By the Numbers

The Newsletter of the Statistical Analysis Committee of the Society for American Baseball Research

Volume 5, Number 3 September, 1993

Greetings

Welcome to all old and new committee members. We have added several new members to the committee roster in the last month or so. I hope you find the committee and its newsletter to your liking. Feel free to make as large (or small) a contribution to the committee as you'd like. And if there are things you would like to see improved, let us know about that too.

The National Convention in San Diego was a complete success. It was well attended, the presentations and logistics were very good, and I was able to meet several committee members for the first time. I hope that you are planning to attend next year's convention in Arlington.

Besides adding several members to the committee, the convention had one other tangible outcome. **Dave Raglin** has agreed to serve as vice-chairman of the committee. For a variety of reasons, SABR wants each committee to have at least two people responsible for seeing that committee business is properly handled. We will let you know what specific "duties" Dave will be handling in the next newsletter. Feel free to contact Dave at 4246 Raleigh Ave. #401, Alexandria, VA 22304, (703) 370-9497. Thanks Dave for volunteering!

This issue of the newsletter contains four interesting articles. A couple of the articles are rather long. While I prefer shorter articles so that a greater number of researchers are able to get their ideas and work out to the public, I will print long articles when the backlog of material is not too great.

Phil Birnbaum kicks off with a methodological piece on designing and using simulations in baseball research.

Don Coffin writes on the worst pitching staffs of all time in response to an article in The National Pastime.

Lawrence Hadley and Elizabeth Gustafson write a provocative article on the merits (or lack thereof) of revenue sharing in baseball, a very hot topic.

Bill Gilbert writes on the surge in offense on display during the 1993 season.

Although we have not depleted our backlog of material, we always encourage submissions for the newsletter.

Please send material, comments, etc., to my address: **Rob Wood**, 2101 California St. #224, Mountain View, CA 94040. My home number is (415) 961-6574, and my daytime number is (415) 854-7101.

****

Notes from committee members

In this section of the newsletter, I will pass along news/ideas/information I receive from committee members.

There were several interesting presentations at the National Convention dealing with statistics and statistical analysis.

Tony Blengino presented a new method for evaluating starting pitcher performance based upon seasonal data on relative control/power factors. If you'd like a copy of the presentation, contact me or Tony at 413 Brooke Ave., Magnolia NJ 08049, (609) 346-2548.

Mark Pankin presented several subtle aspects of baseball using the Project Scoresheet database and his previous work on Markov models. If you'd like a copy of the paper, contact me or Mark at 1018 N. Cleveland St., Arlington, VA 22201, (703) 524-0937.

Walter Szetela presented research on investigating the relationships between several offensive statistics with an eye on identifying the "best" statistic. If you'd like a copy of the paper, contact me or Walter at P.O. Box 853, Point Roberts, WA 98281.

Doug Pappas presented research based upon computer simulations which indicate that the variability in team win totals is much higher than previously believed -- a fact which makes predicting pennant winners a difficult endeavor indeed. If you'd like a copy of the paper, contact me or Doug at 100 E. Hartsdale Ave., #6EE, Hartsdale, NY 10530.

Frank Forthoffer presented a new methodology to determine who are the best hitters. His system is similar to linear weights, but uses more detailed play-by-play data. If you'd like a copy of the paper, contact me or Frank at 2325 Hogan Way, Oceanside, CA 92056.

Frank D'Amico has compiled a record of league home run totals, averages, and ratios for every year from 1876-1935 (Babe Ruth's last season). For example, we find that Babe's 1920 home run total was fully 14.6% of all the home runs hit in the league that season! In addition, Frank has compiled Ruth's annual home run totals vs. team totals (in 1920 Ruth hit more home runs than 14 of the 15 other major league teams). If you'd like a copy of the reports, contact me or Frank at 4460 Overland Ave., Suite 52, Culver City, CA 90230.

Pete DeCoursey writes to state that, though his article "Arm Warnings" which appeared in the previous newsletter was based on Craig Wright's earlier work, the opinions expressed in the article were Pete's own.

