University of East Anglia

Wage Dispersion, National Diversity and Performance

- Evidence from the National Basketball Association

Ritchie Craig Woodard



A thesis presented for the degree of

DOCTOR OF PHILOSOPHY

in Economics

Submitted to

The University of East Anglia

Faculty of Social Science

School of Economics

Norwich - June, 2020

"This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with the author and that use of any information derived therefrom must be in accordance with current UK Copyright Law. In addition, any quotation or extract must include full attribution."

Primary Supervisor:
Dr. Peter Dawson
Secondary Supervisor:
Dr. Georgios Papadopoulos

Candidate:
Ritchie Craig Woodard
Univeristy of East Anglia
Student Number: 4413121

Abstract

This thesis investigates the effects of wage dispersion and national diversity on performance, using a detailed dataset from the National Basketball Association (NBA). Sports data offers a unique opportunity to test labour market theories. The richness of the data, well-defined indicators of success, and the ability to track workers throughout their careers, are useful traits for labour market research. In this investigation, we exploit this to explore the effects of wage dispersion and national diversity on team-level performance, and gain further insights through player-level analysis.

Wage dispersion is a well-researched topic within the academic literature, both within sports economics and more generally; however, the evidence has been mixed with no clear consensus. We contribute to this literature by providing evidence that multiple measures of dispersion, namely the standard deviation and polarisation, have opposing effects on performance. We demonstrate how these are correlated with other commonly-used measures of dispersion, which may explain the mixed evidence previously found. Our findings suggest that an ideal wage structure to maximise performance should have a large spread of wages, but avoid clusterings, particularly at the extremes of the distribution. Furthermore, these team-level performance effects are driven by the top earners in a team, with opposing effects found for the performance of low earners.

Research on national diversity is a relatively more recent branch of the academic literature. There has been some debate regarding how best to capture national diversity, with measures of fractionalisation, polarisation and of cultural distance considered. We explore the relationships between these variables and show how measures of fractionalisation and polarisation can be combined with measures of cultural distance. We find evidence that national fractionalisation has positive effects on performance, while polarisation again has negative effects.

We conclude that polarisation is an underutilised measure in the current literature.

Access Condition and Agreement

Each deposit in UEA Digital Repository is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the Data Collections is not permitted, except that material may be duplicated by you for your research use or for educational purposes in electronic or print form. You must obtain permission from the copyright holder, usually the author, for any other use. Exceptions only apply where a deposit may be explicitly provided under a stated licence, such as a Creative Commons licence or Open Government licence.

Electronic or print copies may not be offered, whether for sale or otherwise to anyone, unless explicitly stated under a Creative Commons or Open Government license. Unauthorised reproduction, editing or reformatting for resale purposes is explicitly prohibited (except where approved by the copyright holder themselves) and UEA reserves the right to take immediate ‘take down’ action on behalf of the copyright and/or rights holder if this Access condition of the UEA Digital Repository is breached. Any material in this database has been supplied on the understanding that it is copyright material and that no quotation from the material may be published without proper acknowledgement.

Contents

Abstract 1

Contents 2

List of Tables 5

List of Figures 8

Glossary of Variables 9

Acknowledgements 11

1 Introduction 12
- 1.1 Wage Dispersion . . . 13
- 1.2 National Diversity . . . 14
- 1.3 Sports Labour Markets . . . 15
- 1.4 Measuring Performance . . . 17
- 1.5 Structure of the Investigation . . . 17

2 Literature Review 19
- 2.1 Wage Dispersion . . . 20
  - 2.1.1 Theory . . . 20
  - 2.1.2 General Evidence . . . 22
  - 2.1.3 Sports Evidence . . . 23
- 2.2 National Diversity . . . 25
  - 2.2.1 Theory . . . 25
  - 2.2.2 General Evidence . . . 26
  - 2.2.3 Sports Evidence . . . 26
- 2.3 Measuring Performance . . . 28
  - 2.3.1 Team Performance . . . 28
  - 2.3.2 Individual Player Performance . . . 30
- 2.4 Summary . . . 33

3 Dispersion, Diversity and Polarisation 35
- 3.1 Measures of Wage Dispersion . . . 36
- 3.2 Measures of National Diversity . . . 42
- 3.3 Summary . . . 47

4 Data and Methodology 48
- 4.1 Data Collection . . . 49

4.2 Data Structure and Econometric Techniques . . . 50
4.3 Dependent Variables . . . 53
4.3.1 Basic Statistics and Notation . . . 53
4.3.2 Measures of Team Performance . . . 55
4.3.3 Measures of Individual Player Performance . . . 58
4.4 Covariates . . . 64
4.4.1 Descriptive Statistics and Correlations - Team-level . . . 67
4.4.2 Descriptive Statistics and Correlations - Player-level . . . 69
4.5 Summary . . . 71

5 Wage Dispersion and Team Performance **73**
5.1 Introduction . . . 74
5.2 Descriptive Statistics and Correlations . . . 74
5.3 Regression Results . . . 78
5.3.1 Robustness Checks . . . 84
5.4 Summary and Discussion . . . 87

6 National Diversity and Team Performance **89**
6.1 Introduction . . . 90
6.2 Descriptive Statistics and Correlations . . . 90
6.3 Regression Results . . . 96
6.3.1 Robustness Checks . . . 101
6.4 National Diversity, Wage Dispersion and Team Performance . . . 103
6.4.1 Robustness Checks . . . 109
6.5 Summary and Discussion . . . 111

7 Wage, Nationality and Player Performance **114**
7.1 Introduction . . . 115
7.2 Descriptive Statistics and Correlations . . . 115
7.3 Wage Dispersion Regression Results . . . 119
7.3.1 Wage Dispersion Robustness Checks . . . 123
7.4 National Diversity Regression Results . . . 128
7.4.1 National Diversity Robustness Checks . . . 131
7.5 Wage Dispersion, National Diversity and Individual Player Performance . . . 134
7.6 Summary and Discussion . . . 139

8 Conclusion **141**
8.1 Choice of NBA . . . 142
8.2 Wage Dispersion . . . 142
8.3 National Diversity . . . 143
8.4 Further Insights from Player Analysis . . . 144
8.5 Limitations and Further Work . . . 145
8.6 Concluding Remarks . . . 146

References **i**

Appendices **ix**
4.A Alternative Estimates of Possession . . . x

4.B Wins Produced (WP) - (Berri et al., 2006) xiii
4.C Player Efficiency Rating (PER) - (Hollinger, 2002, 2005) xviii
4.D Box Plus/Minus (BPM) - (Myers, 2014) xx
4.E The Best Players in the NBA xxi
4.F Importance of Team-level Controls in Determining Team Performance xxii
4.G 'Hausman-type' Test - Control Variables xxv
4.H Importance of Team-level Controls in Determining Wage Dispersion and National Diversity xxvi
4.I Importance of Player-level Controls in Determining Player Performance xxviii
4.J Importance of Player-level Controls in Determining Wage Dispersion and National Diversity xxx
5.A 'Hausman-type' Test - Wage Dispersion xxxii
5.B Alternative Measures of Polarisation xxxiii
5.C Robustness Check - Alternative Measures of Polarisation xxxvi
6.A 'Hausman-type' Test - National Diversity xxxvii
6.B Wage Distributions of Example Teams xxxviii
7.A Wage Dispersion and Individual Player Performance - Calculating the Magnitude of the Predicted Effects xli

List of Tables

3.1 Descriptive Statistics of the Distributions in Figure 3.3 40
3.2 Descriptive Statistics of the Distributions in Figure 3.4 41
3.3 Examples of Linguistic Branches 45
3.4 Examples of Linguistic Distance 46

4.1 Notation and Per-Game Means of Individual Player Statistics 54
4.2 Means of Per-Game Team Statistics 54
4.3 Descriptive Statistics of Team Performance Measures 56
4.4 Correlations Between Team Performance Variables 57
4.5 Marginal Effects of Player Statistics on Team Win Percentage 59
4.6 Descriptive Statistics of Individual Player Performance Measures 61
4.7 Correlation Between Individual Player Performance Measures 64
4.8 Descriptive Statistics of Team-level Control Variables 68
4.9 Correlations Between Team-level Control Variables 69
4.10 Descriptive Statistics of Player-level Control Variables 70
4.11 Correlations Between Player-level Control Variables 71

5.1 Descriptive Statistics of Wage Dispersion Variables 75
5.2 Correlations Between Wage Dispersion Variables 77
5.3 Correlations Between Wage Dispersion and Control Variables 78
5.4 Regression Results of Wage Dispersion Measures (No Controls) 79
5.5 Regression Results of Wage Dispersion Measures (Some Controls) 80
5.6 Regression Results of Wage Dispersion Measures (Full Controls) 81
5.7 Regression Results with Multiple Wage Dispersion Measures 82
5.8 Robustness Checks (Relaxing Controls and Definition of the Team) 84
5.9 Robustness Check (Performance Measures) 85
5.10 Robustness Check (Alternative Models) 86

6.1 Descriptive Statistics Relating to the Structure of National Diversity 90
6.2 Descriptive Statistics of National Diversity Considering Linguistic Distance . 92
6.3 Correlations Between Linguistic Distance Measures 93
6.4 The Most Nationally Diverse Teams and their Roster Structures 94
6.5 Correlations Between National Diversity Variables 95
6.6 Correlations Between National Diversity Measures and Control Variables . . 95
6.7 Regression Results of National Diversity Measures (No Controls) 96
6.8 Regression Results of National Diversity Measures (Some Controls) 97
6.9 Regression Results of National Diversity Measures (Full Controls) 98
6.10 Regression Results for Multiple National Diversity Measures 100

6.11 Robustness Check (Relaxing Controls and Definition of the Team) 101
6.12 Robustness Check (Performance Measures) 102
6.13 Robustness Check (Alternative Models) 103
6.14 Correlations Between Wage Dispersion and National Diversity Variables . . 103
6.15 Regression Results of Wage Dispersion and National Diversity Measures . . 104
6.16 Description of Example Teams 105
6.17 Predicted Effects of Example Teams 106
6.18 Effect of Adding Foreign Players from Differing Numbers of Additional Nationalities 108
6.19 Robustness check (Relaxing Controls and Definition of the Team) 110
6.20 Robustness check (Performance Measures) 110
6.21 Robustness checks (Alternative Models) 111

7.1 Wages at each Wage Rank 116
7.2 Correlations Between Player-level Wage and Team-level Wage Dispersion . . 117
7.3 Descriptive Statistics of Player-level National Diversity Measures 117
7.4 Correlations Between Player- and Team-level National Diversity 118
7.5 Player Performance and Wage Dispersion (No Controls) 119
7.6 Player Performance and Wage Dispersion (Player- and Team-level Controls) 120
7.7 Player Performance and Wage Dispersion with Interactions 121
7.8 Wage Dispersion Robustness Check (Performance Measures) 124
7.9 Wage Dispersion Robustness Check (Reduced Sample) 125
7.10 Wage Dispersion Robustness Check (Alternative Specifications) 126
7.11 Wage Dispersion Robustness Check (Alternative Models) 127
7.12 Player Performance and National Diversity (No Controls) 128
7.13 Player Performance and National Diversity (Player- and Team-level Controls) 129
7.14 Player Performance and National Diversity (Team- and Player-level Measures) 130
7.15 National Diversity Robustness Check (Performance Measures) 132
7.16 National Diversity Robustness Check (Reduced Sample) 133
7.17 National Diveristy Robustness Check (Alternative Models) 133
7.18 Wage Dispersion, National Diversity and Player Performance (Performance Measures) 135
7.19 Wage Dispersion, National Diveristy and Player Performance (Alternative Models) 138

4.A.1 Determining x, z and tmRB xi
4.A.2 Alternative Estimates of Possessions xii
4.B.1 Estimating Win Percentage using Offensive and Defensive Ratings xiii
4.B.2 Estimating the Marginal Effect of Personal Fouls xv
4.B.3 Estimating the Marginal Effect of Blocks xvi
4.E.1 The Best 5 Players by each Player Performance Measure xxi
4.F.1 Team-level Control Variables and Performance xxii
4.G.1 Hausman-Type Test using the Mundlak (1978) Approach - Team-level Control Variables xxv
4.H.1 The Importance of Team-level Control Variables xxvi
4.I.1 Player-level Control Variables and Performance xxviii
4.J.1 The Importance of Player-level Control Variables xxx

4.J.2 The Importance of Player- and Team-level Control Variables xxxi
5.A.1 Hausman-Type Test using the Mundlak (1978) Approach - Wage Dispersion xxxii
5.B.1 Descriptive Statistics of Alternative Polarisation Measures xxxiv
5.B.2 Correlation Between Alternative Measures of Wage Polarisation xxxv
5.C.1 Robustness Check (Alternative Polarisation Measures) xxxvi
6.A.1 Hausman-Type Test using the Mundlak (1978) Approach - National Diversity xxxvii
6.B.1 Wage Distributions of Example Teams xxxviii
6.B.2 Description of Example Teams xxxix
6.B.3 Predicted Effects of Example Teams xxxix
7.A.1 Calculating the Predicted Effects of the Interacted Wage Dispersion Measures on Player Performance xli

List of Figures

3.1 The Arbitrariness of Setting Bounds for the 'Middle' 36
3.2 Foster and Wolfson (2010) Polarisation 38
3.3 Increased Dispersion and Polarisation 40
3.4 Decreased Dispersion and Increased Polarisation 41
3.5 Fractionalisation and Polarisation Examples 44
3.6 Fractionalisation and Polarisation Simulations 45

4.1 Histograms of Team Performance Measures 56
4.2 Player Performance Measures and Minutes Played 62
4.3 Histograms of Standardised Player Performance Measures 63

5.1 Example Wage Distributions 76
5.2 Standard Deviation of Wage and Wage Polarisation 77

6.1 National Fractionalisation and Polarisation 91
6.2 National Fractionalisation and Polarisation: Considering Linguistic Distance 93
6.3 National Fractionalisation and Polarisation: Spitting the sample 99
6.4 Effect of Adding Foreign Players from Differing Numbers of Additional Nationalities 109

7.1 Histogram of Average Linguistic Distance (Considering Two Languages) Between NBA Players and their Teammates 118
7.2 Wage Rank and the Standard Deviation of Wage 122
7.3 The Natural Logarithm of Wage and Wage Polarisation 122

5.B.1 *FW* Polarisation Measures (Natural Log of Wage versus Absolute Wage) . . xxxiv
6.B.1 Histograms of Example Team Wage Distribution xl

Glossary of Variables

Variable	Description
abWagePolFW	Polarisation of Wage within Teams (for absolute values of Wage) -Foster and Wolfson (2010)
BPM	Box Plus/Minus - (Myers, 2014)
CTSize	Core Team Size (Includes Players which Play at least 20 Games and at least 12 Minutes per Game)
DRTG	Team Defensive Rating
FRGN	Proportion of Team's Roster Born Outside of the USA
HighFrac	Dummy Variable: Equal to 1 if NatFrac>0.5, and 0 otherwise
LingDist(1-4)	Average Linguistic Distance between Players within Teams (considering 1-4 languages)
NatFrac	National Fractionalisation within Teams
NatFrac(1-4)	National Fractionalisation within Teams incorporating Linguistic Distance (considering 1-4 languages)
NatNum	Number of Nationalities within Teams
NatPol	National Polarisation within Teams
NatPol(1-4)	National Polarisation within Teams incorporating Linguistic Distance (considering 1-4 languages)
ORTG	Team Offensive Rating
PDG	Points Difference per Game
PER	Player Efficiency Rating - Hollinger (2002, 2005)
Player	Player Identifier
PExp	NBA Experience of Player (in Years)
PNatPer	Proportion of the Team which share the same Nationality as a given Player
PTenureMP	Minutes Played for current Team prior to the current Season
SRS	Simple Ratings System
Size	Roster Size (including all Players to play any minutes for a Team)
TExp	Average NBA Experience of Players within Teams (in Years)
TExpSD	Standard Deviation of NBA Experience of Players within Teams (in Years)
TPayroll	Expenditure on Wages by Teams
Trade	Dummy Variable: Equal to 1 if a Player is Traded during a Season, and 0 otherwise
TTenureMP	Average Number of Minutes Played for Team by current Roster, prior to the current Season

TWage	Average Wage within Teams
Team	Team Identifier
VORP	Value Over Replacement Player - Myers (2013)
Wage	Player Wages in Millions of Dollars ($m)
WageCV	Coefficient of Variation of Wage within Teams
WageGini	Gini Coefficient of Wage within Teams
WageHHI	Herfindahl-Hirschman Index of Wage within Teams
WageKurt	Kurtosis of Wage within Teams
WagePol(FW)	Polarisation of Wage within Teams - Foster and Wolfson (2010)
$\text{WagePolDER}_{\alpha}$	Polarisation of Wage within Teams - Duclos et al. (2004)
WageRank	Player's Ranking within their Team in terms of Wage (multiplied by minus one so higher values indicate higher ranking)
WageSD	Standard Deviation of Wage within Teams
WageSkew	Skewness of Wage within Teams
Year	Year Identifier
WL	Regular Season Win Percentage
WP	Wins Produced - Berri et al. (2006)
WP48	Wins Produced per 48 minutes - Berri et al. (2006)
WS	Win Shares

Acknowledgements

Firstly, I would like to thank my father, Craig Woodard, for his love, support and faith in me throughout my studies. Special thanks also go to my mother, Linda Woodard, and the rest of my family and friends, for their invaluable support.

I would also like to express my sincere gratitude to my supervisors, Dr. Peter Dawson and Dr. Georgios Papadopoulos. They have been an amazing supervisory team, offering me inspiration, calmness and valuable expertise. More widely, I thank the UEA and the School of Economics for the opportunity to pursue this PhD, and for offering a friendly and engaging working environment in which to do so.

I would also like to thank my colleague, Paul Matthias Gorny, for sharing the PhD experience with me. It would not have been the same without you.

Lastly, a thank you to Linda Woodard, Charlie Flint and Aayushi Awasthy for proof reading parts of this thesis.

Chapter 1

Introduction

Chapter Overview

This opening chapter introduces and motivates the topics of this thesis, namely wage dispersion and national diversity. These are pressing issues which affect the cohesion of teams, both in a sporting context and in general work environments. The investigation utilises a detailed dataset from the National Basketball Association (NBA), spanning from 2007-18, to analyse the effects of these phenomena on team- and player-level performance. The findings of this study are certainly of direct relevance for sports teams, but we also believe our results have further implications and can be applied to a much wider range of scenarios. We begin by discussing the general issues surrounding wage dispersion and national diversity in turn. We then consider the sports industry more specifically, before highlighting some of the issues related to measuring the performance of individual players within teams. The chapter ends by describing the structure for the remainder of our investigation.

1.1 Wage Dispersion

The first contribution of this study is to investigate how the dispersion of wages[1] within teams can influence performance, through avenues such as competition, cooperation and conflict. In most industries, firms elect for a hierarchical structure with wage levels largely based on a worker's relative rank within an organisation, which indeed has several benefits (Baker et al., 1994; Gibbs, 1995; Mirrlees, 1976). Productive effort can be elicited through competition amongst workers, as each aims to achieve the highest rank possible and earn the associated higher wage. However, production within teams also usually requires at least some element of interdependence, and therefore workers must cooperate with one another. The firm must decide whether to offer a dispersed wage distribution to encourage competition amongst workers (Tournament Theory - Lazear & Rosen, 1981), or to offer a more egalitarian structure to encourage cooperation (Cohesion Theory - Levine, 1991); with both strategies having their own pros and cons. If every rank received the same wage, it could be considered perfectly rational for each worker to contribute the minimum effort required to retain a position within the firm. When firms offer different wage levels, they can differentiate among their workforce and reward the most productive workers. The goal of the firm is to offer a wage structure which encourages productive and cooperative effort, as well as discouraging destructive competition and sabotage.

Another dimension to this decision comes when we consider multiple teams competing against one another. If workers are free to move between the teams, they can join whichever one offers them the highest wage. For a given wage budget, teams choosing to disperse their wage distribution can attract the most productive workers, by offering higher wages for their top positions than a more egalitarian firm is able to. However, they will also be offering some wages lower than the egalitarian firm, which may only attract lower skilled workers to fill these remaining positions. A further interesting aspect arises if workers care, not only about their absolute wage level, but their relative wage level within their team and possibly within the industry. Feelings of jealously may arise if workers are paid vastly different wages to one another which may discourage cooperative effort (Pfeffer & Langton, 1993). Such feelings may not be felt by workers on the same absolute wage level, as part of a more egalitarian structure. Alternatively, workers may feel motivated by a dispersed wage structure which gives the opportunity of high wages for workers that reach the top positions (Bloom, 1999). The structure of wage distributions can also influence the formation of group identities within teams. Workers may consider others on similar wages to themselves as a type of in-group, with whom they are better able to build relationships and cooperate with, while considering those on vastly different wages to themselves as an out-group (Duclos et al., 2004). If these types of issues are important determinants of team performance, then firms can benefit by structuring their wage distributions in an optimal way. This can offer a competitive advantage by increasing team performance within a given wage budget.

While there is a relatively large body of work investigating the effects of wage dispersion on team performance, the results have been mixed, both theoretically and empirically. While some reconciliations have been proposed, we question their ability to make predictions ex-ante and suggest that confirmation bias may be an issue when applying them ex-post. This study adds to the existing literature in several ways. We argue there may be confounding factors

[1] For the purposes of this thesis, we will use the terms wage, income and salary synonymously, as these variations are used throughout the literature.

that affect both wage dispersion and team performance which have not been controlled for in previous studies. Omitting such variables could lead to biased estimates of the relationship between wage dispersion and team performance. We also argue and show that wage distributions can be dispersed in various ways which can be captured using different dispersion measures. We argue such measures should not be used interchangeably. Additionally, a single dispersion measure may therefore not capture all the relevant aspects of a wage distribution. So, we not only consider different dispersion measures, but consider them simultaneously in the same model. Finally, but perhaps our most important contribution, is the introduction of polarisation measures into this literature. Particularly the indexation of such measures is a relatively recent trend, and has mostly been considered on a country-wide scale (Duclos et al., 2004; Foster & Wolfson, 2010). We believe these issues are as important within teams.

Our empirical analysis suggests that the standard deviation and polarisation of wages within a team are the two most important elements when determining the effects of wage dispersion on performance, and that they should be considered simultaneously. Holding the other constant, a high standard deviation of wage seems to encourage effort from players and increase team output, while conversely, a high polarisation of wage seems to have a negative impact through decreased cooperation and increased conflict. As we explain in Chapter 3, while some changes in the wage structure will influence these measures in the same direction, others can increase the standard deviation of wages while decreasing wage polarisation. Our results therefore suggest a wage structure which is dispersed, but avoids clusterings at the extremes, will maximise performance.

1.2 National Diversity

National diversity in the workplace is becoming increasingly common. As globalisation continues this current trend, it is relevant to consider the effects of this national diversity on production within teams. The effects on country-level outcomes such as growth, governance and the incidence of civil war have received attention in the literature (Alesina et al., 2003; Collier & Hoeffler, 1998; Easterly & Levine, 1997; Horowitz, 2001). However there has been less focus on micro-level outcomes, such as the productivity of teams within the workplace. This investigation will explore whether the productivity of workplace teams is influenced by the level of national diversity within them.

Nationally heterogeneous teams may experience some additional difficulties which do not trouble nationally homogeneous teams, such as language barriers and conflicting cultural norms, which can hinder the production process (Ingersoll et al., 2017; Lazear, 1999). However, opportunities may also arise in the form of cross-cultural learning; generating positive spillover effects which can improve productivity (Hansen et al., 2006; Lazear, 1999). And naturally, teams which are open to becoming nationally diverse are able to recruit from a larger pool of potential workers which may offer them a competitive advantage (Kahane et al., 2013). Further research is required to determine the extent to which these potential effects influence performance; whether the positive or negative aspects dominate; and whether national diversity can be managed to maximise the benefits, while minimising any costs.

This investigation will explore whether the level of national diversity within a team significantly affects its performance, and also whether different types of diversity have differing effects. Diversity can be measured in several ways; by considering the structure of diversity,

the cultural distance, or a combination of the two. We add to the debate on national fractionalisation and polarisation measures, giving further empirical evidence which suggests the measures are distinguishable and capture different phenomena. In addition, we find them to have opposite performance effects to one another when regressed simultaneously, which has unique implications for how the structure of national diversity affects team performance. We also consider measures of linguistic distance, which are commonly used in the literature to proxy cultural distance. We attempt to improve these proxies by assuming that a nation's cultural makeup depends on more than just its most common language, but also on its other commonly spoken languages. Lastly, we show how these measures can be combined to capture both the structure of, and cultural distance within, a team's national diversity.

1.3 Sports Labour Markets

Kahn (2000) describes sports as a labour market laboratory which offers a unique opportunity to test labour market theories, stating "[t]here is no research setting other than sports where we know the name, face, and life history of every production worker and supervisor in the industry." (Kahn, 2000, p.75). Personal and sensitive information, such as wages and nationalities of participants in sports labour markets, are widely available and in greater detail than in other labour market settings. This provides an ideal opportunity for our investigation. We utilise a comprehensive dataset which covers every player to compete in the National Basketball Association (NBA) between 2007-18. This includes their wage, nationality, performance statistics and other relevant information such as their level of experience.

There is a wide range of literature analysing sports labour markets using economic principles. The high wages commanded by sportspeople can be explained by their large 'personal scale of operations' (Mayer, 1960; Rosen & Sanderson, 2001), which allows a relatively small number of workers to service a vast audience without diminishing the product. In the past, sportspeople were not always able to demand a wage equal to their marginal productivity due to bargaining restrictions (Scully, 1974), but since the 1970s, free agency has become the norm and there has been an alignment of wages and marginal productivity, as predicted by Rottenberg (1956). More recent work has looked to estimate a sporting production function using in-game statistics (e.g. Carmichael & Thomas, 1995; Carmichael et al., 2001; Schofield, 1988), while others look at the efficiency with which teams employ their resources (e.g. Coates et al., 2016; Dawson et al., 2000; Kahane et al., 2013). This study better fits the latter description.

There has been some discussion about whether sports teams are profit or utility maximisers (Sloane, 1971). Neale (1964) also notes the 'peculiar' economics of sports, such as the Louis-Schmeling Paradox. This is the observation that while firms in other industries could maximise profits by monopolising the industry, the profitability of sports teams benefits from close competition with their rivals. If a sports team were to 'monopolise' their league, this would decrease the uncertainty of outcome (see Fort & Maxcy, 2003; Jennett, 1984; Peel & Thomas, 1988; Zimbalist, 2002) and could discourage spectators and decrease profitability. Like many North American sports leagues, teams in the NBA operate under a salary cap to encourage competitiveness and increase the uncertainty of outcome. While there are some exceptions, making it what is known as a 'soft cap', the expenditure of teams is still restricted to maintain competitiveness within the league. We therefore consider teams to be win maximisers within league-imposed restrictions which maintain competitiveness. The salary cap can be thought of as a budget within which the team attempts to maximise their performance.

A recent NBA Press Release highlights the level of national diversity within the NBA, and also hints that teams containing international players tend to perform well (NBA Global, 2018). The opening rosters of the 2017-18 season featured 108 international players from a record 42 countries and territories, while the 2018 playoff rosters contain a record 62 international players from a record 33 countries. This continues a recent trend of increased national diversity in the NBA, and is the fourth consecutive season that at least 100 international players have featured in the opening day rosters and that all 30 teams had at least one international player. So, with teams developing a taste for diversifying the national composition of their rosters, this may signify that teams see some potential benefit in doing so.

Writing for the New York Times, Abrams (2010) notes a phenomenon in the NBA which highlights a mechanism through which national diversity can improve production methods. The so-called 'Euro-Step' is a technical move which cleverly uses the two steps allocated to players after dribbling. The player drives past the defender before taking a lateral step in the opposite direction. It comes to the brink of being a travelling violation, which initially confused some defenders and helped the attacker gain a competitive advantage. The traditional focus of American players and coaches had been on athleticism, but the influx of European players has challenged this status-quo. This move, and variations of it, are now commonly used in the NBA and this is not the only example. International players such as Argentina's Manu Ginóbili and Luis Scola, and Germany's Dirk Nowitzki, each have their own signature moves and other European-style moves are becoming increasingly popular in the NBA.

While sports labour markets adhere to many of the same economic principles and assumptions of other labour markets, one may question the extent of the external validity of the results from an investigation such as this one to general labour markets. Despite the similarities, there are unique aspects to sports labour markets which could threaten the applicability of the findings to other contexts. Perhaps the most obvious of these is the high wages of sports stars compared to other industries. A little less obvious is the negotiation between players (or agents) and teams, rather than an advertised position with an associated wage/salary. Even when some negotiation is possible in other industries, it is generally to a lesser extent than in sports labour markets. This is linked to the high-level and creative nature of the skills required to become a professional athlete. Also, as the best rewarded professional basketball league in the world (Fraschilla, 2017; Statista, 2020), the NBA has the ability to attract the best foreign talents to the league. 2019 data from the United States Census Bureau (2020) suggests that roughly 14% of the general US population are foreign-born. Though somewhat comparable, this figure is higher in our dataset at around 20%.[2] The findings of this investigation will therefore be most applicable to highly-skilled, highly-paid, creative professions, perhaps with a higher than average proportion of foreign-born workers, which work in teams to produce output. Examples of such professions could include academic departments within universities, teams of software engineers and trading teams. While more general applicability may be plausible, the threats to external validity discussed here should be taken into consideration.

The NBA offers a setting with clearly defined teams, which exhibit variation in both wage dispersion and national diversity. The data at both the team- and player-level is available, detailed and verifiable. These teams also have clearly defined performance objectives and

[2] Both in our dataset and the US population in general, the proportion of non-natives (foreign-born) is gradually increasing over time (United States Census Bureau, 2011, 2020). The figure for the proportion of foreign-born players in the final season of our data (2017-18) is 21.3%.

detailed performance statistics, which are again available, detailed and verifiable. While the performance outcomes of the teams are easily measurable, this is not as easy for the individual players, as we discuss in the next section.

1.4 Measuring Performance

This study investigates the effects of wage dispersion and national diversity, both separately and simultaneously, on team- and player-level performance. While the ultimate aim of the team is to maximise their overall team productivity, an investigation into the effects on individual players adds nuance and an appreciation of how these overall effects manifest within the team. We may question whether a given wage structure encourages effort from all players or whether its positive effects on some players compensates for the negative effects on others. Similarly, we may ask whether cross-cultural learning is equally beneficial to all players, or whether larger or smaller national groups benefit the most.

There appears to be a gap in the current academic literature on these issues, with most studies solely focussing on the aggregate team-level effects. We suspect that the main reason for this is the difficulty in measuring individual player performance within teams. Team-level performance is usually easily measurable, especially within sports which set out clear criteria for winning contests and ultimately being successful in competitions. However, success at the team-level is determined not only by the individual ability of each player, but requires cooperation between team members to achieve. This interdependence of production highlights the importance of both player- and team-level factors, and particularly the interaction between the two. It also presents a number of issues to overcome in order to study individual player performance within teams, which we aim to overcome in investigation.

While we believe we are the first to investigate the effects of wage dispersion and national diversity on individual players within teams, current work helps us to overcome the issue of measuring player performance within teams. This work has been conducted partly within the academic literature, but also to a significant extent by sports analysts and statisticians outside of academia. We utilise a number of these measures in our investigation, as well as applying multilevel modelling (Gelman & Hill, 2006) to account for player- and team-level factors simultaneously.

1.5 Structure of the Investigation

The remainder of this thesis is structured as follows. Chapter 2 reviews the current literature relating to the effects of wage dispersion and national diversity on performance. We discuss the theories, general evidence and then sport-specific evidence of each topic in turn. This chapter concludes by discussing how performance measures have developed over time. Next, Chapter 3 discusses the construction of our variables of interest, specifically standard deviation, fractionalisation and polarisation for both continuous and discrete variables, and linguistic distance. We show how these measures capture distinct information about the distributions of wages and nationalities within a team, and how they compare to one another and other dispersion measures. Chapter 4 begins by describing the dataset used, its structure and the relevant econometric techniques utilised for our investigation. It then describes the dependent variables selected to capture team and individual player performance, and finally discusses potential covariates which are controlled for in our analysis.

Next comes the three empirical Chapters. Firstly, Chapter 5 investigates the effects of wage dispersion at the team-level. Then Chapter 6 investigates the effects of national diversity, again at the team-level. At the end of Chapter 6, we also consider wage dispersion and national diversity simultaneously in the same regression model to test robustness and check for any interaction effects. Lastly, Chapter 7 investigates the effects of both wage dispersion and national diversity at the player-level, separately and then simultaneously. Each of these chapters begins with some descriptive analysis, before presenting regression results, followed by robustness checks. Each then ends with a discussion of our main findings. Chapter 8 concludes with an overall summary of our investigation.

Chapter 2

Literature Review

Chapter Overview

This literature review is split into three main sections. The first two discuss the literature relating to our two main topic areas, firstly wage dispersion and then national diversity. The relevant theories, as well as evidence from both general labour markets and then specifically within sport, are considered in turn. Within the current literature, most of the evidence on these topics considers the aggregate team-level effects, without much focus on how these effects manifest at the player-level. To our knowledge, there is little to no academic literature that looks to estimate the effects of wage dispersion and/or national diversity on individual player performance within teams. This seems to be due to the difficulties in measuring individual player performance within teams, caused by the interdependent nature of production in many team-based environments. Therefore the third and final section of this literature review is dedicated to discussing how performance can be measured. The primary objective of sports teams is to win games, but we also consider other measures of team performance which will later be used for robustness checks. This then leads into the discussion of how the individual contribution of players towards team performance can be measured.

2.1 Wage Dispersion

There are two main competing hypotheses regarding the effect of wage dispersion on team performance: Tournament Theory (Lazear & Rosen, 1981) and Cohesion Theory (see Levine, 1991). This section discusses these two opposing theories and some reconciliation attempts from the current literature which attempt to resolve the disagreement between the two. We argue that, although these reconciliation attempts are difficult to apply empirically, they enrich the discussion ex-post. Evidence in favour of each theory is presented and analysed with these reconciliations in mind. The theories are then discussed in the context of the sports market and an explanation is offered as to why it is not clear which will dominate. Evidence from the sports industry is then presented.

2.1.1 Theory

The work by Lazear and Rosen (1981), known as Tournament Theory, explains how higher levels of wage dispersion can induce effort and therefore increase the productivity of workers. Lazear and Rosen (1981) were among the first to analyse rank-order payment schemes, in which wages are linked to the relative ranking of workers, rather than absolute output. Their mathematical model shows that, under risk-neutrality, firms will want to increase the level of wage dispersion, ceteris paribus, to induce worker effort due to increased incentives to achieve the highest rank possible. The extent to which the level of wage dispersion spreads will be limited by the increasing marginal costs of effort for the workers. Due to these increasing marginal costs, there will exist some equilibrium spread. Lazear and Rosen (1981) suggest that, even under risk aversion, there are circumstances where workers prefer rank-order payment systems. An example of such a condition is that when workers have larger endowments, they are more likely to prefer rank-order payment systems.

The opposing theory comes from Levine (1991) who qualitatively discusses the advantages of a compressed wage structure in encouraging cohesion within teams. These advantages come in the form of increased cooperation and a lack of sabotage in the workplace. Also, cohesive teams will effectively sanction free-riders when bonuses are based on team production – such a bonus scheme will also decrease wage dispersion in itself. The compression of wages towards the middle is discussed in terms of paying an 'efficiency wage' to low-earners, meaning they earn above their marginal productivity. Of course, to be sustainable, this efficiency wage must be funded by paying top-earners below their marginal productivity. Levine (1991) argues that if all firms were to follow such a wage structure then efficiency would be improved. However, this efficiency wage is underprovided in competitive markets. Firms wishing to recruit the top talents in their industry are unwilling to fund the required efficiency wage by paying their top-earners below their marginal productivity, for fear of losing the top-talents to firms who do not employ such a strategy. This situation is best described as a Prisoner's Dilemma, where the social optimum would be for all firms to offer more egalitarian wage structures, but there is the incentive to 'cheat' to secure the top talents.

To evidence these claims, Levine (1991) points to the Japanese labour market which tends to be egalitarian in terms of wage and status, especially in the first few years of employment at a firm. Labour mobility is low due to a spiralling effect of firm-specific investment in employees, which encourages and is encouraged by longer tenures. Although this low labour mobility can cause inefficient job-worker matches, this is more than compensated by cohesive efficiency amongst the workforce. Though a further point considered by Levine (1991) is that more

cohesive workforces are disadvantageous to firms when the goals of employees oppose those of the firm. Examples include workers collectively bargaining to increase wages and assisting each other to pilfer.

Overall, there appears to be a challenge between inducing worker effort through a dispersed wage structure, without damaging the harmony within the team which brings its own, often intangible, benefits. While no definitive consensus has been reached in the literature, there have been some attempts to reconcile these competing hypotheses regarding the effects of wage dispersion. Firstly, the Hawks and Doves model (Lazear, 1989) considers that when production is interdependent and depends on relative rankings, there are two ways for workers to increase their pay: either through increased effort, which will have positive effects on production, or through sabotage, which will have negative effects on production. The main feature of this model is that workers are of different types, and these types can be distinguished by their costs of sabotage. Hawks have lower costs of sabotage while doves' costs are higher; these costs could be physical, moral or time-related. Therefore, when wage dispersion is high, hawks are likely to compete non-cooperatively, while doves compete cooperatively. This analysis indicates that personality is an important consideration in the employment process, but information on personality types is not often found in empirical datasets. While the personality or type of each worker can be induced by their reaction to wage dispersion, this then suffers from confirmation bias. Also, as the costs of sabotage to workers may be considered subjective to each worker, and liable to interact and change in different situations, it becomes very difficult to use this theory to make predictions and give recommendations ex-ante. Consequently, while this theory is useful for the discussion of results ex-post, its ad hoc interpretation is problematic when looking to develop a general model.

Another reconciliation attempt comes in the form of Equity Theory. Although this theory has its roots in social psychology (see Leventhal, 1976), it has been used in economics to address the issue of wage dispersion (see Trevor et al., 2012). The theory suggests that workers accept differences in wage compared to their colleagues, only if it is justified by some criteria, usually linked to productivity. It suggests that less productive workers do not expect wage equality with their more productive colleagues, but do expect it with their similarly less productive colleagues. Equity Theory is a promising area of research, but it too is not without its problems. The existence of spillover effects means that players' productivity can vary dependant on the behaviour of their teammates and opposition (Gould & Winter, 2009; Kendall, 2003). Due to the complexity and interdependence of production, in many cases it is difficult to disentangle and assign marginal productivity easily and accurately. For example, low productivity may be an indication of low-skilled workers or of high levels of sabotage between workers. With complementarities between workers, there may be disputes between who deserves credit for additional joint production. Even for academics and firms, these marginal productivities are difficult to measure and monitor, so it may be unreasonable to assume that workers have full information. This would involve knowing their own productivity and worth, as well as those of their colleagues. This would surely be a prerequisite for workers to form justification criteria for higher wages and to accept them. Misinformation and mismeasurement could cause conflict, which may be the case in reality, on occasions when workers and firms disagree on the wage level which is justified. While Equity Theory has much potential, especially in settings where marginal productivity is easily observable, it becomes less precise and more subjective in interdependent settings.

Therefore, while these reconciliation attempts are interesting for the discussion of the effects of wage dispersion, their applicability in an empirical setting is questionable. As personality types are usually unknown, and marginal productivities difficult to measure accurately, creating the relevant variables is often not possible. Also, depending on the method of construction, these variables could produce results which suffer from confirmation bias. That is, if we use the relationship between wage and performance to define a variable, then we should not be surprised when this variable shows such a relationship between wage and performance. Despite their lack of practical applicability in an empirical setting, there is a strong argument that these issues are relevant and they certainly enrich the discussion. Therefore, the empirical results of our investigation will be discussed in terms of Tournament Theory and Cohesion Theory, but also with these reconciliations in mind.

2.1.2 General Evidence

There are several examples in the literature of empirical evidence in favour of Tournament Theory (see: Eriksson, 1999; Heyman, 2005; Lee et al., 2008; Main et al., 1993). Furthermore, there is some evidence that the positive effects of wage dispersion are stronger when a high degree of monitoring is employed and for blue-collar workers as opposed to white-collar workers (Lallemand et al., 2004). White collar workers tend to be intrinsically motivated so prefer a more egalitarian wage structure compared to blue-collar workers. Lallemand et al. (2004) also note that, as per the predictions of Lazear (1989), white-collar workers can be expected to display more hawkish behaviour. This is based on the intuition that hawks are better able to win contests, and are therefore more likely to reach the higher levels of employment hierarchies. While contests can be won purely by productive efforts, hawks, who can sabotage effectively, are the more likely victors. This is certainly the case if characteristics such as productive capacity and work ethic are otherwise identically distributed among hawks and doves.

Additionally, there are also a number of empirical studies which support Cohesion Theory, finding a negative relationship between wage dispersion and performance (see: Bloom, 1999; Fredrickson et al., 2010; Grund & Westergård-Nielsen, 2008). A particularly interesting example is provided by Pfeffer and Langton (1993), which investigates the links between wage dispersion and productivity amongst academics. This study suggests that wage differences are tolerated if they are justified by certain criteria (such as productivity in terms of academic publications and experience) in line with Equity Theory. However productivity, satisfaction and collaboration are all improved in universities with lower wage dispersion. This study therefore supplies some evidence in support of both Cohesion Theory and Equity Theory. The authors advocate linking wage to performance, but also applying an element of compression by making the increases in wages less than the increase in marginal productivity. However, the optimal degree of compression is not clear.

In the broader conflict literature, a distinction is made between vertical and horizontal inequalities, and their impacts on economic development. Vertical inequalities occur between individuals, whereas horizontal inequalities occur between groups, often split along ethnic/cultural, political, religious or socio-economic lines (Østby, 2008). Empirically, there is some suggestion that horizontal inequalities (e.g. Cederman et al., 2011; Stewart et al., 2008) are a more potent trigger for civil conflict than vertical inequalities (e.g. Brzezinski, 2013; Cramer, 2003). Additionally, Stewart et al. (2008) suggest that, when addressing horizontal inequalities, indirect

policies aimed at helping numerous deprived groups are preferable to direct policies as the latter may be provocative and detrimental to social cohesion. Østby (2008) demonstrates that income polarisation can be used as an alternative to horizontal inequality, along social lines. And, indeed, there is also evidence that income polarisation may have detrimental effects on social cohesion and economic growth (Brzezinski, 2013; Esteban & Ray, 2011; Ezcurra, 2009). Investigating whether wage polarisation has a similar effect on the productivity of smaller work groups is one of the contributions of the current study. Section 3.1 explains the theory behind and construction of the measures of wage polarisation used in this investigation.

As the previous discussion shows, the effects of wage dispersion seem unclear, both theoretically and empirically. However, thus far the evidence presented has been for general labour markets. Despite sports labour markets behaving like general labour markets in many key aspects, they remain a unique industry. Next, we will discuss evidence from the sports industry to investigate whether any of these theories dominate in this particular setting.

2.1.3 Sports Evidence

Within sports teams, players compete for coveted starting positions while facing competition from their teammates, players from other teams, players from other leagues and from the minor leagues. The player who wins the contest for the starting position is paid more than the back-ups, with the upper bound of the starter's salary constrained by profits or the available wage budget of the team. The lower bound of the back-ups is constrained by opportunity costs. The wage structures implemented in the sports industry should be considered rank-order tournaments (Scully, 1995). As such, small advantages in marginal productivity can result in substantial increases in pay. This may seem to initially indicate that Tournament Theory is relevant to characterise these markets.

Ideally, for a strict application of Tournament Theory, we would observe effort directly rewarded with a prize. Unfortunately, such instances occur rarely, if at all, in empirical data. One may argue that when a player is issued with a new contract this condition is met, but these instances are still problematic. The timing of when a player increases his effort in anticipation of a new contract is ambiguous and not defined. It could be considered to start immediately after signing the previous contract or could be a much shorter time just before that contract expires. It would also not be clear whether a player has any prior certainty that his effort will be rewarded with a large wage increase in a new contract. Unless there are stipulations to be triggered in an existing contract, the player must exert effort before the prize is offered. If a team generally offers a more egalitarian wage structure, this suggests a smaller prize for 'winning the tournament', whereas a higher wage dispersion would suggest larger prizes. The level of wage dispersion in a team therefore indicates potential prizes, hence the rationale for analysing measures of current wage dispersion in this way using Tournament Theory.

Another important factor is the level of cooperation needed amongst players in the starting positions, and the existence of spillover effects, both on-field and in training. Tournament Theory assumes that production is independent and additive, but this is not the case in the sports industry (or indeed many other industries). There are therefore strong indications that Cohesion Theory could play a significant role in the productivity of sportspeople. Furthermore, the competitive nature of reaching the top-level in sport may suggest that top-level sportspeople display hawkish behaviour to win their contests. However, the interdependence

of production encourages dove-like behaviour as there is likely a high personal cost to sabotaging teammates. As previously mentioned, this interdependence of production also makes measuring marginal productivities difficult, which affects both the practicality and applicability of Equity Theory. It is not clear which effects will dominate, and this is reflected by some mixed results in the literature.

Evidence in favour of both Tournament Theory (Pierce, 2017; Simmons & Berri, 2011) and Cohesion Theory (Depken, 2000; Depken & Lureman, 2018) can be found in the sports literature. There are also suggestions of a quadratic relationship between wage dispersion and performance, but again the evidence here is mixed. Vasilescu (2008) finds evidence of an inverted U-shaped relationship, with a turning point at the upper end of the distribution, suggesting Tournament Theory is more relevant. While Coates et al. (2016) also find an inverted U-shaped relationship, their turning point is to the left of the feasible region of their dispersion measure, suggesting Cohesion Theory is more relevant. There is even some evidence that the relationship is actually U-shaped (Franck & Nüesch, 2011). This would imply that teams should either have low dispersion to encourage teamwork or high dispersion to attract the superstars, as teams with average levels of dispersion will benefit from neither. Franck and Nüesch (2011) support their theory by supplying evidence that higher wage dispersions are associated with more individualistic playing styles, while lower wage dispersions are associated with more cooperative play.

Frick et al. (2003) analyse the four major leagues in America and again find mixed results, with positive effects for the NBA and National Hockey League (NHL), but negative effects for the National Football League (NFL) and Major League Baseball (MLB). Furthermore, the effects only show statistical significance for the NBA (positively) and NFL (negatively). There is an argument that the relationship depends on the sport analysed and the level of cooperation/interdependence that each requires. However, we still see some disagreement between studies which analyse the same sports (e.g. Depken (2000); Vasilescu (2008) - baseball; Depken and Lureman (2018); Frick et al. (2003) - hockey). Additionally, we see some evidence that the definition of the team matters. Bucciol et al. (2014), using data from the Italian soccer league, find evidence that dispersion within the roster can have positive effects, but dispersion amongst the starters can be negative.

It should also be noted that the way wage dispersion is measured is not consistent amongst these studies. The most commonly used measure is the Gini Coefficient (e.g. Coates et al., 2016; Franck & Nüesch, 2011; Frick et al., 2003; Simmons & Berri, 2011), with the Coefficient of Variation also considered for robustness (e.g. Coates et al., 2016; Franck & Nüesch, 2011). But other measures are also considered, such as the standard deviation (Pierce, 2017) and a Herfindahl-Hirschman index (HHI) to capture wage disparity (Depken, 2000; Depken & Lureman, 2018). While these measures broadly capture similar effects, the differences in measurement may be important. Also, in the literature to date, these wage dispersion measures have been considered individually. Considering multiple measures of dispersion simultaneously in the same model may also reveal more information about how changes in the wage distribution can affect performance. Furthermore, we argue that an important measure of dispersion is so far missing from the literature: wage polarisation. Section 3.1 discusses the development of polarisation measures, what they measure and how they differ from other measures of inequality.

2.2 National Diversity

Much like the literature on wage dispersion, in addition to whether a relationship exists between national diversity and performance, there is also theoretical and empirical argument about whether the expected effects should be positive or negative. However, especially when compared to issues of wage dispersion, the literature concerning national diversity is scarce. This is partly due to the relative novelty of the issue, but also due to measurement difficulties (particularly in terms of a cultural distance metric) and availability of personal data relating to nationality. Naturally, the sports industry offers a useful vehicle to overcome these problems but, outside of this, the general evidence comes primarily at the country-level. In the following, we again discuss the theoretical implications of national diversity, before presenting some general evidence and then some evidence from the sports industry.

2.2.1 Theory

Nationally heterogeneous teams benefit from a range of cultural backgrounds which aids the 'search for best practice' (Lazear, 1999). However, language barriers and differences in social/procedural norms may disadvantage these teams, affecting performance negatively through reduced cooperation and increased conflict. The simplest way to measure national diversity is to count the number of nationalities within a team. However, the effects of national diversity can also be affected by the structure of this diversity within a team. It may not simply be about the number of nationalities within a team, but also the proportion of the team made up by each of these nationalities. To capture this, measures of fractionalisation can be utilised. These measures take a value of zero for homogeneous teams, and can theoretically take a value up to one for a team with infinitely many players each of a different nationality. Importantly, the measure is not only increasing in the number of nationalities within a team, but also in how equally dispersed team members are into these groupings.

Furthermore, fractionalisation may not be the only relevant measure concerning the structure of diversity, especially when considering potential conflict caused by national diversity within a team. Montalvo and Reynal-Querol (2005) argue that polarisation is a more appropriate measure to capture the possibility of conflict than fractionalisation. While both measures will equal zero for homogeneous teams and are initially highly linearly correlated at lower values of both measures, polarisation is maximised when a team consists of two groups of equal size. Adding more nationalities from this point will increase fractionalisation while decreasing polarisation, meaning the measures are negatively correlated at higher values of fractionalisation. For intermediate values, there is no strong correlation between the two measures. Montalvo and Reynal-Querol (2005) also shows how fractionalisation can be derived from a discrete Gini Coefficient by replacing the absolute difference in incomes with a binary measure, equal to zero if members are of the same national group and one if members belong to different groups. Similarly, basic discrete polarisation measures can be derived in a similar manner from the general form of continuous polarisation axiomatically derived by Duclos et al. (2004).[3]

A downside of both these measures is that they treat each national group as being equally different to every other group. In reality, some nationalities may be considered more culturally similar to others. Some studies consider cultural distance, which is often proxied by the

[3] More information on the derivation and relationship between these measures of fractionalisation and polarisation are given in Section 3.2.

linguistic distance between the primary languages of each country. Languages are regarded as a good proxy of culture as they hint at shared histories amongst cultures, as well as influencing the cultural mindset of nations (Chen, 2013). By considering the structure of each languages' linguistic tree[4], it is possible to calculate the proportion of branches shared with other languages and, from this, create a measure of linguistic distance between zero and one. When nations share the same language they are considered to be culturally identical and have a linguistic distance of zero. When the languages of two nations share no common roots, then their linguistic distance will be one. Many will share some linguistic roots, but not all, giving them a linguistic score between zero and one. Another strategy is to use assimilation programs such as the Automated Similarity Judgment Program (ASJP) (Bakker et al., 2009), which looks for similarities between commonly used words in different languages and generates a score between 0-100. Similar languages will receive low scores and highly different languages will receive high scores.

These measures of linguistic distance can also be incorporated into fractionalisation measures.[5] A number of studies use measures of linguistic fractionalisation using the linguistic tree method of calculating linguistic distance (see: Baldwin & Huber, 2010; Desmet et al., 2009; Fearon, 2003, - though note that these are non-sport, country-level examples).

2.2.2 General Evidence

In the literature, fractionalisation measures are most commonly used to look at ethnic and/or linguistic fractionalisation[6] at the country-level, and the effects this has on civil wars and economic performance. While several studies find ethnolinguistic[7] fractionalisation to have undesirable effects on trust, government performance, civil wars and other economic outcomes (Easterly & Levine, 1997; La Porta et al., 1999; Leigh, 2006), others argue that the effect may be non-monotonic (Collier & Hoeffler, 1998; Horowitz, 2001). Few studies have investigated the effects of national fractionalisation within work teams, which is one of the contributions of our study.

Empirically, Montalvo and Reynal-Querol (2005) give evidence that ethnic polarisation, not ethnic fractionalisation, is a significant determinant of civil war; though they do still find that ethnic fractionalisation is associated with low economic growth. Baldassarri and Bearman (2007) argue that, though this is a rare occurrence, some political issues will individually 'take-off' when they are relatively more polarised than others, creating ideological polarisation. In addition, important and non-trivial social issues can become polarised as people selectively seek to discuss these issues with like-minded people, thereby fuelling further polarisation. Again, few studies have investigated the effects of national polarisation within work teams, as this study does.

2.2.3 Sports Evidence

Brandes et al. (2009) look at the number of nationalities within teams in the Bundesliga (the top German soccer league), but do not find a statistically significant effect overall. They additionally look at the number of nationalities within smaller sub-groups of the team, namely

[4] Using a dataset such as the one developed by Ethnologue (Lewis et al., 2014)

[5] The process of doing this is detailed in Section 3.2.

[6] Alesina et al. (2003) find that ethnic and linguistic fractionalisation, more so than religious, are important determinants of economic success, while Fearon (2003) finds high, but not perfect, correlation between measures of ethnic and linguistic fractionalisation.

[7] Some projects have combined these aspects of ethnicity and language to create ethnolinguistic groups (e.g. World Christian Encyclopaedia, Encyclopaedia Britannica and Atlas Nadorov Mira).

the defense, midfield and strikers. A statistically significant negative effect on performance is found only for the number of nationalities in the defensive block, leading the authors to conclude that the effects of national diversity depend on the nature of the underlying tasks of sub-groups within a team.

Using National Hockey League (NHL) data, Kahane et al. (2013) sort players into five major geographical groups (one North American and four European), based on their country of birth. They then create a Herfindahl-Hirschman index (HHI) based on these groups, and another measure which considers the relative share of European players within the team. Controlling for other factors, they find positive coefficients for both variables and conclude that teams can benefit from including European players alongside Americans, but suggest that the Europeans should be from the same country to limit communications problems. Though they do not explicitly use the measures, these results seem to suggest that polarisation between Americans and a European group is preferable, while high fractionalisation is undesirable.

Carried out parallelly to our study, a working paper by Addesa et al. (2017) also considers national diversity through measures of fractionalisation and polarisation. This study is conducted with data from Serie A (the top Italian soccer league) at the match-level. Their findings suggest that both fractionalisation and polarisation have negative effects on team performance, whether included separately or simultaneously. They claim that in their results, "[t]here is no notable difference between fractionalisation and polarisation, as both indices have the same direction" (p.12), and that the interpretation of including them simultaneously is ambiguous as the measures are correlated. But this correlation is by no means linear, with positive correlation at low fractionalisation, no correlation at mid-values and negative correlation at high fractionalisation.

Another recent working paper by Brox and Krieger (2019) considers national diversity among on-field players, at the match-level, using data from the Bundesliga (the top German soccer league). The study captures national diversity only through measures of fractionalisation, but also investigates whether a quadratic relationship with performance may exist. They find evidence of an inverted U-shaped relationship, and that the optimal level of diversity may be task-specific and depend on the importance of verbal communication.

Linguistics distance is considered by Ingersoll et al. (2017) in their investigation using data from the Champions League (the top-level European international club tournament). They find that increased linguistic distance within teams has positive effects on team performance. As a robustness check, a measure of fractionalisation which considers each language as a different group is also considered. Again, a positive effect between diversity and performance is found.

Overall, there are a number of different measures used in the literature to capture national/cultural diversity. The main ones relate to measures of fractionalisation and polarisation, which capture the structure of national diversity within teams, and measures of linguistic distance, which proxy cultural distance. Some studies have combined these two measures, but this has only been considered at the country-level. There is scope to further investigate and disentangle the effects of fractionalisation and polarisation, improve the proxies of cultural distance, and combine these measures to explore the effects on performance within sports. These are some of the aims of this investigation.

2.3 Measuring Performance

In team sports, much like other team-based environments, output is much easier to measure at the team-level. Players are organised into teams with the ultimate aim of producing team output, which is measured with clear criterion for success. The most basic on-field objective of sports teams is to win games. This can be further distilled down to elements such as points scored or points allowed, and particularly the efficiency of these two factors, to add further nuance. However, this team-level production usually requires not only individual inputs, but a cooperation of complementary inputs, for high performance to be achieved. In addition, not all aspects of this teamwork are equally visible, or their contribution towards the overall performance of the team easily quantifiable. It is exactly this interdependence of production which makes the effects of team-level characteristics on both team- and player-level performance so interesting. Even with a plethora of individual (and team) statistics related to various aspects of sports, it is difficult to apportion team performance to individuals within teams. Because of this, and the lack of literature on the effects of our variables of interest on individual player performance, much of this section of the literature review will focus on evaluating and quantifying player performance.

Despite the aforementioned difficulties, there have been a number of attempts made at creating an index of individual player performance. Many of these are created by sports analyst outside of academic peer-reviewed journals, alongside some academic literature on the subject. Of the non-academic work, Berri and Bradbury (2010) note the lack of formal vetting through an expert peer-review system and discuss some of the conflict between so called '*metricians*' and academics. However, some interesting and insightful work has been done outside of the academic process, though it is oftentimes unstructured and developed a little haphazardly over time. Kubatko et al. (2007) provide a useful overview which brings together and summarises these two strands. While the focus is generally on analysing the importance of statistics in determining performance at the team-level, it is mentioned how some of the measures discussed can also be calculated at the player-level, as well as discussing some specific player performance measures. Understanding how team performance (in essence, win percentage) is generated at the team-level is important in order to understand how individual players contribute and how their performance can be measured. Both of these elements are discussed in this section.

2.3.1 Team Performance

While the success of a team is primarily judged by their ability to win games and therefore their winning percentage, there are other statistics available which add further detail on how these wins are achieved. When analysing basketball statistics, many authors (e.g. Berri et al., 2006; Kubatko et al., 2007) stress the importance of considering the number of possessions a team has, and then how efficient they are at converting their possessions into points. A distinction is made between possessions and plays, in that the end of a possession results in the other team possessing the ball, while a play may end, and therefore a new one starts, with the same team in possession. This can occur when a play produces a missed field goal attempt which is then recovered via an offensive rebound, for example. The play ends with the field goal attempt, but the possession is extended via the offensive rebound because the same team remains in possession. Accordingly, a possession may consist of multiple plays but each play is a part (or can be the entirety) of a single possession. This distinction is important, firstly to understand why much of the work in this area assumes equal possessions for each team

within a game. By definition, a team in possession can only gain an additional possession after their opposition has had a possession.[8] And secondly to understand the effects that different statistics, such as rebounds, have on possessions (not plays).

Possessions are not explicitly recorded, and while they can be tediously counted using play-by-play game logs, they can be also estimated reliably from other game statistics. Slight variations of similar equations are used to estimate possessions by different authors (e.g. Hollinger, 2005; Oliver, 2004); but Kubatko et al. (2007) show that these equations produce similar results[9] and they present a generalised equation without imposing the restrictions of other authors, estimated via regression analysis. Berri et al. (2006) also discuss the importance of possessions and show how possessions acquired and possession employed can be calculated and how, naturally, the two are equivalent. A team will employ all of the possession they acquire in some manner. How possessions are estimated by these various methods is discussed further in Section 4.3.1 and Appendix 4.A.2.

An appreciation of possessions is needed to understand the importance of efficiency in basketball, as well as the influence of tempo. Teams may choose to play at a high tempo which gives them more opportunities to score points, but also offers the same opportunities to their opponents. Berri (1999) notes that the tempo of a team is oftentimes a deliberate strategy by coaches and, therefore, in his measure of player performance chooses to assign tempo effects to players in a team by the proportion of minutes played rather than their individual scoring attempts. He argues that players should not be punished for following coaching instructions and/or shouldering point scoring responsibilities due to being a relatively efficient shooter. And equally, or even more convincingly, players should not be rewarded for avoiding potential scoring opportunities because they are not efficient shooters.

