

Rearranging the Starters

by Aryeh Bak

(Note: This article was written in 2000, but has not been previously published.)

In response to a comment made by Rob Neyer in his daily column on ESPN.com, I am writing this article to explore whether or not an opportunity exists for a major league baseball team to benefit by arranging its starters in an order other than from best to worst. Neyer wrote the following on August 17, 2000:

"And I really don't see many (any?) managers altering their rotations to get their No. 1 starter to face their opponents' No. 1 starters. Or whoever. Why would they? You don't improve your place in the standings by matching up starters."

Obviously, to arrive at any sort of meaningful conclusion, we would need to make some generalizations. This study will cover two teams, the Mets and the Red Sox (for no reason other than the fact that I am a Mets fan and that Mr. Neyer is covering the Red Sox this season). I have analyzed the effect of different rotations for the duration of a five-game series (each of five starting pitchers pitches one game), which for the purpose of this study will constitute our entire season (i.e., there is no series following ours with which to be concerned). I will assume that starters pitch complete games every single time they start, and that the five starters for each team never get injured. In other words, we're back to baseball before 1920!

I also want to assume that both teams (the Mets and Red Sox) have equal, and exactly average, lineups, stadiums, and bullpens. To that end, I will be using Michel Wolverton's support-neutral win-loss percentages as my

statistic. The support-neutral statistics evaluate starting pitchers by measuring the probability that each outing will lead to the starter getting a win (SNW), the starter getting a loss (SNL), and the starter's team winning (SNVA), given league average support from the offense and the bullpen. (More information can be found at baseball.prospectus.com.)

I will use the five pitchers from each team with the most innings pitched throughout the 2000 season through August 18, 2000. Throughout the rest of this article, whenever I say "Mets" or "Red Sox," I refer actually to Mets (or Red Sox) starting pitchers with 20 other guys who are all of exactly average ability. In other words, baseball John Does.

Here are my pitchers (statistics through games of 8/17/00):

Pitcher	ERA	Support-Neutral Wins	Support-Neutral Losses	Support-Neutral Winning Pct
Red Sox				
Pedro Martinez	1.59	14.8	2.0	.882
Jeff Fassero	4.97	6.5	5.6	.538
Pete Schourek	5.05	6.2	5.9	.512
Tim Wakefield	5.03	4.3	4.3	.495
Ramon Martinez	6.14	5.7	7.4	.438
Mets				
Al Leiter	3.11	10.7	6.2	.633
Mike Hampton	3.38	10.6	7.7	.580
Glendon Rusch	4.42	8.6	7.0	.552
Rick Reed	4.45	8.0	6.7	.544
Bobby Jones	5.08	6.4	6.5	.496

The most important and most difficult assumption that I

had to make regarded computing the likelihood that Pitcher A would win over Pitcher B, given the support-neutral winning percentage of each pitcher. The best method that I could come up with was suggested to me by fellow actuary John Ehrhardt. He calls it the "coin flip" method. In his words:

"...you could reduce this to a coin flipping analogy (my probability professor in college once told me that all probability problems could be solved by flipping coins!):

"Pitcher A has a coin that turns up heads 75% of the time.

Pitcher B has a coin that turns up heads 60% of the time.

"They flip their coins repeatedly until one of them gets heads and one gets tails. The one with heads wins.

"On the first pair of tosses, A has a 30% chance of winning ($75\% \times 40\%$) and B has a 15% chance ($25\% \times 60\%$) and there's a 55% chance of a tie ($75\% \times 60\% + 25\% \times 40\%$). With continued flipping to resolve ties, A will continue to win 30% and B 15%. A will win twice often as B or 67% of the time.

"If you reduce Pitcher B to 40%, then the percentages become A- 45%, B-10%, Tie-45%, resulting in A winning almost 82% of the time."

To put it formulaically, John is suggesting that the likelihood that A is victorious over B is equal to:

where a represents A's winning percentage, and b is B's winning percentage (on a support-neutral basis, of course), or, in other words, that wins are distributed binomially.

The next logical step is to see what the Mets' winning percentage against the Red Sox would be for a particular game based on each combination of starting pitchers.

Likelihood of Mets Winning a Game Based on Starters

		<u>Mets</u>				
		<u>Al Leiter</u>	<u>Mike Hampton</u>	<u>Glendon Rusch</u>	<u>Rick Reed</u>	<u>Bobby Jones</u>
<u>Red Sox</u>	Pedro Martinez	0.187	0.156	0.142	0.138	0.116
	Jeff Fassero	0.597	0.543	0.514	0.506	0.458
	Pete Schourek	0.622	0.568	0.540	0.532	0.484
	Tim Wakefield	0.638	0.585	0.557	0.549	0.501
	Ramon Martinez	0.689	0.639	0.613	0.605	0.558

I have assumed that I have advanced knowledge that the Red Sox manager (being a simpleton) has decided that no matter what, he would be pitching his starters in order, from 1 through 5. As manager of the Mets, I have 120 different combinations of ways that I can arrange my starters (5!), and will attempt to maximize my wins.

