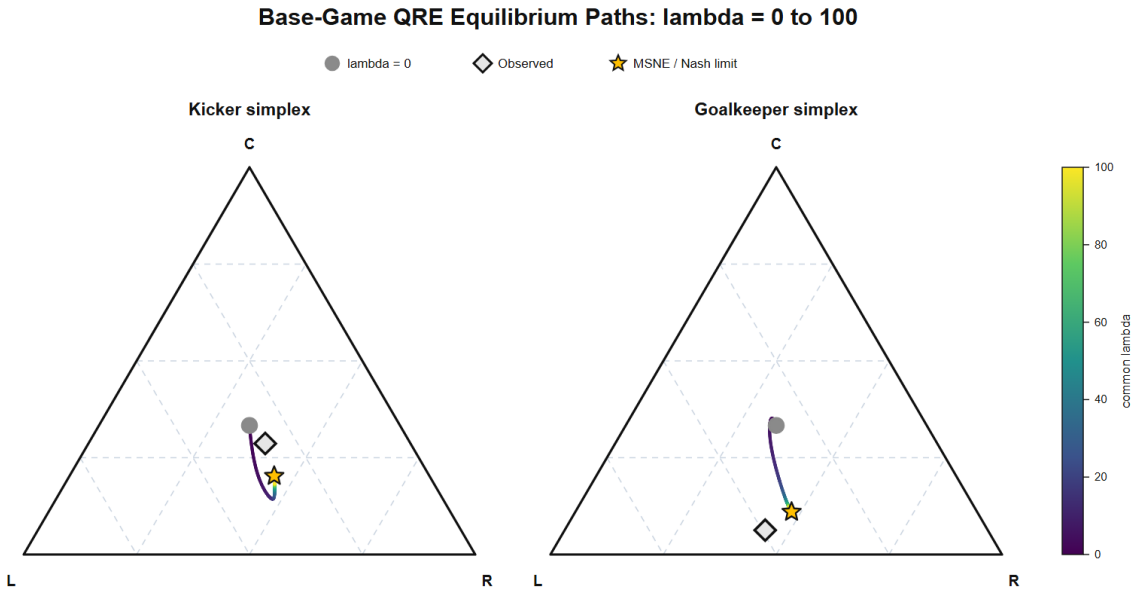


Figure 16 – Base game QRE equilibrium paths

We see that each player starts at complete randomisation when $\lambda = 0$, and approaches the Mixed Strategy Nash Equilibrium as rationality increases. This is consistent with what we know about the QRE model. Looking at the colour scale, we see that majority of the movement away from complete randomisation happens at low levels of λ and that the equilibrium does not move much at higher levels of λ.

Allowing separate λ parameters substantially increases the range of possible equilibria. We will call these separate rationality parameters λ_K and λ_G, for the kicker and goalie respectively. Under the common λ specification, each player has only one equilibrium path for the entire game as seen on Figure 16. With separate λs each player now has multiple equilibrium paths, each dependent on their opponent's λ value.

For example, we fix the kicker's rationality parameter to 0, 1, 3, 5, 10, and 100. We then vary the goalie's rationality and plot the corresponding equilibrium mixes for each value of λ_K and λ_G. This gives 6 separate goalie equilibrium paths as λ_G is varied.

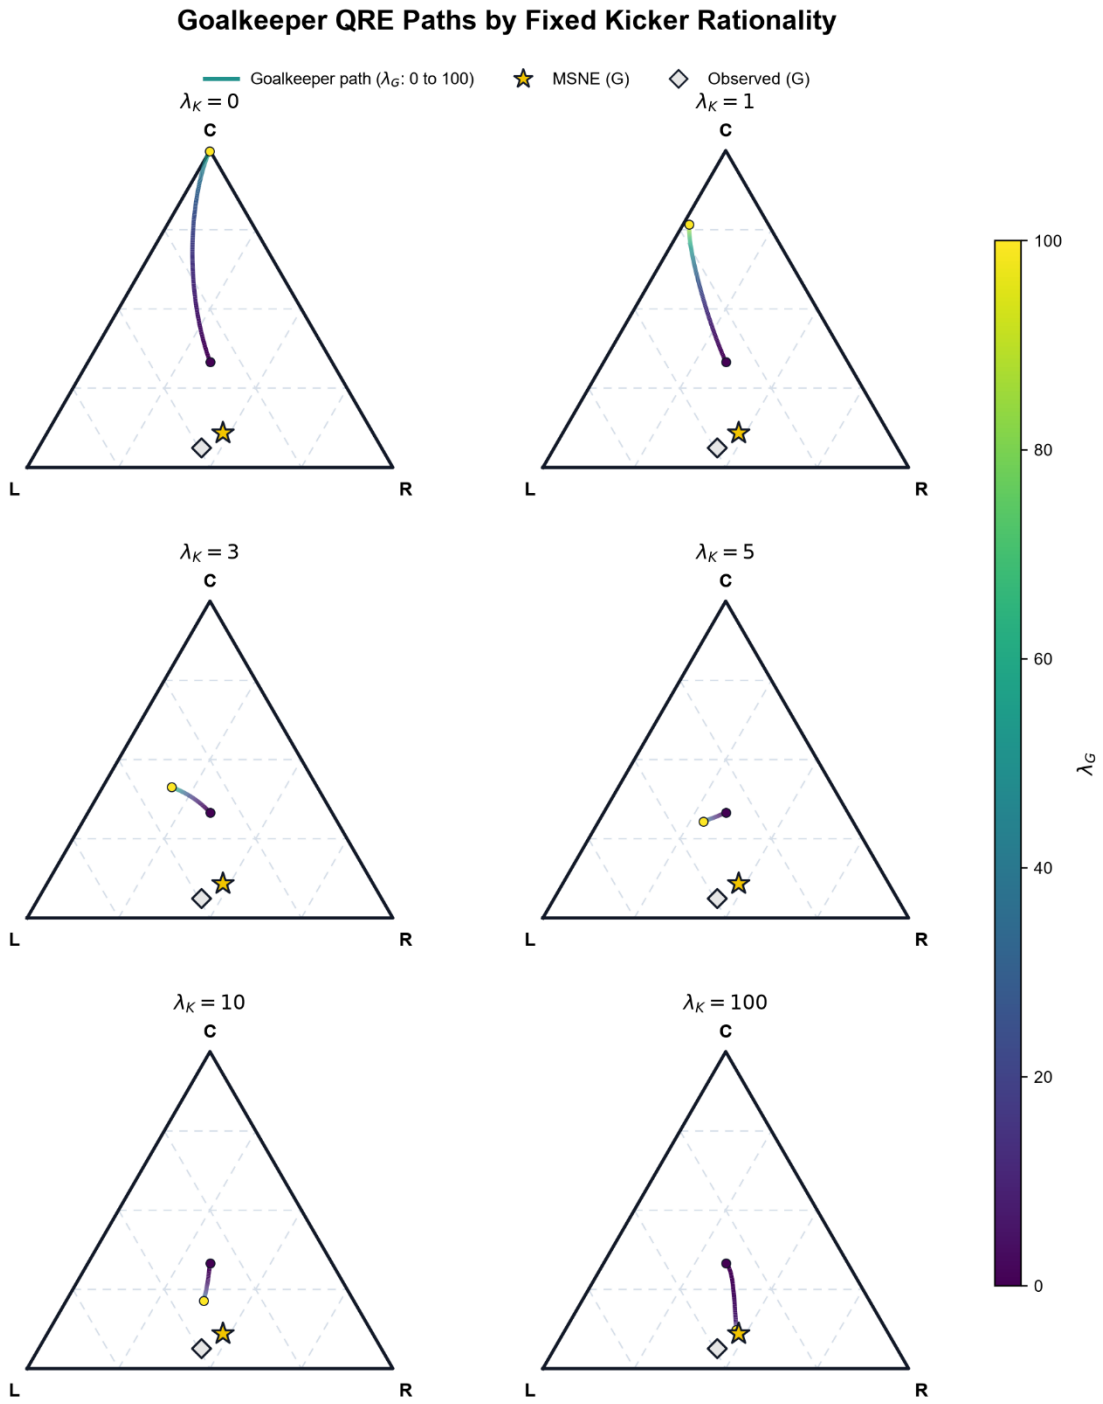


Figure 17 – Base game goalie equilibrium path against fixed kicker

We see that when $\lambda_K = 0$ the goalie increases their rate of centre play as they become more rational. This is because playing centre is the goalie's best pure strategy, we see this from

Figure 3 where the barycentre of the goalie best response simplex lies in the green region. Therefore, as goalie increasingly considers their utility differences, they put more probability on centre. However, this can only happen when the kicker is not responding to the goalie's equilibrium strategy. When λ_K is increased we see the goalie's equilibrium path move towards the MSNE.

By taking every value of λ_k between 0 and 100 and tracing out the goalie's response, we can trace the full range of the goalkeeper's mixed strategies under QRE

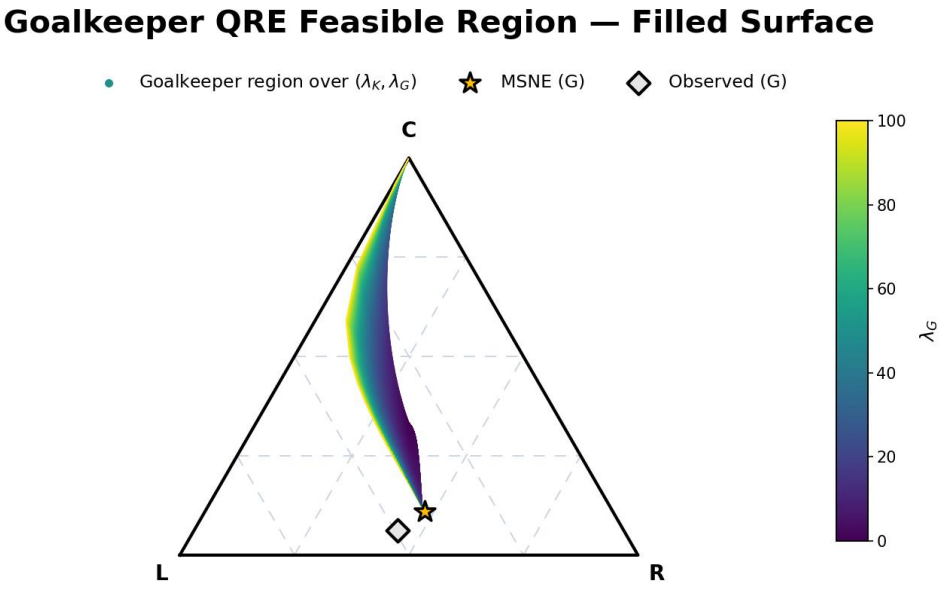


Figure 18 – Base game goalkeeper feasibility region

The region above shows all the possible goalie equilibrium mixtures for all possible pairs of rationality parameters, meaning that the goalie will never pick a probability mix that lies outside the coloured area. Forming this region of feasibility allows us to assess the game as a whole rather than looking at each individual equilibrium.

Repeating for the kicker, the following feasibility area is generated.

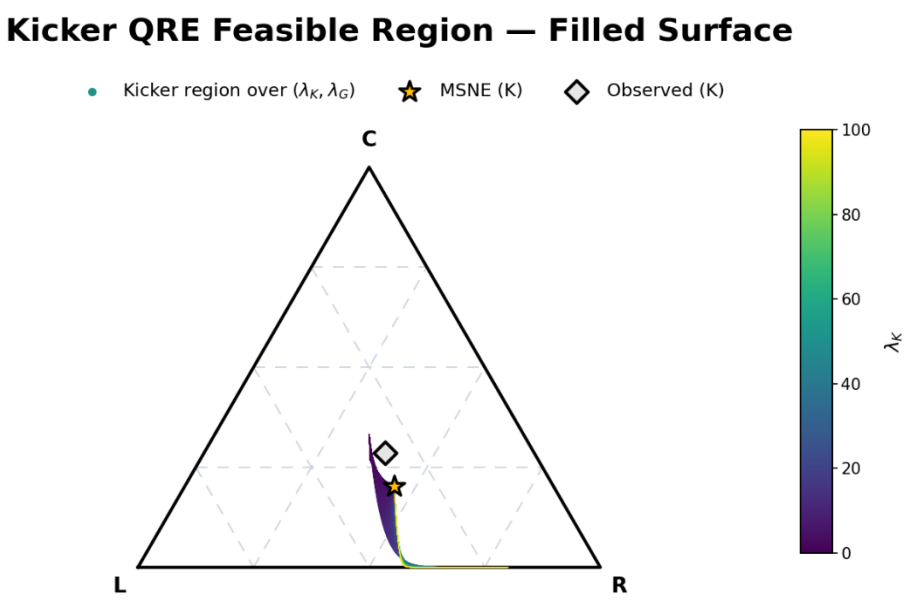


Figure 19 – Base game kicker feasibility region

Figure 18 and Figure 19 then show all possible equilibrium mixes for the no action bias game under the logit QRE equilibrium structure. We see on Figure 18 the point in the feasible region closest to the goalkeeper's observed behaviour is the MSNE prediction. Any reduction in either player's rationality parameter therefore moves the goalkeeper's predicted strategy further from the data. Thus, unlike in the asymmetric matching pennies example, QRE does not generate a beneficial comparative static adjustment in the goalkeeper's behaviour. As under MSNE, it can improve the fit of the kicker's strategy, but this improvement is always outweighed by the corresponding deterioration in the goalkeeper's fit.

Consequently, as Figure 20 below shows, the log likelihood is strictly increasing in each rationality parameter. Maximum Likelihood Estimation therefore places both parameters at the upper bound of the search range, where the QRE predictions converge towards the standard MSNE solution. Therefore, for this game Logit best response alone cannot improve the model's overall fit relative to MSNE.