Commonly used measures of general team performance in terms of offensive and defensive efficiency are *Offensive Rating* and *Defensive Rating* respectively. These measures are simply points scored/allowed divided by the possessions of the team/opponent (or, equivalently, possessions employed/acquired). This is often multiplied by 100 to give a magnitude which is easier to interpret. Due to their construction, a higher Offensive Rating and a lower[10] Defensive Rating indicate better team performance. Kubatko et al. (2007) note that there tends to be very little correlation between these measures. It is not necessarily the case that those teams which score more efficiently also concede more per possession. Naturally, teams tend to cluster around moderate values of these ratings, but any of the combinations of extremes of the two variables is not unreasonable.

Some sections of the academic literature and sports analysts community break these overall team performance measures down into more statistic-specific measures which they refer to as the *Four Factors*. These are *Turnovers per Possession*, *Effective Field Goal Percentage*, *Free Throw Rate* and *Offensive Rebounding Percentage*. We do not see the gain of breaking the efficiency measures down further at the team-level, and at the player-level we use the statistics

[8] In the NBA, the team that wins the initial 'tip off' gains the first possession of the game and also gets the first possession of the fourth quarter. The opposing team gets the first possessions of the second and third quarters, so each team starts two quarters in possession.

[9] Kubatko et al. (2007) find a minimum correlation of 0.95 between the various estimates and actual possessions using their data.

[10] As noted in our later analysis, we chose to multiply Defensive Rating by minus 1 for our investigation, so that better defensive performance is indicated by higher values. We feel this is more logical and eases interpretations and comparisons between performance measures.

incorporated in these factors (and more) individually when constructing player performance measures. The reason for mentioning them here is for completeness and to demonstrate the prevalence of the idea of possessions and efficiency in the statistical analysis of basketball performance. These factors either account specifically for possessions, as in the first measure (*Turnovers per Possession*), for attempts in the next two measures (*Effective Field Goal Percentage* and *Free Throw Rate*), and rebounding opportunities in the final measure (*Offensive Rebounding Percentage*). All therefore make an appreciation of tempo and/or efficiency. Teams want to minimise *Turnovers per Possession* and maximise the other three. These factors are not all considered equally important and the relative importance of each has been estimated. Oliver (2004) suggests relative weights of 40% on *Effective Field Goal Percentage*, 25% on *Turnovers per Possession*, 20% on *Offensive Rebounding Percentage* and 15% on *Free Throw Rate*. Küpfer (2005) suggests roughly 45%, 27%, 14% and 14% respectively; while more recently Jacobs (2017) suggests roughly 46%, 35%, 12% and 7% respectively. There seems to be consistency in the relative ordering of the importance of these variables, but some changes in magnitude are seen in different studies over time.

2.3.2 Individual Player Performance

We use a number of player performance measures in this study, the construction and technical details of which are discussed in Section 4.3.3 and associated appendices. The purpose of the following is to give detail on how these measures have developed over time and comment on some of the intermediary measures, as well as other measures which we do not utilise in our analysis but feel are worthy of mention. Player performance measures can broadly be considered to capture either efficiency (i.e. production per minute or typically per 48 minutes - the length of a basketball game) or overall contribution (production over the course of a season); the merits of each will be discussed.

One type of individual player performance measures are known as *Linear Weights*, which simply add the positive contributions of players and subtract the negative ones. A simple example is the NBA Player Efficiency Rating, which attaches a weight of one to each statistic, adding the positive statistics and subtracting the negative ones. While the weighting system is simplistic, the measure does a reasonable job, as it accounts for efficiency by considering misses, rebounds, turnovers and other such statistics. However more sophisticated weighting systems have been employed for measures such as *TENDEX* (Heeren, 1988), *Points Created* (Bellotti, 1992), and the relatively more recent *Player Efficiency Rating* (PER) (Hollinger, 2005). It is this final measure that we consider in our analysis due to its relative novelty, with further discussion in Section 4.3.3 and Appendix 4.C.

While measures such as PER give a good insight into the individual performance of players, they do not translate obviously back into a measure of team performance such as win percentage. They are also prone to outliers. There are some examples of players playing very few minutes for a team (typically in unimportant matches at the tail-end of the season - informally known as 'dead rubbers') but contributing positively in this time. While it can be argued that this type of efficiency should be rewarded, measures such as PER can return a score which is many times larger than those which are generally considered to be top performers in the league.[11] This presents an issue of whether individual player performance should be judged

[11] Informally, we may say that it does not pass the 'laugh-test'. According to Oliver (2004), this is a common expression 'in the business' for when concepts seem sensible, but the results vastly contradict conventional wisdom. It is, of course, arguable whether conventional wisdom should trump statistical results, but ideally we should aim to be within reasonable bounds to satisfy both.

purely on efficiency, or whether a player's overall contributions to their team should also be considered. While the importance of efficiency has been stressed in this analysis, and by that of others, it is also reasonable to argue that better players will be picked by their coaches to play more minutes and in the toughest matches. So, alongside efficiency, players should also be judged on their consistency and overall contribution to the team.

One such measure of the overall contribution of players is *Win Shares* (WS). The idea of allocating wins to players in the form of win shares was pioneered by James (2001) with his application to baseball. The first element to discuss is the idea of Pythagorean winning percentage, which can also be attributed to James (2001). This is the observation that the link between points scored and points allowed can be used to estimate wins with a Pythagorean-like equation (see James, 2001; Kubatko et al., 2007, for further details). The basis of the Win Share methodology in baseball is to take the runs allocated to each player and convert these into a projection of wins using the Pythagorean method. This idea is further developed to apply to basketball by the website Basketball-Reference[12] using elements from Oliver (2004). Scoring and defending in basketball is considered more of a team effort than it is in the individualistic nature of baseball. *Win Shares* in basketball are therefore calculated using *Points Produced* and *Defensive Ratings* of each player as developed by Oliver (2004). Further details are again given in Section 4.3.3.

When considering individual player performance, *Plus/minus* statistics are also often used. These look at play-by-play data and calculate the team's points differential when a given player is in the game. This statistic is also considered in two variants *Net Plus/Minus* and *Adjusted Plus/Minus*. *Net Plus/Minus* measures the *Plus/minus* statistic when the player is in the game relative to when he is out of the game. This adds an extra appreciation of how well the team manages to play without the player. *Adjusted Plus/Minus* additionally accounts for the quality of the teammates that the player plays alongside and the opponents he plays against when he is in the game, as well as whether the game is played at home or away. The motivation for this is that regular starters (as well as deep reserves) often play a major proportion of their minutes together, so their scores may be influenced by the externalities of their teammates (and indeed their opponents). The idea of *Adjusted Plus/Minus* was pioneered by Jeff Sagarin and Wayne Winston, and originally called $WINVAL$ (Kubatko et al., 2007).[13] It was later expanded upon using regression techniques to incorporate other statistics as outlined by Rosenbaum (2004) and further iterations have been conducted since (e.g. Myers, 2013; Paine, 2012). The measure *Box Plus/Minus* (BPM) used in this study is one of the more recent iterations and should be considered an efficiency measure as it is weighted by minutes played. A further measure *Value Over Replacement Player* ($VORP$) is an adaptation of BPM which considers overall contribution. Both are collected from Basketball-Reference and discussed in further detail in Section 4.3.3 and Appendix 4.D.

An earlier paper by Berri (1999) does a good job of incorporating many of the aspects discussed in this literature review into a single measure of team performance. While this procedure is vastly simplified and improved by Berri et al. (2006), it is informative and interesting to discuss the earlier work and its development here. The resulting measure of player performance,

[12] https://www.basketball-reference.com/about/ws.html

[13] It should perhaps be noted $WINVAL$ is commonly thought to fail the 'laugh test' (Berri et al., 2006; Oliver, 2004), but also that the Dallas Mavericks started using it in 2000 with some success (Wasserman, 2013).

Wins Produced (WP), looked to improve on *Linear Weight* measures by utilising regression analysis and correct for three problems which he identifies from other similar studies as: the choice of functional form, utilisation of opponents' statistics, and the issue of team tempo.

Rather than using the Cobb-Douglas reduced form equations approach of preceding studies (e.g. Hofler & Payne, 1997; McCormick & Clement, 1992; Scott et al., 1985; Zak et al., 1979), Berri (1999) estimates a system of equations using statistics which contribute to both points scored and points allowed, while appreciating the collinearity between them. The equations include a number of statistics and ratios for both the team in question and their opponent, including points per shot, free throw percentage, offensive/defensive rebounds, assits-to-turnovers ratio, and field goal and free throw attempts. Marginal effects on win percentage (through points scored/allowed) were then calculated to decompose the ratios back into their constituent statistics, such as splitting points per shot into two and three point successes and misses; and assists-to-turnover ratio to assists and turnovers individually. The marginal effect of those statistics which are tracked at the player-level could then be attributed to the individual players, all except for the total number of scoring attempts. To account for tempo, scoring attempts were attributed based on the proportion of the team's minutes that a player plays, rather than for each player's accumulation of the statistic. The rationale for this decision was to account for these changes of possession, while not punishing specific players for shouldering scoring responsibilities, either as a result of team instruction or because they are the relatively more efficient shooters in the team, and are therefore expected to attempt more difficult shots. Next, rather than attributing opponents' statistics equally to players as Scott et al. (1985) do, statistics were also attributed to players where possible (steals are an opponents' turnover and blocks are missed field goals[14]). The remaining opponent statistics were attributed to players based on the proportion of the team's minutes that a player plays (in the same manner as scoring attempts). Finally, adjustments are made to eliminate positional biases and to convert the relative performance calculated from the marginal effects into a measure of absolute performance.

Berri et al. (2006) note two issues with the previous model. Firstly, they argue that assists do not belong in the model as, once scoring efficiency is accounted for, the pass leading to that level of efficiency is irrelevant. Instead they make post-estimation adjustments to redistribute production to those providing assists. Secondly, the previous model does not account for what they term as 'team rebounds', which occur when possession is gained in a way that cannot be attributed to an individual player, such as when the ball goes out of bounds. Once 'team rebounds' are estimated (since they are not recorded) and added to total rebounds, Berri et al. (2006) note that the absolute effects of total rebounds and turnovers are statistically the same, i.e. the effect of a change in possession. Additionally, they also simplify the methodology by considering a regression of win percentage on only offensive and defensive ratings (that is points scored per possession employed and points allowed per possession acquired). It is then possible to calculate the marginal effects of points scored/allowed and possessions employed/acquired and, from these, calculate the marginal effects of the different statistics based on their effects on these factors. Marginal effects of statistics tracked at the player- and team-level are then attributed to players in a similar way as previously described, with the major difference being that tempo is now accounted for implicitly, through an appreciation

[14] Berri (1999) concedes that not all blocks would have resulted in successful field goals, so the marginal effect of blocks may be slightly overstated.

of possessions. Finally, corrections for position bias[15] and a transformation from relative to absolute production is performed, before reallocating production to account for assists. Overall, this newer methodology is simpler, as well as being more accurate due to the reduction in potential collinearity between statistics which existed in the previous model.

In this study, we constructed WP using both methodologies and find more accurate[16] results using the newer methodology. Because of this and its relative simplicity, we use this measure for our analysis and henceforth focus on the newer methodology in our more detailed description in Section 4.3.3 and Appendix 4.B. WP is a measure of overall contribution, but we also consider *Wins Produced per 48 minutes* ($WP48$) which is an efficiency measure.[17]

In summary, player performance within teams in a difficult metric to measure, but a number of attempts have been made. Some of these have been created within the academic literature but a greater number have developed outside of the official peer-reviewed structure of academic literature by sports analysts. Broadly, they can be classified as either efficiency measures or measures of overall contribution. The main difference between these two types is whether the number of minutes played is considered in the calculation. Consequently, it is possible to switch between the types by either multiplying or dividing by the number of minutes played. There is some argument about which of these elements is more important, efficiency or overall contribution. We favour measures of overall contribution because they incorporate the coach's decision to select players which may capture other unobserved factors, as well as avoiding anomalies of players appearing efficient when they have played little and therefore contributed little to the team overall. But, as we appreciate that a strong argument can also be made for favouring efficiency[18], we therefore find it prudent to consider measures of both types. Our main analysis considers overall contribution, while efficiency measures are considered to test the robustness of our findings. The construction and descriptive statistics of the selected measures are discussed in Section 4.3.3 and related appendices.

2.4 Summary

Overall, while there is a lack of academic literature that analyses the effects of wage dispersion and national diversity on individual player performance, there is some evidence of the impact on team performance. However, much of this evidence is mixed and lacks consensus. This investigation hypothesises that this lack of consensus is due to the broad definitions of dispersion and diversity considered, and explores whether different types of dispersion/diversity could have differing effects. We expect teams with a higher standard deviation of wage and lower wage polarisation to perform better, while the same can be said of nationally heterogeneous (or fractionalised), but non-polarised structures. While the current literature tells

[15] Different positions perform different jobs for their teams, but Berri et al. (2006) argue that each of these jobs are necessary and complementary. Frontcourt players tend to get more rebounds and commit less turnovers, but this is at least in part due to their positional responsibilities. Adjustments are therefore made to correct these positional biases.

[16] By accuracy here we mean that the sum of WP for all players in the team is closer to the actual wins produced by the team.

[17] This conversion from an overall contribution to efficiency measures could equally have been done for *Win Shares*, but *Wins Produced* was chosen as it is our favoured measure and is produced using the current dataset.

[18] An argument for efficiency may suggest that a player's performance should not be downgraded too much for being in a good team with other good players if they are contributing positively in the time that they are on the court. And also that team performance in basketball relies heavily on efficiency and therefore so should individual performance. However, it is efficiency of possessions which has been stressed thus far, not efficiency of time per se. There are obvious parallels, especially due to shot clock rules, but the two concepts are distinct.

us little about the effects on individual player performance, and how these effects may vary depending on a player's position within the wage or nationality distribution of the team, it does give insights into how best to objectively measure player performance for the purposes of this investigation. These measures of individual player performance are utilised to explore whether any additional nuances relating to the effects of wage dispersion and national diversity on performance can be identified through payer-level analysis. Before the empirical analysis, Chapter 3 gives a theoretical overview of the measures of wage dispersion and national diversity utilised in this investigation, with a particular focus on polarisation. Chapter 4 then discusses our dataset and the methodology used in the three empirical chapters: 5, 6 and 7.

Chapter 3

Dispersion, Diversity and Polarisation

Chapter Overview

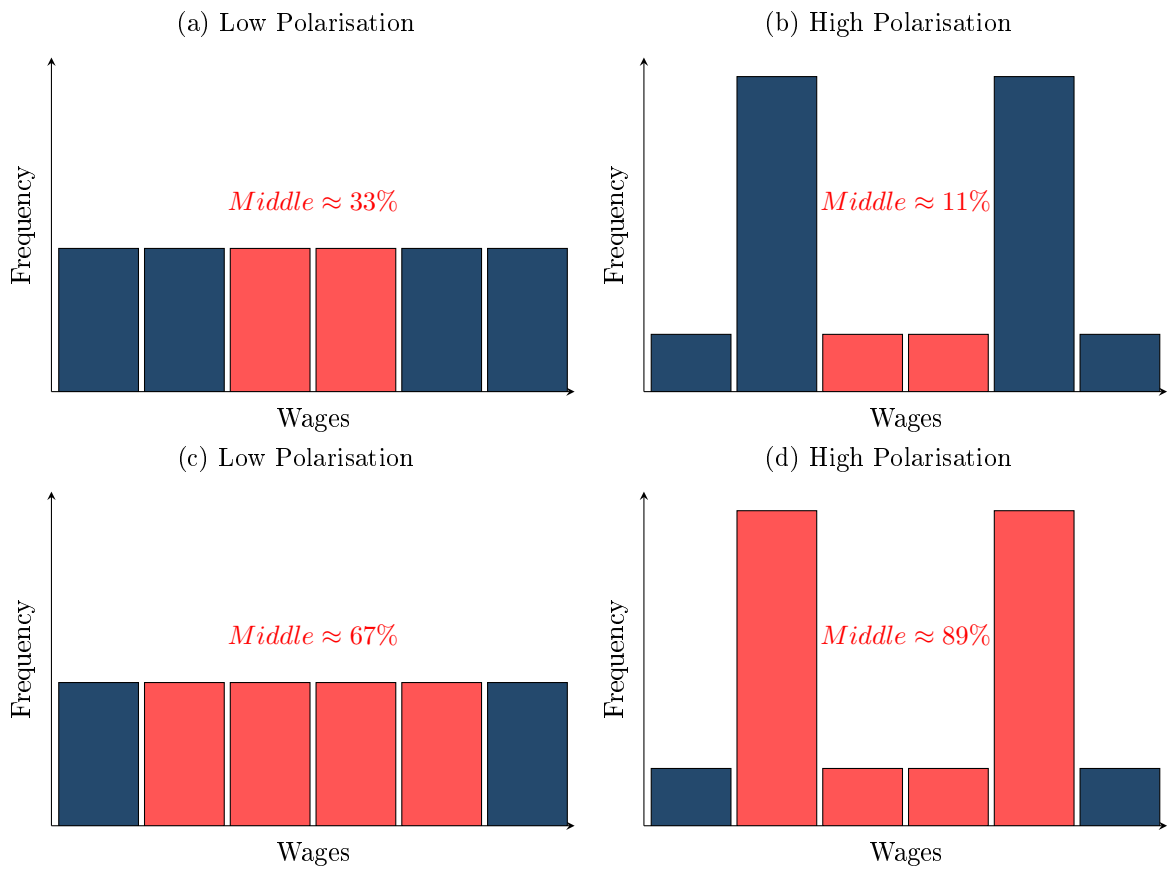
The purpose of this chapter is to give detail on the construction of the variables of interest for this investigation. In terms of wage dispersion, we consider several measures of dispersion, but the main focus is on measures of polarisation. This measure is yet to be considered in the context of wage dispersion within work teams. We compare and contrast it to other measures of dispersion to demonstrate that it captures a unique aspect of dispersion. To capture national diversity, we consider measures which reflect the structure of the diversity and the cultural distance within the diversity, as well as some which combine both of these elements. Again there is a focus on polarisation, but also on fractionalisation measures and measures of linguistic distance.

3.1 Measures of Wage Dispersion

When investigating issues of inequality and disparity, of wages as well as other variables, there are several measures to consider. The most commonly used include the standard deviation, the Gini Coefficient and the Coefficient of Variation; with other information about the distribution, such as skewness and kurtosis, being considered on occasion. A relatively recent addition to these measures is polarisation. Unlike traditional measures of inequality, which consider distributions around global means, polarisation also considers clustering around local means.

The work of Wolfson (1994) and later Foster and Wolfson (2010) consider polarisation to consist of two elements: spread and bipolarity. Spread captures dispersion in the more traditional sense, measuring how spread out the wage distribution is. Bipolarity, on the other hand, captures the effects of clusters occurring at either end of the distribution. Both spread and bipolarity are needed simultaneously to maximise polarisation. The authors also consider the idea of polarisation to be analogous to (or perhaps the dual of) measuring the middle class. While previous studies have made the distinction between a declining middle class (or polarisation) and wage inequality, the standard approach to measuring the middle class had been to set arbitrary bounds of the middle class, either in income space (see Blackburn & Bloom, 1985) or population space (see Levy, 1987). Wolfson (1994) demonstrates, on both theoretical and empirical grounds, how such an approach can produce contradictory results depending on the arbitrary range of the middle class chosen. This is done with a clever example, similar to the one provided in Figure 3.1.

Figure 3.1: The Arbitrariness of Setting Bounds for the 'Middle'



In both the top and bottom panels of Figure 3.1, we move from a low polarisation scenario to one of relatively high polarisation. This is done by making transfers on either side of the

distribution towards the middle of each side. Hopefully the increase in polarisation will be obvious to the reader, but in any case, this same example will be used later in this section (Figure 3.4) to prove that it represents an increase in polarisation, but not in inequality as measured by the Gini Coefficient. But now we focus on the folly of attempting to 'measure the middle' between arbitrary bounds to capture polarisation. In the top panels, (a) and (b), we assume the 'middle' to be the middle third of the distribution in wage space. After the transfers, we see the 'middle' has decreased in size from a third of the population in panel (a), to a ninth in panel (b). This would imply, correctly, that polarisation has increased. In the bottom panels, (c) and (d), we assume the 'middle' is the middle two-thirds of wage space. In this case, the 'middle' appears to increase from two-thirds in panel (c), to eight-ninths of the population in panel (d), implying that polarisation has decreased. Not only is this incorrect, but it is contradictory to our previous conclusion for the exact same change in the distribution. We would encounter the same problem if we chose to measure the wages in middle population space. This example proves that the conclusions reached by measuring the middle using arbitrary bounds can depend on where these bounds are set. As there is no reason to set the bounds at any given points, measuring the middle in this way is not a reliable way to measure polarisation.

Given this arbitrariness of measuring the middle in previous empirical results, Wolfson (1994) proposed an index of polarisation which incorporates information from the entire distribution. The resultant index is closely related to the Gini Coefficient and the Lorenz Curve. It can be written as:

$$P = 2(2T - Gini)/\frac{m}{\mu} \tag{3.1}$$

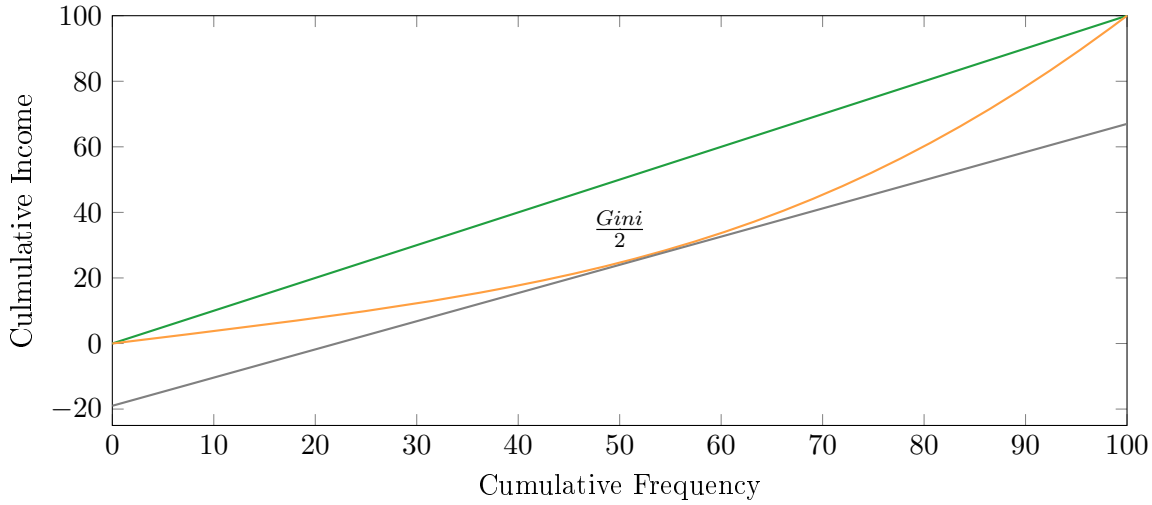
Where T is the area of a trapezoid between the 45° line of equality and a tangent line to the Lorenz Curve at the median (50th percentile); m is the median wage and μ is the mean. Figure 3.2a shows the 45° line, the Lorenz curve and median tangent line of the low polarisation scenario from our previous example. Figure 3.2b shows the area of the Trapezoid, T. The polarisation measure is the area of this trapezoid minus the area under the Lorenz curve, normalised for skewness and multiplied by four. Therefore, the areas highlighted in Figure 3.2c sum to a quarter of the polarisation measure. The reason for multiplying by four may seem arbitrary but, much like why the Gini Coefficient is defined as twice the area between the Lorenz Curve and 45° line, is so that the resultant index ranges between zero and one.

We create this polarisation measure using the Distributive Analysis Stata Package (DASP) (Abdelkrim & Duclos, 2007). As can be seen in Figure 3.2, this methodology considers percentage changes in income and frequency, so it is important to construct the measure using the natural logarithm of wage. This bounds the resultant index so that it ranges between zero and one. Foster and Wolfson (2010) do note that there is no theoretical reason why absolute changes cannot be considered, but the measure will not be bounded and will therefore be more difficult to interpret. We therefore consider percentage changes for much of our analysis, but do consider absolute changes as a robustness check.

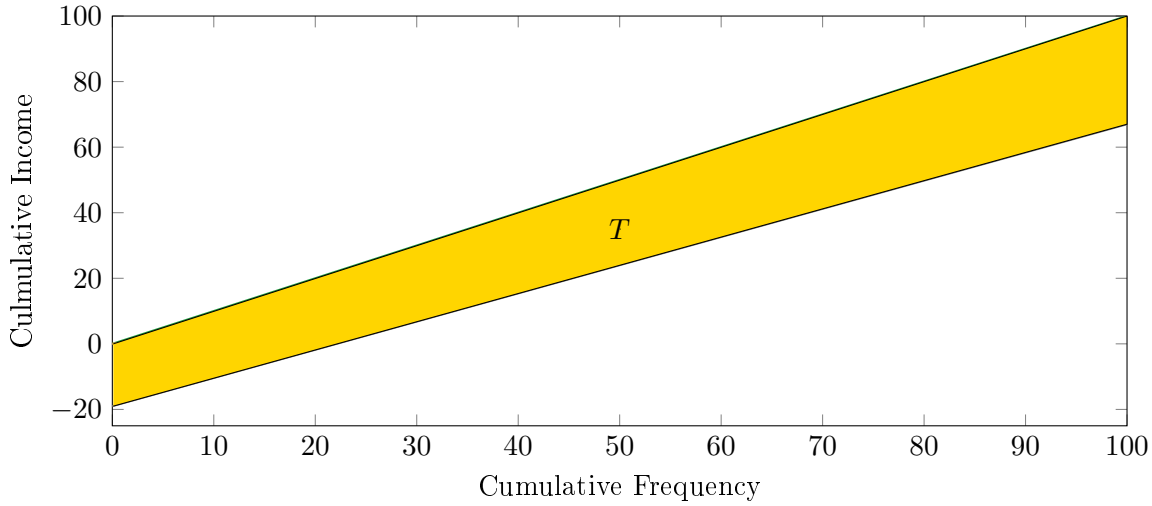
The relationship with the Gini Coefficient and Lorenz Curve helps establish the intuition as to when and why measures of inequality and polarisation conflict. Firstly, when there is an equalising transfer across the median, then both the Lorenz Curve and median tangent curve are affected causing both inequality and polarisation measures to move in the same direction. However, when there is an equalising transfer which occurs wholly on one side of the

Figure 3.2: Foster and Wolfson (2010) Polarisation

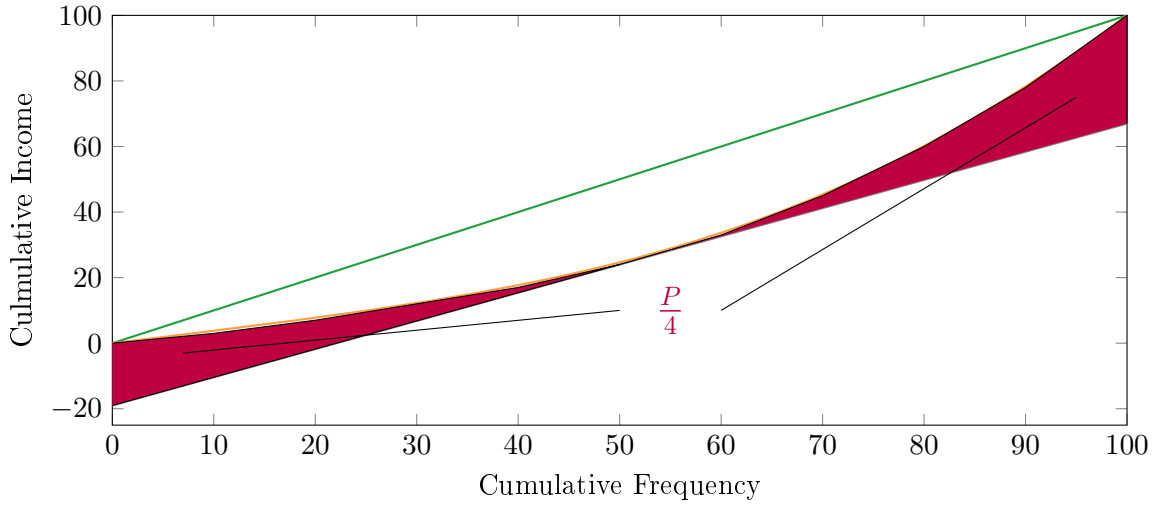
(a)



(b)



(c)



median, then the Lorenz Curve is affected but the median tangent curve is not. It is in these circumstances when inequality and polarisation measures will disagree. These ideas and this measure is developed further by Foster and Wolfson (2010) who discuss ideas of spread and bipolarity. Increased spread is the idea of the rich getting richer and the poor getting poorer, which can be caused by the transfer across the median mentioned above. Increased spread increases polarisation as there is a movement away from the middle. For this reason, it also increases inequality and standard deviation.

Bipolarity captures the idea of more concentrated poles at either end of the distribution. This can be increased by transfers which happen wholly on one side of the distribution. Such transfers could move some extreme incomes closer to the middle, while simultaneously moving other mid-range incomes away from the middle to create more concentrated frequencies on either side. As some are moving towards the middle and some are moving away, spread does not capture this phenomenon, hence the need to also consider bipolarity. Increases in polarisation caused by increased bipolarity are captured by placing more weight on changes which occur near the middle of the distribution. Foster and Wolfson (2010) go on to show that, by considering those above and those below the median as two groups, the polarisation index developed can be expressed as a function of the between-group inequality minus the within-group inequality, with between-group inequality showing spread and within-group inequality showing bipolarity. Empirical evidence is also presented from Canada and the US to show that the problems of using arbitrary bounds of the middle class is not just a theoretical concern, but can lead to contradictory empirical findings. They find evidence of higher levels of both polarisation and inequality in the US compared to Canada.

Important contributions to the polarisation literature were also made by Esteban and Ray (1994) and Duclos et al. (2004). Between them, these papers take an axiomatic approach to measuring polarisation. These papers talk in terms of an alienation-identification framework, rather than spread and bipolarity, but the broad principle of what polarisation should measure remains similar. Using their axiomatic approach, they derive the following formula for measuring polarisation:

$$P_\alpha(f) = \frac{\int\int f(x)^{1+\alpha} f(y)|y-x|dydx}{2\mu^{1-\alpha}} \tag{3.2}$$

Where $f(x)$ and $f(y)$ are density functions and μ is the sample mean. The measure also depends on the parameter α, which Duclos et al. (2004) call the 'polarisation-sensitivity' parameter and note that it should be bounded between 0.25 and 1. This formula also shares a distinct relationship with the Gini Coefficient and will coincide with it when $\alpha = 0$. But the authors prove that to measure polarisation in a meaningful way, not only should $\alpha > 0$, it should be above the lower bound of 0.25.

Zhang and Kanbur (2001) also discuss the differences between inequality and polarisation. Their empirical work applies different polarisation measures which have been developed, including those of Esteban and Ray (1994) and Wolfson (1994), to measure income polarisation in China. They argue that, in this particular empirical setting, the distinction between measures of inequality and polarisation is not as important as theorised. They find that a decomposition of inequality measures gives more relevant information regarding future policy determination. This emphasises the point that while inequality and polarisation measures will react differently to some changes in distributions, they will react similarly to others. It

may also indicate that when issues of spread and bipolarity do not interact as theorised then current measures of polarisation may not be fit for purpose. Or at least, that polarisation on its own may not tell the entire story.

The current literature regarding polarisation has mostly focussed on country-level effects (Alesina & Ferrara, 2005; Humphreys et al., 2002; Keefer & Knack, 2002; Montalvo & Reynal-Querol, 2005). However, we theorise that these issues are just as relevant at the firm-level, or in our case the team-level. In the final part of this section, we will discuss the issues with this scenario in mind. Specifically, we suggest that the dynamics of wage dispersion may be too complex to be fully captured by a single measure and it is therefore important to consider measures simultaneously. This is because there will be some changes in the distribution which cause measures to react in the same way, but some which cause them to react differently. By only considering a single measure, we may capture some of the relevant changes, but if we do not capture all of them, then we may come to incorrect conclusions and/or predictions.

Figure 3.3: Increased Dispersion and Polarisation

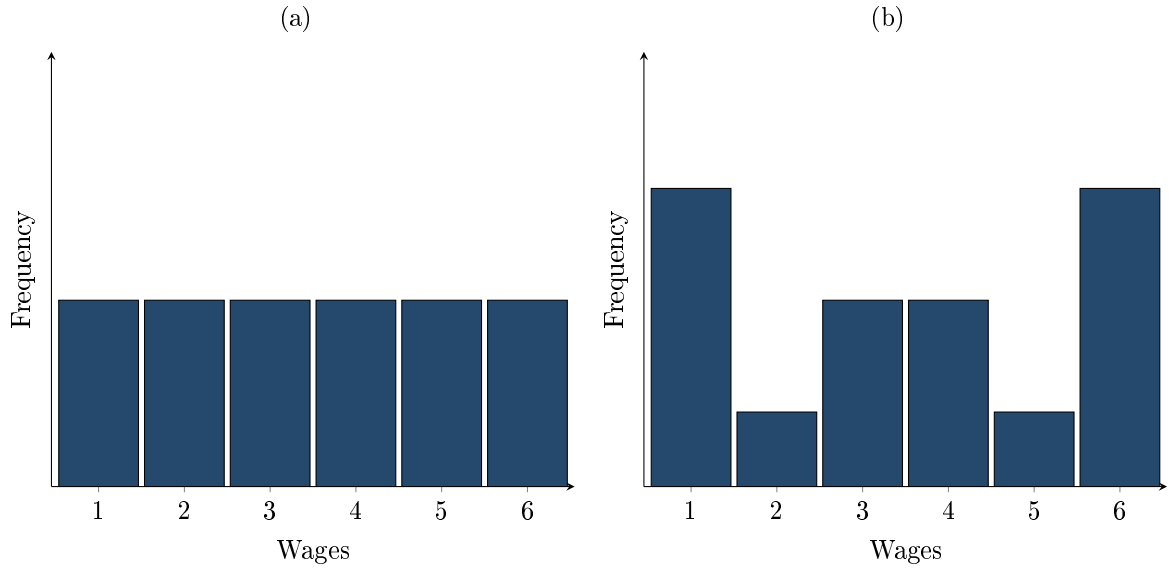


Table 3.1: Descriptive Statistics of the Distributions in Figure 3.3

	3.3a	3.3b	% change
Total	315	315	0
Mean	3.5	3.5	0
Standard Deviation	1.717	1.962	14.269
Gini Coefficient	0.278	0.313	12.590
Coefficient of Variation	0.491	0.561	14.257
HHI	0.0138	0.0146	5.797
Polarisation	0.267	0.337	26.217

Consider Figure 3.3 which contains two different distributions: 3.3a and 3.3b. Figure 3.3a shows a distribution with an even spread across different wage levels, while Figure 3.3b contains concentrated clusterings at either extreme. The difference between the two distributions can be described by transferring 1 from those that earn a wage of 2, to those which earn 5. This transfer does not change the total or average wage expenditure, but creates more clustering at the extremes of 1 and 6. As discussed earlier, a transfer such as this, across the middle from the relatively poor to the relatively rich, will increase polarisation. This is confirmed

by the accompanying table, Table 3.1, where polarisation is measured using the method of Foster and Wolfson (2010). The table also shows that all other reasonable measures of wage dispersion have also increased as a result of this type of transfer. The standard deviation, Gini Coefficient, Coefficient of Variation and HHI have also all increased – all measures would agree that wage dispersion has increased. It is worth noting that the measures increase to varying extents, as we see with the different percentage changes. However, qualitatively, all measures would agree that distribution 3.3b is more dispersed than distribution 3.3a.

Figure 3.4: Decreased Dispersion and Increased Polarisation

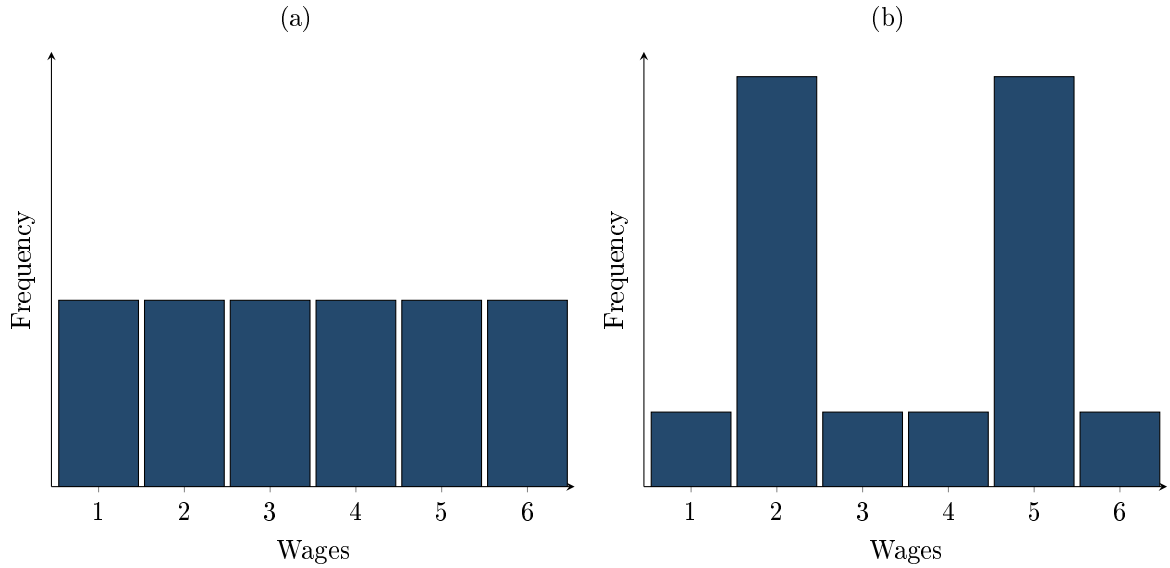


Table 3.2: Descriptive Statistics of the Distributions in Figure 3.4

	3.4a	3.4b	% change
Total	315	315	0
Mean	3.5	3.5	0
Standard Deviation	1.717	1.581	-7.921
Gini Coefficient	0.278	0.2425	-12.770
Coefficient of Variation	0.491	0.452	-7.943
HHI	0.0138	0.0134	-2.899
Polarisation	0.267	0.320	19.850

Now we turn our attention to Figure 3.4, which again contains two distributions, this time: 3.4a and 3.4b. Figure 3.4a is identical to Figure 3.3a and has an even spread across different wage levels; while Figure 3.4b again contains more concentrated clusterings, but this time not quite at the extremes. The difference between Figures 3.4a and 3.4b could be described by making transfers on either side of the distribution. On the left side, we transfer 1 from some of those earning 3 to some of those earning 1, creating a clustering at 2. Similarly, on the right side, we transfer 1 from those earning 6 to those earning 4, to create a clustering at 5. So again, these transfers are not altering the total or average wage expenditure. Table 3.2 shows that the standard deviation, Gini Coefficient, Coefficient of Variation and HHI would all agree, qualitatively at least, that these transfers decrease dispersion and would conclude Figure 3.4a is more dispersed than Figure 3.4b. However, as can be seen by the change in the polarisation measure, Figure 3.4b is more polarised than 3.4a. Although there are less extreme observations in Figure 3.4b, there are two clearly identifiable groups in the distribution.

We believe that situations such as the one given in Figure 3.4b, are more likely to cause the conflict within teams that is predicted by Cohesion Theory (Levine, 1991), than the situation in Figure 3.4a, for example. Polarisation measures are specifically designed to capture the likelihood of conflict and we believe it is this conflict that causes the negative effects of too high dispersion which are associated with Cohesion Theory. We therefore, in short, argue that polarisation measures are superior at capturing the effects of Cohesion Theory. Meanwhile, we argue that having clearly defined groups is not necessarily expected to increase effort, as Tournament Theory (Lazear & Rosen, 1981) suggests more dispersion should. In fact, having two main wage groupings (as in Figure 3.4b), rather than steady incremental increases (as in Figure 3.4a), is more likely to discourage effort. Resultantly, we predict that polarisation measures will be bad at capturing the conditions which are associated with Tournament Theory. The other measures of dispersion used in the literature seem better at capturing these conditions, particularly the standard deviation.

Furthermore, the current discussion, and the examples of Figure 3.3 and Figure 3.4, explain why the omission of polarisation from the current literature may have led to the mixed results that we see. If we were only to consider the effect on the distribution captured by the Gini Coefficient, then it would be almost impossible to distinguish between moving from the distribution of Figure 3.4b to that of 3.3/3.4a, rather than from Figure 3.3/3.4a to 3.3b. Both involve a 0.0353 increase in the Gini Coefficient which is a similar percentage change in each case. Of course, the predicted effect may be slightly different in the two scenarios if we assumed a quadratic relationship, but this would depend on the turning point and how far away from it the observations were. But a major difference is that by moving from Figure 3.4b to 3.4a we decrease polarisation, while moving from Figure 3.3a to 3.3b we increase polarisation. Therefore, if these changes in polarisation are important and we omit them from a regression model, we may have instances where other variables seem to be having a negative effect on performance, due to their correlation with polarisation. Ultimately, we argue that for a given level of dispersion in the traditional sense (e.g. standard deviation), the level of polarisation can still vary to a high degree and therefore the overall effect on performance can vary to a high degree, and vice-versa. We argue that polarisation captures something distinctly different to other measures of wage dispersion and that this difference is important in disentangling the true effects of wage dispersion on performance.

3.2 Measures of National Diversity

This section describes the variables used in this study to measure national diversity. The most basic measure used in the literature simply counts the number of nationalities within a team (e.g. Brandes et al., 2009). However, this simple measure takes little account of the structure of the diversity. For example, a team split 50-50 between two nationalities would be treated the same as an almost homogeneous team which had just one player of a different nationality to his teammates. We therefore favour measures of fractionalisation and polarisation. Montalvo and Reynal-Querol (2005) discuss these measures in detail, but a brief overview is given here.

Fractionalisation can be considered as a discrete measure of the Gini Coefficient for national groups. Equation 3.3 initially shows the discrete form of the Gini Coefficient (DG) for a team; where y and π would conventionally be the income and the proportion of the team that is represented by individuals i and j up to N, respectively. However, in our context, we consider

$|y_i - y_j|$ to be a binary measure of the distance between nationalities, rather than of income differences, to get our measure of fractionalisation ($Frac$). In other words, fractionalisation considers national groups and the distance between them, rather than income differences. The binary measure of national distance equals zero if the players are of the same nationality and equals one if the players are of different nationalities. $Frac$ is minimised when a team is homogeneous, then it is increasing in the number of different nationalities and decreasing in the size of each national group, ceteris paribus. Therefore it is maximised when a team consists of infinitely many, infinitesimally small groups of nationalities.

$$\begin{aligned} DG &= \sum_{i=1}^{N}\sum_{j=1}^{N} \pi_i \pi_j |y_i - y_j| \\ |y_i - y_j| &= 0 \qquad \text{if } i = j \\ |y_i - y_j| &= 1 \qquad \text{if } i \neq j \\ DG = \sum_{i=1}^{N}\sum_{j \neq i} \pi_i \pi_j &= \sum_{i=1}^{N} \pi_i \sum_{j \neq i} \pi_j = \sum_{i=1}^{N} \pi_i (1 - \pi_i) = Frac \end{aligned} \tag{3.3}$$

A measure of polarisation for national groups can be derived in a similar way, beginning with the measure of income polarisation ($P(\alpha, k)$) axiomatically derived by Duclos et al. (2004). This is presented in the first line of Equation 3.4. Again, we assign a distance of zero when players share the same nationality and one when they are of different nationalities. We then rearrange and set $\alpha = 1$ and $k = 4$ to obtain a measure of discrete polarisation (Pol). While α and k may vary for continuous polarisation without violating the axioms of Duclos et al. (2004), these are the only values for the parameters which are sensible for discrete polarisation and bound it between zero and one (Montalvo & Reynal-Querol, 2005, p.800).[19] This measure will again equal zero for homogeneous teams, but it will be maximised when the team consists of two equally sized groups. However, it will also be minimised in the same situation that maximises fractionalisation, when there are infinitely many, infinitesimally small groups.

$$\begin{aligned} P(\alpha, k) &= k \sum_{i=1}^{N}\sum_{j=1}^{N} \pi_i^{1+\alpha} \pi_j |y_i - y_j| \\ |y_i - y_j| &= 0 \qquad \text{if } i = j \\ |y_i - y_j| &= 1 \qquad \text{if } i \neq j \\ P(\alpha, k) = k \sum_{i=1}^{N}\sum_{j \neq i} \pi_i^{1+\alpha} \pi_j &= k \sum_{i=1}^{N} \pi_i^{1+\alpha} \sum_{j \neq i} \pi_j = k \sum_{i=1}^{N} \pi_i^{1+\alpha} (1 - \pi_i) \\ P(1, 4) &= 4 \sum_{i=1}^{N} \pi_i^2 (1 - \pi_i) = Pol \end{aligned} \tag{3.4}$$

When a team consists of only one nationality, both fractionalisation and polarisation will equal zero. When a team consists of only two nationalities, then the measures will have a perfectly linear relationship, with Pol being twice as large as $Frac$.[20] Beyond this, when a team consists of three or more nationalities, there is no linear relationship between the two measures. The relationship remains positive at lower levels of $Frac$ and Pol, but as the num-

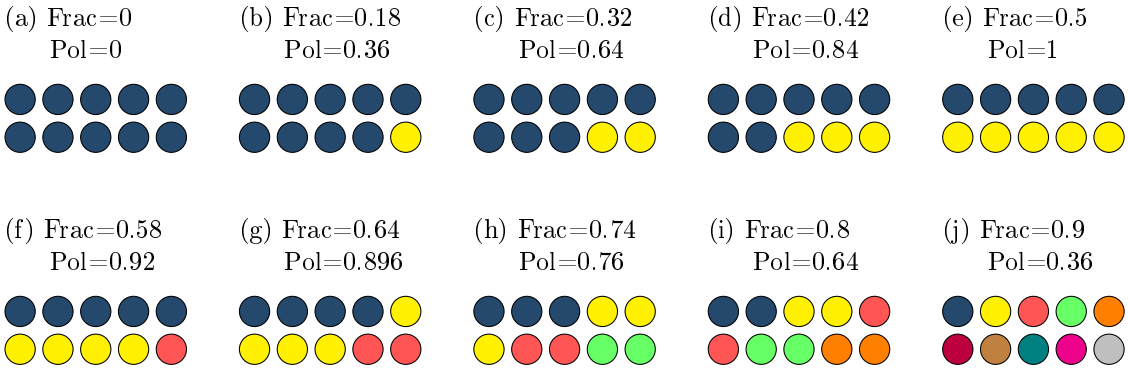
[19] α is a measure of polarisation sensitivity, while k is a scaling constant. While α may vary within given bounds when considering continuous measures, Montalvo and Reynal-Querol (2005) explain that α must equal one for discrete polarisation. Setting the scaling constant, k, equal to four then bounds the measure between zero and one.

[20] This can be seen by substituting π_j for $1 - \pi_i$, as $\pi_i + \pi_j = 1$.

ber of nationalities and hence fractionalisation increases, this relationship weakens and then becomes negative. A similar relationship between the variables is also noted by Montalvo and Reynal-Querol (2005).

To aid the reader's understanding, some explanatory examples are given in Figure 3.5. Each dot represents a player and each colour represents a nationality. Therefore, panel (a) is a homogeneous team, with every player sharing the same nationality. In panel (b), we introduce a second nationality which increases both fractionalisation and polarisation, with the latter twice the former. This continues as we add further players from the second nationality in panels (c)-(e) until we have a perfectly polarised team with polarisation of 1 and fractionalisation of 0.5.

Figure 3.5: Fractionalisation and Polarisation Examples

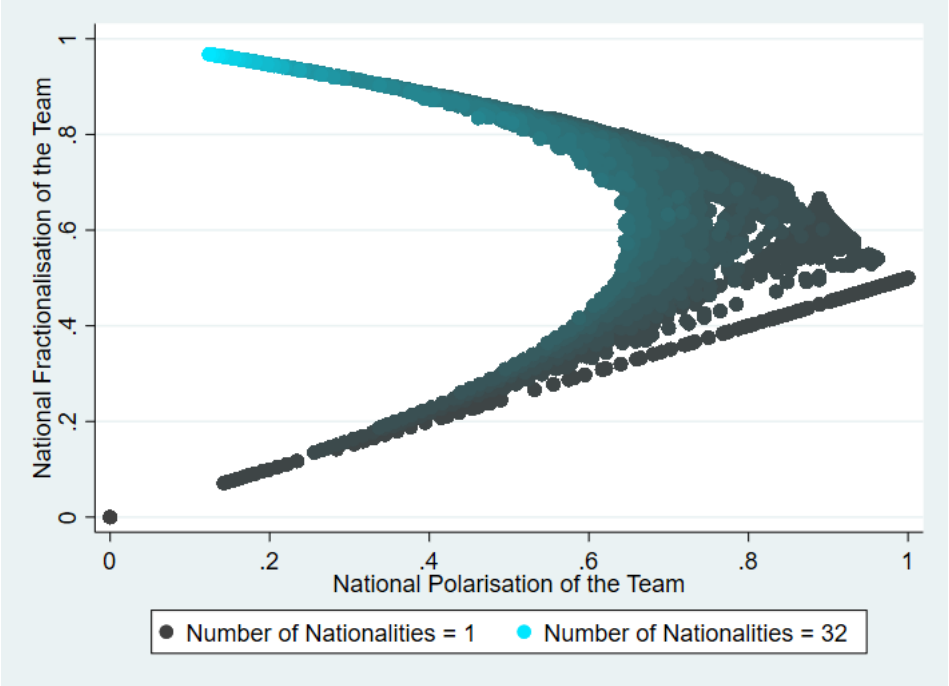


A third nationality is added in panel (f) of Figure 3.5. While this increases fractionalisation, it decreases polarisation. As another player of the third nationality is added in panel (g), a fourth nationality added in (h), and a fifth in (i), fractionalisation continues to increase while polarisation decreases. In panel (j), the team is perfectly fractionalised, or at least it has the maximum level of fractionalisation that a team of ten players can achieve, with every player of a different nationality. In this final panel, polarisation has dropped to a relatively low value (as low as it was in panel (b) with an almost homogeneous team, bar one player). These examples show that, while there is some correlation between these measures when there is a small number of national groups, they clearly capture distinct phenomena.

To give further intuition of the relationship between fractionalisation and polarisation over the full range of possible values, we simulate a dataset of teams consisting of randomly assigned players from randomly assigned nations. The size of these teams are of a similar magnitude to those in our NBA dataset.[21] The maximum number of nationalities in a team was allowed to be higher in our simulations to display higher levels of fractionalisation, though the maximum was still less than the number of nations competing within the league in any given year. Full details and code of the simulations is available upon request. Figure 3.6 displays the results of these simulations, where the same relationship as previously described is apparent. With two nationalities, the linear relationship extends to a polarisation of 1, which gives a fractionalisation of 0.5. However, this same value of fractionalisation can be achieved with a lower polarisation when there are more national groups. With many small national groups, fractionalisation approaches a value of one, while polarisation approaches zero.

[21] 700,000 player observations were created and randomly assigned to years and teams. Iterative reallocations of excess players were then conducted to ensure the size of these groups are comparable to our NBA dataset. Nationalities were also randomly assigned, though the probabilities of each nationality being assigned varies to ensure a range from homogeneous to highly fractionalised teams.

Figure 3.6: Fractionalisation and Polarisation Simulations



Other studies consider measures of cultural distance between nations, usually proxied by the linguistic distance between their primary languages. In this study, we use linguistic data from the Enthnologue Global Database (Lewis et al., 2014) to calculate the linguistic distance between countries, by considering the similarities and differences in the branches of each language's linguistic tree. In this database, each modern-day language is an end branch which can be traced back through up to 15 interior branches. To aid the reader's understanding, three examples are given in Table 3.3 for English, German (standard) and Chinese (Mandarin).

Table 3.3: Examples of Linguistic Branches

Language	Linguistic Branches
English	Indo-European → Germanic → West → English
German (standard)	Indo-European → Germanic → West → High German → German → Middle German → East Middle German
Chinese (Mandarin)	Sino-Tibetan → Chinese

From this data, we then calculate linguistic distance between each language as shown in Equation 3.5; where $Shared_{k,l}$ is the number of linguistic branches shared by language k and language l, and $Max_{k,l}$ is the maximum number of branches that either of the languages has.

$$LingDist_{k,l} = 1 - \frac{Shared_{k,l}}{Max_{k,l}} \tag{3.5}$$

Comparing a language to itself, the number of shared branches will equal the maximum, so a linguistic distance of zero will be given. When two languages share no linguistic branches then they will have a linguistic distance of one. Any languages which share some, but not all, branches will have a linguistic distance between zero and one. Table 3.4 shows the linguistic distances between our three example languages.

Table 3.4: Examples of Linguistic Distance

	English	German	Chinese
English	0		
German	0.4286	0	
Chinese	1	1	0

The principle diagonal is zeros, as each language has a distance of zero to itself. English and German share three branches out of a possible seven, so have a distance of 0.4286. English and German each share no linguistic branches with Chinese, so a distance of one is given for each of these pairs. By considering the primary language of each country to be the one which is most commonly spoken by its population, we can create a measure of linguistic distance between countries by substituting subscripts k and l for i and j of the associated countries. We can then make pairwise comparisons between players within a team and take the average to find the average linguistic distance of each player to his teammates. We call this variable, calculated for each player, $LingDist1$. We can then take the average amongst players within a team.[22] A nationally homogeneous team will have a linguistic distance of zero as each pairwise comparison will give a distance of zero. The variable will be maximised at one when a team consists of players, each with a different nationality and when each of these nationalities share no linguistic branches. All other teams will lie within these bounds.

However, there are two major problems with this measure of linguistic distance which we additionally attempt to resolve, both individually and simultaneously. Firstly, players from countries which speak the same primary language are treated as though they are from the same country. They will generate a linguistic distance of zero and will be considered as equally different to all other nationalities. Secondly, the measures do not take account of all aspects regarding how the linguistic distance within a team is structured. A highly fractionalised team could have the same average linguistic distance as a highly polarised team. These nuances are not captured when considering only average linguistic distance.

To address the first concern, we can consider a country's cultural makeup to consist of more than one language and weigh each of these by its proportion of speakers. We consider up to four primary languages, but only consider languages that are spoken by at least 10% of the country's population. Equation 3.6 summarises our approach, where $LingDistM_{i,j}$ is the linguistic distance between country i and country j, when up to M primary languages are considered. For each primary language k of country i, we compare it to each primary language l of country j, and vice-versa. We identify the primary languages which are most similar (i.e. those which have the smallest distance) and then multiply this distance by the proportion of speakers ($\Lambda_{k,i}$) of the corresponding language in the corresponding country. Finally, we divide through by the sum of all these proportions (to account for countries with many multilingual speakers whose proportions would sum to more than one). The resulting measure is still bounded between zero and one: equalling zero when all the primary languages of the two nations are the same, and one when none of the primary languages of the countries share any linguistic branches. Under this approach, even if large proportions of two countries share a common first language, we still account for differences in their second, and possibly third and fourth languages, which capture cultural differences between the nations.

[22] As will subsequently be discussed, taking a team average in this way is equivalent to a measure of fractionalisation which incorporates linguistic distance.

$$LingDistM_{i,j} = \frac{\sum_{k=1}^{M} [\Lambda_{k,i} \min_l(LingDist_{k,l})] + \sum_{l=1}^{M} [\Lambda_{l,j} \min_k(LingDist_{l,k})]}{\sum_{k=1}^{M} \Lambda_{k,i} + \sum_{l=1}^{M} \Lambda_{l,j}} \tag{3.6}$$

The second problem can be solved by combining our measure of linguistic distance with our measures of fractionalisation and polarisation. Rather than considering the distance between groups to be binary as we did in Equations 3.3 and 3.4 for our basic measures of $Frac$ and Pol, we can consider them to be continuous by using our measures of linguistic distance. This is shown for fractionalisation in Equation 3.7 and for polarisation in Equation 3.8.

$$FracM = \sum_{i=1}^{N} \sum_{j=1}^{N} \pi_i \pi_j LingDistM_{i,j} \tag{3.7}$$

$$PolM = 4 \sum_{i=1}^{N} \sum_{j=1}^{N} \pi_i^2 \pi_j LingDistM_{i,j} \tag{3.8}$$

Equation 3.7 shows that combining linguistic distance with fractionalisation is equivalent to taking the average linguistic distance of players within a team. This becomes clear when considering that the proportion of a team comprising of each nationality is simply the sum of the proportions of the team that each player of that nationality represents. However, combining linguistic and polarisation in Equation 3.8 captures a slightly different aspect. While both measures are increasing in the argument of linguistic distance between players, $FracM$ will be maximised when a team consists of players all of different nationalities which speak different languages which, in turn, do not share any linguistic branches. Whereas $PolM$ will be maximised when a team consists of two evenly sized national groups, which share no linguistic branches in their M primary languages.

As with wage dispersion, we hypothesise that diversity/dispersion in the more traditional sense will have positive effects on team performance. However, polarisation may foster conflict, sabotage and other non-cooperative behaviours through creating an in- versus out-group mentality within the team. Teams would therefore maximise their performance with a nationally fractionalised team which avoids national polarisation.

3.3 Summary

This chapter has provided a theoretical overview of the measures used in our investigation to capture wage dispersion and national diversity. The values that these measures take in our dataset are explored in the relevant empirical chapters 5, 6 and 7. The main focus has been on measures of polarisation for both continuous and discrete variables, and how these differ from other commonly used measures. These polarisation measures will be considered alongside more traditional measures of dispersion/diversity during the later empirical analysis. In general, we expect the more traditionally used measures to be positively associated with performance, while polarisation has negative effects. Next, the dataset used and other prerequisites are discussed in Chapter 4, prior to the empirical analysis.

Chapter 4

Data and Methodology

Chapter Overview

This chapter covers the remaining prerequisites before presenting our main empirical work. First we describe the dataset, discuss its structure and make note of the econometric techniques utilised in our investigation. These include mainly panel data techniques for the team-level analysis and multilevel modelling for the player-level analysis. We then consider the measures of team and individual player performance which will be the dependent variables in empirical analysis of Chapters 5, 6 and 7. In each case, we will justify the selection of our main measure and specify which measures will be used to conduct robustness tests. Lastly, we make note of potential covariates which will be considered as control variables in our analysis.

4.1 Data Collection

This study utilises data from the National Basketball Association (NBA) from 2007-18 to investigate the impact of wage dispersion and national diversity on performance. Our full dataset includes 1,262 distinct players appearing for the 30 NBA teams[23] over 11 seasons, resulting in 5,932 player-team-year observations. This includes all players which appear for each team in each season, meaning players can appear multiple times within each season, once for each team they played for, due to trades. The way this is dealt with depends on the level of analysis (team versus player) and the econometric model being utilised. This is discussed further in Section 4.2. Most of the data for this study was sourced through public access websites. Player information was extracted from a particularly useful site called Basketball-Reference[24] (BR). This includes the age (calculated as age on the opening day of the given season), experience (number of seasons since debut NBA season), tenure (number of seasons previously played for player's current team), nationality (based on place of birth); as well as games played, games started and minutes played of each player in each season. Additionally, player- and team-level performance statistics are extracted from the same website; these are described in Section 4.3.

The most difficult information to source was the data for player wages. Some individual wage data is available through BR; however, this data is incomplete. Therefore, wage data was collected from three other sources to increase the coverage of the dataset. Due to players joining teams at different stages of the season and being traded between teams during the season, the snapshot of the teams' roster provided by each source does not necessarily cover every player that played for each team at some point during each season, and therefore does not completely satisfy wage data for every observation of our dataset. As well as allowing the sources to fill in the gaps of one another, collecting wage data from multiple sources also improves the reliability of the data by comparing across sources where possible. The other three sources used were Fox Sports [25] (FOX) and two other databases complied by Mark Deeks[26] (MD) and Patricia Bender[27] (PB) respectively. Once combined, the four datasets offer wage data for every player in each season of the dataset (2007-18). The only missing data was for some instances of player observations which occur twice or more frequently within a season, i.e. when players are traded between teams. In these situations, where wage data is still missing, it is assumed that the player's contract is traded, so that the player received the same wage as they did at their previous team in the current season.