I have discovered that if I pitched my starters in order, 1-5, I would be expected to win 2.377 games out of every five. If, however, I arranged my starters 5, 1, 2, 3, 4, I would be expected to win 2.443 games out of every five, the optimal selection of my 120 choices. This difference (.066 per 5 games, or .013 per game) might seem insignificant. If, however, that effect is accumulated throughout a 162-game season, it translates to 2.15 wins more than expected, easily enough to propel a team into the post-season.

Obviously, the numerous generalizations and simplifications need to be factored in before these results can be used in actual seasonal situations.

For one, the order of opposing starting pitchers can be

hard to predict due to off days, rainouts and injuries.

Another factor to consider is that for each slot greater than 1 in which you pitch your ace, there is a 20% chance that that starter will pitch one less game than he would have otherwise. I did include that effect in my study, and it had a minimal effect on my increase in wins.

Instead of my 5,1,2,3,4 rotation adding an extra 2.15 wins, it adds an extra 2.10 wins to what I could have expected by using a traditional 1-5 rotation. (This .05 win "loss" is attributable to one extra Bobby Jones start, at the expense of a Mike Hampton start.) Finally, I have assumed average lineups, stadiums, and bullpens. In reality, of course, that is not the case, and it can significantly affect these results.

While there are many difficulties in applying this sort of analysis to real-life situations, scenarios certainly do exist under which it could and should be used. Additionally, a refinement of the analysis could clarify situations where it should be applied. One obvious area where starting pitching rotations should be manipulated is in the post-season. A very obvious situation is heading into the sixth game of the World Series. A manager will either be in a position where he has to win two games, or win at least one of two games. Given his team's needs, the manager should be able to determine how best to apply his starting pitchers.

Finally, I should note that the selection of the Red Sox as one of the two teams in this study turned out to be a serendipitous one. Any "advantage" that can be gained by manipulating starting pitchers will increase along with the standard deviations of the starting pitchers' winning percentages. The greatness of Pedro Martinez, for instance, on the Red Sox makes it that much more advantageous for the Mets to "sacrifice" the game when he

is opposing them, by starting their fifth starter. If all starting pitchers were of equal ability, then no advantage could be gained by arranging your starters in an order other than from best to worst. As the deviations of the starters that you are opposing increases, the ability to gain by manipulating your own starters increases.

Bill James' response:

1) The statistical method you have used to derive an expected win outcome, given a known winning percentage from pitcher A and a known winning percentage from pitcher B, is one that I myself developed in the 1970s (which is not to say that it was not independently developed by someone else at some other time—it very well could have been). I am very confident that the method does work. Pete Palmer, in response to a letter from me, tested the theory against real data about 1979, and confirmed that it does work.

2) This exact question has also been studied by Dallas Adams and Tom Tippett, among others.

3) I am assuming that the .882 Winning Percentage for Pedro Martinez actually applies not when Martinez *starts*, but merely *when he is on the mound*. It makes a huge practical difference to the application of the data, because, as you have pointed out, "as the deviations of the starters that you are opposing increases, the ability to gain by manipulating your own starters increases." Since Martinez doesn't pitch about one-third of the innings in the games when he starts, the real winning percentage when he starts is more like .760 than .882. This fact very substantially reduces the potential benefit of the strategy.

4) I can't see how in the world you can reach the conclusion that you have reached from the data you have given. You can't; in fact; the

data you have given here basically screams that Neyer is correct.

An accurate synopsis of your data, I think, is that even if one assumes that the starting pitchers never come out of the game, even if one assumes that starting rotations are locked in place for long periods of time, even if one assumes that the “individual pitcher winning percentages” can be known in advance of the season (when in fact they are almost 100% unknown) ... even if one assumes all of these unrealistic advantages, the benefit derived from pursuing this strategy is so small that it could not possibly approach the opportunity costs of abandoning other strategic options, such as matching left-handed pitchers against teams weak against left-handers. In practice, in a real world in which starters come out after six innings, starting rotations must be adjusted 50 times a year for other factors, and the level of performance from each pitcher is largely unknown, it is inconceivable that there could be any benefit whatsoever derived from attempting to manipulate the rotation in this way.

5) You have also flipped the question very suddenly at the end of your article, by arguing that these advantages should be exploited in post-season play. Well, yes, that is true—but real-life managers DO make these kind of “matchup adjustments” in post-season play, every year, and Neyer never suggested that they should not. You’ve got one issue in the premise, and an entirely unrelated issue suddenly sprung upon the reader in the conclusion.

6) Sorry to sound negative. . .the article is intelligent, and the method is essentially sound as far as it goes. I appreciate the research. It’s just that, in my opinion, your conclusion is very obviously contradicted by your research.