Garry Gelade*

Evaluating the ability of goalkeepers in English Premier League football

Abstract: This paper examines the performance of goalkeepers in the English Premier League. A commonly used metric to assess goalkeeper ability is the saves-to-shots ratio. It is shown that goalkeepers playing in weak teams have to defend against on-target shots that have a relatively high probability of scoring as compared to their counterparts in higher-performing teams. This tends to produce a downwards bias in the save-to-shots ratio of goalkeepers in weak teams, an effect which may lead to their ability being underestimated. A logistic regression model is used to adjust goalkeepers' saves-to-shots ratios for differences in save difficulty. The adjusted ratio is more stable across seasons, suggesting that it is a more reliable estimate of true goalkeeper ability than the unadjusted ratio.

Keywords: football; goalkeeper; performance.

***Corresponding author: Garry Gelade,** Business Analytic Ltd, 1, Circus Lodge, Circus Road, St Johns Wood, London, UK, e-mail: garry@business-analytic.co.uk

1 Introduction

How good is a goalkeeper, and how much does he contribute to his team? Keeping goal in football requires a number of different skills. In addition to making saves from on-target shots, a goalkeeper is responsible, amongst other things, for defusing threats by claiming possession of the ball in dangerous situations, distributing the ball to the outfielders to initiate an attack, and communicating with his defenders. However, it is generally accepted that the single most important goalkeeping skill is preventing goals by making saves. By and large the number of shots that a team concedes depends on the actions of their own defenders and the opposition attack rather than on the actions of the goalkeeper, and it is only relatively rare goalkeeping errors such as fumbling, or failing to claim a ball, taking up an inappropriate position, or not communicating with his defenders, that might afford the opposition additional opportunities to shoot. Typically, therefore goalkeepers are assessed by their saves-to-shots ratio (*SSR)* defined as the number of saves made divided by the number of shots on target. This metric has a strong link with winning football matches; in a sample of 686 seasonal performances from five first-tier European leagues between 2006 and 2012, the correlation between the number of wins per season (counting a draw as half a win) and the SSR was 0.58.

Previous research on goalkeeping in football can be divided into two categories. First, studies in laboratory or other non-game settings have examined the perceptual, physical and biomechanical aspects of goalkeeping performance. For example Savelsbergh et al. (2002) found that successful goalkeepers are more accurate in predicting the height and direction of the penalty kick, wait longer before initiating a response and spend longer periods of time fixating on the non-kicking leg compared with their less-successful counterparts; Dessing and Craig (2010) found that goalkeepers tend to misjudge the flight of curving shots; Wood and Wilson (2010) found that moving goalkeepers were more of a distraction to penalty takers than stationary goalkeepers; and Sørensen, Thomassen and Zacho (2008) explored the biomechanical profiles of elite and non-elite goalkeepers.

The second area of research is in-game studies. These studies have tended to focus on penalty kicks. For instance, using data from the French and Italian leagues, Chiappori, Levitt and Groseclose (2002) found that penalty takers and goalkeepers play the mixed-strategy Nash equilibrium, while in a sample of televised football matches from national leagues and championship competitions, Bar-Eli et al. (2007) found evidence for a goalkeeper bias to action (moving rather than staying still) which leads to performance inefficiencies. In a study of FIFA World Cup penalty shoot-outs, Roskes et al. (2011) found that goalkeepers under pressure (i.e., when their teams were behind) showed a bias to rightward motion when defending against penalties.

Such studies have produced a wealth of knowledge about the mental and physical aspects of goalkeeper performance in penalty situations. Nevertheless, relatively few goals in elite football are scored from penalties; between 2006 and 2012 in the English Premier League, only 7.7% of goals scored were penalty goals, and the rest were scored from regular play. It is therefore surprising that comparatively few studies of have used in-game data to examine goalkeeper performance in regular play. There are probably several contributing factors. First, goal attempts from penalties are much less variable than goal attempts from regular play; only two players are involved, the shooting distance is constant and the ball is at rest when struck. A key variable, the location where the shot was taken from is therefore eliminated from analysis. Second, the interlocking strategies of the penalty taker and the shooter are theoretically interesting from a game-theoretic and decision-making point of view, and the data are amenable to mathematical modelling. Third, and perhaps most importantly, large amounts of rich in-game data during regular play have only recently become available through the activities of commercial sports data companies such as OptaSports and ProZone.