In the resulting dataset, there is exact and complete agreement across the available data of the four sources for 65.51% of the observations. There is very strong correlation, close to 1, between the sources. Furthermore, there is no variation greater than 10% between the sources for 82.92% of the observations suggesting that some small difference may exist due to rounding. The average agreement per observation is 79.50%. A decision was therefore taken to use the mode wage across the four sources where available and the mean where it is not

[23] There was one permanent team relocation in 2008 (the Seattle Supersonics became Oklahoma City Thunder), a temporary relocation ended at the beginning of our sample in 2007 (the Hornets returned from Oklahoma to New Orleans after Hurricane Katrina) and a team-name change in 2012 (the New Jersey Nets became the Brooklyn Nets). We see no reason why these changes should be correlated with our explanatory variables and bias our estimates in any way. Therefore we ignore these changes when coding the teams.

[24] `https://www.basketball-reference.com/`

[25] `https://www.foxsports.com/nba`

[26] `http://www.shamsports.com/`

[27] `http://www.eskimo.com/~pbender/`

(due to no agreement between sources or an equal split). The minimum pairwise correlation between all four sources was greater than 0.97, so we feel comfortable with combining these sources in this way.

From the data collected, several variables have been constructed, including our variables of interest to investigate wage dispersion and national diversity. Several control variables are also constructed to capture any heterogeneity among teams that is correlated with both our main variables and team performance. These variables are examined in more detail in Section 4.4, where we discuss how they are constructed, how they are correlated and why they are important controls to consider. Appendix 4.F and 4.I regress these control variables on the performance measures at the team- and player-level, respectively, to show their relevance in determining performance. Appendix 4.H and 4.J also regress them on our variables of interest, to show their partial correlations with our wage dispersion and national diversity measures. Throughout this investigation, we signify the natural logarithm of a variable by adding the prefix ln- to the variable name, and the quadratic form by adding the suffix -^{2}. This study is also particularly interested in the dispersion of variables within teams and adds the suffix -SD to the variable name when considering the standard deviation of that variable. In the remainder of this section, we begin by describing our data structure, its identifiers and the econometric techniques utilised. We then discuss the measures of team performance used as dependent variables and the justification of our choice of control variables.

4.2 Data Structure and Econometric Techniques

At the team-level, we consider a team's roster to include every player who featured for the team in a given season and aggregate the statistics accordingly. As a robustness check, we also consider the roster to consist of just the 'core team' for which we adopt the definition of Simmons and Berri (2011), which includes only those players who played at least 20 games and at least 12 minutes per game. The individual player data is aggregated for the 30 NBA teams over eleven seasons (2007-2018) giving a strongly-balanced, *short* panel data structure with 330 observations. The variable $Team$ denotes the team identifier and the variable $Year$ denotes the seasons (time identifier).

It is important to consider team effects, as teams are heterogeneous and will vary in many aspects including financial resources, popularity and geographical location. It is possible that apart from team performance, this team-level heterogeneity may also affect aspects of wage dispersion and national diversity. Some of these will be captured by our control variables, but there may be some elements of team heterogeneity which are more difficult to capture. Therefore, panel data methods which control for potential correlation between this time-invariant, team-specific heterogeneity and our explanatory variables, namely the Fixed Effects (FE) model, are used. We conducted Hausman-type tests using the Mundlak (1978) approach which formally justifies this decision, though we do briefly consider the results from Random Effects (RE) and multilevel models (discussed subsequently) for robustness. We use robust standard errors which we cluster by $Team$ to account for potential serial correlation across teams and heteroskedasticity (Cameron & Miller, 2015).

There are also several reasons to control for time effects, with each performance measure being calculated per season and dependent on the relative competition. The wage cap is applied, though with exceptions, to teams across the NBA and is adjusted each season. Therefore, although our performance measures are not expected to change on average over time, it is

important to control for possible changes in explanatory variables over time and investigate their effect on team performance, holding year effects constant. We therefore include *Year* dummies, but again remove them to test the robustness of our results.

Overall, at the team-level, identification is based on the assumption that time-varying unobservable factors affecting team performance are not correlated with the explanatory variables of interest (wage dispersion and national diversity). While this is plausible to some extent, it is not strictly guaranteed. If there is some persistence in team performance, then teams which perform well in one season are more likely to do well in the following season. If our variables of interest are influenced by the previous season's performance, which in turn influences this season's performance, this potential autocorrelation would mean that this season's dispersions are correlated with the previous season's errors. If this is the case, strict exogeneity cannot be assumed. Examples of how recent success may influence current wage dispersion and national diversity are that similarly talented players may be willing to accept a lower wage to play for a more successful team, and more successful teams may find it easier to recruit talented players from abroad. Caution must therefore be used when assuming time-invariance of the team-specific fixed effects and the strictly causal nature of our estimates.

At the player-level, the data has an inherent multilevel structure, with players nested within teams over multiple seasons. Our analysis considers both player- and team-level variables, as well as interactions between the two. Therefore, to gain relevant statistical inferences, our choice of modelling should account for the design of data collection and the variables being considered (Gelman & Hill, 2006). We therefore favour multilevel modelling (MLM) for our main regression analysis at the player-level, which brings a number of advantages. MLM has appeared and reappeared in academic work over time, and has seen a re-emergence due to the vast increase in computing power (Greenland, 2000). It has been utilised to investigate a range of topics, notably in education research (Borman & Dowling, 2010; Hermans et al., 2008; P. W. Hill & Rowe, 1996; Johnes, 2006), and health research (Goldstein et al., 2002; Li et al., 2005; Riva et al., 2007), among other areas. While the player-level analysis of this investigation is novel, MLM has been used within sports economics too, for example, to investigate driver and constructor performance in Formula One (Bell et al., 2016) and attendance in soccer (Ferreira & Bravo, 2007).

As the name indicates, MLM allows for multiple levels of nesting, so in our case we can control for both team- and player-level heterogeneity simultaneously. If, as may reasonably be expected, player-level effects may also vary by team due to unobservable factors (such as management style, team morale and player attitudes for example), then classical regression analysis only gives us some blunt tools to account for this. In the absence of MLM, one is left with the option of 'complete pooling' or 'no pooling' approaches (Gelman & Hill, 2006). Complete pooling would include both player and team dummies in an FE-type model, which ignores the between variation (in our case, between player and between team variation), and focuses on the within (season) variation. Conversely, no pooling, which would exclude these dummies in an RE-type model, tends to overstate the between variation (Gelman & Hill, 2006). The MLM model allows for partial pooling and provides a compromise between the two extremes. When there is very little within variation, MLM reduces towards the no pooling (RE) model; when there is very large within variation, it reduces towards complete pooling (FE) model. It adds more weight to the within variation when more of it exists, as well as when there are more observations for a player.

If we, for example, set a panel data structure at the individual player-level over time and account for team-level differences through dummies, this would not account for correlated error terms of individuals within team-groups, and simply express the average difference as a constant term. MLM allows for partial pooling and clustering at each level of analysis without fear of overfitting (Gelman & Hill, 2006). In essence, it takes the coefficients from the regression at the first level and these become the inputs at the second level.[28] This also allows for the inclusion of team-level indicators along with team-level predictors without the collinearity problems which would occur with traditionally used alternatives. We are therefore able to account for the true structure of the data and coherently incorporate both player- and team-level effects simultaneously through the MLM model.

We also account for position and year effects through dummies. These are not accounted for using additional levels as this would lead to overfitting. We already track the players over time in the player level, so the addition of year dummies simply captures heterogeneity specific to a given season. The position dummies account for any differences which may occur between positions such as differences in average performance or wages. We do not expect these dummies to play a large role in the analysis of wage dispersion and national diversity and they are removed during robustness checks. Again, we use robust standard errors which we cluster by *Player* to account for potential serial correlation across players and heteroskedasticity (Cameron & Miller, 2015).

An additional advantage of MLM is that it allows for multiple player observations within each year, as we have in our data, and which is not compatible with panel data modelling. Consequently, when utilising MLM, we retain the entire dataset of 5,932 player-team-year observations. In our robustness checks, we also consider panel data regressions (FE and RE) with the variable $Name$ as the player identifier, $Year$ as the time identifier and include team and position dummies. As panel data structures can only include a single player observation within each panel (Year), we keep only the player-year observations for the team which the player played the most minutes for. This reduces the sample size to 5,213 but allows for the panel data structure. We will also again consider our definition of the 'core team' to consider only players which play regularly[29] as an additional robustness check.

One issue with our utilisation of multilevel modelling is the lack of a clear, universally accepted goodness-of-fit measure. Due to computations at different levels, the concept of an R^2 measure to report the proportion of the variation explained by a model, cannot be applied in the usual way. Some studies (Edwards et al., 2008; Nakagawa & Schielzeth, 2013; Snijders & Bosker, 1994) have sought to address this by proposing alternative measures with similar properties. However, this research is currently focused on MLM with a single additional level and does not extend to two additional levels as our data has. Furthermore, a clear consensus on a 'pseudo R^2' is yet to be reached amongst the alternatives proposed. Consequently, we report only the log likelihood of the MLM regressions. The robustness tests which use FE and RE, for which an R^2 can be calculated, give an indication of the goodness-of-fit.

It is also possible to use MLM for the team-level analysis with a single team level, and we do

[28] In our usage, we only consider the most basic form of MLM which includes only the constant of the regression at each level. These are known as intercept-only models, which consider an error term at each level (Rodrîguez, 2019).

[29] The 'core team' is considered to consist of players which played at least 20 games and at least 12 minutes per game, as defined by Simmons and Berri (2011).

this as a robustness check. However, as we have a strongly balanced panel, there is no partial pooling to be done based on the number of observations. The Hausman-type tests (using the Mundlak (1978) approach) also indicate that there is sufficient team-level heterogeneity to warrant preferring the FE model, so the two approaches will be similar. But we show the comparison as a robustness check.

4.3 Dependent Variables

This section will discuss the measures utilised in this study to capture team performance and the individual performance of players within teams. As established, a major problem with analysing individual player performance within teams is the interdependence of production, which creates difficulties in attributing credit (and/or blame) to individual players. Simple and obvious proxies exist such as points scored directly by players. However, such a statistic is overly simplistic and ignores multiple aspects through which players can contribute towards the team's overall performance. With modern-day data analytics, there are a number of additional statistics which one may consider to create a more sophisticated measure of individual player performance, but this then raises the issue of how these statistics should be weighted. Section 2.3 gives an overview of the existing measures created and utilised, both within the academic literature and by sports analysts outside of academia. This section presents more detail on the technicalities of the measures including formulae and describing how the measures are constructed for the current study.

The measures of team performance considered in this section are *Regular Season Win Percentage* (WL), the *natural logarithm of Regular Season Win Percentage* ($lnWL$), *Points Difference per Game* (PDG), a *'Simple Ratings System'* which accounts for schedule strength (SRS), *Offensive Rating* ($ORTG$) and *Defensive Rating* ($DRTG$). The measures of individual performance considered are *Wins Produced* (WP), *Win Shares* (WS), *Value Over Replacement Player* ($VORP$), *Wins Produced per 48 minutes* ($WP48$), *Player Efficiency Rating* (PER) and *Box Plus/Minus* (BPM). The primary measure used in this investigation will be WL at the team-level and WP, created from the current dataset using the methodology of Berri et al. (2006), at the player-level; while the other measures will be used to check robustness and to capture different aspects of performance. An overview of all these measures is given here, with some technical details in the supporting appendices. We also look at the distributions and descriptive statistics of these measures, discussing correlations between them and how they are influenced by the game changing over time.

4.3.1 Basic Statistics and Notation

Firstly, Table 4.1 lists the individual statistics of players which are used to construct the measures of individual player performance, alongside their respective notation and mean. In addition, when referring to the team summation of a statistic, we remove the prefix $p-$ (for player) and add the prefix $t-$ for the team or $o-$ for the opponent. The means shown in Table 4.1 are per-game averages of all players in our dataset. Our analysis will focus on per-game averages to account for variations in the number of games played in each season.[30] This includes deep reserves who play infrequently, which is why the means may seem low.

[30] In nine of the eleven seasons in our dataset, each team played 82 games, which is the current standard in the NBA. However, the 2011-12 season was shortened to 66 games due to a lockout caused by disagreements between owners and players, mainly relating to revenue sharing and the structure of the salary cap. In the 2012-13 season, a game between the Indiana Pacers and Boston Celtics was cancelled due to the 2013 Boston Marathon bombing. As the play-off position of both teams was already decided, the decision was made not to reschedule the match. Therefore these teams only played 81 games in the 2012-13 season.

Table 4.1: Notation and Per-Game Means of Individual Player Statistics

Notation	Description	Mean
pFGS	Successful Field Goals	2.978
pFGA	Field Goals Attempt	6.660
p3PS	Successful 3 pointer	0.609
p3PA	3 point Attempt	1.742
p2PS	Successful 2 pointer	2.369
p2PA	2 point Attempt	4.917
pFTS	Successful Free Throw	1.391
pFTA	Free Throw Attempt	1.855
pORB	Offensive Rebound	0.886
pDRB	Defensive Rebound	2.558
pAST	Assist	1.751
pSTL	Steal	0.615
pBLK	Block	0.393
pTOV	Turnover	1.125
pPF	Personal Foul	1.730
pPTS	Points Scored	7.955

Table 4.2 shows the per game averages at the team-level and, in addition, the estimated possessions employed/acquired per game ($tPOSS$) of each team. As discussed in Section 2.3, as teams receive an equal number of possessions per game, efficiency is an important factor when analysing basketball performance. As teams tend to play at different tempos, possessions must therefore be taken into consideration. Possessions are not specifically recorded, but can be estimated accurately using a number of methods. Berri et al. (2006) describe how possessions employed and possessions acquired can be estimated separately, but will give the same result as each possession acquired must be employed in some capacity.[31] This symmetry indicates reliability in the estimation method.

Table 4.2: Means of Per-Game Team Statistics

	Mean	Std Dev.	Min.	Max.
tFGS	37.726	1.710	33.227	43.073
tFGA	82.909	2.913	75.841	90.317
t3PS	7.676	2.036	3.683	15.317
t3PA	21.498	5.357	11.293	42.317
t2PS	30.050	1.930	23.390	35.256
t2PA	61.411	4.573	41.902	74.305
tFTS	17.770	1.922	12.244	24.110
tFTA	23.393	2.478	16.573	31.085
tORB	10.758	1.254	7.646	14.585
tDRB	31.732	1.867	27.195	36.537
tAST	21.899	1.852	17.354	30.378
tSTL	7.576	0.855	5.537	10.012
tBLK	4.849	0.755	2.488	8.167
tTOV	14.330	1.108	11.182	17.720
tPF	20.378	1.450	16.585	24.756
tPTS	100.897	5.019	86.955	115.890
tPOSS	96.513	3.044	89.530	103.515

Possessions are also estimated using a variety of other slightly different methodologies, as summarised by Kubatko et al. (2007). We favour the methodology of Berri et al. (2006) due

[31] Only very slight differences exist due to an estimation of the proportion of free throws that result in a change of possession. This is discussed further in Appendix 4.A.2.

to the aforementioned symmetry between possessions employed and acquired. They can also be estimated (with a little additional work) from our dataset without relying on estimates from previous data. However, each of the methods produce very similar results. This is shown in Appendix 4.A.2, where equations, summary statistics and correlations of these different estimates of possessions are presented.

4.3.2 Measures of Team Performance

For the majority of our analysis, we use win percentage over the course of the regular season, given by WL, as our primary measure of team performance. This is commonly used in the sports literature (e.g. Depken, 2000; Frick et al., 2003; Kahane et al., 2013), as ultimately a team with a higher win percentage is usually considered to be a more successful team. However, we also consider alternative measures of team performance to test the robustness of our results. We consider the natural logarithm of win percentage ($lnWL$) to capture percentage changes in win percentage, as it could be argued that win percentage is more difficult to improve at higher values.

It may also be argued that win percentage is a relatively blunt measure altogether, as it treats narrow victories/losses the same as large ones. Therefore, we also consider points difference per game (PDG), which is the average difference between points scored and points conceded, divided by the number of games played. Again, this is a common robustness check (see Kahane et al., 2013). Another interesting variable we consider for robustness is a simple ratings system (SRS) constructed by Basketball-Reference. This measure accounts for average points difference but also for the strength of the schedule that a team faces. As precise detail on how this measure is constructed is not given, we retain it only as a robustness check.

In addition, we also consider team Offensive Rating ($ORTG$) and Defensive Rating ($DRTG$). These measures account for the tempo that teams play at by considering the estimated number of possessions that a team has. The formulas for these are given in Equations 4.1 and 4.2, where $tPTS$ is points scored by the team, $oPTS$ is points scored by the opponent (or equivalently points allowed by the team), and $tPOSS$ is possessions[32]; all of which are measured per game at the team-level. A slight alteration to the traditional calculation of $DRTG$ is made, which is to multiply $oPTS$ by minus one. This is done so that higher performance is indicated by higher values, as is the case with our other variables.

$$ORTG = \frac{tPTS}{tPOSS} \tag{4.1}$$

$$DRTG = \frac{-oPTS}{tPOSS} \tag{4.2}$$

Table 4.3 shows descriptive statistics of our measures of team performance, while Figure 4.1 presents their distributions. There are no draws in the NBA[33], so each game will end with one team winning and the other losing. Therefore the average win percentage must be 50%, as shown in Table 4.3. The most successful team within a single season in our dataset was the 2015-16 Golden State Warriors with a win percentage of 89%. The worst performing team was the Charlotte Bobcats with a 10.6% winning percentage in the 2011-12 season. There

[32] Specifically, possessions employed are used to construct Offensive Rating and possessions acquired are used to construct Defensive Rating. But as the two estimates of possession are virtual identical, this distinction is ignored here to simplify notation.

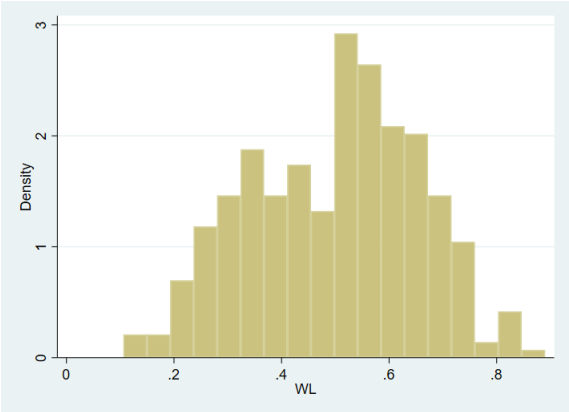
[33] In the rare scenario that the scores are level at the end of regulation time, additional five minute overtime periods are played until a winner is decided.

is only a slight negative skew to the winning percentages in our data, as Figure 4.1a shows. Conversely, $lnWL$ is highly negatively skewed as is apparent in Figure 4.1b.

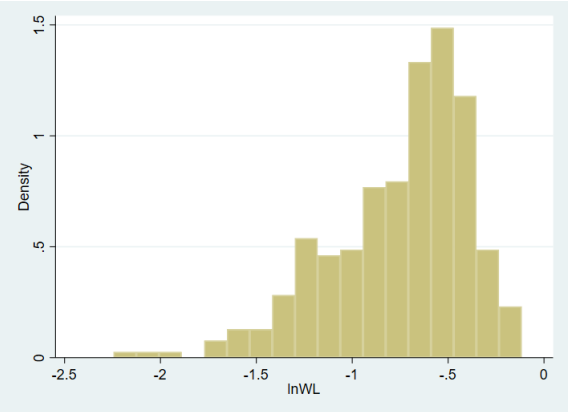
Table 4.3: Descriptive Statistics of Team Performance Measures

	Mean	Std Dev.	Skewness	Min.	Max.
WL	0.500	0.157	-0.146	0.106	0.890
lnWL	-0.752	0.363	-0.954	-2.244	-0.117
PDG	0.001	4.722	-0.129	-13.900	11.600
SRS	-0.002	4.596	-0.136	-13.960	11.350
ORTG	104.619	3.599	-0.031	92.553	113.464
DRTG	-104.600	3.081	0.282	-111.855	-95.887

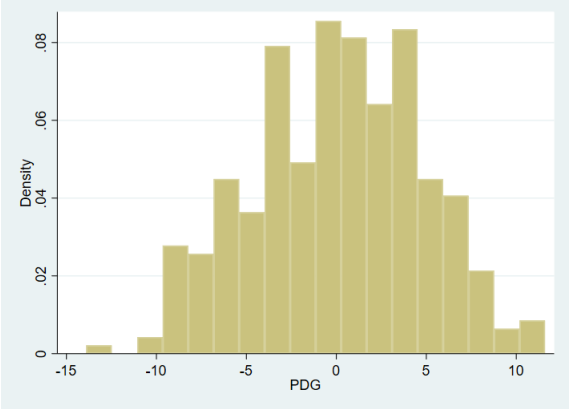
Figure 4.1: Histograms of Team Performance Measures



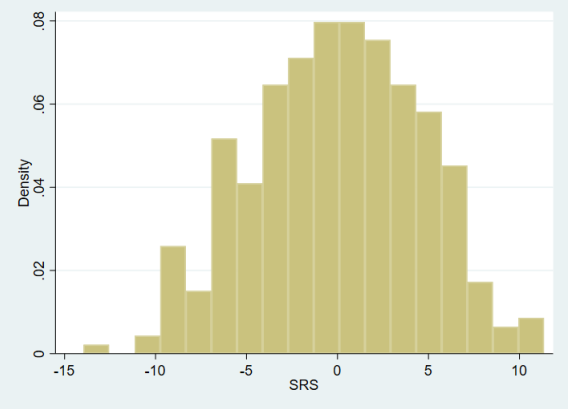
(a)



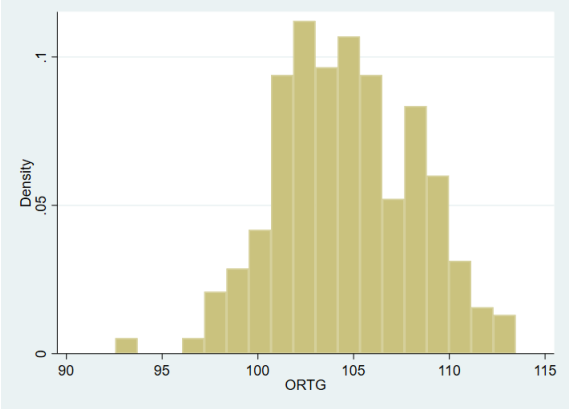
(b)



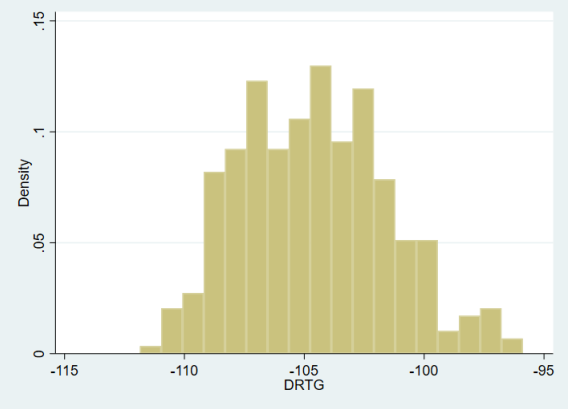
(c)



(d)



(e)



(f)

The overall statistics of PDG and SRS, displayed in Table 4.3, are very similar. Both averages are virtually zero, with slight differences occurring only due to rounding of per-game

averages. However, Figures 4.1c and d reveal slight differences in the distributions. The ratings system appears to reallocate amongst the distribution based on schedule strength. *SRS* is even more clustered around zero.

As *ORTG* is points scored per possession and *DRTG* is points allowed per possession, naturally their averages virtually cancel out over our dataset, with a slight difference again due to per-game rounding.[34] *ORTG* is slightly negatively skewed, while *DRTG* is relatively more positively skewed, though neither is particularly high. There are more instances of above average *DRTG* with a longer right tail, while *ORTG* has a longer left tail. In 4.1e, slight outliers can be seen in the left tail which are the 2011-12 Charlotte Bobcats (who also had the worst regular season win percentage in our sample) and the 2014-15 Philadelphia 76ers. Of interest for our investigation, both of these teams had a reasonably low standard deviation of wage, while the 2011-12 Bobcats had a high national polarisation and the 76ers of 2014-15 had a reasonably high wage polarisation given their low wage standard deviation. Though of course this evidence is only anecdotal. The right tail of *DRTG* in Figure 4.1f is driven by a few observations, though interestingly not the best performing team within a season, the 2015-16 Golden State Warriors. With an *ORTG* of 112.74 and *DRTG* of -102.18 in this season, Golden State rated well on both these measures, particularly on offence, but not at the top of either. This suggests that a combination of good offence and good defence is needed to be successful.

Table 4.4: Correlations Between Team Performance Variables

	WL	lnWL	PDG	SRS	ORTG	DRTG
WL	1.000					
lnWL	0.976***	1.000				
PDG	0.972***	0.949***	1.000			
SRS	0.969***	0.947***	0.997***	1.000		
ORTG	0.759***	0.749***	0.781***	0.790***	1.000	
DRTG	0.659***	0.634***	0.677***	0.662***	0.070	1.000

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.4 shows the pairwise correlations between our measures of team performance. All correlations are positive and all are significant at a 1% level, except pairwise between *ORTG* and *DRTG*. As Kubatko et al. (2007) observes, offensive and defensive ratings are virtually uncorrelated and this is confirmed in Table 4.4 with a correlation coefficient of just 0.07. Naturally, the correlation between WL and $lnWL$ is very high, but this is also the case between these measures, PDG and SRS. The correlation is not as high between these first four measures and *ORTG* and *DRTG*. This should not be a surprise, as each of these alone is not sufficient for a team to be successful. Team performance is determined by a combination of both these ratings, which is elaborated on further in Section 4.3.3 when discussing player performance and particularly in our construction of individual player *Wins Produced* in Appendix 4.B. However, we retain offensive and defensive rating in our analysis to investigate whether the factors we identify influence team offence, defence or both.

Overall, we use the win percentage of each team (WL) as our primary measure of performance of teams over the regular season. The ultimate sporting objectives of teams is to win contests

[34] Small discrepancies in the estimations of possessions employed and possessions acquired, explained in Appendix 4.A.2, may also play a very small role in these slight differences.

and maximise their win percentage across the season. The other measures discussed will be used to test the robustness of our results and capture additional nuances relating to team performance.

4.3.3 Measures of Individual Player Performance

This section describes the measures of individual player performance considered, namely *Wins Produced* (WP), *Win Shares* (WS), *Value Over Replacement Player* ($VORP$), *Wins Produced per 48 minutes* ($WP48$), *Player Efficiency Rating* (PER) and *Box Plus/Minus* (BPM). The first three of these are measures of overall contribution of players towards the performance of teams. The last three are efficiency measures which capture how efficiently a player contributes given his on-court opportunities. It is possible to convert efficiency measures into measures of overall contribution by multiplying by the minutes played of each player, or to convert the other way through division - however, in some cases this will distort the scale of the measure. These links are clear in the variables chosen, particularly between WP and $WP48$, but also between BPM and $VORP$, as will be demonstrated in our discussion.

The primary measure of individual player performance considered in our investigation is Wins Produced (WP), constructed using the methodology of Berri et al. (2006). An outline of this method is described here, with additional detail, including full regression results and calculations, presented in Appendix 4.B. After regressing win percentage on various specifications of statistics, Berri et al. (2006) realised that the marginal effects of all of the relevant performance statistics can be obtained from a simple regression of $ORTG$ and $DRTG$ on team win percentage. From this, one can calculate the marginal effect of scoring or allowing an additional point, and employing or acquiring an additional possession. These can then be used to find the marginal effects of scoring, possession and what Berri et al. (2006) call crime and rejection statistics. Here we describe these marginal effects conceptually; a more technical explanation is given in Appendix 4.B.

For scoring statistics, a successfully made three point field goal increases points scored by three while increasing possessions employed by one. A missed three point attempt increases possessions employed without changing points scored. A similar story holds for two point field goals. Scoring by the opponent has a similar effect, but on $DRTG$ rather than $ORTG$, as they lead to points for the opponent and possessions being acquired by the team. Free Throws are slightly different as, in many cases, multiple free throws are awarded so not every attempt will end a possession employed. However, the proportion that ends a possession can be estimated, so a free throw is one point scored minus the probability that it ends the possession. For possession statistics, offensive rebounds extend a possession and prevent the preceding shot from increasing possessions employed by one, so are worth the negative of a possession employed. Committing a turnover increases the possessions employed by one, while defensive rebounds and opponent's turnovers (including steals) increase possessions acquired by one.[35]

Crimes and rejections are personal fouls and blocks. These are not explicitly accounted for in the model, but can still be calculated. Personal fouls replace opponent's free throw successes, but not every personal foul leads to a successful free throw. However, the proportion that do can be estimated using regression analysis, as described in Appendix 4.B. The marginal

[35] It is also relevant to account for what Berri et al. (2006) call 'team rebounds' which are an accounting device for when missed attempts go out of bounds, for example. Team rebounds are discussed further in our discussion about estimating possessions in Appendix 4.A.2. They are also worth a possession acquired.

effect of a personal foul is then the probability that it leads to a successful free throw for the opponent multiplied by the marginal effect of that free throw on win percentage. Blocks are a little more complicated still, but again regression analysis is used to solve the problem. Blocks prevent potentially successful shots of the opponent and their effect on this can again be estimated via regression analysis as described in Appendix 4.B. But field goal successes of the opponent are already accounted for, so when analysing blocks, we consider the number of blocks compared to the team average. Blocks therefore do not impact the evaluation of team performance, but reallocate the blame for (or credit for preventing) opponent field goals amongst players. The marginal effects of the relevant statistics are presented in Table 4.5. The estimates of Berri et al. (2006) are presented in an additional column for comparison.

Table 4.5: Marginal Effects of Player Statistics on Team Win Percentage

	Marginal Effect	Berri et al. (2006)
POINTS AND POSSESSION		
Points Scored	+0.032	+0.033
Points Allowed	-0.032	-0.033
Possessions Employed	-0.034	-0.034
Possessions Acquired	+0.034	+0.034
SCORING STATISTICS		
Successful 3pt Field Goal	+0.063	+0.066
Opponent Successful 3pt Field Goal	-0.063	-0.066
Successful 2pt Field Goal	+0.031	+0.033
Opponent Successful 2pt Field Goal	-0.031	-0.032
Successful Free Throw	+0.018	+0.018
Opponent Successful Free Throw	-0.017	-0.018
Missed Field Goal	-0.034	-0.034
Missed Free Throw	-0.015	-0.015
POSSESSION STATISTICS		
Offensive Rebound	+0.034	+0.034
Turnover	-0.034	-0.034
Defensive Rebound	+0.034	+0.034
Team Rebound	+0.034	+0.034
Opponent turnover	+0.034	+0.034
Steal	+0.034	+0.034
CRIMES AND REJECTIONS		
Personal foul	-0.018	-0.018
Block	+0.019	+0.021

These marginal effects are then applied at the player-level. Statistics are attributed directly to individuals where possible (scoring: $p3PS$, $p2PS$, $pFTS$, $pFGM$ and $pFTM$; possession: $pORB$, $pDRB$, $pSTL$ and $pTOV$; and crimes and rejections: pPF and $pBLK$). However, this is not always possible, mainly for defensive statistics ($o3PS$, $o2PS$ and $oTOV - tSTL$ - opponent's turnovers which are not steals); but also rebounds and turnovers which cannot be attributed to individual players ($tmRB$ and $tmTOV$). These remaining statistics are attributed to players based on the proportion of the team's minutes that the player played. This appears to be a reasonable approach to apportioning these statistics, especially when considering the cooperative nature of zone defence which is becoming more prevalent in the NBA.[36]

[36] Interestingly, the increased popularity of zone defence looks to be the result of European influence in the NBA (Parker, 2015).

After multiplying these marginal effects by the relevant statistics, the positional averages[37] are then subtracted to give relative performance. On average, each team wins half their games so each player's expected wins produced is half the games played by the team multiplied by the proportion of team minutes played by the player. So finally, expected wins produced is added to give absolute wins produced. We make one additional adjustment to normalise this measure to an 82 game season. As not all seasons in our data consist of 82, we divide by the number of games played and multiply by 82 to allow for cross-season comparisons.

Win Shares (WS) is another similar measure of overall contribution offered by Basketball Reference. Win Shares are made up of two components: Offensive Win Shares (OWS) and Defensive Win Shares (DWS), and is just a simple summation of the two. Both measures calculate a marginal contribution towards the points for/against the team and then divide this by a *marginal points per win*. This is calculated based on league points per game and the ratio of the team's pace to that of the league. The *marginal offensive contribution* of each player is defined as the points produced by the player, minus a league-wide expectation of points produced, given the number of offensive possessions of the player. The *marginal defensive contribution* is a league-wide expectation of the points against, given the portion of the team's minutes played by the player and the defensive possessions faced, minus the player's defensive rating. Higher values for each of these marginal contributions indicate better performance. The reason for dividing the marginal contributions by *marginal points per win* is to account for how these points are converted into wins, which will vary by team. This measure is recorded in absolute number of win shares, so we again make the same adjustment as we did for WP to normalise it to an 82 game season. While negative values are less common, they are possible and should be interpreted as performance so poor that it takes away from the win shares generated by teammates. In our dataset, there are 1,127 instances of negative marginal offensive contributions, while only 10 for marginal defensive contributions. Overall, there are 472 negative Win Shares out of the 5,213 observations, equating to roughly 9%.

Box Plus/Minus (BPM) is a per-minute, efficiency measure which has developed in various steps over time. Firstly, a version of *Adjusted Net Plus/Minus* (APM) is calculated, which differs slightly in that it is calculated from box scores rather than play-by-play data, where each observation is a period of the game where no substitutions occurred. Otherwise a similar methodology to Paine (2012) and Myers (2013), as pioneered by Rosenbaum (2004), is used to create an APM measure. However, in addition, Ridge regression (aka Tikhonov regularisation) techniques (Golub et al., 1999; Hoerl & Kennard, 1970; Tikhonov & Arsenin, 1977) are used, as suggested by Sill (2010), in an attempt to solve the problem of over-fitting (and multicollinearity). Though this biases the estimates towards zero, it reduces the magnitude of the aforementioned problems and reduces the size of the error term. This improves its ability to make out-of-sample predictions which is important for the construction of this measure.[38] The resulting intermediate statistics is known as $RAPM$.

The general idea of BPM is to take these historical (2001-14) $RAPM$ statistics as a target, weight them to adjust for the upward bias of players with a small sample size and regress them on the factors outlined by Equations 4.D.1 and 4.D.2 of Appendix 4.D to produce weights for

[37] These are calculated based on the number of minutes each player plays in each position.

[38] Sill (2010) also argues that it is important to NBA teams which wish to use these types of measures to predict the future performance of their current players and prospective recruits.

various individual statistics including a measure relating to minutes played.[39] These weights are then used to produce BPM. This methodology was conducted by Basketball-Reference and the statistics are collected from their website. The mean of BPM is just less than -2 which is related to its construction and the weighting system used (Myers, 2014).

Table 4.6: Descriptive Statistics of Individual Player Performance Measures

	Mean	Std Dev.	Skewness	Min.	Max.
WP	2.308	3.857	2.137	-5.624	29.516
WS	2.335	2.878	1.765	-2.100	20.300
VORP	0.547	1.257	2.748	-1.600	12.400
WP48	0.057	0.188	0.945	-2.031	3.077
PER	12.619	6.786	1.470	-41.700	133.800
BPM	-2.114	4.749	-1.709	-59.000	54.400

Value Over Replacement Player is simply an adjustment of BPM to covert it from a per-minute efficiency rating to a measure of overall contribution. The process of doing so is given in Equation 4.3. Essentially, the value of a replacement player is assumed to be minus two (Myers, 2014) so this is subtracted. Then the efficiency score is multiplied by the player's proportion of the team minutes, before lastly making an adjustment in the case that it was not an 82 game season as usually assumed. The resulting measure is an estimate of the additional points differential per 100 possessions that a player contributes above that of a standard replacement player. As a replacement player is estimated to contribute a points difference per 100 possessions of minus two, it may not be surprising that $VORP$ does not have a low minimum, but still around 35% of the players in our dataset received negative $VORP$ scores, though only marginally negative. $VORP$ has some nice qualities such as it can be interpreted linearly, in that a player with a $VORP$ of four can be considered twice as good as a player with a $VORP$ of two, or at least this seems to be the case in terms of market values (Myers, 2014). It can also be multiplied by 2.7 to estimate the contribution to wins.

$$VORP = [BPM - (-2.0)] \cdot \left(\frac{pMP}{tMP}\right) \cdot \left(\frac{tG}{82}\right) \tag{4.3}$$

Player Efficiency Rating (PER) is a per-minute, *Linear Weights* measure developed by Hollinger (2002, 2005). It is designed to sum up all a player's positive accomplishments, subtract the negative accomplishments, and give a per-minute rating of a player's performance. The positive accomplishments include three pointers, assists, steals and blocks among others, while the negative accomplishments include field goal misses, free throw misses and turnovers; all of which are weighted according to importance, league averages and percentage successes. A 'pace' adjustment is then made for team possessions per 48 minutes, and finally standardised by the league average. Full equations are given in Appendix 4.C along with accompanying discussion. Summary statistics of the measure are given in Table 4.6. The resulting measure can be positive or negative and ranges between -41.7 and 133.8 in our dataset. Again, positive ratings are more common, with only 152 of the observations (roughly 2.5%) being negative. The mean across our dataset is around 12.7. It should be noted that the highest two scores (133.8 and 129.1) are anomalies and are over 40 higher than the 3rd highest score of 88.3. They are both scored by players (Naz Long of the Utah Jazz in 2017-18

[39] The inclusion of minutes played is controversial to some and is later corrected for in $VORP$.

and DeAndre Liggins of the Miami Heat in 2013-14) who played less than one minute in one game but, due to making positive contributions in this time, were judged to be extremely efficient. This is the danger of only accounting for efficiency, as important as it is.

Figure 4.2: Player Performance Measures and Minutes Played

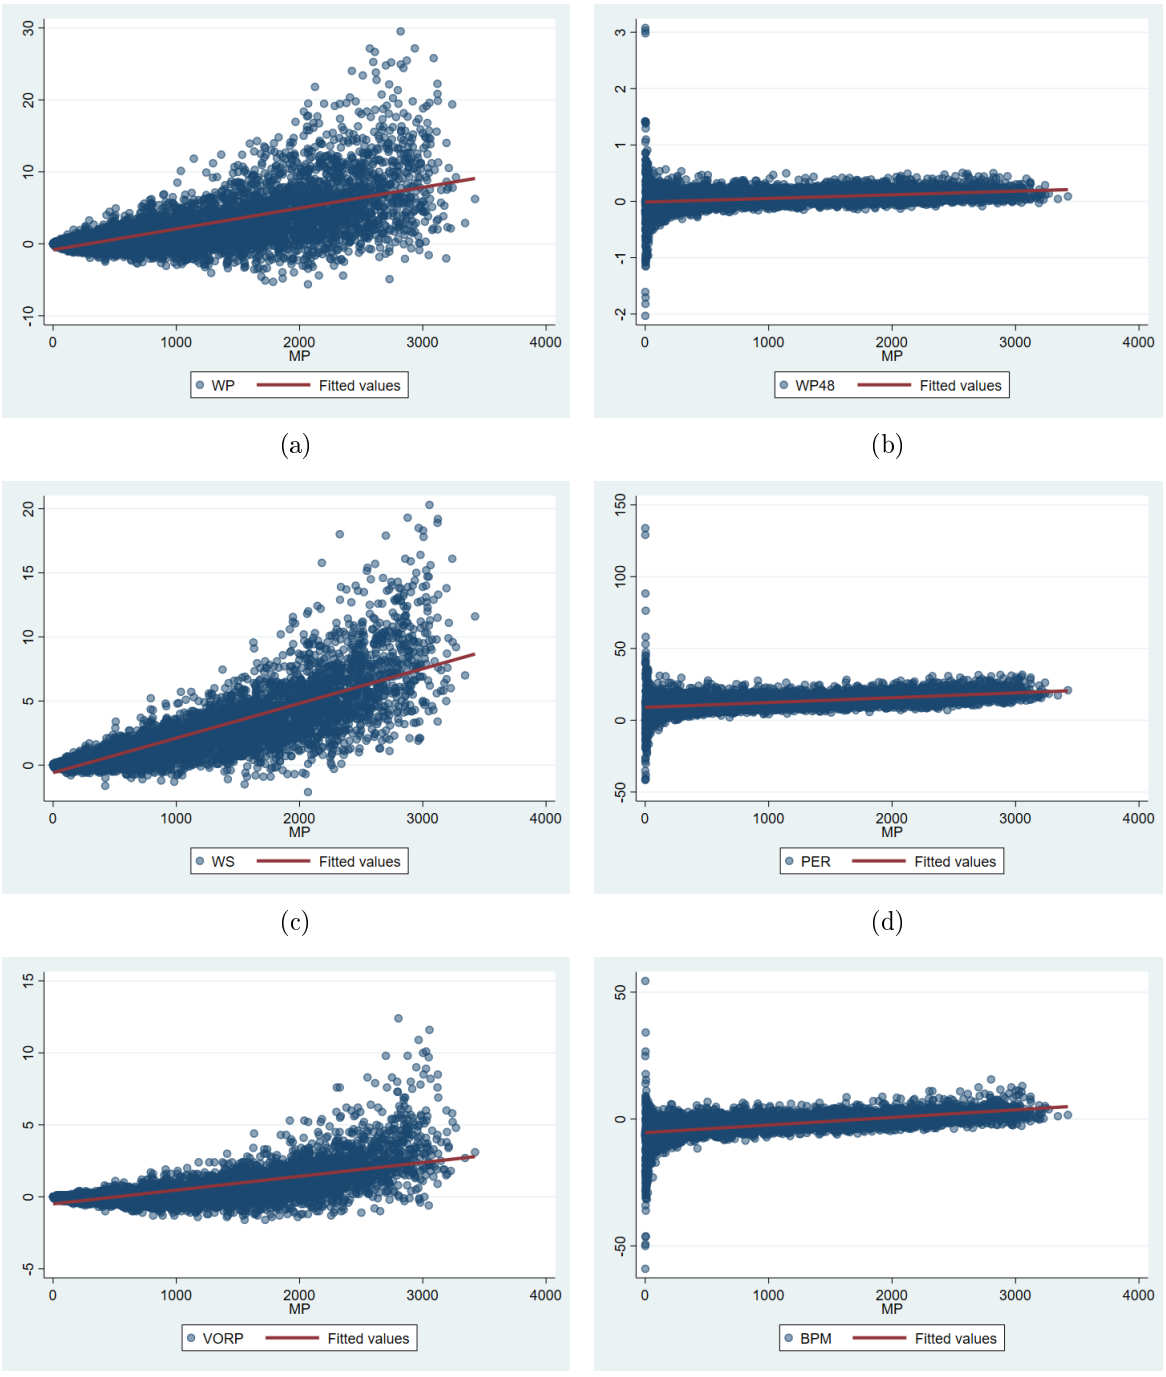
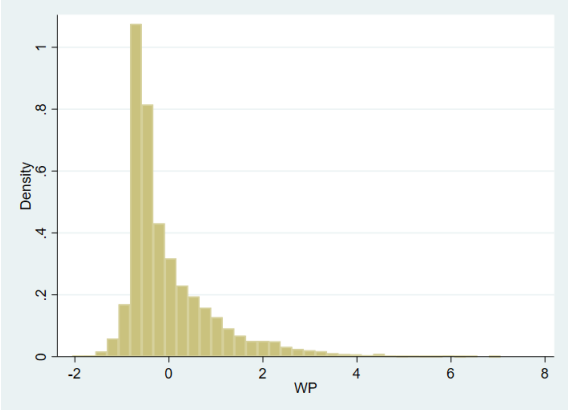


Figure 4.2 presents scatterplots of the individual player performance measures against minutes played. The overall contribution measures in the left column of Figure 4.2 (a, c and e) show a general positive trend with minutes played, though the variation is heteroscedastic in nature. There is more variation in player performance for players who play more minutes, which seems natural. The efficiency measures are presented in the right column of 4.2 (b, d and f) and also show a general upward trend, though the slope is significantly less steep. Most of the variation in these variables is at very low levels of minutes played, with the most extreme values occurring at close to zero. This again shows the fallacy of accounting only

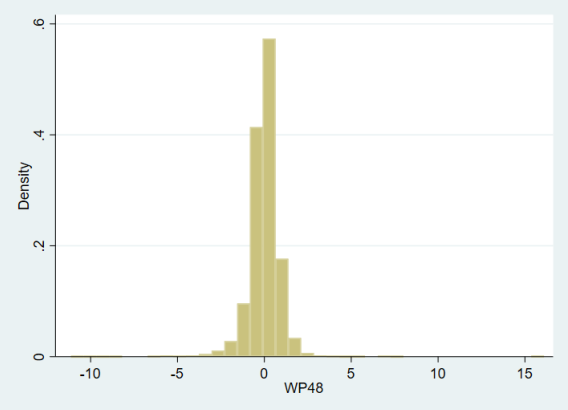
for efficiency, with some players appearing very efficient for short periods of time, but not showing this consistently - arguably just through a lack of opportunity, but more likely due to an inability to do so.

The standardised distributions of the measures of player performance are presented in Figure 4.3. As the variables are measured on different scales, standardised distributions were chosen for comparability.

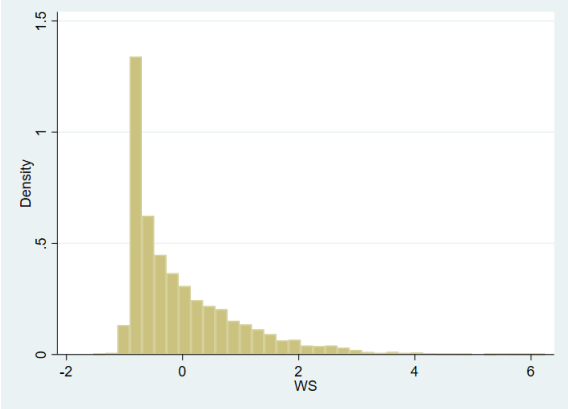
Figure 4.3: Histograms of Standardised Player Performance Measures



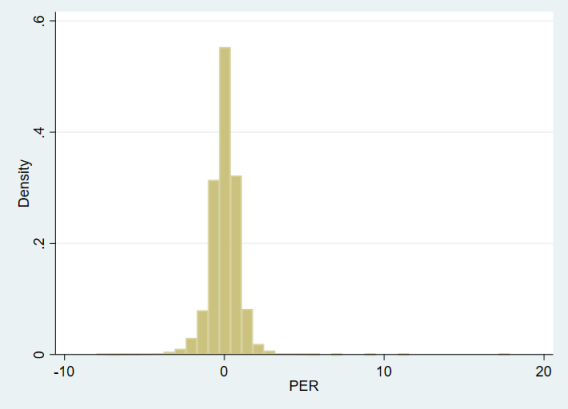
(a)



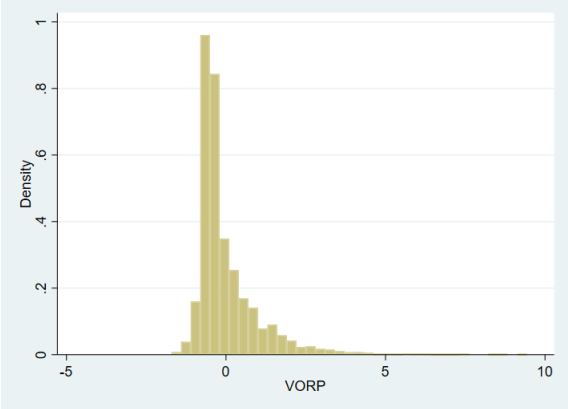
(b)



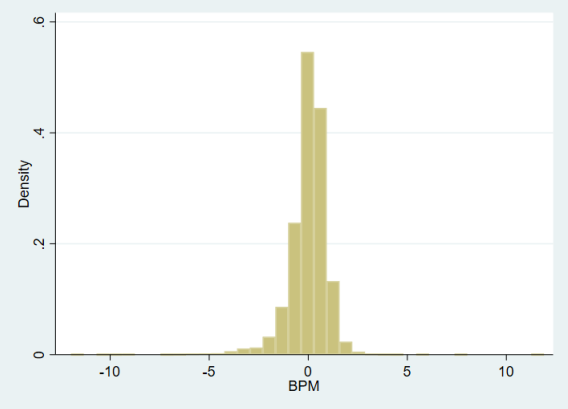
(c)



(d)



(e)



(f)

Maybe not surprisingly, the overall contribution measures, presented in the left column of 4.3 (a, c and e) are highly positively skewed with a large majority clustered around zero. The efficiency measures (b, d and f) look to be more symmetrically distributed, though there are some relatively extreme values, particularly for PER and $WS48$.

Overall, there seems to be a question of whether efficiency or overall contribution should be valued more when looking to quantify individual player performance. WP, WS and $VORP$ capture overall contribution, while $WS48$, PER and BPM capture efficiency.

Table 4.7: Correlation Between Individual Player Performance Measures

	WP	WS	VORP	WP48	PER	BPM
WP	1.000					
WS	0.871***	1.000				
VORP	0.802***	0.908***	1.000			
WP48	0.498***	0.409***	0.352***	1.000		
PER	0.486***	0.544***	0.483***	0.829***	1.000	
BPM	0.530***	0.608***	0.584***	0.696***	0.788***	1.000

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4.7 confirms that there are much higher correlations within the groups of efficiency measures and overall contribution measures, respectively. While all the correlations are positive and statistically significant at a 1% level, it is clear that the correlations within these groups are stronger than across them. By considering all these measures of individual player performance, we can account for the various aspects of individual performance, most predominately overall contribution and efficiency.

4.4 Covariates

This section discusses covariates between our dependent variables and variables of interest which will be used as control variables. We explain how and why these variables may affect our statistical inferences and therefore why it is relevant to control for them in our regression analysis. Descriptive statistics and correlations amongst these control variables are presented in Section 4.4.1 for the team-level variables, and Section 4.4.2 for the player-level variables. Correlations between our variables of interest and these covariates are presented in the relevant empirical Chapters: 5, 6, 7. The relationship between these covariates and our dependent variables are explored through baseline regressions, which are presented at the team-level in Appendix 4.F and at the player-level in Appendix 4.I. Regression analysis is also utilised to show the importance of the covariance between these variables and our variables of interest in Appendix 4.H at the team-level and Appendix 4.J at the player-level. In this section, we also discuss the potential reverse causality between wage and performance which would cause the wage to be endogenous in our latter regression analysis. Ultimately, we decide to retain wage due to its importance as a control variable (and for use in interaction terms), and because its inclusion does not seem to significantly affect the statistical inferences of the other regressors.

While the heterogeneity between teams and players is partly addressed by utilising the econometric techniques discussed in Section 4.2, such as panel data techniques and multilevel modelling, we also supplement this through the inclusion of relevant control variables. This removes biases on the coefficients of interest, caused by potential correlations between the remaining heterogeneity and the main explanatory variables. Some of these variables have been considered in previous studies (with wage expenditure being the most commonly used), but we also introduce some additional variables. The control variables used in our analysis can be classified, by the performance-affecting element they control for, into four groups:

player quality/financial resources, in-season roster stability, between-season roster stability, and player experience. In what follows, we discuss why these variables are important determinants of performance, but also why they are likely to be correlated with the measures of wage dispersion and national diversity, and why it is therefore relevant to control for them to isolate the effects of our variables of interest.

The first element that we control for is the financial resources of the team. Although NBA teams operate under a salary cap, there are exceptions in these rules, which makes it what is known as a 'soft cap'. We expect teams which are willing and able to spend more on wages to attract higher quality players and therefore increase their win percentage. We also expect teams which spend more on wages to have higher wage dispersion within their teams. This is due to the Theory of Superstars (Rosen, 1981), which observes that the top talents, or superstars, are paid in excess of their on-field productivity because of the advertising and commercial opportunities they bring. Therefore, teams which can recruit these superstars are likely to have higher wage budgets as well as more wage dispersion. If we do not control for the level of wage expenditure, then we may see a correlation between higher wage dispersion and higher performance, simply because these are both characteristics of richer teams which can afford to sign superstars. By controlling for wage expenditure, we find the effect of wage dispersion for a given wage budget.

In terms of national diversity, teams which spend more on wages may also be more successful in attracting the top talents, wherever these players may hail from. It may take lucrative offers to tempt players to move away from their homeland and home league to pursue a career in the NBA. It is therefore the top international talents that are more likely to move to the NBA rather than the mediocre players. On the other hand, many of the top talents and household names in the NBA such as Kobe Bryant, LeBron James and Stephen Curry are all from the USA. Whether teams have low national diversity through pursuing the top or mediocre domestic talents, it is relevant to control for the level of wage expenditure to ensure that this is not conflating our statistical inferences regarding national diversity. At the team-level, we use the variable $TPayroll$, which is the total wage expenditure of each team in each year. At the player-level, we consider each players annual income, $Wage$, as well as considering their rank in terms of wage within their team, $WageRank$.

Other studies have investigated the relationship between wage and individual player performance, but using wage as the dependent variable and identifying the individual statistics which influence it. Typically, anywhere from around 50% to approaching 75% of the variation in pay can be explained by individual performance statistics (often lagged) alongside factors such as age/experience (e.g. Berri et al., 2006; Simmons & Berri, 2011; Stiroh, 2007). This methodology is commonly used to investigate (particularly racial) discrimination of wages in the NBA (e.g. Groothuis & Hill, 2013; J. R. Hill, 2004; Kahn & Sherer, 1988). While this strand of literature is relevant for the current study in terms of the link between national diversity and race, the general topic of wage determination also brings into question the endogeneity of wages through simultaneity and therefore the direction of the causality. While wage is determined by past performance and wage in turn influences current performance, the potential correlation of the dependent variable and the lags of a control variable is liable to present potential issues. However, the importance of wage as a control variable and for the interaction terms used in this investigation may outweigh these endogeneity concerns de-

pending on their magnitude. In the absence of relevant instruments, we conduct a number of robustness checks throughout our investigation. These include the exclusion of control variables and considering the results from different econometric models[40], as well as conducting Hausman tests.

Next, we control for the level of in-season roster stability. Teams may be prevented from fielding their strongest team consistently throughout the season due to several reasons including injuries, trades and possibly suspensions. We expect teams with higher levels of in-season roster stability to perform better. However, if being unable to field players means that teams are forced to sign free agents or promote from their development teams, then the presumably lower wages of these players may affect the levels of wage dispersion within the team. If the probability of these players being born in the USA is higher than the league average then it may also have an influence on national diversity within the roster. This relationship with both performance and our variables of interest may bias our estimates of the relationships between them, so we therefore control for the level of in-season roster stability.

We do this at the team-level using the variable $CTSize$ which measures the size of the 'core team' defined as those players which played at least 20 games and at least 12 minutes per game throughout the season. This definition of the 'core team' is identical to that used by Simmons and Berri (2011). A larger 'core team' indicates less in-season roster stability. We also consider the size of entire roster in a variable called $Size$. However, this includes players which do not play often and we find $CTSize$ to have a more robust relationship with performance. At the player-level, we consider a dummy variable, $Trade$, equal to one if the player was traded during the season.

We also control for the level of between-season roster stability. When players have not played together long, they may be unfamiliar with each other's playing style and struggle to cooperate effectively. This could be further exasperated in nationally diverse teams where culture and language barriers exist. We may also expect international players to have shorter tenures if they join after playing in other national leagues. It should also be noted that one of the rules of the soft wage cap which the NBA operates under, is that teams may exceed the cap under certain conditions to resign their own free agents. This means that teams with higher between-season roster stability may spend more on wages and may have a more dispersed wage structure. We also speculate that between-season roster stability, wage dispersion and national diversity can all affect team morale which in turn affects performance, so these issues may be interlinked. The level of between-season roster stability is captured by measuring the average tenure of the team; that is, how long each player has played for his current team, prior to the current season. We calculated this in terms of seasons, games and minutes played, and after multiple trial specifications, we favour the variable measuring minutes played. The team-level average of this variable is called $TTenureMP$, while the player-level variable is called $PTenureMP$.

Another variable we control for is the level of experience within the team. Scully (1995) notes that back-loaded pay schedules are a characteristic of labour markets which is particularly pronounced in the sports industry. If wages are closely correlated with experience, and players recognise and accept this, then using player wages as an indication of their quality is only

[40] Namely Random Effects (RE), Fixed Effects (FE) and Multilevel Modelling (MLM).

justified when controlling for their level of experience. Also, similar to the argument for including tenure, there are exceptions in the soft wage cap for veterans which may affect wage dispersion. Levels of experience may also differ between international and domestic players, and may therefore be affected by the level of national diversity within teams.

We calculate experience as the number of years since the player's NBA debut. Again, we also had experience in games and minutes played, but favour years in this instance as this is what affects a player's veteran status. This choice also decreases the expected correlation between experience and tenure due to the different choice of measurement. We call this variable $TExp$ for the team-level average and $PExp$ for the player-level variable.

Additionally, we also control for the dispersion of experience within the team. If experience is highly correlated with wages as theorised, then by only measuring the dispersion of wage and not controlling for the dispersion of experience, we will be capturing the effects of both simultaneously. It will therefore not be clear whether a result is driven by differences in wages within a team or differences in experience. That is, whether a result is caused by high wage dispersion or a mix of youth and experience within the team. A similar story can be told in terms of national diversity, where it could be unclear whether a result is driven by national diversity or a mix of youth and experience. As we wish to isolate the effects of our variables of interest, it is relevant to control for the dispersion of experience. We do this by controlling for the standard deviation of experience within the team, $ExpSD$. The expected effect of the standard deviation of experience on performance is unclear and presents the debate whether it is beneficial to have a mix of youth and experience or to have players of similar age, presumably preferably in their prime. Importantly for our latter analysis, we also consider experience polarisation, but do not find this to be important.

Finally, we note that we also considered the age of players within our dataset. However, alongside experience and tenure, it did not explain any additional variation. Therefore, though we trialled its inclusion in different specifications, we chose not to include it as a control, retaining experience and tenure instead.

4.4.1 Descriptive Statistics and Correlations - Team-level

Table 4.8 shows some basic summary statistics of our team-level control variables, where we also consider the 'between' variations across teams and the 'within' variation for teams over seasons. The average team expenditure on wages by NBA teams ($TPayroll$) in our dataset was around \$76.5m, which is roughly \$10m over the average wage cap for the period (roughly \$66.5m). This shows that while the wage cap limits spending in some way, it is a soft cap so does not set a strict limit. The standard deviation for $TPayroll$ given in Table 4.8 is \$17.82m. This is the standard deviation of the total expenditure between team-year observations, not the standard deviation of wages within each team later used as a key variable of interest in the empirical analysis of wage dispersion. The variance is higher within the year dimension than between the team dimension, which may be expected due to the wage cap which is adjusted yearly. However, the between variation is still large and there is enough variation to justify controlling for this variable in our analysis. $TWage$ represents the mean wage of players within each team, which on average is \$4.46m.

NBA rules for the period allow a maximum roster of 15 active plus two inactive players, with a minimum of 13 (12 active plus one inactive) players, at any one time. Additionally, since

the 2005 Collective Bargaining Agreement (CBA), the league guarantees a league average of at least 14 players and is subject to a surcharge for not meeting this obligation (Coon, 2005, 2011, 2017). Our measure of the roster size differs slightly in that we consider all players which play for each team throughout the season, while the restrictions above apply to the roster size at any given point during each season. It is therefore logical that the average of our measure of the roster size is higher, at almost 18. The smallest roster was that of the 2009-10 Orlando Magic with 12 players, which is just within the restrictions mentioned. The average size of the core team (those playing more than 20 games and more than 12 minutes per game) is roughly 12, which seems reasonably for a 5-a-side game played with frequent substitutions and given the roster restrictions mentioned.