Cappy Gagnon writes to share an anecdote concerning the late Allan Roth. Roth never missed a meeting of the Los Angeles chapter of SABR which now bears his name. He once confided in Cappy about an incident which showed his insight and also his frustration with a certain baseball "expert". Before the 1982 NLCS, Roth and a national league manager were asked by ABC to handicap the Braves-Cardinals match-up. The "expert" said that the series would be a contest between the power of the Braves and the speed of the Cards -- a triumph for simple conventional wisdom. Roth, the first statistician employed by a major league team, applied a little more research and analysis to the question. He discovered that in the head-to-head matches that season, the Braves battery had allowed fewer steals than the Cards permitted, and the Cards out-homered the Braves. When faced with the evidence, the pasta-eating diet-drinking manager replied something to the effect that Roth had "never played in the majors, so what could you know?"

****

BASEBALL SIMULATION MADE EASY

by Phil Birnbaum

There have been several attempts to put together a simulation of a baseball game, and all seem to suffer from the same problem -- they're complicated to build and explain.

First, they involve a lot of research; you have to find out hundreds of different facts, like how often runners go from first to third. Second, you've got to write the program, incorporating all your research. And finally, you need to have your facts and your program verified by other sabermetricians before you can expect any of the simulation's results to be considered valid.

The simulation I'll be describing here has the advantage that it's not complicated at all. It doesn't require any calculation of probabilities of game events, so we won't need to argue about whether runners go from first to third 18% of the time, or just 15%. It's also simple conceptually, so you may not need to look at my program at all -- you can just program my method yourself, and verify my programming by reproducing my results.

We'll start by trying to simulate a team full of average batters. By average batters, I mean that they all have batting lines equal to the league average. I used 1988 American League, so that means all nine hitters bat something like this, rounded to 550 at-bats:

	ab	h	2b	3b	HR	BB	K	avg
AL, 1988	550	143	26	3	14	51	88	.259

I've rounded off these numbers for printing here, but not in my simulation.

The typical way to simulate a batting order, the way that's been used in all the sabermetric work I've seen, is to assign a number between 1 and the number of plate appearances -- 601 in the above batting line -- to each event in the batting order. Then, you just repeatedly choose random numbers, and match them to the events.

In our case, we have 100 singles, so maybe the numbers 1-100 represent singles. There are 26 doubles, and so the numbers 101-126 represent doubles. Etc.

To simulate a full game, the typical simulation would pick a random number for the first batter, then the second batter, and so on, until we have three outs and the inning ends. Then clear the bases and keep going, until we have nine innings, at which point we've simulated half (one team's offense) of a major league game.

But with this simulation, we run into that problem of resolving events that aren't explicitly contained in the batting line. How often should we have wild pitches? Balks? Errors? How often should a runner on first advance to third on a single, or score on a double? How frequent are double plays? How often will a runner steal second? Third?

These issues can be resolved with some work and research. We can figure out how often runners on second score on singles, and how often they go only to third. We can find out how often an out is turned into an error, and how often the batter advances past first on the error. We can look up how often runners are picked off first, and how often a runner on third is out trying to score on a passed ball. We can look up all these things, and we can adjust the simulation accordingly. We can have the simulation check before every batter to see if a runner steals, or is caught, or picked off. We can check after every out to see if the out is really an error, and if so, what kind of error and how the runners move. We can draw a random number after every strikeout to see if the catcher drops the third strike and the runner takes first.

Or, we can do it an easier way. Instead of simulating a baseball offense, to correspond to the "real" batting line above, let's say you wanted to simulate a single serving of granola, based on a full box. That is, you're given a box of granola as your model, and you need to produce a bowlful of granola ingredients that models the ingredients in the box. Would you count every ingredient in the granola, and classify them according to shape and size? Would you figure out that exactly 26% of the granola consists of raisins, and that of those, 34% are plump, juicy, "first to third" raisins, while another 10% are wimpy "infield" raisins? Would you calculate that 13% consists of almonds, and 50% oats? Would you then pick a random number, find that it corresponds to a date chunk of a certain size, and then construct such a chunk and add it to your bowl?

Of course not -- you'd just shake the box thoroughly, reach in, and pull out a bowlful, and that bowlful would be a pretty good representative sample of what was in the box. And so that's exactly how this simulation works -- it pulls actual plate appearances out of the box of actual baseball events.

Instead of one box, though, we'll create 24 boxes, and label them for the base/out situation. We'll label box 1 "nobody-out-nobody-on", and box 10 "one-out-runner-on-first", and box 24 "two-outs-bases-loaded", and so on. (We use 24 boxes instead of one so we don't wind up pulling out a double play with two outs, or a 5-4 force out with nobody on base.)