A likelihood heatmap showing each possible λ combination confirms this.

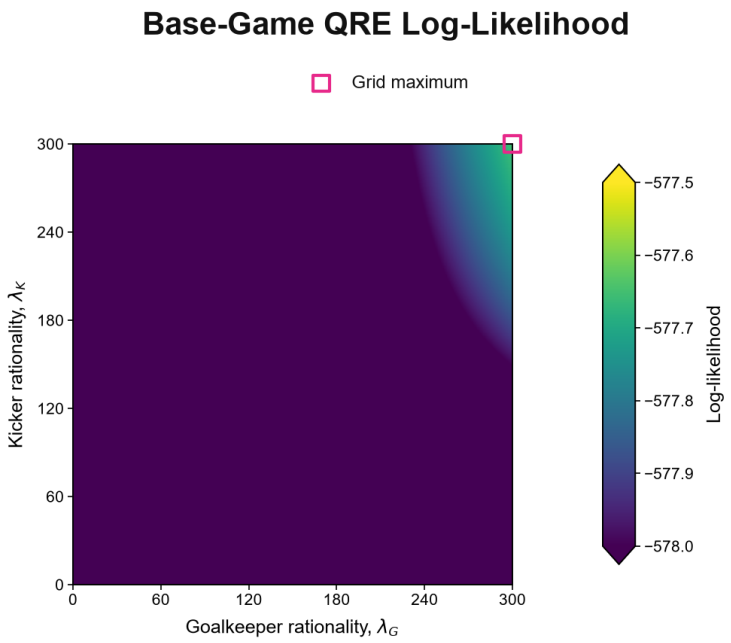


Figure 20 – Base game QRE likelihood heatmap

The heatmap extends the range of searched λs up to 300 and we see that the model still fits best when players act as rationally as the maximum λ allows. The constant colour gradient shows that this ceiling effect is a constant trend rather than an artifact of numerical noise.

Looking at the region in Figure 18, we can see that manipulating the rationality parameters for both or either player will not bring the quantal response equilibrium closer to the real world observed data. Manipulating the rationality parameters can only change where the specific equilibrium falls within the shaded area, and because the closest point in the green area to the observed data is the MSNE, we can conclude that for the base penalty kick game, the initial Mixed Strategy Nash Equilibrium calculated by Azar and Bar-Eli (Azar & Bar-Eli, 2011) is the closest that we can get to the real world data without the inclusion of behavioural biases.

This research examines whether two extensions to the standard penalty kick model can better explain the observed behaviour of kickers and goalkeepers. Firstly, was the inclusion of action bias to the MSNE model discussed earlier and secondly by using a Quantal Response equilibrium model instead of a Mixed Strategy Nash Equilibrium model. Either of these methods in isolation do not add substantial explanatory power. Whilst adding action bias to the MSNE model did improve fit at the 5% significance level, the fit improvement was only though the adjustment of the kicker's mix. The goalie's mix only improved under heterogeneous action bias, which was shown to not provide significant improvement over the homogeneous model. The base QRE model does not provide any improvement in fit as the MLE search always hits the upper λ search limit.

These findings contrast with the broad pattern in the experimental literature. (Ochs, 1995) finds that QRE outperforms MSNE in several but not all variants of matching pennies. Even where QRE does not improve fit, it generally produces predictions that differ from MSNE. Other studies show that QRE can fit observed behaviour more closely when enriched with additional behavioural assumptions. (Goeree et al., 2003), for example, incorporate risk-averse preferences, while (Leng et al., 2023) allow for altruism and spite. (McKelvey et al., 2000) similarly demonstrate the importance of allowing response precision to vary across players. Although these studies examine games that differ from the penalty kick game, they suggest that QRE often provides a flexible framework through which payoff and behavioural heterogeneity can generate economically intuitive changes in predicted play.

The present results are therefore striking. Introducing action bias into the MSNE model changes the goalkeeper's marginal strategy only when the equilibrium moves to reduced support. This occurs because the kicker's indifference conditions pin down the goalkeeper's mixing probabilities while all three actions remain in support. Once the centre is removed from the goalkeeper's support, the marginal probabilities assigned to the remaining actions can move closer to their observed values, but the model assigns zero probability to outcomes such as (Centre, Centre). The improvement in the marginal fit therefore comes at the cost of a substantial distortion in the predicted joint distribution.

QRE relaxes the indifference conditions and permits equilibrium probabilities to adjust smoothly while retaining full support. Nevertheless, QRE alone offers no improvement over the baseline MSNE model in this application. Any potential improvement in the fit of the kicker's behaviour is more than offset by a deterioration in the fit of the goalkeeper's behaviour, so MLE drives both rationality parameters towards the MSNE limit.

This exercise gives the surprising result that Logit QRE provides no improvement over the base MSNE. However, from the discussion of previous literature it has been established that the QRE solution concept provides a better foundation for payoff changes to intuitively move the equilibrium in a desired direction. The next question is therefore whether QRE provides a better foundation than MSNE for action bias to influence behaviour while preserving full support and generating smoother equilibrium adjustments. The next section examines this possibility.

(6.2) QRE with action bias

Action bias shifts the goalkeeper's payoffs and, consequently, the players' equilibrium logit responses. Whereas MSNE yields a unique equilibrium for each level of action bias, QRE may have multiple equilibria. Figure 21 shows how the resulting set of feasible predictions changes with the degree of action bias for all λ combinations from 0-100.

As action bias increases, we see the goalkeeper's probability mix move away from playing centre. The larger the action bias, the stronger and quicker the move away from centre is. When action bias is introduced into the model, the expected utility of playing centre reduces for the goalie. As the goalie becomes more rational, they begin to consider utility differences more than at lower levels of rationality. This is why we see the goalie's strategy move away from uniform randomisation as rationality increases. When action bias is low, such as 0.1, the incentive to move away from playing centre isn't very strong. So, the best response curves, and therefore the area traced, do not move sharply away from playing centre. As action bias increases, the difference in utility between diving (playing left or right) and not diving (staying centre) increases. This means that playing centre is a much more unattractive option for the goalie. This leads to the goalie reducing the amount they play centre even at lower levels of rationality. If the area of feasibility overlaps with the observed data we know that the selected level of action bias can be used in the model to match the predicted strategy to the real world data.

This differs from the MSNE model where the goalie's equilibrium mix only changes when moving to a reduced support equilibrium. Once logit best responding is introduced through bounded rationality, we see that the goalie is incentivised to reduce their rate of centre play even at levels of action bias below the critical value. This provides much more flexibility to the model and leads to more intuitive comparative statics changes as ε is varied.

Observing the goalie's feasibility areas as ε varies gives:

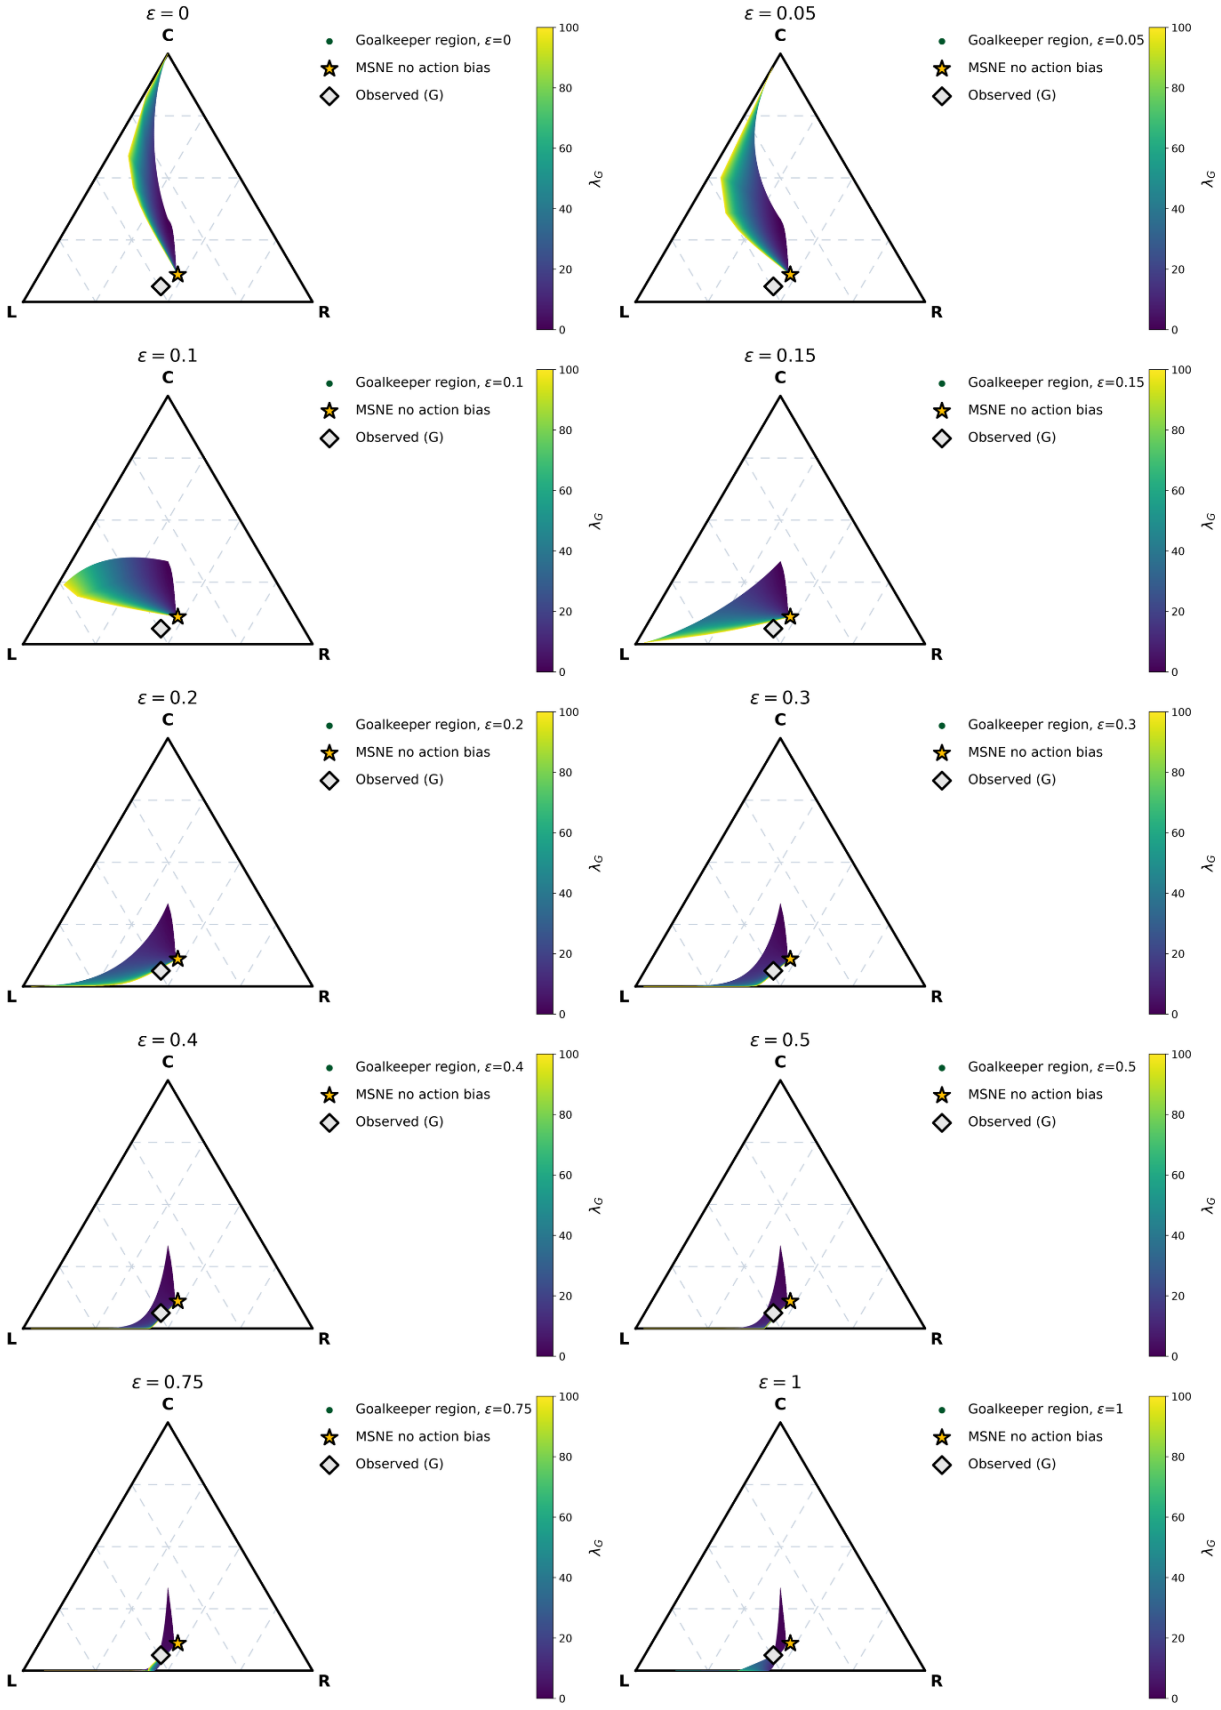


Figure 21- Goalkeeper feasibility region as ε varies

(6.3) Parameter estimation

Running an MLE over λ_K, λ_G, and ε, we take the fitted parameters and present them in Table 33 below. Table 34 presents the predicted strategy mixes at the fitted equilibrium, and Figure 22 visualises the equilibrium on a simplex.

