Contextual Batting and Bowling in Limited Overs Cricket

by

James Thomson

B.Sc., Dalhousie University, 2018

Project Submitted in Partial Fulfillment of the
Requirements for the Degree of
Master of Science

in the
Department of Statistics and Actuarial Science
Faculty of Science

SIMON FRASER UNIVERSITY
Fall 2020

Declaration of Committee

Name: **James Thomson**

Degree: **Master of Science**

Thesis title: **Contextual Batting and Bowling in Limited Overs Cricket**

Committee: **Chair:** Joan Hu
Professor, Statistics and Actuarial Science

Harsha Perera
Co-Supervisor
Lecturer, Statistics and Actuarial Science

Tim Swartz
Co-Supervisor
Professor, Statistics and Actuarial Science

Paramjit Gill
Examiner
Associate Professor, Data Science, Mathematics, Statistics
University of British Columbia (Okanagan Campus)

Abstract

Cricket is a sport for which many batting and bowling statistics have been proposed. However, a feature of cricket is that the level of aggressiveness adopted by batsmen is dependent on match circumstances. It is therefore relevant to consider these circumstances when evaluating batting and bowling performances. This project considers batting performance in the second innings of limited overs cricket when a target has been set. The runs required, the number of overs completed and the wickets taken are relevant in assessing the batting performance. We produce a visualization for second innings batting which describes how a batsman performs under different circumstances. The visualization is then reduced to a single statistic "clutch batting" which can be used to compare batsmen. An analogous analysis is then provided for bowlers based on the symmetry between batting and bowling, and we define a statistic "clutch bowling".

Keywords: ball-by-ball data, Duckworth-Lewis-Stern resource table, one-day cricket, Twenty20 cricket

Dedication

To Nora and my parents for supporting me through the ups and downs of 2020

Acknowledgements

My time at Simon Fraser has been filled with incredible opportunities. I've spent the last two and a half years working in two sports, first football and later cricket. Every project and opportunity has been made possible by the support of my amazing supervisors Harsha and Tim.

My first foray into sports statistics was the Big Data Bowl for the NFL in 2019, which I have Dani Chu, Matthew Reyers and Lucas Wu to thank for. Later we were fortunate enough to work for the NFL, thanks to Dr. Michael Lopez of the NFL and Skidmore College.

Working on cricket statistics followed, and while doing my thesis I was fortunate to work with Robert Nguyen of the University of New South Wales who has introduced me to many more opportunities. Thank you for all of the Zoom talks about Australian cricketers and pushing me to be a better sports statistician.

Thank you to all of the incredible professors in the statistics department including Dr. Derek Bingham, Dr. Richard Lockhart, Dr. Rachel Altman, Dr. Jinko Graham, and Ian Berkovitz. And thank you to the staff members Charlene Bradbury, Sadika Jungic, Jay Young and Kelly Jay.

I couldn't have done any of this alone, so thank you to my family and Nora Whelan for always being so supportive. It's been an amazing two and a half years, I can't wait to see what comes next!

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction to Cricket

1.1 The History and Rules of Cricket

Cricket originated in England in the late 16^{th} century, and then spread around the world to the British colonies, and later to other countries. The administration and rules of the game are governed by the International Cricket Council (ICC) which has 104 Member nations, consisting of 12 Full Members and 92 Associate Members. The Full Members are qualified to play in all three formats of the game including Test matches - considered to be the highest level of cricket, where matches can go on for five days - while both Full and Associate Members participate in the ’limited overs’ versions of the game which are called One-Day International (ODI) and Twenty20 International (T20I) matches. The 12 Full Members are Afghanistan, Australia, Bangladesh, England, India, Ireland, New Zealand, Pakistan, South Africa, Sri Lanka, West Indies, and Zimbabwe and they enjoy voting rights on who can aspire to Full Membership and play Test matches. However, for this project we are only analyzing limited over matches (ODI and T20 Internationals, Big Bash League (BBL), and Indian Premier League (IPL)).

In cricket, an ’over’ consists of 6 deliveries from the bowler to the batsmen. Sometimes extra deliveries may be added to an over by the umpire to compensate the batsman for unfair deliveries, which may be ’wides’ (ball is not within reasonable reach of the batsman) or ’no balls’ (illegal delivery by the bowler). An ’inning’ is the batting session of one team, and in limited over matches there are two innings, one for each team. One-Day matches have 50 overs per inning and Twenty20 matches have 20 overs per inning.

The game is played on a circular or oval grass field, but there are no specific dimensions required by the ICC, only minimum and maximum sizes for the field are defined. As shown in Figure 1.1 the field itself is split into four parts; an Outfield, an Infield, a Close-Infield and the Pitch. Unlike the field itself which can have some variation in size, the dimensions of the Pitch, where the bowler delivers the balls to the batsman, are clearly defined by the ICC. Figure 1.2 shows the dimensions of the Pitch and the applicable nomenclature.

Each team has 11 players. One team bats first (the first inning of the match) and tries to score as much as they can within the stipulated number of overs (50 or 20) or until they lose all 10 wickets, i.e. having 10 of their 11 batsmen given out (dismissed), similar to the 3 outs in baseball. Thereafter, the other team (who were previously fielding) bats, and that is the second inning of the match. The team batting second tries to beat the score of the team that batted first. The score is measured in 'runs' which may be visualized as 'points'.

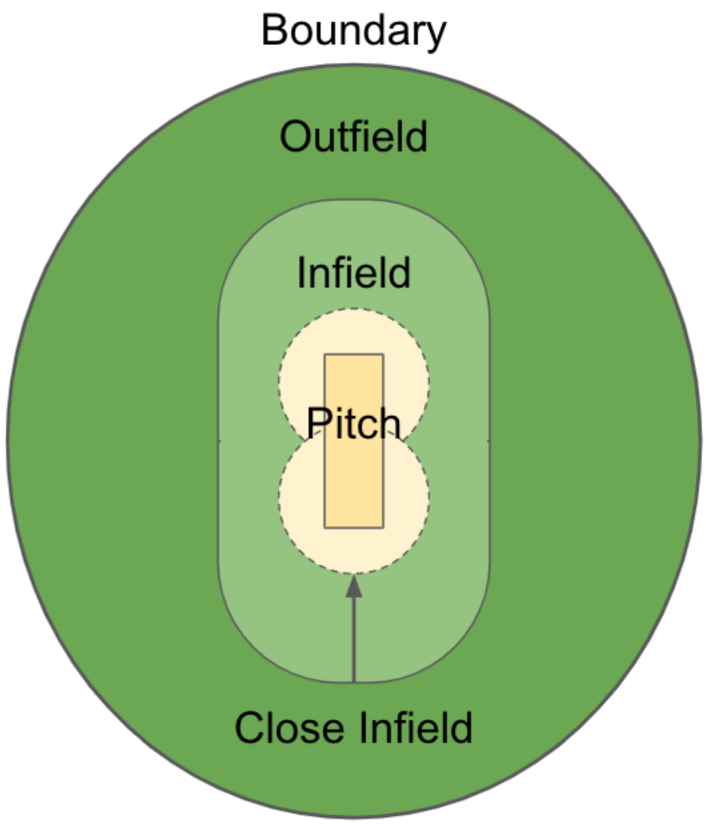

Figure 1.1: The field can be oval or circular, only the pitch in the centre has a set of specific dimensions as shown in Figure 1.2

The core way for the batting side to score runs is for the batsman to hit the ball that is delivered to him in such way and with such placement that the two batsmen can successfully run from one Popping Crease to the opposite Popping Crease of the Pitch (Fig 1.2) before any fielder can collect the ball and throw it back to where the stumps (wickets) are, where a fielding side member will position himself to receive the ball and touch the stumps with the ball. If the batsmen have completed the run from one Popping Crease to the opposite Popping Crease they get one run. In the time it takes for the fielder to collect and throw the ball to the stumps the batsmen may cross over on the Pitch once, twice, thrice or more