Table 4.8: Descriptive Statistics of Team-level Control Variables

Variable	Panel	Mean	Std Dev.	Min.	Max.
TPayroll ($m)	Overall	76.486	17.824	45.268	137.723
	Between		5.355	67.919	86.846
	Within		17.026	43.745	134.649
TWage ($m)	Overall	4.458	1.111	1.636	8.378
	Between		0.389	3.447	5.308
	Within		1.043	2.210	8.095
Size	Overall	17.976	2.647	12	28
	Between		1.061	15.909	19.909
	Within		2.433	12.430	26.612
CTSize	Overall	12.042	1.746	8	18
	Between		0.522	10.909	13.182
	Within		1.669	8.042	17.224
TTenureMP^{+}	Overall	2.537	1.327	0.326	8.014
	Between		0.909	1.333	5.724
	Within		0.980	-0.852	5.438
TExp	Overall	4.853	1.308	1.696	8.895
	Between		0.689	3.774	6.495
	Within		1.119	2.150	8.448
TExpSD	Overall	3.970	0.806	1.867	6.351
	Between		0.462	3.109	5.009
	Within		0.665	2.242	6.609

$^{+}$in thousands of minutes

The average tenure of the players in our sample was just under 2,537 minutes, with a high overall standard deviation of just over 1,327 minutes. From the other tenure variables in our dataset, we calculate that these minutes were played in an average of roughly 90 games, over an average of a little under one and half seasons. Given a maximum number of 82 games possible in a season and a maximum of 48 minutes per game, these numbers seem reasonable, if a little lower than expected. There appears to be a large amount of variation in this variable both between and within, suggesting tenure in terms of minutes played both amongst the different teams and for the teams over different seasons. It should be noted that the minimum and maximum given in Table 4.8 are for the average tenure of the teams, not for a single player. There are indeed rookies and new signings with tenures of zero minutes and players with tenures longer than the maximum team average of just over 8,000. Almost 10% of the sample (575 out of 5,932) have tenures longer than this average of team averages.

The average NBA experience of players within each team is almost five years. Given the average career span of professional athletes, this figure seems reasonable. Again, to be clear, the results here are the team averages. The maximum of almost nine years is the maximum team average, while the highest experience of a single player in our data was Kevin Garnett with 20 years of NBA experience by the 2015-16 season in which he played for the Minnesota Timberwolves. The standard deviation of experience is comparatively lower than that of tenure at just over a year. This variation is again apparent both between and within the panels. This is the variation between the team averages, whereas the average standard deviation of experience within observations is just under four years.

Table 4.9: Correlations Between Team-level Control Variables

	TPayroll	TWage	Size	CTSize	TTenureMP	TExp	TExpSD
TPayroll	1.000						
TWage	0.788***	1.000					
TSize	0.154***	-0.208***	1.000				
CTSize	0.086	-0.059	0.440***	1.000			
TTenureMP	0.162***	0.371***	-0.330***	-0.189***	1.000		
TExp	0.267***	0.439***	-0.180***	-0.131**	0.553***	1.000	1.000
TExpSD	0.292***	0.283***	0.009	0.028	0.425***	0.663***	1.000

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

When considering the correlations among our team-level control variables presented in Table 4.9, the largest correlation is between the total and average expenditure of teams on wages ($TPayroll$ and $TWage$). The correlation is not perfect as rosters are different sizes, but it is no surprise that the correlation is high and significantly positive. The second highest correlation is between the average experience ($TExp$) and the standard deviation of this experience within teams ($TExpSD$) at 0.663. Due to the draft system in American sports, teams are likely to take on rookies each year, so it is logical that teams with more experienced players in their rosters will have higher average experience, as well as higher standard deviation of experience. As expected, there is also positive correlations between the average experience of teams and their payroll and average expenditure, with correlation coefficients of 0.267 and 0.439, respectively. While these are relatively low, we believe that the validation of our main results is increased by also controlling particularly for experience and the dispersion of this experience. We also see correlation between experience and tenure as would be expected, but this is low enough to suggest they are capturing distinct effects. The size of the core team is relatively less correlated to the over control variables. This could suggest that the size of the core team is associated with an element of luck in terms of injuries/form/suspensions.

4.4.2 Descriptive Statistics and Correlations - Player-level

Table 4.10 shows the descriptive statistics of the player-level control variables. Firstly, we consider the wage of each player ($Wage$). The mean across all players is \$4.42m.[41] Naturally, we expect increased labour demand for high performing players, as well as lower supply, resulting in a higher wage, and vice versa for low wages (Scully, 1974). Player's wages will also have an effect on the wage dispersion of the team. Players with high or low wages are therefore more likely to be in teams with higher wage dispersion. Also, due to the work conducted on

[41] This varies slightly to the mean of the average wage of players within teams ($TWage$), presented in Table 4.8, due to varying team sizes.

racial discrimination (e.g. Groothuis & Hill, 2013; J. R. Hill, 2004; Kahn & Sherer, 1988), we may also expect some correlation between nationality and wage. So, when analysing the effects of wage dispersion and national diversity on the performance of players, it would seem to make sense to control for the individual wages of the players, despite the aforementioned endogeneity concerns. Additionally, interactions between wage and particular wage dispersion might be informative if the effects depend on wage.

We also consider the rank of each player's wage within their team ($WageRank$), the descriptive statistics of which are also presented in Table 4.10. We have multiplied the wage rankings by minus one, so that a higher value indicates players with a 'better' wage ranking within their team.[42] This variable should not suffer from endogeneity to the same degree as $Wage$. Although we may expect some causation between a player's past performance and their wage ranking within their team, a player's rank is also dependent on their teammates. We therefore experiment with controlling for $Wage$ and $WageRank$, and also find an interesting result when controlling for them simultaneously.

Table 4.10: Descriptive Statistics of Player-level Control Variables

	Mean	Std Dev.	Min.	Max.
Wage ($m)	4.424	5.119	0.005	34.683
WageRank	-9.636	5.484	-27	-1
Trade	0.234	0.424	0	1
PTenuremp^{+}	2.472	4.989	0	48.673
PExp	4.819	4.151	0	20

$^{+}$in thousands of minutes

Next, we consider the variable $Trade$. This is a dummy variable which takes the value 1 if the player appears for multiple teams within a season and zero otherwise. This helps capture when a player's season is disrupted by switching teams within a season. Especially when considering the player performance measures which capture overall contribution, it is important to consider when a player was only with a team for part of the season. Table 4.10 shows that 23.4% of the observations (1,390 out of 5,932) in our sample were for players which played for more than one team during a season. However, by definition, there are multiple observations of these players within a season. The percentage of player-year observations which played for more than one team within a season is 12.9% (671 out of 5,213).

We also consider tenure (as other studies have done; e.g. Brandes et al. (2009)) in terms of minutes played ($PTenuremp$) to capture any effects of players remaining with the same team for a long period of time. Accounting for experience in years and tenure in minutes played also allows us to capture longevity in both of these measurements. Table 4.10 shows that, on average, the players in our sample had played 2,472 minutes for their respective teams. Rookies and players moving to a new team will have a tenure of zero. This accounts for 48.3% of the observations in our sample (2,863 out of 5,932). The player with the most minutes for a single team in our dataset is Dirk Nowitzki who had played 48,673 minutes for the Dallas Mavericks prior to the 2017-18 season.

[42] For clarity, -1 is the highest rank in each team-year combination and therefore the highest paid player in a team, with the rest in descending order.

Lastly, we also control for the experience of each player ($PExp$). In line with other studies (e.g. Berri & Simmons, 2009; Groothuis & Hill, 2008; Ingersoll et al., 2017; Simmons & Berri, 2011), we find it suitable to also include a quadratic term. Due to the back-loaded pay schedules of sportspeople (Scully, 1995), we expect young players to get paid below, and older players to get paid above, their marginal productivity due to risk aversion on the side of the teams. Also, due to rookie and veteran clauses in NBA contracts (Coon, 2005, 2011, 2017), the influence of wages and where a player lies within their team's wage ranking may be dependent on their level of experience within the NBA. We also might expect native players to have more NBA experience as they are less likely to have developed in the European leagues, for example. Table 4.10 displays that, on average, the players in our sample have almost two and a half years of NBA experience. Rookies have zero years of experience and these account for 15.6% of the sample (925 out of 5,932). The most experienced player in our sample was Kevin Garnett who garnered 20 years of NBA experience by the 2015-16 season when he retired.

Table 4.11: Correlations Between Player-level Control Variables

	Wage	WageRank	Trade	PTenureMP	PExp
Wage	1.000				
WageRank	0.783***	1.000			
Trade	-0.097***	-0.163***	1.000		
PTenureMP	0.530***	0.439***	-0.135***	1.000	
PExp	0.439***	0.449***	0.010	0.349***	1.000

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.11 displays the correlations between our player-level control variables. All of the correlations, except between $Trade$ and $PExp$, are statistically significant at a 1% level. Unsurprisingly, the highest pairwise correlation is between $Wage$ and $WageRank$ at 0.783. Although this correlation is high, it is not too close to 1, suggesting additional insight may be gained by considering both of these measures. The negative correlations between $Trade$ and the other variables (except the insignificant correlation with $PExp$) suggests that those players traded during the season are amongst the lower paid players in a team and have lower tenures. This is likely an indication that teams do not wish to disrupt their 'core team' during the season and therefore mid-season trades happen, at least more frequently, for fringe players. Both tenure and experience appear to be positively related to wage, as expected. And, again, despite a positive relationship between the two, tenure and experience seem to be capturing distinct effects with a correlation coefficient of 0.349, which is significantly lower than 1.

Overall, we feel these variables represent a sensible set of controls to help identify the effects of wage dispersion and national diversity for the purposes of our investigation. Through regression analysis, we show these variables are statistically significant determinants of both our measures of performance and variables of interest, at the team-level in Appendices 4.F and 4.H; and at the player-level in Appendices 4.I and 4.J.

4.5 Summary

This chapter has covered the prerequisites to the empirical analysis conducted in Chapters 5, 6 and 7. It has discussed the dataset, econometric techniques, measures of performances (dependent variables) and covariates (control variables) utilised in this investigation. The

team-level variables are used in Chapters 5 and 6, to investigate the effects of wage dispersion and national diversity, respectively, on team performance. Section 6.4, at the end of Chapter 6, also considers a combined model which considers wage dispersion and national diversity simultaneously. Chapter 7 then explores whether additional insights can be gained through player-level analysis, firstly for wage dispersion, then for national diversity, and finally for both combined.

Chapter 5

Wage Dispersion and Team Performance

Chapter Overview

This Chapter investigates the effects of wage dispersion on team performance, with a particular focus on the standard deviation and polarisation of wages. In the current literature, both the theoretical and empirical evidence regarding wage dispersion has produced mixed results, so we contribute a novel reconciliation. Using data from the National Basketball Association (NBA), for the period 2007-18, we compare multiple different measures of wage dispersion and explore their effects on team performance. We find that a higher standard deviation of wages within teams has positive effects on team performance. However, particularly when holding this constant, increased wage polarisation can have negative effects through an increased possibility of conflict within the team. This suggests that wage distributions should be managed in a specific way to maximise team success. The results of this study are robust to controlling for a comprehensive set of control variables and to multiple different measures of team performance. This evidence may explain the mixed results in the current literature through omitted variable bias and/or model misspecification.

5.1 Introduction

In the current literature, several different measures are used to investigate the relationship between wage dispersion and performance. While some of these measures capture similar phenomena, they react in different ways to some changes in the distribution of wage. If these differences affect the dynamics of the relationship between wage dispersion and performance, then using different measures may give different results. We consider measures already used in the current literature including the standard deviation, Gini Coefficient, Coefficient of Variation and the Herfindahl-Hirschman index (HHI) of wages within teams to see if we gain qualitatively consistent results amongst the measures. Additionally, we use measures of skewness and kurtosis to see if these measures offer any further information about the dynamics of wage dispersion within teams. However, our major contribution in this area is the inclusion of polarisation measures. We believe these measures capture an important aspect of the relationship between wage dispersion and performance, specifically conflict and the negative effects predicted by Cohesion Theory (Levine, 1991). Whereas traditionally used measures of spread such as the standard deviation of wage will encourage effort via Tournament Theory (Lazear & Rosen, 1981).

Section 5.2 describes these wage dispersion variables and looks at their summary statistics. Next, we look at the correlations among these wage dispersion variables to see whether they are capturing similar or unique information about the distribution of wages within teams in our sample. This also indicates if it is suitable to include the measures simultaneously. We also look at the correlations between these wage dispersion measures and our control variables. Regression results are then presented in Section 5.3, followed by robustness checks in Section 5.3.1. Section 5.4 discusses and summarises.

5.2 Descriptive Statistics and Correlations

Table 5.1 shows the summary statistics of the wage dispersion measures currently used in the literature and our polarisation measure[43], as well as the level of skewness ($WageSkew$), kurtosis ($WageKurt$). As expected, we find the wage data to be positively skewed – a common characteristic of wage data. However, the minimum value for this variable shows us that not every team had a positively skewed wage distribution. There are only two instances of negatively skewed wage distributions in our data, and these were both for the Minnesota Timberwolves, in 2010-11 and 2011-12. But with win percentages of 17% and 26% respectively, this suggests the strategy was not particularly successful, but of course a full regression analysis is required to confirm this. On average, teams in our sample have excess kurtosis in their wage distributions, implying fatter tails than that of a normal distribution. But again, this is not true of all distributions and we see a large variance in the level of kurtosis. Given our discussion of polarisation and 'measuring the middle', excess kurtosis may seem to indicate a shrinking 'middle' and therefore polarisation. However, kurtosis is a unimodal measure and does not capture local clusterings as polarisation measures do. It is possible for a distribution to have fatter tails without necessarily increasing the degree of clustering at the extremes, so it is unlikely to capture the dynamics leading to increased probability of conflict predicted by polarisation.

[43] Constructed using the methodology of Foster and Wolfson (2010)

Table 5.1: Descriptive Statistics of Wage Dispersion Variables

Variable	Panel	Mean	Std Dev.	Min.	Max.
WageSD ($m)	Overall	4.880	1.581	1.554	10.166
	Between		0.755	3.328	6.609
	Within		1.395	1.751	9.906
WageGini	Overall	0.528	0.0797	0.215	0.724
	Between		0.0443	0.423	0.592
	Within		0.0667	0.320	0.703
WageCV	Overall	1.060	0.216	0.397	1.796
	Between		0.120	0.825	1.249
	Within		0.181	0.579	1.614
WageHHI	Overall	0.122	0.0245	0.0679	0.222
	Between		0.0121	0.100	0.152
	Within		0.0215	0.0593	0.192
WageSkew	Overall	1.349	0.477	-0.746	2.902
	Between		0.202	1.014	1.827
	Within		0.434	-0.462	2.627
WageKurt	Overall	4.040	1.628	1.549	11.604
	Between		0.623	3.130	5.761
	Within		1.508	1.218	10.448
WagePol	Overall	0.209	0.0749	0.0500	0.516
	Between		0.0397	0.136	0.299
	Within		0.0639	0.0481	0.454

The average standard deviation of wages within our team-year observations is roughly $4.88m. Given that the average wage of NBA players in our dataset is around $4.42m, this may seem large, but is also characteristic of positively skewed data, with a few observations to the extreme right of the distribution. There is slightly higher 'within' variation, which should not be surprising due to the yearly changes in the wage cap, but the 'between' variation is only slightly lower implying differences between teams also exist. Table 5.1 also shows an average Gini Coefficient of 0.528, Coefficient of Variation of 1.06 and HHI of 0.122 within our sample. We do not feel it necessary to include the individual formulas for each of these variables here, as they are commonly used measures. Each of these show a good amount of variation, again with the 'within' variation being a little larger than the 'between' variation in each case. The purpose of this study is to investigate whether these measures have positive effects on performance in line with Tournament Theory (Lazear & Rosen, 1981), or negative effects in line with Cohesion Theory (Levine, 1991). However, we propose that these measures are more closely linked to the former and that association with the latter is only due to joint correlation with polarisation measures. Therefore, we expect to see positive relationships, especially when controlling for polarisation.

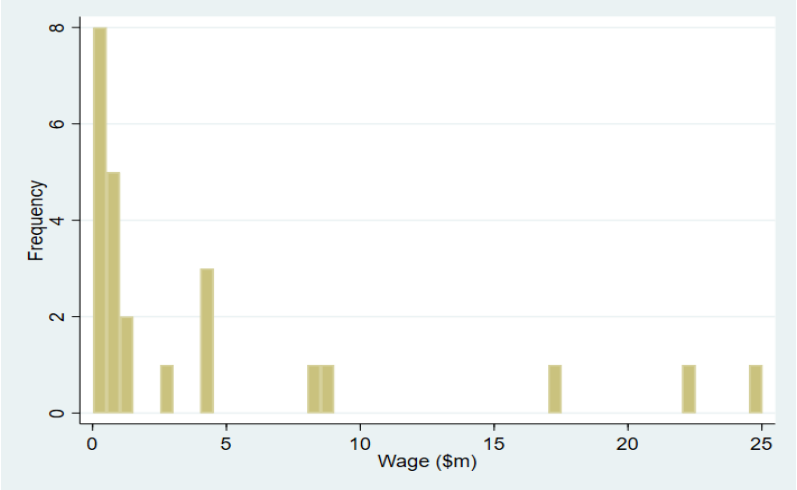
The primary measure of polarisation used in this investigation is constructed using the methodology of Foster and Wolfson (2010). We call this variable $WagePol$[44]. Table 5.1 also shows that $WagePol$ ranges between 0.05-0.516, with a mean of 0.209 and standard deviation of 0.075. As this variable is bounded between zero and one, it is relatively easy to interpret. A team will have a maximum polarisation of 1 when half the team earns zero and

[44] For simplicity we call this measure $WagePol$ in this chapter, though it is referred to as $lnFW$ in Appendix 5.B and 5.C to distinguish it from alternative measures of wage polarisation when comparing them.

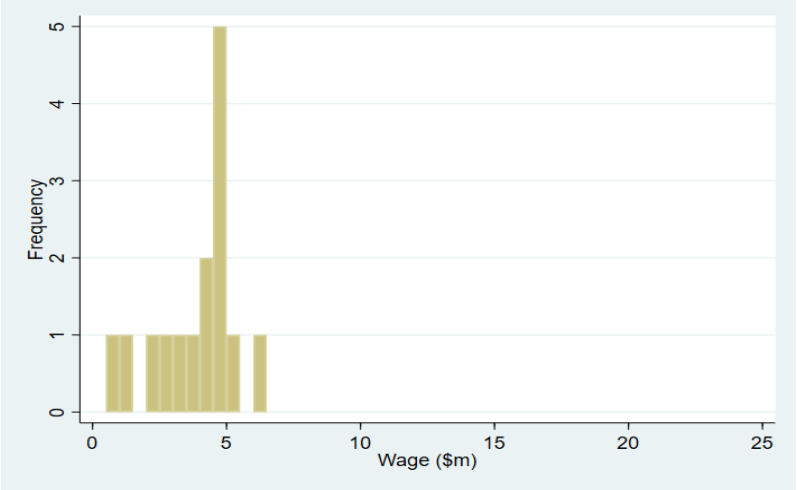
the other half earn twice the mean. This scenario is not possible in our data due to minimum salary rules. However, the maximum of 0.52 still seems quite low. This value of polarisation is from the 2016-17 Dallas Mavericks, and Figure 5.1a shows a histogram of their wages. We see clustering at the lower end of the distribution but not at the upper end. This is likely why the measure does not come close to the higher values of polarisation in our sample. When we compare Figure 5.1a to Figure 5.1b, which shows the lowest polarisation value in our sample of 0.05 by the 2011-12 Minnesota Timberwolves, it is clear that the distribution has a much lower upper bound and its main clustering is closer to the middle of the distribution. These examples of the extremes should display to the reader how polarisation incorporates both spread and bipolarity and how, in our sample, clusterings are more likely to occur at the lower end of the distribution.

Figure 5.1: Example Wage Distributions

(a) High Wage Polarisation of the 2016-17 Dallas Mavericks



(b) Low Wage Polarisation of the 2011-12 Minnesota Timberwolves



However, this raises the question of how highly correlated measures such as the standard deviation of wage and wage polarisation are in our sample. Standard deviation is essentially a measure of spread, which is one of the components of polarisation. Although spread is slightly different in each instance, standard deviation could be considered at least a good proxy for the spread component of polarisation. However, the measures differ firstly because the standard deviation does not account for the degree of clustering in the distribution and secondly because the standard deviation is not bounded as polarisation is, dealing with absolute changes

rather than percentage changes. The 2011-12 Timberwolves (Figure 5.1b) have the lowest polarisation and one of the lowest standard deviations of wage ($1.56m compared to the sample minimum of $1.55m); while the 2016-17 Mavericks (Figure 5.1a) have the highest polarisation but their standard deviation of wage is over $3m less than the sample maximum ($7.15m compared to $10.17m). Table 5.2 shows correlations between our wage dispersion measures. It reveals significant but reasonably low correlation between the standard deviation of wage and many of the other wage dispersion measures including polarisation, suggesting there is variation to exploit.

Figure 5.2: Standard Deviation of Wage and Wage Polarisation



Figure 5.2 shows a scatterplot between $WageSD$ and $WagePol$ within our dataset. Although the upward sloping trend line again demonstrates the positive correlation between the variables, the scatterplot displays that this correlation is weak and there is plenty of variation to exploit.

Table 5.2: Correlations Between Wage Dispersion Variables

	WageSD	WageGini	WageCV	WageHHI	WageSkew	WageKurt	WagePol
WageSD	1.000						
WageGini	0.665***	1.000					
WageCV	0.616***	0.962***	1.000				
WageHHI	0.598***	0.684***	0.759***	1.000			
WageSkew	0.269***	0.565***	0.716***	0.547***	1.000		
WageKurt	0.088	0.353***	0.533***	0.361***	0.918***	1.000	
WagePol	0.381***	0.804***	0.702***	0.346***	0.258***	0.157***	1.000

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The highest correlation in Table 5.2 is between the Gini Coefficient and Coefficient of Variation. This should not be surprising as both measures are specifically designed to measure similar phenomena. These measures are also highly correlated with the HHI measure. This is also not surprising theoretically, but given that studies using the HHI measure tend to give

support for Cohesion Theory (see Depken, 2000; Depken & Lureman, 2018), while studies using other inequality measures find evidence for Tournament theory (Pierce, 2017; Simmons & Berri, 2011), a negative coefficient here might have offered some explanation for this. The correlation coefficient between the Gini Coefficient and our polarisation measure is reasonably high at 0.804. Section 3.1 discussed how these measures can sometimes agree, but how they will also react differently to some changes in distribution. This result suggests that the measures agree with respect to many of the differences between wage distributions in our sample, but this does not mean that the remaining differences are not significant. Whether the differences in the distributions captured by these different dispersion measures have significant effects on team performance, will be revealed during our regression analysis.

Our final piece of analysis, before we turn to the regression results, looks at the pairwise correlation between our wage dispersion variables and our control variables. Table 5.3 shows that there are no exceptionally high correlation coefficients between our wage dispersion and control variables, but many are still highly statistically significant. The standard deviation of wage is the most correlated with a reasonably large positive correlation with $TPayroll$, as predicted by Theory of Superstars (Rosen, 1981) and to a smaller extent with experience as predicted by back-loaded pay schedules (Scully, 1974). There is also a positive association between the standard deviation of experience and the standard deviation of wage, as we predicted, giving evidence that this is an important control to isolate the effects of only wage dispersion itself. The other correlations are reasonably low, but are generally highest with the standard deviation of experience, again displaying the potential importance of this control.

Table 5.3: Correlations Between Wage Dispersion and Control Variables

	TPayroll	CTSize	TTenureMP	TExp	TExpSD
WageSD	0.800***	0.0134	0.308***	0.406***	0.418***
WageGini	0.381***	0.0893	0.0769	0.154***	0.327***
WageCV	0.298***	0.0983*	0.0441	0.116**	0.306***
WageHHI	0.193***	-0.195***	0.276***	0.248***	0.321***
WageSkew	0.0134	0.0433	-0.0706	-0.0372	0.114**
WageKurt	-0.0747	0.0555	-0.145***	-0.124**	0.0428
WagePol	0.272***	0.0804	0.00363	0.0550	0.236***

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Of course, the correlations in Table 5.3 only tell part of the story. The analysis conducted in Appendix 4.F shows that the control variables influence team performance, while Appendix 4.H shows that they also have statistically significant partial correlations with our key wage dispersion measures, $WageSD$ and $WagePol$. Still, these bivariate correlations are useful to inform the reader of the potential degree of collinearity in the model, and to show the correlations between the other measures of wage dispersion and the control variables.

5.3 Regression Results

For the majority of our analysis, we use regular season win percentage (WL) as our measure of team performance, while controlling for year and team effects using a Fixed Effects (FE) model. These assumptions are later relaxed to check the robustness of our results. First, we regress our different measures of wage dispersion separately on our team performance

measure, to identify the extent to which a relationship between the variables exists. We then incrementally add controls into the model and find $WageSD$ and $WagePol$ to show the most robust results. Next we consider different specifications of the model which include: quadratic terms, considering measures simultaneously, and interactions. Our results suggest that the omittance of $WageSD$ and/or $WagePol$, as well as other relevant controls, can cause other dispersion measures to show contradicting results. This omitted variable bias/model misspecification offers a potential explanation for the mixed results in the current literature. Section 5.3.1 conducts additional robustness tests before an overall discussion of our results in Section 5.4.

Table 5.4: Regression Results of Wage Dispersion Measures (No Controls)

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL	(6) WL	(7) WL
WageSD	0.0544***						
	(0.00663)						
WageGini		0.0851					
		(0.153)					
WageCV			0.00974				
			(0.0623)				
WageHHI				1.383**			
				(0.559)			
WageSkew					-0.0158		
					(0.0216)		
WageKurt						-0.0110*	
						(0.00539)	
WagePol							-0.260
							(0.153)
Constant	0.259***	0.456***	0.490***	0.330***	0.523***	0.549***	0.553***
	(0.0399)	(0.0908)	(0.0775)	(0.0829)	(0.0490)	(0.0416)	(0.0443)
R^2	0.199	0.00157	0.000156	0.0483	0.00248	0.0149	0.0144
Years	✓	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓	✓
Controls	✗	✗	✗	✗	✗	✗	✗

N = 330
Robust standard errors clustered by Team in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Column (1) of Table 5.4 shows a positive coefficient for the standard deviation of wage on team performance which is statistically significant at a 1% level. This result suggests evidence in favour of Tournament Theory (Lazear & Rosen, 1981). We also see positive coefficients for the Gini Coefficient, Coefficient of Variation and HHI in columns (2), (3) and (4) respectively, though only statistical significance for $WageHHI$ at a 5% level. This is contradictory to other studies which have utilised the HHI of wage which find negative effects (Depken, 2000; Depken & Lureman, 2018). Conversely, Columns (5), (6) and (7) show negative coefficients for wage skewness, kurtosis and polarisation respectively, which provides evidence in favour of Cohesion Theory (Levine, 1991). However, only the coefficient for $WageKurt$ is marginally statistically significant at a 10% level. Overall, the results presented in Table 5.4 suggest that the effect of wage dispersion on team performance depends on the definition of dispersion. Measures of spread, such as the standard deviation of wage, and to some extent inequality,

appear to encourage effort from players and increase performance in line with Tournament Theory. However, other measures, particularly those relating to the distribution of wages relative to the center and the extremes of the distribution, appear to aggravate the problems associated with Cohesion Theory, discouraging effort and potentially increasing the likelihood of conflict. This adds weight to the idea that polarisation captures a unique element of wage dispersion which is not captured by other measures currently considered in the literature.

Table 5.5: Regression Results of Wage Dispersion Measures (Some Controls)

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL	(6) WL	(7) WL
WageSD	0.0289***						
	(0.00717)						
WageGini		-0.0913					
		(0.125)					
WageCV			-0.0156				
			(0.0489)				
WageHHI				0.373			
				(0.475)			
WageSkew					-0.000414		
					(0.0166)		
WageKurt						-0.00310	
						(0.00399)	
WagePol							-0.341***
							(0.104)
Constant	0.319	-0.240	-0.235	-0.273	-0.237	-0.188	-0.304
	(0.300)	(0.249)	(0.248)	(0.254)	(0.259)	(0.269)	(0.247)
R^2	0.389	0.359	0.358	0.361	0.358	0.359	0.382
Years	✓	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓	✓
Controls	∼	∼	∼	∼	∼	∼	∼

∼ indicates that lnTPayroll and lnCTSize have been included, but not the full set of control variables
N = 330
Robust standard errors clustered by Team in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5.5 shows the results when including two control variables, the level of wage expenditure ($TPayroll$) and 'core team' size ($CTSize$). It is common in the literature to only use wage expenditure as a main control (e.g. Depken, 2000; Depken & Lureman, 2018; Pierce, 2017) or expenditure and team size (e.g. Frick et al., 2003). The results of Table 5.5 are largely similar to those in Table 5.4, but with three key differences. Firstly, the coefficients on $WageGini$ and $WageCV$ become negative, though they are statistically insignificant in both specifications. This displays that the sign of the inequality measures may be sensitive to different control specifications, indicating that the results are not robust. Secondly, the coefficient on $WageHHI$ remains positive but becomes statistically insignificant, showing that the result is not robust to the inclusion of some control variables. Thirdly, the coefficient on $WagePol$ remains negative but becomes statistically significant at a 1% level. This adds further weight to our theory that wage polarisation captures the negative effects predicted by Cohesion Theory, though we should be sure to control for all confounding factors before confirming this result. Additionally, the coefficient on $WageKurt$ remains negative but becomes statistically

insignificant. Otherwise, the coefficient on the standard deviation of wage remains positive and highly significant at a 1% level. The key variables influencing team performance here appear to be $WageSD$ positively and $WagePol$ negatively.

Table 5.6: Regression Results of Wage Dispersion Measures (Full Controls)

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL	(6) WL	(7) WL
WageSD	0.0212***						
	(0.00764)						
WageGini		0.0327					
		(0.105)					
WageCV			0.0252				
			(0.0409)				
WageHHI				0.352			
				(0.427)			
WageSkew					0.0119		
					(0.0147)		
WageKurt						0.00120	
						(0.00335)	
WagePol							-0.189*
							(0.101)
Constant	0.856**	0.527*	0.527*	0.478	0.485	0.507	0.437
	(0.315)	(0.309)	(0.308)	(0.310)	(0.307)	(0.311)	(0.310)
R^2	0.486	0.471	0.472	0.474	0.473	0.471	0.478
Years	✓	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓

N = 330
Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5.6 shows the robustness of the results of Table 5.5, to the inclusion of the full set of team-level control variables described in Section 4.4. These controls are the average tenure of the team in minutes played ($TTenureMP$), average experience in years ($TExp$) and the standard deviation of experience ($TExpSD$), further to $lnTPayroll$ and $lnCTSize$. Still, the only statistically significant coefficients are those of $WageSD$ in Column (1), positive at a 1% level, and $WagePol$ in Column (7), negative but now at a 10%. The fall in significance of $WagePol$ seems to be due to its correlation with the standard deviation of experience ($ExpSD$) as shown in Table 5.3. Table 5.3 also shows that $ExpSD$ is correlated with $WageSD$ with the correlation being positive in both instances. Removing $ExpSD$ from the specifications increases the significance level of $WagePol$ to a 5% level, while also decreasing the significance of $WageSD$ to a 5% level. This is because $ExpSD$ has a negative effect on performance as shown in Appendix 4.F, so its inclusion helps isolate the positive effects of $WageSD$, while its exclusion causes its negative effects to be conflated with those of $WagePol$. This suggests that the standard deviation of wage, wage polarisation and standard deviation of experience have distinct effects which should each be controlled for to isolate the effects of wage dispersion.[45]

[45] This is considered in our next set of regression results, in Column (2) of Table 5.7, by including all three in the same model.

The coefficients on *WageGini* and *WageCV* in Table 5.6 become positive again, but remain statistically insignificant as they have throughout all the specifications presented. The coefficients on *WageSkew* and *WageKurt* become positive when the full set of controls are included, having previously been negative. Though only *WageKurt* appeared significantly negative at a 10% level when including no controls in Table 5.4. Therefore, across our specifications, our most robust results are a positive effect on performance relating to a high standard deviation of wage, and a negative effect associated with higher wage polarisation.

Next, we investigate the possibility of a quadratic relationship as found in other studies (e.g. Coates et al., 2016; Franck & Nüesch, 2011; Vasilescu, 2008), and also whether multiple measures of dispersion should be included simultaneously in the same model, to capture all of the effects of wage dispersion on team performance. We find no evidence of a quadratic relationship between *WageSD* and team performance, with the results shown in Column (1) of Table 5.7. Though the coefficients indicate an inverted U-shape, only the positive linear coefficient is statistically significant, albeit only at a 10% level. We find similar results for the other dispersion measures and therefore conclude no evidence of a quadratic relationship. This result is consistent, with or without the inclusion of the control variables.

Column (2) of Table 5.7 includes the standard deviation of wage and wage polarisation simultaneously alongside the full set of controls (which include the standard deviation of experience). The coefficient on *WageSD* remains positive and significant at a 1% level, as it has been when included individually. The coefficient on *WagePol* remains negative and its statistical significance increases, compared to Table 5.6, also to a 1% level. This demonstrates that *WageSD* and *WagePol* are capturing distinct aspects of wage dispersion, each with opposing effects on team performance. Furthermore, considering them simultaneously helps isolate the effects of each.

Table 5.7: Regression Results with Multiple Wage Dispersion Measures

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL	(6) WL
WageSD	0.0491*	0.0287***	0.0344***	0.0257**		0.0313**
	(0.0251)	(0.00862)	(0.0101)	(0.0103)		(0.0118)
WageSD^2	-0.00260					
	(0.00236)					
WagePol		-0.313***		-0.368**	-0.523***	-0.248
		(0.112)		(0.180)	(0.181)	(0.286)
WageGini			-0.268**	0.0874	0.429**	
			(0.129)	(0.206)	(0.180)	
WageSD × WagePol						-0.0134
						(0.0503)
Constant	0.809**	0.832**	1.019***	0.775**	0.358	0.818**
	(0.340)	(0.305)	(0.308)	(0.319)	(0.313)	(0.314)
R^2	0.490	0.503	0.494	0.503	0.492	0.503
Years	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓

N = 330
Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In column (3) *WageGini* is included alongside *WageSD* which causes the coefficient on *WageGini* to become negative and statistically significant at a 5% level. When including *WageCV* or *WageHHI* in place of *WageGini*, their coefficients too become negative but are not statistically significant. This suggests that it is the spread of wage distributions which drive the positive coefficient of inequality measures found in previous specifications. When controlling for this spread, increased inequality appears to have a statistically significant negative effect on performance. So, given these results, it is relevant to investigate whether, when holding the standard deviation of wage constant, it is increases in wage inequality or increases in wage polarisation which have negative effects on performance. We argue that it is polarisation which is driving these negative coefficients for four reasons.

Firstly, as both measures are bounded between zero and one and the coefficient of *WagePol* in column (2) is larger than that of *WageGini* in column (3) of Table 5.7. Secondly, as the coefficient on polarisation is statistically significant at a higher level (1% compared to 5%) suggesting that the effect is more consistently found in the data. Thirdly, because when both *WagePol* and *WageGini* are including simultaneously in the model, along with *WageSD*, it is *WagePol* which remains statistically significant with a negative coefficient, while the coefficient on *WageGini* becomes positive and statistically insignificant. This result is shown in Model (4) of Table 5.7. And fourthly, as shown in column (5) of Table 5.7, when including *WagePol* and *WageGini* simultaneously without *WageSD*, the coefficient on *WageGini* becomes positive and significant at a 5% level, while the coefficient on *WagePol* remains negative and significant at a 1% level. Therefore, given these results, we conclude that it is increases in wage polarisation which have negative effects on team performance and the standard deviation of wage which has positive effects. The effects found using inequality measures look to be driven by their correlation with these key variables and can show different effects depending on the specification used.

Theoretically, a team would be perfectly polarised if half the team earned zero and the other half earned any other identical, positive amount, regardless of how large this amount is. We hypothesise that the further this amount is from zero, the more likely conflict is to occur and the more damaging it will be. Though it does not capture it perfectly, this difference between polarised groups can be proxied by the standard deviation of wage. Therefore, we interact wage polarisation with the standard deviation of wage. Column (6) of Table 5.7 shows the inclusion of this interaction term to our model. It is appropriate to compare the results of column (6) to those of column (2), which is the same specification just without the interaction term. The coefficient on *WageSD* increases in size but only retains statistical significance at a 5% level. This suggests that, theoretically, the positive effects of standard deviation could be larger if wage polarisation was zero. But this is both an out-of-sample prediction and an impossibility, as some standard deviation must always cause some degree of polarisation. The negative coefficient on the interaction term, though it is statistically insignificant, suggests that the more polarised a wage distribution becomes, the smaller the positive effect of higher wage standard deviation. The coefficient on *WagePol* becomes smaller and is no longer statistically different to zero. This would suggest that when standard deviation is zero, wage polarisation has no significant impact on team performance. This is logical as a perfectly polarised wage distribution with a standard deviation of zero would theoretically mean that the poles have no distance between them. Again, the negative interaction term suggests that the higher the standard deviation of wage (a proxy of the distance between the poles), the

more negative the effect on team performance, though this effect is not statistically significant.

While these additional findings are of some interest, they lack statistical significance. We therefore favour the specification without the interaction term, presented in column (2) of Table 5.7, as our key finding. This specification suggests that the positive effects of high wage dispersion associated with Tournament Theory (Lazear & Rosen, 1981) are best captured by the standard deviation of wage, while the negative effects associated with Cohesion Theory (Levine, 1991) are captured by wage polarisation. While these measures may be correlated, this correlation is relatively low (around 0.38 - see Table 5.2). For a given standard deviation of wage, teams can and should decrease wage polarisation as much as possible to maximise team performance. Similarly, for a given wage polarisation, teams should increase the standard deviation of the wage distribution as much as possible. The resultant distribution can have a large spread of wages, but should avoid clusterings, especially at the extremes. Next, we test the robustness of this result to relaxing some of the assumptions of our model and to using different measures of team performance. Appendix 5.C also shows the robustness of our result to the use of different measures of wage polarisation.

5.3.1 Robustness Checks

Table 5.8 shows the robustness of our results to adjustments such as excluding team and year effects and changing the definition of the team. Columns (1) and (2) reiterate that the results relating to $WageSD$ and $WagePol$ are robust to the inclusion of additional controls, as already demonstrated. The coefficients are larger when excluding the controls in column (1), but the coefficients remain statistically significant at a 1% level in column (2) when including the full set of control variables. As the control variables significantly influence both team performance and wage dispersion, as demonstrated in Appendices 4.F and 4.H, the change in the magnitudes of the coefficients is expected, but the main result is robust.

Table 5.8: Robustness Checks (Relaxing Controls and Definition of the Team)

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL
WageSD	0.0637***	0.0287***	0.0282***	0.0271***	0.0301***
	(0.00734)	(0.00862)	(0.00748)	(0.00921)	(0.00729)
WagePol	-0.581***	-0.313***	-0.287**	-0.210	-0.523***
	(0.133)	(0.112)	(0.104)	(0.129)	(0.158)
Constant	0.336***	0.832**	1.035***	0.917***	0.749**
	(0.0404)	(0.305)	(0.192)	(0.330)	(0.276)
R^2	0.266	0.503	0.474	0.452	0.507
Years	✓	✓	✗	✓	✓
Teams	✓	✓	✓	✗	✓
Controls	✗	✓	✓	✓	✓
CoreTeam	✗	✗	✗	✗	✓

N = 330
Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Column (3) excludes the year dummies, while column (4) excludes team effects by using a Random Effects model, as opposed to the Fixed Effects models used so far. Appendix 4.F shows that both team and year effects significantly influence team performance, while Appendix 4.H shows only the team effects are significant determinants of wage dispersion. We therefore expect dropping the team effects to have a larger impact on the model than dropping

the year effects, and this is what we see. The magnitude of the coefficients does not change much between columns (2) and (3), and there are only small changes in the standard errors when excluding the year dummies, though this is enough to reduce the statistical significance of *WagePol* to a 5% level. The differences between columns (2) and (4) are larger, particularly on the coefficient on wage polarisation which becomes marginally insignificant. However, a Hausman-type test using the Mundlak (1978) approach (which is valid for robust, clustered standard errors), presented in Appendix 5.A, supports the inclusion of team effects using a Fixed Effects model due to correlation between time-invariant, team-specific heterogeneity and the regressors.

In column (5) of Table 5.8, we change the definition of the team to only include the players classified as part of the 'core team' using the definition of Simmons and Berri (2011). This includes only players who play at least 20 games and at least 12 minutes per game in a given season.[46] We then recreate the variables of interest and controls for this definition of the team and re-estimate the model. Comparing this to column (2) which is the same specification but for the entire roster, we see the same qualitative result in terms of signs and significance levels but with even larger coefficients. This result suggests that the effects may not qualitatively depend on the definition of the team as Bucciol et al. (2014) suggest, though the effects appear to be stronger when considering just the 'core team'.

Next, Table 5.9 shows the robustness of our main result to using different measures of team performance as the dependent variable. Column (1) again displays the results with win percentage as the dependent variable for comparison. Columns (2), (3) and (4) show the robustness of our result to using the natural logarithm of win percentage ($lnWL$), points difference per game (PDG) and a simple ratings system (SRS) respectively. Our main finding shows strong robustness to these different measures of team performance, with variations in magnitude due to the scale of the dependent variables. The levels of statistical significance are also consistent, with the negative effect of *WagePol* decreasing in significance to a 5% level when considering PDG and $lnWL$ in columns (2) and (3), being the only small difference. Though it returns to a 1% level when accounting for the strength of team's schedule in column (4).

Table 5.9: Robustness Check (Performance Measures)

	(1) WL	(2) lnWL	(3) PDG	(4) SRS	(5) ORTG	(6) DRTG
WageSD	0.0287***	0.0627***	0.874***	0.863***	0.486*	0.401**
	(0.00862)	(0.0186)	(0.285)	(0.268)	(0.238)	(0.151)
WagePol	-0.313***	-0.678**	-8.694**	-8.458***	-6.551**	-2.419
	(0.112)	(0.287)	(3.162)	(3.052)	(2.738)	(2.212)
Constant	0.832**	-0.334	14.307*	13.960*	116.349***	-101.678***
	(0.305)	(0.668)	(8.312)	(7.977)	(6.686)	(4.978)
R^2	0.503	0.474	0.469	0.469	0.526	0.400
Years	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓

N = 330
Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[46] This removes 33% (1958 of 5932) of the distinct player-team-year observations.

Columns (5) and (6) show the effects of *WageSD* and *WagePol* on offensive and defensive ratings, respectively. The results here are qualitatively similar, though the coefficients show lower levels of significance. As established in Section 4.3.2, these ratings individually do not tell the whole story with regards to overall team performance, though they do offer additional insight. *WageSD* has a significant effect on both ratings, though its magnitude is higher for *ORTG*, while its significance level is higher for *DRTG* due to a smaller standard error. Therefore, *WageSD* seems to have positive effects on *DRTG* more consistently. On the other hand, the coefficient on *WagePol* is larger and is statistically significant at a 5% level for *ORTG*, while it is statistically insignificant for *DRTG*. Therefore, the negative effects of wage polarisation on team cohesion appear to damage team performance through reduced offensive production perhaps through discouraging cooperation in favour of individualistic offensive play.

Table 5.10: Robustness Check (Alternative Models)

	(1) WL	(2) WL	(3) WL
WageSD	0.0287***	0.0270***	0.0274***
	(0.00862)	(0.00809)	(0.00865)
WagePol	-0.313***	-0.196	-0.239**
	(0.112)	(0.123)	(0.120)
Constant	0.832**	1.050***	0.891***
	(0.305)	(0.201)	(0.313)
R^2	0.503	0.431	-
log likelihood	313.3	-	270.8
Years	✓	✗	✓
Teams	✓	✗	✓
Controls	✓	✓	✓
Model	FE	RE	MLM

N = 330
Robust standard errors clustered by Team in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

Lastly, in Table 5.10, we consider our choice of regression model. Column (1) shows the 'complete pooling' Fixed Effects (FE) approach taken throughout this Chapter. Column (2) shows the 'no pooling' Random Effects (RE) approach which excludes Team and Year dummies. These represent the extremes of the traditionally used regression techniques. Column (3) shows the results of the multilevel modelling (MLM) approach. Across all three models, the effects of *WageSD* are consistently positive and statistically significant at a 1% level. While the negative coefficient on *WagePol* is consistently negative, it is significant at a 1% level for FE, at a 5% level for MLM, but marginally statistically insignificant for RE (with a p-value of 0.111). This suggests that the effects of wage polarisation on team performance are driven by the within variation of teams over seasons. The 'no pooling' approach tends to overweight the variation between teams (Gelman & Hill, 2006), which seems to be lower for the effects of wage polarisation. Overall, there is broad consensus between the models of the positive effects of a higher standard deviation and the negative effects of a higher wage polarisation. But this additional analysis reveals that the effects are wage polarisation are driven by changes within teams over seasons.

5.4 Summary and Discussion

The aim of this chapter has been to investigate how different types of dispersion within wage distributions can affect team performance. This is a well-researched topic, however both the theoretical and empirical evidence have presented mixed results. Theoretically, there is disagreement between two well-established theories: Tournament Theory (Lazear & Rosen, 1981) and Cohesion Theory (Levine, 1991), which make opposing predictions, but both have received empirical support. Using evidence from the NBA, for the period 2007-18, we attempt to reconcile this mixed evidence by contributing to the literature in four distinct ways. Firstly, we develop a rigorous set of control variables which significantly influence both wage dispersion and team performance. Omitting such variables could lead to biased estimates and potentially misleading conclusions. We demonstrate how the apparent effects of some wage dispersion measures can change sign when some of these control variables are omitted. This evidence suggests that omitted variable bias is a potential problem which has so far been overlooked.

Secondly, in the existing literature, different measures of wage dispersion tend to be used synonymously despite capturing distinct, albeit in some cases similar, information about the distribution of wage. These measures can be highly correlated in some dimensions, but this does not mean they necessarily correlate in the aspects which are important in determining team performance. We therefore argue that the choice of dispersion measure is important, so consider a range of measures in this study. Furthermore, it is assumed in the literature that this issue can be captured by a single measure of wage dispersion, either in a linear, quadratic or logarithmic form. We argue that the reality is more complex than this and that the distinct information captured by the various dispersion measures can affect performance in different ways. Therefore, our third contribution is the inclusion of multiple wage dispersion measures simultaneously in the model to help unravel this complexity and isolate the effects of each, while holding other types of dispersion constant.

Our fourth, and arguably most important, contribution is the introduction of polarisation measures to the discussion. These measures were developed to capture the likelihood of conflict within countries, but we argue that they are equally important at the firm-, or in our case team-, level. These polarisation measures consider clusterings, not just around a global mean, but also around local means and are increasing in two arguments: spread and bipolarity. It is particularly this bipolarity which is not captured by other dispersion measures, but that can be an important determinant of conflict.

Piecing these contributions together, we arrive at a model which includes both the standard deviation of wage and wage polarisation simultaneously, and finds them to have opposing effects on team performance. Increases in the standard deviation can have positive effects on team performance, as workers increase their effort levels to reach the highest wage level possible, as predicted by Tournament Theory. However, particularly when holding this constant, increases in wage polarisation can create divisions within the team, raising the probability that conflict will occur, as in the predictions of Cohesion Theory. This can lead to competitive effort becoming non-cooperative and even result in sabotage. Therefore, to maximise team performance, wage distributions should be structured so as to maximise the standard deviation, but minimise the degree of polarisation. Such a wage distribution can have a large spread of wages, but should avoid clusterings, especially at the extremes.

While the mechanisms discussed thus far are consistent with the relevant literature and theories, we cannot rule out other potential mechanisms which would lead to similar findings. One such example would be that a team with the suggested ideal wage structure could also have a spread, but not polarised, distribution of abilities. Given the higher number of games in an NBA season, as well as the high frequency of in-game substitutions, this distribution of abilities may be beneficial to maintaining consistent performance throughout entire games and entire seasons. Highly polarised teams may suffer when key players are substituted and/or injured. Assuming efficient labour markets, distributions of wages and abilities would be synonymous, making it difficult to distinguish between these mechanisms. To do so would require a measure of ability which is distinct from wages and for wages not to be fully determined by ability (i.e. inefficient labour markets). Wage dispersion theories also suggest that performance is affected by the level of wage dispersion in a team. If this is true, it would also be important to be able to distinguish between raw ability (the input) and performance (the output), as the later would be influenced by the level of wage dispersion. Given these difficulties, it is important to note that other mechanisms may also be plausible.

The evidence presented is from the NBA. While we believe these findings could apply consistently over different sports, further work is required to test this. Theoretically, the main structure of sports labour markets would remain consistent, with players competing for a limited number of starting positions in each team, but also being required to cooperate to maximise team production. With these dynamics in place, it is reasonable to assume that the same relationship between wage dispersion and team performance could be found. However, differences could be apparent in sports with less frequent substitutions, such as soccer, where the dynamics between the starters may be considered more important than the dynamics amongst the entire team. Or perhaps in sports such as American Football, where distinct Defensive, Offensive and Special Team units are found within the team roster. It would also be interesting to investigate whether this effect is consistent within sports which operate within a hard salary cap and ones which operate with no salary cap at all.

Another suggestion for further work would be to test this finding with data from general labour markets outside the sports industry. Detailed individual wage data is not as readily available for these industries so it may be difficult to find sufficiently rich data to construct these variables. But again, it would not be surprising if the theory leading to the main finding of this chapter also applied to these firms. If so, the implications could be an important consideration. Firms in general labour markets tend to have hierarchical or tiered systems, with promotions required to achieve the next level. This can create clearly distinct groups within the organisation. The results here would advise creating intermediary groups to populate the 'middle'. It would also be interesting to test whether these dynamics exist across entire large organisations, or whether workers create smaller subgroups comprising of their immediate superiors and subordinates with which to compare their wages.

Chapter 6

National Diversity and Team Performance

Chapter Overview

Globalisation has led to a general increase in national diversity within work teams. This is true within many industries, and especially within sports. With the public nature of production, a limited number of top-level performers and the possibility of large rewards for relatively small comparative improvements in production, it is of little surprise that sports teams look to identify talent from around the globe. The substantial wages on offer and the prestige attached to certain leagues around the world are often enough to tempt workers to relocate to a new country. This raises the question of whether national diversity in itself has an effect on team performance. As teams look to benefit from a wider talent pool by recruiting players from a variety of countries, the effects of national diversity on cooperation, and potentially on conflict, within the team should be considered. Using data from the National Basketball Association (NBA), we utilise panel data techniques to estimate team performance as a function of national diversity, while including several relevant control variables which are found to be important in isolating these effects. Measures of fractionalisation and polarisation are used to capture the structure of national diversity within teams and are found to be simultaneously important. These measures have previously been used to analyse economic outcomes at the country-level, but this study is amongst the first to apply them to the context of national diversity within work teams. We also experiment with measures of linguistic distance to proxy cultural distance between nations which are incorporated into the fractionalisation and polarisation measures. At the end of this chapter, we also present a combined model, which considers the effects of wage dispersion (from Chapter 5) together with the effects of national diversity.

6.1 Introduction

National Diversity can relate to the number of national groups within team, the size of these groups, as well as the cultural distance between these national groups. Considering all of these aspects makes it a difficult phenomenon to capture. Existing literature has addressed these aspects, but usually individually and in different scenarios. This chapter considers measures of fractionalisation and polarisation which have previously been applied to diversity and conflict at the country-level. These measures are also combined with measures of linguistic distance, used as a proxy for cultural distance, to incorporate the notion of distance between nations within different structures.

Section 6.2 discusses the descriptive statistics of and correlations between these measures of national diversity. Section 6.3 then presents regression results before robustness checks are carried out in Section 6.3.1. Section 6.4 combines these results relating to national diversity with those from Chapter 5 relating to wage dispersion to present a model which considers both simultaneously. The chapter concludes by discussing the results in Section 6.5, suggesting that national diversity can have positive effects, though its structure should be managed. Also, that these results relating to national diversity are robust to the inclusion of wage dispersion variables, and vice versa.

6.2 Descriptive Statistics and Correlations

Table 6.1 presents descriptive statistics relating to the general structure of national diversity within NBA teams from 2007 to 2018. *FRGN* relates to the proportion of each team born outside the USA, while *NatNum* is the number of nationalities within a team. *NatFrac* and *NatPol* are basic measure of fractionalisation and polarisation respectively, which consider an equal cultural distance of 1 between countries, as discussed in Section 3.2.

Table 6.1: Descriptive Statistics Relating to the Structure of National Diversity

Variable	Panel	Mean	Std Dev.	Min.	Max.
FRGN	Overall	0.194	0.102	0	0.563
	Between		0.0649	0.0582	0.370
	Within		0.0800	-0.0515	0.401
NatNum	Overall	4.312	1.735	1	10
	Between		1.0550	2	7.182
	Within		1.390	0.221	8.676
NatFrac	Overall	0.329	0.154	0	0.766
	Between		0.0968	0.107	0.564
	Within		0.121	-0.0726	0.639
NatPol	Overall	0.493	0.169	0	0.715
	Between		0.101	0.205	0.656
	Within		0.137	-0.0434	0.778

While less than 20% of the players competing in the NBA during our sample period were born outside the USA[47], our dataset contains players from over 60 different countries. Importantly, there are also large variations in the structure of national diversity within teams to exploit.

[47] While the average in Table 6.1 relates to average proportion of each team, the overall proportion in the dataset is virtual identical (19.443% and 19.454%, respectively).

This ranges from nationally homogeneous teams consisting of only one nationality, to more heterogeneous (or fractionalised) teams with a maximum of ten different nationalities. There are thirteen instances of homogeneous teams which unsurprisingly consist exclusively of players from the USA. The team with the highest proportion of foreign players was the 2012-13 Spurs with 56.25% from nine nations. There are two instances of teams which consist of ten nationalities within rosters of 17 and 19 players respectively. Both of these observations belong to the San Antonio Spurs from the 2013-14 and 2014-15 seasons. In the 2010-11 season, the Toronto Raptors had nine nationalities within a roster of 22, as did the 2012-13 Spurs. While there are a number of examples of teams with eight nationalities.

Table 6.1 also shows our basic measures of fractionalisation and polarisation have a minimum of zero for the homogeneous teams, and reach a maximum fractionalisation of 0.77 and maximum polarisation of 0.71. The most fractionalised team in our sample was the 2012-13 San Antonio Spurs with seven USA players, two French, one Australian, one Argentinian, one Belgian, one Brazilian, one Canadian, one Kiwi and one from the US Virgin Islands; resulting in nine nationalities in a roster of 16. Meanwhile the most nationally polarised team was the 2007-08 Toronto Raptors with eleven players from the USA, two Slovenian, two Spanish, an Argentinian and an Italian. Both examples have higher than average proportions of players born outside the USA, while the highly fractionalised example has many more national groups than the polarised one. However, $FRGN$ and $NatNum$ do not cover all aspects relating to fractionalisation and polarisation.

Figure 6.1: National Fractionalisation and Polarisation

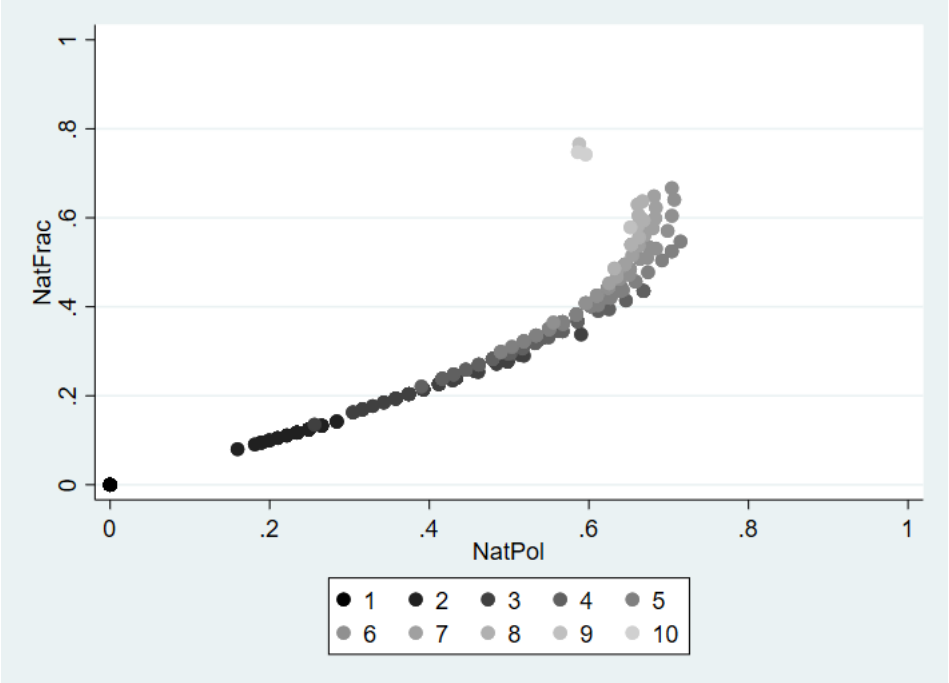


Figure 6.1 shows a scatterplot of these two measures, $NatFrac$ and $NatPol$, with the number of nationalities indicated by different shades. As discussed in Section 3.2, fractionalisation and polarisation are perfectly correlated when considering two or less national groups. This correlation gradually decreases as the number of national groups increases and then becomes negative when teams become highly fractionalised with many national groups. However, compared to our simulations presented in Figure 3.6 of Chapter 3, there are at best just a few

examples of high national polarisation and not many examples of high national fractionalisation. We will therefore mainly be exploiting variation in the measures for teams with higher numbers of national groups. The median number of nationalities within teams is four, so there may not be sufficient variation to clearly distinguish between $NatFrac$ and $NatPol$ in a large portion of our data.

Table 6.2: Descriptive Statistics of National Diversity Considering Linguistic Distance

Variable	Panel	Mean	Std Dev.	Min.	Max.
NatFrac1	Overall	0.227	0.112	0	0.533
	Between		0.0634	0.0933	0.355
	Within		0.0925	-0.0595	0.487
NatFrac2	Overall	0.158	0.0896	0	0.401
	Between		0.0492	0.0560	0.231
	Within		0.0754	-0.0655	0.405
NatFrac3	Overall	0.147	0.0848	0	0.402
	Between		0.0461	0.0535	0.223
	Within		0.0716	-0.0767	0.403
NatFrac4	Overall	0.147	0.0849	0	0.402
	Between		0.0461	0.0536	0.224
	Within		0.0718	-0.0775	0.403
NatPol1	Overall	0.324	0.133	0	0.583
	Between		0.0716	0.173	0.429
	Within		0.113	-0.0383	0.573
NatPol2	Overall	0.222	0.113	0	0.517
	Between		0.0607	0.104	0.338
	Within		0.0963	-0.0623	0.501
NatPol3	Overall	0.207	0.107	0	0.517
	Between		0.0569	0.0996	0.293
	Within		0.0916	-0.0791	0.502
NatPol4	Overall	0.206	0.107	0	0.517
	Between		0.0567	0.0995	0.293
	Within		0.0918	-0.0805	0.502

Table 6.2 shows the descriptive statistics of our fractionalisation and polarisations measures which account for cultural distance proxied by linguistic distance. The numbered suffixes indicate the number of language considered for each country, as described in Section 3.2. Compared to the basic measures presented in Table 6.1, the average of the corresponding measure which considers one language is roughly two thirds as large. As the number of languages considered increases, the measures get smaller logarithmically. Considering more languages captures more similarities between countries and therefore reduces the measure of distance between them.[48]

The pairwise correlation between these measures which consider linguistic distance is high as shown in Table 6.3. However, there is some variation in fractionalisation when moving from considering one language to two or more, but little variation when moving from two to three or four. In terms of polarisation, the correlation coefficient is only lower when comparing just

[48] This is true on average though with notable exceptions such as the distance between USA and Canada. In the basic measure the distance between these countries is one, as it is for all countries. Considering one language the distance is zero, but considering two languages the distance is roughly 0.08.