Estimation results:

ε	λ_K	λ_G	Log likelihood
0.971603	0.000	2.905	-566.803176

Table 33 –Fitted QRE parameters

The predicted mixes at these parameter values are:

	Left	Centre	Right
Kicker	0.3333	0.3333	0.3333
Goalie	0.4930	0.0629	0.4441

Table 34 – Predicted QRE marginal mixes

Plotting gives:

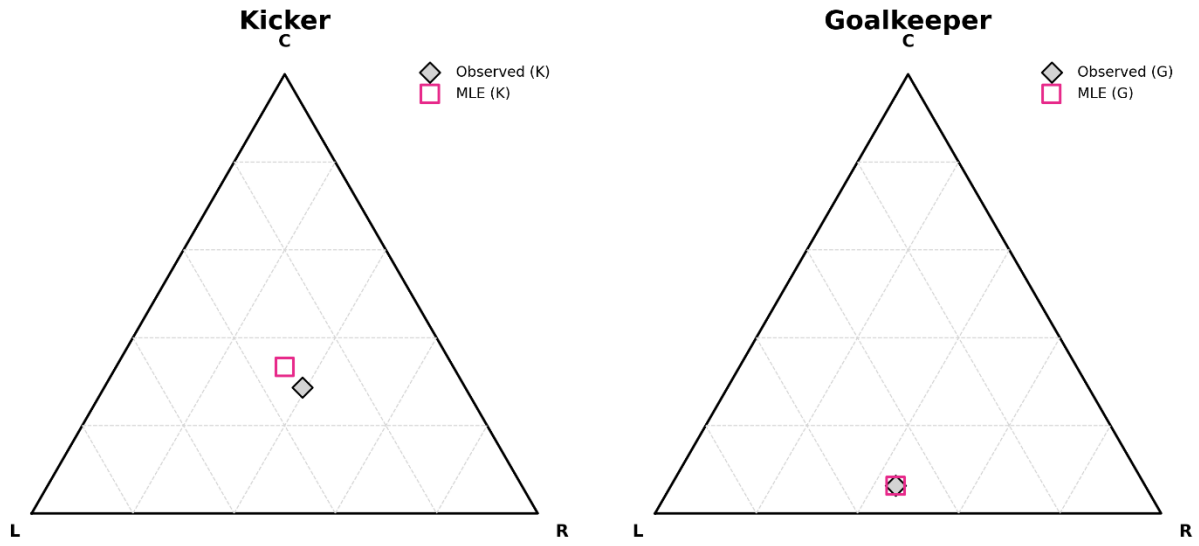


Figure 22 – Fitted QRE plot

The action bias estimate is $\varepsilon = 0.971603$, from the behavioural interpretation established earlier, this implies the goalie behaves as if conceding whilst staying in the centre has approximately twice the baseline utility cost of conceding after diving. Under MSNE, a bias of this magnitude produces the reduced support equilibrium discussed above. Under QRE, however, play retains full support and continues to respond to changes in ε, even beyond the value at which the MSNE shifts to reduced support.

We see that the goalie's strategy is replicated near perfectly whilst the λ_K= 0 parameter means the kicker mixes evenly across all three directions.

These results may seem strange at first, especially the kicker not considering utility differences at all. To understand these results, we can analyse them from both a computational/modelling perspective as well as a behavioural perspective.

Beginning with the mechanics of this equilibrium, we observe the following area plots to understand the imputed λ values.

The plots show the set of equilibria as λ_K and λ_G vary, given ε is set to the MLE value of 0.9716. The colour scaling shows the opponent's λ at each point. This shows the player's equilibrium path relative to their opponent's rationality, allowing us to analyse why the MLE returns λ_K= 0

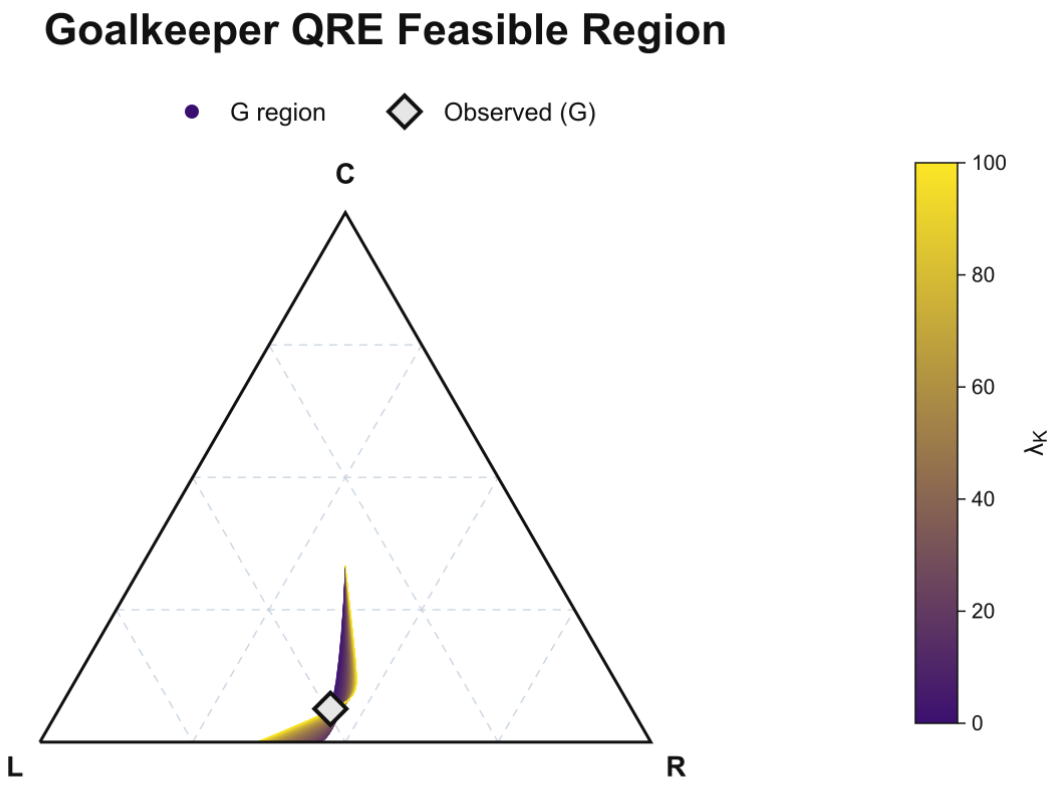


Figure 23 – Goalkeeper feasibility region at MLE ε

The yellow area shows how the goalie responds to a highly rational kicker, while the purple shows how the goalie will respond given the kicker is less rational. From this we observe that the goalie will always reduce centre play relative to complete randomisation, for any value of

$\lambda_K \geq 0$ and $\lambda_G > 0$. This means that the model can reproduce the goalkeeper's aggregate strategy relatively well regardless of the kicker's rationality.

However, the kicker may increase or decrease their rate of centre play as they increase rationality. The direction of this centre change depends on what the rationality parameter of the goalie is set to. For the kicker plot below, the colour scaling has been restricted to $0 \leq \lambda_G \leq 10$, this is done to make the phenomenon easier to observe.

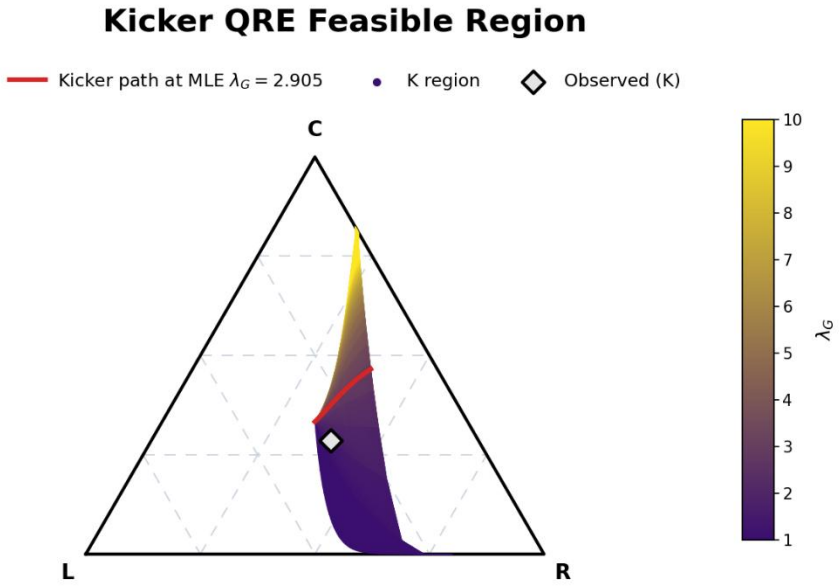


Figure 24 – Kicker feasibility region at MLE ε

When goalkeeper rationality is high, the kicker's probability of choosing Centre increases as λ_K increases. When goalkeeper rationality is low, the kicker's probability of choosing Centre decreases as $\lambda_K$increases. We see that at the λ_G MLE value, the kicker would increase their rate of centre play as λ_K increases, meaning any λ_K value above 0 would lead to a worse fit. Therefore, the MLE estimates $\lambda_K = 0$. Goalkeeper behaviour can be matched well for any λ_K, this is because all goalkeeper BR lines overlap with the observed data. Recall that a player's λ corresponds to how far along the BR line their mix lies on. A value of λ_G smaller than the MLE value would therefore produce a goalkeeper mix that lies too close to the centre of the simplex and therefore does not replicate the observed data. Therefore, the MLE cannot adjust λ_G to get a more favourable QRE curve for the kicker.

This MLE result also gives an interesting behavioural observation. Because the kicker's rationality parameter is set to zero, we know that they do not change their played strategy based on their expected utility differences. Following this, we can extrapolate that kickers do not adjust their strategy based on who their opponent is. This result reinforces the assumption

of “Identical goalkeepers” presented by (Chiappori et al., 2002). The logic presented in that paper is “If goalies are indeed homogeneous, then a given kicker's strategy will be independent of the goalie he is facing”. The results of this Maximum Likelihood Estimation reinforce this identical goalie assumption through the λ_k = 0 parameter result. However, this result shows that kickers treat goalies as identical and does not necessarily confirm that goalies themselves are identical.

Running an LR test between the no action bias model and the action bias model returns:

$$LR = 2[-566.803176 - (-576.249682)] = 18.93012, \qquad p = 0.0000138$$

The LR result shows that action bias within the QRE model provides significant explanatory power at the 1% level. This QRE estimation gives evidence that goalies face action bias that implies goalies treat staying in the centre and conceding as approximately twice as bad as diving and conceding. The $\lambda_K = 0$ result implies kickers have the same equilibrium mix regardless of the goalie they are facing. This is the same implication as the ‘identical goalkeepers’ assumption presented by (Chiappori et al., 2002). Before concluding on these results, we first test to see if behavioural heterogeneity in the form of a heterogeneous ε parameter provides any additional explanatory power.

(6.4) Heterogeneity in ε

(6.4.1) Complete information QRE

To see if QRE can identify the correlations within the data better than the MSNE framework, goalkeeper heterogeneity has been included, and the model has been estimated under the complete information setup as was done for the MSNE model. However, because the QRE framework always returns a strictly mixed equilibrium strategy there is no need to manually set ε_{high} like was done for the heterogeneous MSNE. Therefore, the MLE searches over: $\lambda_k, \lambda_{G,} \varepsilon_{low}, \varepsilon_{high}, and\ \pi$

λ_k	λ_G	ε_{low}	ε_{high}	π	Log Likelihood
0.000	2.905	Not identified	0.9716	1.000	-566.803176

Table 35 – Heterogeneous QRE fitted parameters

We see that the Heterogeneous QRE model collapses to the homogeneous QRE result. Because λ_k still returns zero, any variation in the goalie is ignored. This leads to the HQRE model only able to fit the goalie marginal, and many combinations of ε_{low}, ε_{high}, and π can reproduce the same pooled goalkeeper mix. This is confirmed in section (6.6)

The heterogeneous QRE model collapses to the homogeneous result. Because $\lambda_K = 0$, the kicker chooses each action with equal probability and does not respond to variation in the goalkeeper's predicted strategy. The estimated mixture also places all weight on one goalkeeper type, leaving the other type's action bias parameter unidentified. Consequently, introducing the specified two type structure does not improve the maximized likelihood. This shows that the aggregate data do not separately identify this form of goalkeeper heterogeneity; it does not establish that individual goalkeepers are behaviourally identical, though it does reinforce the idea that goalies are treated as identical by the kicker. This result provides additional suggestive evidence in favour of the 'identical goalkeepers' assumption, though it must be emphasised that this result does not mean goalkeepers *are* identical but rather that they are seemingly treated as identical by their opponent.

(6.4.2) Incomplete information QRE

An MLE search under the incomplete information assumption returned the same marginal (and therefore joint) distribution of actions as the homogenous estimation. The parameters of the ε distribution were unidentified as multiple combinations of ε_{low}, ε_{high}, and π can reproduce the same pooled goalkeeper strategy. Thus, introducing incomplete information heterogeneity does not improve the model's fit or explanatory power. This is essentially the same result as the complete information case. The reasoning behind this will be shown later in section (6.6)

(6.5) Common λ restriction

To see if one common rationality parameter explains the data as well as allowing separate rationality parameters an MLE was run using one common λ for both players. This was estimated with homogeneous action bias as well as the heterogeneous action bias setups defined earlier.

Single λ Models	ε	ε_{low}	ε_{high}	π_{high}	λ	**Log likelihood**
Homogeneous action bias	3.871271	-	-	-	0.662255	−567.153318
Heterogeneous action bias, complete information	-	3.871263	3.871271	0.995904	0.662255	−567.153318

Table 36 – Common λ fitted parameters

We see a similar pattern as to when each player has their own λ parameter. Allowing ε to vary as two types does not provide any explanatory power to the model. The incomplete information case was also estimated and returns the same joint and marginal distribution as the homogenous model. The Heterogeneous action bias parameters are not uniquely defined here as ε_{low}, ε_{high}, and π generate the same pooled goalie strategy and therefore the same maximised likelihood.