In one of the few published studies on in-game goalkeeper performance Oberstone (2010) examined a commercially developed index of goalkeeper performance (the Opta Index) and found that it could be accurately predicted from six game actions, of which shots faced inside and outside the box and goals conceded were the most important predictors. Although the *SSR* was not explicitly part of the model, the connection is clear. Secondly, Pollard, Ensum and Taylor (2004) examined scoring as a function of shot location, finding that shots originating close to the goal and centrally were more likely to result in a goal than shots originating further away and to one side of the goal. This finding suggests a potential problem with the *SSR* as an indicator of goalkeeper ability; goalkeepers in teams with weak defences are likely to be exposed to a higher proportion of challenging shots than their counterparts in teams with strong defences. The ability of goalkeepers in weak teams may therefore be under-estimated. Similar considerations were advanced by Schuckers (2011) and formed the basis for a study of goal-tending in the National Hockey League and the development of an adjusted measure of goal-tender performance.

This paper highlights the shortcomings of the traditional *SSR* in football and develops a more accurate measure of goalkeeper ability (the "standardized *SSR*") which controls for differences in save difficulty. First, performance data from the English Premier League is used to model the probability of scoring from on-target goal attempts. The model is then used to predict the expected number of shots saved by each goalkeeper. By comparing the expected and actual numbers of saves we can construct a measure of goalkeeper performance adjusted for difficulty.

2 Data source and measures

The data used in this paper was provided by OptaSports, and includes all the on-target goal attempts in two consecutive seasons (2010–2011, 2011–2012) of the English Premier League. After excluding penalty shots there were 6596 shots defended by 59 goalkeepers, playing in 23 teams. For goalkeepers facing at least 50 shots, the mean *SSR* was 71% (SD=4.6%). Each shot in the dataset is associated with variables characterising its path, trajectory and style which are described next.

2.1 Shot path

The variables describing shot path are illustrated in Figure 1A and B. A path was encoded by four variables, one pair describing the start-point of the shot, and one pair describing the end-point. The first pair (X, Y) is shown in Figure 1A. X represents the distance between the shot start-point and the goal line, and Y represents the distance between the start-point and the pitch midline.

The second pair (W, Z) is shown in Figure 1B. W denotes the distance between the goal midline and ball at the end of its path, and Z denotes the height of ball at the end of its path.

All distances are measured in meters. Note that because the path variables X and W are unsigned, paths that are symmetrical about the playing area midline receive the same encoding.

2.2 Shot trajectory

Shot trajectory was represented by six binary descriptors: Bounce (yes/no); Swerve (yes/no); Deflection (yes/no);

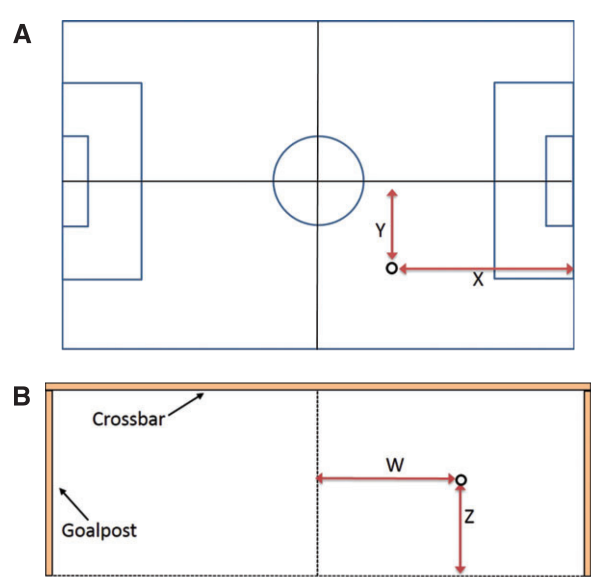

Figure 1 Shot path variables.
(A) Shot start-point variable. (B) Shot end-point variables.