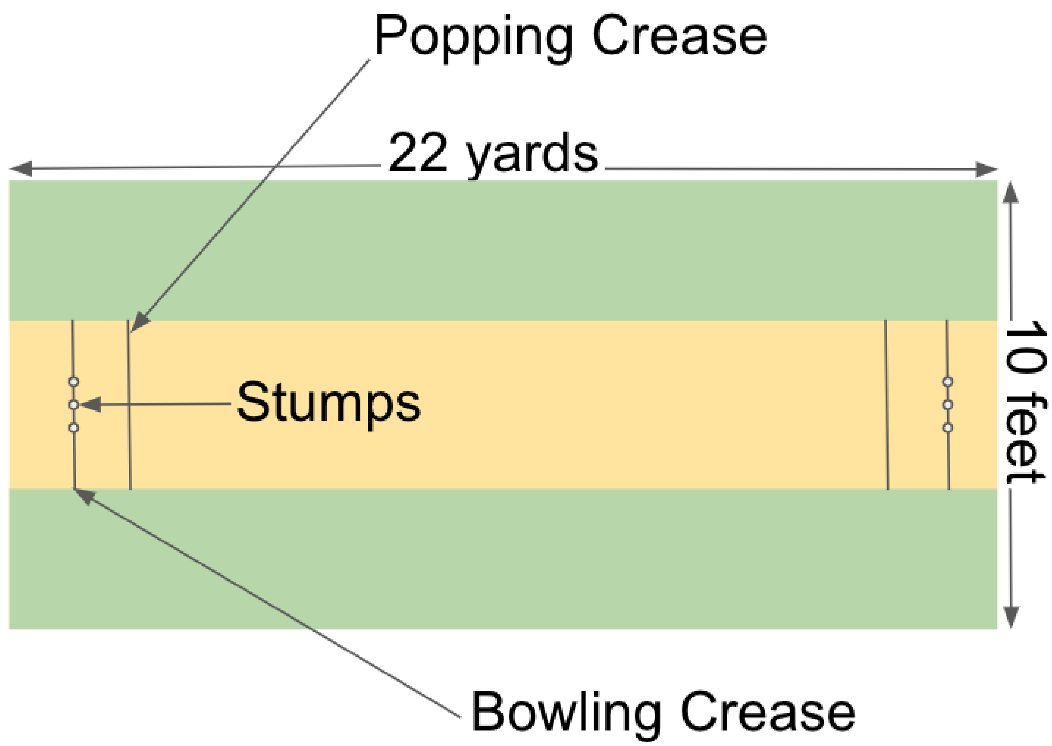

Figure 1.2: The Pitch

and will get the appropriate number of runs. If a fielding side member manages to hit one of the stumps with the ball while the batsmen are still running and the batsman on the side where the stumps have been hit has not reached the Popping Crease he will be given run-out (dismissed). If the batsman hits the ball along the ground all the way to the boundary at the edge of the Outfield he will get 4 runs, but if the ball travels all the way in the air from his bat to the boundary he will get 6 runs.

There are several ways in which a batsman can be dismissed (given out). A ball delivered to him hits his stumps (bowled out), or the ball hits his bat or gloves and goes up in the air and is caught by a fielder before the ball touches the ground (caught-out), or the batsman's legs obstruct the delivered ball from hitting the stumps (leg-before-wicket or LBW), or he himself hits the stumps (hit-wicket), or he is run-out as explained previously. Being stumped-out is similar to being run out in that the wicket-keeper is able to touch the ball to the stumps while the batsman is out of the Popping Crease.

A good bowler will limit the number of runs scored by the batsmen by delivering balls that are difficult to hit, or delivering balls that gets past the batsman and hits the stumps, or trapping the batsman in such a way that his legs obstruct the ball from hitting the stumps (LBW), or deceiving the batsman with a delivery that hits his bat or gloves and goes up in the air where a fielder can catch it.

1.2 Weather Delays and Resources

Cricket is the second most popular sport in the world in terms of the fan base and is just second to soccer. However it failed to catch on in North America the way football, hockey, basketball, and baseball have. Despite its popularity, cricket has not yet embraced sports statistics for player and team evaluations in the way that sports which are popular in North America have embraced statistics. The one exception in the case of cricket is its adoption of a method known as the Duckworth-Lewis-Stern (DLS) method which is dealt with below.

Unlike with other sports the size and nature of the grass field and pitch in which cricket is played precludes the possibility of covering the field against rain, or playing a match in the rain where the field and pitch would soon be a quagmire. Therefore matches can be interrupted or not even started because of rain. This is where the DLS method of employing resources (wickets and overs) comes in and has proven to be of immense value.

The number overs and wickets remaining in an inning determine the number of resources available to a batting team. Two British statisticians, Frank Duckworth and Tony Lewis (1998), defined the exponential decay relationship in Equation 1.1 between the resources u overs remaining and w wickets lost to the average runs scored by a team.

$$Z(u,w) = Z_0(w)[1 - exp\{-b(w)u\}] \tag{1.1}$$

The decay constant $b(w)$ and the asymptotic average score $Z_0(w)$ from the last $10 - w$ wickets in unlimited overs are both functions of w, however the definitions of these functions are confidential due to proprietary rights. The average score of an N overs inning with 0 wickets lost is given by

$$Z(N,0) = Z_0[1 - exp\{-bN\}] \tag{1.2}$$

and so the ratio in Equation 1.3 gives the proportion of resources available to a second inning team at the beginning of u overs and w wickets lost.

$$P(u,w) = Z(u,w)/Z(N,0) \tag{1.3}$$

The results of the 2014-2015 DL method in Table 1.1 can be easily used by officials to calculate a target for a second inning team for interruptions in 1 or both innings.

In 2014 Duckworth and Lewis passed caretakership of the method to Australian statistician Steven Stern, and the method was renamed the Duckworth-Lewis-Stern (DLS) method. T20 cricket had increased rapidly in popularity, requiring a similar table of its own. Initially the DL table was truncated at 20 overs, and by dividing each value by 56.6 - the resources

	Wickets Lost										
Overs Available	0	1	2	3	4	5	6	7	8	9	10
50	100.0	93.4	85.1	74.9	62.7	49.0	34.9	22.0	11.9	4.7	0.0
40	89.3	84.2	77.8	69.6	59.5	47.6	34.6	22.0	11.9	4.7	0.0
30	75.1	71.8	67.3	61.6	54.1	44.7	33.6	21.8	11.9	4.7	0.0
20	56.6	54.8	52.4	49.1	44.6	38.6	30.8	21.2	11.9	4.7	0.0
10	32.1	31.6	30.8	29.8	28.3	26.1	22.8	17.9	11.4	4.7	0.0
5	17.2	17.0	16.8	16.5	16.1	15.4	14.3	12.5	9.4	4.6	0.0
1	3.6	3.6	3.6	3.6	3.6	3.5	3.5	3.4	3.2	2.5	0.0
0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 1.1: Abbreviated version of the Duckworth-Lewis resource table (2014-2015 Standard Edition). The table entries indicate the percentage of resources remaining in a match with the specified number of wickets lost and overs available.

available at 20 overs remaining with no wickets taken - produced the table that began with 100% resources at the start of a T20 match. This received criticism from Perera and Swartz (2013) due to the differences between ODI and T20 matches, namely; differing powerplay lengths, indications of heavy hitter preference in T20 team composition, and that contextually 20 overs remaining in a 50 over match is not the same as beginning a new match with 20 overs remaining. At the same time, first inning scoring in ODI matches had changed to frequently exceed 300 runs, which the original DL method had not been created to fairly assess such matches. Stern (2016) addressed these concerns with a modification to the exponential function to address both the higher scoring in ODI and to include T20 data to update the DL method, and the DLS method replaced the DL method for setting second inning targets, and continues to do so to this day.

1.3 Motivation for this Project

With the DLS table we know the proportion of resources available at any time in an ODI or T20 match, and throughout the second inning we know the batting teams' required runs. This ratio defines the context these batsmen and bowlers face, and the difference in this ratio before and after the delivery gives us their contextual performance. Contextual performance in cricket is currently not a reported statistic, but without context our evaluation of these players is not as complete as it could be. We suspect some players who have outstanding reputations may not be seen as accomplished players under pressure, and vice versa.

In the interest of having enough data to properly analyze the performance of batsmen and bowlers, their data is combined and multiple contexts describe batting and bowling in both ODI and T20 data. We think this is acceptable given the DLS table takes both T20 and ODI data to produce the proportion of resources available at each over and wicket.

1.4 Organization of the Project

In Chapter 2 we defined the context based on required runs and resources available. In Chapter 3 we introduce a single statistic to describe contextual urgency of a given at bat using the resources of the DLS table and the required runs with examples. Even though the DLS table was built to determine resources available to a second inning team after a weather delay, we will not consider games with weather delays in the analysis, which make up only a small proportion of the total matches in the dataset. In Chapter 4 we observe some surprising results with respect to batsmen and bowlers who are considered elite in T20 and/or ODI. In Chapter 5 we discuss the findings and potential other ways to utilize context and resources in performance evaluation.

Chapter 2

Introduction to Contextual Batting and Bowling

Imagine the following scenario: It is the second innings of a limited overs cricket match (either one-day or Twenty20). There are 3 overs remaining, 7 wickets have been taken and the batting team requires R additional runs to win the match. During the next over, the batsman of interest scores 5 runs. How has the batsman performed?

Clearly, the answer to the question is that it depends on the number of runs R required to win the match. If $R = 10$, the batsman has done his job, and his team is in a good position to win the match. However, if $R = 40$, the batsman has underperformed, and the chance that his team will win has diminished considerably.

The evaluation of batting performance is therefore contextual. Yet, context is not considered when using traditional batting statistics. This project attempts to incorporate context in the evaluation of batting and bowling. The basic idea is that prior to every ball bowled, there is a ratio of runs required to resources available that describes the contextual urgency of the second innings chase. After the ball is bowled, the ratio changes. Therefore, performance is measured according to the change in the ratio. A batsman has performed well if there is a decrease in the ratio.