Project Scoresheet tells us there were 89,713 events in the American League in 1988. We'll take each event, each real event, that actually occurred in 1988 play, and put it in its respective box, depending on when it happened. We'll throw in the entire event -- that is, instead of just throwing in "single", we'll throw in "infield single, runner takes third on a throwing error by the left fielder".

When we're done, each box will be full of real events. For instance, we'll find that we've got 5,293 events in box 2, nobody-out-and-runner-on-first. Some of those events are singles, runner to second. Some are doubles, runner scores. Some are passed balls. Some are hit batsman. Some are stolen bases, runner to third when the catcher threw the ball into center field. Some are errors on foul fly balls, batter remains at bat.

Now, to simulate a game, all we have to do is reach into the proper box, stir the contents around a bit, and pull out an event. If it's a double, we don't have to pick a random number to determine which runners score, because it's written right on the event. If it's an out, we don't have to worry about whether it's a double play. If it's a strikeout, we don't have to worry about whether the catcher dropped the third strike. All those contingencies are included in the description of the event. And most importantly, all the events are included in the boxes in the exact proportion in which they occurred in 1988. We don't need to worry about making sure that exactly 25.9% of all the "at-bat" events are "hit" events, because they're already in that proportion inside the boxes. By using the exact real-life events, we've saved ourselves the exhausting task of figuring out all these proportions ourselves.

Okay, let's try simulating a sample inning:

Reach into box marked ...	and pull out
No outs, nobody on	Double.
No outs, runner on second	Groundout, 1-3, runner advances.
One out, runner on third	Walk.
One out, first and third	Sac fly, runner scores.
Two outs, runner on first	Groundout, 5-3. End of inning. One run, one hit one man left.

Does this make sense? All we had to do was draw random events, and keep track of the outs, runners, and runs.

All right, let's use the simulation to answer some real baseball questions: How many runs does this average team score? Actually, we don't have to use the simulation for this at all -- we can just plug the 1988 AL batting line into the runs created formula, which gives us 4.43 runs per 27 outs (not including intentional walks; I removed intentional walks from the simulation and all calculations because the simulation isn't smart enough to figure out when an intentional walk should occur).

I ran the simulation anyway, 91,330 games (about 9 hours worth, while I slept). The simulation produced 4.41 runs per 9-inning game.

Our average team has runners on first and third with one out. How many runs, on average, will they score in the inning?

To simulate this situation, I just started drawing from the first-and-third-one-out box instead of the leadoff box, and ran until three outs, as before, counting the number of runs that scored. After 3000 remainders-of-innings, an average of 1.16 runs each had been scored.

This same question was answered by a previous non-simulation study (I'd name it, but I can't find it now. I think it was in Sabermetric Review, July 1987). That study gave a figure of 1.19, and suggested a 65.1% chance of at least one run scoring in the inning. The simulation gives 65.4%.

How does RBI performance vary by position in the batting order? Well, that depends on the players involved, of course; cleanup hitters drive in more runs not just because they bat higher in the order, but because they're not Alfredo Griffin. But we can get some idea of how much of the difference is caused by the batting slot by running the simulation. I ran 7,500 nine-inning games, keeping track of actual statistics for each spot in the order, per 162 games. The results below represent the same, league-average hitter in each of the nine lineup spots:

	ab	h	2b	3b	HR	RBI	BB	K	avg
Leadoff	686	176	31	4	17	73	63	107	.257
2nd	664	174	31	4	17	75	61	105	.261
3rd	649	171	30	4	17	82	61	106	.263
4th	632	163	29	3	16	86	61	101	.259
5th	619	160	28	3	16	81	57	96	.259
6th	603	156	29	4	15	77	56	96	.259
7th	587	153	27	3	15	75	55	94	.261
8th	570	150	25	4	15	73	54	91	.263
9th	553	143	26	3	13	71	53	88	.258

As expected, RBI totals are larger in the middle of the order, and lower at the top and bottom. The differences, though, are smaller than I would have expected.

One interesting thing to note is the extent to which each of the nine hitters varies from the league average. While the 1988 American League batted .259, our leadoff man hit just .257, while the third-place hitter's average was the highest of the group, at .263.

