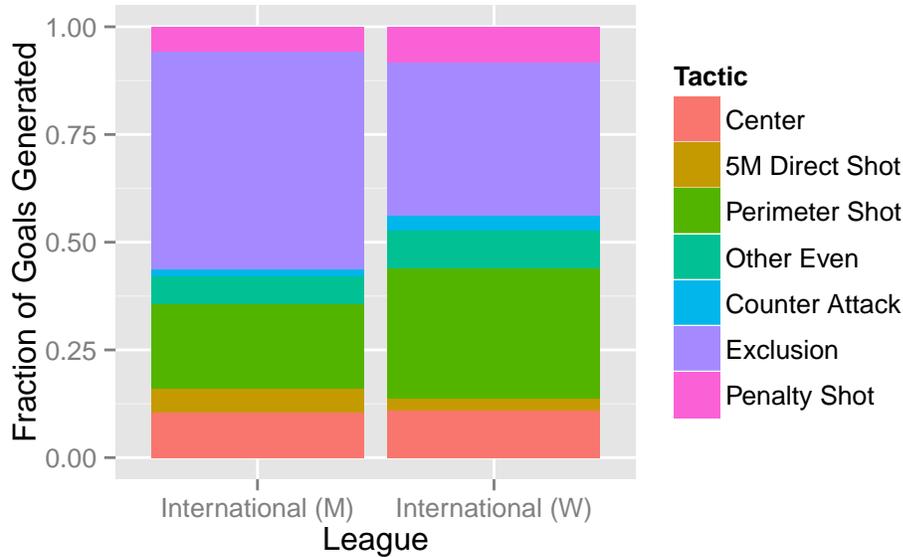


Exclusions: Efficiency over Frequency

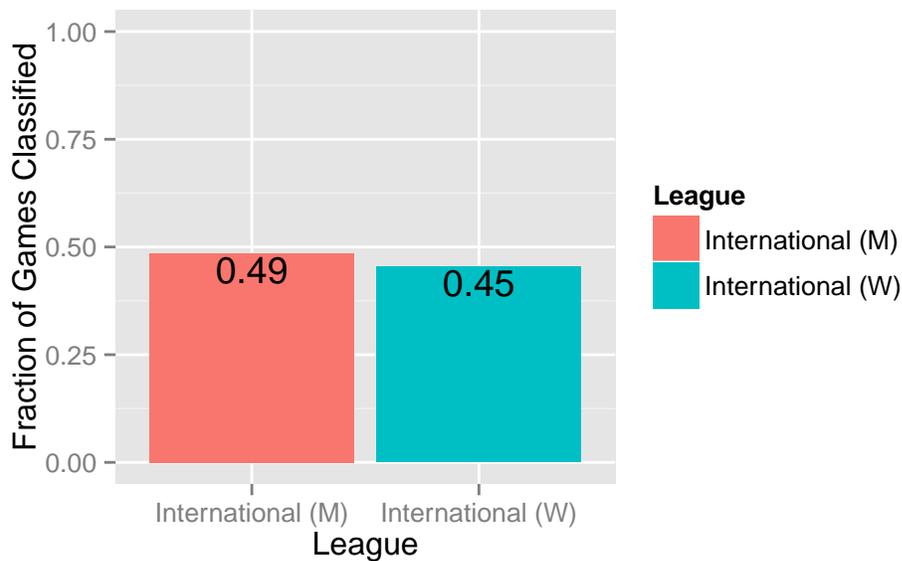
John Mayberry, Water Polo Analytics Group

Exclusions play an undisputedly important role in water polo. To illustrate how important a role they play, the chart below shows a breakdown of how goals were scored in 68 Men's and 44 Women's international water polo contest spanning the 2012 Olympics (Men's only), 2013 World Championships, and 2014 European Championships.

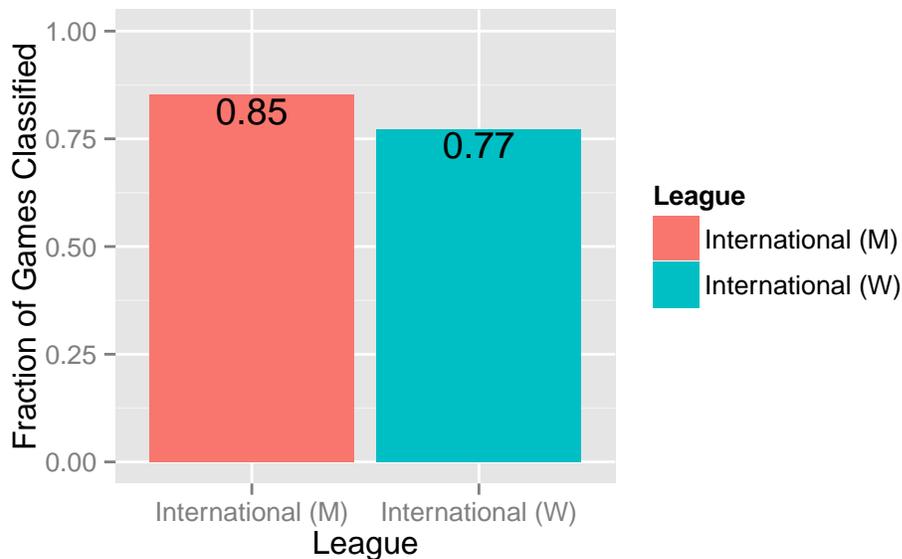


While the distributions of goals scored is significantly different between the two leagues, the two graphs have one thing in common: exclusions generate more goals than any other tactic, accounting for 51% of all goals in Men's and 35% in Women's water polo.