Table 6.3: Correlations Between Linguistic Distance Measures

	NatFrac1	NatFrac2	NatFrac3	NatFrac4	NatPol1	NatPol2	NatPol3	NatPol4
NatFrac1	1.000							
NatFrac2	0.910***	1.000						
NatFrac3	0.896***	0.984***	1.000					
NatFrac4	0.895***	0.982***	0.999***	1.000				
NatPol1	0.927***	0.865***	0.854***	0.852***	1.000			
NatPol2	0.781***	0.934***	0.919***	0.915***	0.864***	1.000		
NatPol3	0.769***	0.915***	0.939***	0.938***	0.850***	0.975***	1.000	
NatPol4	0.768***	0.911***	0.938***	0.939***	0.847***	0.970***	0.999***	1.000

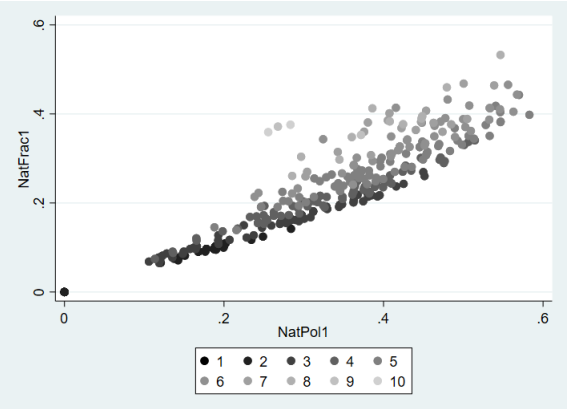
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

one language to two or more. Therefore, overall, it seems that it may be worthwhile to consider both one or two languages when considering linguistic distance, especially as the latter of the two accounts for the proportion of speakers of each language (which is not accounted for when considering one language). However, there appears to be little advantage to considering more than two languages as it causes little change in our proxies. Therefore much of the forthcoming analysis is limited to discussing these measures when considering one or two languages.

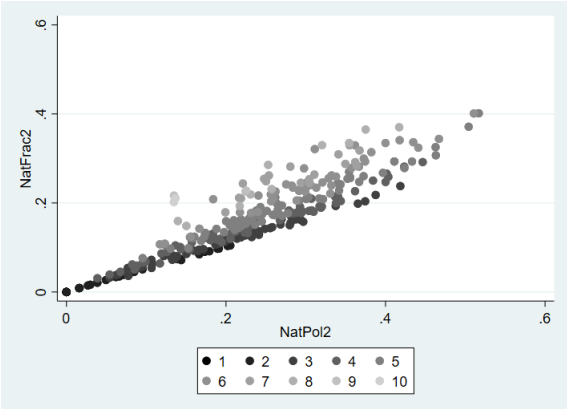
The scatterplots presented in Figure 6.2 show that considering linguistic distance can change the relationship between fractionalisation and polarisation. The linear relationship between the variables still exists when there are one or two national groups, though additionally there are teams with two national groups which have zero fractionalisation and polarisation. This is because the nations have zero linguistic distance. When considering one language, these teams are the 2010-11 Detroit Pistons with a roster of 12 players from the USA and one from the UK, and the 2015-16 Philadelphia 76ers, as well as both the 2012-13 and 2013-14 Miami Heat, which each consisted of 17 Americans and one Canadian. The primary language of all three countries is English, resulting in the zero linguistic distance.

Figure 6.2: National Fractionalisation and Polarisation: Considering Linguistic Distance

(a) National Fractionalisation and Polarisation considering one language



(b) National Fractionalisation and Polarisation considering two languages



However, none of these teams have zero fractionalisation and polarisation when considering two languages, as the USA has a different secondary language (Spanish) to the UK and Canada (French). The only team to have two national groups and measure zero for both

$NatFrac2$ and $NatPol2$ is the 2007-08 Los Angeles Clippers consisting of 19 players from the USA and one Puerto Rican[49]. While the USA and Puerto Rico have different primary languages, the two most spoken languages in both countries are Spanish and English. Although the proportion of speakers of each language in each country differ, our measure fails to capture this when comparing identical sets of languages, as the proportions are multiplied by zero distance. While one can question whether these nuances are completely accurate, linguistic distance seems to act as a reasonable proxy of cultural distance.

Table 6.4: The Most Nationally Diverse Teams and their Roster Structures

Variable	Team	Value	Roster Structure
$FRGN$	2012-13 San Antonio Spurs	0.5625	7 USA, 1 Argentina, 1 Australia, 1 Belgium, 1 Brazil, 1 Canada, 1 France, 1 New Zealand, 1 US Virgin Islands
$NatNum$	2013-14 San Antonio Spurs	10	9 USA, 2 France, 1 Argentina, 1 Australia, 1 Belgium, 1 Brazil, 1 Canada, 1 Italy, 1 New Zealand, 1 US Virgin Islands
	2014-15 San Antonio Spurs	10	8 USA, 1 Argentina, 1 Australia, 1 Belgium, 1 Brazil, 1 Canada, 1 France, 1 Italy, 1 New Zealand, 1 US Virgin Islands
$NatFrac$	2012-13 San Antonio Spurs	0.766	7 USA, 1 Argentina, 1 Australia, 1 Belgium, 1 Brazil, 1 Canada, 1 France, 1 New Zealand, 1 US Virgin Islands
$NatFrac1$	2013-14 Minnesota Timberwolves	0.533	10 USA, 1 Cameroon, 1 Martinique, 1 Montenegro, 1 Puerto Rico, 1 Russia, 1 Senegal, 1 Spain
$NatFrac2$	2008-09 Memphis Grizzlies	0.401	10 USA, 1 Iran, 1 Saint Vincent and the Grenadines, 1 Serbia, 1 Spain
$NatPol$	2007-08 Toronto Raptors	0.715	11 USA, 2 Slovenia, 2 Spain, 1 Argentina, 1 Italy
$NatPol1$	2008-09 Orlando Magic	0.583	12 USA, 1 Guadeloupe, 1 Poland, 1 Saint Vincent and the Grenadines, 1 Turkey
$NatPol2$	2008-09 Memphis Grizzlies	0.517	10 USA, 1 Iran, 1 Saint Vincent and the Grenadines, 1 Serbia, 1 Spain

For the remainder of the distribution, we see less correlation based purely on the number of nationalities within the team. The two teams with ten nationalities rank 38^{th} and 49^{th} for $NatFrac1$, and 107^{th} and 94^{th} for $NatFrac2$. They also rank much lower in terms of polarisation than they did in the basic measure. Additionally, for a given level of fractionalisation or polarisation, we see much more variation in the other when considering linguistic distance, though the measures are still highly correlated. When considering one language, the most fractionalised team is the 2013-14 Minnesota Timberwolves, while the most polarised is the 2008-09 Orlando Magic. When considering two languages, both measures agree the most diverse team is the 2008-09 Memphis Grizzlies. The most diverse teams, as judged by each of the measures discussed so far, are summarised in Table 6.4. When considering linguistic distance, it appears to dominate the effects of the structure to some extent, though

[49] Puerto Rico is considered as a distinct country from the USA in the Ethnologue database (Lewis et al., 2014). We therefore follow this classification.

fractionalisation tends to be higher when there are large linguistic distances between many groups while polarisation is higher when there are large linguistic distance between few groups.

Table 6.5 displays the pairwise correlations of the measures of national diversity discussed thus far. While Table 6.4 shows that most of the measures suggest a different team is the most diverse, Table 6.5 shows that some high correlations still exist. The high correlations between $FRGN$, $NumNat$ and $NatFrac$ of 0.939, 0.994 and 0.946, display that the value of basic fractionalisation is mainly driven by the proportion of non-USA players in the team, which in turn has a high correlation with the number of nationalities in the team. However, this is not so much the case for $NatFrac1$ and $NatFrac2$ where the linguistic distance between the nations plays a larger role. The lowest correlations exist between the measures of polarisation which consider linguistic distance ($NatPol1$ and $NatPol2$), suggesting these measures capture something distinctly different to the other measures which have been previously studied.

Table 6.5: Correlations Between National Diversity Variables

	FRGN	NatNum	NatFrac	NatFrac1	NatFrac2	NatPol	NatPol1	NatPol2
FRGN	1.000							
NatNum	0.939***	1.000						
NatFrac	0.994***	0.946***	1.000					
NatFrac1	0.866***	0.845***	0.882***	1.000				
NatFrac2	0.767***	0.749***	0.786***	0.910***	1.000			
NatPol	0.893***	0.850***	0.932***	0.850***	0.768***	1.000		
NatPol1	0.656***	0.650***	0.697***	0.927***	0.865***	0.776***	1.000	
NatPol2	0.538***	0.531***	0.576***	0.781***	0.934***	0.655***	0.864***	1.000

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

To end this section, we consider the pairwise correlations between our measures of national diversity and the controls discussed in Section 4.4. These are presented in Table 6.6.

Table 6.6: Correlations Between National Diversity Measures and Control Variables

	TPayroll	CTSize	TTenureMP	TExp	TExpSD
FRGN	0.028	0.020	0.126**	-0.092*	0.007
NatNum	0.055	0.139**	0.035	-0.147***	0.013
NatFrac	0.030	0.030	0.103*	-0.100*	0.002
NatFrac1	-0.017	0.016	0.016	-0.096*	-0.019
NatFrac2	-0.018	0.006	-0.027	-0.109**	-0.005
NatPol	0.033	0.024	0.044	-0.096*	-0.018
NatPol1	-0.023	-0.001	-0.045	-0.068	-0.030
NatPol2	-0.020	-0.002	-0.079	-0.083	-0.011

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The most consistently significant correlations are between the average NBA experience within teams ($TExp$) and the national diversity measures. The negative coefficients show that teams with more national diversity tend to have less NBA experience.[50] This relationship is expected with some international players starting their careers in basketball leagues of other countries.

[50] Less correlation is found between diversity and age, and the effects on performance appear to be better captured by experience, as discussed in Section 4.4.

There is also some relationship between national diversity and $TTenureMP$, though, for the most part, this correlation is not statistically significant. It also appears that rosters with a larger core team ($CTSize$) also have more nationalities in them. A similar relationship exists when considering the size of the entire roster but we retain the size of the core team for consistency. We consider national diversity within the core team during the robustness checks and find similar results.

It is also noteworthy that the sign of the correlation between $TPayroll$ and the measures of national diversity is not consistent between the basic measures and those which consider linguistic distance. It appears that teams which spend more on wages have players from more nationalities in their teams but less linguistic distance, though these correlations are not statistically significant.

6.3 Regression Results

For the regression analysis of the effects of national diversity on team performance, we take a similar approach as was taken for wage dispersion in Chapter 5. The majority of the analysis is conducted with regular season win percentage (WL) as the measure of team performance and with team and year effects included. Table 6.7 shows the results of including our measure of national diversity without controlling for any additional covariates.

Table 6.7: Regression Results of National Diversity Measures (No Controls)

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL	(6) WL	(7) WL	(8) WL
FRGN	0.114							
	(0.116)							
NatNum		-0.00343						
		(0.00708)						
NatFrac			0.0691					
			(0.0784)					
NatFrac1				0.134				
				(0.102)				
NatFrac2					0.146			
					(0.139)			
NatPol						0.0715		
						(0.0721)		
NatPol1							0.141	
							(0.0836)	
NatPol2								0.131
								(0.108)
Constant	0.480***	0.514***	0.480***	0.471***	0.477***	0.468***	0.455***	0.471***
	(0.0318)	(0.0352)	(0.0333)	(0.0299)	(0.0303)	(0.0402)	(0.0316)	(0.0308)
R^2	0.00447	0.00119	0.00378	0.00848	0.00672	0.00518	0.0140	0.00881
Years	✓	✓	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✗	✗	✗	✗	✗	✗	✗	✗

N = 330

Robust standard errors clustered by Team in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Column (1) of Table 6.7 suggests that having a high proportion of players born outside the USA may have a positive effect on performance. However, column (2) suggests that having more nationalities within a team is also harmful for team performance. These results appear contradictory, which was unexpected especially considering the highly positive correlation between the variables displayed in Table 6.5. Though it should be noted that neither of these coefficients are statistically different from zero. The measures of fractionalisation and polarisation show positive coefficients whether linguistic distance is considered or not, though again no statistical significance is shown. These results are presented in columns (3)-(8) of Table 6.7. So perhaps there is weak evidence to suggest that high national diversity in a general sense may have positive effects on performance. With the positive coefficient on $FRGN$ and the negative coefficient on $NatNum$, perhaps a polarised structure would allow for a high proportion of foreign players and a low number of nationalities simultaneously. But the evidence is far from conclusive at this stage before including additional covariates to the model.

Table 6.8 shows the results of again including the national diversity measures individually, but this time we control for the wage expenditure of the team ($TPayroll$) and the size of the core team ($CTSize$).

Table 6.8: Regression Results of National Diversity Measures (Some Controls)

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL	(6) WL	(7) WL	(8) WL
FRGN	0.198*							
	(0.101)							
NatNum		0.00795						
		(0.00613)						
NatFrac			0.125*					
			(0.0678)					
NatFrac1				0.177**				
				(0.0828)				
NatFrac2					0.176			
					(0.110)			
NatPol						0.0903		
						(0.0605)		
NatPol1							0.127*	
							(0.0635)	
NatPol2								0.107
								(0.0807)
Constant	-0.374	-0.315	-0.371	-0.353	-0.340	-0.355	-0.318	-0.304
	(0.249)	(0.254)	(0.249)	(0.253)	(0.253)	(0.246)	(0.252)	(0.254)
R^2	0.371	0.364	0.370	0.372	0.367	0.366	0.369	0.364
Years	✓	✓	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓	✓	✓
Controls	∼	∼	∼	∼	∼	∼	∼	∼

∼ indicates that lnTPayroll and lnCTSize have been included, but not the full set of control variables
N = 330
Robust standard errors clustered by Team in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The results in Table 6.8 are qualitatively the same as in Table 6.7 for all of the measures except the coefficient on $NatNum$ changes from negative to positive. This appears to be driven

by the positive correlation between the number of nationalities in the team and the size of the core team.[51] Larger core teams represent disruption in the roster through injuries, trades and suspensions which has a negative effect on performance. Larger teams are also more likely to have more nationalities representing them, so the negative coefficient found in Table 6.7 appears to be driven by this correlation. For the other measures of national diversity in Table 6.8, we see consistent signs and increased statistical significance for the coefficients at a 10% level for $FRGN$, $NatFrac$ and $NatPol1$; and at a 5% level for $NatFrac1$.

Next we include the full set of controls with the results presented in Table 6.9. The signs of the coefficients are consistent with those in Table 6.8, but the full set of controls reveals higher statistical significance. Columns (1) and (2) suggest that teams with a higher proportion of foreign players and teams with more nationalities perform better. This would suggest that national fractionalisation has positive effects on team performance and indeed this is what we see in columns (3) and (4) at a 1% level; and column (5) at a 5% level.

Table 6.9: Regression Results of National Diversity Measures (Full Controls)

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL	(6) WL	(7) WL	(8) WL
FRGN	0.242***							
	(0.0806)							
NatNum		0.0135***						
		(0.00483)						
NatFrac			0.152***					
			(0.0545)					
NatFrac1				0.204***				
				(0.0674)				
NatFrac2					0.222**			
					(0.0918)			
NatPol						0.0981*		
						(0.0525)		
NatPol1							0.130**	
							(0.0547)	
NatPol2								0.125*
								(0.0717)
Constant	0.369	0.435	0.374	0.402	0.406	0.402	0.443	0.451
	(0.290)	(0.288)	(0.290)	(0.293)	(0.298)	(0.292)	(0.297)	(0.303)
R^2	0.491	0.488	0.489	0.491	0.486	0.481	0.483	0.479
Years	✓	✓	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓	✓

N = 330

Robust standard errors clustered by Team in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

We also see statistically significant positives coefficients for the polarisation measures in columns (6)-(8), though at lower significance levels than their fractionalisation counterparts. It is therefore unclear whether both national fractionalisation and polarisation have positive effects on team performance or whether the effect of one variable is driven by the correlation

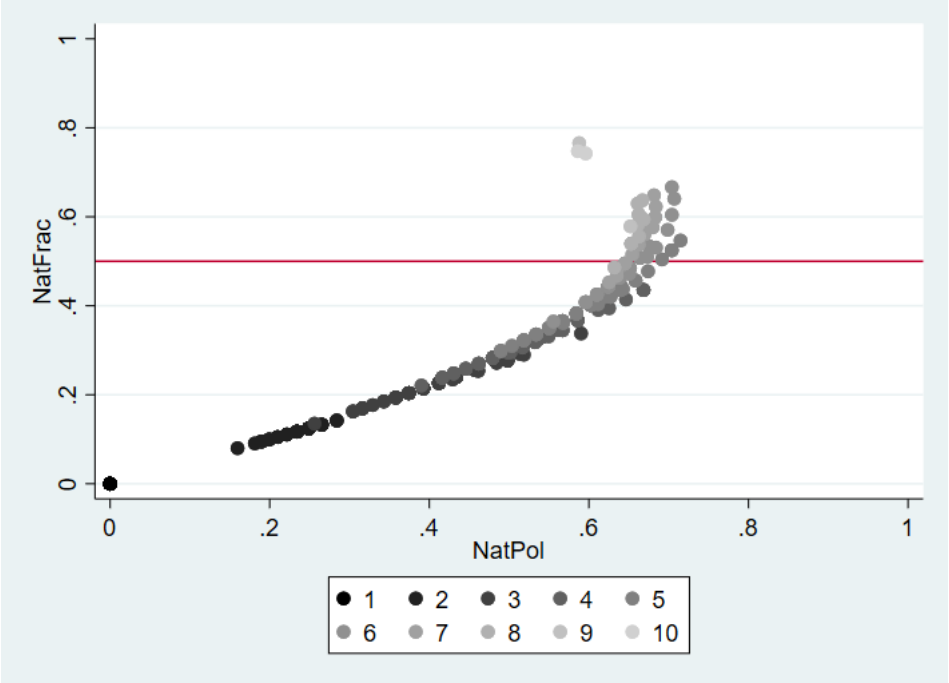
[51] The results are similar when including the size of the entire roster in place of the size of the core team, but we retain the core team measure for consistency throughout our investigation.

with the other, particularly at low levels of fractionalisation. This is investigated further in Table 6.10.

In column (1) of Table 6.10, $FRGN$ and $NumNat$ are included simultaneously. Although the signs of the coefficients are consistent with our previous discussion, neither is statistically significant. This appears to be due to the high correlation between the variables with much larger standard errors compared to including them separately. In columns (2)-(4) the measures of national fractionalisation are included simultaneously alongside their polarisation counterparts.[52] The coefficients on fractionalisation retain their positive sign and statistical significance. Though the standard errors again increase in size, the magnitude of the coefficients also increase. The coefficients on polarisation become negative but are statistically insignificant, except in the specification where two primary languages are considered for linguistic distance in column (4). When including linguistic distance in the measures, the coefficients are larger. However, we should be wary that the value of these measures becomes smaller when considering more languages. We therefore have only tentative evidence that including linguistic distance intensifies these relationships. We speculate that the correlation between fractionalisation and polarisation at low levels of fractionalisation may still be preventing us from isolating the individual effects of each. We explore this further in columns (5)-(7).

To isolate the individual effects of national fractionalisation and polarisation, we first create a dummy variable called $HighFrac$ which equals 1 when $NatFrac$ is greater than 0.5 and 0 otherwise. This splits the sample as shown in Figure 6.3.

Figure 6.3: National Fractionalisation and Polarisation: Spitting the sample



The value of 0.5 was chosen as it is the point where measures of fractionalisation and polarisation can begin to disagree whether a team becomes more 'diverse' theoretically (see simulations in Figure 3.6 and discussion in Section 3.2). It is also the point where the mea-

[52] Various combinations of the fractionalisation and polarisation were trialled, though these did not reveal any additional insights. Therefore only the results for the counterpart pairs are presented (i.e. $NatFrac$ and $NatPol$; $NatFrac1$ and $NatPol1$; and $NatFrac2$ and $NatPol2$).

Table 6.10: Regression Results for Multiple National Diversity Measures

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL	(6) WL	(7) WL
FRGN	-0.223						
	(0.264)						
NatNum	0.00119						
	(0.0156)						
NatFrac		0.349***					
		(0.118)					
NatPol		-0.186					
		(0.112)					
NatFrac1			0.433**				
			(0.162)				
NatPol1			-0.199				
			(0.128)				
NatFrac2				0.587**			
				(0.216)			
NatPol2				-0.302*			
				(0.162)			
HighFrac=0 × NatFrac					0.374		
					(0.336)		
HighFrac=1 × NatFrac					0.453***		
					(0.144)		
HighFrac=0 × NatPol					-0.197		
					(0.265)		
HighFrac=1 × NatPol					-0.281*		
					(0.145)		
HighFrac=0 × NatFrac1						0.337	
						(0.265)	
HighFrac=1 × NatFrac1						0.949***	
						(0.232)	
HighFrac=0 × NatPol1						-0.111	
						(0.191)	
HighFrac=1 × NatPol1						-0.656***	
						(0.218)	
HighFrac=0 × NatFrac2							0.587*
							(0.332)
HighFrac=1 × NatFrac2							1.270***
							(0.297)
HighFrac=0 × NatPol2							-0.268
							(0.227)
HighFrac=1 × NatPol2							-0.921***
							(0.288)
Constant	0.373	0.410	0.389	0.386	0.417	0.405	0.409
	(0.289)	(0.292)	(0.291)	(0.292)	(0.301)	(0.293)	(0.296)
R^2	0.491	0.494	0.494	0.492	0.494	0.498	0.497
Years	✓	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓

N = 330

Robust standard errors clustered by Team in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

sures begin to become less linearly dependent in our data. Perhaps unfortunately, this is at about the 87^{th} percentile of the data, leaving only 43 of 330 teams with fractionalisation greater than 0.5. However, considering a much lower cut-off begins to conflate the instances where fractionalisation and polarisation are highly correlated with the instances where they are not.

In columns (5)-(7) of Table 6.10, we interact $HighFrac$ with the measures of national fractionalisation and polarisation in our regression specifications. The results reveal that when fractionalisation is low, the model struggles to differentiate between fractionalisation and polarisation due to their correlation at these values. Although the coefficient on $NatFrac$ is positive and the coefficient on $NatPol$ is negative when $HighFrac$ equals zero, both are statistically insignificant in most specifications (with a significantly positive coefficient at a 10% level only on $NatFrac2$ in column (7)). However, when $NatFrac$ is greater than 0.5 and $HighFrac$ therefore equals one, the coefficients on the fractionalisation measures are positive and highly significant at a 1% level. The coefficients on polarisation are negative and statistically significant at a 10% level in column (5) and at a 1% level in columns (6) and (7). We also considered splitting the sample based on values of $NatFrac1$ and $NatFrac2$ for columns (6) and (7), but find splitting the sample based on $NatFrac$ to be preferable, both theoretically and empirically.

6.3.1 Robustness Checks

For our robustness checks, we use the simpler specification of our model without the interaction term with $HighFrac$. We include the measures of fractionalisation and polarisation which incorporate linguistic distance considering two languages, as these measures appeared to show the greatest significance in our analysis thus far. Columns (1) and (2) of Table 6.11 reiterate that we do not find significant effects without including additional control variables. We even see the signs of the coefficients switch, but with high standard errors, these are statistically insignificant.

Table 6.11: Robustness Check (Relaxing Controls and Definition of the Team)

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL
NatFrac2	-0.0836	0.587**	0.468**	0.609***	0.561***
	(0.366)	(0.216)	(0.204)	(0.216)	(0.141)
NatPol2	0.192	-0.302*	-0.226	-0.336**	-0.288***
	(0.286)	(0.162)	(0.159)	(0.167)	(0.0997)
Constant	0.470***	0.386	0.540***	0.420	0.415
	(0.0308)	(0.292)	(0.145)	(0.294)	(0.266)
R^2	0.00908	0.492	0.459	0.451	0.496
Years	✓	✓	✗	✓	✓
Teams	✓	✓	✓	✗	✓
Controls	✗	✓	✓	✓	✓
CoreTeam	✗	✗	✗	✗	✓

N = 330
Robust standard errors clustered by Team in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

When excluding year effects in column (3), the coefficient on $NatFrac2$ is positive and significant at a 5% level, while the coefficient on $NatPol2$ is negative though insignificant. In

column (4), the team effects are excluded and we see a similar result, but the coefficient on $NatFrac2$ has increased to a 1% level, while the coefficient on $NatPol2$ is significant at a 5% level. It appears that the inclusion of year effects is more important for determining the effects of national diversity, unlike for wage dispersion where teams effects were more important. Finally, in column (5), we find that the effects are more highly significant when considering national diversity within the 'core team', with both coefficients statistically significant at a 1% level.

In Table 6.12, we see the robustness of our results to different measures of team performance. The results are consistent in the first four columns, with the highest significance for both coefficients in column (2), when considering the natural logarithm of regular season win percentage. Interestingly, the coefficient on $NatFrac2$ is a little less than twice as large as the coefficient on $NatPol2$. This is true fairly consistently across our different performance measures. When introducing national diversity into a previously nationally homogeneous team then, for a given linguistic distance, introducing players from a second nationality will increase polarisation twice as much as fractionalisation. These results may therefore model the initial difficulties of introducing national diversity into teams. However, as more nationalities are added into the team, the rise in fractionalisation becomes higher than that of polarisation, and then the relationship becomes negative. Therefore, the positive effects of national diversity may be felt more strongly by national diverse teams.

Table 6.12: Robustness Check (Performance Measures)

	(1) WL	(2) lnWL	(3) PDG	(4) SRS	(5) ORTG	(6) DRTG
NatFrac2	0.587**	1.435***	17.89**	17.32**	7.063	11.04*
	(0.216)	(0.516)	(7.127)	(6.932)	(6.550)	(5.497)
NatPol2	-0.302*	-0.759**	-9.887*	-9.724*	-4.403	-5.494
	(0.162)	(0.369)	(5.492)	(5.350)	(4.841)	(4.366)
Constant	0.386	-1.336*	0.730	0.643	110.5***	-109.6***
	(0.292)	(0.663)	(7.635)	(7.329)	(6.670)	(4.520)
R^2	0.492	0.468	0.457	0.456	0.510	0.408
Years	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓

N = 330
Robust standard errors clustered by Team in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

When considering offensive and defensive ratings in columns (5) and (6), we see consistency in the signs of the coefficients but with lower significance levels. Interestingly, the positive effect of national fractionalisation seem to found more consistently when analysing the defensive side of the game, with the positive coefficient on $NatFrac2$ statistically significant at a 10% level.

Table 6.13 shows the effects of national diversity on team performance estimated using different regression techniques. The 'complete pooling' Fixed Effects (FE) approach used in this chapter is presented in column (1), while column (2) presents the results of the 'no pooling' Random Effects (RE) approach, and column (3) presents the partial pooling Multilevel modelling (MLM) approach. The positive effects of $NatFrac2$ on team performance are consistent

Table 6.13: Robustness Check (Alternative Models)

	(1) WL	(2) WL	(3) WL
NatFrac2	0.587**	0.504**	0.609***
	(0.216)	(0.206)	(0.210)
NatPol2	-0.302*	-0.270	-0.336**
	(0.162)	(0.165)	(0.163)
Constant	0.386	0.559***	0.420
	(0.292)	(0.150)	(0.285)
R^2	0.492	0.425	-
log likelihood	309.7	-	269.0
Years	✓	✗	✓
Teams	✓	✗	✓
Controls	✓	✓	✓
Model	FE	RE	MLM

N = 330
Robust standard errors clustered by Team in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

across all three models, significant at a 5% level for FE and RE, and at a 1% level for MLM. This suggests that this effect is driven by both variation of teams over seasons (within) and variation across teams (between). The MLM approach captures both of these effects resulting in a larger coefficient and higher statistical significance. The negative coefficient on $NatPol2$ is consistent across the models, but only statistically significant at a 10% level for FE, at a 5% level for MLM and is marginally statistically insignificant for RE (with a p-value of 0.102). The negative effects of $NatPol2$ seems to be stronger within team observations rather than between them. Though there are efficiency gains when also considering the between variation, as the coefficient on $NatPol2$ is statistically significant at a higher level when using the MLM approach in column (3), compared to the FE approach in column (1) which ignores this between variation.

6.4 National Diversity, Wage Dispersion and Team Performance

In this section, we present a full model which combines our results relating to both wage dispersion (from Chapter 5) and national diversity (explored earlier in this chapter). First, we discuss some correlations, before presenting the model and conducting robustness checks.

Table 6.14: Correlations Between Wage Dispersion and National Diversity Variables

	NatFrac	NatFrac1	NatFrac2	NatPol	NatPol1	NatPol2
WageSD	-0.017	-0.051	-0.043	-0.007	-0.047	-0.031
WagePol	-0.029	-0.09	-0.084	-0.056	-0.125**	-0.109**

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6.14 presents the pairwise correlations between the wage dispersion and national diversity measures, which were previously found to significantly influence team performance. Only the correlations across these groups of measures are presented, as the correlations within them have previously been discussed. All of the pairwise correlations are negative suggest-

ing that more nationally diverse teams have lower wage dispersion. Though most of these correlations are small and insignificant. However, we do see significantly negative coefficients between $WagePol$ and the measures of national polarisation which consider linguistic distance ($NatPol1$ and $NatPol2$). This suggests that teams with large cultural differences tend to have less wage polarisation.

Column (1) of Table 6.15 repeats the final specification from the investigation into the effects of wage dispersion from Chapter 5. The national diversity results are repeated in column (2) for the basic measures, and in column (3) for the measures which incorporate linguistic distance considering two languages. In columns (4) and (5), the combined models are presented, which consider both wage dispersion and national diversity.

Table 6.15: Regression Results of Wage Dispersion and National Diversity Measures

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL	(6) WL
WageSD	0.0287***			0.0290***	0.0287***	0.0289***
	(0.00862)			(0.00807)	(0.00807)	(0.00823)
WagePol	-0.313***			-0.299**	-0.300**	-0.247*
	(0.112)			(0.118)	(0.117)	(0.143)
NatFrac		0.349***		0.363***		
		(0.118)		(0.105)		
NatPol		-0.186		-0.206**		
		(0.112)		(0.0902)		
NatFrac2			0.587**		0.602***	0.605***
			(0.216)		(0.203)	(0.200)
NatPol2			-0.302*		-0.325**	-0.271
			(0.162)		(0.150)	(0.233)
WagePol × NatPol2						-0.270
						(0.695)
Constant	0.832**	0.410	0.386	0.742**	0.708**	0.691**
	(0.305)	(0.292)	(0.292)	(0.283)	(0.284)	(0.280)
R^2	0.503	0.494	0.492	0.525	0.522	0.522
Years	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓

N = 330
Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The results are broadly similar and remain consistent with our previous conclusions. However, we see a slight increase in significance for our measures of national polarisation and a slight decrease for wage polarisation. We try interacting these terms in column (6) but do not find significant results. We also trial other combinations of interactions but again find no significant results, suggesting there is no interaction between the effects of wage dispersion and national diversity within teams. However, there does seem to be some correlation between national polarisation and wage polarisation. Controlling for each has a slight effect on the other, but does not change our main conclusions.

Due to the insignificance of any interaction terms between wage dispersion and national diversity, our main conclusions are best summarised by column (5) of Table 6.15. This model predicts that a one standard deviation increase in $WageSD$ ($1.581m) is expected to increase

a team's regular season win percentage by 0.045 or 4.5 percentage points. A one standard deviation increase in $WagePol$ (0.075) is predicted to decrease a team's win percentage by 0.022 or 2.2 percentage points. Therefore, teams should consider how changes in their wage structure affect both of these elements of wage dispersion to assess the overall impact. Changes which ideally increase the standard deviation of wage while decreasing wage polarisation are expected to have the most beneficial effects on a team's regular season win percentage. These changes would decrease clustering at the extremes of the distribution but maintain a spread across different wage levels.

To give a more intuitive interpretation of these results, it is useful to consider the marginal effects of specific managerial decisions in the hiring process. An issue with considering the marginal effect of adding a given player to a team is that the effects on our variables of interest, and therefore the predicted effects on performance, are sensitive to the initial values of our variables of interest.[53] As such, we consider a team of 18 players (the mean roster size of our sample), with each player receiving the mean wage of each wage rank (see Appendix 6.B for full details on these mean wages for each wage rank and the subsequent changes described). This is our base case, labelled Team A, which would have a $WageSD$ of 4.736 and $WagePol$ of 0.191. A description of Team A and our subsequent examples are summarised in Table 6.16. Table 6.17 then summarises the values of $WageSD$ and $WagePol$, and the predicted effect based on our main specification; as well as how each differs from our base case in terms of win percentage and absolute wins (assuming an 82-game season).

Table 6.16: Description of Example Teams

Team	Description (change from base case)
A	Base Case - each player receives the mean wage of the relevant wage rank
B	\$0.5m is transferred from those ranked 13^{th}-15^{th} to those ranked 4^{th}-6^{th}
C	\$0.5m is transferred from those ranked 4^{th}-6^{th} to those ranked 13^{th}-15^{th}
D	\$0.5m is transferred from those ranked 1^{st}-3^{rd} to those ranked 7^{th}-9^{th}, and from those ranked 10^{th}-12^{th} to those ranked 16^{th}-18^{th}
E	Those previously ranked 1^{st} and 18^{th} are paid an average of these wage ranks
F	Those previously ranked 5^{th} and 14^{th} are paid an average of these wage ranks

From our starting point of Team A, we can firstly consider similar changes to those discussed in Chapter 3. For Team B, we consider a transfer across the middle of the distribution, by transferring \$0.5m from those in the second bottom sixth of the distribution (those ranked 13^{th}-15^{th}) and giving the same amount to each player ranked in the second top sixth of the distribution (ranked 4^{th}-6^{th}).[54] Initially, such a transfer from backups to potential starters may seem like an efficient use of resources. However, our model would predict Team B to achieve 0.579 less wins over the course of a season. This is driven by the larger increase in wage polarisation due to clustering at the top and bottoms of the distribution. This causes a net negative effect on performance despite a small increase in the standard deviation of wages, as the increase in wage polarisation is relatively larger. This is further emphasised by Team C

[53] For example, adding a highly paid player to a previously egalitarian team would have a large effect on the standard deviation of wages, whereas adding that same player to a team which already contains some superstars might have a smaller effect on the standard deviation and add to the polarisation of wage.

[54] \$0.5m was chosen as the transfer amount for demonstrational purposes as it is a non-negligible amount which can also be taken from the lower ranked players without causing them to earn negative wages.

which considers the opposite change from the base case: a \$0.5m transfer from those ranked 4^{th}-6^{th} to those ranked 13^{th}-15^{th}. Team C is predicted to outperform Team A by 0.397 wins. These examples suggest that it is important to consider team cohesion, and funnelling wages away from the less well paid towards the already better paid players may not always be a good strategy. An alternative explanation may be that the marginal improvement in talent of directing the spending lower down the wage distribution may be larger, meaning that Team C use their wage distribution more efficiently to maximise their overall playing talent.

Table 6.17: Predicted Effects of Example Teams

Team	WageSD	WagePol	Pred. Effect	Δ from base case	Δ in wins
A	4.7360	0.1910	0.0802	-	-
B	4.8430	0.2250	0.0732	-0.0071	-0.579
C	4.6450	0.1660	0.0851	0.0048	0.397
D	4.5380	0.2180	0.0664	-0.0138	-1.133
E	3.8050	0.1720	0.0589	-0.0213	-1.748
F	4.6520	0.1550	0.0886	0.0083	0.683

We may also consider transfers on either side of the distribution. Team D represents transfers, compared to the base case, of \$0.5m from the top sixth of the distribution (ranked 1^{st}-3^{rd}) to the third sixth (ranked 7^{th}-9^{th}), and from the fourth sixth (ranked 10^{th}-12^{th}) to those in the bottom sixth (ranked 16^{th}-18^{th}). Creating equality on either side of the distribution in this way decreases $WageSD$ while increasing $WagePol$ and therefore harms performance. Team D would be predicted to win 1.133 less games than Team A.[55]

However, as a cautionary tale against interpreting these examples as simply indicating that more overall equality improves performance, firstly a strictly egalitarian team would be predicted to perform notably worse than any of the hypothetical teams mentioned. Specifically, with $WageSD$ and $WagePol$ both equalling zero, such a team would be predicted to win 6.579 games less than Team A. Indeed, these findings do not discourage variation in wages. Consider Team E, which differs from Team A only in that the previously 1^{st} and 18^{th} ranked players are instead paid a wage which is the mean of their previously wages. Replacing these wages at either extreme of the distribution with the average of the two decreases both $WageSD$ and $WagePo$l, but the decrease in $WageSD$ is relatively larger, causing a predicted decrease of 1.748 wins. This is because the average of these wages is still reasonably large, as both would rank joint 4^{th} with the new wage. It is therefore causing clustering at the top end of the top end of the distribution, which mitigates some of the effect of the decreased spread. However, if we were to replace the wages of the 5^{th} and 14^{th} ranks (the 25^{th} and 75^{th} percentile values) with an average of the two previous wages, we would expect to see an increase in performance. This distribution is considered in Team E. It maintains the extreme values, but the new clustering is nearer to the middle of the distribution. As a result, $WageSD$ decreases slightly, but the larger middle causes $WagePol$ to decrease relatively more. Overall, Team E is predicted to outperform Team A by 0.683 wins over the course of an 82-game season.

Overall, the marginal impact of adding a player to a team depends on the existing values of wage standard deviation and polarisation. The examples described consider changes from

[55] Note that the decrease in $WageSD$ is caused more by the changes in the top half of the distribution as these are further away from the mean due to the positive skew of the wage distribution. Whereas the change in $WagePol$ is larger for the change in the bottom half of the distribution, but each contributes roughly equally to the overall fall in performance.

a team which pays the sample average to each wage rank. This example team and each of the examples given would be predicted to outperform a strictly egalitarian team - variation in wage is important to induce effort. It is beneficial for a team to have players towards the extremes of the distribution in order to maximise the standard deviation of wage, but this should be balanced by incorporating a large 'middle' to limit polarisation.

When assessing the impact of national diversity, the model in column (5) of Table 6.15 predicts that a one standard deviation increase in $NatFrac2$ (0.09) is expected to increase a team's win percentage by 0.054 or 5.4 percentage points. A one standard deviation increase in $NatPol2$ is predicted to decrease the regular season win percentage of teams by 0.037 or 3.7 percentage points. Again, the overall impact will depend on how changes in national diversity affect the overall national structure of the team. As displayed in Figure 6.2b, it is possible for teams to increase $NatFrac2$ without increasing $NatPol2$, and this is generally achieved by increasing the number of national groups and avoiding two dominant national groups. Examples of fully nationally fractionalised teams do not exist in our dataset. While the evidence suggests that such teams would perform well, one should be cautious in making such out-of-sample predictions.

Again, it is useful to provide a more intuitive interpretation of these results in terms of the hiring decision. As with our results regarding wage dispersion, the marginal effect of adding an additional player to a team is dependent on the initial values of our variables of interest, so it is important to establish a base case. In this instance, we choose a nationally homogeneous team as our base case. This nationally homogeneous team will have a value of zero for both $NatFrac$ and $NatPol$. To simplify these interpretations, we consider the specification in column (4) of Table 6.15, which does not consider linguistic distance between national groups. The base case should be considered to be 18 (the sample mean roster size) Americans, just because the only instances of nationally homogeneous teams in our data are all American teams.

From here, we can consider the marginal effect of replacing subsequent American players with foreign players from differing numbers of nationalities. We consider adding from 1 up to 10 foreign players, as 10 is the maximum number of foreign players in a team in our sample. Again for simplicity, we assume that the same number of players are added from each additional nationality. Therefore, for one additional nationality we can add from 1 to 10 players, but only multiples of two are added for two additional nationalities, multiples of three for three, etc., up to five additional nationalities. This is done because, otherwise, different values of $NatFrac$ and $NatPol$ could be achieved depending on the proportion of foreign players of each nationality. For example, if we considered five foreign players from two additional nationalities, we would get different values of $NatFrac$ and $NatPol$ depending on whether we considered this ratio between the national groups to be 4:1 or 3:2. We therefore consider evenly sized additional national groups for simplicity. Table 6.18 summarises the predicted effects for different numbers of foreign players from differing numbers of nationalities. Firstly the predicted effect on win percentage is reported, followed by the change in absolute wins compared to the base case (a nationally homogeneous team). The predicted change in absolute wins is also plotted in Figure 6.4.

Given that a nationally homogeneous team would have values of zero for both $NatFrac$ and $NatPol$, Table 6.18 shows that adding one additional nationality is predicted to decrease performance. This ranges from a predicted decrease of 0.422 wins over a 82-game season

Table 6.18: Effect of Adding Foreign Players from Differing Numbers of Additional Nationalities

No. of Add. Nat.	Number of Foreign Players 1	2	3	4	5	6	7	8	9	10
1	-0.005 -0.422	-0.010 -0.794	-0.014 -1.116	-0.017 -1.389	-0.020 -1.612	-0.022 -1.786	-0.023 -1.910	-0.024 -1.984	-0.025 -2.009	-0.024 -1.984
2		-0.003 -0.262		0.006 0.458		0.021 1.745		0.039 3.180		0.053 4.347
3			0.005 0.408			0.039 3.200			0.082 6.706	
4				0.019 1.521				0.084 6.875		
5					0.037 3.006					0.133 10.927

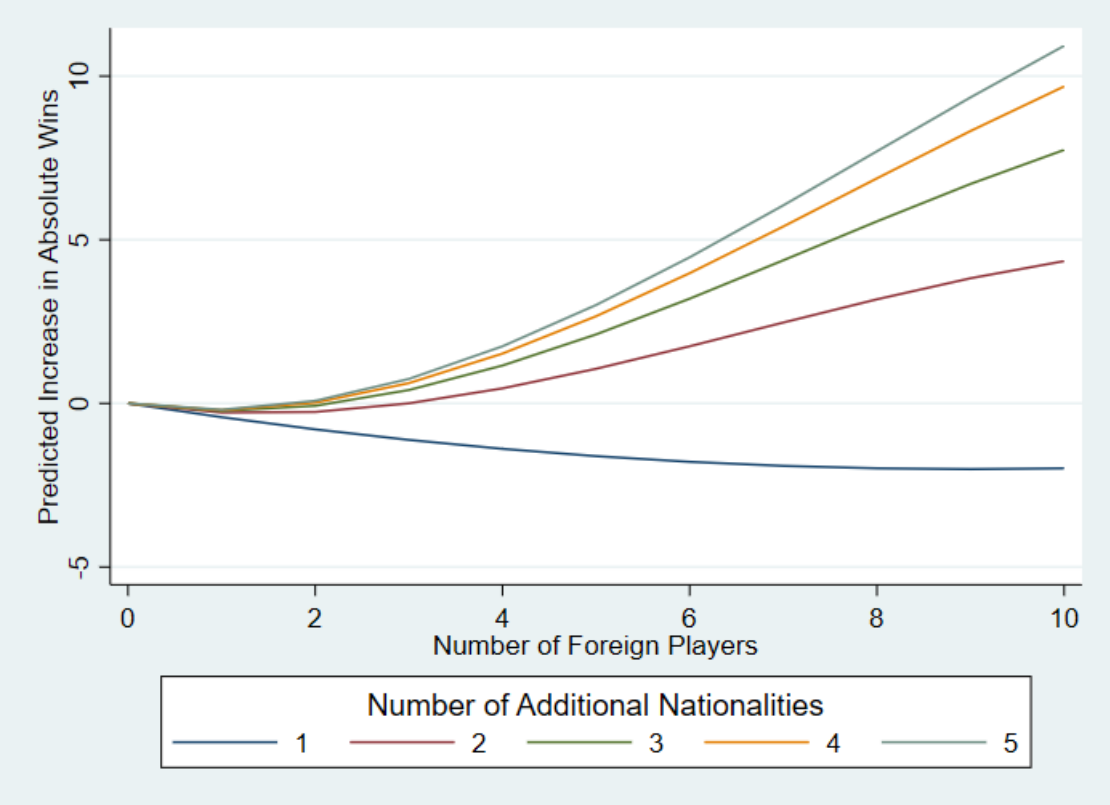
compared to a nationally homogeneous team for one foreign player, to a decrease of just over 2 wins for nine foreign players from one additional nationality which would represent a perfectly polarised team (nine Americans and nine players from another nationality). As shown in Section 3.2, when there are just two nationalities (American and one additional), $NatPol$ will always be twice as large as $NatFrac$. As the positive coefficient on $NatFrac$ is less than twice as large as the negative coefficient on $WagePol$, the negative effects of national polarisation are always predicted to dominate for a team consisting of two nationalities.

When considering two additional nationalities (so three in total), a negative effect is still predicted when there are only two foreign players (one from each additional nation). Such a team is predicted to win 0.262 less games over an 82-game season compared to a nationally homogeneous team. However, as more American players are replaced by foreign players from these two additional nations, the positive effects of $NatFrac$ outweigh the negative effects of $NatPol$ and increasingly better performance is predicted. A team with 10 foreign players, split into five from each additional nation, and therefore eight Americans, is predicted to win 4.347 more games than a nationally homogeneous team over an 82-game season. In general, the more additional nationalities are considered, the more positive the effect on performance. This is also shown in Figure 6.4. The effect (compared to a nationally homogeneous team) of increasing the number of foreign players from only one additional nationality is downward sloping, though at a decreasing rate. For two of more nationalities, the effects generally increase with the number of foreign players, and with a steeper slope as the number of additional nationalities increases. Though it should be noted that each line begins as downward sloping[56], they quickly begin to slope upwards and show positive predicted performance compared to a nationally homogeneous team. However, these general interpretations of the results again come with some cautionary words.

It is possible for a team with more nationalities to be more polarised than a team with less nationalities, depending on the proportions of players in each. For example, a team with four nationalities overall split 8:8:1:1 would be highly polarised, despite being relatively

[56] To create Figure 6.4, for demonstational purposes, each additional national group is modelled to account for an equal proportion of the number of foreign players even if this is not a whole number. This gives a line of best fit and explains why the effect of one foreign player in total from three additional nationalities, for example, can be read from the graph.

Figure 6.4: Effect of Adding Foreign Players from Differing Numbers of Additional Nationalities



fractionalised. Such a team would still be predicted to outperform a nationally homogeneous team by 2.599 wins. But Table 6.18 shows that a team with three nationalities (two additional nationalities) split 8:5:5 would be predicted to win a further 1.748 games, as we have seen (4.347 more wins than a nationally homogeneous team). Therefore, while more nationalities in a team is generally beneficial, the size of each national group is also important. These results suggest that having many equally small national groups is the most beneficial structure to maximise performance. Table 6.18 and Figure 6.4 suggests that a team with five additional nationalities (six in total) split 8:2:2:2:2:2, with eight Americans, is predicted to outperform a nationally homogenous team by almost 11 wins. Such a team would have a $NatFrac$ of 0.741 and a $NatPol$ of 0.658. The only team to have comparable values for both measures was the San Antonio Spurs for three consecutive season from 2012 to 2015, in which they had win percentages of 70.7%, 75.6% and 67.1%, respectively. These are the three observations that were on the backwards bending section of Figure 6.1. It seems that these observations are contributing to our results and suggest that a highly fractionalised national structure may have helped drive this successful period for the San Antonio Spurs, where they were once crowned NBA champions and once runner up, having twice won the Western Conference.

6.4.1 Robustness Checks

Table 6.19 shows the robustness of our results to relaxing some of the controls and changing the definition of the team. In column (1), we see the results regarding wage dispersion show robustness when removing the control variables. Although the national diversity measures have the same signs, they are not statistically significant when the controls are removed, even when controlling for wage dispersion. Column (2) again shows that, in the presence of relevant controls, dispersion/diversity in a more traditional sense has positive effects on performance, while polarisation of both wage and nationality has negative effects.

Table 6.19: Robustness check (Relaxing Controls and Definition of the Team)

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL
WageSD	0.0647***	0.0287***	0.0279***	0.0277***	0.0305***
	(0.00733)	(0.00807)	(0.00700)	(0.00877)	(0.00720)
WagePol	-0.566***	-0.300**	-0.266**	-0.203	-0.505***
	(0.136)	(0.117)	(0.109)	(0.129)	(0.147)
NatFrac2	0.239	0.602***	0.461**	0.652***	0.577***
	(0.226)	(0.203)	(0.198)	(0.211)	(0.143)
NatPol2	-0.0448	-0.325**	-0.234	-0.384**	-0.305***
	(0.199)	(0.150)	(0.152)	(0.162)	(0.101)
Constant	0.301***	0.708**	1.003***	0.816***	0.662**
	(0.0384)	(0.284)	(0.175)	(0.305)	(0.263)
R^2	0.276	0.522	0.488	0.472	0.535
Years	✓	✓	✗	✓	✓
Teams	✓	✓	✓	✗	✓
Controls	✗	✓	✓	✓	✓
CoreTeam	✗	✗	✗	✗	✓

N = 330
Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Column (3) suggests that the wage dispersion results are robust to the removal of the year effects. However, while the positive coefficient on $NatFrac2$ is still positive and statistically significant at a 5% level, $NatPol2$ is insignificant, though it keeps its negative sign. When excluding team effects in column (4), it is the national diversity measures which appear more robust. And while the effects of $WageSD$ remain robust, the negative effects of $WagePol$ lose their statistical significance, though again retaining their sign. Finally, column (5) shows that the results of wage dispersion and national diversity are robust, and even seem to be intensified, when considering just the core team of players.

Table 6.20: Robustness check (Performance Measures)

	(1) WL	(2) lnWL	(3) PDG	(4) SRS	(5) ORTG	(6) DRTG
WageSD	0.0287***	0.0629***	0.876***	0.865***	0.487**	0.401**
	(0.00807)	(0.0169)	(0.269)	(0.253)	(0.228)	(0.157)
WagePol	-0.300**	-0.649**	-8.408**	-8.206**	-6.552**	-2.111
	(0.117)	(0.297)	(3.365)	(3.247)	(2.747)	(2.328)
NatFrac2	0.602***	1.466***	18.29**	17.71**	7.374	11.15*
	(0.203)	(0.503)	(6.794)	(6.589)	(6.416)	(5.658)
NatPol2	-0.325**	-0.808**	-10.51**	-10.33**	-4.926	-5.617
	(0.150)	(0.355)	(5.109)	(4.959)	(4.695)	(4.453)
Constant	0.708**	-0.627	10.87	10.70	115.3***	-104.2***
	(0.284)	(0.614)	(7.732)	(7.439)	(6.572)	(4.832)
R^2	0.522	0.494	0.487	0.487	0.529	0.420
Years	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓

N = 330
Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6.20 confirms the robustness of our results, when considering wage dispersion and national diversity simultaneously, to different measures of team performance. The results are consistent across columns (1)-(4) adding reliability to our findings. Columns (5) and (6) suggest that a higher standard deviation has benefits for both the offensive and defensive side of the game. We have evidence that a lower wage polarisation improves offensive performance, while a higher national fractionalisation improves defensive performance. The signs of the other coefficients are in line with our main findings but lack statistical significance in these final two specifications.

Lastly, Table 6.21 shows the robustness of our main results across different modelling techniques; Fixed Effects (FE), Random Effects (RE) and Multilevel Modelling (MLM).

Table 6.21: Robustness checks (Alternative Models)

	(1) WL	(2) WL	(3) WL
WageSD	0.0287***	0.0273***	0.0279***
	(0.00807)	(0.00771)	(0.00814)
WagePol	-0.300**	-0.180	-0.233*
	(0.117)	(0.123)	(0.122)
NatFrac2	0.602***	0.540***	0.638***
	(0.203)	(0.209)	(0.198)
NatPol2	-0.325**	-0.314*	-0.367**
	(0.150)	(0.166)	(0.151)
Constant	0.708**	1.040***	0.779***
	(0.284)	(0.182)	(0.288)
R^2	0.522	0.448	-
log likelihood	319.9	-	277.3
Years	✓	✗	✓
Teams	✓	✗	✓
Controls	✓	✓	✓
Model	FE	RE	MLM

N = 330
Robust standard errors clustered by Team in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

The negative effects of polarisation, both in terms of wages and nationality, seem to be driven by the variation within team observations over seasons, with higher statistical significance seen for $WagePol$ and $NatPol2$, in column (1) of Table 6.21, for the FE model. The positive effects of the more traditionally used measures of both wage dispersion and national diversity, $WageSD$ and $NatFrac$, are more consistent across the three models in terms of statistical significance. This suggests that it is driven by both between and within variation.

6.5 Summary and Discussion

Through our investigation into the effects of national diversity on team performance, we have gained a number of insights. Firstly, the relationship between national diversity and team performance is difficult to identify without a suitable set of controls. Before including such controls, we did not identify any statistically significant relationships between our measures of national diversity and team performance. Additionally, some of the qualitative results were

misleading and changed sign with the inclusion of relevant controls, particularly those relating to the size of the team.

In general, we find a positive relationship between team performance and multiple different measures of national diversity. However, further analysis reveals that, in particular, national fractionalisation has positive effects on performance, while national polarisation has negative effects. Nationally fractionalised teams contain a range of cultural perspectives which aids the 'search for best practice' (Lazear, 1999). Increasing this type of diversity offers the potential for cross-cultural learning and positive spillover effects. For a given number of nationalities, fractionalisation is maximised when each national group is equally small which removes the existence of a single or small number of dominant national groups.

Conversely, national polarisation is high when there are a small number of large groups and is maximised when there are two national groups of even size. Such a structure fosters an in-versus out-group mentality which is a feature of theories relating to polarisation. A polarised team is less likely to cooperate and more likely to clash and cause conflict, thereby hindering the process of cultural mixing. So while it is not necessarily optimal for a team to be homogeneous, the key is to introduce heterogeneity in a way which does not cause high polarisation.

Fractionalisation can be maximised and polarisation minimised simultaneously when a team consists of many relatively small national groups. Our results therefore imply that such a structure is optimal for team performance. This also has implications for teams first introducing diversity. When adding an additional nationality to a homogeneous team, both fractionalisation and polarisation will increase, with the latter increasing twice as fast. Our regression results suggest that the positive effects of fractionalisation are a little less than twice the size of the negative effects of polarisation. This therefore models the initial difficulties of introducing national diversity. However, as the number of nationalities and fractionalisation increase, the positive effects of fractionalisation are more likely to dominant the negative effects of polarisation.

At even higher levels of fractionalisation (above 0.5), polarisation begins to decrease as fractionalisation increases. We find evidence that the positive effects of fractionalisation, as well as the negative effects of polarisation, are even stronger above this threshold. Our results suggest that a 'truly diverse' team with many small national groups would perform best. Potentially even with each player having a different nationality. However we must, of course, be wary when extrapolating empirical results to make out-of-sample predictions. Despite the large number of nationalities which offer national diversity within the NBA, it is still a league which is largely dominated by players from the USA. We have few data points which enter the region of high fractionalisation where the relationship with polarisation becomes negative. Perhaps further increases in national diversity in future datasets will allow for a more robust test of this theory.

It is important to also consider other plausible mechanisms through which the observed relationship between national diversity and team performance could exist, beyond those considered thus far. Particularly in terms of the positive effects of national fractionalisation, it is plausible that teams which scout globally have a higher number of nationalities within their teams, as well as being able to acquire the best talents which contribute towards better performance. Though it should be noted that wage expenditure is controlled for, which should

act as a proxy for ability, this mechanism may be particularly plausible if national discrimination in wages exists.[57] As the best rewarded professional basketball league in the world (Fraschilla, 2017; Statista, 2020), it is possible that NBA teams are able to acquire talent from foreign leagues at a relatively lower wage compared to domestic talent. Especially if not all NBA teams are scouting globally, or in particular countries, as they may well be for domestic talent. Regardless of the mechanism, our results suggest that nationally fractionalised, but not polarised, teams perform best.

We also experiment with proxies of cultural distance and incorporate these into fractionalisation and polarisation measures. Cultural distance is proxied by linguistic distance, but unlike other studies, we allow the cultural makeup of a nation to be determined by more than just its most popular spoken language. Measures are created which allow for between one and up to four different languages, weighted by the proportion of speakers of the language in that nation. This approach distinguishes different nations with the same most popularly spoken language, as well as capturing similarities and differences in other popular languages within a nation which may also influence, or be influenced by, the nation's culture (Chen, 2013).

We show how cultural distance can be incorporated into fractionalisation and polarisation to create measures that capture both the structure of the national diversity and the cultural distance within it. Incorporating linguistic distance into fractionalisation is equivalent to taking the average linguistic distance across players in the team. However, incorporating linguistic distance into polarisation captures a distinct aspect which, for a given linguistic distance, is maximised when two evenly sized national groups exist within the team. Though linguistic distance has been incorporated into fractionalisation measures in previous studies, this study is believed to be the first to show how it can be incorporated into polarisation measures. When using these measures we find some evidence that they intensify the effects of both fractionalisation and polarisation, though the qualitative effects are similar.

Overall, this study suggests that national diversity can have positive effects on team performance, but it is important to manage the structure of the diversity correctly. It also discusses the importance of cross-cultural mixing for the potential benefits of national diversity to be realised and suggests that polarisation can be detrimental to this process.

Furthermore, we find evidence that these results are robust to the inclusion of wage dispersion measures in the regression specification. When regressing team performance on wage dispersion and national diversity measures simultaneously, there is little change in the estimated effects of each, compared to considering them separately. Although some correlation between wage dispersion and national diversity may seem theoretically plausible, there is little evidence of this empirically in our data. The only statistically significant correlation was between wage polarisation and measures of national polarisation which consider linguistic distance. When included simultaneously, the coefficient on wage polarisation decreases and that on national polarisation increases. However, these changes are small and do not alter our main conclusions.

[57] The evidence on this topic, and on racial pay discrimination, is mixed; though there is some suggestion that it is lessening over time (e.g. Groothuis & Hill, 2013; J. R. Hill, 2004; J. R. Hill & Groothuis, 2017; Kahn & Sherer, 1988).

Chapter 7

Wage, Nationality and Player Performance

Chapter Overview

This Chapter investigates the effects of both wage dispersion and national diversity on the performance of individual players within teams. The previous chapters have found that, in terms of both wage and nationality, dispersion/diversity in the traditional sense has positive effects on overall team performance, however polarisation in these dimensions has negative effects. This suggests an ideal structure which encourages heterogeneity but avoids groupings which can cause an in- versus out-group dynamic, discouraging cooperation and increasing conflict within teams. The current chapter aims to investigate whether these aggregate team effects are consistent across all individuals or whether the effects are dependent on where they lie within the team's structures. One of the major difficulties encountered in this part of the investigation was assigning the individual contributions of players towards overall team performance. We consider a number of player performance measures including those which measure the overall contribution of players and others which measure efficiency. These are discussed in detail in Section 4.3.3. We find that the effects of wage dispersion depend on a player's wage and their wage ranking in the team. We also find evidence of a trade-off between the player- and team-level effects of national diversity.

7.1 Introduction

In this chapter, we take the results of Chapters 5 and 6, and apply them at the individual player-level to uncover any additional nuances. We explore whether these aggregate effects are consistent across all players, whether they are stronger for certain groups within the team, or even whether some groups react contrarily but are averaged out at the team-level. More specifically, in Chapter 5, we found that a higher standard deviation of wages within teams has positive effects on performance, in line with Tournament Theory (Lazear & Rosen, 1981); while a higher polarisation of wages has negative effects, in line with Cohesion Theory (Levine, 1991). We therefore investigate whether the positive effects of a higher standard deviation of wage on team performance are driven by the increased performances of those who 'win the tournament' or whether those on lower wages are also motivated, as predicted by Tournament Theory, despite not directly benefiting from the structure. It may even be possible that those on low wages are discouraged by the wage structure, but the positive effects on high earners more than compensates for this. In addition, we investigate whether the negative effects of a polarised wage structure are concentrated to those on higher or lower wages; or whether they are felt equally throughout the entire distribution.

We also propose that similar nuances may exist in terms of national diversity. In Chapter 6, we found that national fractionalisation has positive effects on team performance while national polarisation has negative effects, so we ask whether these are consistent across all players of all national groups within the team. When a team is fully fractionalised, each player will be of a different nationality. But when a team is highly, but not fully, fractionalised, we may ask whether it is beneficial for the individual players to be part of one of the larger or smaller national groups. If it is the former then, paradoxically, this may encourage a polarised national structure, either by players choosing to join teams where their nationality is already highly represented, or via teams observing the higher performance of players in the larger national group and increasing the size of a small number of national groups as a result. Additionally, even in the case of equally-sized national groups, we may ask whether there are benefits of being in the 'middle' of this structure rather than in the extremes, in terms of having linguistic similarities with teammates.