The search selects this because with one common lambda and action bias affecting only the goalkeeper's centre payoff, the correlation pattern is highly constrained. It cannot independently adjust arbitrary cells in the observed (3x3) table. Moving the epsilons apart does not improve the joint likelihood; it damages the fitted marginal strategies at least as much as it helps any joint cell residuals.

The heterogeneous specification therefore collapses to the corresponding homogeneous model. This means that the proposed two type structure cannot be identified from the aggregate data. The reason for this interesting result will be discussed in section (6.6), for now we continue with the discussion on λ specification.

An LR test between common λ with homogeneous action bias and without action bias returns

$$LR_{\text{common}} = 2[-567.153318 - (-576.249682)] = 18.192728.\ \ p = 0.0000200$$

This LR test shows action bias improves fit in the common λ model to a statistically significant degree These results provide evidence consistent with goalkeeper action bias on the aggregate level. There is no evidence that the specified form of heterogeneity improves fit. Though individual heterogeneity may remain present, it cannot be distinguished from homogeneity using these model specifications and aggregate data.

This leads to the conclusion that within this dataset, the model estimations are consistent with goalies holding action bias, yet the results from the model estimations suggest that all goalies feel the same amount of disutility when staying in the centre and conceding.

Plotting the common and separate λ models with homogeneous action bias allows us to see how well each model replicates the observed strategy mixes.

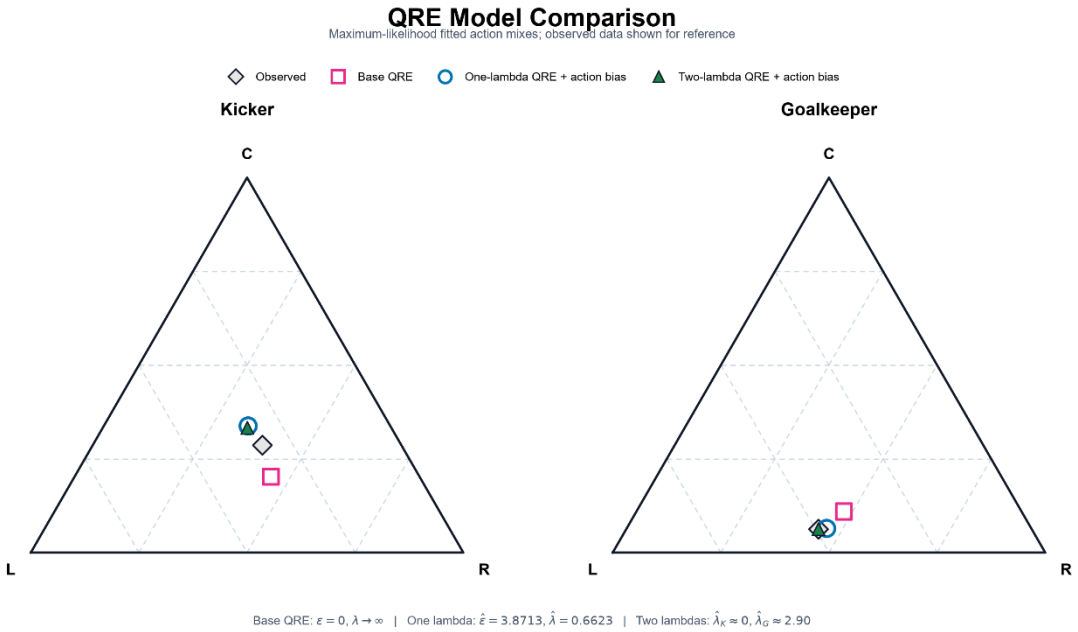


Figure 25 – QRE fitted model comparison

For Figure 25 only the homogeneous implementations of action bias have been plotted. The heterogeneous models return the same mixes as the respective homogeneous model, making plotting their mixes redundant.

We see that the two λ model reproduces mixes slightly closer than the one λ model. The 'Base QRE' mixes are the same for both models as both models push to the λ cap making them indistinguishable without action bias. Including action bias into the one and two λ models allow for a much closer replication of the observed mixes for both λ implementations.

From Figure 25 it can be seen that the two λ specification fits the observed marginals slightly better than the single λ model. This is more obvious on the goalkeeper plot than the kicker plot.

To see if the improvement in fit is worth the extra parameter when moving from a common λ to separate λs, we run the following parameter and likelihood comparisons:

Model	ε	λ_K	λ_G	**Log likelihood**
Common lambda	3.871271	0.662255	0.662255	−567.153318
Separate lambdas	0.971603	0.000000	2.904793	-566.803176

Table 37 – Common vs separate λ fitted parameters

An LR test between these two models returns $p = 0.403$. This shows that allowing two rationality parameters does not improve fit at a statistically significant level, so the common λ restriction cannot be rejected at conventional significance levels.

(6.6) Heterogeneity in ε redundancy

The result that allowing for heterogeneity within ε does not improve model fit or explanatory power is surprising. It's natural to assume that different goalkeepers would have varying levels of mental bias, yet allowing for this assumption is shown to not improve fit to the correlations within the joint distribution. The reason behind this result is actually quite simple and visualised on Figure 26.

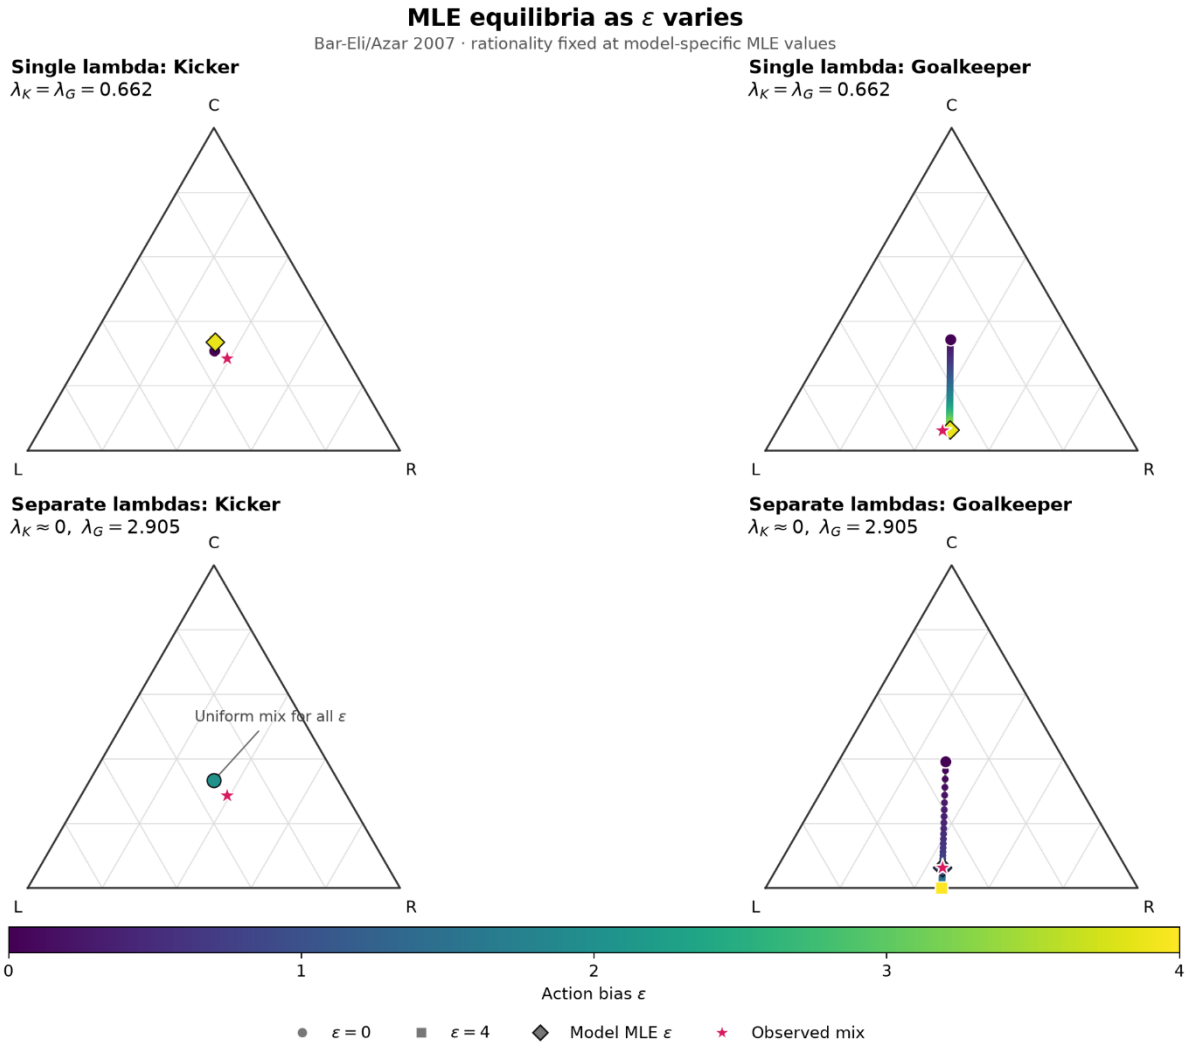


Figure 26 – MLE equilibria as ε varies

Figure 26 shows the equilibria when λ is fixed to the MLE value and ε varies from 0 to 4. We see that they are essentially linear. The separate λ goalie path is exactly linear, whilst the common λ path curves slightly to the right (this bending is so small it is undetectable visually). This means that the pooled mixes from any distribution of ε can be replicated by a single homogeneous action bias parameter. This is why the MLE search returns undefined parameter estimations, as many combinations can replicate the same pooled mix. This makes heterogeneity within ε redundant describing the marginal distributions. This finding is caused

by the payoff specifications used in this analysis and should not be treated as general QRE behaviour. Other forms of heterogeneity may lead to meaningful identification of the marginal and joint distribution.

(7) Azar et al. model comparison

We now have multiple models established. This section aims to summarise key findings and conclude between model and parameter specification. Doing so will allow us to determine what behavioural model best explains aggregate penalty kick behaviour for the (Azar & Bar-Eli, 2011) dataset.

(7.1) Action bias specification

Within the MSNE setup we have explored the following models.

- The base MSNE model presented by (Azar & Bar-Eli, 2011)
- The MSNE model with homogeneous action bias
- The MSNE with two type heterogeneous action bias (under both complete information and incomplete information)

The MLE values for each of these are presented below.

Model	**Main MLE parameters**	**Log likelihood**
Base MSNE	No free behavioural parameter	−576.249682
Homogeneous MSNE + AB	$\varepsilon = 0.064059$	−570.942578
Two type heterogeneous MSNE + AB Complete information assumption	$\varepsilon_{low} = 0.028702$ $\varepsilon_{high} \geq 0.760788$ $\pi_{high} = 0.077042$	−570.107401

Table 38 – MSNE model comparison

The incomplete information case is not presented in Table 38 or discussed further as it has already been shown to be redundant.

Seeing as there are multiple specifications of action bias we compare fit to the observed data versus the trade off in parameter complexity.

Model	**Number of behavioural parameters**	**Log likelihood**	**AIC**	**BIC**	**Interpretation**
Base MSNE	0	−576.249682	1152.499363	1152.499363	Lowest complexity, worst fit

Homogeneous MSNE + AB	1	−570.942578	1143.885156	1147.541148	Best balance of fit and simplicity
Two type heterogeneous MSNE + AB	2	−570.107401	1144.214803	1151.526786	Better fit, but heavier complexity penalty

Table 39 – AIC and BIC testing for MSNE models

The homogeneous model is preferred by both AIC and BIC. Although the two type model has the highest likelihood, its improvement over the homogeneous model is not significant at the 5% level in their direct bootstrap comparison: $LR = 1.670, p = 0.098$; although the result provides weak evidence in favour of the two type model at the 10% level. AIC and BIC testing prefer the more parsimonious homogeneous action bias model. Therefore, we take the homogeneous model as the preferred MSNE specification.

This tells us that within the MSNE framework, there is evidence of action bias, and there is no evidence of heterogeneity within the parameters tested.

When discussing the QRE framework, we have the same specifications of action bias.

- QRE with no action bias
- QRE with homogeneous action bias
- QRE with heterogeneous action bias, under both complete information and incomplete information

Each of these models can also be implemented with one rationality parameter (λ) for both players, or two parameters (λ_G, λ_K), one for each player. The full table comparing all models can be found in the appendix. For now we only present the models that return results useful for interpretation.

Model	**Lambda specification**	**MLE parameters**	**Log likelihood**	**Interpretation**
Base QRE	Both common/ separate lambda	$\lambda \to \infty$ $\varepsilon = 0$	−576.249682	Collapses to the base MSNE.
QRE with homogeneous action bias	Common lambda	$\lambda = 0.662254$ $\varepsilon = 3.871276$	−567.153318	Model preferred under parsimony

QRE with homogeneous action bias	Separate lambdas	$\lambda_K \approx 0$ $\lambda_G = 2.904793$ $\varepsilon = 0.971603$	−566.803176	Best Likelihood fit
QRE with heterogeneous action bias	Both common/ separate lambda	Find in appendix	Find in appendix	Effectively collapses to the respective homogeneous model.

Table 40 – QRE fitted model comparison

The main result is that complete information heterogeneity does not improve the QRE fit. In both the common lambda and separate lambda specifications, the heterogeneous model collapses back to the corresponding homogeneous action bias model. This shows that under both the MSNE and QRE model specifications, there is no evidence for behavioural heterogeneity in action bias, this interesting result and its implications will be discussed further in Section (9.2).

(7.1.1) Conclusion on ε specification

We can now conclude that neither complete nor incomplete information heterogeneity improves model fit within the dataset. This applies to the MSNE and the QRE framework. Therefore, within the dataset analysed, homogeneous action bias is the preferred ε specification.