Swartz (2017) provides a review of the various measures that have been proposed to assess batting and bowling performance in cricket. These measures range from simple statistics such as batting and bowling averages and strike and economy rates to WAR (wins above replacement) type measures that are based on match simulation. Player evaluation metrics also differ in their intent, varying from an economic focus (Karnik 2010) to graphical visualizations (van Staden 2009). However, a commonality of all of the proposed measures is that they do not incorporate context at a ball-by-ball level. For example, a player's batting average is obtained by dividing his total runs scored over all matches by his total number of dismissals. Therefore, batting average fails to account for any of the three contextual features.

The *pressure index* (PI) defined by Shah and Shah (2014) and later modified by Bhattacharjee and Lemmer (2016) captures aspects of context. The pressure indices are calculated during the second innings and they attempt to describe the changing circumstances of matches. A difficulty with both measures is that the second innings always commence with PI = 100.0 (or unity), and therefore the indices do not distinguish between the difficulty of attaining large targets versus the difficulty of attaining small targets at the beginning of the second innings. Both papers mention that the pressure index may be used to evaluate batting performance.

An example of the value added by contextual statistics can be found in American football. Fans may have noticed that recently team are 'going for it' more often on 4^{th} and a few yards instead punting. This shift came about from the NFL's recent embrace of football analytics, and the use of Expected Points (EP), first introduced by Carter and Mahol (1971). The EP takes into account field position and possible outcomes of the play, and has led to this shift in attitude in 4^{th} down situations.

Chapter 3

Contextual Performances

3.1 Data

A key element of our approach is that it requires ball-by-ball data. Ball-by-ball data analysis is not common in cricket as most analyses are based on summary statistics as presented in match scorecards. We have developed a parser of match commentary logs that provides detailed ball-by-ball data including the batsman, the bowler, the over, the number of wickets taken and the outcome of the ball. Commentary logs for high level domestic matches and international matches for teams belonging to the International Cricket Council (ICC) are available from the website www.cricinfo.com. The parser has been carefully verified and we believe that it has close to 100% accuracy. The parser was first used in an application to determine optimal batting orders in ODI cricket (Swartz et al. 2006).

Second innings data were collected for 395 ODI matches and 625 domestic Twenty20 and Twenty20 International matches. The domestic matches consisted of those from the Indian Premier League (IPL) and the Big Bash League (BBL) that took place between April 2015 to October 2019. We excluded all matches that were reduced in length due to delays; this resulted in a loss of 10.8% of the ODI matches and 4.5% of the Twenty20 matches. In this dataset, 169,251 balls were bowled. The percent runs scored and wickets taken for all balls faced/delivered in this dataset is given in Table 3.1. We ignore the rare event of a 3, 5, or 7 being scored since combined they make up fewer than 1% of all outcomes. All extras are credited to the batsmen rather than the team and are included in the runs scored.

Runs	Rate
0	44.5%
1	33.6%
2	5.9%
4	8.8%
6	2.9%
Wicket	3.8%

Table 3.1: Runs and wicket percentages of all deliveries in the dataset

Since the DLS methodology is built only on ODI and T20I data, the inclusion of match data from all One-Day or T20 domestic leagues would not be appropriate. However the Australian Big Bash League and the Indian Premier League are the top domestic leagues T20 in the world, and host to many of the most celebrated batsmen and bowlers in the world. The IPL allows teams to field up to 4 international players in a match, and with the highest salary cap in domestic cricket regularly bring in the top players worldwide, which are of interest to us in this study.

3.2 Contextual batting

As previously mentioned, a batsman's approach in the second innings (his degree of aggressiveness versus cautiousness) is dependent on context. And context is a function of (1) the runs required, (2) the overs remaining and (3) the wickets taken. How then should we quantify context in terms of these three elements in the second innings of a limited overs cricket match?

We begin with the interplay between overs and wickets. A batsmen can be more aggressive when there are fewer overs remaining and can be more aggressive when fewer wickets have been taken. Fortunately, the interplay between overs and wickets is described via the Duckworth-Lewis-Stern (DLS) resource table. Although some details of the construction of the DLS table are propriety, the estimation of resources is based on run scoring from historical matches. In one-day cricket, a batting team begins their innings with 100% of their resources available (i.e. 50 overs and 10 wickets at their disposal). When the team has used up all of their overs or 10 wickets have been taken, the innings are complete and they have 0% of their resources remaining. For intermediate values of overs and wickets, the DLS table gives the appropriate resource percentage. In the case of T20 cricket, a simple transformation of the resources from the one-day table gives the T20 resource percentage. The Duckworth-Lewis method (Duckworth and Lewis 1998, 2004) was introduced in the context of resetting targets in interrupted one-day cricket matches.

For the purposes of our investigation, what is important to note is that DLS resources provide a measure that is proportional to run scoring capability. It therefore follows that at any particular juncture of the second innings, the ratio r of runs required (for victory) to the resources available describes the contextual urgency of the second innings chase.

The ratio of runs required to resources available r is a key statistic in our work. Using the combined ODI/Twenty20 dataset, Figure 3.1 provides a histogram of r based on all of the balls bowled in the second innings. It is good to have a physical understanding of r. In the ODI matches in our dataset, the average number of runs scored in the first innings is 263. Therefore, at the beginning of the second innings of ODI matches, the average value of the ratio is $r = 263/100 = 2.63$. In the combined ODI/Twenty20 dataset, it is also interesting to note that the batting side was never able to win if $r > 9.17$ at any point in the second innings, Further, when $r > 3.33$, the batting side won only 25% of the time, and when $r > 2.80$, the batting side won only 50% of the time. Therefore, we will define highly challenging batting contexts as those for which $r \in (2.80, 3.33)$. Only 33,282 second innings balls were bowled in this challenging scenario.

The upper limit is arbitrary, a 25% win probability is a reasonable cutoff for competitive and anything higher we would consider difficult, and the method begins to penalize batsmen unfairly beyond this point. If a team starts with an $r = 3.33$, that's a target of 333 in ODI and 189 in T20. Figure 3.2 shows the required runs for each combination of overs remaining and wickets taken when $r = 3.33$. These required runs seem high but achievable, so we believe that this range is fair.

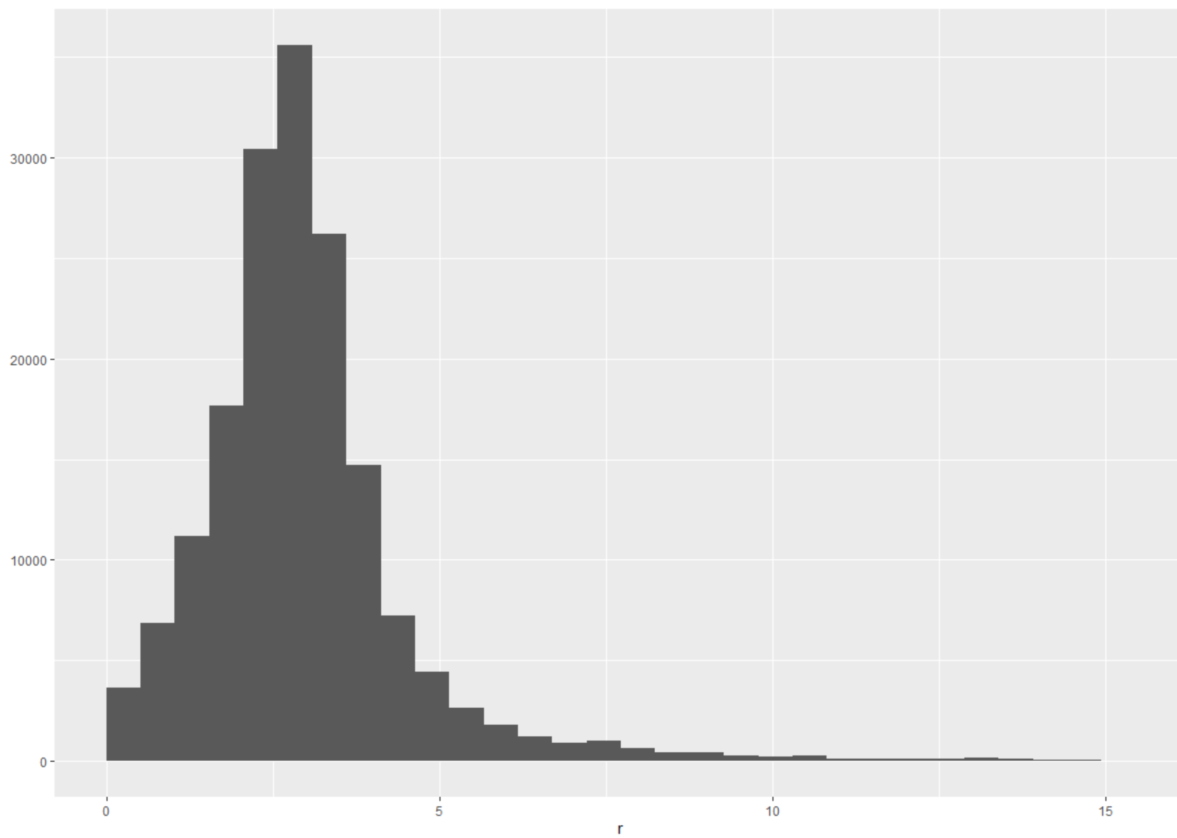

Figure 3.1: Histogram of r based on all second innings balls in the combined ODI/Twenty20 dataset.