The remainder of this Chapter is structured as follows. Section 7.2 looks briefly at the descriptive statistics of the player-level variables relating to wages and nationality. We then explore some correlations between these, as well as with some of the related team-level variables. Sections 7.3 to 7.5 will present regression analysis, where our variables of interest will consist of player-level and team-level variables, as well as interactions between them. Section 7.3 considers wage dispersion, Section 7.4 then considers national diversity, and finally Section 7.5 considers both simultaneously. Section 7.6 will conclude the chapter with an overall discussion of the results.

7.2 Descriptive Statistics and Correlations

The wages of players have previously been discussed as a control variable in Section 4.4 and descriptive statistics relating to wages have been analysed in Section 4.4.2. Measures of wage dispersion have also been discussed extensively in Section 3.1 and analysed in Chapter 5. In this chapter, we will consider interactions between player-level wage and team-level wage dispersion to explore how the overall team effects manifest at the player-level. In addition,

we will also consider the rank of wage (*WageRank*) for players within their teams. If the predictions of Tournament Theory (Lazear & Rosen, 1981) are correct, then we may expect players to be motivated, not only by their absolute wage level, but also by their comparative wage level within their team. Table 7.1 shows the descriptive statistics for *Wage* at different values of *WageRank*.

Table 7.1: Wages at each Wage Rank

WageRank	Mean Wage ($m)	Std Dev.	Min.	Max.	Freq.
-1	16.801	5.265	4.962	34.683	338
-2	12.938	4.373	4.459	27.740	324
-3	10.385	3.760	3.458	22.642	331
-4	7.848	2.831	2.382	19.300	328
-5	6.198	2.317	1.552	15.550	336
-6	5.000	1.809	1.423	13.760	328
-7	4.083	1.461	1.160	11.563	329
-8	3.303	1.259	1.089	10.337	334
-9	2.755	1.090	0.948	7.421	333
-10	2.250	0.987	0.915	7.000	331
-11	1.850	0.827	0.687	6.796	330
-12	1.561	0.706	0.427	6.261	334
-13	1.244	0.541	0.067	3.442	348
-14	1.037	0.436	0.050	2.500	309
-15	0.819	0.414	0.045	2.329	315
-16	0.612	0.407	0.009	2.117	265
-17	0.488	0.367	0.027	1.740	224
-18	0.357	0.331	0.009	1.471	162
-19	0.241	0.242	0.005	1.228	132
-20	0.164	0.183	0.027	0.915	87
-21	0.135	0.173	0.012	0.816	56
-22	0.117	0.127	0.030	0.507	22
-23	0.069	0.064	0.022	0.307	20
-24	0.085	0.086	0.024	0.303	9
-25	0.060	0.036	0.030	0.099	3
-26	0.045	0.018	0.032	0.058	2
-27	0.031	0	0.031	0.031	2
Overall	4.424	5.119	0.005	34.683	5932

For clarity, negative values of *WageRank* are presented in Table 7.1, so -1 is the highest ranked player in a team. At higher wage rankings, we naturally see higher mean wages, but also a higher standard deviation of these wages. There are large variations in wage for a given wage rank, especially for the higher earners. The standard deviation of wages amongst the top earners of each team is over $5m. Although the absolute value of the standard deviations are lower at lower ranks, they are still large relative to the means. This variation can be exploited to determine whether the absolute level or relative ranking of wage affects performance, or possibly both. The frequencies in Table 7.1 become much lower at the lower ranks as rosters of this size are rare. Therefore effects found for the top end of the wage ranking may be considered more relevant as they apply to more players.[58]

Table 7.2 shows the correlation between the player-level wage variables and the team-level measures of wage dispersion. As discussed briefly in Section 4.4, there is a positive relationship between *Wage* and *WageRank* which is statistically significant at a 1% level. However, as the coefficient is not close to one, and due to the variation in wage for different wage

[58] The frequency of the wage ranks is occasionally above 330 due to ties in the rankings.

ranks displayed in Table 7.1, there is variation to be exploited here. Though the correlation coefficient is much smaller, there is also a statistically significant relationship between $Wage$ and $WageSD$ at a 1% level. This confirms that players with high wages tend to be in teams with higher wage dispersion. This indicates that controlling for wage is important in order to isolate the effects of $WageSD$. There is also a highly statistically significant, and this time negative, correlation between $WageRank$ and $WagePol$, suggesting that it may also be necessary to control for the rank of wage.

Table 7.2: Correlations Between Player-level Wage and Team-level Wage Dispersion

	Wage	WageRank	WageSD	WagePol
Wage	1.000			
WageRank	0.783***	1.000		
WageSD	0.173***	-0.012	1.000	
WagePol	-0.010	-0.135***	0.365***	1.000

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

For National Diversity, in addition to the team-level measures of fractionalisation and polarisation discussed in Section 3.2 and used in Chapter 6, we also consider some player-level variables. These include $LingDist1-4$ which are the measures of average linguistic distance between a player and his teammates considering 1-4 languages. The construction of these measures is discussed in Section 3.2. We also consider $PNatPer$ which captures the percentage of the team which shares the same nationality as a given player, i.e. for each player, what percentage of the team consists of him and his countrymen. The descriptive statistics of these variables are presented in Table 7.3.

Table 7.3: Descriptive Statistics of Player-level National Diversity Measures

	Mean	Std Dev.	Min.	Max.
PNatPer	0.671	0.312	0.036	1.000
LingDist1	0.228	0.258	0	0.962
LingDist2	0.158	0.200	0	0.962
LingDist3	0.147	0.189	0	0.962
LingDist4	0.147	0.189	0	0.962

On average, a player can expect 67.1% of their team to consist of their countrymen, though this is greatly skewed by the proportion of US players in the league. For players born outside of the USA, on average 6.12% of their team consists of their countrymen. The minimum $NatPer$ value of 3.6% applies to Marc Gasol of Spain and Beno Udrih of Slovenia who were both part of the large 28-man roster of Memphis Grizzlies in the 2015-16 season, where they were the lone players to represent their respective countries.

The maximum value of $LingDist1$ (0.962) applies to Cheick Diallo of Mali, Omer Asik of Turkey and Omri Casspi of Israel. All three were part of the 26-man roster of the New Orleans Pelicans in 2016-17 and each were their lone countrymen sharing no primary language or linguistic branches with the other seven nations in the team. The primary languages of Mali, Turkey and Israel are Bamanankan, Turkish and Hebrew, respectively. However, the secondary languages of Mali and Turkey are French and Kurdish, which both share Indo-

European branches with the languages of other countries in the team. So the values of *LingDist*2 drop for Cheick Diallo and Omer Asik (to 0.817 and 0.879 respectively). Meanwhile, the secondary language of Israel is Arabic which still shares no linguistic branches with languages of other countries in the team. Therefore the values of 0.962 for $LingDist2-4$ belong solely to Omri Casspi. These examples are provided to give the reader additional insight regarding the measures, but they are not extreme outliers. There are many other examples of players with average linguistic distances to their teammates above 0.9, for example. This is displayed in Figure 7.1 which shows a histogram of the *LingDist*2 variable across our entire sample of NBA players. Due to the similarities between the measures when considering two or more languages, we focus on just *LingDist*1 and *LingDist*2 for the remainder of the analysis.

Figure 7.1: Histogram of Average Linguistic Distance (Considering Two Languages) Between NBA Players and their Teammates

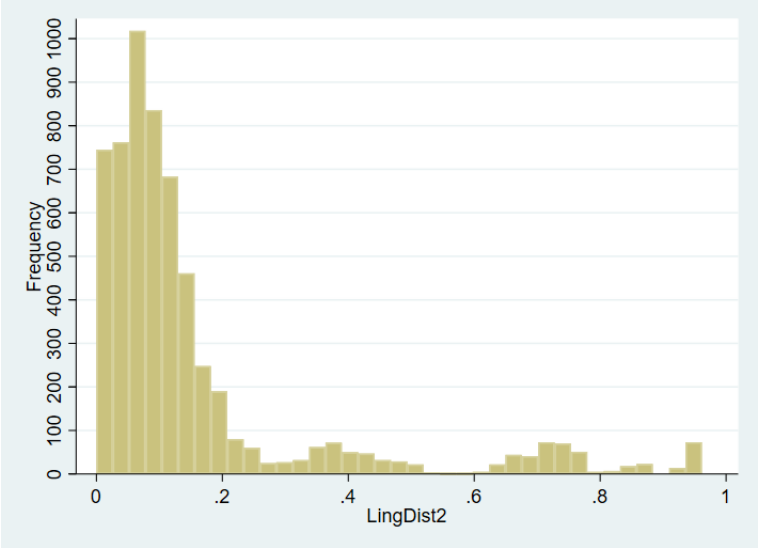


Table 7.4 shows the pairwise correlations between our player- and team-level measures of national diversity. All of these correlations are statistically significant at a 1% level and all are positive except between *PNatPer* and the other measures. But as *PNatPer* is, in a sense, an inverse measure of diversity, the negative correlations here are expected. The high statistical significance of these correlations is also expected, considering that the team-level variables are constructed using the proportions of the team made up from each national group and the linguistic distance between them. However, the correlation coefficients are sufficiently far from one to suggest collinearity is not a problem.

Table 7.4: Correlations Between Player- and Team-level National Diversity

	PNatPer	LingDist1	LingDist2	NatFrac	NatPol	NatFrac2	NatPol2
PNatPer	1.000						
LingDist1	-0.888***	1.000					
LingDist2	-0.807***	0.888***	1.000				
NatFrac	-0.488***	0.376***	0.346***	1.000			
NatPol	-0.455***	0.362***	0.338***	0.933***	1.000		
NatFrac2	-0.383***	0.388***	0.441***	0.783***	0.766***	1.000	
NatPol2	-0.278***	0.332***	0.412***	0.570***	0.648***	0.933***	1.000

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Considering both the player- and team-level variables, as well as interactions between them, should allow additional insight into the effects of wage dispersion and national diversity on performance. Analysing the effects on player performance should indicate the channels through which these issues influence team performance. The regression analysis conducted in the next sections will be analysed in such a way to identify these channels.

7.3 Wage Dispersion Regression Results

Table 7.5 shows the results of regressing individual player performance on the wage dispersion variables, without including any control variables at this stage. Specifically, we consider Wins Produced (WP) as our main measure of individual player performance. Column (1) shows a positive relationship between the wage of players and their performance which is statistically significant at a 1% level. We find that a logarithmic relationship fits best here.

Table 7.5: Player Performance and Wage Dispersion (No Controls)

	(1) WP	(2) WP	(3) WP	(4) WP	(5) WP	(6) WP
lnWage	0.604***		0.230***			
	(0.0483)		(0.0595)			
WageRank		0.162***	0.110***			
		(0.0132)	(0.0182)			
WageSD				0.148***		0.187***
				(0.0368)		(0.0395)
WagePol					-1.198**	-2.123***
					(0.491)	(0.534)
Constant	2.610***	4.627***	3.958***	2.329***	3.220***	2.614***
	(0.257)	(0.315)	(0.350)	(0.310)	(0.299)	(0.318)
log likelihood	-14829.4	-14817.1	-14812.1	-14956.2	-14964.6	-14948.7
Players	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Years	✓	✓	✓	✓	✓	✓
Position	✓	✓	✓	✓	✓	✓
Player Controls	✗	✗	✗	✗	✗	✗
Team Controls	✗	✗	✗	✗	✗	✗

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

In column (2), there is also a positive relationship between the rank of a player's wage and his performance, again statistically significant at a 1% level. This is perhaps unsurprising given the positive correlation between $Wage$ and $WageRank$. However, when including these variables simultaneously in column (3), both retain their positive signs and statistical significance at a 1% level. While the regression coefficients of both these variables reduce in size when including them simultaneously, it appears that each is capturing a distinct relationship with player performance. This suggests that players are motivated by both their wage level and their ranking within the wage structure of the team, which adds further evidence that these structures are tournament-like in nature. Columns (4)-(6) seem to confirm the evidence from our team-level regression in Chapter 5. $WageSD$ has positive effects on performance while $WagePol$ has negative effects, whether included separately or simultaneously.

Table 7.6 shows the same specifications as in Table 7.5, but with both the player- and team-level controls included. In columns (1)-(3), we see the same effects as described previously for *lnWage* and *WageRank*. Although the regression coefficients reduce in size, we still find evidence that both the wage level and the wage rank of players significantly affect player performance at a 1% significance level when included separately. When included simultaneously, the positive effect is statistically significant at a 1% level for *lnWage* and at a 5% level for *WageRank*. In columns (4)-(6), we see consistent results for *WageSD* with a positive relationship which is statistically significant at a 1% level; but this is not the case for *WagePol* which becomes statistically insignificant. This suggests that, when controlling for player- and team-level covariates, wage polarisation does not have a statistically significant effect on player performance. However, given the statistically significant negative relationship identified in Table 7.5 and in Chapter 5, one may suspect that the effects of wage polarisation are just not consistent across all players.

Table 7.6: Player Performance and Wage Dispersion (Player- and Team-level Controls)

	(1) WP	(2) WP	(3) WP	(4) WP	(5) WP	(6) WP
lnWage	0.354***		0.226***	0.359***	0.362***	0.360***
	(0.0384)		(0.0577)	(0.0383)	(0.0396)	(0.0395)
WageRank		0.0897***	0.0407**			
		(0.0113)	(0.0179)			
WageSD				0.126***		0.123***
				(0.0436)		(0.0448)
WagePol					0.539	0.117
					(0.489)	(0.501)
Constant	4.854***	4.824***	4.847***	7.091***	5.100***	7.099***
	(1.640)	(1.638)	(1.639)	(1.789)	(1.658)	(1.791)
log likelihood	-14534.0	-14536.7	-14531.7	-14529.4	-14533.5	-14529.4
Players	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Years	✓	✓	✓	✓	✓	✓
Position	✓	✓	✓	✓	✓	✓
Player Controls	✓	✓	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓	✓	✓

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

Interactions between player-level wage variables and team-level wage dispersion variables are explored in Table 7.7. In column (1), all four variables (*lnWage*, *WageRank*, *WageSD* and *WagePol*) are considered simultaneously but without any interactions. This tells the same story as the previous regressions, with positive coefficients which are statistically significant at a 1% level for *lnWage* and *WageSD*, and at a 5% level for *WageRank*; while *WagePol* remains statistically insignificant. In column (2), we include an interaction term between *WageRank* and *WageSD* to test whether the spread of wages has different effects on player performance depending on where the player ranks in wage structure. In line with Tournament Theory (Lazear & Rosen, 1981), we may expect that the 'winners' of the tournament prefer a larger spread. The positive coefficient on the interaction between *WageRank* and *WageSD*

confirms this hypothesis and is statistically significant at a 5% level. When including this interaction term, the coefficient on $WageRank$ becomes insignificant. This is also logical because if, hypothetically, a team had zero standard deviation of wage, there would be no difference in terms of wage between the different wage ranks. Interpreting the linear term of $WageSD$ is less meaningful, as it would apply to a wage rank of zero which is not possible (the highest rank possible is -1).

Table 7.7: Player Performance and Wage Dispersion with Interactions

	(1) WP	(2) WP	(3) WP	(4) WP
lnWage	0.245***	0.194***	0.451***	0.453***
	(0.0602)	(0.0595)	(0.0931)	(0.0922)
WageRank	0.0362**	-0.00807	0.0299*	-0.0294
	(0.0182)	(0.0279)	(0.0179)	(0.0288)
WageSD	0.119***	0.223***	0.126***	0.264***
	(0.0452)	(0.0750)	(0.0454)	(0.0787)
WagePol	0.0238	-0.0256	0.467	0.550
	(0.501)	(0.502)	(0.429)	(0.426)
WageRank × WageSD		0.0117**		0.0151***
		(0.00514)		(0.00547)
lnWage × WagePol			-0.726***	-0.967***
			(0.281)	(0.303)
Constant	6.969***	6.336***	6.935***	6.106***
	(1.797)	(1.841)	(1.795)	(1.842)
log likelihood	-14527.5	-14524.3	-14524.8	-14519.8
Players	✓	✓	✓	✓
Teams	✓	✓	✓	✓
Years	✓	✓	✓	✓
Position	✓	✓	✓	✓
Player Controls	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

Figure 7.2 explores the interaction between $WageRank$ and $WageSD$ further. These effects are plotted from the coefficients in column (2) of Table 7.7. The ranges of the variables used are similar to the sample ranges.[59] In Figure 7.2a, we see the effects of $WageSD$ for different values of $WageRank$. When the standard deviation is low, there is little difference in the performance of players of different wage ranks.[60] As the standard deviation of wage increases, we see the performance of players with high wage ranks increase, while the performance of those with low wage ranks decreases. Figure 7.2b shows the effects of $WageRank$ for different values of $WageSD$ which tells the same story. While there is some evidence that players who are low in the wage ranking of their team may perform better when the standard deviation of wage is low (holding $lnWage$ constant), the performance of players with a high wage rank are higher when the standard deviation of wage is higher. From our investigation in Chapter

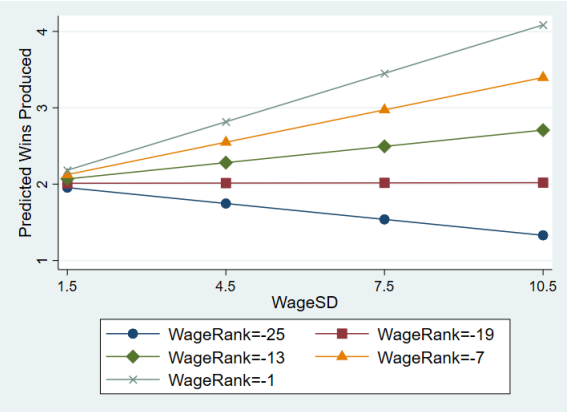
[59] We focus on the top 25 wage ranks due to the low frequency of wage ranks lower than this.

[60] Note that this regression specification holds $lnWage$ constant. When excluding $lnWage$, we see some differences between wage ranks when standard deviation is low, but these are not large.

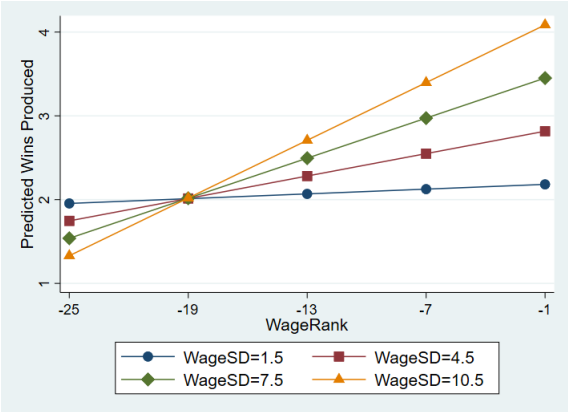
5, we know that the overall effect on team performance is positive. This player-level analysis suggests that this is driven by the top earners in the team performing better, not by increased performance throughout the team.

Figure 7.2: Wage Rank and the Standard Deviation of Wage

(a) Effects of the Standard Deviation of Wage for Different Wage Ranks



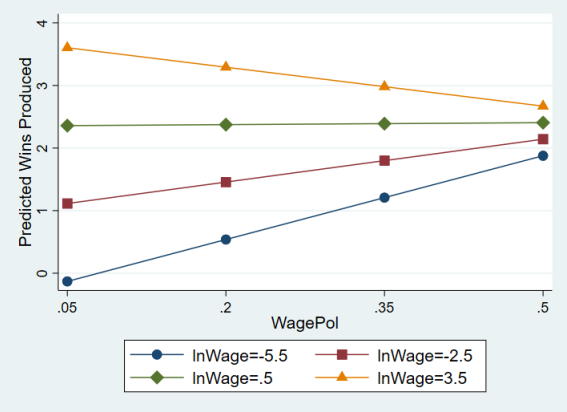
(b) Effects of Wage Rank for Different Standard Deviations of Wage



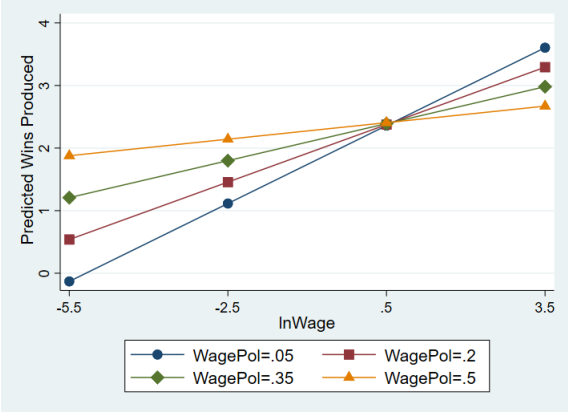
In column (3) of Table 7.7, we include an interaction between $lnWage$ and $WagePol$ to test whether the effects of wage polarisation depend on the player's wage. The negative coefficient on the interaction term suggests that the effect of wage polarisation affects players on higher wages more negatively and this is statistically significant at a 1% level. In this specification, the coefficient on $lnWage$ remains positive and statistically significant at a 1% level. The coefficient also becomes larger suggesting that the positive relationship between wage and performance would be maximised if there was hypothetically no wage polarisation, and performance is higher when wage polarisation is low. The coefficient on $WagePol$ is positive but statistically insignificant, though the statistical significance of this coefficient is not particularly informative anyway. It is the effect of wage polarisation when $lnWage$ is zero, which is not to say that wage itself is zero. $Wage$ is measured in millions of dollars, so $lnWage$ will equal zero when a player is paid \$1m. Therefore, wage polarisation may have positive effects for players who earn less than \$1m and who therefore have a negative value for $lnWage$. This is an unexpected result and is explored further in Figure 7.3.

Figure 7.3: The Natural Logarithm of Wage and Wage Polarisation

(a) Effects of Wage Polarisation for Different Wages



(b) Effects of Wage for Different Wage Polarisations



Again, the ranges of the variables used to construct Figure 7.3 are similar to the sample ranges. In Figure 7.3a, we see that players on high wages significantly outperform those on lower wages when wage polarisation is low. As wage polarisation increases, the performance of the high earners decreases, while the performance of low earners (those with negative value of $lnWage$) increases. This second part was unexpected, but can be explained by the influence of polarisation on cooperation and conflict. If teams with high wage polarisation are less willing to cooperate, then this could lead to individualistic play by players. Evidence from Franck and Nüesch (2011) links high wage dispersion to more individualistic play, but here we suggest it might specifically be linked to high wage polarisation. Those on low wages may be unwilling to facilitate the performance of high earners, instead choosing to maximise their own performance rather than that of the team. When wage polarisation is low, there is more cohesion within the team and players are willing to sacrifice their individual performance for the good of the team. Figure 7.3b shows the effects of $lnWage$ for different values of $WagePol$ which can be explained through a similar line of reasoning. Low earners appear to perform best when wage polarisation is high, while the top earners perform best when wage polarisation is low. And, as revealed in Chapter 5, we find evidence that wage polarisation has negative effects on overall team performance.

Lastly, column (4) of Table 7.7 includes both of these interaction terms simultaneously. This reveals that both of these interaction effects are robust to the inclusion of the other. We do trial interacting each of the player-level wage variables with the other team-level wage dispersion variables, but do not find significant results for these. Therefore, Tournament Theory (Lazear & Rosen, 1981) seems to be linked to the wage ranking of players and the standard deviation of wage within teams, while Cohesion Theory (Levine, 1991) is linked to the wage of players and wage polarisation within the team.

7.3.1 Wage Dispersion Robustness Checks

Table 7.8 tests the robustness of our main result to different measures of individual performance. Columns (1)-(3) show the robustness of our findings across the different measures of overall contribution of players. The only notable difference among the results of these three columns is that the positive coefficient on $WagePol$ is statistically significant at a 5% level in column (2) for WS. However, this is the predicted effect when $lnWage$ equals zero, not when $Wage$ equals zero. The former occurs at around the 26^{th} percentile of the data. Given that the highly statistically significant negative coefficient on the interaction term of $lnWage$ and $WagePol$ is larger than the coefficient on $WagePol$, then all three of these measures agree that wage polarisation has positive effects on low earners and negative effects on high earners. There is just some disagreement about when these changes become statistically different from zero. The other coefficients are consistent in their signs and statistical significance.

Columns (4)-(6) suggest that our results are not robust to considering the efficiency of player performance. There appears to be some consistency in the positive effect of $WageSD$, though this is statistically insignificant when considering PER in column (5). There is also some evidence of the positive effects for the interaction between $WageRank$ and $WageSD$ on the efficiency measures, though the evidence here is weak. The coefficients on the other variables either lack statistical significance, change sign, or both. As discussed in Section 4.3.3 and displayed in Figure 4.2 (and also suggested in Appendix 4.E), the efficiency measures are susceptible to award very high values for players which play infrequently, perhaps in 'dead

Table 7.8: Wage Dispersion Robustness Check (Performance Measures)

	(1) WP	(2) WS	(3) VORP	(4) WP48	(5) PER	(6) BPM
lnWage	0.453***	0.534***	0.130***	-0.0171*	-0.166	0.551**
	(0.0922)	(0.0690)	(0.0323)	(0.00963)	(0.338)	(0.217)
WageRank	-0.0294	-0.00605	-0.00955	0.00214	0.0969	0.0427
	(0.0288)	(0.0184)	(0.00888)	(0.00195)	(0.0639)	(0.0481)
WageSD	0.264***	0.197***	0.0840***	0.0104***	0.132	0.254***
	(0.0787)	(0.0510)	(0.0273)	(0.00343)	(0.118)	(0.0872)
WagePol	0.550	0.713**	0.219	-0.110***	-2.472*	-0.704
	(0.426)	(0.315)	(0.138)	(0.0413)	(1.480)	(1.027)
WageRank × WageSD	0.0151***	0.0100***	0.00500***	0.000605*	0.00977	0.000822
	(0.00547)	(0.00343)	(0.00174)	(0.000351)	(0.0114)	(0.00837)
lnWage × WagePol	-0.967***	-1.036***	-0.395***	0.0422*	1.840**	0.432
	(0.303)	(0.220)	(0.111)	(0.0243)	(0.932)	(0.586)
Constant	6.106***	5.603***	1.468***	0.0663	13.58***	2.570
	(1.842)	(1.281)	(0.558)	(0.0996)	(3.035)	(2.010)
log likelihood	-14519.8	-12357.7	-7610.5	2332.7	-18530.4	-16305.8
Players	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Years	✓	✓	✓	✓	✓	✓
Position	✓	✓	✓	✓	✓	✓
Player Controls	✓	✓	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓	✓	✓

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

rubber' matches at the end of the season when playoff places have already been decided. To test whether these observations are affecting our results and possibly leading them to be unrepresentative, we consider a reduced sample of only those players which play over 1,000 minutes in a season. This figure is between the median number of minutes played of 894 and the mean of 1,083. It also removes the apparent outliers from the scatterplots in Figure 4.2. The regression results for the efficiency measures when considering this reduced sample are presented in Table 7.9.

However, Table 7.9 shows that there is little change in the results for the efficiency measures when considering the reduced sample compared to the entire sample. The coefficients on $lnWage$ and the interaction between $lnWage$ and $WagePol$ in columns (1) and (2) for $WP48$ and PER, which were previously in contradiction to the results found for the overall contribution measures, become statistically insignificant. But otherwise, we do not find robust evidence for the effects of wage dispersion on the efficiency of player performance. Columns (4)-(6) show that this is also the case when excluding the interactions terms. There is some evidence that performance is improved by a higher standard deviation of wage, in agreement with our results for the overall contribution measures, but we do not find statistically significant evidence across all of the efficiency measures, particularly for PER. The possible explanations for this are either that our findings relating to wage dispersion do not apply to the efficiency of players, just their overall contribution; or that efficiency is not a good measure of performance. Given that Appendix 4.F suggests that teams with a small 'core team' perform better, perhaps teams should prefer players who perform consistently well and

often, and therefore score highly in terms of overall contribution.

Table 7.9: Wage Dispersion Robustness Check (Reduced Sample)

	(1) WP48	(2) PER	(3) BPM	(4) WP48	(5) PER	(6) BPM
lnWage	-0.0111	0.305	0.150	-0.00873*	0.345*	0.159
	(0.00698)	(0.266)	(0.190)	(0.00460)	(0.178)	(0.116)
WageRank	-0.00144	-0.0352	-0.0134	0.000245	-0.0184	-0.0154
	(0.00129)	(0.0493)	(0.0342)	(0.000891)	(0.0319)	(0.0225)
WageSD	0.00737***	0.00204	0.128*	0.00413***	-0.0313	0.131***
	(0.00230)	(0.0879)	(0.0665)	(0.00159)	(0.0583)	(0.0450)
WagePol	-0.0258	0.195	-0.456	-0.0310	0.233	-0.386
	(0.0344)	(1.315)	(0.902)	(0.0198)	(0.764)	(0.546)
WageRank × WageSD	0.000490**	0.00505	-0.000450			
	(0.000226)	(0.00898)	(0.00629)			
lnWage × WagePol	-0.00428	0.0137	0.0459			
	(0.0193)	(0.787)	(0.582)			
Constant	0.185***	15.55***	2.892*	0.204***	15.72***	2.863*
	(0.0640)	(2.219)	(1.580)	(0.0620)	(2.163)	(1.522)
log likelihood	3411.9	-6605.0	-5586.4	3409.6	-6605.2	-5586.4
Players	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Years	✓	✓	✓	✓	✓	✓
Position	✓	✓	✓	✓	✓	✓
Player Controls	✓	✓	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓	✓	✓

N = 2,778
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

Table 7.10 shows the robustness of our results regarding the effects of wage dispersion on the performance of players, in terms of WP, to relaxing some of the elements controlled for. In columns (1) and (2), we individually exclude the team-level and then player-level controls. When excluding the player controls in column (2), the signs of the coefficients on $WageRank$ and $WagePol$ change sign, but given that both of these are not statistically different from zero, this does not alter our main conclusions. Despite some changes in magnitude, the remaining coefficients show robustness to the removal of these controls.

In columns (3) and (4) of Table 7.10, we exclude the Position and then the Year dummies. These exclusions make little change from our main results (last given in column (1) of Table 7.8), again showing the robustness of our findings. In columns (5) and (6), we do not account for the Team level, and then the Player level, in our multilevel model. When we do not control for team-level heterogeneity (beyond the team-level control variables) in column (5), the coefficients are of a slightly lower magnitude than in the full specification, but the results are generally robust. When player-level heterogeneity is not controlled for (beyond the player-level control variables), the coefficients are of a much larger magnitude, both in a positive and negative direction. It is therefore important to use multilevel modelling to control for this player-level heterogeneity so as not to overstate the predicted effects. This is explored further in Table 7.11, where we demonstrate why multilevel modelling is the most efficient way to capture this heterogeneity.

Table 7.10: Wage Dispersion Robustness Check (Alternative Specifications)

	(1) WP	(2) WP	(3) WP	(4) WP	(5) WP	(6) WP
lnWage	0.396***	0.554***	0.457***	0.485***	0.433***	0.813***
	(0.0899)	(0.101)	(0.0928)	(0.0934)	(0.0930)	(0.116)
WageRank	-0.0191	0.0287	-0.0292	-0.0373	-0.0286	-0.0671*
	(0.0288)	(0.0299)	(0.0290)	(0.0286)	(0.0291)	(0.0363)
WageSD	0.314***	0.268***	0.258***	0.264***	0.245***	0.523***
	(0.0762)	(0.0851)	(0.0786)	(0.0783)	(0.0795)	(0.0905)
WagePol	-0.0281	0.537	0.553	0.669	0.492	1.675***
	(0.402)	(0.448)	(0.428)	(0.420)	(0.430)	(0.510)
WageRank × WageSD	0.0156***	0.0141**	0.0149***	0.0158***	0.0140**	0.0421***
	(0.00544)	(0.00598)	(0.00546)	(0.00543)	(0.00550)	(0.00622)
lnWage × WagePol	-0.977***	-1.190***	-0.937***	-0.953***	-0.958***	-2.032***
	(0.304)	(0.332)	(0.303)	(0.305)	(0.308)	(0.394)
Constant	1.804***	7.942***	5.526***	8.044***	5.802***	4.972***
	(0.458)	(1.960)	(1.833)	(1.454)	(1.889)	(1.844)
log likelihood	-14539.8	-14754.0	-14531.8	-14530.2	-14556.6	-15483.6
Players	✓	✓	✓	✓	✓	✗
Teams	✓	✓	✓	✓	✗	✓
Years	✓	✓	✓	✗	✓	✓
Position	✓	✓	✗	✓	✓	✓
Player Controls	✓	✗	✓	✓	✓	✓
Team Controls	✗	✓	✓	✓	✓	✓

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

In Table 7.11, we compare the results obtained using Multilevel Modelling (MLM) to those of 'complete pooling' using Fixed Effects (FE) and 'no pooling' using Random Effects (RE). In columns (1)-(3), the sample is reduced from 5,932 to 5,213, to allow for a panel data structure. When players appear for multiple teams within a given season, we keep only the observation for the team for which he played the most minutes. This results in 719 observations being dropped, but allows the data to be set as panel data, with *Player* denoting the player identifier and $Year$ denoting the seasons (time identifier). The variables are then recalculated for the reduced sample.

In column (1), MLM is run as before and we see consistent results to considering the full sample. The only small change of note is that the interaction between $WageRank$ and $WageSD$ drops in significance from a 1% level to a 5% level. When using the 'complete pooling' FE approach in column (2), we see qualitatively much the same results in terms of the signs of the coefficients, but with much lower significance levels. Only the linear term of $WageSD$ remains statistically significant but now at a 10% level.[61] Therefore, this model offers some evidence of the effects of Tournament Theory through $WageSD$, but not for the other effects. In column (3), we use the 'no pooling' RE approach and we see results consistent with MLM in column (1) with larger coefficients and higher significance levels. A Hausman test using the Mundlak (1978) approach is invalid for such an unbalanced panel, so the clustering of the

[61] When excluding the insignificant interaction term between $WageSD$ and $WageRank$, the coefficient on $WageSD$ decreases a little (to 0.154) but is then statistically significant at a 1% level.

standard errors is removed to allow for a standard Hausman test to be conducted. This test return a Wald Statistic of 105.2 (p-value=0.000) suggesting that the RE approach may be biased. However, the MLM approach controls for the player- (and team-) level heterogeneity without the loss of efficiency of the FE approach. Instead of ignoring the between variation, both the within and between variation is exploited through partial pooling.

Table 7.11: Wage Dispersion Robustness Check (Alternative Models)

	(1) WP	(2) WP	(3) WP	(4) WP	(5) WP	(6) WP
lnWage	0.482***	0.127	0.626***	0.883***	0.267	1.032***
	(0.108)	(0.125)	(0.119)	(0.249)	(0.277)	(0.258)
WageRank	-0.0152	-0.0346	-0.0175	-0.0966**	-0.0852	-0.103**
	(0.0319)	(0.0386)	(0.0326)	(0.0471)	(0.0540)	(0.0469)
WageSD	0.266***	0.172*	0.307***	0.212**	0.109	0.176*
	(0.0831)	(0.0972)	(0.0875)	(0.0938)	(0.108)	(0.0927)
WagePol	0.417	-0.660	0.983*	3.914**	3.836*	5.117***
	(0.532)	(0.694)	(0.573)	(1.797)	(2.238)	(1.835)
WageRank × WageSD	0.0124**	0.00601	0.0182***	0.0132*	0.0112	0.0127*
	(0.00563)	(0.00664)	(0.00590)	(0.00741)	(0.00877)	(0.00757)
lnWage × WagePol	-0.943***	-0.571	-1.170***	-2.663**	-2.704**	-2.610**
	(0.363)	(0.401)	(0.405)	(1.103)	(1.332)	(1.160)
Constant	7.064***	6.341**	8.125***	13.21***	10.11***	13.79***
	(2.085)	(2.559)	(1.732)	(2.737)	(3.133)	(2.202)
N	5213	5213	5213	3546	3546	3546
log likelihood	-13017.8	-11534.8	-	-9164.0	-8048.4	-
R^2	-	0.186	0.242	-	0.202	0.179
Players	✓	✓	✗	✓	✓	✗
Teams	✓	✓	✗	✓	✓	✗
Years	✓	✓	✗	✓	✓	✗
Position	✓	✓	✗	✓	✓	✗
Player Controls	✓	✓	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓	✓	✓
Core Team	✗	✗	✗	✓	✓	✓
Model	MLM	FE	RE	MLM	FE	RE

Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

In columns (4)-(6), we further restrict the sample to only those considered to be part of the 'core team'[62]. This reduces the size of the dataset by a further 1,667 to 3,546. We then use the MLM approach in column (4), FE in (5) and RE in (6), and see the same general pattern as described for columns (1)-(3) when comparing the models. We also see consistency in the signs across the models. However, when comparing the MLM models in columns (1) and (4), for example, although the coefficients are generally larger for the core team regressions, the significance levels tend to be lower. In general, there is less wage dispersion in the entire team than in the core team, so this explains the larger coefficients. However, the lower significance levels suggest the effects may not be as strong for wage dispersion within the core team.

Overall, we find evidence that both a player's wage and their wage ranking have positive relationships with individual player performance. Furthermore, a higher standard deviation

[62] Defined by Simmons and Berri (2011) as those which played at least 20 games and at least 12 minutes per game.

of wage strengthens the relationship between wage rank and player performance. Those at the top end of the wage ranking perform better when the standard deviation of wage is high, while those lower in the wage ranking perform worse. When the standard deviation of wage is low, there is evidence that the wage rank of a player has little relationship with their performance. We also find evidence of differing effects of wage polarisation on player performance, but we find this to be more closely linked to a player's wage, rather than wage rank. While a positive (logarithmic) relationship between wage and player performance exists even when wage polarisation is high, the relationship is stronger when wage polarisation is low. Therefore, high earners perform better in teams which are less polarised in terms of wage. We find these results to be generally robust, but apply to the performance of players in terms of overall contribution more than efficiency.

7.4 National Diversity Regression Results

To begin our investigation into the effects of national diversity on individual player performance, we first regress individual player-level and then team-level measures of national diversity separately on Wins Produced (WP), without including control variables. The results are presented in Table 7.12.

Table 7.12: Player Performance and National Diversity (No Controls)

	(1) WP	(2) WP	(3) WP	(4) WP	(5) WP	(6) WP
PNatPer	-0.223					
	(0.246)					
LingDist2		-0.115				
		(0.354)				
NatFrac			0.521*		0.847	
			(0.303)		(0.891)	
NatPol				0.391	-0.316	
				(0.275)	(0.806)	
NatFrac2						1.901
						(1.492)
NatPol2						-1.494
						(1.202)
Constant	3.097***	2.993***	2.801***	2.782***	2.847***	2.998***
	(0.304)	(0.287)	(0.290)	(0.302)	(0.310)	(0.284)
log likelihood	-14966.8	-14967.2	-14965.5	-14966.0	-14965.4	-14966.1
Players	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Years	✓	✓	✓	✓	✓	✓
Position	✓	✓	✓	✓	✓	✓
Player Controls	✗	✗	✗	✗	✗	✗
Team Controls	✗	✗	✗	✗	✗	✗

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

The results for the player-level variables are shown in columns (1) and (2) of Table 7.12. Column (1) suggests that players from smaller national groups perform better, while column (2) suggests that players with a larger average linguistic distance to their teammates perform worse. This is, in a sense, contradictory, as players from smaller national groups tend to have

larger linguistic distances to their teammates, shown by the reasonably high and statistically significant negative correlation between *PNatPer* and *LingDist*2, displayed in Table 7.4. However, neither coefficient is statistically different from zero.

Columns (3)-(6) show the results for the team-level variables. The results show a similar pattern to the team-level regressions in Chapter 6. When included separately, in columns (3) and (4), both *NatFrac* and *NatPol* show positive coefficients, but when included simultaneously in column (5), there is a positive coefficient on *NatFrac* and a negative coefficient on *NatPol*. Column (6) shows a similar result for *NatFrac*2 and *NatPol*2 (i.e. when accounting for linguistic distance between national groups considering two primary languages per country). However, the only statistically significant coefficient is the positive one for *NatFrac* when included separately in column (3).

Table 7.13: Player Performance and National Diversity (Player- and Team-level Controls)

	(1) WP	(2) WP	(3) WP	(4) WP	(5) WP	(6) WP
PNatPer	0.112					
	(0.225)					
LingDist2		-0.433				
		(0.325)				
NatFrac			0.563**		1.157	
			(0.286)		(0.818)	
NatPol				0.392	-0.564	
				(0.261)	(0.744)	
NatFrac2						2.416*
						(1.376)
NatPol2						-1.716
						(1.109)
Constant	4.815***	4.998***	4.468***	4.524***	4.537***	4.645***
	(1.648)	(1.637)	(1.630)	(1.637)	(1.639)	(1.630)
log likelihood	-14533.9	-14533.2	-14531.7	-14532.6	-14531.4	-14532.1
Players	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Years	✓	✓	✓	✓	✓	✓
Position	✓	✓	✓	✓	✓	✓
Player Controls	✓	✓	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓	✓	✓

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

Player- and team-level controls are included in Table 7.13. The results in columns (1) and (2) suggest that players perform better when they are part of a larger national group and when the average linguistic distance to their teammates is smaller. These results are no longer contradictory to one another and suggests that players may perform better in homogeneous teams. However, firstly, these results are not statistically significant. And secondly, the coefficient on *NatFrac* in column (3) remains positive when including the controls and increases in significance to a 5% level, suggesting that fractionalised teams perform better. In columns (4) and (5), we see the same pattern of results as before the controls were included. There is still a positive coefficient on *NatPol* when it is included separately in column (4), but this

again becomes negative when including $NatFrac$ and $NatPol$ simultaneously in column (5), though these coefficients are statistically insignificant. Column (6) tells a similar story when considering linguistic distance with $NatFrac2$ and $NatPol2$, with just the positive coefficient on $NatFrac2$ statistically significant at a 10% level. Overall, we appear to have a trade-off between the effects of the player- and team-level variables, though the results are still lacking in statistical significance to make meaningful inferences.

Table 7.14: Player Performance and National Diversity (Team- and Player-level Measures)

	(1) WP	(2) WP	(3) WP	(4) WP	(5) WP	(6) WP	(7) WP
PNatPer	0.520**	0.408	0.516**	0.355			-0.0353
	(0.261)	(0.263)	(0.262)	(0.250)			(0.375)
LingDist2					-0.839**	-0.758**	-0.800
					(0.352)	(0.368)	(0.551)
NatFrac	0.911***		1.477*		1.361*		
	(0.333)		(0.817)		(0.820)		
NatPol		0.620**	-0.541		-0.507		
		(0.306)	(0.746)		(0.746)		
NatFrac2				3.135**		2.890**	2.845**
				(1.415)		(1.383)	(1.424)
NatPol2				-2.056*		-1.696	-1.661
				(1.107)		(1.110)	(1.133)
Constant	4.047**	4.189**	4.116**	4.391***	4.628***	4.718***	4.747***
	(1.650)	(1.663)	(1.660)	(1.643)	(1.638)	(1.629)	(1.658)
log likelihood	-14529.7	-14531.3	-14529.4	-14531.1	-14528.7	-14530.2	-14530.2
Players	✓	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓	✓
Years	✓	✓	✓	✓	✓	✓	✓
Position	✓	✓	✓	✓	✓	✓	✓
Player Controls	✓	✓	✓	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓	✓	✓	✓

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

In Table 7.14, we consider player- and team-level measures of national diversity simultaneously. Columns (1)-(3) consider the size of player's national group alongside the team-level measures of fractionalisation and polarisation, separately then simultaneously. Column (1) gives further evidence to the trade-off previously alluded to. The coefficient on $PNatPer$ is positive and statistically significant at a 5% level, while $NatFrac$ is also positive and significant at a 1% level. This suggests that, for a given level of national fractionalisation in the team, players perform better when they are part of a larger national group. Also, for a given size of their national group within the team, players benefit when the remainder of the team is more nationally fractionalised. It appears, therefore, that players benefit most when their teammates provide national diversity, from which they can benefit from the positive spillover effects. It should be noted that the coefficient on $NatFrac$ is significantly larger than that on $PNatPer$. So if, for example, we consider a homogeneous versus a completely fractionalised team, then the homogeneous team would benefit from $PNatPer$ being equal to one, while the nationally fractionalised team would have $NatFrac$ close to one. The size of these coefficients suggest that performance will be higher in the nationally fractionalised team, which is in line

with the results of Chapter 6. Column (2) shows a similar relationship when including $NatPol$ in the place of $NatFrac$. However, this is due to the positive correlation between $NatFrac$ and $NatPol$. Column (3) shows that, when they are included simultaneously, $NatFrac$ has a positive relationship with player performance at a 10% significance level, while $NatPol$ shows a negative relationship with player performance, though it is not statistically significant.

Column (4) of Table 7.14 includes $NatFrac2$ and $NatPol2$ alongside $PNatPer$. The positive coefficient on $NatFrac2$ is statistically significant at a 5% level, while the negative coefficient on $NatPol2$ becomes statistically significant at a 10% level. However, the positive coefficient on $PNatPer$ becomes statistically insignificant. As this occurs when including the team-level variables which account for linguistic distance, the positive coefficient on $PNatPer$ may be caused by its correlation with linguistic distance. In columns (5) and (6), we repeat the specifications of columns (3) and (4), but with $LingDist2$ replacing $PNatPer$. The results are similar, but with $LingDist2$ remaining statistically significant at a 5% level when included alongside $NatFrac2$ and $NatPol2$. Though the coefficient on $NatPol2$ in column (6) is insignificant, this is only marginal, with a p-value of 0.127. Lastly, in column (7), we consider $PNatPer$ and $LingDist2$ simultaneously to check whether they are capturing distinct or similar effects. As both coefficients becomes insignificant in this final specification, we conclude they are capturing similar effects and the positive effects of $PNatPer$ are driven by its negative correlation with $LingDist2$. We also trial some interaction terms between the player- and team-level measures of national diversity but do not find evidence of such relationships.

Overall, we find evidence that players perform better when there is a smaller linguistic distance between them and their teammates, but also when there are larger average linguistic distances ($NatFrac2$) within the team. Though this sounds, in a sense, contradictory, it can also be considered as a trade-off. The team overall benefits from being nationally fractionalised, though each player individually could maximise their own performance by being in a large national group within a fractionalised team. If players are focussed on improving just their individual performance, then they would prefer other players in the team to provide national diversity, while remaining part of a large national group themselves, and still benefiting from the positive spillover effects that national diversity brings. There is also weak evidence of a negative relationship between national polarisation which accounts for linguistic distance (considering two primary languages) between national groups. Despite this relationship often showing marginal statistically insignificant results, we chose to retain this variable in our main specification due to its effects at the team-level, as found in Chapter 6.

7.4.1 National Diversity Robustness Checks

Table 7.15 tests the robustness of our main results relating to the effects of national diversity on player performance across different individual player performance measures. Column (1) shows the previously reported results for the effects on WP, columns (2) and (3) show the effects on the other measures of overall contribution (WS and $VORP$), while columns (4)-(6) show the effects on the efficiency performance measures ($WP48$, PER and BPM). The negative coefficients on $LingDist2$ and $NatPol2$, and the positive coefficient on $NatFrac2$ are consistent across all measures, though the levels of statistical significance are higher for the overall contribution measures. While statistically insignificant for WP, the negative coefficient of $NatPol2$ is statistically significant at a 10% level for WS and at a 5% level for $VORP$. Though the signs of the coefficients are consistent for the efficiency measures of performance, only $NatFrac2$ is statistically significant at a 10% level for $WP48$ and at a 1% level

Table 7.15: National Diversity Robustness Check (Performance Measures)

	(1) WP	(2) WS	(3) VORP	(4) WP48	(5) PER	(6) BPM
LingDist2	-0.758**	-0.419*	-0.224**	-0.0333	-1.212	-0.977
	(0.368)	(0.242)	(0.0991)	(0.0256)	(0.876)	(0.639)
NatFrac2	2.890**	2.298**	1.395***	0.132*	2.568	4.529***
	(1.383)	(0.911)	(0.420)	(0.0690)	(2.461)	(1.686)
NatPol2	-1.696	-1.405*	-0.826**	-0.0865	-1.179	-2.138
	(1.110)	(0.739)	(0.346)	(0.0564)	(1.980)	(1.334)
Constant	4.718***	4.248***	1.013**	0.0446	13.46***	-1.720
	(1.629)	(1.103)	(0.479)	(0.0855)	(2.546)	(1.831)
log likelihood	-14530.2	-12378.5	-7621.0	2317.4	-18540.7	-16313.0
Players	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Years	✓	✓	✓	✓	✓	✓
Position	✓	✓	✓	✓	✓	✓
Player Controls	✓	✓	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓	✓	✓

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

for BPM. This suggests that the positive relationship between average linguistic distance within the team and individual player performance is our most consistent and statistically robust result. However, in Table 7.16, we again check whether the apparent insignificance of the other coefficients, when considering the efficiency measures, may be due to the extreme measurements of players which play infrequently.

Table 7.16 considers a reduced sample of only those players which played at least 1,000 minutes, reducing the sample size from 5,932 to 2,778 distinct player-team-year observations. We consider the three efficiency performance measures twice, with $PNatPer$ alongside $NatFrac2$ and $NatPol2$ in columns (1)-(3), and then with $LingDist2$ replacing $PNatPer$ in columns (4)-(6). In general, when considering the reduced sample, we see coefficients of a similar magnitude, but levels of statistical significance closer to those of the overall contribution measures. The statistical significance of the relationships is most consistently strong for $WP48$ (even stronger than for WP in column (1) of Table 7.15), while the relationship between the team-level measures and BPM are also highly statistically significant. The relationships between national diversity and PER are weakest, though the positive effects of $NatFrac2$ are still statistically significant at a 10% level.

The final robustness check of this section considers our choice of Multilevel Modelling (MLM), comparing it to the 'complete pooling' Fixed Effects (FE) approach which includes dummies for players, teams, years and positions; and the 'no pooling' Random Effects (RE) approach which excludes all of these dummies. In columns (1)-(3), repeated player observations within years are removed, retaining only the observation for which the player played the most minutes. This allows for a panel data structure for the FE and RE approaches. The coefficients on the team-level variables for the MLM model in column (1) are close in magnitude to the FE estimates in column (2), suggesting that these results are driven by the variation within players over time. However, the coefficient for $LingDist2$ is positive using the FE approach,

Table 7.16: National Diversity Robustness Check (Reduced Sample)

	(1) WP48	(2) PER	(3) BPM	(4) WP48	(5) PER	(6) BPM
PNatPer	0.0179*	0.240	-0.0474			
	(0.0100)	(0.388)	(0.250)			
LingDist2				-0.0384**	-0.623	-0.543
				(0.0175)	(0.616)	(0.377)
NatFrac2	0.140***	3.195*	3.985***	0.130***	3.135*	4.413***
	(0.0514)	(1.879)	(1.282)	(0.0498)	(1.818)	(1.248)
NatPol2	-0.0886**	-2.148	-2.658***	-0.0730*	-1.945	-2.697***
	(0.0397)	(1.411)	(1.009)	(0.0393)	(1.383)	(1.003)
Constant	0.130**	15.36***	0.122	0.147**	15.60***	0.149
	(0.0581)	(2.013)	(1.403)	(0.0574)	(1.986)	(1.371)
log likelihood	3404.9	-6608.3	-5586.4	3406.0	-6608.0	-5585.5
Players	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Years	✓	✓	✓	✓	✓	✓
Position	✓	✓	✓	✓	✓	✓
Player Controls	✓	✓	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓	✓	✓

N = 2,778
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

as opposed to negative for MLM and RE. This suggests it is driven by the variation between players, which is ignored by the FE approach. This is one of the advantages of MLM, as it exploits both within and between variation through partial pooling.

Table 7.17: National Diveristy Robustness Check (Alternative Models)

	(1) WP	(2) WP	(3) WP	(4) WP	(5) WP	(6) WP
LingDist2	-0.807**	0.163	-0.486	-0.946	0.584	-0.257
	(0.393)	(2.092)	(0.405)	(0.596)	(2.195)	(0.611)
NatFrac2	3.486*	3.073	2.473	7.857***	5.975*	5.577*
	(1.810)	(2.458)	(1.846)	(2.904)	(3.585)	(2.918)
NatPol2	-1.926	-2.135	-1.317	-4.251**	-4.203*	-3.005
	(1.383)	(1.641)	(1.410)	(2.094)	(2.346)	(2.082)
Constant	4.808***	4.727**	6.574***	11.45***	9.901***	12.57***
	(1.824)	(2.244)	(1.316)	(2.493)	(2.877)	(1.774)
N	5213	5213	5213	3546	3546	3546
log likelihood	-13026.1	-11537.5	-	-9166.4	-8050.2	-
R^2	-	0.185	0.307	-	0.201	0.319
Players	✓	✓	✗	✓	✓	✗
Teams	✓	✓	✗	✓	✓	✗
Years	✓	✓	✗	✓	✓	✗
Position	✓	✓	✗	✓	✓	✗
Player Controls	✓	✓	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓	✓	✓
Core Team	✗	✗	✗	✓	✓	✓
Model	MLM	FE	RE	MLM	FE	RE

Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

Columns (4)-(6) of Table 7.17 show the same specification but when only the 'core team'[63] is considered. The results here suggest a lower level of statistical significance for the effects of $LingDist2$ and higher significance for the team-level variables, $NatFrac2$ and $NatPol2$.

Overall, we find our results to be robust to different measures of player performance in terms of overall contribution and arguably even more significant in terms of efficiency when considering a reduced sample. When excluding players which play less than 1,000 minutes in a season, we see particularly significant results for two of our three measures of player efficiency. We also find that the positive effects of national fractionalisation and negative effects of national polarisation dominate the negative effects of a high linguistic distance at the player-level when considering the 'core team'.

7.5 Wage Dispersion, National Diversity and Individual Player Performance

In this section, we consider a combined model which incorporates our findings of the relationships between both wage dispersion and national diversity, and the performance of individual players within teams. This is to test whether not controlling for wage dispersion in our previous results is biasing our estimates of the effects of national diversity, and vice versa. Table 7.18 shows the results of the combined specification for the different measures of individual player performance.

Columns (1)-(3) of Table 7.18 show that our findings for both wage dispersion and national diversity are robust to the inclusion of the other when considering the overall contribution of players. The relationships between wage dispersion and player performance retain their high levels of statistical significance for $lnWage$, $WageRank$ and the two interaction terms: between $WageRank$ and $WageSD$; and $lnWage$ and $WagePol$. The coefficient on $WageRank$, for when $WageSD$ is zero, remains statistically insignificant which, alongside the positive coefficient on the interaction with $WageSD$, suggests the relationship between wage rank and player performance is weak when the standard deviation of wage is low, and stronger when it is high. More specifically, when $WageSD$ is at the sample average of \$4.895m and the other variables in the model are held constant, the model in column (1) of Table 7.18 predicts that a one standard deviation increase in $WageRank$ will increase wins produced by 0.247. When $WageSD$ is at the sample minimum of \$1.554m, the model predicts this difference to be negative but small at -0.028 wins produced; while at the sample maximum of \$10.166m, the model predicts the difference to be 0.681 wins produced. The effect of the standard deviation of wage on player performance for different wage ranks tells a similar story. At the mean of $WageRank$ (just under 10^{th}), a one standard deviation increase in $WageSD$, holding the other variables in the model constant, is predicted to increase a player's production by 0.199 wins. For the top earner in each team, the predicted increase in production is 0.404 wins; while for the lowest wage rank, it is predicted to decrease performance by 0.211 wins produced. The overall effect of a higher standard deviation of wage is positive, offering support for Tournament Theory (Lazear & Rosen, 1981). But this additional analysis suggests the effects depend on where a player ranks in the team in terms of wage. More effort is encouraged from the 'winners' of the tournament who rank higher in the team's wage structure. While there is some evidence that the lowest earners in the team may be discouraged.

[63] Defined by Simmons and Berri (2011) as those which played at least 20 games and at least 12 minutes per game.

Table 7.18: Wage Dispersion, National Diversity and Player Performance (Performance Measures)

	(1) WP	(2) WS	(3) VORP	(4) WP48	(5) PER	(6) BPM
lnWage	0.465***	0.542***	0.134***	-0.0166*	-0.149	0.569***
	(0.0926)	(0.0691)	(0.0324)	(0.00960)	(0.338)	(0.216)
WageRank	-0.0283	-0.00533	-0.00921	0.00220	0.0981	0.0432
	(0.0288)	(0.0183)	(0.00888)	(0.00195)	(0.0639)	(0.0478)
WageSD	0.271***	0.204***	0.0878***	0.0106***	0.136	0.266***
	(0.0787)	(0.0512)	(0.0274)	(0.00341)	(0.117)	(0.0868)
WagePol	0.545	0.716**	0.223	-0.114***	-2.468*	-0.601
	(0.439)	(0.324)	(0.142)	(0.0416)	(1.486)	(1.024)
WageRank × WageSD	0.0150***	0.00997***	0.00498***	0.000593*	0.00956	0.000957
	(0.00546)	(0.00342)	(0.00173)	(0.000351)	(0.0115)	(0.00831)
lnWage × WagePol	-0.980***	-1.047***	-0.401***	0.0420*	1.829**	0.396
	(0.303)	(0.220)	(0.111)	(0.0242)	(0.932)	(0.583)
LingDist2	-0.757**	-0.428*	-0.223**	-0.0324	-1.182	-1.001
	(0.364)	(0.236)	(0.0977)	(0.0256)	(0.870)	(0.637)
NatFrac2	3.232**	2.590***	1.503***	0.147**	2.730	5.069***
	(1.396)	(0.921)	(0.428)	(0.0689)	(2.469)	(1.681)
NatPol2	-1.855*	-1.521**	-0.875**	-0.100*	-1.340	-2.484*
	(1.126)	(0.751)	(0.355)	(0.0562)	(1.985)	(1.329)
Constant	6.057***	5.516***	1.405**	0.0695	13.61***	2.216
	(1.846)	(1.278)	(0.559)	(0.100)	(3.062)	(2.039)
log likelihood	-14515.2	-12352.4	-7601.9	2335.4	-18529.0	-16299.0
Players	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Years	✓	✓	✓	✓	✓	✓
Position	✓	✓	✓	✓	✓	✓
Player Controls	✓	✓	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓	✓	✓

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

In column (1) of Table 7.18, we also see a statistically insignificant coefficient on $WagePol$, alongside a highly statistically significant coefficient on the interaction term between $WagePol$ and $lnWage$. This suggests that wage polarisation does not affect the performance of players earning around \$1m (when $lnWage$ equals zero), but there are positive effects for lower earners and negative effects for higher earners which appear to be cancelling out at a wage of around \$1m. When $lnWage$ is at the sample average of 0.740 (which indicates a wage of \$2.096m, which is just below the sample median of \$2.398m), the predicted effect of a one standard deviation change in $WagePol$ is -0.014 wins produced. When $lnWage$ is at the sample minimum of -5.293, the predicted effect is 0.447; while at the sample of maximum for $lnWage$ of 3.546, the predicted effect is -0.229 wins produced. A similar analysis can be conducted to find the effects of increasing a player's wage for given levels of wage polarisation. When $WagePol$ is at the sample average of 0.215, the model in column (1) of Table 7.18 predicts that a one standard deviation increase in $lnWage$ will increase player production by 0.364 wins. When $WagePol$ is at the sample minimum of 0.05, the predicted effect is 0.596; while at the sample maximum of 0.516, the predicted effect is negative but small at -0.058 wins

produced. Overall, we find evidence that the effects of wage polarisation are similar to those predicted by Cohesion Theory (Levine, 1991). Furthermore, the effects can vary depending on a player's wage. While there is some evidence that low earners may perform better in teams with a high wage polarisation, the performance of the top earners is maximised when wage polarisation is low. The theories regarding polarisation make links to conflict and a lack of cooperation (Duclos et al., 2004; Foster & Wolfson, 2010; Montalvo & Reynal-Querol, 2005), so these results could be driven by players adopting individualistic styles of play when wage polarisation is high. Low earners simply aim to maximise their individual performance, rather than acting cooperatively for the good of the team, when wage polarisation is high. When wage polarisation is low, there is more cooperation throughout the team so the top earners, and the team overall, perform better.