(7.2) Lambda specification

Comparing the common vs separate lambda models, we see that allowing for separate lambdas improves likelihood fit when action bias is included. But both models have the same likelihood in the no action bias model due to both models estimating $\lambda \to \infty$.

Once action bias is included in the model the separate λ model imputes a slight likelihood improvement over the common λ model, though this improvement was not shown to be significant enough to reject the common λ restriction.

To further compare the common versus separate lambda models we run the following BIC test:

Model	**Number of free parameters**	**Log likelihood**	**BIC**
Common lambda QRE + AB	2	−567.153318	**1145.618620**

Separate lambda QRE + AB	3	−566.803176	1150.574328

Table 41 – Common vs separate λ BIC

BIC testing prefers the common lambda restriction, reinforcing the result from earlier that the common lambda restriction cannot be rejected under the LR testing done in section (6.5). Therefore, we conclude that the single λ QRE model with homogeneous action bias is the preferred QRE specified model.

(7.3) Comparison between preferred MSNE and preferred QRE

We now have the Homogeneous Action bias as our preferred Nash Equilibrium model and the Common Lambda Homogeneous Action Bias QRE model as our preferred behavioural QRE model. We then run statistical tests on these two models to conclude on fit improvement versus parameter complexity

Model	**AIC**	**BIC**
Homogeneous MSNE + AB	1143.885156	1147.541148
Common lambda QRE + AB	**1138.306636**	**1145.618620**

Table 42 – Preferred QRE vs MSNE

We see that the AIC and BIC testing both prefer the QRE model.

A standard LR test should not be done to compare these two models as the MSNE model corresponds to the QRE limit $\lambda \to \infty$. This places the null on a boundary, violating the regularity conditions behind the ordinary chi-square approximation. Instead, a bootstrap test is used to conclude on statistical fit between these two models.

Quantity	**Result**
MSNE + AB log likelihood	−570.942578
Common lambda QRE + AB log likelihood	−567.153318
Observed LR statistic	7.578520
Bootstrap replications	499
LR values ≥ observed LR	2
Bootstrap p-value	**0.006**

Table 43 – QRE vs MSNE bootstrap test

The parametric bootstrap likelihood ratio test rejected the homogeneous MSNE with action bias in favour of the homogeneous common lambda QRE with action bias at the 1% significance level. We can therefore take the single λ homogeneous action bias model as the overall preferred model for this dataset.

We have now estimated and fit tested all relevant models against each other. Figure 27 below presents the conclusions of the model comparisons.

(7.4) Overall model comparison and narrative

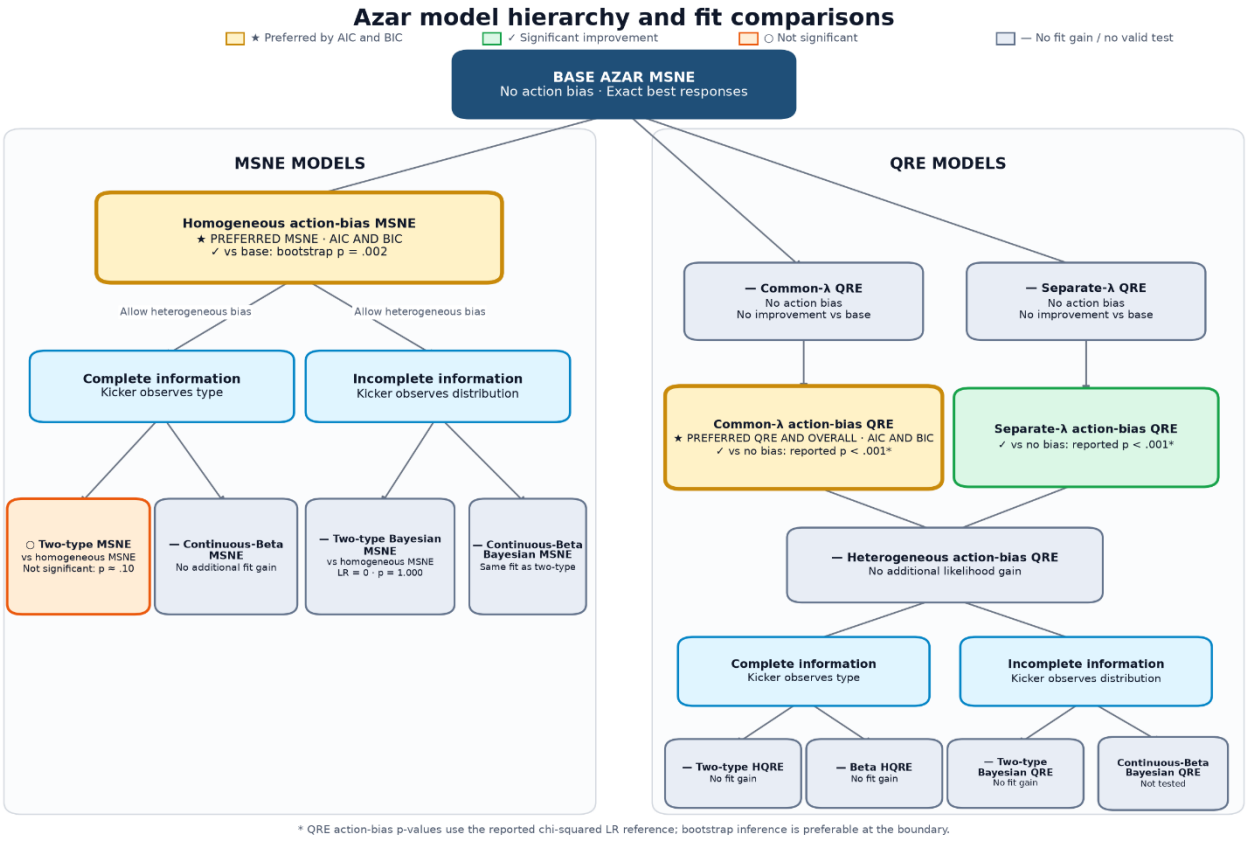


Figure 27 – Model comparison flowchart

From reading Figure 27 it becomes easy to follow the narrative of this research. Amongst the MSNE models, the homogeneous action bias implementation provides the best improvement in fit relative to parameter complexity. Adding heterogeneity into action bias does not provide significant explanatory power for any of the distribution or information types. Allowing logit best responding through QRE does not improve fit at all over the base MSNE as the models estimate the maximal amount of rationality allowed by the search parameters. Including action bias improves fit to a statistically significant degree for both the one and two λ models. Allowing

for heterogeneity in ε does not provide any improvement in fit over homogeneous action bias, similar to the MSNE models. The common λ model is overall preferred by AIC and BIC.

This allows us to make the following generalised conclusions on whether action bias provides explanatory power when modelling the Azar et al. dataset. Inclusion of action bias improves model fit to a statistically significant degree for both the MSNE and QRE model specifications. Heterogeneity within ε does not provide statistically significant improvement over a homogeneous action bias parameter. QRE with action bias improves fit over the MSNE with action bias to a statistically significant degree.

(8) Singh and Mukherjee comparison

(8.1) Model estimations

However, these conclusions are only made for the Azar et al. dataset, the research conducted in the replication study conducted by (Singh & Mukherjee, 2026) concludes that action bias is not present in their dataset.

The observed directions and stopped kicks in this dataset are:

Observed directions	Goalkeeper			
Kicker	Left	Centre	Right	Total
Left	49 (16.8%)	15 (5.1%)	35 (12.0%)	99 (33.9%)
Centre	23 (7.9%)	7 (2.4%)	28 (9.6%)	58 (19.9%)
Right	44 (15.1%)	11 (3.8%)	80 (27.4%)	135 (46.2%)
Total	116 (39.7%)	33 (11.3%)	143 (49.0%)	292 (100.0%)

Table 44 - (Singh & Mukherjee, 2026) Joint frequency table

Stopped kicks	Goalkeeper			
Kicker	Left	Centre	Right	Total
Left	(22/49) 44.9%	(0/15) 0.0%	(0/35) 0.0%	(22/99) 22.2%
Centre	(6/23) 26.1%	(5/7) 71.4%	(5/28) 17.9%	(16/58) 27.6%
Right	(0/44) 0.0%	(0/11) 0.0%	(33/80) 41.2%	(33/135) 24.4%
Total	(28/116) 24.1%	(5/33) 15.2%	(38/143) 26.6%	(71/292) 24.3%

Table 45 - (Singh & Mukherjee, 2026) Stopped kicks

We see for this dataset, staying in the centre leads to the lowest marginal stop probability for the goalie as opposed to the Azar et al. dataset where staying in the centre had the highest marginal stop probability.

Singh and Mukherjee compare the observed stopping rate when the goalkeeper stays in the centre (5/33 = 15.2%) with the stopping rate when the goalkeeper dives (66/259 = 25.5%). A Yates-corrected chi-squared test does not reject equality between these rates: $(\chi^2(1) = 1.1826,\ p = 0.2768)$.

This result leads the researchers to not reject equal scoring probabilities between staying in the centre and diving. This test is what the researchers use to conclude that their dataset gives evidence of no goalkeeper action bias. This fails to tell the full story however, as their data

does not indicate MSNE play. Figure 28 below shows the best response simplexes for the dataset collected by Singh and Mukherjee.

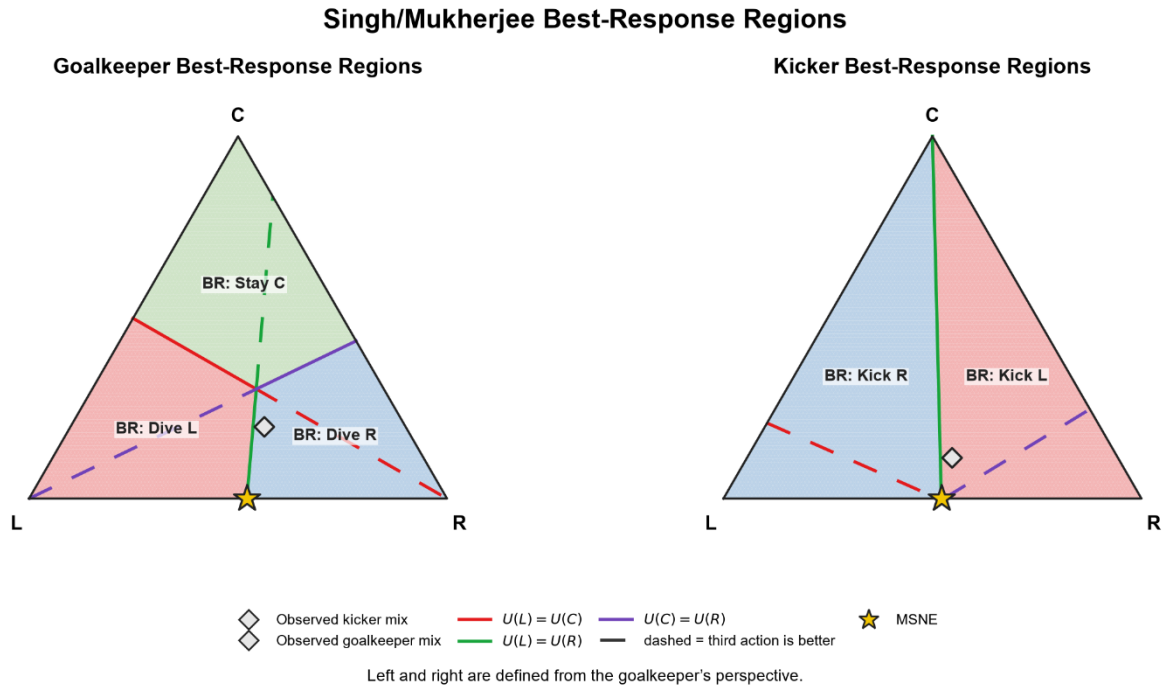


Figure 28 - (Singh & Mukherjee, 2026) Best response regions

We see that the kicker has no BR: Kick C region within the simplex. This is because kicking to the centre is never a best response. Even when the goalie only mixes over L and R, the kicker's payoff is strictly worse when kicking to the centre compared to the sides. We can show this algebraically.

Consider the situation where the goalie only mixes over left and right

	L	C	R
Goalie	p	0	1-p

The kicker's utility equations become

$$U_k(\text{Left}) = 0.551p + (1-p)$$

$$U_k(\text{Centre}) = 0.739p + 0.821(1-p)$$

$$U_k(\text{Right}) = p + .588(1-p)$$

Equating $U_k(\text{Right}) = U_k(\text{Left})$

$$0.551p + (1-p) = p + .588(1-p)$$

$$p = .4785$$

Therefore, the goalie's equilibrium mix is: $G = (0.478513, 0, 0.521487)$.

The kicker's equilibrium probabilities then become: $K = (0.478513, 0, 0.521487)$

Both players then have a reduced support mix where neither includes centre. Given the dataset has centre observations for both the kicker and the goalkeeper we can conclude that the aggregate observations are incompatible with treating the penalty kicks as a homogeneous MSNE game.

The inclusion of action bias will not move the equilibrium to full support. Action bias changes the goalkeeper's payoffs but leaves the kicker's best response regions unchanged. Because kicking to the centre is never a best response in the Singh and Mukherjee game, no value of the action bias parameter can generate a full support MSNE.

This reduced support result motivates the use of QRE modelling instead of MSNE.

Figure 29 below presents the fitted parameter models for the MSNE, common λ QRE, separate λ QRE, both with action bias and without action bias.