Some of this is due to random chance, of course, but there is another factor operating. The real-life league batting average goes up (and down) substantially in certain situations. For example, the AL hits for a higher average with no outs and a runner on first:

American League, 1988	AB	H	avg
No outs, nobody on	19280	4880	.253
No outs, runner on 1st	4009	1169	.292

The higher batting averages by some of the batting slots may be due to the fact that those hitters come up more often in situations like this, where batting averages are generally higher. Note that the simulation takes care of this "automatically", due to the fact that we split the events by base-out situation.

Above, in the RBI chart, I listed batting performance per 162 games. If we want, we can instead calculate performance per 550 at-bats instead of per 162 games. Here's that breakdown for RBI:

slot	1	2	3	4	5	6	7	8	9
RBI/550ab	59	62	69	75	72	70	70	70	71

This is interesting; the last five slots are virtually identical, while the leadoff and number 2 batters are lower and the cleanup hitters are higher. I would have expected a gradual decline from the fifth to ninth slot.

If we wanted, of course, we could get this simulation to keep track of all sorts of other things. When a second place batter hits a home run, how many runners, on average, are there on base? Will the second place hitter make the last out of the game more often than the sixth place hitter? Which position in the order comes up most often with the bases loaded? These questions can be easily investigated with the simulation.

So far, we've been talking about a league average team, which isn't really very useful. We can run the average team all we want, but it won't provide us any useful information about a specific player or a specific team. What we need is to be able to simulate a team made up of any offensive players we choose.

We already know that an average team will score an average 4.41 runs per game. What if we want to know about, say, a team of Ruben Sierras? How many runs would a lineup of nine Sierras produce, assuming they hit as Sierra did in 1989?

We obviously can't just simulate a game the same way as before, because the league boxes don't contain events in Ruben Sierra proportions, but in league proportions. Our boxes contain 26 doubles per 601 plate appearances -- Ruben Sierra hits 31.5 (or did, in 1989). It looks like in order to get the boxes to simulate Ruben Sierra, we'll need to add a few doubles to our boxes.

Well, we can't really do that, because we don't know what kind of doubles to add. How many of the added doubles should have a runner on first scoring? In how many should Sierra be out trying to take third base? How many should roll through the outfielder's legs for a two-base error? By trying to add doubles, we run into the same problems as with the random-number simulation -- having to provide the complex details of the general event.

So let's do this -- instead of adding doubles to the boxes, we'll subtract some of everything else. It doesn't matter how big the boxes are, as long as they contain all the events in the proper proportion. When we pull a bowlful of granola out of the box, it doesn't matter if the box is full, or half-full, or even if it's one of those single servings where you pour the milk right into the box -- as long as the stuff inside is the same recipe.

Let's go back to granola, but for now, we'll simplify it just by considering raisins (hits) and flakes (outs). In the league boxes, we have raisins in the proportion of, say, 100 raisins per 300 flakes, for a .250 batting average. Now, if we need to change the boxes to bat .300, we need 129 raisins per 300 flakes. But raisins are hard to add, because we don't know exactly what kind of raisins. So we'll remove some flakes. If we remove 67 of every 300 flakes, we'll leave those same 100 raisins per 233 flakes, for a .300 batting average.

To adjust for the categories other than hits and outs, we'll do the same sort of thing for every event in the boxes. Sierra hits a lot of triples, so maybe we leave in all the triples. He hits lots of home runs, but not as many as triples, relative to the league, so we'll probably have to take out some home runs. And outs, we'll have to remove lots and lots of outs, because Ruben Sierra makes a lot fewer outs than the league.

The procedure by which all this is calculated is in the appendix [available from the author or from the editor], which you'll want to read if you decide to create this simulation yourself. Or you can trust me that what I actually did remove does leave the boxes containing Sierra-like proportions of events. The calculation is not particularly complicated.

One more thing before we go on -- when I say we removed the events from the boxes, I was simplifying a bit. What we do instead of (for example) removing 41% of the strikeouts, is this: when we pick an event out of a box, if it's a strikeout, we ignore it 41% of the time (determined by a random number) and pick another event. This has the same effect as literally removing the strikeouts, but you don't have to worry about which boxes to remove them from, or whether you removed too many catcher-drops-third-strike and not enough runner-picked-off-after-third-strike-on-throw-by-catcher, etc.

So having described the process, I should prove that my programming worked. Here's the result of running 6,365 games of a Ruben Sierra lineup:

	ab	h	2b	3b	HR	BB	K	SB	CS	avg
simulated	634	195	35	14	29	45	82	?	?	.308
actual	634	194	35	14	29	44	82	8	2	.306

What do you think? I forgot to print out the simulated steals and caught stealings, but I'd say it looks OK anyway.