The calculations for the predicted effects discussed are presented and summarised in Appendix 7.A. While these effects seem small, firstly there are a number of factors held constant. Importantly, we have discussed the effects of increases in $lnWage$ and $WageRank$ while holding the other constant. When $WageSD$ and $WagePol$ are at the sample means, a one standard deviation increase in $lnWage$ which also increases $WageRank$ by one standard deviation, would be predicted to increase an average player's production by roughly 0.601 wins (0.364 + 0.247). The predicted effects of $WageSD$ and $WagePol$ are also per player. So, for a sample average roster size of roughly 18, a one standard deviation increase in $WageSD$, holding other variables constant, is expected to increase team performance by roughly 3.647 wins[64] or increase team winning percentage by roughly 4.5 percentage points[65], which is very similar to the predicted effect found at the team-level (discussed in Section 6.4). It is much more complex to get a estimate of the overall team effect of $WagePol$ as, unlike $WageRank$ which takes the same values for every team of 18 players, wages (and therefore $lnWage$) can take many different values. For the purposes of demonstration, if we assumed every player in an 18-man roster is paid the sample average of their wage rank (presented in Table 7.1), then the predicted effect of a one standard deviation increase in $WagePol$ is a reduction of 0.5 wins or a decrease in team winning percentage of 0.6 percentage points[65]. This is much smaller than the effect discussed in Section 6.4 because, even if a hypothetical team paid each wage rank the sample mean, as the team gets more polarised in terms of wage, there would be more clustering at both ends of the distribution and, therefore, each wage rank would no longer get the sample mean of that rank. Still, the example demonstrates the negative effects of wage polarisation across the team.

In terms of national diversity, the coefficient on $NatPol2$ becomes statistically significant, at a 10% level in column (1) of Table 7.18, when additionally controlling for wage dispersion. In Section 6.4, we saw a similar result at the team-level when including both wage dispersion and national diversity in a combined model, with the statistical significance of $NatPol2$ increasing from a 10% to a 5% level in that instance (see Table 6.15). This provides further evidence of a link between wage polarisation and national polarisation, suggesting each should be controlled for when analysing the effects of the other. The effects of $LingDist2$ and $NatFrac2$ remain statistically significant at a 5% level.

The model in column (1) of Table 7.18 predicts that a one standard deviation increase in $LingDist2$ is expected to decrease the production of a player by 0.152 wins. However, this

[64] This is calculated by summing the marginal effects at each value of $WageRank$ between -18 and -1.
[65] Assuming an 82 game season.

will also have an effect on the team-level variables. A one standard deviation increase in $NatFrac2$ is predicted to increase performance by 0.286, while a one standard deviation increase in $NatPol2$ is expected to decrease production by 0.207 wins. Multiplying these team-level predictions by 18 (the sample mean of roster size), brings these predictions to 5.148 wins or 6.28 percentage points for $NatFrac2$ and 3.726 wins or 4.544 percentage points for $NatPol2$.[66] These are a little higher than the team-level predictions discussed in Section 6.4, though the average linguistic distance of each player needs to be factored in. It is not possible to increase the average linguistic distance of the team without increasing the average linguistic distance of at least some players to their teammates; so again a trade-off is indicated. Again, this relationship is difficult to predict as there are several ways that the structure of national diversity can vary. Adding an additional player, even with a given linguistic distance to his teammates could change the average linguistic distance of his teammates in vastly different ways. However, the model does suggest that an ideal structure of nationalities is one which is highly fractionalised, but with low national polarisation.

In columns (4)-(6) of Table 7.18, we see the model does not fit the efficiency measures so well, but there are some consistent results. The relationship between $WageRank$, $WageSD$ and performance seems robust, with an insignificant effect of $WageRank$ when $WageSD$ equals zero; and some evidence, for $WP48$ in column (4), that the positive effects of $WageRank$ increase with $WageSD$. The coefficient on $WageSD$ is also positive and statistically significant for $WP48$ and BPM at a 1% level, suggesting that a higher standard deviation of wage also has positive effects on player efficiency. The relationship between $lnWage$, $WagePol$ and player performance does not seem robust in terms of player efficiency, with some even contradictory results such as the positive coefficients on the interaction term between $lnWage$ and $WagePol$ (though in these cases the coefficient on $WagePol$ is significantly negative). This could be evidence of a completely different relationship between wage polarisation and player efficiency than was found for overall contribution. However, it rather appears that these contradictory results are driven by the outliers of the efficiency measures, as the evidence in Table 7.19 suggests.

The relationships between national diversity and player efficiency, in Table 7.18, all have consistent signs to our previous discussion; though the team-level variables have higher levels of statistical significance. This suggests that the effects of national diversity on player efficiency are mainly driven by team-level diversity, without significant differences depending on how diverse a player is from the central culture of the team (i.e. how large his linguistic distance to his teammates is).

The final piece of regression analysis in this chapter considers the use of different econometric models to estimate the relationships discussed. In columns (1)-(3) of Table 7.19, we reduce the sample by removing additional observations of players who play for multiple teams within a season, keeping only the observation for the team which they played the most minutes for. This reduces the sample size to 5,213 but allows for a panel data structure to be set. The results obtained using the 'complete pooling' Fixed Effects (FE) approach (including player, team, year and position dummies), and the 'no pooling' Random Effects (RE) approach (including no dummies), are presented in columns (2) and (3), respectively. Each shows consistency to the multilevel (MLM) approach presented in column (1) in terms of the signs of the coefficients.

[66] Assuming an 82 game season.

However, the results for FE show much lower levels of statistical significance as a result of ignoring the variation between players and only considering the variation of players over time (within variation). The results obtained using RE in column (3) show some higher levels of statistical significance, but do not control for either player- or team-level heterogeneity. The Mundlak (1978) approach to the Hausman test is not valid for highly unbalanced panels, but a standard Hausman test can be conducted once the clustering of standard errors is removed. This test produces a Wald statistic of 648 (p-value=0.000) which suggests there are high levels of heterogeneity and therefore the RE estimates may be biased.

Table 7.19: Wage Dispersion, National Diveristy and Player Performance (Alternative Models)

	(1) WP	(2) WP	(3) WP	(4) WP	(5) WP	(6) WP
lnWage	0.499***	0.134	0.638***	0.902***	0.257	1.034***
	(0.108)	(0.125)	(0.120)	(0.251)	(0.279)	(0.261)
WageRank	-0.0154	-0.0350	-0.0176	-0.0975**	-0.0843	-0.104**
	(0.0318)	(0.0386)	(0.0326)	(0.0468)	(0.0539)	(0.0467)
WageSD	0.267***	0.171*	0.306***	0.224**	0.114	0.182**
	(0.0831)	(0.0970)	(0.0877)	(0.0939)	(0.108)	(0.0927)
WagePol	0.448	-0.622	1.036*	4.043**	3.782*	5.367***
	(0.543)	(0.698)	(0.582)	(1.807)	(2.265)	(1.848)
WageRank × WageSD	0.0124**	0.00625	0.0181***	0.0134*	0.0114	0.0129*
	(0.00560)	(0.00664)	(0.00588)	(0.00739)	(0.00873)	(0.00755)
lnWage × WagePol	-0.965***	-0.599	-1.182***	-2.734**	-2.686**	-2.632**
	(0.364)	(0.401)	(0.407)	(1.113)	(1.338)	(1.168)
LingDist2	-0.784**	0.366	-0.472	-0.966	0.318	-0.273
	(0.387)	(2.093)	(0.404)	(0.597)	(2.227)	(0.612)
NatFrac2	3.661**	2.934	2.651	8.567***	6.150*	6.370**
	(1.809)	(2.452)	(1.862)	(2.935)	(3.619)	(2.951)
NatPol2	-2.078	-2.139	-1.417	-4.682**	-4.206*	-3.493*
	(1.385)	(1.635)	(1.425)	(2.107)	(2.372)	(2.105)
Constant	6.762***	5.902**	8.036***	12.86***	9.510***	13.95***
	(2.097)	(2.570)	(1.778)	(2.720)	(3.088)	(2.207)
N	5213	5213	5213	3546	3546	3546
log likelihood	-13013.9	-11532.8	-	-9158.3	-8045.7	-
R^2	-	0.186	0.243	-	0.203	0.181
Players	✓	✓	✗	✓	✓	✗
Teams	✓	✓	✗	✓	✓	✗
Years	✓	✓	✗	✓	✓	✗
Position	✓	✓	✗	✓	✓	✗
Player Controls	✓	✓	✓	✓	✓	✓
Team Controls	✓	✓	✓	✓	✓	✓
Core Team	✗	✗	✗	✓	✓	✓
Model	MLM	FE	RE	MLM	FE	RE

Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

In columns (4)-(6), only the 'core team'[67] is considered, further reducing the sample to 3,546. The difference between the MLM, FE and RE approaches are similar to the previous descrip-

[67] Defined by Simmons and Berri (2011) as those which played at least 20 games and at least 12 minutes per game.

tion and we still see consistent coefficients across these models, except for the coefficients on $LingDist2$, though they are consistently not statistically different from zero.

Comparing the 'core team' results in columns (4)-(6) of Table 7.19 to those in columns (1)-(3), there is some consistency but also some significant differences. Firstly, the negative coefficient on $WageRank$ becomes statistically significant at a 5% level for the MLM approach in column (4) and the RE approach in column (6). This suggests that when $WageSD$ is low and for a given wage level, players perform worse at higher wage ranks within the 'core team'. This may be caused by some sense of injustice or jealousy compared to others. When excluding $lnWage$, we no longer see this statistical significant effect. The positive coefficient on $WagePol$ also becomes statistically significant, though the magnitude of the coefficient on the interaction term between $lnWage$ and $WagePol$ also increases significantly in magnitude. In terms of national diversity, we see a similar effect as found in Section 7.4.1, with the team-level measures seeming more important than the player-level measures when considering the 'core team'.

7.6 Summary and Discussion

This chapter has investigated the effects of both wage dispersion and national diversity on the performance of players within teams. We find evidence that both the wage of players, and their wage ranking within a team, have positive relationships with player performance. Furthermore, we find evidence that each of these influence different elements of the relationship between wage dispersion and player performance. In Chapter 5, it was found that a higher standard deviation of wage has positive effects on performance. This gives evidence for the ideas of Tournament Theory (Lazear & Rosen, 1981), that a higher wage dispersion encourages players to compete to reach the top of the wage ranking. The evidence in this chapter suggests that not all players are motivated by a higher standard deviation of wage, but the increase in team performance is driven by the 'winners' of the tournament, i.e. those which rank highly in terms of wage. Those that rank near the bottom of the wage distribution perform worse when the standard deviation of wage is high, so not all players benefit from such a wage structure, but the team does overall.

In Chapter 5, we also found that wage polarisation has negative effects on team performance, which fits more closely to Cohesion Theory (Levine, 1991). In this chapter, we find the player-level effects of wage polarisation are linked to a player's wage. Our evidence even suggests that players on low wages perform better when wages are polarised. We attribute this to a polarised wage structure encouraging individualistic, rather than cooperative, play. This idea has some precedence in the theoretical literature, in terms of increased sabotage and lack of cooperation in teams with dispersed wage structures (Levine, 1991), and how polarisation can lead to conflict and a lack of cooperation (Montalvo & Reynal-Querol, 2005). When wage polarisation is low, players on higher wages perform better. This suggests that, in less polarised wage structures, players on lower wages may be willing to sacrifice their individual performance (and/or are less willing to sabotage) in order to facilitate the performance of others for the good of the team. Though, a more detailed analysis, which qualifies individualistic versus cooperative play, is required to confirm this hypothesis.

An alternative explanation could be that a polarised wage structure only creates competition between low earners and high earners, and not within these groups. For a given standard

deviation of wage, a more polarised wage structure would have more clustering at either end of the distribution. This creates clear groups of higher earners and low earners, with little in-between, and smaller wage variation within each of these groups. This may encourage higher levels of minimum effort to remain within a wage grouping, but discourage higher maximum levels of effort due to the smaller differences in wages within the groups. This would explain the higher performance of those with low wage rankings as they strive to stay within their current wage grouping, as well as the relative lower performance of the top wage ranks, as they have little incentive to further increase their performance once they are among those in the group of high earners. In Section 7.5, there is some evidence that the effects of wage dispersion may differ within the 'core team', compared to considering the entire roster, and there is scope for future work to investigate these types of dynamics of smaller groups within teams.

This investigation, and other related studies (Addesa et al., 2017; Brandes et al., 2009; Brox & Krieger, 2019; Lazear, 1999), highlight that national diversity can have positive effects on performance, through a range of cultural backgrounds having the potential to improve productivity; but equally warn of the possible negative effects associated with language barriers and conflicting cultural norms, which can harm performance. These previous studies have sought to capture which of these effects dominate, and/or whether a quadratic relationship exists. Chapter 6 provides an alternative explanation that the structure of diversity matters. We find evidence that national fractionalisation maximises the benefits and, while holding fractionalisation constant, national polarisation maximises the costs. This chapter finds evidence of these effects at the player-level, and also suggests that players with larger average linguistic distances to their teammates may face additional difficulties which have negative effects on their performance.

Our most robust result, in terms of national diversity, across the specifications of this chapter, is the positive relationship between national fractionalisation and player performance (particularly when accounting for linguistic distance). We also find evidence that players with a smaller linguistic distance to their teammates perform better. Together these findings suggest a type of trade-off. Players with a large linguistic distance to their teammates offer diversity and aid the 'search for best practice' (Lazear, 1999), but it is those who are linguistically closer to their teammates on average which benefit most from the resulting positive spillover effects. This suggests that there are additional difficulties faced by players which are linguistically further from the 'central culture'[68] of the team. These difficulties may relate to language barriers, conflicting cultural norms or perhaps cultural adjustment.[69] We also find some evidence that, for a given level of national fractionalisation, players in teams which are more nationally polarised perform worse. Though, this is not a robust result at the player-level across all the specifications of this chapter, the evidence is stronger when considering the 'core team' and when also controlling for wage dispersion in the full combined model of Section 7.5.

[68] One may think of a 'central culture' as either the culture of the largest national group, or some 'synthetic culture' which is a weighted average of the players' individual cultures.

[69] Though players with larger average linguistic distances to their teammates need not be non-Americans, they typically are in our dataset. While we control for each player's NBA experience and tenure with their current team, we do not capture whether a player is also adjusting to the cultural environment of a new country, except partially and indirectly through linguistic distance.

Chapter 8

Conclusion

Chapter Overview

This chapter summarises the main findings from our investigation into the effects of wage dispersion and national diversity on performance, as well as highlighting the key contributions made to the existing literature. We begin by justifying our choice of data, before discussing the findings relating to wage dispersion and national diversity in turn. Next, we consider the further insights gained, and the difficulties overcome, when conducting a player-level analysis to supplement the team-level findings. We also identify limitations of the investigations and make suggestions for future work, before finishing with some concluding remarks.

8.1 Choice of NBA

This investigation has analysed the effects of wage dispersion and national diversity on performance, using a detailed dataset from the National Basketball Association (NBA), spanning the seasons 2007 to 2018 inclusive. This is initially done at the team-level, and then further insights are sought through an analysis of player performance within teams. Wage dispersion and national diversity are pressing issues within the NBA, but also in other sports and, more broadly, within almost every work-based team environment. Despite some peculiarities, sports labour markets operate under similar economic principles to general labour markets (Neale, 1964; Rosen & Sanderson, 2001; Rottenberg, 1956; Scully, 1974). They therefore offer a unique opportunity to test labour market theories in what Kahn (2000) describes as a labour market laboratory; where detailed information about workers is available, firms operate under similar conditions with clear performance objectives, and performance is easily observable. This makes the data ideal for the current investigation. In particular, basketball is a sport where production can be highly interdependent and therefore requires high levels of cooperation within teams. Although this presents difficulties when measuring individual player performance, the interdependence of production makes the issues of wage dispersion and national diversity, and their effects on performance within teams, all the more interesting.

8.2 Wage Dispersion

The NBA operates under a 'soft' wage cap which limits the expenditure of teams on wages. More generally, firms operate within a wage budget and look to maximise performance within this. By dispersing their wage structure, teams can offer higher wages to some of their workers and therefore attract the top talents in the industry. However, to stay within the wage budget, they then must also offer lower wages to other workers. These wage differences can affect the incentives of players and how the players interact with one another. Chapter 2 discusses the academic literature on the topic of wage dispersion and identifies two main competing hypotheses: Tournament Theory (Lazear & Rosen, 1981) and Cohesion Theory (Levine, 1991). Tournament Theory predicts that a higher wage dispersion will increase player effort to compete for the higher wages on offer, while Cohesion Theory predicts that higher wage dispersion can lead to unproductive effort through sabotage and a lack of cooperation. Both theories have received empirical backing within the sports industry and more generally. The literature also offers some reconciliations of these contradicting theories in Equity Theory (see: Leventhal, 1976; Trevor et al., 2012) and the 'Hawks and Doves' model (Lazear, 1989), though the practical, ex ante application of these is questionable.

We contribute to this literature in several ways. Firstly, we note that a number of measures are considered in the literature under the broad term of dispersion. We consider a number of these in our analysis to see if they are capturing differing effects. Secondly, we introduce a new measure into this branch of the literature: namely, polarisation. While the theories regarding polarisation have been applied to income at the country-level (Duclos et al., 2004; Foster & Wolfson, 2010), we find them to also be relevant to the topic of wage dispersion within teams. Chapter 3 gives a theoretical overview of polarisation, demonstrating how it captures unique aspects of a wage distribution which are not captured by other dispersion measures. We speculate that the omission of polarisation measures in the current literature may explain some of the mixed evidence found, due to other measures only partially covarying with polarisation. Thirdly, given that we show different measures of dispersion capture

unique aspects of a distribution, we consider multiple measures of dispersion simultaneously. Our empirical analysis supports these contributions with evidence to suggest that both the standard deviation and polarisation of wage have differing effects and simultaneously affect performance. We also show that other measures of dispersion appear to have differing effects on performance when other measures of wage dispersion are included simultaneously, and depending on the specification of other control variables included. Whereas the positive effects of wage standard deviation and the negative effects of wage polarisation show robustness to such changes.

Specifically, the empirical analysis of Chapter 5 finds that a high standard deviation of wage encourages effort and has positive effects on the team overall. However, this effect is not consistent across all players, as revealed in Chapter 7. Those at the lower end of their respective team's wage distribution perform worse when the standard deviation of wage is high, while those at the top end are driving this overall positive effect. While this aspect of wage dispersion is linked to a player's relative position within a team's wage distribution, the effects of wage polarisation are more closely linked to a player's absolute wage. Chapter 5 also finds that high wage polarisation has negative effects on team performance overall. Polarisation theories (Duclos et al., 2004; Foster & Wolfson, 2010; Montalvo & Reynal-Querol, 2005) suggest that this is due to polarisation causing an in- versus out-group mentality. Chapter 7 additionally reveals that, when wage polarisation is high, we see some evidence of lower paid players increasing their performance. However, this appears to come at the detriment of players on high wages, therefore resulting in a negative effect on the team overall. This suggests that, when wage polarisation is high, players seem less willing to cooperate for the good of the team. Overall, our results suggest that an ideal wage structure, to maximise team performance, should have a high spread across different wage-levels and avoid clusterings, especially at the extremes of the distribution.

8.3 National Diversity

Globalisation has led to a general increase in national diversity within workplaces and this is also true within the NBA. Our dataset contains players from over 60 different countries who have played in the NBA at some point during the sample period used in this investigation. Chapter 2 discusses how increased national diversity can lead to a range of cultural backgrounds which aids the 'search for best practice' (Lazear, 1999). However, national diversity can also present problems in the form of language barriers and conflicting cultural norms. National diversity can also be a difficult concept to capture, with various measures being utilised in the literature. Measures of fractionalisation are often used (Easterly & Levine, 1997; La Porta et al., 1999; Leigh, 2006) to investigate issues such as ethnolinguistic fractionalisation at the country-level on macroeconomic outcomes. There is also some suggestion that polarisation may again be playing a role (Montalvo & Reynal-Querol, 2005). Other studies focus on cultural distance, often proxied by linguistic distance (Bakker et al., 2009; Ingersoll et al., 2017), and there are also studies which combine fractionalisation and linguistic distance (e.g. Baldwin & Huber, 2010). Investigations of these issues within the sports industry are relatively novel and have only been considered in relatively recent studies (Addesa et al., 2017; Brandes et al., 2009; Brox & Krieger, 2019; Ingersoll et al., 2017; Kahane et al., 2013).

Our contribution to this literature is to investigate the effects of national fractionalisation and polarisation on performance within teams, both separately and simultaneously. Through

an analysis of the equations, simulations and through demonstrative examples, presented in Chapter 3, we show how these measures are linearly correlated when a team consists of two or less national groups, but as the number of national groups increases, this positive correlation decreases, and even becomes negative at high levels of national fractionalisation. Although there is high correlation between these measures of national fractionalisation and polarisation in our NBA dataset, due to most teams having a majority USA players, we find evidence that the measures have distinct and opposing effects on performance when considered simultaneously.

We also make a contribution regarding the use of linguistic distance as a proxy of cultural distance. Conventionally, the culture of a country is assumed to be captured by their single, primary language. However, we argue that other languages commonly spoken in countries may also influence and/or be indicators of their culture (Chen, 2013). As detailed in Chapter 3, we construct measures of linguistic distance which consider up to four commonly spoken languages per country, weighted by the proportion of speakers. The evidence from our dataset suggests that additional insights can be gained by considering two languages, but considering more than this causes little change in the proxy. We also show how these measures of linguistic distance can be combined with fractionalisation (Baldwin & Huber, 2010), and additionally combine them with measures of polarisation.

In the empirical analysis of Chapter 6, we find evidence that national fractionalisation has positive effects on performance. While national polarisation appears to also have positive effects when included separately, when holding national fractionalisation constant, we find it to have negative effects on performance. Therefore, the positive effect initially found seems to be driven by the correlation between national fractionalisation and national polarisation. These results suggest that an ideal national structure to maximise performance is a highly fractionalised one, which avoids polarisation between two dominant national groups. However, we also note that examples of teams with this type of national structure are rare in our dataset, with many teams consisting of a majority of players from the USA. We are therefore wary that such out-of-sample predictions require further research to test.

Additionally, in Chapter 7, for a given level of national diversity in the team, we find that players with a smaller average linguistic distance to their teammates tend to perform better. This suggests that, while the team benefits from national fractionalisation through the positive spillovers that national diversity brings, players who are further from the central national culture of the team may benefit the least. This presents a trade-off between seeking to maximise individual player or team performance. Though we do not find statistically significant evidence of this in our study, future research may investigate whether this apparent trade-off disappears when fractionalisation is high and there is, therefore, no dominant culture within the team. If true, this may also explain the resistance encountered when a formally homogeneous team becomes more diverse.

8.4 Further Insights from Player Analysis

Another contribution of this investigation is that the effects of wage dispersion and national diversity have also been analysed at the player-level. The results of this analysis have been discussed in the preceding discussion, but this additional level of analysis is new to this section of the literature and offer important insights for future investigations. A key benefit of sports

data is its richness and detail, which makes this level of analysis possible. This study has exploited this to gain additional insights into the effects that wage dispersion and national diversity have on individual players within teams. Our main findings are that the effects of wage dispersion depend on a player's wage, as well as where they rank within the team's wage structure. We also find some evidence to suggest that the players who provide the most national diversity to a team are the ones which benefit the least from it.

There are, of course, additional difficulties encountered when conducting this further level of analysis, which this investigation has sought to overcome. One of these difficulties is the measurement of individual player performance within teams. This is particularly difficult to measure when production is interdependent. However, it is precisely when production is interdependent that the issues of wage dispersion and national diversity become more salient, so there is certainly a trade-off here. Measures of individual player performance in basketball have been developed in the academic literature and, even more so, by sports analysts outside of the academic framework. Though there are issues regarding the lack of formal vetting through an expert peer-review system (Berri & Bradbury, 2010), some interesting work has been conducted by so-called 'metricians'. In our analysis, we consider a number of these performance measures (from both academic and non-academic work), including some which capture the overall contribution of players, and others which capture the efficiency of performance when players are given opportunities to play. While we favour measures of overall contribution for the bulk of our analysis (considering efficiency for robustness checks), there are pros and cons to both approaches. As these measures continue to develop, the insights gained from player-level analysis could be further enhanced.

Another difficulty is in the structure of player-level data. While team-level data is usually highly balanced, there are higher levels of attrition when considering players. Players may also be traded and play for multiple teams during a season, which does not fit a panel-data structure well. Furthermore, we may expect heterogeneity at both the team- and player-level which should be accounted for. This study utilises Multilevel Modelling (MLM) (Gelman & Hill, 2006) to address these econometric difficulties. These models account for heterogeneity at multiple levels of analysis and allow for multiple player observations for different teams within a season. They are therefore ideally suited to an analysis of players within teams.

8.5 Limitations and Further Work

In this study, we proxy the cultural distance between countries using linguistic distance. This is a common approach in the literature (e.g. Alesina et al., 2003; Bakker et al., 2009; Baldwin & Huber, 2010; Desmet et al., 2009; Ingersoll et al., 2017). We sought to build on this by constructing a measure of linguistic distance which accounts for multiple commonly-spoken languages within countries. While this appears to act as a reasonable proxy, there is scope for further improvement. One problem our proxy suffers from is that countries which speak the same languages will be assigned a distance of zero, even if these languages are spoken by vastly different proportions of the population. For example, the two most commonly spoken languages in both the UK and Canada are English and French, and so the countries are assigned a distance of zero. Although the shared languages do indicate similarities, it seems incorrect to say there are no cultural differences between the countries. An improved proxy could account for these differences in the proportion of speakers even when the same languages are spoken in each country. It could also consider other factors which indicate cultural

differences/similarities such as geographical distance, ethnicity, religion and/or trade volume between the countries. These factors could be considered instead of, or alongside, linguistic distance.

In this investigation, we classify the nationality of players based on their place of birth. However, while we control for the level of NBA experience and tenure with the current team, we do not capture whether players grew up in a different country from their country of birth. We also do not account for the parentage or dual nationality of players. For example, Ime Udoku is a Nigerian-American who, due to his Nigerian parentage, represents the Nigerian basketball team internationally. However, he was born in the USA, so is therefore classified as American in our dataset. Conversely, Wally Szczerbiak is a Spanish-American who was born in Spain (to an American father who was born in Germany to Ukrainian parents), and is therefore classified as Spanish. It is difficult, if not impossible, to capture all of the factors which may reflect a player's cultural background. However, future studies could consider the dual nationalities of players, their parentage, the national team they choose to represent, whether they competed in the National Collegiate Athletic Association (NCAA), and/or whether they have competed in any other national leagues besides the NBA.

Another natural extension to the work of this investigation is to apply the theories and results, relating to wage dispersion and nationality, to other sports and workplaces. While the robustness of our results have been tested within the current NBA dataset and there is theoretical reason to believe they would apply to a range of workplace teams, further research is required to test this. Sports, and other industries, which do not operate within a wage cap offer the opportunity to test whether our results relating to wage dispersion are robust when competing teams have vastly different wage budgets. Furthermore, other datasets (and possibly the NBA in the future) consist of higher levels of national fractionalisation, which would offer the opportunity to test our prediction that full national fractionalisation is the optimal structure for teams to maximise performance. European Soccer offers both of these. Although Financial Fair Play limits the wage expenditure of teams to a degree, teams operate under vastly different wage budgets within the same leagues. Many leagues also have much higher levels of national fractionalisation within teams. However, collecting verifiable wage data and measuring individual player performance are two potential problems of using European Soccer data.

Lastly, another potential extension to the work of this thesis is to consider the role of the coach/management team. It is possible that good coaches structure their rosters considering these issues, either explicitly or intuitively. It is also reasonable to assume that better coaches are able maximise the benefits and minimise costs relating to wage dispersion and national diversity through shrewd man management. This would be a difficult aspect to capture, but other studies have looked to capture the quality of coaches by considering their historical performance records. Also, both the wage and nationality of the coach could have implications for how they manage these issues and therefore their effects on performance.

8.6 Concluding Remarks

A major theme of this investigation has been a focus on measures of polarisation. We have shown how polarisation can be measured for both discrete and continuous variables, demonstrated that it captures distinct information about distributions, and find it to capture dis-

tinct elements of both wage dispersion and national diversity which affect performance. In economics, and in academic work more generally, many different measures of dispersion are considered when analysing a range of issues. For example, the Gini Coefficient is a commonly used measure, particularly in economics, and it is widely acknowledged to capture distinct information about a distribution (namely inequality) compared to other measures, such as the standard deviation (which measures spread). Polarisation captures yet another distinct element about distributions. Although the indexation of polarisation is relatively new (Duclos et al., 2004; Foster & Wolfson, 2010; Montalvo & Reynal-Querol, 2005), we believe they are currently underutilised measures in the literature, and additional insights into a range of topics could be gained by considering polarisation.

Polarisation could even, perhaps, be taken a step further. If a researcher was interested in clusterings, not just on either side of the median or between two discrete groups, but instead at more intervals, then it may be possible to adapt polarisation measures to include a cubic or higher power term, rather than a squared one. This is an idea left to future research.

References

Abdelkrim, A., & Duclos, J.-Y. (2007). DASP: Distributive Analysis Stata Package. PEP, World Bank, UNDP and Universitè Laval. Retrieved from http://www.dasp.ecn.ulaval.ca/

Abrams, J. (2010). An NBA move that crossed an ocean. New York Times. Retrieved from https://www.nytimes.com/2010/11/18/sports/basketball/18moves.html

Addesa, F., Rossi, G., & Bove, V. (2017). Cultural diversity and team performance in the Italian Serie A. In 2rd International Conference Sport Economics & Sport Management. Retrieved from http://www.eprints.leedsbeckett.ac.uk/3808/

Alesina, A., Devleeschauwer, A., Easterly, W., Kurlat, S., & Wacziarg, R. (2003). Fractionalization. Journal of Economic growth, 8(2), 155–194.

Alesina, A., & Ferrara, E. L. (2005). Ethnic diversity and economic performance. Journal of Economic Literature, 43(3), 762–800.

Baker, G., Gibbs, M., & Holmstrom, B. (1994). The wage policy of a firm. The Quarterly Journal of Economics, 109(4), 921–955.

Bakker, D., Müller, A., Velupillai, V., Wichmann, S., Brown, C. H., Brown, P., ... Holman, E. W. (2009). Adding typology to lexicostatistics: a combined approach to language classification. Linguistic Typology, 13(1), 169–181.

Baldassarri, D., & Bearman, P. (2007). Dynamics of political polarization. American Sociological Review, 72(5), 784–811.

Baldwin, K., & Huber, J. D. (2010). Economic versus cultural differences: forms of ethnic diversity and public goods provision. American Political Science Review, 104(4), 644–662.

Bell, A., Smith, J., Sabel, C. E., & Jones, K. (2016). Formula for success: multilevel modelling of Formula One driver and constructor performance, 1950-2014. Journal of Quantitative Analysis in Sports, 12(2), 99–112.

Bellotti, R. S. (1992). The points created basketball book, 1991-92. Night Work Publishing Company.

Berri, D. J. (1999). Who is 'most valuable'? Measuring the player's production of wins in the National Basketball Association. Managerial and Decision Economics, 20(8), 411–427.

Berri, D. J., & Bradbury, J. C. (2010). Working in the land of the Metricians. Journal of Sports Economics, 11(1), 29–47.

Berri, D. J., Schmidt, M., & Brook, S. (2006). The wages of wins: Taking measure of the many myths in modern sport. Stanford University Press.

Berri, D. J., & Simmons, R. (2009). Race and the evaluation of signal callers in the National Football League. Journal of Sports Economics, 10(1), 23–43.

Blackburn, M., & Bloom, D. E. (1985). What is happening to the middle class? American Demographics, 7(1), 18–25.

Bloom, M. (1999). The performance effects of pay dispersion on individuals and organizations. Academy of Management Journal, 42(1), 25–40.

Borman, G., & Dowling, M. (2010). Schools and inequality: a multilevel analysis of Coleman's equality of educational opportunity data. Teachers College Record, 112(5), 1201–1246.

Brandes, L., Franck, E., & Theiler, P. (2009). The effect from national diversity on team production. Schmalenbach Business Review, 61(2), 225–246.

Brox, E., & Krieger, T. (2019). Birthplace diversity and team performance. ZEW-Centre for European Economic Research Discussion Paper(19-020).

Brzezinski, M. (2013). Income polarization and economic growth. National Bank of Poland working paper(147).

Bucciol, A., Foss, N. J., & Piovesan, M. (2014). Pay dispersion and performance in teams. PloS one, 9(11), e112631.

Cameron, A. C., & Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. Journal of Human Resources, 50(2), 317–372.

Carmichael, F., & Thomas, D. (1995). Production and efficiency in team sports: an investigation of Rugby League Football. Applied Economics, 27(9), 859–869.

Carmichael, F., Thomas, D., & Ward, R. (2001). Production and efficiency in Association Football. Journal of Sports Economics, 2(3), 228–243.

Cederman, L.-E., Weidmann, N. B., & Gleditsch, K. S. (2011). Horizontal inequalities and ethnonationalist civil war: A global comparison. American Political Science Review, 105(3), 478–495.

Chen, M. K. (2013). The effect of language on economic behavior: evidence from savings rates, health behaviors, and retirement assets. American Economic Review, 103(2), 690–731.

Coates, D., Frick, B., & Jewell, T. (2016). Superstar salaries and soccer success: the impact of designated players in Major League Soccer. Journal of Sports Economics, 17(7), 716–735.

Collier, P., & Hoeffler, A. (1998). On economic causes of civil war. Oxford Economic Papers, 50(4), 563–573.

Coon, L. (2005). 2005 Collective Bargaining Agreement. NBA Salary Cap FAQ. Retrieved from `http://www.cbafaq.com/salarycap05.htm`

Coon, L. (2011). 2011 Collective Bargaining Agreement. NBA Salary Cap FAQ. Retrieved from `http://www.cbafaq.com/salarycap11.htm`

Coon, L. (2017). 2017 Collective Bargaining Agreement. NBA Salary Cap FAQ. Retrieved from `http://www.cbafaq.com/salarycap.htm`

Cramer, C. (2003). Does inequality cause conflict? Journal of International Development: The Journal of the Development Studies Association, 15(4), 397–412.

Dawson, P., Dobson, S., & Gerrard, B. (2000). Estimating coaching efficiency in professional team sports: evidence from English Association Football. Scottish Journal of Political Economy, 47(4), 399–421.

Depken, C. A. (2000). Wage disparity and team productivity: evidence from Major League Baseball. Economics Letters, 67(1), 87–92.

Depken, C. A., & Lureman, J. (2018). Wage disparity, team performance, and the 2005 NHL collective bargaining agreement. Contemporary Economic Policy, 36(1), 192–199.

Desmet, K., Weber, S., & Ortuño-Ortín, I. (2009). Linguistic diversity and redistribution. Journal of the European Economic Association, 7(6), 1291–1318.

Duclos, J.-Y., Esteban, J., & Ray, D. (2004). Polarization: concepts, measurement, estimation. Econometrica, 72(6), 1737–1772.

Easterly, W., & Levine, R. (1997). Africa's growth tragedy: policies and ethnic divisions. The Quarterly Journal of Economics, 112(4), 1203–1250.

Edwards, L. J., Muller, K. E., Wolfinger, R. D., Qaqish, B. F., & Schabenberger, O. (2008). An R2 statistic for fixed effects in the linear mixed model. Statistics in Medicine, 27(29), 6137–6157.

Eriksson, T. (1999). Executive compensation and tournament theory: empirical tests on Danish data. Journal of Labor Economics, 17(2), 262–280.

Esteban, J., & Ray, D. (1994). On the measurement of polarization. Econometrica: Journal of the Econometric Society, 819–851.

Esteban, J., & Ray, D. (2011). Linking conflict to inequality and polarization. American Economic Review, 101(4), 1345–74.

Ezcurra, R. (2009). Does income polarization affect economic growth? the case of the european regions. Regional Studies, 43(2), 267–285.

Fearon, J. D. (2003). Ethnic and cultural diversity by country. Journal of Economic Growth, 8(2), 195–222.

Ferreira, M., & Bravo, G. (2007). A multilevel model analysis of professional soccer attendance in chile 1990-2002. International Journal of Sports Marketing and Sponsorship, 8(3), 49–66.

Fort, R., & Maxcy, J. (2003). Competitive balance in sports leagues: an introduction. Journal of Sports Economics, 4(2), 154–160.

Foster, J. E., & Wolfson, M. C. (2010). Polarization and the decline of the middle class: Canada and the US. The Journal of Economic Inequality, 8(2), 247–273.

Franck, E., & Nüesch, S. (2011). The effect of wage dispersion on team outcome and the way team outcome is produced. Applied Economics, 43(23), 3037–3049.

Fraschilla, F. (2017). Top basketball leagues in the world outside the NBA. ESPN. Retrieved from `https://www.espn.co.uk/nba/story/_/id/21691644/fran-fraschilla-ranks-world-top-basketball-leagues-nba`

Fredrickson, J. W., Davis-Blake, A., & Sanders, W. G. (2010). Sharing the wealth: social comparisons and pay dispersion in the CEO's top team. Strategic Management Journal, 31(10), 1031–1053.

Frick, B., Prinz, J., & Winkelmann, K. (2003). Pay inequalities and team performance: empirical evidence from the North American Major Leagues. International Journal of Manpower, 24(4), 472–488.

Gelman, A., & Hill, J. (2006). Data analysis using regression and multilevel/hierarchical models. Cambridge University Press.

Gibbs, M. (1995). Incentive compensation in a corporate hierarchy. Journal of Accounting and Economics, 19(2-3), 247–277.

Goldstein, H., Browne, W., & Rasbash, J. (2002). Multilevel modelling of medical data. Statistics in medicine, 21(21), 3291–3315.

Golub, G. H., Hansen, P. C., & O'Leary, D. P. (1999). Tikhonov regularization and total least squares. SIAM Journal on Matrix Analysis and Applications, 21(1), 185–194.

Gould, E. D., & Winter, E. (2009). Interactions between workers and the technology of production: evidence from professional baseball. The Review of Economics and Statistics, 91(1), 188–200.

Greenland, S. (2000). Principles of multilevel modelling. International Journal of Epidemiology, 29(1), 158–167.

Groothuis, P. A., & Hill, J. R. (2008). Exit discrimination in Major League Baseball: 1990-2004. Southern Economic Journal, 574–590.

Groothuis, P. A., & Hill, J. R. (2013). Pay discrimination, exit discrimination or both? Another look at an old issue using NBA data. Journal of Sports Economics, 14(2), 171–185.

Grund, C., & Westergård-Nielsen, N. (2008). Age structure of the workforce and firm performance. International Journal of Manpower, 29(5), 410–422.

Hansen, Z., Owan, H., & Pan, J. (2006). The impact of group diversity on performance and knowledge spillover–an experiment in a college classroom (Vol. No. w12251; Tech. Rep.). National Bureau of Economic Research.

Heeren, D. (1988). The basketball abstract. Prentice Hall.

Hermans, R., Tondeur, J., van Braak, J., & Valcke, M. (2008). The impact of primary school teachers' educational beliefs on the classroom use of computers. Computers & education, 51(4), 1499–1509.

Heyman, F. (2005). Pay inequality and firm performance: evidence from matched employer–employee data. Applied Economics, 37(11), 1313–1327.

Hill, J. R. (2004). Pay discrimination in the NBA revisited. Quarterly Journal of Business and Economics, 81–92.

Hill, J. R., & Groothuis, P. A. (2017). Is there a wage premium or wage discrimination for foreign-born players in the NBA? International Journal of Sport Finance, 12(3), 204–221.

Hill, P. W., & Rowe, K. J. (1996). Multilevel modelling in school effectiveness research. School effectiveness and school improvement, 7(1), 1–34.

Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: biased estimation for nonorthogonal problems. Technometrics, 12(1), 55–67.

Hoffer, A. J., & Freidel, R. (2014). Does salary discrimination persist for foreign athletes in the NBA? Applied Economics Letters, 21(1), 1–5.

Hofler, R. A., & Payne, J. E. (1997). Measuring efficiency in the National Basketball Association. Economics Letters, 55(2), 293–299.

Hollinger, J. (2002). Pro basketball prospectus, 2002. Brassey's Incorporated.

Hollinger, J. (2005). Pro basketball forecast, 2005-06. Potomac Books.

Horowitz, D. L. (2001). Ethnic groups in conflict. University of California Press.

Humphreys, M., Posner, D. N., & Weinstein, J. M. (2002). Ethnic identity, collective action, and conflict: an experimental approach. Unpublished. Retrieved from `http://http://www.sscnet.ucla.edu/polisci/wgape/papers/1_Posner.pdf`

Ingersoll, K., Malesky, E., & Saiegh, S. M. (2017). Heterogeneity and team performance: evaluating the effect of cultural diversity in the world's top soccer league. Journal of Sports Analytics, 3(2), 67–92.

Jacobs, J. (2017). Introduction to Oliver's four factors. Squared Statistics: Understanding basketball analytics. Retrieved from `https://www.squared2020.com/2017/09/05/introduction-to-olivers-four-factors/`

James, B. (2001). The new Bill James historical baseball abstract. Simon and Schuster.

Jennett, N. (1984). Attendances, uncertainty of outcome and policy in Scottish League Football. Scottish Journal of Political Economy, 31(2), 176–198.

Johnes, J. (2006). Measuring efficiency: A comparison of multilevel modelling and data envelopment analysis in the context of higher education. Bulletin of Economic Research, 58(2), 75–104.

Kahane, L., Longley, N., & Simmons, R. (2013). The effects of coworker heterogeneity on firm-level output: assessing the impacts of cultural and language diversity in the National Hockey League. Review of Economics and Statistics, 95(1), 302–314.

Kahn, L. M. (2000). The sports business as a labor market laboratory. Journal of Economic Perspectives, 14(3), 75–94.

Kahn, L. M., & Sherer, P. D. (1988). Racial differences in professional basketball players' compensation. Journal of Labor Economics, 6(1), 40–61.

Keefer, P., & Knack, S. (2002). Polarization, politics and property rights: links between inequality and growth. Public Choice, 111(1-2), 127–154.

Kendall, T. D. (2003). Spillovers, complementarities, and sorting in labor markets with an application to professional sports. Southern Economic Journal, 389–402.

Khan, E. (2013). Advanced NBA stats for dummies: how to understand the new hoops math. Bleacher Report. Retrieved from https://www.bleacherreport.com/articles/1813902-advanced-nba-stats-for-dummies-how-to-understand-the-new-hoops-math

Kubatko, J., Oliver, D., Pelton, K., & Rosenbaum, D. T. (2007). A starting point for analyzing basketball statistics. Journal of Quantitative Analysis in Sports, 3(3).

Küpfer, E. (2005). Team similarity. APBRMetrics forum. Retrieved from http://www.sonicscentral.com/apbrmetrics/viewtopic.php?p=118&sid=ff925dbe120f6c6f86e596d0d6dabbb8

Lallemand, T., Plasman, R., & Rycx, F. (2004). Intra-firm wage dispersion and firm performance: evidence from linked employer-employee data. Kyklos, 57(4), 533–558.

La Porta, R., Lopez-de Silanes, F., Shleifer, A., & Vishny, R. (1999). The quality of government. The Journal of Law, Economics, and Organization, 15(1), 222–279.

Lazear, E. P. (1989). Pay equality and industrial politics. Journal of Political Economy, 97(3), 561–580.

Lazear, E. P. (1999). Culture and language. Journal of Political Economy, 107(S6), S95–S126.

Lazear, E. P., & Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. Journal of Political Economy, 89(5), 841–864.

Lee, K. W., Lev, B., & Yeo, G. H. H. (2008). Executive pay dispersion, corporate governance, and firm performance. Review of Quantitative Finance and Accounting, 30(3), 315–338.

Leigh, A. (2006). Trust, inequality and ethnic heterogeneity. Economic Record, 82(258), 268–280.

Leventhal, G. S. (1976). The distribution of rewards and resources in groups and organizations. In Advances in Experimental Social Psychology (Vol. 9, pp. 91–131). Elsevier.

Levine, D. I. (1991). Cohesiveness, productivity, and wage dispersion. Journal of Economic Behavior & Organization, 15(2), 237–255.

Levy, F. (1987). The middle class: Is it really vanishing? The Brookings Review, 5(3), 17–21.

Lewis, M. P., Simons, G. F., & Fennig, C. D. (2014). Ethnologue: Languages of the World, seventeenth edition. Dallas, Texas: SIL International. Retrieved from http://www.ethnologue.com

Li, F., Fisher, K. J., Brownson, R. C., & Bosworth, M. (2005). Multilevel modelling of built environment characteristics related to neighbourhood walking activity in older adults. Journal of Epidemiology & Community Health, 59(7), 558–564.

Main, B. G., O'Reilly III, C. A., & Wade, J. (1993). Top executive pay: tournament or teamwork? Journal of Labor Economics, 11(4), 606–628.

Mayer, T. (1960). The distribution of ability and earnings. The Review of Economics and Statistics, 189–195.

McCormick, R. E., & Clement, R. C. (1992). Intrafirm profit opportunities and managerial slack: evidence from professional basketball. Advances in the Economics of Sport, 1, 3–35.

Mirrlees, J. A. (1976). The optimal structure of incentives and authority within an organization. The Bell Journal of Economics, 105–131.

Montalvo, J. G., & Reynal-Querol, M. (2005). Ethnic polarization, potential conflict, and civil wars. American Economic Review, 95(3), 796–816.

Mundlak, Y. (1978). On the pooling of time series and cross section data. Econometrica: Journal of the Econometric Society, 69–85.

Myers, D. (2013). ASPM and VORP. Chosen Stats. Retrieved from `http://www.godismyjudgeok.com/DStats/aspm-and-vorp/`

Myers, D. (2014). About box plus/minus (BPM). Basketball Reference. Retrieved from `https://www.basketball-reference.com/about/bpm.html`

Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R2 from generalized linear mixed-effects models. Methods in Ecology and Evolution, 4(2), 133–142.

NBA Global. (2018). 2018 NBA playoffs feature record 62 international players. ABS-CBN Sports. Retrieved from `https://www.sports.abs-cbn.com/nba/news/2018/04/13/2018-nba-playoffs-feature-record-62-international-players-40250`

Neale, W. C. (1964). The peculiar economics of professional sports. The Quarterly Journal of Economics, 78(1), 1–14.

Oliver, D. (2004). Basketball on paper: rules and tools for performance analysis. Potomac Books, Inc.

Østby, G. (2008). Polarization, horizontal inequalities and violent civil conflict. Journal of Peace Research, 45(2), 143–162.

Paine, N. (2012). Statistical plus/minus. Basketball History. Retrieved from `http://www.bballhistory.wordpress.com/statistical-plusminus/`

Parker, C. (2015). How the NBA learned to love zone defence. The Guardian. Retrieved from `https://www.theguardian.com/sport/2015/nov/25/nba-learned-love-zone-defense`

Peel, D., & Thomas, D. (1988). Outcome uncertainty and the demand for football: an analysis of match attendances in the english football league. Scottish Journal of Political Economy, 35(3), 242–249.

Pfeffer, J., & Langton, N. (1993). The effect of wage dispersion on satisfaction, productivity, and working collaboratively: evidence from college and university faculty. Administrative Science Quarterly, 382–407.

Pierce, R. (2017). Salary dispersion and team performance in the National Basketball Association. Economics Student Theses and Capstone Projects. 51. Retrieved from `https://www.creativematter.skidmore.edu/econ_studt_schol/51/`

Piette, J., Anand, S., & Zhang, K. (2010). Scoring and shooting abilities of NBA players. Journal of Quantitative Analysis in Sports, 6(1).

Riva, M., Gauvin, L., & Barnett, T. A. (2007). Toward the next generation of research into small area effects on health: a synthesis of multilevel investigations published since July 1998. Journal of Epidemiology & Community Health, 61(10), 853–861.

Rodrîguez, G. (2019). Multilevel models. Princeton University. Retrieved from `https://data.princeton.edu/pop510/egm`

Rosen, S. (1981). The economics of superstars. The American Economic Review, 71(5), 845–858.

Rosen, S., & Sanderson, A. (2001). Labour markets in professional sports. The Economic Journal, 111(469), F47–F68.

Rosenbaum, D. T. (2004). Measuring how NBA players help their teams win. 82Games.com, 4–30. Retrieved from `http://www.82games.com/comm30.html`

Rottenberg, S. (1956). The baseball players' labor market. Journal of Political Economy, 64(3), 242–258.

Schofield, J. A. (1988). Production functions in the sports industry: an empirical analysis of professional cricket. Applied Economics, 20(2), 177–193.

Scott, F. A., Long, J. E., & Somppi, K. (1985). Salary vs. marginal revenue product under monopsony and competition: the case of professional basketball. Atlantic Economic Journal, 13(3), 50–59.

Scully, G. W. (1974). Pay and performance in Major League Baseball. The American Economic Review, 64(6), 915–930.

Scully, G. W. (1995). The market structure of sports. University of Chicago Press.

Sill, J. (2010). Improved NBA adjusted +/- using regularization and out-of-sample testing. In Proceedings of the 2010 MIT Sloan sports analytics conference.

Simmons, R., & Berri, D. J. (2011). Mixing the princes and the paupers: pay and performance in the National Basketball Association. Labour Economics, 18(3), 381–388.

Sloane, P. J. (1971). The economics of professional football: the football club as a utility maximiser. Scottish Journal of Political Economy, 18(2), 121–146.

Snijders, T. A., & Bosker, R. J. (1994). Modeled variance in two-level models. Sociological methods & research, 22(3), 342–363.

Statista. (2020). Average annual player salary in the sports industry in 2018/19, by league. Author. Retrieved from `https://www.statista.com/statistics/675120/average-sports-salaries-by-league/`

Stewart, F., Brown, G. K., & Langer, A. (2008). Policies towards horizontal inequalities. In Horizontal inequalities and conflict (pp. 301–325). Springer.

Stiroh, K. J. (2007). Playing for keeps: pay and performance in the NBA. Economic Inquiry, 45(1), 145–161.

Tikhonov, A. N., & Arsenin, V. I. (1977). Solutions of ill-posed problems (Vol. 14). Winston, Washington, DC.

Trevor, C. O., Reilly, G., & Gerhart, B. (2012). Reconsidering pay dispersion's effect on the performance of interdependent work: reconciling sorting and pay inequality. Academy of Management Journal, 55(3), 585–610.

United States Census Bureau. (2011). The foreign-born population in the United States. United States Census Bureau. Retrieved from `https://www.census.gov/newsroom/pdf/cspan_fb_slides.pdf`

United States Census Bureau. (2020). Foreign born: 2019 current population survey detailed tables. United States Census Bureau. Retrieved from https://www.census.gov/data/tables/2019/demo/foreign-born/cps-2019.html

Vasilescu, O. (2008). Winning and wage inequality in Major League Baseball. Available at SSRN 1099394. Retrieved from https://dx.doi.org/10.2139/ssrn.1099394

Wasserman, J. (2013). Unmasking the forefathers of advanced NBA stats. Bleacher Report. Retrieved from https://www.bleacherreport.com/articles/1728780-unmasking-the-forefathers-of-advanced-nba-stats

Wolfson, M. C. (1994). When inequalities diverge. The American Economic Review, 84(2), 353–358.

Zak, T. A., Huang, C. J., & Siegfried, J. J. (1979). Production efficiency: the case of professional basketball. Journal of Business, 379–392.

Zhang, X., & Kanbur, R. (2001). What difference do polarisation measures make? An application to China. Journal of Development Studies, 37(3), 85–98.

Zimbalist, A. S. (2002). Competitive balance in sports leagues: An introduction. Sage CA: Thousand Oaks, CA.

Appendices

4.A Alternative Estimates of Possession

The importance of considering possessions when determining performance in basketball is discussed in Chapter 2 and Section 4.3.1. While possessions statistics are not explicitly recorded by the NBA or other sources, they can be estimated using other statistics which are recorded. We primarily use the methodology of Berri et al. (2006) to calculate estimates of possessions employed ($tPOSSemp$) and possessions acquired ($tPOSSacq$), but also consider other methodologies summarised by Kubatko et al. (2007) for comparison. In our main analysis we use an average of $tPOSSemp$ and $tPOSSacq$, which we refer to as simply $tPOSS$.

Berri et al. (2006) note that possessions can be employed in three ways: through a field goal attempt ($tFGA$), a free throw attempt ($tFTA$) or by turning the ball over ($tTOV$). However, not all scoring attempts lead to a change in possession so offensive rebounds ($tORB$) must be subtracted. This logic gives Equation 4.A.1. Free throw attempts are multiplied by the proportion z, as multiple free throws are often awarded together, with only the final attempt resulting in open play being resumed if missed. z is therefore the proportion of free throws which are final attempts and would result in open play if missed. This can and subsequently will be estimated.

$$tPOSSemp = tFGA + z \cdot tFTA + tTOV - tORB \tag{4.A.1}$$

Possessions are acquired when the opponent successfully makes a field goal ($oFGS$) or free throw ($oFTS$), through defensive rebounds ($tDRB$) or team rebounds ($tmRB$), and by forcing the opponent to turnover the ball ($oTOV$). This gives Equation 4.A.2. Again a calculation is needed, this time to determine the proportion of opponent's free throw successes which are on the final attempt and therefore lead to a possession being acquired. Team rebounds ($tmRB$) must also be estimated as they are more an accounting device than an official statistic and as such are not recorded. They account for instances such as missed shots which go out of bounds and therefore result in a change of possession without a rebound being recorded.

$$tPOSSacq = oFGS + x \cdot oFTS + oTOV + tDRB + tmRB \tag{4.A.2}$$

So, to estimate possessions, we must assigned values to z, x and $tmRB$. Fortunately, all three can be estimated from a single regression. Berri et al. (2006) give a more complex and in-depth explanation for this[70], but we noticed that the same outcome can be achieved by setting $tPOSSemp$ equal to $tPOSSacq$, then rearranging the known statistics to the left-hand side and those containing unknown parameters to the right, as shown in Equation 4.A.3. Berri et al. (2006) call the summation of the left-hand side variable Field Goal Attempt Difference ($FGAD$), due to their alternative reasoning.

$$\begin{aligned} tFGA - oFGS - oTOV - tORB - tDRB + tTOV &= tmRB + x \cdot oFTS - z \cdot tFTA \\ &= FGAD \end{aligned} \tag{4.A.3}$$

We can then regress $FGAD$ on $oFTS$ and $tFTA$ to give x and z respectively, while the constant plus the error term for each team will give $tmRB$. These results are presented in Table

[70] This additional reasoning is reassuring in that the equality of $tPOSSemp$ and $tPOSSacq$ is not just by assumption. One of the nice features of this methodology is that this equality is a result rather assumption.

4.A.1. Imputing these values, we find that on average teams employ and acquire around 96.5 possessions a game[71], as shown in Table 4.A.2.

Table 4.A.1: Determining x, z and tmRB

	(1) FGAD
oFTS	0.456***
	(0.0157)
tFTA	-0.429***
	(0.0116)
Constant	4.571***
	(0.326)
N	330
R^2	0.840

Standard errors in parentheses
$^{+}$ $p < 0.20$, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Other methods of estimating possessions take historical actual values of possessions which have been tediously counted, then regress known statistics to find their marginal effects on possessions. Slight variations of similar equations are used to estimate possessions by different authors (e.g. Hollinger, 2005; Oliver, 2004); but Kubatko et al. (2007) show that these equations produce similar results (minimum 0.95 correlation with actual possessions using their data) and they present a generalised equation without imposing the restrictions of other authors, estimated via regression analysis.

Kubatko et al. (2007) use data from 2002-2006 to estimate the regression coefficients for the effects of each statistic on possessions. We use the coefficients found to create an estimate of possessions ($tPOSSa$) as presented in Equation 4.A.4, where the joint prefix $to-$ represents the average of a team and its opponents. Kubatko et al. (2007) note that taking the average in this way seems to be advantageous when estimating possessions, and improved the power of the regression model to explain in actual possessions.

$$\begin{aligned}\widehat{tPOSSa} = 3.2258 + 0.964toFGA - 0.3452toFGM + 0.4637toFTA - 0.2073toFTM \\ - 0.6227toORB + 0.3643otDRB + 0.9767toTOV\end{aligned} \tag{4.A.4}$$

A field goal attempt, which while holding misses constant is a successful attempt, almost always results in a new possession (with the exception of flagrant fouls) hence why the coefficient on $toFGA$ in Equation 4.A.4 is close to one. A missed field goal will increase both attempts and misses, so overall will increase possessions on average, though less then successful field goals ($0.964 > 0.6188$). This is because of the possibility of an offensive rebound extending the possession. For a full discussion of the other variables and regression methodology see Kubatko et al. (2007).

[71] Though the overall averages are identical very slight differences can occur between possessions employed and acquired for individual teams due to the parameters estimates. But these differences are small, with the largest difference occurring on the sixth decimal place.

We also estimate possessions with two of the additional approaches suggested by Kubatko et al. (2007) which are presented in Equations 4.A.5 and 4.A.6.

$$\widehat{tPOSSb} = \frac{(toFGA + 0.44 \cdot toFTA - toORB + toTOV)}{2} \tag{4.A.5}$$

$$\widehat{tPOSSc} = \frac{toFGA + 0.4 \cdot toFTA - 1.07 \cdot \left(\frac{tORB \cdot (tFGA - tFGM)}{tORB + oDRB}\right) + tTOV}{2} \tag{4.A.6}$$

Overall, we get similar estimates of possessions using each of these methodologies. Table 4.A.2 shows that all the measures created agree that on average teams had around 96-97 possessions per game in our sample. The standard deviations, maxima and minima of the measures are also similar.

Table 4.A.2: Alternative Estimates of Possessions

	Mean	Std Dev.	Min.	Max.
tPOSSemp	96.527	3.045	89.544	103.529
tPOSSacq	96.527	3.045	89.544	103.529
tPOSSa	96.120	3.085	88.633	104.753
tPOSSb	96.772	3.041	89.908	104.170
tPOSSc	96.373	2.978	89.097	103.603

4.B Wins Produced (WP) - (Berri et al., 2006)

The methodology for calculating Wins Produced (WP), as developed by Berri et al. (2006), begins with a simple regression of Offensive Rating ($ORTG$) and Defensive Rating ($DRTG$) on winning percentage (WL) to determine the contribution of offence and defence, respectively. This equation is shown in Equation 4.B.1.

$$WL = \beta_0 + \beta_1 ORTG + \beta_2 DRTG + \epsilon \tag{4.B.1}$$

The results of running this regression for our dataset are presented in Table 4.B.1. Note that $DRTG$ has a positive coefficient as we multiply opponents points by minus one, so that higher values of $DRTG$ indicate better performance. Other studies do not make this adjustment, so have a negative coefficient for $DRTG$. The ratings explain almost 95% of the variation in WL. The coefficients of both ratings are of a similar magnitude which suggests that offensive and defensive performance both contribute towards performance fairly equally.