Strong (yes/no); Weak (yes/no); Far post (vs. near post). In all cases, the first characteristic listed was encoded as 1 and the second 0. Shots originating and ending on different sides of the field (i.e., shots that crossed the playing area mid-line) were classified as "Far post" shots, while those originating and ending on the same side of the field were classified as "Near post" shots.

2.3 Shot style

Shot style was represented by three binary descriptors: Assist (yes/no), Open play (vs. set play); Volley (yes/no). Again, the first mentioned characteristic was encoded as 1 and the second as 0.

Table 1 reports the univariate statistics for the study variables.

3 Model development and results

The dependent variable for the model was shot outcome, encoded 1 if a goal was scored and 0 if it was saved. Shot outcome was modelled by a logistic regression equation, in which the independent variables were: the four descriptors of shot path and their higher order and interaction terms; the six trajectory descriptors; and the three style descriptors. During model development, a number of interaction terms were examined; most contributed little or nothing to increased predictive accuracy, and only interactions amongst the shot path variables were retained. The model results for the final model are shown in Table 2.

Table 1 Univariate statistics for on-target shots.

Variable	Mean	SD
Start distance from goal line *X* (m.)	14.78	8.2
Start distance from mid-line *Y* (m.)	6.76	5.2
End distance from mid-line *W* (m.)	1.58	1.0
End height *Z* (m.)	0.79	0.6
	Percent of shots	
Goals scored	28.7%	
Swerve	13.5%	
Bounce	20.1%	
Deflection	6.6%	
Far post	47.8%	
Strong	11.4%	
Weak	7.7%	
Assisted	77.4%	
Fast break	5.0%	
Open play	72.3%	
Volley	9.4%	

Note: N=6596.

The regression model predicted the correct outcome of 81% of shots and the Nagelkerke pseudo-R^2 was 0.463. The likelihood ratio test for the overall model was highly significant (χ^2=2510.6, df=10, p~=0) indicating that the model was significantly better than the intercept-only model. The Hosmer-Lemeshow test for differences between observed and expected frequencies was only just non-significant (χ^2=14.5, df=8, p=0.07); however the Hosmer-Lemeshow statistic is sample-size dependent (Paul, Pennell, and Lemeshow 2013), and in large samples, small departures from the model can be deemed significant.

To assess the performance of the model on out-of-sample data, a 10-fold cross-validation was conducted (Geisser 1975). In this procedure, the data are divided into ten equally-sized and mutually exclusive random samples which serve as test (out-of-sample) data sets, with the corresponding non-selected cases serving as the training sets. Goodness of fit in the test sets was comparable with the fit of the overall data: the average Nagelkerke R^2 was 0.453 (minimum=0.451, maximum=0.460); the average of correct classifications was 81% (minimum=75%, maximum=83%); and the Hosmer-Lemeshow test was non-significant in six of the ten test samples. It can therefore be concluded that the overall fit of the model is adequate, and that it generalizes well to new data.

The signs of the X and Y coefficients indicate that on-target shots originating close to the goal, and close to the pitch midline are more likely to score. This is consistent with expectation; shots originating close to the goal give the goalkeeper less time to react, and shots from close to the midline give the striker a large target to aim at. A similar result was found by Pollard et al. (2004), although those authors modelled all shots rather than on-target shots, and parameterized the distance from the pitch midline in terms of an angle rather than a distance. The present model finds small but significant quadratic start-location terms, and an XY interaction; Pollard, Ensum and Taylor however did not find any significant higher order terms. Ignoring the small effects of the quadratic term, the odds ratio for X indicates that increasing the distance from goal by one meter decreases the odds of scoring by approximately 23.4%. Similarly, increasing the distance from the midline by one meter decreases the odds of scoring by approximately 32.5%.

The start-point portion of the model is illustrated in the probability density maps of Figure 2A and B.