The motivation for our approach is based on well established concepts in limited overs cricket; the runs required in the chase and the resources available. For every ball that a batsman faces, we therefore have his ratio r_0 before the ball is bowled and his ratio r_1 after the ball is bowled. For example, if there is a target of 250 runs at the beginning of the second innings in an ODI match, then we have a starting ratio $r_0 = 250/100 = 2.5$. If the batsman scores a run on the first ball, then referring to the DLS table, $r_1 = 249/99.8 = 2.495$. If

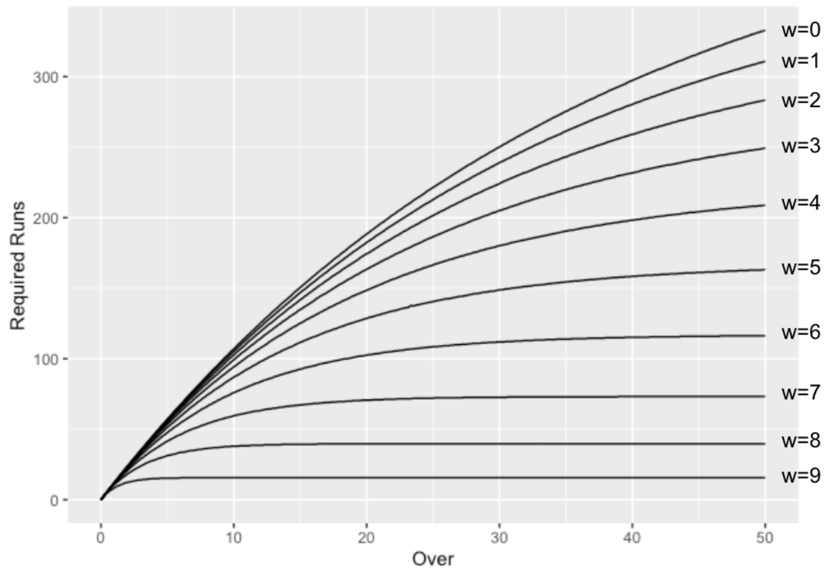

Figure 3.2: The required runs at an $r = 3.33$ for u overs and w wickets lost

instead, no runs were scored on the first ball, $r_1 = 250/99.8 = 2.505$. And in the case where a wicket was obtained on the first ball, $r_1 = 250/93.4 = 2.677$.

Therefore, the ordered pairs $(r_0, r_0 - r_1)$ over all balls that a batsman has faced describes the batsman's performance with respect to the contextual difficulty of the chase. On a particular ball, the quantity r_0 describes the difficulty of the chase with larger values of r_0 corresponding to more challenging chases. The quantity $r_0 - r_1$ describes the contribution by the batsman based on the outcome of the ball where $r_0 - r_1 > 0$ corresponds to improving his team's situation.

However, there is a difficulty with the interpretation of $r_0 - r_1$. Towards the end of the second innings when resources are limited, it is possible that the ratios r_0 and r_1 can be relatively large. In this case, $r_0 - r_1$ can vary greatly with respect to a given ball. In fact, r_1 is undefined if the ball in question is the last ball of the innings or if it results in the 10th wicket (since the remaining resources are nil). To adjust for this, and to compare apples to apples, we make two modifications. First, we introduce the arbitrary cutoff that when resources are less than 10%, we do not include the batting outcome. Second, we introduce the statistical technique of standardization. We disregard the rare events corresponding to scoring three runs and five runs, and for a given ball, we calculate the 7 outcome possibilities of 0-6 runs scored, a wicket is taken, or an extra is given, given by

$r_0 - r_1(0)$ – the result of $r_0 - r_1$ if 0 runs are scored
$r_0 - r_1(1)$ – the result of $r_0 - r_1$ if 1 runs are scored
$r_0 - r_1(2)$ – the result of $r_0 - r_1$ if 2 runs are scored
$r_0 - r_1(4)$ – the result of $r_0 - r_1$ if 4 runs are scored
$r_0 - r_1(6)$ – the result of $r_0 - r_1$ if 6 runs are scored
$r_0 - r_1(w)$ – the result of $r_0 - r_1$ if a wicket falls
$r_0 - r_1(e)$ – the result of $r_0 - r_1$ if an extra occurs

We then define our variation for each delivery as

$$\begin{aligned} V \; = [\; & (r_0 - r_1(0) - 0)^2 + \; (r_0 - r_1(1) - 0)^2 + \; (r_0 - r_1(2) - 0)^2 + \; (r_0 - r_1(4) - 0)^2 + \\ & (r_0 - r_1(6) - 0)^2 + \; (r_0 - r_1(w) - 0)^2 + (r_0 - r_1(e) - 0)^2 \;] \; / \; 6 \end{aligned}$$

and replace the observed $r_0 - r_1$ with the standardized quantity

$$s(r_0 - r_1) = \frac{r_0 - r_1}{\sqrt{V}} \; . \tag{3.1}$$

Therefore, using (1), the plot of $(r_0, s(r_0 - r_1))$ over all balls that a batsman has faced describes the batsman's performance with respect to the contextual difficulty of the chase. The points are then smoothed to provide a trend for the batsman. This smoothed curve is referred to as the batsman's *contextual batting function*. It describes performance over a range of contextual circumstances. When one batsman's curve dominates (i.e. lies above) another batsman's curve, the first batsman is the better batsman in all contexts.

Note that in the case of extras such as byes, leg-byes, wide-balls and no-balls, we credit the extra runs to the batsman. Although a case may be made that these extra runs are not a function of batting performance, they occur while the batsman is on-strike. Perhaps the batsman should receive credit for the extras as the bowler takes the strengths of the batsman into account during the delivery. In Section 3.3, we propose an analogous visualization for bowlers; in this case, it is evident that extras ought to be charged against bowlers. Therefore, we retain symmetry in the visualization by also giving credit to batsmen for extras. Extras occur at the rate of 5.1% in Twenty20 cricket (Davis, Perera and Swartz 2015).

Consider Figure 3.3 which displays the points $(r_0, s(r_0 - r_1))$ and the contextual batting function (ODI and T20) for the high profile batsman Steve Smith who was the former captain of Australia. The function is provided over the range of contexts $r_0 \in (0, 4)$. We observe that Smith bats infrequently in some contexts, and has not batted at all when $r_0 > 4$. This is partly explained by noting that Australia is a strong cricketing nation and rarely falls behind by huge margins during matches. We also note that some contexts (e.g. $r_0 < 2.80$) correspond to more comfortable chases which are not as interesting. It appears that for most contexts, Smith's contextual batting function lies above the par line $s(r_0 - r_1) = 0$ which suggests that he is improving his team's situation in these chases. Figure 3.3 also illustrates a difficulty in visualization; when the points are plotted, it is difficult to compare the contextual batting function with the par line. We also observe that wickets are very damaging when assessing contextual batting performance; the large negative points in Figure 3.3 correspond to wickets and these points pull the contextual batting curve downwards.

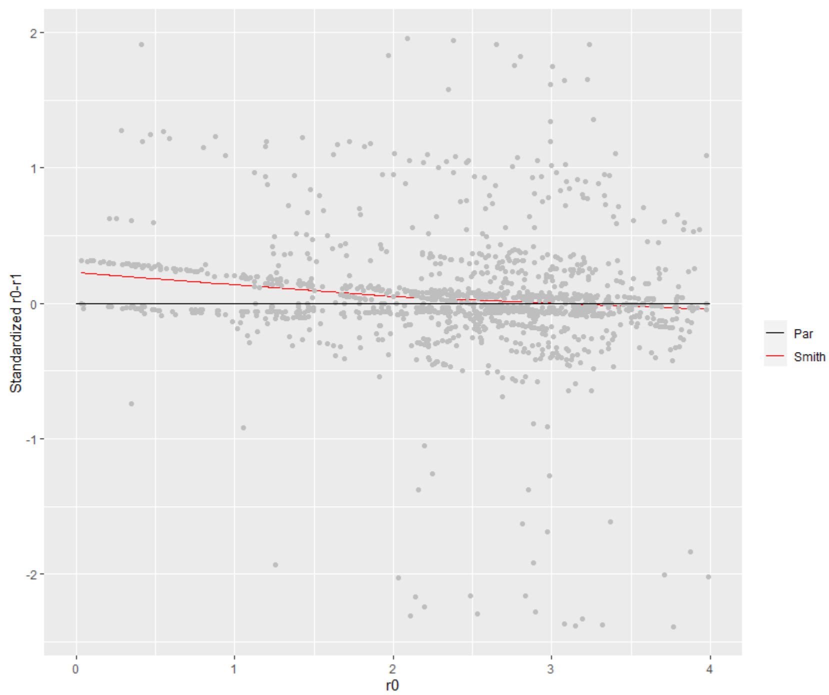

Figure 3.3: The points $(r_0, s(r_0 - r_1))$ and the resulting contextual batting function for Steve Smith over the contextual range $r_0 \in (0, 4)$.