Table 4.B.1: Estimating Win Percentage using Offensive and Defensive Ratings

	(1) WL
ORTG	3.140***
	(0.0667)
DRTG	3.091***
	(0.0865)
Constant	0.448***
	(0.120)
N	330
R^2	0.945

Standard errors in parentheses
** $p < 0.05$, *** $p < 0.01$

Equation 4.B.2 shows the regression equation again but with the formulae for $ORTG$ and $DRTG$ substituted in.

$$WL = \beta_0 + \beta_1 \frac{tPTS}{tPOSSemp} + \beta_2 \frac{-oPTS}{tPOSSacq} + \epsilon \tag{4.B.2}$$

This helps with the next step which is to derive the marginal effect of an additional point scored, point allowed, possession employed and possession acquired. These are shown in Equations 4.B.3, 4.B.4, 4.B.5 and 4.B.6 respectively.

$$\frac{\partial WL}{\partial tPTS} = \beta_1 \frac{1}{avtPOSSemp} \tag{4.B.3}$$

$$\frac{\partial WL}{\partial oPTS} = \beta_2 \frac{-1}{avtPOSSacq} \tag{4.B.4}$$

$$\frac{\partial WL}{\partial tPOSSemp} = \beta_1 \frac{-avtPTS}{avtPOSSemp^2} \tag{4.B.5}$$

$$\frac{\partial WL}{\partial tPOSSacq} = \beta_2 \frac{avoPTS}{avtPOSSacq^2} \tag{4.B.6}$$

We can then calculate these marginal effects using our data and use them to find the marginal effects of each statistic, as described in Section 4.3.3. For completeness, Equation 4.B.7 shows the regression equation again, with the formulae for $tPOSSemp$ and $tPOSSacq$ additionally substituted in.

$$\begin{aligned} WL = \beta_0 + \beta_1 \frac{tPTS}{tFGA + 0.429 \cdot tFTA + tTOV - tORB} + \\ \beta_2 \frac{oPTS}{oFGS + 0.456 \cdot oFTS + oTOV + tDRB + tmRB} + \epsilon \end{aligned} \tag{4.B.7}$$

From Equation 4.B.7, the marginal effects of the different statistics can be calculated. Equations 4.B.8 to 4.B.15 present these marginal effects for the scoring statistics. For field goal successes, these are simply the number of points scored/allowed multiplied by the marginal effect of a point scored/allowed, plus the marginal effect on possessions employed/acquired. For free throws, the effect on possessions is not one-for-one, as shown by the coefficients in each denominator in Equation 4.B.7, but the method is the same. Misses are the same except no points are scored, so we only have the effect on possessions.

$$\frac{\partial WL}{\partial t3PS} = \beta_1 \left(\frac{3}{avtPOSSemp} - \frac{avtPTS}{avtPOSSemp^2} \right) \tag{4.B.8}$$

$$\frac{\partial WL}{\partial o3PS} = \beta_2 \left(\frac{3}{avtPOSSacq} - \frac{avoPTS}{avtPOSSacq^2} \right) \tag{4.B.9}$$

$$\frac{\partial WL}{\partial t2PS} = \beta_1 \left(\frac{2}{avtPOSSemp} - \frac{avtPTS}{avtPOSSemp^2} \right) \tag{4.B.10}$$

$$\frac{\partial WL}{\partial o2PS} = \beta_2 \left(\frac{2}{avtPOSSacq} - \frac{avoPTS}{avtPOSSacq^2} \right) \tag{4.B.11}$$

$$\frac{\partial WL}{\partial tFTS} = \beta_1 \left(\frac{1}{avtPOSSemp} - 0.429 \cdot \frac{avtPTS}{avtPOSSemp^2} \right) \tag{4.B.12}$$

$$\frac{\partial WL}{\partial oFTS} = \beta_2 \left(\frac{1}{avtPOSSacq} - 0.456 \cdot \frac{avoPTS}{avtPOSSacq^2} \right) \tag{4.B.13}$$

$$\frac{\partial WL}{\partial tFGM} = \beta_1 \left(-\frac{avtPTS}{avtPOSSemp^2} \right) \tag{4.B.14}$$

$$\frac{\partial WL}{\partial tFTM} = \beta_1 \left(-0.429 \cdot \frac{avtPTS}{avtPOSSemp^2} \right) \tag{4.B.15}$$

The possession statistics are presented in Equations 4.B.16 to 4.B.22. Each of these are explicitly part of Equation 4.B.7 except $tSTL$, but this is an opponent turnover ($oTOV$). For clarity, $tmTOV$ is part of $tTOV$. So, mathematically, the calculation of these marginal effects

should be clear.

$$\frac{\partial WL}{\partial tORB} = \beta_1 \frac{avtPTS}{avtPOSSemp^2} \tag{4.B.16}$$

$$\frac{\partial WL}{\partial tTOV} = \beta_1 \frac{-avtPTS}{avtPOSSemp^2} \tag{4.B.17}$$

$$\frac{\partial WL}{\partial tmTOV} = \beta_1 \frac{-avtPTS}{avtPOSSemp^2} \tag{4.B.18}$$

$$\frac{\partial WL}{\partial tDRB} = \beta_2 \frac{avoPTS}{avtPOSSacq^2} \tag{4.B.19}$$

$$\frac{\partial WL}{\partial tmRB} = \beta_2 \frac{avoPTS}{avtPOSSacq^2} \tag{4.B.20}$$

$$\frac{\partial WL}{\partial oTOV} = \beta_2 \frac{avoPTS}{avtPOSSacq^2} \tag{4.B.21}$$

$$\frac{\partial WL}{\partial tSTL} = \beta_2 \frac{avoPTS}{avtPOSSacq^2} \tag{4.B.22}$$

The crimes and rejections element of performance takes a little more work, as tPF and $tBLK$ are not explicitly part of the regression equation. However, the marginal effect of tPF can be found by estimating the effects of personal fouls on opponents' free throw successes ($oFTS$). This is presented in Table 4.B.2, while controlling for team and year effects.

Table 4.B.2: Estimating the Marginal Effect of Personal Fouls

	(1) oFTS
tPF	1.079***
	(0.0335)
Constant	-3.917***
	(0.698)
N	330
R^2	0.885
Year	✓
Team	✓

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The marginal effect of a personal foul is then the marginal effect of personal fouls on opponents' free throw successes, multiplied by the marginal effect of opponents' free throws on win percentage. This is displayed in Equation 4.B.23. When analysing overall performance, the marginal effect of tPF will be included in place of $oFTS$.

$$\frac{\partial WL}{\partial tPF} = 1.079 \cdot \beta_2 \left(\frac{1}{avtPOSSacq} - 0.456 \cdot \frac{avoPTS}{avtPOSSacq^2} \right) \tag{4.B.23}$$

The marginal effect of blocks is yet one more step more complex, as blocks cannot replace another element[72], as tPF does to $oFTS$. Firstly, to find the marginal effect of a block, we regress opponents' two point successes on blocks, while holding attempts constant. The results of this regression are presented in Table 4.B.3.

Table 4.B.3: Estimating the Marginal Effect of Blocks

	(1) o2PS
o2PA	0.491***
	(0.0191)
tBLK	-0.614***
	(0.0729)
Constant	2.253*
	(1.182)
N	330
R^2	0.806
Year	✓
Team	✓

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The coefficient on $tBLK$ in Table 4.B.3 suggests that, for a given number of attempts, each block reduces the number of two point successes[73] an opponent has by 0.614. The marginal effect of a block is then this value multiplied by the marginal effect of $o2PS$ as shown in Equation 4.B.24.

$$\frac{\partial WL}{\partial tBLK} = -0.614 \cdot \beta_2 \left(\frac{2}{avtPOSSacq} - \frac{avoPTS}{avtPOSSacq^2} \right) \tag{4.B.24}$$

As blocks are not explicitly in the regression equation, and do not replace an existing element, they are not considered at the team-level. At the player-level, they are assigned by the number of blocks a player produces compared to the team average (or, more precisely, compared to the proportion of the team blocks that the player would be expected to make given his minutes played). Therefore considering blocks does not alter the estimate of team performance, but simply reallocates production amongst team members.

A similar approach is taken for assists. At the team-level, once scoring is taken into account, assists do not add any additional value. However, at the player-level it seems appropriate that players are credited for their assists. Berri et al. (2006) find that the marginal effect of an assist is 0.022 wins. This is found by regressing the productivity of players without considering assists on a number of factors which are calculated throughout their book, including

[72] Opponents misses have not been explicitly mentioned or included as they are implicitly included through defensive rebounds. Opponents misses which are recovered via an opponents offensive rebounds simply extend the possession and do not affect the team through points or possessions.

[73] As noted by Berri et al. (2006), while blocks do not just occur for two attempts, information on the proportion which do is unknown. Therefore an assumption is made that all blocks apply to two point attempts.

the player's previous performance, the coaches ability and the stability of the roster. As we do not have information on all of these factors, we use the value of 0.022 calculated by Berri et al. (2006) to reallocate production to credit assists.

Putting all this together, we then allocate these marginal effects on a player-by-player basis where possible. This can be done for statistics which can be attributed to individual players which are $p3PS$, $p2PS$, $pFTS$, $pFGM$, $pFTM$, $pORB$, $pTOV$, $pDRB$, $pSTL$, pPF; and then $pBLK$ and AST as described. The remaining statistics which cannot be attributed at the player-level are $o3PS$, $o2PS$, $tmTOV$, $tmRBD$ and $oTOV$ ($-tSTL$). These are attributed by the proportion of the teams' minutes that each player has played. Once apportioned, these statistics are multiplied by their marginal effects, summed together, and then the positional average is subtracted to account for positional biases[74], which then give us the marginal effect of each player. The final step, to convert this to a measure of absolute performance, is to add the average number of wins that each player would be expected to produce given their minutes played. As each game will result in one winner and one loser, on average teams will win 50% of their matches. In an 82 game season, this would mean 41 wins. As five players are on the court at any one time, a player who plays all the possible minutes would be expected to produce 8.2 wins. Therefore, the expected wins of a player can be calculated based on the proportion of minutes played.

This section is designed to give an overview of the construction of the WP measure calculated for this investigation. A more comprehensive overview is given by Berri et al. (2006). A final adjustment made for this investigation is to normalise the number of wins produced to account for the number of games played in a season. To do this, we divide by the number of games played and multiply by 82. This will cause no change in the nine seasons in which each team played 82 games, but adjust the 2011-12 season, as well as the Indiana Pacers and Boston Celtics in the 2012-13 season, to account for the number of games played.

[74] This is to account for the differing roles of each position on the court. Each role is necessary and therefore each player's productivity is calculated compared to the average at their position. When a player plays more than one position, their production is split between the positions based on the minutes played in each.

4.C Player Efficiency Rating (PER) - (Hollinger, 2002, 2005)

The Player Efficiency Rating (PER) is a per-minute performance measure developed by Hollinger (2002, 2005) who is an ESPN columnist. It is a linear weights measure, though its formulation is more complex than the NBA Efficiency measure, for example. It is a commonly used measure, but has also faced some criticism (see: Berri & Bradbury, 2010; Berri et al., 2006; Hoffer & Freidel, 2014; Khan, 2013; Piette et al., 2010). One such criticism is that it focuses on offensive performance and only accounts for two defensive statistics: blocks and steals. As an efficiency measure it is also prone to give arguably misleading values for those players which play very few minutes, as discussed in Section 4.3.3.

Equation 4.C.1 shows the formula for calculating Unadjusted Player Efficiency Ratings ($uPER$). The Equations in 4.C.2 are separated out to simply the formula. The notation of the statistics has been adjusted to be in line with the rest of this investigation. Additionally, the prefix $lg-$ represents the league average. The weights in the equation are more subjective rather than having robust statistical backing. For example, assists ($pAST$) are multiplied by $\frac{2}{3}$, as Hollinger (2002) argues that a scoring play involving an assist requires a player getting open, the pass being made and the shot being successful. As the player making the assist is responsible for one of these three actions, he is attributed a third of the two points, or $\frac{2}{3}$, for each assist. Perhaps coincidently, Berri et al. (2006) find an assist to worth roughly the same: 0.22 wins or 0.67 points.

$$\begin{aligned} uPER = & \frac{1}{pMP} \cdot \Bigg[p3PS + \frac{2}{3} \cdot pAST \\ & + \left(2 - factor \cdot \frac{tAST}{tFGS}\right) \cdot pFGS \\ & + \left(pFTS \cdot 0.5 \cdot \left(1 + \left(1 - \frac{tAST}{tFGS}\right)\right) + \frac{2}{3} \cdot \frac{tAST}{tFGS}\right) \\ & - VOP \cdot pTOV \\ & - VOP \cdot pDRB\% \cdot (pFGA - pFGS) \\ & - VOP \cdot 0.44 \cdot (0.44 + (0.56 \cdot pDRB\%)) \cdot (pFTA - pFTS) \\ & + VOP \cdot (1 - pDRB\% \cdot pORB) \\ & + VOP \cdot pSTL \\ & + VOP \cdot pDRB\% \cdot pBLK \\ & - PF \cdot \left(\frac{lgFT}{lgPF} - 0.44 \cdot \frac{lgFTA}{lgPFS} \cdot VOP\right) \Bigg] \end{aligned} \tag{4.C.1}$$

Where

$$\begin{aligned} factor &= \frac{2}{3} - \frac{0.5 \cdot \frac{lgAST}{lgFGS}}{2 \cdot \frac{lgFGS}{lgFT}} \\ VOP &= \frac{lgPTS}{lgFGA - lgORB + lgTOV + 0.44 \cdot lgFTA} \\ DRB\% &= \frac{lgTRB - lgORB}{lgTRB} \end{aligned} \tag{4.C.2}$$

Once $uPER$ has been calculated, it is adjusted to account for the pace that the team plays at to create Adjusted Player Efficiency Rating ($aPER$). This is shown in Equation 4.C.3 where $uPER$ is multiplied by the ratio of the league pace ($lgPace$) over the team's pace ($tPace$).

Equation 4.C.4 shows that these pace measurements are dependent on possessions.

$$aPER = \frac{lgPace}{tPace} \cdot uPER \tag{4.C.3}$$

Where

$$Pace = 48 \cdot \left(\frac{tPoss + oPoss}{2 \cdot (tMP/5)} \right) \tag{4.C.4}$$

A final adjustment is made, as shown in Equation 4.C.5, to normalise the measure to the league average of the measure to allow for cross-season comparisons. This last step also, famously, means that a value of 15 represents an average player.

$$PER = aPER \cdot \left(\frac{15}{lgaPER} \right) \tag{4.C.5}$$

4.D Box Plus/Minus (BPM) - (Myers, 2014)

As discussed in Section 4.3.3, BPM is calculated based on regression coefficients of individual performance statistics obtained from a regression on historical *Adjusted Plus/Minus* statistics. These statistics are known as *Ridge-Regressed Adjusted Plus/Minus* ($RAPM$) after the methodology used, and are from the period 2001-14. The regression equation is given in Equation 4.D.1.

$$\begin{aligned} RAPM = {} & a_1 ReMPG + a_2 pORB\% + a_3 pDRB\% + a_4 pSTL\% + a_5 pBLK\% + a_6 pAST\% \\ & - a_7 pUSG\% \cdot pTOV\% + a_8 pUSG\%(1 - pTOV\%)[2(pTS\% - tTS\%) + a_9 pAST\% \\ & + a_{10}(3PAr - Lg3PAr) - a_{11}] + a_{12}\sqrt{AST\% \cdot TRB\%} + \epsilon \end{aligned} \tag{4.D.1}$$

Where $ReMPG$ is MPG with 4 games of 0 MPG added to account for players who only played a few games with high MPG, perhaps at the end of the season while a playoff team is resting its stars. BPM is then calculated from these regression coefficients as shown in Equation 4.D.2.

$$\begin{aligned} RAPM = {} & 0.123 ReMPG + 0.120 pORB\% - 0.151 pDRB\% + 1.256 pSTL\% + 0.532 pBLK\% \\ & - 0.306 pAST\% - 0.921 pUSG\% \cdot pTOV\% + 0.711 pUSG\%(1 - pTOV\%) \\ & [2(pTS\% - tTS\%) + 0.017 pAST\% + 0.298(3PAr - Lg3PAr) - 0.213] + \\ & 0.726\sqrt{AST\% \cdot TRB\%} \end{aligned} \tag{4.D.2}$$

The BPM measure is designed to capture the points contributed by each player to their team. For further detail on the measure, see Myers (2014).

4.E The Best Players in the NBA

Table 4.E.1: The Best 5 Players by each Player Performance Measure

Year	Player	Team	Pos.	Wage ($m)	MP	WPRank	WSRank	VORPRank	WP48Rank	PERRank	BPMRank	Av. Rank
2014-15	DeAndre Jordan	LAC	C	11.440	2820	1	54	121	41	351	215	130.5
2010-11	Dwight Howard	ORL	C	16.647	2935	2	25	76	62	90	149	67.333
2016-17	DeAndre Jordan	LAC	C	21.166	2570	3	74	153	40	272	215	126.167
2010-11	Kevin Love	MIN	PF	3.638	2611	4	85	164	47	143	276	119.833
2007-08	Dwight Howard	ORL	C	6.061	3088	5	50	233	91	200	631	201.667
2012-13	Kevin Durant	OKC	SF	16.670	3119	16	4	20	158	57	48	50.5
2013-14	Kevin Durant	OKC	SF	17.833	3122	24	3	11	229	45	29	56.833
2008-09	LeBron James	CLE	SF	14.411	3054	27	1	2	219	28	9	47.667
2016-17	Russell Westbrook	OKC	PG	26.540	2802	31	47	1	175	41	6	50.167
2008-09	Chris Paul	NOP	PG	4.574	3002	37	6	5	246	44	13	58.5
2012-13	LeBron James	MIA	SF	17.545	2877	49	2	6	259	29	12	59.5
2009-10	LeBron James	CLE	SF	15.780	2966	53	5	3	297	34	10	67
2007-08	LeBron James	CLE	SF	13.041	3027	127	18	4	645	51	13	143
2008-09	Steven Hill	OKC	PF	0.148	2	3856	4833	2945	3	3	64	1950.667
2007-08	Gerald Green	HOU	SF	1.441	4	3870	4536	2945	4	4	5	1894
2017-18	Georgios Papagiannis	POR	C	0.185	4	3898	4536	2945	8	5	2	1899
2013-14	DeAndre Liggins	MIA	SG	0.078	1	4019	4833	2945	1	2	4	1967.333
2017-18	Naz Long	UTA	SG	0.167	1	4022	4536	2945	2	1	1	1917.833
2014-15	Sim Bhullar	SAC	SF	0.030	3	4121	4833	2945	39	9	3	1991.667
2016-17	Danuel House	WAS	SG	0.543	1	4128	4833	2945	5	3223	5887	3503.5

4.F Importance of Team-level Controls in Determining Team Performance

Table 4.F.1 shows our preferred specification of the control variables. Models (1) - (6) show the coefficients when running this specification for our six different performance measures: win percentage (WL), the natural logarithm of win percentage ($lnWL$), points difference per game (PDG), a simple ratings system (SRS), offensive rating ($ORTG$), and defensive rating ($DRTG$). As we would expect, the results suggest that teams with larger payrolls perform better than those which do not spend as much on wages. What might be unexpected is relatively low significance shown for $lnTPayroll$, only at a 5% level in columns (1) and (2); 10% in columns (3),(4) and (6); and insignificant in column (5). However, as discussed in Section 4.4.1, Table 4.9 shows significant correlation between $TPayroll$ and the other control variables, especially experience ($TExp$) and the standard deviation of experience ($TExpSD$). When a simple regression of payroll and team performance is run, the coefficient is positive and significant at a 1% level in all specifications as would be expected. The coefficients estimated are also much larger for each of the performance measures.

Table 4.F.1: Team-level Control Variables and Performance

	(1) WL	(2) lnWL	(3) PDG	(4) SRS	(5) ORTG	(6) DRTG
lnTPayroll	0.163**	0.457**	4.045*	4.024*	2.070	2.091*
	(0.0739)	(0.170)	(2.087)	(2.001)	(1.665)	(1.179)
lnCTSize	-0.377***	-0.862***	-11.30***	-10.99***	-8.208***	-3.481***
	(0.0436)	(0.0909)	(1.293)	(1.267)	(1.051)	(0.919)
TTenureMP	0.0496**	0.116***	1.656**	1.603**	1.731***	-0.0102
	(0.0191)	(0.0404)	(0.616)	(0.592)	(0.436)	(0.444)
TTenureMP^2	-0.00677***	-0.0165***	-0.210**	-0.204**	-0.209***	-0.0133
	(0.00209)	(0.00449)	(0.0768)	(0.0752)	(0.0555)	(0.0611)
TExp	0.0539***	0.117***	1.549***	1.474***	0.821***	0.795***
	(0.0102)	(0.0225)	(0.360)	(0.346)	(0.241)	(0.246)
TExpSD	-0.0394***	-0.0936***	-1.259**	-1.234***	-0.704**	-0.613*
	(0.0137)	(0.0319)	(0.461)	(0.440)	(0.282)	(0.304)
Constant	0.522*	-1.016	4.475	4.184	111.7***	-106.9***
	(0.304)	(0.692)	(7.898)	(7.535)	(6.413)	(4.515)
R^2	0.471	0.446	0.439	0.438	0.506	0.388
Years	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓

N = 330
Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Therefore, the explanation here is one of multicollinearity between the control variables which is causing the lack of statistical significance. An interpretation of this is that if additional wages are spent, but are not used to keep experienced, non-injury-prone players together playing within the same team, then these wages are inefficiently spent and have less of an effect on performance. However, the aim here is to specify a reasonable set of controls as a baseline model, not to pin down the exact marginal effects of each of these controls. Therefore, we do

not find this multicollinearity between the control variables a problem and believe it is still relevant to include them as control variables.

In terms of the other control variables, there were few surprises in the coefficients found. We find that teams with a larger core team perform worse. This corresponds to our theory that the size of the core team indicates the number of injuries within the team, or otherwise how often the team can consistently field their preferred line-up. If there is any insight to be gained here, it is that deliberate squad rotation is not a beneficial strategy, and there appears to be advantages to fielding consistent teams. It should be noted that there was no evidence of a quadratic relationship here, but a logarithmic one.

The effect of tenure in minutes played is a little more unexpected. Not so much for the positive coefficient on the linear term and negative coefficient on the quadratic term, but due to the turning point. We find the turning point to be between 3,523-4,140 minutes across the specifications presented, which is near the middle of the range of 326 – 8,014 minutes of average team tenure within our sample. The interpretation with experience held constant is a little different, but there is no evidence for interacting the terms; and even when excluding experience and/or the standard deviation of experience, the turning point only increases to just over 5,000 minutes for win percentage. When excluding all the other controls and regressing win percentage on $TTenureMP$ and $TTenureMP2$, controlling only for year and team effects, the turning point increases to 6,828 minutes. Only three teams have a longer average tenure than this. Therefore, this result may suggest that there are some negative effects of a team remaining together for too long and some teams may exceed this limit. However, as increases in tenure are often accompanied by increases in experience, for example, this relationship is likely to be one which increases but at a decreasing rate for most of our sample range.

We believe that experience, and particularly the standard deviation of experience, are potentially important controls to isolate the effects of wage dispersion, due to backloaded pay schedules. Despite the positive correlation between payroll ($TPayroll$) and average team experience ($TExp$), displayed in Table 4.9, we find that higher experience has positive effects on performance, even when holding wage expenditure constant. In terms of the standard deviation of experience, we find this to have a negative effect on performance, when holding the other control variables constant. This suggests that a mix of youth and experience may not be a beneficial strategy, but that teams comprised of similarly experienced players perform best. Furthermore, there is no evidence in favour of including an interaction between experience and the standard deviation of experience. As our main analysis finds polarisation to have a statistically significant effect on performance, we also consider including the polarisation of experience, but find this variable to be insignificant.

In addition, we also test the joint significance of our year and team effects for our main measure of team performance, WL. We find the team effects to have a F-statistic of 1487.08, which makes them jointly significant at a 1% level. The year effects are jointly significant at a 5% level, with a F-statistic of 21.25. These results confirm that it is appropriate to control for team and year effects to establish the relationship between our control variables and team performance. Furthermore, Appendix 4.G additionally conducts a Hausman-type test using the Mundlak (1978) method which gives evidence to support our use of Fixed Effects.

The regression results presented in Appendix 4.H show that, in addition to significantly af-

fecting team performance, these control variables also significantly influence the levels of wage dispersion and national diversity within teams. This is done by regressing the main variables of interest from Chapter 5 (the standard deviation of wage, $WageSD$, and the polarisation of wage, $WagePol$) and Chapter 6 (national fractionalisation, $NatFrac2$ and national polarisation, $NatPol2$) on these control variables. This justifies their inclusion in our analysis to reduce the size of the error terms and prevent any bias in the estimated effects of our variables of interest. We believe these controls set a good baseline model to investigate the effects of both wage dispersion and national diversity on team performance.

4.G 'Hausman-type' Test - Control Variables

We conduct a 'Hausman-type' test using the Mundlak (1978) approach (which is valid for robust, clustered standard errors) to determine whether a Fixed-Effects (FE) or Random-Effects (RE) regression should be preferred. The same specification as in column (1) of Tables 4.F.1 is run, but this time using RE and additionally including the team-specific means of the variables.

Table 4.G.1: Hausman-Type Test using the Mundlak (1978) Approach - Team-level Control Variables

	(1) WL
m_lnTPayroll	0.111 (0.195)
m_lnCTSize	-0.313 (0.229)
m_TTenureMP	0.0684 (0.0672)
m_TTenureMP^2	-0.00553 (0.0101)
m_TExp	-0.0532** (0.0261)
m_TExpSD	0.0636 (0.0558)
R^2	0.481
Years	✓
Teams	✗
Controls	✓

N = 330
Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5.A.1 reports just the coefficients of the team-specific means, as these are of interest for this test and the coefficients of the other variables are identical to the FE estimates. If these team-specific means are jointly significant, then there is evidence of correlation between time-invariant, team-specific heterogeneity and the regressors, so we should favour FE. An F-test of joint significance returns a test statistic of 14.37, and a p-value of 0.0258. We therefore have evidence to favour FE regressions significant at a 5% level.

4.H Importance of Team-level Controls in Determining Wage Dispersion and National Diversity

Having established that these team-level control variables significantly affect team performance, it is relevant to test whether they also significantly influence wage dispersion and national diversity. Table 4.H.1 shows the results of regressing each of our main variables of interest: the standard deviation of wage in column (1), wage polarisation in (2), national fractionalisation in (3) and National Polarisation in (4); on the control variables, while holding the other relevant variables of interest constant. The p-values of joint significance are reported in the statistics of Table 4.H.1.

Table 4.H.1: The Importance of Team-level Control Variables

	(1) WageSD	(2) WagePol	(3) NatFrac2	(4) NatPol2
lnPayroll	3.975***	0.0431	-0.0298**	0.0314*
	(0.436)	(0.0411)	(0.0127)	(0.0168)
lnCTSize	-0.463	0.0119	0.00745	-0.0176
	(0.273)	(0.0274)	(0.0102)	(0.0148)
TTenureMP	0.573***	-0.0113	-0.00605	0.00575
	(0.150)	(0.0111)	(0.00650)	(0.00763)
TTenureMP^2	-0.0505***	-0.000824	0.00149	-0.00154
	(0.0153)	(0.00174)	(0.00105)	(0.00130)
TExp	0.0952	-0.0188***	-0.00447*	0.00491*
	(0.0657)	(0.00541)	(0.00221)	(0.00288)
TExpSD	0.0221	0.0244***	0.00482*	-0.00489
	(0.104)	(0.00716)	(0.00242)	(0.00295)
WagePol	4.326***			
	(0.873)			
WageSD		0.0237***		
		(0.00393)		
NatPol2			0.728***	
			(0.0218)	
NatFrac2				1.208***
				(0.0414)
Constant	-137.9***	-0.0775	0.110*	-0.0652
	(19.28)	(0.161)	(0.0597)	(0.0816)
R^2	0.726	0.280	0.886	0.884
Years	0.1684	0.2465	0.0009	0.0005
Teams	0.0000	0.0000	0.0000	0.0000

N = 330
Robust standard errors clustered by Team in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The results show that all our control variables are statistically significant determinants of at least one of the measures of wage dispersion and/or national diversity. All except $lnCTSize$, though we choose to retain it as a control, due to its strong relationship with performance. Additionally, we find the year effects to be jointly insignificant in determining the wage dispersion variables, but significant at a 1% level in determining $NatFrac2$ and $NatPol2$. The

team effects are jointly highly significant in all models.

When these variables affect both team performance and wage dispersion/national diversity, then excluding them from our analysis has the potential to bias our estimates. If they affect only team performance, but not wage dispersion/national diversity, then excluding them will not bias our estimates, but including them can reduce the size of the error term and increase the precision of our estimates. We therefore elect to retain all the controls for our main analysis of the effects of wage dispersion and national diversity on team performance. We believe these control variables set a good baseline model to investigate the effects of wage dispersion and national diversity on team performance. The resultant coefficients seem reasonable and match theoretical predictions. Controlling for these variables allows us to hold constant factors which may be correlated with wage dispersion and national diversity, and help isolate the relationships between these and team performance.

4.I Importance of Player-level Controls in Determining Player Performance

This section displays the importance of the control variables used in the regression analysis of Chapter 7 in determining player performance, which includes both individual- and team-level variables. Due to the interdependence of production, we argue that it is not only the characteristics of the player themselves which affects their performance, but also the characteristics of the team in which the players compete. The team-level controls used are $lnTPayroll$, $lnCTSize$, $TTenureMP$, $TExp$ and $TExpSD$. These are the ones used in Chapters 5 and 6, and explored further in Appendices 4.F and 4.H. This section therefore focuses on the player-level control variables: $lnWage$, $Trade$, $PTenureMP$ and $PExp$.

Table 4.I.1: Player-level Control Variables and Performance

	(1) WP	(2) WS	(3) VORP	(4) WP48	(5) PER	(6) BPM
lnWage	0.354***	0.405***	0.0789***	0.0116***	0.797***	0.812***
	(0.0384)	(0.0257)	(0.0115)	(0.00357)	(0.119)	(0.0764)
Trade	-1.414***	-1.437***	-0.371***	-0.0197***	-0.874***	-1.019***
	(0.0788)	(0.0588)	(0.0264)	(0.00555)	(0.191)	(0.125)
PTenureMP	0.0920***	0.0913***	0.0449***	0.000258	0.108***	0.0139
	(0.0188)	(0.0113)	(0.00623)	(0.000668)	(0.0246)	(0.0163)
PExp	0.336***	0.306***	0.156***	0.00582**	0.265***	0.239***
	(0.0405)	(0.0300)	(0.0135)	(0.00228)	(0.0737)	(0.0533)
$PExp^2$	-0.0339***	-0.0304***	-0.0142***	-0.000742***	-0.0377***	-0.0237***
	(0.00324)	(0.00237)	(0.00110)	(0.000153)	(0.00499)	(0.00364)
lnTPayroll	0.272	0.321	0.179*	0.00396	-0.494	-0.335
	(0.363)	(0.246)	(0.105)	(0.0195)	(0.581)	(0.411)
lnCTSize	-1.616***	-1.636***	-0.653***	-0.00221	0.832*	-0.0623
	(0.322)	(0.221)	(0.0957)	(0.0139)	(0.462)	(0.325)
TTenureMP	0.0398	0.0459	0.0181	0.00708	0.174	0.227*
	(0.111)	(0.0767)	(0.0360)	(0.00492)	(0.177)	(0.128)
$TTenureMP^2$	0.00114	-0.00471	-0.00212	-0.0000883	-0.00834	-0.0122
	(0.0177)	(0.0122)	(0.00584)	(0.000669)	(0.0244)	(0.0181)
TExp	0.0752	0.0926**	0.0306*	0.00221	0.0440	0.108*
	(0.0499)	(0.0366)	(0.0160)	(0.00279)	(0.0945)	(0.0636)
TExpSD	-0.0145	-0.0361	-0.0257	-0.00351	-0.120	-0.172**
	(0.0684)	(0.0472)	(0.0214)	(0.00385)	(0.127)	(0.0869)
log likelihood	-14534.0	-12382.6	-7628.4	2315.0	-18542.1	-16318.8
Players	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Years	✓	✓	✓	✓	✓	✓
Position	✓	✓	✓	✓	✓	✓

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

Table 4.I.1 shows our preferred specification of the control variables for the player-level re-

gressions for each of the player performance measures. Columns (1)-(3) show the effects on the overall performance measures, while columns (4)-(6) show the effects on the efficiency measures. For the reader's convenience, the player-level variables are separated from the team-level variables by the dashed line. As expected, we find wage to have a positive relationship with player performance, and this is consistent and statistically significant at a 1% level across all of the performance measures. The negative coefficient on $Trade$ suggests that the players who play for multiple teams within a season perform worse. This was expected for the overall contribution measures as these players would have less games to contribute at each team. However, the coefficients remain negative for the efficiency measures too, suggesting that the players who are traded during the season are among the worse performers. This also adds evidence to our earlier suggestion that players perform better in teams which remain consistent throughout a season. The negative effects of $Trade$ are statistically significant at a 1% level consistently across the performance measures.

We find a positive relationship between performance and tenure in terms of minutes played, and again this is statistically significant at a 1% level across all the measures. For player experience, we find some evidence of an inverted U-shaped relationship, but the statistical significance of this is not consistent across all models. However, Table 4.11 shows that player experience is correlated with both $lnWage$ and $PTenureMP$. When excluding either of these variables, the significance of both the linear and quadratic term increases in significance across the measures. We therefore elect to retain the quadratic form for $PExp$. The observant reader may notice that the quadratic form of tenure is used at the team-level, but the quadratic form of experience is used at the player-level instead. This decision was driven by the data after many trials. We do not see a contradiction in that, although the performance of sportspeople has a quadratic relationship with experience, the range of team average experience is closer to a linear relationship. And also that, although there may be some negatives to a team being together too long, individually players benefit from staying with a team for longer, perhaps with player turnover around them.

The signs of the team-level variables are largely consistent with our previous analysis, though the significance levels are generally much lower. However, given that the relationship between these variables and our variables of interest has already been established, we choose to retain them as control variables.

4.J Importance of Player-level Controls in Determining Wage Dispersion and National Diversity

Table 4.J.1 shows the results of regressing our variables of interest on our player-level controls. We also include the other relevant variable of interest for each topic: wage dispersion and national diversity. The results shows that each of these player-level controls has a statistically significant relationship with at least one of our variables of interest. Particularly $lnWage$ has strong relationships with the variables of interest. Therefore, despite some endogeneity issues (discussed in Section 4.2), this analysis suggests that $lnWage$ is a relevant control variable and its exclusion may hinder the models ability to identify the effects of our variables of interest on performance, separately from the effects of wage.

Table 4.J.1: The Importance of Player-level Control Variables

	(1) WageSD	(2) WagePol	(3) NatFrac2	(4) NatPol2
lnWage	-0.0157	-0.0119***	-0.000216	0.000677
	(0.0162)	(0.00101)	(0.000403)	(0.000505)
Trade	-0.121***	0.00769***	-0.000203	0.000985
	(0.0334)	(0.00207)	(0.000977)	(0.00124)
PTenureMP	0.0200***	0.000498*	0.000229	-0.000340
	(0.00679)	(0.000273)	(0.000176)	(0.000215)
PExp	0.0395**	0.000592	-0.000397	0.000254
	(0.0164)	(0.000860)	(0.000400)	(0.000490)
$PExp^2$	0.000114	-0.00000356	0.00000681	-0.000000896
	(0.00116)	(0.0000607)	(0.0000286)	(0.0000350)
WagePol	4.975***			
	(0.209)			
WageSD		0.0179***		
		(0.000720)		
NatPol2			0.736***	
			(0.00393)	
NatFrac2				1.183***
				(0.00793)
log likelihood	-9126.3	7539.4	12486.5	11079.9
Players	✓	✓	✓	✓
Teams	✓	✓	✓	✓
Years	✓	✓	✓	✓
Position	✓	✓	✓	✓

N = 5,932
Robust standard errors clustered by Player in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

Table 4.J.2 shows the results of regressing the variables of interest on both the player- and team-level controls simultaneously. Below the dashed line of Table 4.J.2, we again see that the team-level controls significantly influence our variables of interest. With the inclusion of player-level controls, the team-level controls show statistical significance at even higher levels. Above the dashed line, we see that, even once team-level variables have been controlled for,

there are still statistically significant relationships between the player-level controls and our variables of interest.

Table 4.J.2: The Importance of Player- and Team-level Control Variables

	(1) WageSD	(2) WagePol	(3) NatFrac2	(4) NatPol2
lnWage	0.0164	-0.0128***	-0.000218	0.000604
	(0.0109)	(0.000969)	(0.000377)	(0.000479)
Trade	-0.0430	0.00477**	-0.0000440	0.000775
	(0.0270)	(0.00203)	(0.000939)	(0.00119)
PTenureMP	-0.00232	0.000976***	-0.0000529	0.0000467
	(0.00368)	(0.000269)	(0.000142)	(0.000178)
PExp	0.000908	0.00298***	0.000317	-0.000409
	(0.00999)	(0.000797)	(0.000357)	(0.000447)
$PExp^2$	-0.000115	-0.000130**	-0.0000212	0.0000224
	(0.000671)	(0.0000555)	(0.0000252)	(0.0000316)
lnTPayroll	4.609***	0.0361***	-0.0188***	0.0193***
	(0.117)	(0.00794)	(0.00334)	(0.00439)
lnCTSize	-0.332***	0.0220***	0.00403	-0.00783**
	(0.0733)	(0.00596)	(0.00254)	(0.00324)
TTenureMP	0.384***	-0.0156***	-0.00420***	0.00330**
	(0.0280)	(0.00289)	(0.00124)	(0.00153)
$TTenureMP^2$	-0.0365***	0.000858**	0.00162***	-0.00177***
	(0.00360)	(0.000413)	(0.000215)	(0.000275)
TExp	0.0866***	-0.0175***	-0.00554***	0.00572***
	(0.0141)	(0.00132)	(0.000473)	(0.000627)
TExpSD	0.118***	0.0231***	0.00331***	-0.00206**
	(0.0213)	(0.00178)	(0.000666)	(0.000832)
WagePol	3.457***			
	(0.165)			
WageSD		0.0222***		
		(0.000930)		
NatPol2			0.734***	
			(0.00379)	
NatFrac2				1.195***
				(0.00771)
log likelihood	-7217.5	7804.3	12706.5	11261.6
Players	✓	✓	✓	✓
Teams	✓	✓	✓	✓
Years	✓	✓	✓	✓
Position	✓	✓	✓	✓

N = 5,932

Robust standard errors clustered by Player in parentheses.

* p<0.10, ** p<0.05, *** p<0.01

5.A 'Hausman-type' Test - Wage Dispersion

We conduct a 'Hausman-type' test using the Mundlak (1978) approach (which is valid for robust, clustered standard errors) to determine whether a Fixed-Effects (FE) or Random-Effects (RE) regression should be preferred. The same specification as in column (2) of Table 5.7 and as used during the robustness checks is run, but this time using RE and additionally including the team-specific means of the variables.

Table 5.A.1: Hausman-Type Test using the Mundlak (1978) Approach - Wage Dispersion

	(1) WL
m_lnTPayroll	-0.226 (0.257)
m_lnCTSize	-0.304 (0.196)
m_TTenureMP	0.130** (0.0642)
m_TTenureMP^2	-0.0131 (0.00993)
m_TExp	-0.0777** (0.0304)
m_TExpSD	0.0479 (0.0369)
m_WageSD	0.0288 (0.0305)
m_WagePol	1.042*** (0.364)
R^2	0.548
Years	✓
Teams	✗
Controls	✓
Wage Dispersion	✓

N = 330
Robust standard errors clustered by Team in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5.A.1 reports just the coefficients of the team-specific means, as these are of interest for this test and the coefficients of the other variables are identical to the FE estimates. If these team-specific means are jointly significant, then there is evidence of correlation between time-invariant, team-specific heterogeneity and the regressors, so we should favour FE. An F-test of joint significance returns a test statistic of 32.92 and a p-value of 0.0001. We therefore have strong evidence to favour FE regressions.

5.B Alternative Measures of Polarisation

In our main regression analysis, we focus on a single measure of polarisation, constructed using the methodology outlined by Foster and Wolfson (2010). In the main analysis, we simply call this measure $WagePol$, but we now switch to $lnFW$ to clearly distinguish it from the alternative measures of polarisation discussed in this appendix. These include an additional measure constructed using a variation of the methodology of Foster and Wolfson (2010), as well as four measures using the methodology of Duclos et al. (2004). We will now use FW and DER to indicate the respective methodologies. Again, we utilise DASP Abdelkrim and Duclos (2007) to create these measures. Although the formulas were given and discussed Chapter 3, they are repeated here for the readers convenience and to aid discussion.

$$FW = 2(2T - Gini)/\frac{m}{\mu} \tag{5.B.1}$$

As discussed in Section 3.1, though the FW measure given in Equation 5.B.1 does not depend on any parameters, it can be constructed using either the natural logarithm of wages ($lnFW$) or the absolute values ($abFW$). The difference graphically would be whether the axes are given in absolute or percentage terms. Foster and Wolfson (2010) suggest using the logarithmic form to ensure the measure is bounded between zero and one, but note that using the absolute values is also theoretically permissible. We therefore create $lnFW$ (previously $WagePol$ in our previous analysis) using the natural logarithm of wage as our main measure of wage polarisation, but now additionally consider $abFW$ created using the absolute wage values.

$$DER_{\alpha} = \frac{\int\int f(x)^{1+\alpha} f(y)|y-x|dydx}{2\mu^{1-\alpha}} \tag{5.B.2}$$

Section 3.1 also discussed the DER measure and how it depends on the parameter α, which can vary between 0.25 and 1. We therefore construct this measure for α at each of the bounds and for two additional values in-between. We therefore use values of 0.25, 0.5, 0.75 and 1, giving us the variables $DER_{0.25}$, $DER_{0.5}$, $DER_{0.75}$ and DER_{1}, respectively. There is no clear theoretical reason to favour a particular value of α, so all these measures are considered.

Table 5.B.1 shows that the FW measure constructed with absolute wage values, $abFW$, is not bounded and instead ranges between 0.144-2.422. The two FW measures agree on which distributions have highest and lowest polarisation, however there is disagreement on how the measures rank the distributions in-between. Figure 5.B.1 presents a scatterplot of the two measures. While there is a positive correlation between them, they tend to disagree more at higher levels of wage polarisation. This is because $lnFW$ considers percentage changes, while $abFW$ considers absolute wage differences, which will put more weight on larger differences when they exist.

In terms of the DER measures, the summary statistics in Table 5.B.1 seem to show that the higher the value of α, the lower the value of the measure and the lower the standard deviation of the measure. The measure with an α equal to 0.5 seems to be the closest to our primary FW measure, but has a vastly different standard deviation. To investigate this further, Table 5.B.2 looks at how these polarisation measures are correlated.

Table 5.B.2 shows high, but not perfect correlation between the FW and low-α DER measures (0.25 and 0.5). Looking at the first column, the correlation with $lnFW$ drops significantly as

Table 5.B.1: Descriptive Statistics of Alternative Polarisation Measures

Variable	Panel	Mean	Std Dev.	Min	Max
lnFW	Overall	0.209	0.0748	0.0500	0.516
	Between		0.0397	0.136	0.299
	Within		0.0637	0.0479	0.454
abFW	Overall	0.820	0.384	0.144	2.422
	Between		0.183	0.454	1.317
	Within		0.339	0.0857	2.149
$DER_{0.25}$	Overall	0.346	0.0358	0.180	0.423
	Between		0.0194	0.294	0.376
	Within		0.030	0.232	0.419
$DER_{0.5}$	Overall	0.229	0.0153	0.150	0.261
	Between		0.00765	0.207	0.242
	Within		0.0133	0.173	0.265
$DER_{0.75}$	Overall	0.153	0.00787	0.126	0.171
	Between		0.00332	0.147	0.161
	Within		0.00716	0.132	0.173
DER_{1}	Overall	0.103	0.00703	0.0791	0.119
	Between		0.00326	0.0953	0.109
	Within		0.00626	0.0840	0.124

Figure 5.B.1: *FW* Polarisation Measures (Natural Log of Wage versus Absolute Wage)

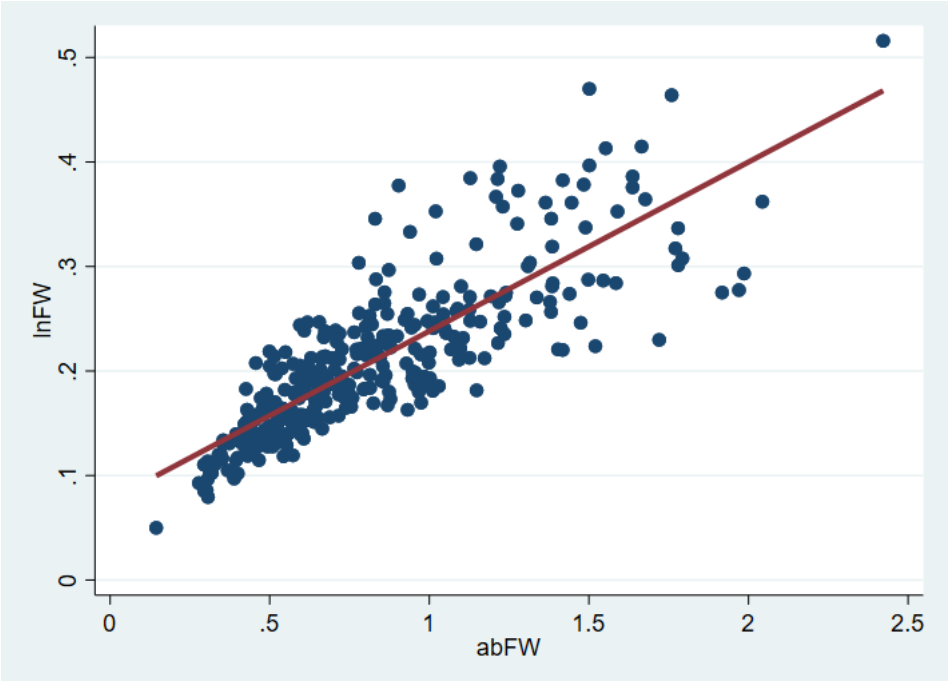


Table 5.B.2: Correlation Between Alternative Measures of Wage Polarisation

	lnFW	abFW	$DER_{0.25}$	$DER_{0.5}$	$DER_{0.75}$	DER_1
lnFW	1.000					
abFW	0.830***	1.000				
$DER_{0.25}$	0.846***	0.780***	1.000			
$DER_{0.5}$	0.840***	0.715***	0.924***	1.000		
$DER_{0.75}$	0.513***	0.309***	0.420***	0.732***	1.000	
DER_1	-0.039	-0.227***	-0.265***	0.112*	0.758***	1.000

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

α changes from 0.5 to 0.75, and even becomes negative when α equals 1. This result suggests that the value of α can have a drastic effect on how these measures captured polarisation and could potentially lead to contradicting results. We test the robustness of our main results from Chapter 5 to these alternative measures of wage polarisation in Appendix 5.C.

5.C Robustness Check - Alternative Measures of Polarisation

Table 5.C.1 provides details on the robustness of our main result from Chapter 5 to the different polarisation measures. Discussion of how these measures are constructed, their descriptive statistics and correlations, can be found in Appendix 5.B. Both FW measures are statistically significant, at a 1% for $lnFW$ but only at a 5% level for $abFW$; while $WageSD$ remains significant also at a 1% level. This suggests that the construction of the FW measure using the natural log of wage seems to better capture the negative effects of wage polarisation. For the DER measures with a low α, the coefficients are statistically significant at a 5% level for $DER_{0.25}$, and at a 10% level for $DER_{0.5}$, while again $WageSD$ remains significant at a 1% level. As α increases the significance levels of both $WageSD$ and the polarisation measure drops, but the signs of the coefficients remain consistent. It appears that, for high levels of α, the DER measures focus too strongly on spread and not enough on bipolarity (or too much on alienation and not enough on identification as Duclos et al. (2004) would say). As a result the measures struggle to capture a distinct effect from $WageSD$.

Table 5.C.1: Robustness Check (Alternative Polarisation Measures)

	(1) WL	(2) WL	(3) WL	(4) WL	(5) WL	(6) WL
WageSD	0.0287***	0.0302***	0.0327***	0.0268***	0.0188**	0.0175**
	(0.00862)	(0.00902)	(0.0101)	(0.00904)	(0.00697)	(0.00731)
lnFW	-0.313***					
	(0.112)					
abFW		-0.0521**				
		(0.0199)				
$DER_{0.25}$			-0.589**			
			(0.268)			
$DER_{0.5}$				-1.131*		
				(0.562)		
$DER_{0.75}$					-1.432	
					(0.993)	
DER_1						-0.841
						(1.272)
R^2	0.503	0.497	0.495	0.495	0.491	0.487
Years	✓	✓	✓	✓	✓	✓
Teams	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓

N = 330
Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

These results confirm that increases in the standard deviation of wage seem to promote competition in a positive way, while increases in polarisation, for a given standard deviation, seem to promote conflict which damages performance. DER polarisation measures with higher alphas do not seem to capture this phenomenon quite as well.

6.A 'Hausman-type' Test - National Diversity

We conduct a 'Hausman-type' test using the Mundlak (1978) approach (which is valid for robust, clustered standard errors) to determine whether a Fixed-Effects (FE) or Random-Effects (RE) regression should be preferred. The same specification as in column (4) of Table 6.10 and as used during the robustness checks is run, but this time using RE and additionally including the team-specific means of the variables.

Table 6.A.1: Hausman-Type Test using the Mundlak (1978) Approach - National Diversity

	(1) WL
m_lnTPayroll	0.0744 (0.197)
m_lnCTSize	-0.279 (0.292)
m_TTenureMP	0.0704 (0.0686)
m_TTenureMP^2	-0.00692 (0.0104)
m_TExp	-0.0598** (0.0270)
m_TExpSD	0.0836 (0.0658)
m_NatFrac2	0.168 (0.956)
m_NatPol2	-0.325 (0.808)
R^2	0.500
Years	✓
Teams	✗
Controls	✓
National Diversity	✓

N = 330
Robust standard errors clustered by Team in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6.A.1 reports just the coefficients of the team-specific means, as these are of interest for this test and the coefficients of the other variables are identical to the FE estimates. If these team-specific means are jointly significant, then there is evidence of correlation between time-invariant, team-specific heterogeneity and the regressors, so we should favour FE. An F-test of joint significance returns a test statistic of 18.99 and a p-value of 0.0149. We therefore have evidence to favour FE regressions at a 5% level.

6.B Wage Distributions of Example Teams

Table 6.B.1 gives further clarification of the wage distributions of the hypothetical teams created in Section 6.4. These hypothetical teams were created to give the reader further intuition of the size of the predicted effects of wage dispersion on team performance.

Table 6.B.1: Wage Distributions of Example Teams

Wage Rank	Wage (in $m) A	B	C	D	E	F
1	16.801 (-)	16.801 (-)	16.801 (-)	16.301 (-$0.5m)	8.579 (1st & 18th)	16.801 (-)
2	12.938 (-)	12.938 (-)	12.938 (-)	12.438 (-$0.5m)	12.938 (-)	12.938 (-)
3	10.385 (-)	10.385 (-)	10.385 (-)	9.885 (-$0.5m)	10.385 (-)	10.385 (-)
4	7.848 (-)	8.348 (+$0.5m)	7.348 (-$0.5m)	7.848 (-)	7.848 (-)	7.848 (-)
5	6.198 (-)	6.698 (+$0.5m)	5.698 (-$0.5m)	6.198 (-)	6.198 (-)	3.6175 (5th & 14th)
6	5 (-)	5.5 (+$0.5m)	4.5 (-$0.5m)	5 (-)	5 (-)	5 (-)
7	4.083 (-)	4.083 (-)	4.083 (-)	4.583 (+$0.5m)	4.083 (-)	4.083 (-)
8	3.303 (-)	3.303 (-)	3.303 (-)	3.803 (+$0.5m)	3.303 (-)	3.303 (-)
9	2.755 (-)	2.755 (-)	2.755 (-)	3.255 (+$0.5m)	2.755 (-)	2.755 (-)
10	2.25 (-)	2.25 (-)	2.25 (-)	1.75 (-$0.5m)	2.25 (-)	2.25 (-)
11	1.85 (-)	1.85 (-)	1.85 (-)	1.35 (-$0.5m)	1.85 (-)	1.85 (-)
12	1.561 (-)	1.561 (-)	1.561 (-)	1.061 (-$0.5m)	1.561 (-)	1.561 (-)
13	1.244 (-)	0.744 (-$0.5m)	1.744 (+$0.5m)	1.244 (-)	1.244 (-)	1.244 (-)
14	1.037 (-)	0.537 (-$0.5m)	1.537 (+$0.5m)	1.037 (-)	1.037 (-)	3.6175 (5th & 14th)
15	0.819 (-)	0.319 (-$0.5m)	1.319 (+$0.5m)	0.819 (-)	0.819 (-)	0.819 (-)
16	0.612 (-)	0.612 (-)	0.612 (-)	1.112 (+$0.5m)	0.612 (-)	0.612 (-)
17	0.488 (-)	0.488 (-)	0.488 (-)	0.988 (+$0.5m)	0.488 (-)	0.488 (-)
18	0.357 (-)	0.357 (-)	0.357 (-)	0.857 (+$0.5m)	8.579 (1st & 18th)	0.357 (-)

Change from base case (Team A) in parentheses
Where two rankings are listed in parentheses, the reported wage is a mean of these rankings

The column headers represent each example team, A-F. For each wage rank, the wage (in $m) is reported, as well as the change from the base case in parentheses. This is either a transfer of ±$0.5m, or a mean of the two rankings reported in the parentheses. It should be

noted that the actual wage rankings may change after the transfers, but the wages are not rearranged to reflect this, as it would simply add unnecessary complication.

Table 6.B.2 is a repetition of Table 6.16, which describes the transfers, for the reader's convenience.

Table 6.B.2: Description of Example Teams

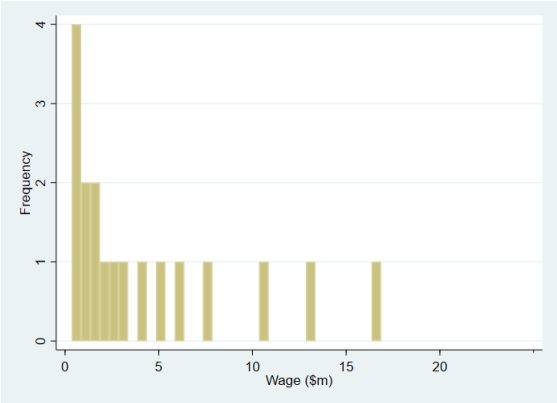
Team	Description (change from base case)
A	Base Case - each player receives the mean wage of the relevant wage rank
B	\$0.5m is transferred from those ranked 13^{th}-15^{th} to those ranked 4^{th}-6^{th}
C	\$0.5m is transferred from those ranked 4^{th}-6^{th} to those ranked 13^{th}-15^{th}
D	\$0.5m is transferred from those ranked 1^{st}-3^{rd} to those ranked 7^{th}-9^{th}, and from those ranked 10^{th}-12^{th} to those ranked 16^{th}-18^{th}
E	Those previously ranked 1^{st} and 18^{th} are paid an average of these wage ranks
F	Those previously ranked 5^{th} and 14^{th} are paid an average of these wage ranks

Figure 6.B.1 presents histograms for the wage distributions of each of the example teams, A-F. While Table 6.B.3 is a repetition of Table 6.17, which gives values of $WageSD$ and $WagePol$ for each distribution, as well as the predicted effects on team performance, again for the reader's convenience.

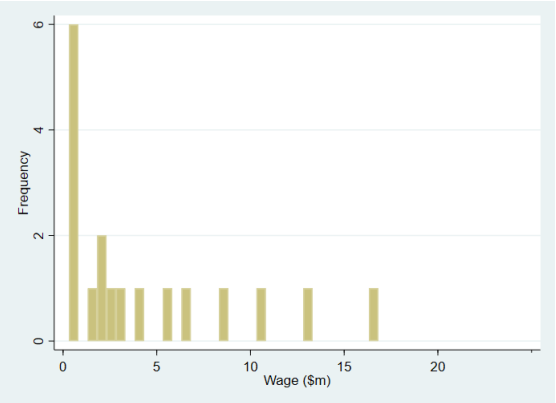
Table 6.B.3: Predicted Effects of Example Teams

Team	WageSD	WagePol	Pred. Effect	Δ from base case	Δ in wins
A	4.7360	0.1910	0.0802	-	-
B	4.8430	0.2250	0.0732	-0.0071	-0.579
C	4.6450	0.1660	0.0851	0.0048	0.397
D	4.5380	0.2180	0.0664	-0.0138	-1.133
E	3.8050	0.1720	0.0589	-0.0213	-1.748
F	4.6520	0.1550	0.0886	0.0083	0.683

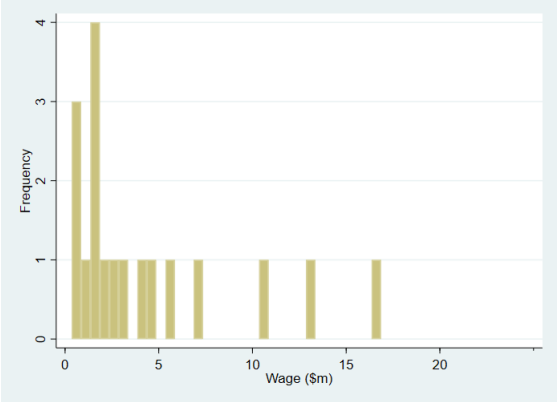
Figure 6.B.1: Histograms of Example Team Wage Distribution



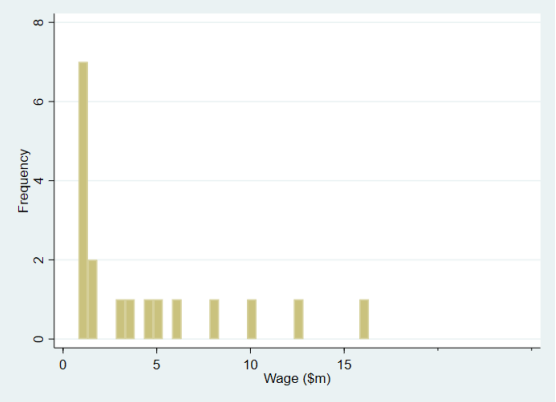
(a) Team A



(b) Team B



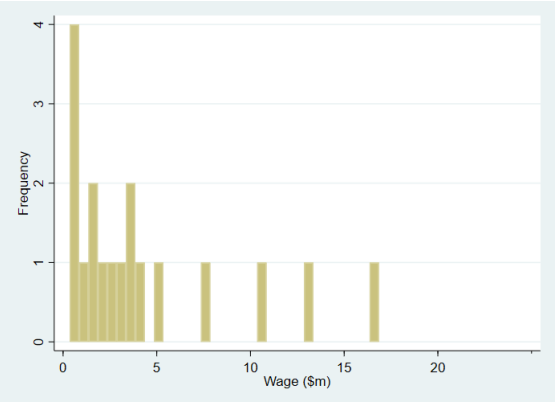
(c) Team C



(d) Team D



(e) Team E



(f) Team F

7.A Wage Dispersion and Individual Player Performance - Calculating the Magnitude of the Predicted Effects

Table 7.A.1 summarises the calculations of the predicted effects of the wage dispersion variables on individual player performance, when they are interacted in Table 7.18 of Section 7.5. For each variable listed in the first column of Table 7.A.1, we then present its coefficient from the regression; its standard deviation; the variable that it is interacted with in the regression specification; the coefficient on this interaction term; then either the minimum, mean or maximum of the variable it is interacted with; and finally the predicted effect.

Table 7.A.1: Calculating the Predicted Effects of the Interacted Wage Dispersion Measures on Player Performance

Variable	Reg. Coef.	Std Dev.	Interacted with	Reg. Coef.	Min. Mean Max.	Predicted Effect
lnWage	0.465	1.431	WagePol	-0.98	0.05 0.215 0.516	0.596 0.364 -0.058
WageRank	-0.0283	5.484	WageSD	0.015	1.554 4.895 10.166	-0.028 0.247 0.681
WageSD	0.271	1.575	WageRank	0.015	-27 -9.636 -1	-0.211 0.199 0.404
WagePol	0.545	0.078	lnWage	-0.98	-5.293 0.74 3.546	0.447 -0.014 -0.229

For example, the predicted effect of a one standard deviation increase in $lnWage$ when $WagePol$ is at its minimum is 0.596 wins produced, as presented in the first row. This is calculated by multiplying the coefficient on $lnWage$ (0.465) by its standard deviation (1.431), then adding to this the product of its standard deviation (again 1.431), the coefficient on the interaction term with $WagePol$ (-0.98) and the sample minimum of $WagePol$ (0.05). This calculation gives the predicted effect of 0.596.[75] The same logic can be applied to the other variables for the minimum, mean and maximum of the variable they are interacted with.

[75] The predicted effect given is calculated before rounding. For the figures presented, the calculation gives 0.595.