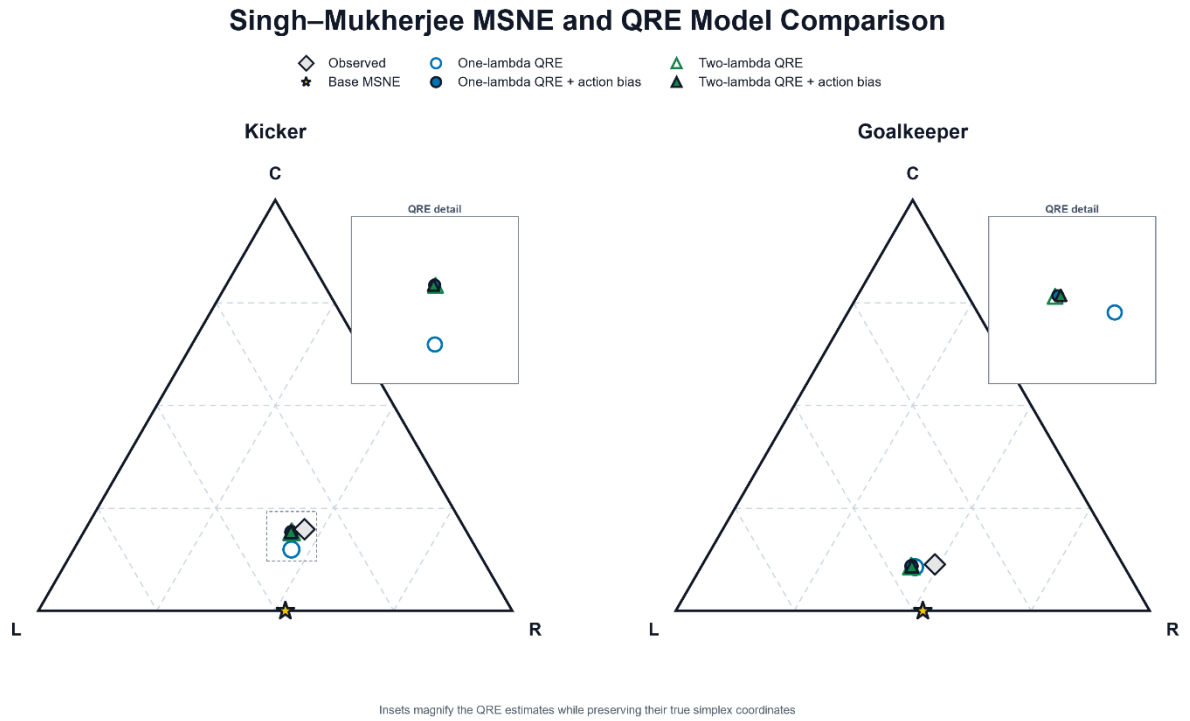


Figure 29 - (Singh & Mukherjee, 2026) QRE model comparison

We see that each of the four QRE models estimated result in almost the same predicted mixes. Even when zoomed in the points still overlap significantly. This shows that all fit improvement over the MSNE is being driven by the removal of perfect rationality and that the specific form of QRE model used is negligible when discussing explanatory power. This will be discussed in more detail in the following section.

(8.2) Conflicting results between the Azar et al. and Singh and Mukherjee estimations.

Before discussing the behavioural differences implied by both datasets, we first compare the model fit rankings.

MLEs were run for all MSNE and QRE models and the BIC value for each is reported in Table 46 below. To avoid a zero probability error, the (Singh & Mukherjee, 2026) MSNE results are estimated using a QRE model with λ fixed to 3000, this mimics the sharp best responding of the MSNE model whilst placing positive probability on all possible actions. However, the parameter counts will be counted as if the MSNE was used when running the BIC calculations. This allows the BIC values to be directly compared against the equivalent Azar et al. models.

Therefore, the Singh & Mukherjee MSNE BIC results describe a numerical approximation to MSNE rather than the exact MSNE values as this would return $-\infty$.

Model type	**Bar-Eli/Azar 2007**	**Singh/Mukherjee 2026**
Base MSNE	1152.499363	44405.975189
Homogeneous MSNE + AB	1147.541148	44411.651943
Two type heterogeneous MSNE + AB	1151.526786	44417.328697
Common lambda QRE no AB	1158.155355	**1186.828275**
Common lambda QRE + AB	**1145.618620**	1188.005202
Separate lambda QRE no AB	1163.811347	1188.007483
Separate lambda QRE + AB	1150.574328	1193.681170
HQRE two type action bias	1161.886312	1205.028862

Table 46 – BIC results (Azar & Bar-Eli, 2011) and (Singh & Mukherjee, 2026) datasets

The BIC preferred model for each dataset is reported below.

Dataset	**Preferred model by BIC**	**Family**	**Log likelihood**	**BIC**
(Azar & Bar-Eli, 2011)	Common lambda QRE + action bias	QRE	-567.153318	1145.618620
(Singh & Mukherjee, 2026)	Common lambda QRE, no action bias	QRE	-590.575761	1186.828275

Table 47 – Preferred BIC model (Azar & Bar-Eli, 2011) and (Singh & Mukherjee, 2026)

We see that each dataset has a different preferred model. The QRE family of models is preferred by both datasets, and action bias is only preferred in the Azar et al. dataset. Fit improvement in the Singh and Mukherjee dataset is driven by the allowance of logit best responding, as λ and action bias specification lead to negligible changes to the fitted mixes when varied. This is in direct contrast to the Azar. et al. estimations where bounded rationality alone did not provide any explanatory power over the base MSNE model. The Singh & Mukherjee results therefore support QRE's full support structure but does not provide evidence of goalkeeper action bias.

(9) Behavioural discussion

We can now compare the two datasets and make behavioural inferences as to the structure of the penalty kick game.

(9.1) Perfect versus bounded rationality

Firstly, allowing for bounded rationality improves explanatory power for both datasets analysed. The (Singh & Mukherjee, 2026) scoring probabilities imply that perfectly rational players will exclude centre from their chosen strategy mixes. Allowing for logit best responding allows model estimation to fit the observed data much better.

In the Azar and Bar-Eli data, QRE without action bias converges to the base MSNE and provides no improvement in fit. QRE becomes empirically valuable only when it is combined with action bias. The single λ QRE model with action bias significantly outperforms the homogeneous MSNE action bias model in the parametric-bootstrap test. For both datasets the removal of the indifference condition required by MSNE allow for statistically significant improvements in explanatory power.

This implies that players may be responding to their own payoffs as opposed to acting in a way that keeps their opponent from having a profitable deviation. This aligns with the (Goeree & Holt, 2001) research discussed in the literature review, where the unintuitive comparative statics of the MSNE model were not replicated in real world play. This shows why a model that allows players to respond to their own payoff changes may be beneficial when modelling real world behaviour.

(9.2) Goalkeeper action bias

The Azar et al. dataset shows evidence of goalkeeper action bias. Including action bias into the goalie's utility functions improves overall likelihood for both the QRE and MSNE models. The goalkeeper survey conducted by (Bar-Eli et al., 2007) show that professional goalkeepers report feeling worse when conceding the goal when staying in the centre compared to when diving and also conceding. We have now formally modelled this reported additional disutility and can confirm that the observed data collected in the 2007 study show evidence of players holding some form of action bias. Our model estimations imply that this extra disutility is felt uniformly across goalkeepers in the dataset and allowing for heterogeneity in the size of action bias felt does not provide any additional explanatory power.

However, the data collected by (Singh & Mukherjee, 2026) provides no evidence of action bias. Therefore, we can conclude that the existence of goalkeeper action bias is dataset dependent. A noteworthy distinction between the two datasets is that the Singh and Mukerjee dataset is taken from international World Cup games, whilst the Azar et al. dataset is only reported to be taken from "various leagues and championships worldwide" (Bar-Eli et al., 2007). Assuming these leagues are not at the top international standard we can infer that professional goalkeepers still show signs of action bias, yet the best of the best (World Cup) players do not seem to hold this same mental bias. However, more research using datasets taken from distinct leagues at varying levels must be conducted before concluding that this result is not simply a coincidence.

(9.3) Limitations of behavioural conclusions

As discussed by (Chiappori et al., 2002), there are limitations to the interpretations that can be made when treating the penalty kicks as a universal homogeneous game. Player heterogeneity may lead to the aggregate data being made up of multiple games with unique equilibria as opposed to one equilibrium that can adequately describe the entire dataset.

The analysis conducted in this paper was done using the aggregate data, however multiple forms of heterogeneity were tested for. Firstly, the Azar et al. frequencies and scoring rates were used to try and determine kicker footedness. We found that footedness could not be reliably estimated. Secondly, behavioural heterogeneity within action bias was tested and the results show that under both the MSNE and QRE solution concepts we do not find any evidence for behavioural heterogeneity within action bias.

Chiappori et al. establish their 'Identical Goalkeepers' assumption and the results of this research fail to identify heterogeneity within the goalkeeper parameters tested. The results are observationally consistent with treating goalkeepers as identical from the kicker's perspective, though this evidence can only be treated as suggestive. The MSNE calculated from the aggregate Singh and Mukherjee (2026) payoff matrix excludes centre, yet centre is observed in the aggregate data. This discrepancy may reflect unmodelled heterogeneity; different kicker/goalkeeper matchups may use different equilibrium supports, thereby expanding the range of actions observed after aggregation, rather than indicating genuine departures from equilibrium play.

The main form of heterogeneity presented by Chiappori et al. is kicker footedness. Because footedness may affect the pooled distribution of kick directions and therefore the goalkeeper's expected utility from diving or staying in the centre, its effect on the analysis of goalkeeper

action bias cannot be determined from the aggregate data. This remains a limitation of the present analysis.

All conclusions of this research must be viewed under the limitation of working with aggregate penalty kick data. However, much of the supplementary analysis conducted in this thesis was done to try and mitigate this limitation as much as possible.

(10) Conclusion

The overarching aim of this research has been to assess the existence of goalkeeper action bias by formally incorporating action bias into the goalie's payoffs. Overall, the results show evidence of action bias within the primary dataset analysed. The statistical testing showed that action bias provides significant explanatory power when using both the MSNE and QRE solution concepts. Action bias combined with bounded rationality produced the best fit to the observed data. These findings present evidence that professional footballers may not be acting as perfectly rational utility maximisers and instead could be responding to own payoff effects whilst also having their outcome utilities influenced by mental biases. This contributes to the game theory literature space by highlighting that the current benchmark of MSNE may not be fully capturing players' incentives and behaviour. The results show a trend where bounded rationality better explains player behaviour compared to assuming perfect rationality. The existence of action bias within goalkeepers is not shown to be universal, with suggestive evidence that the most elite players do not have this mental bias. Further research must be conducted before concluding on this behavioural trend. The aggregation issue discussed in this thesis emphasises the need for more penalty kick research to be conducted with match specific and player specific variables accounted for. This area is important for future research because player fixed effects and competition level must be accounted for before treating behavioural conclusions as generalisable. Overall, this research finds evidence of goalkeeper action bias and establishes a framework of analysis that can be used by future researchers to further describe real world competitive play.

(11) References

Almeida, C. H., & Volossovitch, A. (2023). Multifactorial analysis of football penalty kicks in the Portuguese First League: A replication study. *International Journal of Sports Science & Coaching, 1*(18), 160-175. https://doi.org/10.1177/17479541221075722

Almeida, C., Volossovitch, A., & Duarte, R. (2016). Penalty kick outcomes in UEFA club competitions (2010-2015): The roles of situational, individual and performance factors. *International Journal of Performance Analysis in Sport, 2*(16), 508-522. https://doi.org/10.1080/24748668.2016.11868905

Anderson, A., Rosen, J., Rust, J., & Wong, K.-P. (2025). Disequilibrium Play in Tennis. *Journal of Political Economy, 1*(133), 190-251. https://doi.org/10.1086/732529

Azar, O. H., & Bar-Eli, M. (2011). Do soccer players play the mixed-strategy Nash equilibrium? *Applied Economics, 25*(43), 3591-3601. https://doi.org/10.1080/00036841003670747

Bar-Eli, M., Azar, O. H., Ritov, I., Keidar-Levin, Y., & Schein, G. (2007). Action bias among elite soccer goalkeepers: The case of penalty kicks. *Journal of Economic Psychology, 28*(5), 606-621. https://doi.org/10.1016/j.joep.2006.12.001

Chiappori, P., Levitt, S., & Groseclose, T. (2002). Testing Mixed-Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in soccer. *American Economic Review, 4*(92), 1138-1151. https://doi.org/0.1257/00028280260344678

Crawford, V., & Smallwood, D. (1984). Comparative statics of mixed-strategy equilibria in noncooperative two-person games. *Theory and Decision, 3*(16), 225-232. https://doi.org/10.1007/BF00134647

Friedman, D., & Zhao, S. (2021). When are mixed equilibria relevant? *Journal of Economic Behavior and Organization, 191*, 51-65. https://doi.org/10.1016/j.jebo.2021.08.031

Goeree, J., & Holt, C. (2001). Ten Little Treasures of Game Theory and Ten Intuitive Contradictions. *American Economic Review, 91*(5), 1402-1422. https://doi.org/10.1257/aer.91.5.1402

Goeree, J., Holt, C., & Palfrey, T. (2003). Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior, 45*(1), 97-113. https://doi.org/10.1016/S0899-8256(03)00052-6

Hart, S., & Esteva, R. (2014, October). *Penalties: How to hit them, where to put them*. Retrieved from UEFA.com: https://www.uefa.com/uefachampionsleague/news/025a-0ea75351191f-7aaf2e07ac3e-1000--penalties-how-to-hit-them-where-to-put-them/

Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review, 93*(2), 136-153. https://doi.org/10.1037/0033-295X.93.2.136

Leng, A., Lian, Z., Lien, J., & Zheng, J. (2023). Revisiting the Asymmetric Matching Pennies Contradiction in China. *Behavioral Sciences, 13*. https://doi.org/10.3390/bs13090757

McKelvey, R., & Palfrey, T. (1995). Quantal response equilibria for normal form games. *Games and Economic Behaviour, 10*, 6-38. https://doi.org/10.1006/game.1995.1023