3.3 A summary statistic for contextual batting

The contextual batting function describes a complete picture of how a batsman performs against a range of contexts. When comparing two batsman, one batsman is superior if his curve dominates (lies above) the other curve. However, it will not always be the case that one curve dominates the other and therefore interpretation of contextual batting functions is necessary. For example, it could be the case that a batsman is very good at pushing through a win when a win is expected but is unable to produce a huge number of runs when his team is trailing badly. In this case, his contextual batting function would lie above the line $s(r_0 - r_1) = 0$ for smaller values of r_0 and lie below the line $s(r_0 - r_1) = 0$ for larger values of r_0. In fact, many batsman will have functions that take this general shape with crossover points on the line $s(r_0 - r_1) = 0$.

Although the contextual batting function provides a description of batting performance across a range of contexts, the general public is more at ease with simple univariate statistics that are directly comparable. We will therefore propose a single statistic *clutch batting* which is a summarization of the contextual batting function when matches are highly challenging (i.e. $r_0 \in (2.80, 3.33)$). Clutch batting provides overall evaluation of second innings batting in challenging chases.

In Figure 3.4, we illustrate the clutch batting statistic with reference to Steve Smith. Smith's contextual batting function is restricted to the challenging range $r_0 \in (2.80, 3.33)$. This allows us to narrow in on his performance in the more interesting matches where there is a challenging chase. We compare Smith's performance against the par line $s(r_0 - r_1) = 0$ where a player is doing just enough to maintain the difficulty of the chase. Interestingly, Smith does not seem to help his team in the most difficult contexts (e.g. $r_0 > 3.05$) where his team is struggling. We also note that there is a downward slope to his contextual batting curve. This makes sense from a a cricketing perspective since it becomes more and more difficult to overcome a losing position as r_0 increases.

If we denote the contextual batting function as $f(r_0)$, then the clutch batting statistic is defined as

$$\begin{aligned} \mathrm{C_{bat}} &= \left(\int_{2.80}^{3.33} (f(r_0) - 0)\ dr_0 \right) 100 \\ &= \left(\int_{2.80}^{3.33} f(r_0)\ dr_0 \right) 100 \qquad (3.2) \end{aligned}$$

where 100 is a scaling factor that is introduced to make the statistic more appealing. Therefore, (2) involves an area calculation involving the contextual batting function and the par line. Accordingly, $\mathrm{C_{bat}}$ is an overall measure of batting performance in challenging situations where larger values of $\mathrm{C_{bat}}$ denote greater proficiency. In Smith's case, his clutch batting statistic is 0.46 which suggests that overall, Smith is improving his team's situation in challenging chases.

3.4 Contextual bowling performance

In cricket, there is an inherent symmetry between batting and bowling. Whereas a batsman attempts to score runs and avoid wickets, the bowler attempts to limit runs and take wickets. Therefore, the previous development of contextual batting can be modified to provide an analysis of contextual bowling. As before, during the second innings, we study the ratio of the runs required by the batting team (to win the match) to the resources available. Opposite to batsmen, a bowler attempts to increase the ratio on each delivery of the ball.

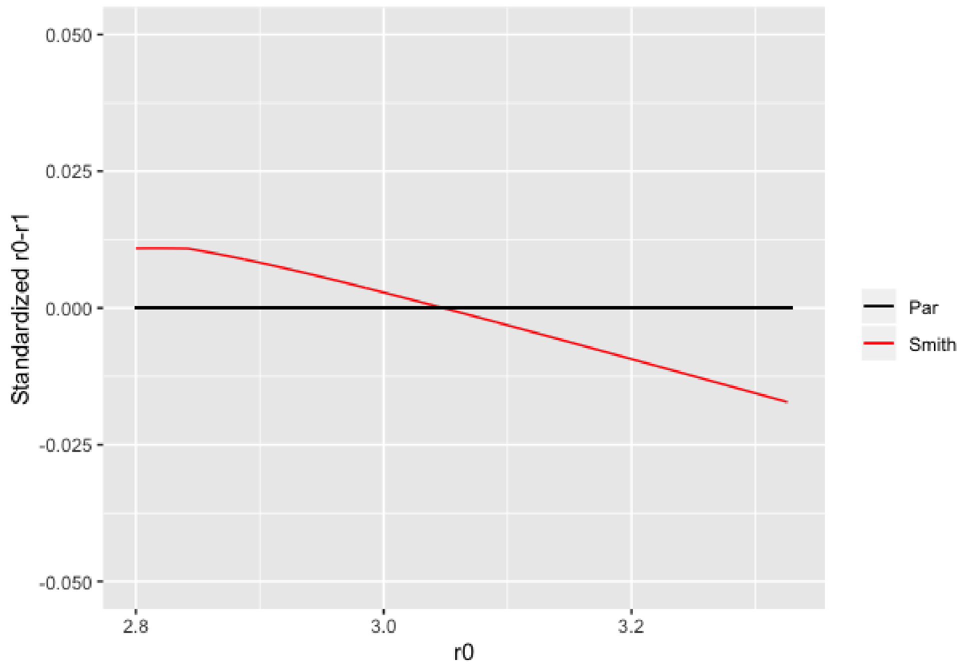

Figure 3.4: The contextual batting plot for Steve Smith over the challenging range $r_0 \in (2.80, 3.33)$.

In clutch batting, we defined $r_0 \in (2.80, 3.33)$ as highly challenging contexts corresponding to a win probability interval $(0.25, 0.50)$ for the batting side. We likewise define $r_0 \in (2.27, 2.80)$ as highly challenging contexts corresponding to a win probability interval $(0.25, 0.50)$ for the bowling side. We therefore define the clutch bowling statistic as

$$\mathrm{C}_{\mathrm{bowl}} = \left(\int_{2.27}^{2.80} -f(r_0) \; dr_0 \right) 100 \tag{3.3}$$

where the negative sign has been introduced since we wish positive values of the statistic to be associated with clutch bowling.

An attractive feature of the clutch batting statistic (2) and the clutch bowling statistic (3) is that batsmen and bowlers can be assessed on the same scale. We note that it is possible to change the bounds of integration in (2) to correspond to difficult situations for the bowler. However, to retain symmetry with the batting statistic, we investigate the same challenging chases from the point of view of the batsman.

Like Steve Smith, we illustrate the clutch bowling statistic for Rashid Khan over the challenging range $r_0 \in (2.27, 2.8)$ in Figure 3.5. This time we see that his entire performance is under the par line $s(r_0 - r_1) = 0$, meaning he is always helping his team by increasing the difficulty of the chase for the batsmen. Rashid Khan has a clutch bowling statistic of 3.26, and as r_0 decreases we see a downward slope, which makes sense from a cricketing perspective since the bowler faces a more challenging context as $r_0 \to 0$.

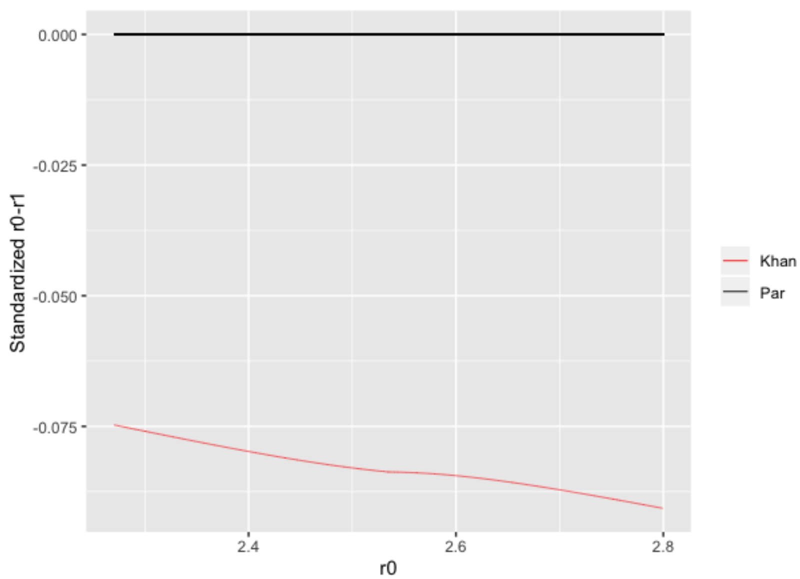

Figure 3.5: The contextual bowling plot for Rashid Khan over the challenging range $r_0 \in (2.27, 2.80)$.

Chapter 4

Data Analysis

4.1 Details of implementation

As previously emphasized, we use combined data from both ODI cricket and T20 cricket during the period April 2015 through October 2019. The synthesis of the two formats is possible due to the introduction of the runs to resource ratio r that describes and standardizes the contextual difficulty of the chase at any point during the second innings in limited overs cricket.