A few more questions: How many runs would that team of Sierras score? 6.65, the simulation said. Runs Created says 7.08, using the SB version of the runs created formula, and 25.5 batting outs per game.

This result perhaps casts doubt on the validity of the simulation -- it appears to be off by almost half a run. However, I feel the simulation result is more accurate than runs created. Bill James suggests, in the 1985 Abstract, that RC has problems with high-performance batting lines. In fact, the RC formula has never, to my knowledge, been tested against very high-scoring teams or individuals, since it's difficult to get a team of Ruben Sierras together to play 162 games so you can find out how many runs they do actually score.

When I ran the simulation against batting lines that are closer to the league average, the results were much, much closer. I hope to provide details in a future paper.

How many times would a team of Ruben Sierras get shut out? Actually, the Sierras would be shut out less often than they would score 17 (or more) runs, 1.4% to 1.6%. That's about twice per season each.

Can you simulate a lineup of nine different players? Sure. All you have to do is reproportion the boxes between hitters; that is, take a whole bunch of home runs out of the box when Jose Oquendo comes up, and put them back in before Pedro Guerrero's at-bat. Here's the results of running 10,643 games of a hypothetical 1988 Blue Jay lineup. All the players' batting stats have been normalized to their actual number of 1988 at-bats, just to make the numbers easier to read, and to show that the simulation did what it was supposed to.

	ab	h	2b	3b	HR	RBI	BB	K	SB	CS	avg
Moseby	472	114	17	7	10	49	67	93	27	8	.241
Gruber	569	160	34	5	16	77	38	92	21	5	.280
Bell	614	166	27	5	23	94	26	65	4	2	.270
McGriff	536	152	34	4	35	101	81	147	5	1	.283
Fernandez	648	187	42	4	5	76	45	64	16	6	.289
Whitt	398	101	11	2	16	60	65	40	4	2	.253
Lee	381	112	16	3	2	46	26	63	3	3	.294
Leach	199	55	13	1	0	22	18	27	0	1	.277
Campusan	142	31	10	2	2	17	9	33	0	0	.217

The simulated stats are real close to the actual '88 stats of the players listed; you can look them up if you want. This lineup scored 4.76 runs per game.

Can we use this simulation to answer the question of which lineups produce the most runs? Well, we can run a whole bunch of lineups, and see which lead to the highest scoring. For example, I ran the same Blue Jay team as above, but in an order where you would expect a lot fewer runs to be scored:

	ab	h	2b	3b	HR	RBI	BB	K	SB	CS	avg
Lee	381	112	16	3	2	35	26	62	3	4	.293
Leach	199	56	13	1	0	18	17	27	0	1	.280
Moseby	472	114	17	6	10	60	70	96	27	7	.242
Campusan	142	31	10	2	2	19	10	33	0	0	.218
Whitt	398	100	11	2	15	60	65	38	5	3	.251
Bell	614	165	28	5	23	92	32	66	4	2	.269
Gruber	569	158	35	5	15	75	39	91	20	4	.279
McGriff	536	152	36	4	34	98	81	150	6	1	.283
Fernandez	648	185	42	4	5	78	46	66	15	5	.285

This lineup scored 4.63 runs per game, or .13 less than the one above. Multiply that by 162 games, and you get 21 runs, or two wins per season.

You could run this again for any other lineup you wanted, or even for all 362,880 possible lineups. This would take a fair chunk of time, though, at several hours a pop.

And even if you had that much time, or that much faster a computer, you'd still have the problem of statistical significance. Even after running a large number of games -- the first lineup played 10,643, the second 7,425 -- the .13 difference still has a standard deviation of .046. That means the difference, though likely to be near .13, could be much higher or lower. In fact, the 95% confidence interval is (.04,.22), which means the difference between the two lineups may be as few as 6.5 runs per season, or as many as 36 runs.

However, the .13 is reasonably close to other studies, which used different methods or a different simulation. A study by Doug Bennion in the August, 1987 Baseball Analyst found that batting Manny Lee cleanup and Jesse Barfield ninth led to a loss of 10-15 runs over a full season. (This is using stats from 1986, when Barfield was good.)