McKelvey, R., Palfrey, T., & Weber, R. (2000). The effects of payoff magnitude and heterogeneity on behavior in 2×2 games with unique mixed strategy equilibria. *Journal of Economic Behavior & Organization, 42*, 523-548. https://doi.org/10.1016/S0167-2681(00)00102-5

Ochs, J. (1995). Games with Unique, Mixed Strategy Equilibria: An Experimental Study. *Games and Economic Behaviour, 10*, 202-217. https://doi.org/doi.org/10.1006/game.1995.1030

Selten, R., & Chmura, T. (2008). Stationary Concepts for Experimental 2x2-Games. *American Economic Review, 98*, 938-966. https://doi.org/10.1257/aer.98.3.938

Singh, R., & Mukherjee, S. (2026). Replication: No action 'Bias' among elite soccer goalkeepers during penalty kicks. *Journal of Economic Psychology, 113*, Article 102886. https://doi.org/j.joep.2026.102886

Walker, M., & Wooders, J. (2001). Minimax play at Wimbledon. *The American Economic Review, 91*, 1521-1538. https://doi.org/10.1257/aer.91.5.1521

(12) Appendix

(12.1) Derivation of ε critical value

The new payoff structure that includes action bias leads to the following utility equations:

$$U_g(L) = -0.704\alpha - 0.902\beta - (1 - \alpha - \beta) = -1 + 0.296\alpha + 0.098\beta$$

$$U_g(C) = -(1 + \varepsilon)[\alpha + 0.4\beta + (1 - \alpha - \beta)] = (1 + \varepsilon)(-1 + 0.6\beta)$$

$$U_g(R) = -\alpha - 0.968\beta - 0.746(1 - \alpha - \beta) = -0.746 - 0.254\alpha - 0.222\beta$$

We know from observing Figure 6 that ε* is the value that the goalkeeper's LCR best response lies on the CR edge of the plot where $U_g(L) = U_g(C) = U_g(R)$ implies $\alpha = 0$. Therefore, it is the value of ε that leads to the goalie being indifferent across all three actions when the kicker only plays centre and right.

The goalie is indifferent across all three actions when:

$$= -1 + 0.296\alpha + 0.098\beta = (1 + \varepsilon)(-1 + 0.6\beta) = -0.746 - 0.254\alpha - 0.222\beta$$

Equating the left and right expressions gives:

$$\alpha = \frac{127}{275} - \frac{32}{55}\beta$$

Equating the left and centre equations gives:

$$0.296\alpha + \varepsilon = (0.502 + 0.6\varepsilon)\beta$$

Substituting α

$$\beta = \frac{\left(\varepsilon + \frac{4699}{34375}\right)}{\frac{18541}{27500} + \frac{3}{5}\varepsilon}$$

When $\alpha = 0$

$$0 = \frac{127}{275} - \frac{32}{55}\beta$$

$$\beta = \frac{127}{160}$$

$$\frac{127}{160} = \frac{\left(\varepsilon + \frac{4699}{34375}\right)}{\frac{18541}{27500} + \frac{3}{5}\varepsilon}$$

$$\varepsilon^* = \frac{31877}{41900} \approx 0.7608$$

Therefore, the value of ε where the goalie is indifferent across all three directions whilst the kicker only plays CR is ≈ 0.7608

(12.2) Unique equilibria testing

The best response plots below allow us to exclude support profiles which are inconsistent with equilibrium. We first identify all seven of the kicker's support options (L,C,R,LC,LR,CR,LCR) and check whether it is compatible with equilibrium given the goalkeeper's possible best response options. We start with the case where $\varepsilon < \varepsilon^*$. The figure below shows the best response correspondences for such cases[3].

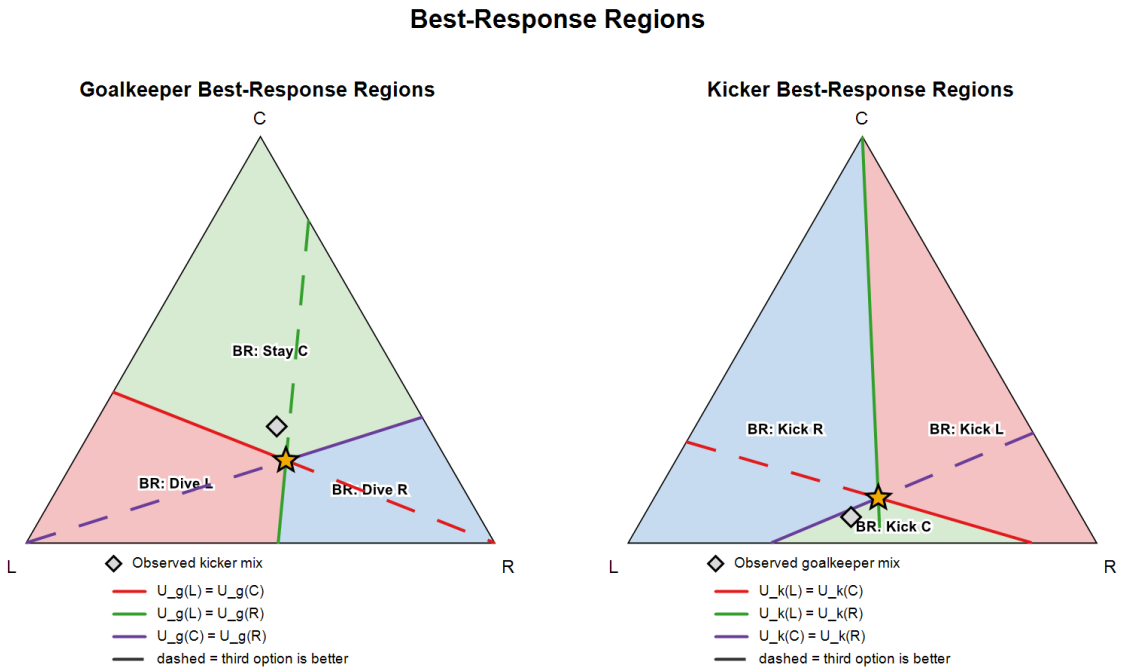
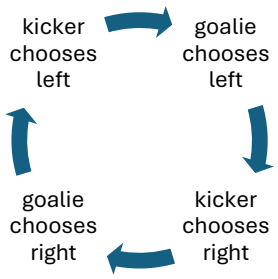


Figure 30 – Best response plot under full support

If an opponent's mix lands anywhere in the coloured regions the player's best response is to exclusively play one pure strategy. For example, if the kicker's strategy support profile is (L) we obtain the following cycle showing that we cannot have an equilibrium.

[3] The plot shows the case where $\varepsilon = 0$ but all other full support cases are qualitatively similar



If the kicker chooses left with certainty the goalie's best response is to exclusively dive left, shown in the red region of the left simplex above; to which the kicker's best response is to kick to kick to the right with certainty, shown by the blue region on the righthand simplex above. This contradicts the starting assumption, showing the kicker support case (L) cannot support an equilibrium.

In the following table, each possible K support is indicated in the first column. The second column gives the supports of all G strategies that are best responses to some K strategy with the column 1 support. Finally, column 3 gives the supports of all K best responses to some G strategy with a support in column 2. If the first column support does not appear in the third column, there cannot be an equilibrium with K's strategy supported as in column 1.

When $\varepsilon < \varepsilon^*$:

Kicker support	**Goalie possible BR options**	**Kicker possible BR to all available goalie BR**	**Possible Equilibrium**
L	L	R	No
C	C	LR	No
R	R	L	No
LC	LC, L, C	R, LR	No
LR	LR, L, R	L, C, R, LC, CR	No
CR	C, R, CR	LR, L	No
LCR	L, C, R, LC, LR, CR, LCR	L, C, R, LC, LR, CR, LCR	Yes

Table 48- Possible support testing for $\varepsilon < \varepsilon^*$

From this we rule out any possible equilibria in all support cases except LCR. It is only when all three indifference lines intersect that we see both players land on a stable equilibrium. Seeing as there is only one intersection point on either plot, we can see that there is only one unique equilibrium.

A qualitative change to the best response correspondences only occurs at $\varepsilon = \varepsilon^*$. This is because varying ε changes the position of the goalkeeper indifference curves and therefore the points of intersection but does not change the which plot edge (reduced support strategy) each of the lines intersect with until ε^*. For example, the red ($U_g(L) = U_g(C)$) line only ever intersects the left (Left-Centre) edge of the simplex meaning the goalie will never respond to a kicker playing CR with anything other than C, R, CR. This holds true for all support cases where varying ε, given $\varepsilon < \varepsilon^*$, will not change any of the goalie's best response options for a given kicker support. This means we repeat analysis for $\varepsilon = \varepsilon^*$ and $\varepsilon > \varepsilon^*$.

When ε=ε*:

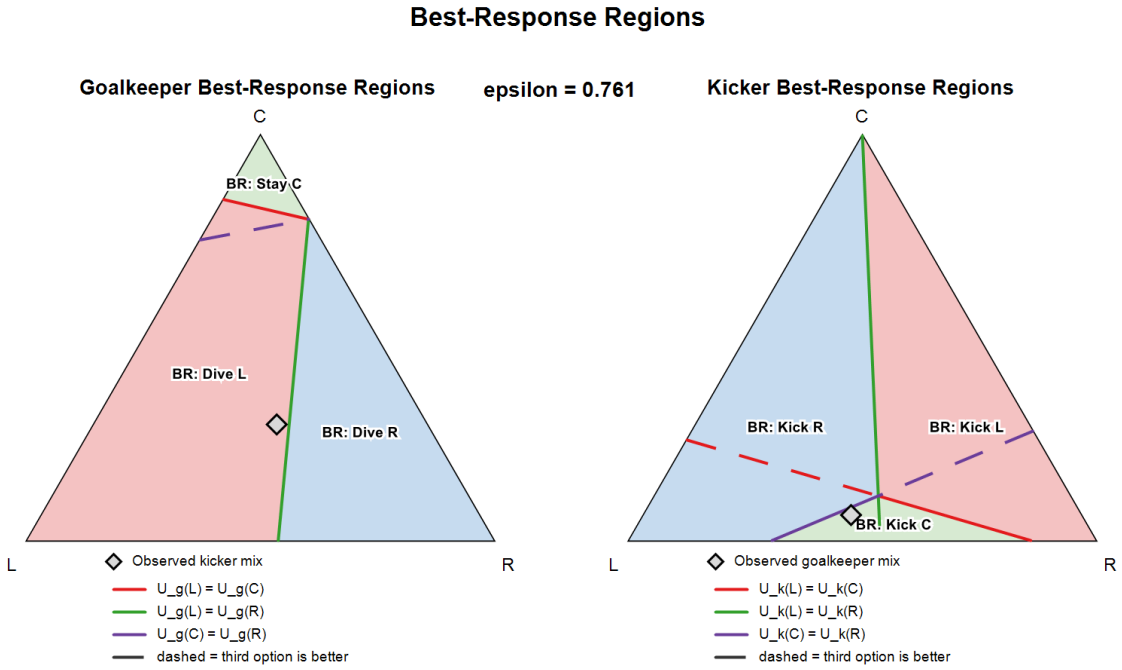


Figure 31 – Best response plot for knifes edge case

Kicker support	Goalie possible BR options	Kicker possible BR to all available goalie BR	Possible Equilibrium
L	L	R	No
C	C	LR	No
R	R	L	No
LC	LC, L, C	R, LR	No
LR	LR, L, R	L, C, R, LC, CR	No
CR	C, R, LC, LR, CR, LCR	L, C, R, LC, LR, CR, LCR	Yes
LCR	L, C, R, LC, LR	L, C, R, LC, LR, CR	No

Table 49 – Possible support testing for $\varepsilon = \varepsilon^*$

When the kicker plays CR, it means playing left is strictly worse for the kicker.

$$U_k(L) \leq U_k(C) = U_k(R)$$

Additionally, the purple line on the right plot lies within the interior of the simplex and intercepts the Centre Right edge, meaning the goalie must play LCR, or CR to support the kicker's CR profile.

$$U_k(C) = U_k(R)$$

$$0.704p + q + (1 - p - q) = p + q + 0.746(1 - p - q)$$

$$0.32p + 0.822q = 0.222$$

$$p = \frac{111 - 411q}{160}$$

$$1 - p - q = \frac{49 + 251q}{160}$$

From the main text we know the kicker is indifferent across all three actions when

$$p = \frac{7620}{18541}$$

$$q = \frac{2041}{18541}$$

$$r = \frac{8880}{18541}$$

Thus, at the knifes edge case the equilibria are:

$$(\alpha, \beta, \gamma) = \left(0, \frac{127}{160}, \frac{33}{160}\right)$$

$$(p, q, r) = \left(\frac{111 - 411q}{160}, q, \frac{49 + 251q}{160}\right)$$

$$0 \leq q \leq \frac{2041}{18541}$$

Beyond the critical value of ε the red and green lines no longer intersect, changing the goalie's best response options once more.