An issue related to clutch batting and clutch bowling is that we limit the analysis of batsmen and bowlers to those who have faced/delivered sufficient balls. This provides us with reliable statistics that are not heavily influenced by the results of only a few batting and bowling outcomes. We set the minimum number of balls faced/delivered to 300. This leaves us with 24 batsmen and 19 bowlers under consideration. With 12 full member nations in the International Cricket Council (ICC), our study is therefore restricted to a small number of prominent batsmen and bowlers on each ICC team.

The shape of the contextual batting and bowling curves are impacted by smoothing. We prefer curves that are not too "wiggly" since we do not believe there are physical reasons why batting and bowling performances should oscillate over the range r of contextual urgency. Smoothing is carried out in the R programming language using the *loess* function (Cleveland 1979). The parameters *span* and *degree* determine the characteristics of smoothing in loess. The parameter span $\in (0, 1)$ controls the smoothing neighbourhood where larger span means that more nearby data influence the fit. The parameter degree > 0 specifies the order of the smoothing polynomials where higher order polynomials permit more wiggle in the fitted curve. In our application, we have set span $= 1$ and degree $= 1$.

Once the loess function has been determined, the calculation of the clutch batting statistic (2) and the clutch bowling statistic (3) require numerical integration to obtain the areas beneath the loess curve. This is done using the *uniroot* and *integrate* functions in R. The function *uniroot* finds the roots of the contextual curves, and *integrate* obtains the corresponding areas between the roots.

To get a sense of the reliability of the clutch batting and clutch bowling statistics, we associate standard errors to the statistics. This is implemented through a bootstrapping procedure where for each batsman and bowler, we resample (with replacement) the balls faced/delivered from their individual dataset. From the resampled data, the clutch batting/bowling statistic is calculated, and this resampling procedures is repeated $M = 10,000$ times. From the M simulated statistics, a standard error is calculated.

4.2 Clutch batting analysis

In Table 4.1, we present the clutch batting statistic 24 prominent batsmen in limited overs cricket who have faced at least 300 balls. For comparison purposes, we also present the batting average, the strike rate and the survival rate (i.e. balls per dismissal) of van Staden (2009) where large values of the three statistics are all indicative of good batting. An immediate reaction is that clutch batting correlates positively but not strongly with batting average (average number of runs scored per wicket) where the sample correlation coefficient is 0.34. The correlation between clutch batting and strike rate is 0.70 and the correlation between clutch batting and survival rate is -0.34. This suggests that clutch batting detects an aspect of performance that resembles features of the strike rate. Clutch batting and the survival rate appear to have little in common. Note that to make fair comparisons, we have calculated the common statistics based on the same timeframe considered in our dataset.

We see from Table 4.1 that the best clutch batsman is Jason Roy of England followed by his countryman Jos Buttler. Whereas Roy and Buttler are known as a solid batsman, they are spectacular in situations when their team is in desperate need of runs. Perhaps this partly explains England's good run of form in recent years. We also observe that the remarkable Virat Kohli of India is also a top clutch batsman. Shai Hope of the West Indies is situated at the bottom of the table. He does not pull his team from the brink in desperate chase situations. We do note that Hope's statistic was based on only 331 balls and has a standard error of 0.95.

Another observation from Table 4.1 is that the bootstrap standard error is large. This is caused by large values of $|s(r_0 - r_1)|$ which impact clutch batting. These impactful observations correspond to scoring sixes and dismissals. We note that the standard errors tend to decrease with greater numbers of at-bats. Given the large standard errors, we can only make broad inferences concerning the differentiation between batsmen.

Our top performing batsman JJ Roy's contextual batting plot is shown in Figure 4.1 (a). His contextual batting curve lies completely above the Par line, which indicates he excels at putting his team in a better position to win in competitive situations. Figure 4.1(b) breaks down JJ Roy's performance by runs and outs. In this competitive range 15.9% his bats are 4s, nearly double the average of 8.8% for batsmen in the entire dataset. He is also getting out less, and scoring fewer 1s and 0s than the average batsman. The lowest performing

Batsman	Country	Balls	C_{bat} (Std Err)	Bat Avg	Strike Rate	Surv Rate
JJ Roy	England	416	2.40 (0.79)	41.5	116.6	31.2
JC Buttler	England	470	2.15 (0.91)	40.6	138.3	31.5
AB de Villiers	South Africa	325	2.04 (0.99)	48.4	146.2	34.7
CH Gayle	West Indies	374	1.81 (1.08)	35.6	130.9	28.7
Q de Kock	South Africa	534	1.68 (0.64)	47.1	110.9	36.2
S Dhawan	India	507	1.50 (0.89)	38.3	116.4	32.7
V Kohli	India	659	1.00 (0.49)	66.7	114.8	49.3
RG Sharma	India	812	0.68 (0.52)	49.5	112.4	39.0
AJ Finch	Australia	608	0.52 (0.51)	44.5	115.8	34.1
SPD Smith	Australia	399	0.46 (0.75)	41.8	101.1	39.3
SR Watson	Australia	319	0.20 (1.17)	24.4	137.3	20.8
BA Stokes	England	317	0.14 (0.71)	46.3	107.3	33.1
S Al Hasan	Bangladesh	329	0.02 (0.84)	41.7	99.6	36.4
LRPL Taylor	New Zealand	369	0.01 (0.64)	52.3	92.2	57.2
KS Williamson	New Zealand	676	0.01 (0.56)	52.4	97.5	44.4
TWM Latham	New Zealand	306	-0.54 (0.69)	43.0	88.8	45.4
E Lewis	West Indies	472	-0.72 (0.99)	27.1	108.5	27.9
PR Stirling	Ireland	312	-0.94 (1.03)	36.6	94.8	37.2
HM Amla	South Africa	303	-0.98 (1.05)	38.3	102.4	49.4
S Sarkar	Bangladesh	339	-1.32 (1.15)	35.4	108.8	27.3
F du Plessis	South Africa	338	-1.32 (1.07)	44.3	106.4	46.4
WU Tharanga	Sri Lanka	316	-1.38 (1.07)	37.8	91.9	37.9
AM Rahane	India	388	-1.41 (1.12)	35.3	107.5	34.3
SD Hope	West Indies	331	-1.75 (0.95)	33.1	78.7	57.8

Table 4.1: Clutch batting C_{bat} and other statistics for 24 batsmen who have faced at least 300 balls in high level limited overs cricket matches. For comparison purposes, batting average, strike rate and survival rate were calculated over the same data collection period.

batsman Shai Hope's contextual batting curve in Figure 4.1 (c) lies completely below the Par line. In Figure 4.1 (d) Shai Hope's runs and outs are shown in competitive situations. 2.2% of his bats are outs, well under the average of 3.8% for all batsmen. However he also has fewer 4s and 6s, and gets more 0s and 1s than the average batsman.

These plots in Figure 4.1 indicate that as the difficulty of the chase increases, 0s and 1s become increasingly negative pulling the batsman's contextual batting curve below the Par line. Batsmen like JJ Roy who score more 2+ runs than the average batsman do well. Batsmen like Shai Hope who score fewer 2+ runs do poorly, even without losing many wickets. Given that we are looking only in situations where batsmen face a difficult chase, batsmen need to be scoring more runs to put their team in a better position to win, and clutch batting statistic reflects this.