Does the simulation handle batter-pitcher matchups? Sort of. As described, the simulation assumes a league-average pitcher; that is, every hitter hits with the same event probabilities every at-bat. You can adjust for a specific pitcher by using the Log5 method.

(The Log5 method is used to answer the question, "If the batter hits a home run (or a double, or a groundout, or whatever) X percent of the time, and the pitcher allows a home run Y percent of the time, and the league as a whole allows home runs Z percent of the time, how often will the batter hit a home run against this pitcher?" See the 1983 Abstract for details.)

To adjust for the opposing pitcher, don't adjust the boxes based on the batter's raw stats. Instead, use the Log5 method to compute the batter's projected stats against this pitcher, then adjust the boxes based on these computed stats.

So the simulation seems to work. But why is it any better than other simulations? Why is it better than table game simulations, or random number simulations?

Well mainly, it's easy to understand and verify. From the way I've described it, and the technical description in the appendix, it may seem the opposite, that this simulation is a lot more complicated and mathematical than other simulations, or even APBA cards. But really, that's not the case. It may be a little more complicated to set up, but intuitively, I think it's pretty simple. You just pull actual, real-life events out of a box, and make sure they're in proportion to the stats of the player in question.

Here, in this article, I've described the simulation in full in only a few pages. Could you describe APBA that simply, describe how you have to construct a card, and what each of the numbers mean, and prove that coupled with the events on the board, they accurately define a player's performance? Wouldn't you have to show that yes, runners go first-to-third the proper proportion of the time, and there is the right number of catcher's interference, and triple plays, and foul pop-ups dropped for errors, and everything? And wouldn't you need to do this for every base-out situation, to show that yes, the simulation matches the actual 1988 AL average of .356 (really!) with runners on second and third with one out?

And even if you could, and did, a simulation is no good for research unless it's simple, because other sabermetricians have to be able to verify the simulation and duplicate your results. You could write the perfect simulation using another method, but you'd have a computer program thousands of lines long, which nobody else could use, and a hundred pages of explanation, which nobody could understand. This simulation has the virtue of being simple to explain, and simple to construct. (You can probably write a simulation almost equivalent to mine just from these past few pages of description.)

Of course, there are some things the simulation doesn't do. It doesn't compensate for faster of slower runners on base. It doesn't adjust for whether the infield is in or back. It doesn't know when to time an intentional walk or a stolen base. Any of these features would have to be added; however, it's at least as easy to add them to this simulation as to any other.

If you have any questions, drop me a line, and I'll try to answer them. An appendix to this article which explains how to adjust the boxes to simulate a particular batter is available upon request.

Phil Birnbaum, 322 Patricia Ave., Willowdale, Ont., Canada, M2R 2M5, (416) 222-9352. *Thanks to John Matthew IV for his assistance with this paper.*

****

WHO *REALLY* HAD THE WORST PITCHING STAFF OF ALL-TIME?

by Don Coffin

In an article in the 1993 issue of *The National Pastime*, John Thom argues that the 1930 Philadelphia Phillies had the worst pitching staff of all time. And, lord, were they bad. They gave up 1199 runs, 7.68 per game[1]--no other team in the 20th century has given up more than 1064 (the St. Louis Browns, in 1936, followed closely by the 1936 A's, who gave up 1045). They had an earned run average of 6.71, worsting the 1936 Browns (6.24) and A's (6.08) by a substantial margin. And, of course, they finished 40 games out of first, despite being fourth in runs scored [944, trailing the St. Louis Cardinals (1044), the Chicago Cubs (998) and the New York Giants (959)]. But the worst pitching staff ever? I don't think so. In fact, I don't think the 1930 Phillies were even the worst staff in the history of the Phillies. That honor (?) belongs to the 1919 team.

In 1930, the other seven teams in the NL gave up an average of 832 runs per team. The Phillies were a mere 44% worse than the average. (Call this ratio M/A, most-to-average.) The best pitching staff in the league in 1930, the Brooklyn Dodgers gave up 738 runs. The

1. Thom mistakenly credits them with giving up 7.71 runs per game; however, the Phillies played 156, not 154 games in 1930.

Phillies were only 62% worse than the best team. (Call this ratio M/L--most-to-least.) Neither of these is the worst performance in the 20th century history of major league baseball.