When $\varepsilon > \varepsilon^*$:

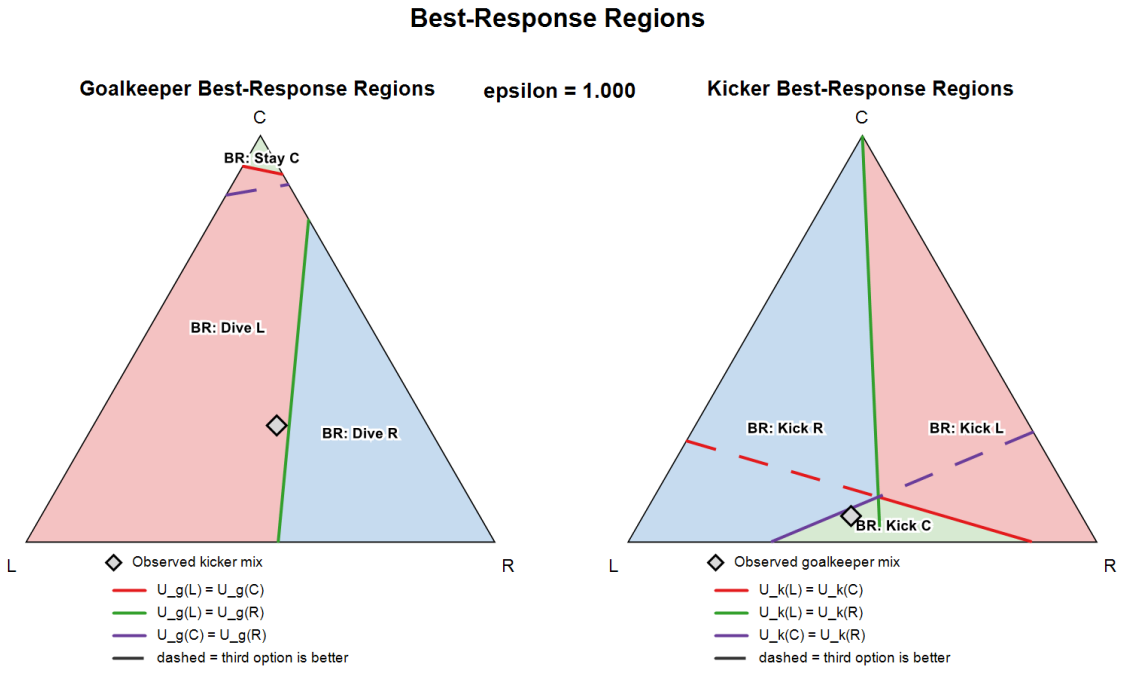


Figure 32 – Best response plot for reduced support case

Kicker support	Goalie possible BR options	Kicker possible BR to all available goalie BR	Possible Equilibrium
L	L	R	No
C	C	LR	No
R	R	L	No
LC	LC, L, C	R, LR	No
LR	LR, L, R	L, C, R, LC, CR	No
CR	L, C, R, LC, LR	L, C, R, LC, LR, CR, LCR	Yes
LCR	L, C, R, LC, LR	L, C, R, LC, LR, CR	No

Table 50 - Possible support testing for $\varepsilon > \varepsilon^*$

Here we see that the goalie no longer has the LCR support option when the kicker plays CR. This is because the red, green, and purple lines no longer all intersect.

Only the green and red intercept the CR edge of the left plot. We know from the tables above that goalie LC never leads to an equilibrium. Therefore, in equilibrium, the goalie will only ever play LR. Thus, instead of a continuum of mixes, the goalie now only has one best response that results in an equilibrium. This equilibrium mix is the point where the purple kicker indifference line intercepts the CR edge of the righthand plot, meaning the only equilibrium mixes played are:

$$(\alpha, \beta, \gamma) = \left(0, \frac{127}{160}, \frac{33}{160}\right)$$

$$(p, q, r) = \left(\frac{111}{160}, 0, \frac{49}{160}\right)$$

Observing the plot also shows why this mix does not change as ε continues to increase. The goalie mixes over LR, the kicker responds by mixing over CR. Neither the point on the left plot where the green line intercepts the CR edge, or the point on the right plot where the purple line intercepts the LR plot move as ε increases past the critical value. Therefore, the supports played, as well as the mixes within those supports are stable (fixed, unmoving).

To summarise:

$$\varepsilon^* = \frac{31877}{41900} \approx 0.7608$$

When $0 \leq \varepsilon < \varepsilon^*$

$$(\alpha, \beta, \gamma) = \left(\frac{127}{275} - \frac{32}{55}\beta(\varepsilon), \beta(\varepsilon), \frac{148}{275} - \frac{23}{55}\beta(\varepsilon)\right)$$

$$(p, q, r) = \left(\frac{7620}{18541}, \frac{2041}{18541}, \frac{8880}{18541}\right)$$

where

$$\beta(\varepsilon) = \frac{\varepsilon + \frac{4699}{34375}}{\frac{18541}{27500} + \frac{3}{5}\varepsilon}$$

When $\varepsilon = \varepsilon^*$

$$(\alpha, \beta, \gamma) = \left(0, \frac{127}{160}, \frac{33}{160}\right)$$

$$(p, q, r) = \left(\frac{111 - 411q}{160}, q, \frac{49 + 251q}{160}\right), \qquad 0 \leq q \leq \frac{2041}{18541}$$

And when $\varepsilon > \varepsilon^*$

$$(\alpha, \beta, \gamma) = \left(0, \frac{127}{160}, \frac{33}{160}\right)$$

$$(p, q, r) = \left(\frac{111}{160}, 0, \frac{49}{160}\right)$$

(12.3) Player vs joint likelihood

The likelihood function used to test the fit of models is:

$$\mathrm{LL}(\theta) = \sum_{i\in\{L,C,R\}} \sum_{j\in\{L,C,R\}} n_{ij} \log\Big(P_{ij}(\theta)\Big)$$

However, for diagnostic purposes separate individual player likelihood functions were also used. These are:

$$\mathrm{LL}_K(\theta) = \sum_i n_i^K \log p_i^K\,(\theta)$$

$$LL_G(\theta) = \sum_j n_j^G \log p_j^G\,(\theta)$$

These expressions apply to the kicker and goalie respectively.

Under the homogeneous setup:

$$\sum_i n_i^K \log p_i^K\,(\theta) + \sum_j n_j^G \log p_j^G\,(\theta) = \sum_{i\in\{L,C,R\}} \sum_{j\in\{L,C,R\}} n_{ij} \log\Big(P_{ij}(\theta)\Big)$$

$$where\ \theta = (\varepsilon)$$

n_i^K is the number of observed kicks in direction i, and n_j^G is the number of observed dives where the goalie chooses direction j

Therefore, the sum of the two player likelihoods is the same as the joint likelihood.

For the heterogeneous action bias setup

$$\theta = \left(\varepsilon_{\text{low}}, \varepsilon_{\text{high}}, \pi\right)$$

$$P_{ij}(\theta) = (1-\pi)p_i(\varepsilon_L)q_j(\varepsilon_L) + \pi p_i(\varepsilon_H)q_j(\varepsilon_H)$$

This leads to the following joint likelihood function:

$$LL_{\text{HetAB}}(\theta) = \sum_{i\in\{L,C,R\}} \sum_{j\in\{L,C,R\}} n_{ij} \log\big[(1-\pi)p_i(\varepsilon_L)q_j(\varepsilon_L) + \pi p_i(\varepsilon_H)q_j(\varepsilon_H)\big]$$

For each player

$$p_i^{pooled} = (1-\pi)p_i(\varepsilon_L) + \pi p_i(\varepsilon_H)$$

$$q_j^{pooled} = (1-\pi)q_j(\varepsilon_L) + \pi q_j(\varepsilon_H)$$

The player specific likelihoods therefore become:

$$LL_K = \sum_i n_i \ \log p_i^{pooled}$$

$$LL_G = \sum_j n_j \ \log p_j^{pooled}$$

Unlike for the homogeneous model, $P_{ij}^{joint} \neq p_i^{pooled} q_j^{pooled}$

Here, P_{ij}^{joint} is the probability of the kicker playing i and the goalie playing j

This occurs because:

$$P_{ij}^{joint} = (1-\pi)p_i(\varepsilon_L)q_j(\varepsilon_L) + \pi p_i(\varepsilon_H)q_j(\varepsilon_H)$$

By contrast:

$$p_i^{pooled} q_j^{pooled} = [(1-\pi)p_i(\varepsilon_L) + \pi p_i(\varepsilon_H)]\big[(1-\pi)q_j(\varepsilon_L) + \pi q_j(\varepsilon_H)\big]$$

(12.4) QRE model results table

Model	Lambda specification	MLE parameters	Log likelihood	Interpretation
Base QRE	Common lambda	$\lambda \to \infty$ $\varepsilon = 0$	−576.249682	Collapses to the base MSNE.
Base QRE	Separate lambdas	$\lambda_K \to \infty$ $\lambda_G \to \infty$ $\varepsilon = 0$	−576.249682	Same limiting fit as the common lambda base QRE.
QRE with homogeneous action bias	Common lambda	$\lambda = 0.662254$ $\varepsilon = 3.871276$	−567.153318	Overall preferred QRE model under parsimony.
QRE with homogeneous action bias	Separate lambdas	$\lambda_K \approx 0$ $\lambda_G = 2.904793$ $\varepsilon = 0.971603$	−566.803176	Best fitting QRE model.
Complete information QRE with action bias	Common lambda	$\lambda = 0.662255$ $\varepsilon_{low} = 3.871263$ $\varepsilon_{high} = 3.871271$ $\pi_{high} \approx 0.995904$	−567.153318	Produces the same pooled equilibrium and maximized likelihood as the common lambda homogeneous model.
Complete information QRE with action bias	Separate lambdas	$\lambda_K \approx 0$ $\lambda_G = 2.904793$ $\varepsilon_{low} = not\ identidied$ $\varepsilon_{high} = 0.971605$ $\pi_{high} = 1.000000$	−566.803176	Produces the same pooled equilibrium and maximized likelihood as the separate lambda homogeneous model.
Incomplete information QRE with action bias	Common lambda	$\lambda = 0.662255$ $\varepsilon_{low} = \text{not identified}$ $\varepsilon_{high} = \text{not identified}$ $\pi_{high} = not\ identified$	−567.153318	Produces the same pooled equilibrium and maximized likelihood as the common lambda homogeneous model.
Incomplete information	Separate lambdas	$\lambda_K \approx 0$ $\lambda_G = 2.904793$	−566.803176	Produces the same pooled equilibrium and maximized

QRE with action bias		$\varepsilon_{low} = not\ identidied$ $\varepsilon_{high} = 0.971605$ $\pi_{high} = 1.000000$		likelihood as the separate lambda homogeneous model.

Table 51 - QRE MLE table

(12.5) Incomplete information result

(12.5.1) Proof of proposition 2

Let $F(\varepsilon)$ represent an arbitrary distribution of ε. Each goalkeeper's strategy is then a function of type: $q(\varepsilon) = \big(q_L(\varepsilon), q_C(\varepsilon), q_R(\varepsilon)\big)$. Here, $q_L(\varepsilon), q_C(\varepsilon), q_R(\varepsilon)$ represent the goalkeeper's probability of going left, centre, and right, respectively. The pooled goalkeeper mix is therefore:

$$\bar{q} = \int q\,(\varepsilon)\, dF(\varepsilon).$$

Here, $\bar{q}$ denotes the goalkeeper's pooled probability mixture across all three actions. Under incomplete information the kicker is responding to the goalkeeper's pooled mix. Therefore, the pooled goalkeeper mix must keep the kicker indifferent to maintain equilibrium; given we are trying to match to the observed data, doing so requires the kicker's mix to be in full support. Given action bias does not affect the kicker's payoffs we know that the pooled goalkeeper mix must be the same as the full support MSNE mix.

Therefore,

$$\bar{q} = (0.410981069,\ 0.110080362,\ 0.478938569)$$

The kicker does not need to make every goalkeeper type indifferent across all actions observed in the pooled strategy. Instead, each action used in the pooled goalkeeper mix must be optimal for at least one goalkeeper type.

$$U_g(L) = -1 + 0.296\alpha + 0.098\beta$$

$$U_g(C) = (1 + \varepsilon)(-1 + 0.6\beta)$$

$$U_g(R) = -0.746 - 0.254\alpha - 0.222\beta$$

We see that action bias only affects the goalkeeper's centre payoffs. It follows that as ε varies $U_g(L)$ and $U_g(R)$ will remain constant for a given kicker mix. Consequently, if both left and right appear with positive probability in the pooled goalkeeper mix, the kicker must choose a mix satisfying $U_g(L) = U_g(R)$.

We define p as the kicker's Bayesian mix. At the Bayesian MSNE $U_g(L|p) = U_g(R|p)$ is always satisfied for all goalkeeper types. Given $U_g(C|p, \varepsilon)$ decreases with ε, there exists a point satisfying $U_g(L|p) = U_g(R|p) = U_g\big(C|p, \varepsilon^{\#}\big)$. Here, $\varepsilon^{\#}$ denotes the value of ε where the goalkeeper is indifferent between all three actions.

Therefore,

$$q_C(\varepsilon) = \begin{cases} 1, & \varepsilon < \varepsilon^{\#} \\ 0, & \varepsilon > \varepsilon^{\#} \end{cases}$$

We know to keep the kicker indifferent the goalkeeper's pooled centre rate must be 0.110080362, and the pooled centre rate must be the population average of all type specific mixes and for a continuous distribution $\Pr(\varepsilon = \varepsilon^{\#}) = 0$

Therefore, $\Pr(\varepsilon < \varepsilon^{\#}) = 0.110080362$.

Therefore, the MLE search can only identify the kicker's mix which satisfies the condition $U_g(L|p) = U_g(R|p)$ and is the closest fit to the observed data. From Figure 3 we know that this point must lie on the green line on the left plot. The best fitting kicker mix is then $\mathrm{p} = (0.2979,\ 0.2817,\ 0.4204)$. Substituting these values into $U_g(L|p) = U_g(R|p) = U_g(C|p, \varepsilon^{\#})$ gives:

$$-0.88420935 = (1 + \varepsilon^{\#})(-0.83097779) = -0.88420935$$

$$\varepsilon^{\#} = 0.064059$$

Therefore, the incomplete information assumption identifies the same kicker mix and pooled goalie mix as the homogeneous action bias model. The only thing we learn from the MSNE incomplete information assumption is that $F(\varepsilon)$ must be such that 11% of goalkeepers have an action bias value less than 0.064059.