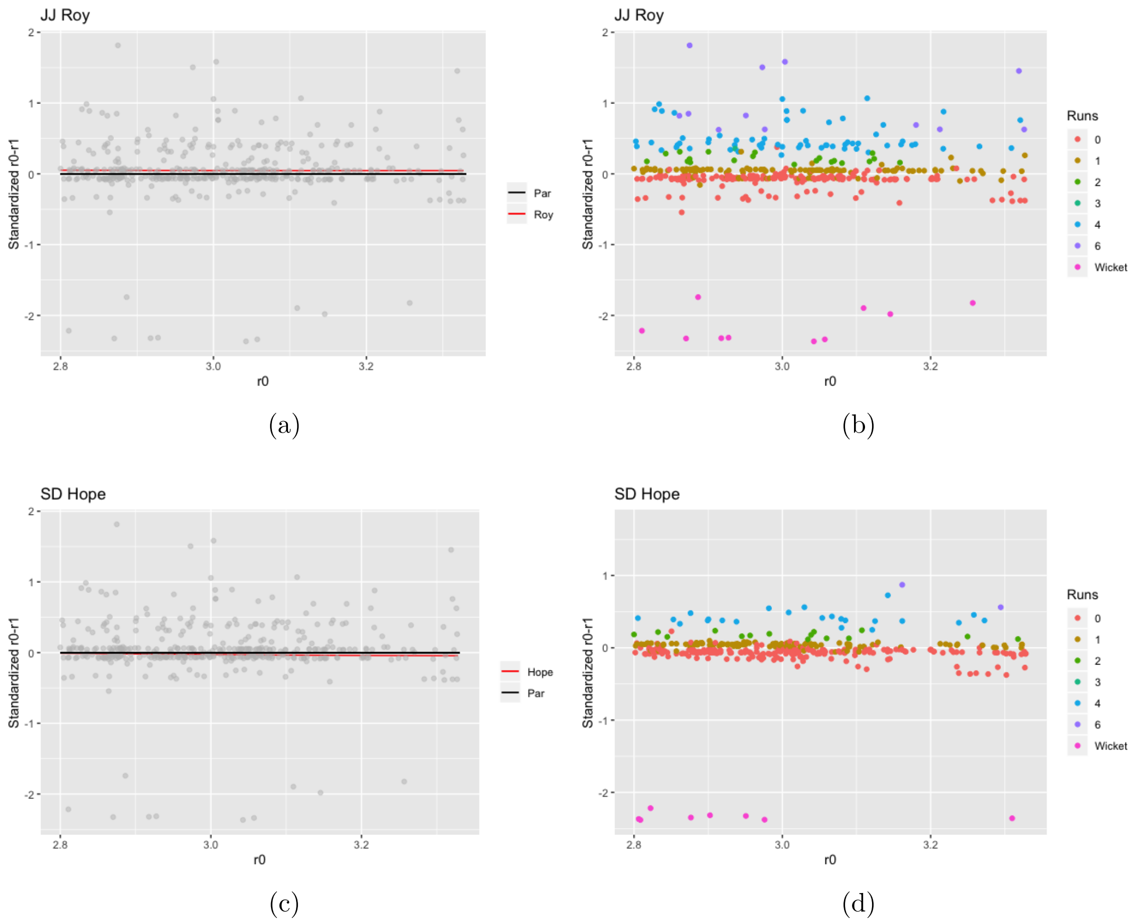

Figure 4.1: The performances of JJ Roy and Shai Hope, both their contextual batting curves and their standardized ball-by-ball performance in the contextual batting range $r_0 \in (2.8, 3.33)$

4.3 Clutch bowling analysis

In Table 4.2, we present the clutch bowling statistic for 19 prominent bowlers in limited over cricket who have delivered at least 300 balls. For comparison purposes, we also present the strike rate, the bowling average and the economy rate where low values of the three statistics are all indicative of good bowling. An immediate reaction is that clutch bowling correlates moderately with bowling strike rate (average number of balls bowled per wicket) where the sample correlation coefficient is -0.51 The correlation between clutch bowling and bowling average is -0.54 and the correlation between clutch bowling and economy rate is -0.72. This suggests that clutch bowling is detecting an aspect of performance that resembles features of the economy rate. Note that to make fair comparisons, we have calculated the common statistics using the same timeframe considered in our dataset. As with clutch batting, we observe that the bootstrap standard error is high and this makes it difficult to differentiate between bowlers. We note that the range (-2.84,3.26) of the clutch bowling statistics in

Table 4.2 is similar to the range (-1.75,2.40) of the clutch batting statistics in Table 4.1. The similarity is a consequence of the symmetry in the definitions of clutch bowling and clutch batting.

Looking at particular players, we observe that both Rashid Khan and Mujheeb Ur Rahman of Afghanistan are strong clutch bowlers. This comes as no surprise as they are the top two T20 bowlers according to the current ICC rankings. Perhaps it is a surprise that the South African fast bowler, Kagiso Rabada sits as the second worst clutch bowler in Table 2. We also obtained the clutch bowling statistic of Lasith Malinga of Sri Lanka ($C_{bowl} = -0.24$) who only delivered 278 balls over the time period. Malinga is of particular interest due to his reputation as an incredible "death" bowler. The clutch bowling statistic suggests that Malinga's reputation is perhaps overrated.

Bowler	Country	Balls	C_{bowl} (Std Err)	Strike Rate	Bowl Avg	Econ
R Khan	Afghanistan	680	3.26 (0.64)	20.9	16.0	5.2
M Ur Rahman	Afghanistan	418	2.07 (1.03)	28.5	21.0	4.7
MJ Henry	New Zealand	330	0.95 (1.13)	33.2	30.7	5.8
MJ Santner	New Zealand	373	0.00 (0.72)	35.9	29.6	5.3
M Nabi	Afghanistan	624	-0.02 (0.60)	35.3	24.9	5.3
PJ Cummins	Australia	304	-0.10 (0.75)	33.0	25.0	5.7
M Rahman	Bangladesh	467	-0.10 (0.68)	28.5	21.8	5.9
TA Boult	New Zealand	364	-0.13 (0.67)	28.1	24.8	6.1
YS Chahal	India	351	-0.14 (0.87)	23.0	22.6	6.3
I Tahir	South Africa	465	-0.19 (0.54)	26.4	21.6	5.7
I Wasim	Pakistan	323	-0.21 (0.88)	35.7	27.3	5.2
S Al Hasan	Bangladesh	591	-0.33 (0.61)	33.6	31.2	5.8
JJ Bumrah	India	341	-0.35 (0.81)	21.0	21.5	5.8
M Mortaza	Bangladesh	620	-0.46 (0.39)	41.7	37.6	5.6
A Zampa	Australia	371	-0.56 (0.51)	29.9	29.5	6.1
TG Southee	New Zealand	362	-0.62 (0.63)	36.7	24.7	6.6
B Kumar	India	534	-0.74 (0.78)	33.9	25.7	6.1
K Rabada	South Africa	346	-1.02 (0.63)	38.8	21.9	5.5
MP Stoinis	Australia	364	-2.84 (0.86)	36.4	33.6	7.2

Table 4.2: Clutch bowling C_{bowl} and other statistics for 19 bowlers who have delivered at least 300 balls in high level limited overs cricket matches. For comparison purposes, strike rate, bowling average and economy rate were calculated over the same data collection period.

Rashid Khan is the highest performing bowler, and his contextual bowling plot is shown in Figure 4.2 (a). It lies completely below the Par line, indicating he puts his team in a better position to win in competitive bowling situations. His performance is broken down by runs given up and wickets taken in Figure 4.2 (b). Rashid Khan is taking wickets with 4.2% of his deliveries, above the average of 3.8%. He is also giving up 4s at 7.5% and 6s at 2.4% of his deliveries, which are both less than the average bowler. The lowest performing bowler Marcus Stoinis' bowling curve in Figure 4.32 (c) lies completely above the Par line.

In Figure 4.2 (d), Marcus Stoinis' runs given up and wickets taken show us that he is taking wickets with 2.4% of his deliveries, less than the average of 3.8%. He's also giving up slightly more 1s, 2s, 4s and 6s than the average bowler.

The plots in Figure 4.2 illustrate how a bowler must give up fewer runs and take more wickets to help their team in competitive situations. Rashid Khan performs exactly how an ideal bowler would in this situation, while Marcus Stoinis is failing to make the batting team's position more difficult. Since 0s and wickets are always negative, keeping the number of 2+ runs scored off a single delivery low and taking more wickets are the goals for bowlers in these competitive bowling situations.

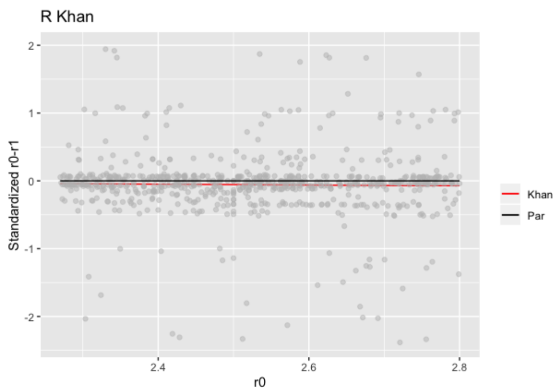

(a)

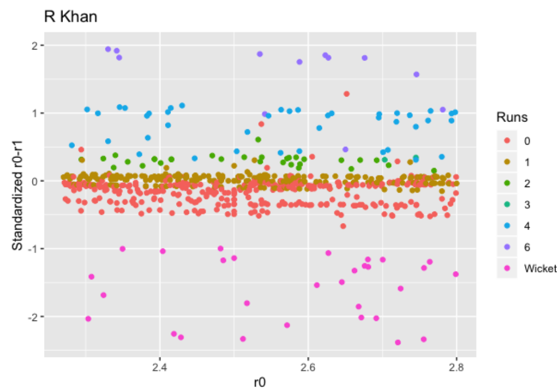

(b)

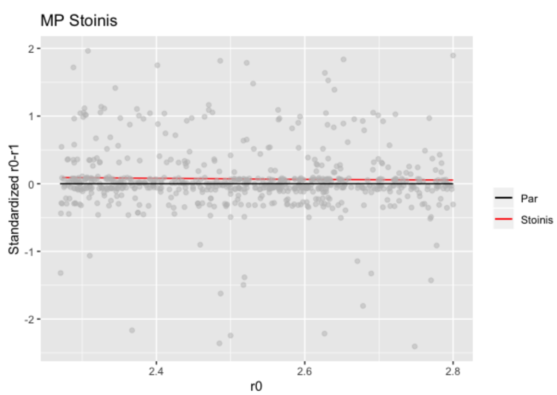

(c)