Figure 2A maps the empirical probabilities of scoring from an on-target shot on a grid defined by the path

Table 2 Logistic regression coefficients.

Variable category	Variable	LR Coefficient (B)	Std. Err.	Exp(B)[a]
Shot path	Start distance from goal line (*X*)	−0.267***	0.013	0.766
	Start distance from midline (*Y*)	−0.393***	0.024	0.675
	X^2	0.001***	0.000	1.001
	Y^2	0.003**	0.001	1.003
	$X*Y$	0.012***	0.001	1.012
	End distance from midline (W)	−0.059 ns	0.156	0.943
	End height (Z)	−2.412***	0.257	0.090
	W^2	0.277***	0.044	1.319
	Z^2	0.970***	0.107	2.638
	$W*Z$	0.224***	0.058	1.251
Trajectory	Swerve	0.769***	0.117	2.158
	Bounce	0.293***	0.091	1.340
	Deflection	1.097***	0.134	2.995
	Far post	0.467***	0.069	1.595
	Strong	1.149***	0.113	3.155
	Weak	−2.120***	0.193	0.120
Style	Assisted	−0.406***	0.084	0.666
	Fast break	0.855***	0.165	2.351
	Open play	0.287***	0.085	1.332
	Volley	−0.021 ns	0.117	0.979
	Constant	2.725***	0.218	

***$p<=0.001$; **$p<=0.01$. Note: Model summary: $-2LL=5402$; Nagelkerke $R^2=0.463$. [a]Odds ratio.

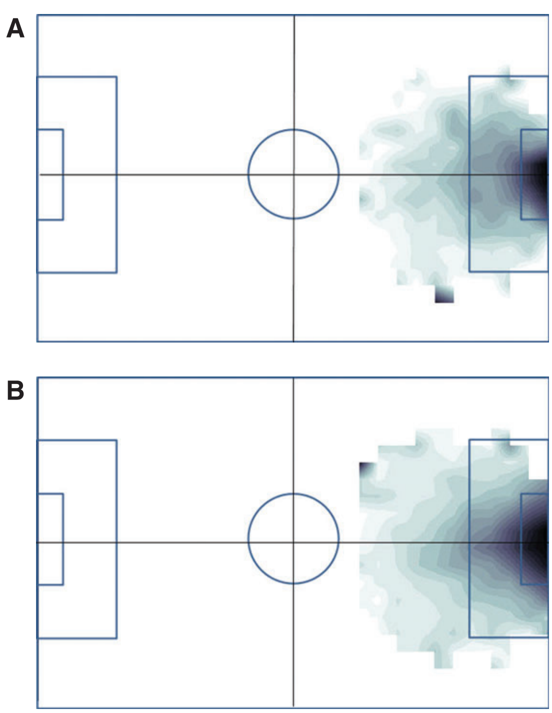

Figure 2 Scoring probabilities by shot start-point.
(A) Empirical probablities of scoring by start-point.
(B) Model probablities of scoring by start-point.
Note: 32×20 grid. Darker areas indicate higher probability of scoring.

start-points. This figure uses a 32 (X)×20 (Y) grid to classify the start-points, with the grid data smoothed by linear interpolation. It is clear that the probability density is approximately symmetrical around the pitch midline.

Figure 2B shows the corresponding density map based on the model predictions for the out-of-sample data calculated during the cross-validation. The map was constructed by computing the expected probability of scoring (*S*) for each shot, and plotting the mean values of *S* within each square on the same 32×20 grid used for Figure 2A, and using the same smoothing procedure. Comparison of the two figures indicates that the empirical and expected density maps are in quite good agreement.

Figure 3A shows the empirical probabilities of scoring mapped on a 20 (Z)×7 (W) grid of smoothed path end-points and the corresponding expectations are shown in Figure 3B.