Two teams--the 1911 Boston Braves and the 1915 Philadelphia Athletics--had higher M/A ratios. The Braves managed an all-time worst 1.56, while the A's were right in there at 1.54. Seven other teams have managed M/A ratios of 1.4 or worse (see Table 1). While the 1930 Phillies manage a third place finish in this category, they trail the leaders substantially, and aren't much worse than the other seven teams on the list.

The Phillies rank only 10th in the National League--and 17th in the majors--in M/L ratio since 1900. At the top (bottom?) in the NL were the 1911 Braves, with a mere 1.88 M/L ratio--they gave up 88% more runs than the best team in the league that year (the New York Giants managed to yield only 542 runs that year). The 1909 Cardinals ran second here, with an M/L ratio of 1.87.[2] And the amazing 1939 Browns managed an M/L ratio of 1.86, or 479 more runs than the Yankees (the 1930 Phillies, by way of comparison, gave up only 461 more runs than did the Dodgers that year). The complete list, down to the Phillies with their meager 1.62 M/L ratio, is in Table 2. In fact, the 1930 Phillies didn't even perform the worst in franchise history--they had only the fourth-worst M/L ratio for the franchise (the 1919 team, at 1.74 did the worst).

One item of interest here is that no expansion team is in the top 17 in M/L ratio. The Rockies, of course, appear to be on a pace to change that this year.

Based on these comparisons, I'd rank the 1930 Phillies as no worse than the fourth worst pitching staff in history. Let's call the roll:

- 1911 Boston Braves, with an M/A of 1.56 and an M/L of 1.89, both all-time worsts.
- 1915 Philadelphia Athletics (1.54 and 1.81).
- 1919 Philadelphia Phillies (1.43 and 1.74)
- 1930 Philadelphia Phillies (1.44 and 1.62).

I would want to consider the 1953 Tigers as well (1.41 and 1.69), but let's give fourth worst to the 1930 Phillies. So the 1930 Phillies were truly, awesomely bad. To have the fourth worst pitching in major league baseball is a real achievement. But they were worsted not only by three other teams, but by one other Phillie squad, the 1919 staff that had a 4.14 ERA against a league average of 2.91 (42% worse than the league, while the 1930 Phillies were only 35% worse than the league--6.71/4.97.)[3]

It's important to note that none of this adjusts for park effects. In *Total Baseball*[4], Pete Palmer calculates park-adjusted ERAs and then normalizes to the league average. Based on his adjusted ERAs, the worst pitching performance of all time was turned in by the 1915 Philadelphia Athletics, with a normalized, park-adjusted ERA of 68, suggesting they were 32% worse than the league average. (The 1911 Braves are one of the worst half-dozen or so teams, with an adjusted ERA of 75.)

Finally, if we look at the numbers in Tables 1 and 2, we cannot but be impressed by the awfulness of both the Phillies' and the Athletics' pitchers at various points in this century. These two teams combine for six of the 10 worst M/A ratios (three each) and eight of the 17 worst M/L ratios (splitting them again, four each). Now, part of this is surely a result of the lack of quality on their pitching staffs. But just as surely, part of it must be a park-effect. For example, the 1930 Phillies gave up 44% more runs than the average of the rest of the league. However, Palmer and Thorn list them at a normalized (park-adjusted) ERA of just 19% worse than the league. Table 3 shows how the Phillie and Athletic teams ranked on normalized, park-adjusted ERAs for their worst years.

The ten teams had an average M/A ratio of 1.40, suggesting that they were about 40% worse than the rest of their leagues. However, their normalized, park-adjusted ERAs averaged 78, suggesting that they were really only about 22% worse than average. What this adjustment tells us is that bad pitching staffs in hitter's parks will look even worse than they are.

Did the 1930 Phillies have a good pitching staff? By no means. I rank them as the fourth-worst of all time. But they were not the worst. People think of them as the worst because they gave up so many runs. But the fact is that the entire league got shelled in 1930. And while the

2. The Cardinals managed this, despite an M/A of 1.34 because the Cubs gave up only 390 runs that year, less than 2/3 the league average.

3. The 1919 Phillies also finished fourth in the league in runs scored, and also finished last--47.5 games behind, in a war-shortened season. Projected over 154 games, they would have been 53 games out at the finish.