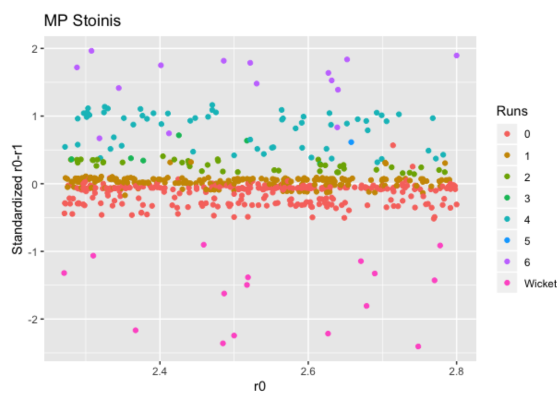

(d)

Figure 4.2: The performances of Rashid Khan and Marcus Stoinis, both their contextual bowling functions and their standardized ball-by-ball performance in the contextual bowling range $r_0 \in (2.27, 2.8)$

4.4 Data Synthesis

One of the bold assumptions that we have made in the paper involves the synthesis of data. In forming the clutch statistics, we used data from both domestic cricket and international cricket. In addition, we combined T20 data and one-day data. Of course, our intention was to provide the best measures of clutch performance statistics through utilizing as much data as possible. To investigate these assumptions, we first investigated batting averages corresponding to the 24 batsmen in Table 4.1 using the same timeframe. We observed that six of the batsmen (Latham, Stirling, Tharanga, Hope, Sarkar and Taylor) did not compete domestically in T20. We therefore excluded these batsmen from the following analysis. We then separated the T20 batting averages by calculating a domestic batting average x and an international batting average y for the remaining 18 batsman. Using a simple linear regression of y versus x, a lack of difference between the two competitions would imply an intercept $\beta_0 = 0.0$ and a slope $\beta_1 = 1.0$. We obtained estimates (standard errors) of 25.98 (15.15) and 0.15 (0.42) for the intercept and slope, respectively. This suggests that there may be slight differences in the scoring patterns between the two competitions. The calculation of the proposed clutch statistics are dependent on the runs to resource ratio r introduced in Chapter 3. As we have combined T20 and ODI datasets, it is important to check that r is invariant to the two formats. In Figure 4.3, we have overlaid the histograms of r calculated for all second innings balls for the two datasets. We observe that the two histograms have roughly the same shape and this suggests that amalgamation of the two datasets may be appropriate. It could be the case that the r_0 values for ODI are slightly larger.

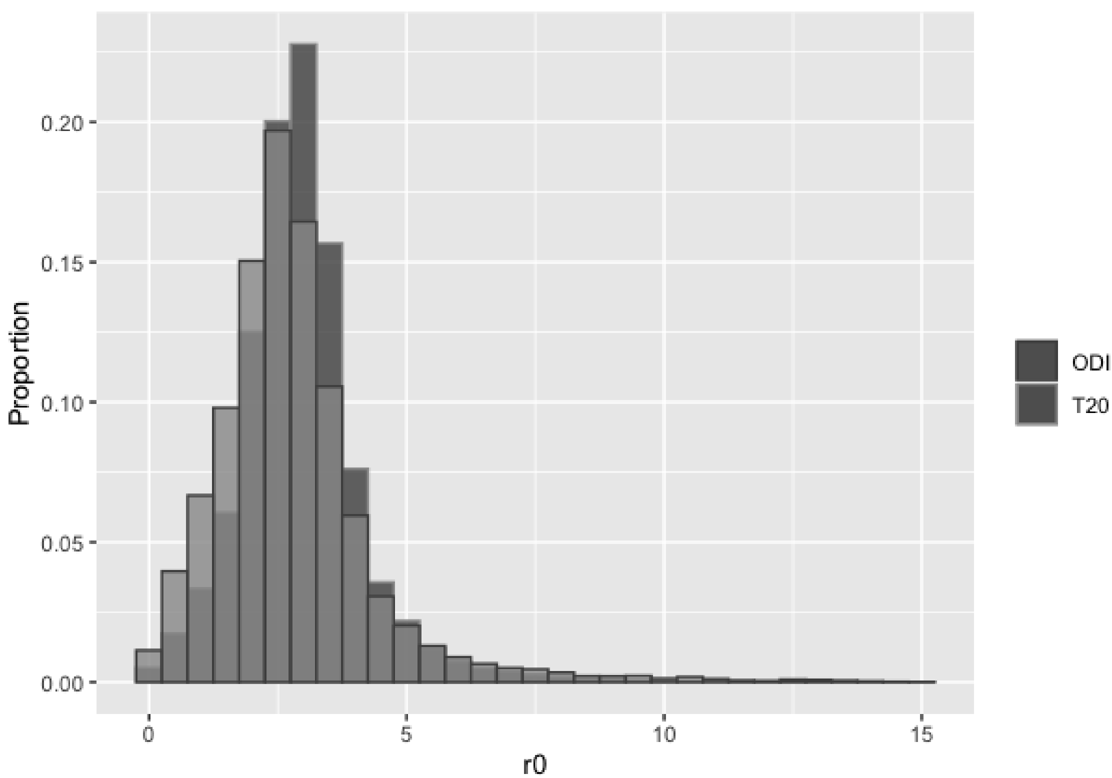

Figure 4.3: Frequency histograms of r based on all second innings balls displayed for the ODI and Twenty20 datasets.

Chapter 5

Conclusions

The proposed clutch batting and clutch bowling statistics are not intended to usurp traditional and popular statistics such as batting average and bowling strike rate. Rather, these statistics were introduced to provide insight on an aspect of performance that had not been previously investigated. Specifically, we are interested in how batsmen and bowlers perform in the second innings when their teams are in difficult situations. Hence, the clutch batting and clutch bowling statistics are contextual.

There are names that don't appear in our results that will perhaps surprise readers, David Warner or Mahendra Dhoni for example. Both see under 300 balls in the challenging context, primarily batting in easier and more difficult contexts. The selection of the r limits could be adjusted to include more players, or the minimum number of balls faced dropped for players of interest. There is also the possibility of directly comparing players using the methods above, calculating the contextual batting or bowing statistic between players could be useful in player evaluation and team selection.

The definition of the functions that calculate resources are confidential, so we are forced to take them at face value. We are unaware of other research in cricket that utilizes resources in this manner, however we think this is an area that offers new insights in player evaluation. The target is only available in second inning cricket, but resources are known for both innings. The ratio of a bowler's economy to resources taken per over could be of interest as well, and like our analysis has a nice symmetry for batsmen evaluation as well.

The data analyses have demonstrated that performance in difficult contexts does not correlate highly with overall performance measures. From the point of view of tactics, the clutch batting and clutch bowling statistics may provide teams with useful information to determine optimal batting and bowling orders in difficult contexts.

There are various ways in which our ideas concerning contextual performance may be explored in future research. For example, one could study different contexts as described by the runs to resource ratio r_0. Also, it is possible to narrow or expand data collection timeframes under consideration. Alternatively, one could define statistics that weight recent performances more highly. Contextual performance in sport is clearly an important and understudied subject area that deserves greater attention.

Bibliography

Bhattacharjee, D. and Lemmer, H.H. (2016). Quantifying the pressure on the teams batting or bowling in the second innings of limited overs cricket matches. *International Journal of Sports Science & Coaching*, 11, 683-692.

Carter, V., & Machol, R. E. (1971). Operations research on football. Operations Research, 19(2), 541-544.

Cleveland, W.S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, 74, 829-836.

Davis, J., Perera, H. and Swartz, T.B. (2015). Player evaluation in Twenty20 cricket. *Journal of Sports Analytics*, 1, 19-31.

Duckworth, F.C. and Lewis, A.J. (1998). A fair method for resetting targets in one-day cricket matches. *Journal of the Operational Research Society*, 49, 220-227.

Duckworth, F.C. and Lewis, A.J. (2004). A successful operational research intervention in one-day cricket. *Journal of the Operational Research Society*, 55, 749-759.

Karnik, A. (2010). Valuing cricketers using hedonic price models. *Journal of Sports Economics*, 11, 456-469.

Perera, H. P., & Swartz, T. B. (2013). Resource estimation in T20 cricket. IMA Journal of Management Mathematics, 24(3), 337-347.

Shah, P. and Shah, M. (2014). Pressure index in cricket. *Journal of Sports and Physical Education*, 1, 9-11.

Stern, S.E. (2016). The Duckworth-Lewis-Stern method: Extending the Duckworth-Lewis methodology to deal with modern scoring rates. *Journal of the Operational Research Society*, 67, 1469-1480.

Swartz, T.B., Gill, P.S., Beaudoin, D. and de Silva, B.M. (2006). Optimal batting orders in one-day cricket. *Computers and Operations Research*, 33, 1939-1950.

Swartz, T.B. (2017). Research directions in cricket. In *Handbook of Statistical Methods and Analyses in Sports*, Editors J.A. Albert, M.E. Glickman, T.B. Swartz and R.H. Konig, Chapman & Hall/CRC Handbooks of Modern Statistial Methods: New York, 445-460.

van Staden, P.J. (2009). Comparison of cricketers' bowling and batting performance using graphical displays. *Current Science*, 96, 764-766.