As expected, shots arriving close to the goal posts and in the corners have an elevated probability of scoring, and again the empirical and predicted density maps are in good agreement. Table 1 shows that the higher-order terms for the end-point parameters in the model have a substantial impact on the probability of scoring. The coefficients for W and W^2 indicate that the minimum probability of scoring occurs when W is close to zero, that is when the ball arrives midway between the goalposts; the coefficients for Z and Z^2 indicate that the minimum probability of scoring occurs when Z is about 1.3, that is when the ball arrives 1.3 m from the ground, around the goalkeepers midriff. Finally, the positive coefficient for the WZ interaction indicates an elevated probability of scoring for balls in the top corners of the goal.

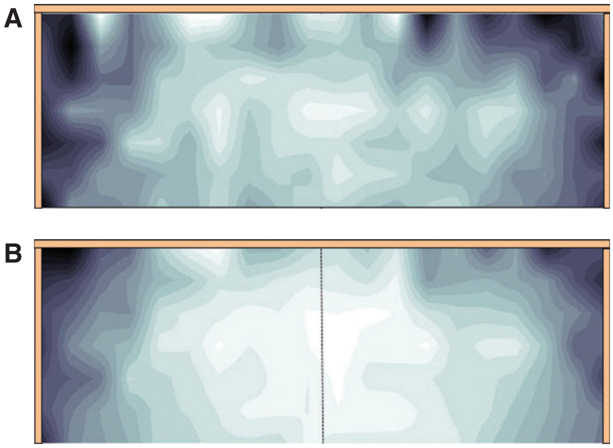

Figure 3 Scoring probabilities by shot end-point.
(A) Empirical probablities of scoring by end-point.
(B) Model probablities of scoring by end-point.
Note: 20×7 grid. Darker areas indicate higher probability of scoring.

The positive coefficient for Far Post indicates that on-target shots to the far post are more dangerous than shots to the near post. This might seem counterintuitive, as the near post is closer to the striker than the far post, and it should therefore be easier to score at the near post. However, conventional wisdom dictates that a goalkeeper should take up a position protecting his near post, so as to force the striker to attempt the more difficult and narrowly angled shot to the far post, and hopefully miss the target. An examination of missed shots provides some support for this conjecture. First, shots were encoded as on or off target and regressed on the style, trajectory and shot location (X, Y) variables; shots to the far post are 7% less likely to be on target than near post shots ($p<0.05$); for shots taken wide of the pitch midline (Y>10 m), the difference is even greater, with far post shots being 21% less likely to be on target than near post shots ($p<0.001$). Next the distance outside the goalpost was calculated for shots that went wide, and regressed on the style, trajectory and start location variables; missed shots to the far post are 2.06 m further away from the goalpost than missed shots to the near post ($p<0.001$). These results indicate that shots to the far post are more difficult than shots to the near post. Nevertheless, conditional upon being on target, a far post shot is more likely to score than a near post shot. (For an interesting game-theoretic discussion of near post/far post strategies, see Moschini 2004.)

The signs of the other model coefficients conform to expectations. Coefficients for the trajectory predictors indicate that shots which swerve, bounce or are deflected are more likely to score presumably because their flight through the air is less predictable and harder to judge. In particular, swerving the ball more than doubles the odds of scoring; this corroborates previous research that the trajectory of a swerving ball is particularly difficult to judge (Craig et al. 2006; Dessing and Craig 2010). Deflection increases the odds of scoring almost threefold.

The style predictor coefficients indicate that shots developed on a fast break have an elevated probability of scoring, presumably because the defending team is likely to have fewer players behind the ball. The same factor might also account for the greater likelihood of scoring from open play as opposed to set play. Assisted shots are considerably less likely to score than unassisted shots, although the reason for this is not clear. Finally, shots taken on the volley are no more likely to score than those that are not.

Next, model probabilities were averaged over each season for each defending team, and the results are shown in Table 3.

The first data column of Table 3 shows the mean probability of the attacking team scoring from an on-target shot; the second data column shows the number of league points scored. The correlation between the model probability of scoring and league points is −0.51 ($p<0.001$). This confirms that goalkeepers playing in teams finishing lower in the league tend to face on-target shots that have a higher probability of scoring (and that are more difficult to save). It is clear that failing to account for this may lead to systematic bias in assessments of goalkeeper performance; specifically, the ability of goalkeepers in weak teams will be underestimated.