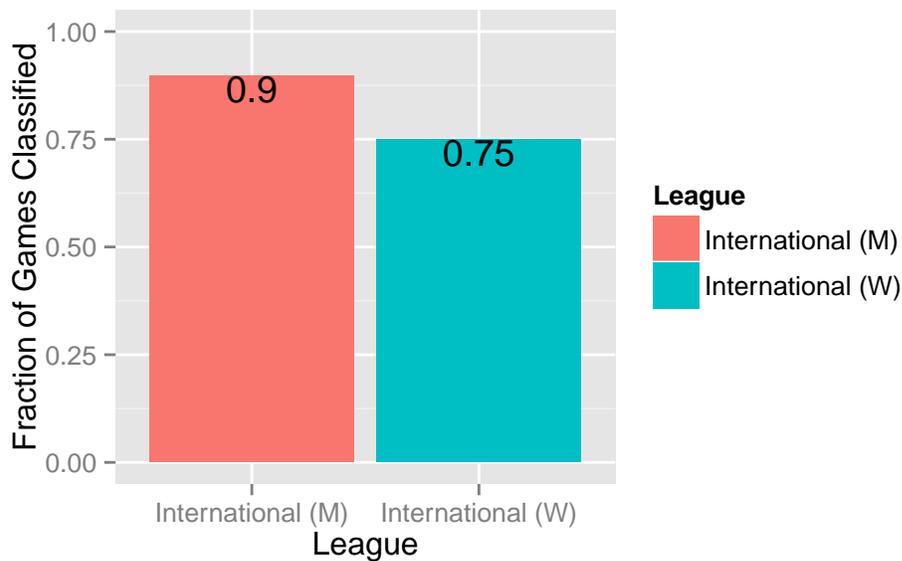
Given the importance of exclusions in generating goals, we might expect that exclusion opportunities should be a good *classifier* of game outcomes: the team with more exclusions should be more likely to win the game and hence, the winner can be classified by looking at exclusion opportunities alone. The data, however, paints a different picture.



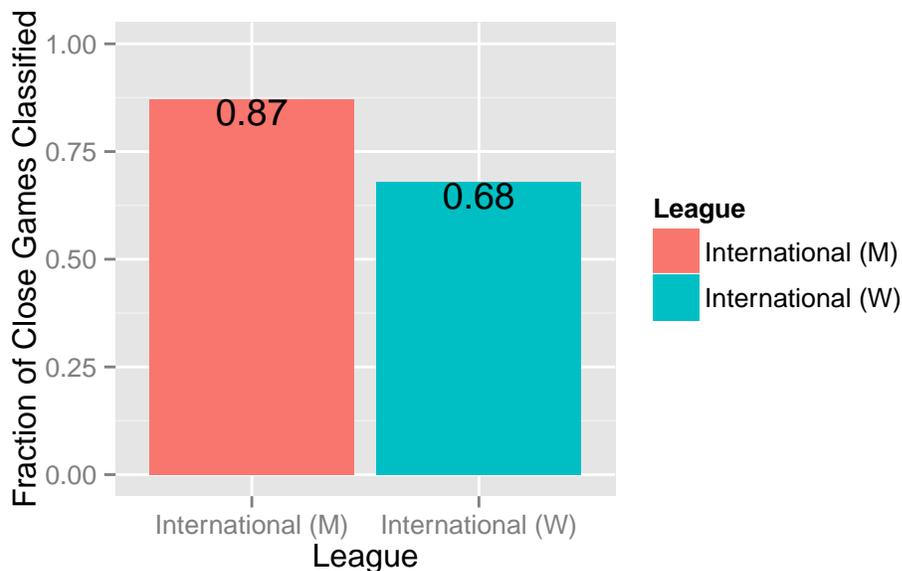
The plot above shows the fraction of games in each league in which the winning team won the exclusion ratio (i.e. received more power play opportunities than the losing team). In both leagues, the fractions are close to (and actually slightly below) 50%! This means that if you had to use exclusion opportunities to predict game outcomes, you should predict that the team with more opportunities will lose. But you would still be wrong about half the time and might as well be flipping a coin to determine the outcome. In contrast, if we look at power play *goals* as a classifier, the picture is completely different.



Now we see that over 80% of all men's games and almost 80% of all women's games are won by the team who scores more power play goals. Since the number of opportunities do not greatly differ, this can only mean that winning teams make more efficient use of the opportunities they are given. We can quantify a team's efficiency by what we call the *Exclusion Conversion Rate* (or ECR), defined as the fraction of exclusion opportunities which the team converted to goals. In other words, you can think of ECR as the probability that a team will score a goal on an exclusion opportunity. ECR correctly classifies 90% of all men's contests and three-quarters of women's as demonstrated by the graph below.



The 15% difference in classifying rates is interesting, but it is not large enough to satisfy the classical ‘p-value < 0.05’ standards of frequentist statistics. In other words, it is unclear if there actually is a difference in the importance of converting exclusions in men’s and women’s water polo in the population of all games or if the difference in our data set can be attributed to the fact that we are looking at only a small subset of all games played in the two leagues and just got an unlucky sample. In any case, it is clear that ECR is important in both leagues. In fact, even if we restrict our attention to close games, which we define as games determined by three goals or less, the importance of winning the ECR battle does not change drastically.



Why does Exclusion Conversion Rate play such an important role in determining the outcome of games? How can we explain the fact that losing teams tend to receive more exclusion opportunities than winning teams when common sense says otherwise? We will examine answers to these questions next time!