4 Measures of goalkeeper performance

From the point of view of the goalkeeper and defending team, it is convenient to think in terms of expected save probabilities, which we denote *ESP*. (*ESP*=1–*S*). The *ESP* of a shot can be interpreted as the probability that it will be saved by an average goalkeeper. Average *ESP*s were calculated for all goalkeepers who defended 50 or more shots separately for each season. *ESP*s varied between 66% and 77%, and a one-way analysis of variance found significant differences between goalkeepers ($F=1.4$, $df=79$, $p=0.006$). This indicates that some goalkeepers are consistently exposed to more challenging shots than others, and that failing to take this into account may lead to biased estimates of performance.

Table 3 Probabilities of conceding a goal and League points.

Defending team	Probability of attacking team scoring	League points
Season 2010		
Manchester United	27.8%	80
Chelsea	24.0%	71
Manchester City	28.6%	71
Arsenal	26.8%	68
Tottenham Hotspur	23.0%	62
Liverpool	25.7%	58
Everton	29.1%	54
Fulham	27.1%	49
Aston Villa	31.3%	48
Sunderland	28.6%	47
West Bromwich Albion	29.7%	47
Newcastle United	31.6%	46
Stoke City	31.6%	46
Bolton Wanderers	30.7%	46
Blackburn Rovers	28.9%	43
Wigan Athletic	32.6%	42
Wolverhampton Wanderers	29.1%	40
Birmingham City	26.2%	39
Blackpool	28.1%	39
West Ham United	27.9%	33
Season 2011		
Manchester City	24.6%	89
Manchester United	24.4%	89
Arsenal	30.3%	70
Tottenham Hotspur	26.8%	69
Newcastle United	28.5%	65
Chelsea	27.1%	64
Everton	30.8%	56
Liverpool	26.6%	52
Fulham	26.5%	52
West Bromwich Albion	26.4%	47
Swansea City	29.2%	47
Norwich City	31.9%	47
Sunderland	28.8%	45
Stoke City	27.7%	45
Wigan Athletic	32.6%	43
Aston Villa	33.3%	38
Queens Park Rangers	28.6%	37
Bolton Wanderers	32.3%	36
Blackburn Rovers	34.1%	31
Wolverhampton Wanderers	27.1%	25

Next, the expected number of shots saved (number of shots defended × *ESP*) was calculated for each goalkeeper/season. The actual number of shots saved minus the expected number is a measure of the goalkeeper's contribution to the team in terms of extra goals saved (or conceded if the difference is negative). Finally, we can calculate a "standardized" saves-to-shots ratio for each goalkeeper;

$$S3R = SSR * \overline{ESP} / ESP \qquad (1)$$

where *S3R* is the standardized saves-to-shots ratio; *SSR* is the actual saves-to-shots-ratio; *ESP* is the expected save probability; and $\overline{ESP}$ is mean expected save probability for all goalkeepers. *S3R* is the saves-to-shots ratio that a goalkeeper would have registered if all the shots he defended were of average difficulty, and it is a normalized measure of performance that can be used as a rating of goalkeeper ability.

5 Evaluating goalkeeper performance

Table 4 reports the raw and model-adjusted performance statistics for each goalkeeper.

In 2010 the goalkeeper contributing most to his team was Ben Foster, who saved almost 10 more goals than expected; in 2011 the most negative contribution came from Paul Robinson, who conceded 13 more goals than expected. Given that the average team in the English Premier League concedes 53 goals in a season, it is clear that variations in goalkeeper performance can have a substantial influence on team results.

To illustrate the use of the *S3R*, consider Joe Hart's performances for Manchester City. Hart's observed *SSR* rose from 71% in 2010 to 75% in 2011. However, when the difference is *ESP*s is taken into account, his *S3R* actually fell, from 78% to 75%. Similarly, Szczesny's observed *SSR* in 2010 was 74% but fell substantially to 67% the following season. This suggests a dramatic drop in performance, but virtually all of the decline was due to a decrease in the expected save probability from 75% to 70%; this indicates a deterioration in Arsenal's defensive play, rather than a loss of form by Szcenzy. In fact, his *S3R* fell by just 1.1%. These examples demonstrate that relying on the uncorrected values of *SSR* can produce misleading conclusions.

6 Conclusions

Examination of mean square errors indicates that the standardized save-to-shots ratio is a more stable measure of goalkeeper performance than the unstandardized metric. For the 15 goalkeepers who played in both seasons, the mean square error between seasons for the

Table 4 Raw and model-adjusted performance statistics.

Goalkeeper	Team	No. shots defended	Actual shots saved	Expected shots saved	Observed SSR	Expected SSR	Goals contributed	S3R
2010 season								
Hart	Man. City	137	107	97.8	78%	71%	9.2	78%
Begovic	Stoke City	122	93	86.5	76%	71%	6.6	77%
Foster	Birmingham	215	169	159.3	79%	74%	9.7	76%
Jääskeläinen	Bolton	165	120	114.6	73%	69%	5.4	75%
Al-Habsi	Wigan	157	112	107.5	71%	68%	4.5	74%
Van der Sar	Man. Utd.	109	81	78.3	74%	72%	2.7	74%
Harper	Newcastle	66	45	44.1	68%	67%	0.9	73%
Schwarzer	Fulham	125	90	89.1	72%	71%	0.9	72%
Cech	Chelsea	127	97	96.5	76%	76%	0.5	72%
Gordon	Sunderland	67	49	49.3	73%	74%	−0.3	71%
Robinson	Blackburn	162	115	115.9	71%	72%	−0.9	71%
Reina	Liverpool	129	95	95.8	74%	74%	−0.8	71%
Howard	Everton	138	97	97.8	70%	71%	−0.8	71%
Green	West Ham	220	158	159.5	72%	73%	−1.5	71%
Mignolet	Sunderland	104	72	72.8	69%	70%	−0.8	70%
Gomes	Tottenham	145	110	111.5	76%	77%	−1.5	70%
Szczesny	Arsenal	57	42	43.0	74%	75%	−1.0	70%
Hahneman	Wolves	65	44	45.2	68%	70%	−1.2	69%
Friedel	Aston Villa	159	105	109.2	66%	69%	−4.2	69%
Kingson	Blackpool	109	74	77.1	68%	71%	−3.1	68%
Hennessey	Wolves	120	81	85.9	68%	72%	−4.9	67%
Carson	West Brom.	138	88	94.5	64%	68%	−6.5	66%
Gilks	Blackpool	94	63	68.7	67%	73%	−5.7	65%
Krul	Newcastle	77	48	53.7	62%	70%	−5.7	64%
2011 season								
Given	Aston Villa	136	97	89.4	71%	66%	7.6	77%
De Gea	Man. Utd.	127	101	94.9	80%	75%	6.2	76%
Stockdale	Fulham	58	44	41.6	76%	72%	2.4	75%
Mignolet	Sunderland	123	92	87.1	75%	71%	4.9	75%
Hart	Man. City	121	96	91.3	79%	75%	4.7	75%
Vorm	Swansea	180	134	128.4	74%	71%	5.7	74%
Howard	Everton	128	92	88.6	72%	69%	3.5	74%
Ruddy	Norwich	193	135	131.0	70%	68%	4.0	73%
Begovic	Stoke	105	78	75.7	74%	72%	2.3	73%
Schwarzer	Fulham	133	101	98.9	76%	74%	2.2	73%
Sørensen	Stoke	62	46	45.1	74%	73%	1.0	73%
Friedel	Tottenham	153	114	111.9	75%	73%	2.1	73%
Krul	Newcastle	154	110	110.1	71%	71%	−0.1	71%
Foster	West Brom.	154	114	114.7	74%	74%	−0.7	71%
Al-Habsi	Wigan	172	115	115.9	67%	67%	−0.9	71%
Cech	Chelsea	119	84	85.8	71%	72%	−1.8	70%
Bogdan	Bolton	114	76	77.9	67%	68%	−1.9	70%
Reina	Liverpool	108	78	80.0	72%	74%	−2.0	69%
Hennessey	Wolves	227	161	166.1	71%	73%	−5.1	69%
Szczesny	Arsenal	124	83	86.4	67%	70%	−3.4	68%
Jääskeläinen	Bolton	98	62	65.6	63%	67%	−3.6	67%
Kenny	QPR	160	106	114.6	66%	72%	−8.6	66%
Robinson	Blackburn	159	94	107.2	59%	67%	−13.2	63%

unstandardized ratio was 2.5 as compared to 1.3 for the standardized ratio, a reduction of almost 50%.

This increased stability across seasons suggests that the *S3R* is a superior estimate of true goalkeeping ability and should be used in preference to the unadjusted saves-to-shots ratio when assessing goalkeeper performance or making recruitment decisions.

References

Bar-Eli, Michael, Ofer H. Azar, Ilana Ritov, Yael Keidar-Levin, and Galit Schein. 2007. “Action Bias among Elite Soccer Goalkeepers: The Case of Penalty Kicks.” *Journal of Economic Psychology* 28:606–621.

Chiappori, P.-A., Steven Levitt, and Timothy Groseclose. 2002. “Testing Mixed-Strategy Equilibria When Players are Heterogeneous: The Case of Penalty Kicks in Soccer.” *American Economic Review* 1138–1151.

Craig, Cathy M., Eric Berton, Guillaume Rao, Laure Fernandez, and Reinoud J. Bootsma. 2006. “Judging Where a Ball will Go: The Case of Curved Free Kicks in Football.” *Naturwissenschaften* 93:97–101.

Dessing, Joost C. and Cathy M. Craig. 2010. “Bending it like Beckham: How to Visually Fool the Goalkeeper.” *PLos One* 5:e13161.

Geisser, Seymour. 1975. “The Predictive Sample Reuse Method with Applications.”*Journal of the American Statistical Association* 70:320–328.

Moschini, GianCarlo. 2004. “Nash Equilibrium in Strictly Competitive Games: Live Play in Soccer.” *Economics Letters* 85:365–371.

Oberstone, Joel. 2010. “Comparing English Premier League Goalkeepers: Identifying the Pitch Actions that Differentiate the Best from the Rest.” *Journal of Quantitative Analysis in Sports* 6:9.

Paul, Prabasaj, Michael L. Pennell, and Stanley Lemeshow. 2013. “Standardizing the Power of the Hosmer–Lemeshow Goodness of Fit Test in Large Data Sets.“ *Statistics in Medicine* 32:67–80.

Pollard, Richard, Jake Ensum, and Samuel Taylor. 2004. “Estimating the Probability of a Shot Resulting in a Goal: The Effects of Distance, Angle and Space.” *International Journal of Soccer and Science* 2:50–55.

Roskes, Marieke, Daniel Sligte, Shaul Shalvi, and Carsten K. W. De Dreu. 2011. “The right side? Under Time Pressure, Approach Motivation Leads to Right-Oriented Bias.” *Psychological Science* 22:1403–1407.

Savelsbergh, Geert J. P., A. Mark Williams, John Van Der Kamp, and Paul Ward. 2002. “Visual Search, Anticipation and Expertise in Soccer Goalkeepers.” *Journal of Sports Sciences* 20:279–287.

Schuckers, M. 2011. “DIGR: A Defense Independent Rating of NHL Goaltenders Using Spatially Smoothed Save Percentage Maps.” In *Proceedings of the MIT Sloan Sports Analytics Conference*.

Sørensen, Henrik, Martin Thomassen, and Morten Zacho. 2008. “Biomechanical Profile of Danish Elite and Sub-Elite Soccer Goalkeepers.” *Football Science* 5:37–44.

Wood, Greg and Mark R. Wilson. 2010 “A moving Goalkeeper Distracts Penalty Takers and Impairs Shooting Accuracy.” *Journal of Sports Sciences* 28:937–946.