4. John Thorn and Pete Palmer (eds.), *Total Baseball*, Warner Books, 1993.

Table 1: The Ten Worst M/A Ratios in the 20th Century		
Team	Year	M/A
Boston Braves	1911	1.56
Philadelphia Athletics	1915	1.54
Philadelphia Phillies	1930	1.44
Philadelphia Phillies	1919	1.43
Philadelphia Athletics	1916	1.42
Philadelphia Athletics	1954	1.42
Philadelphia Phillies	1923	1.41
Washington Senators	1909	1.41
Detroit Tigers	1953	1.41
New York Yankees	1908	1.40
Source: Thorn and Palmer, *Total Baseball*, Warner Books, 1993.		

Table 2: The Seventeen Worst M/L Ratios in the 20th Century		
Team	Year	M/A
Boston Braves	1911	1.88
St. Louis Cardinals	1909	1.87
St. Louis Browns	1939	1.86
Brooklyn Dodgers	1905	1.83
Philadelphia Athletics	1915	1.81
Philadelphia Phillies	1919	1.74
Philadelphia Athletics	1954	1.74
Boston Braves	1906	1.70
Brooklyn Dodgers	1944	1.70
Detroit Tigers	1953	1.69
St. Louis Browns	1910	1.68
Boston Braves	1907	1.67
Kansas City Athletics	1964	1.67
Philadelphia Phillies	1904	1.65
Kansas City Athletics	1955	1.64
Philadelphia Phillies	1945	1.63
Philadelphia Phillies	1930	1.62
Source: Thorn and Palmer, *Total Baseball*, Warner Books, 1993.		

Table 3: Normalized and Park-Adjusted ERAs for the Phillies and the Athletics

Team	Year	ERA*	M/A
Philadelphia Phillies	1904	79	1.34
Philadelphia Athletics	1915	68	1.54
Philadelphia Athletics	1916	73	1.42
Philadelphia Phillies	1919	78	1.43
Philadelphia Phillies	1923	86	1.41
Philadelphia Phillies	1930	81	1.44
Philadelphia Phillies	1945	83	1.30
Philadelphia Athletics	1954	75	1.42
Kansas City Athletics	1955	78	1.39
Kansas City Athletics	1964	81	1.30
Average		78	1.40

*ERA is park-adjusted and then normalized (compared to a league average of 100). A lower normalized, adjsted ERA represents a worse performance.

Source: Thorn and Palmer, *Total Baseball*, Warner Books, 1993.

Phillies were awful, I think that at least three teams would have given up even more runs, had their pitchers been working in the conditions prevailing in the National League in 1920. I think the 1911 Braves would have given up somewhere around 1350 runs under 1930 conditions; the 1915 Athletics, something around 1300 runs; and the 1919 Phillies about 1240.

We always have to put performance in context. The context, in the National League in 1930, is that everyone scored and gave up a lot of runs. I think John Thom failed to recognize this context adequately, and in failing to put the 1930 Phillies in context, he mis-identified them, only slightly, in labeling them the worst pitching staff in history.

Don Coffin, 3400 Broadway, Indiana University NW, Gary, IN 46408; (219) 980-6630

****

SHARING LOCAL MEDIA REVENUE AND THE COASE THEOREM

by Lawrence Hadley & Elizabeth Gustafson

The current hot business topic in Major League Baseball is the sharing of local media revenue between the large and small media-market teams. Almost everybody in the baseball world seems to favor this Robin Hood idea. The opposition seems to be limited to the owners of the large-market teams, and they are often portrayed as short-term profit maximizers unable to appreciate the overall "best interests of baseball."

Underlying the idea of revenue sharing is the notion that large media-market teams have an unfair advantage over small media-market teams in their ability to compete in the labor market for top players. Many believe that revenue sharing would lead to more equal competition in the labor market which would translate into more equal competition on the field. More extreme voices have claimed that without revenue sharing, the small media-market teams will be unable to survive financially. The concern is the ability of these teams to meet the payroll required to keep a competitive team on the field.

I. THE COASE THEOREM

Large and small cities will always be with us. Obviously, revenue sharing will not change the sizes of cities and the relative importance of their media. Interestingly, revenue sharing is also unlikely to change the allocation of players between the large and small market teams. To understand this unusual proposition, it is necessary to understand an economic theorem developed by Ronald Coase[5]. Coase demonstrated that in the case of profit-maximizing behavior, an alteration of property rights will not effect the allocation of resources, but only the distribution of wealth.

In this discussion, the existing property rights are those of each team to 100 percent of their own local media revenue. The proposed alteration of property rights is some form of sharing this revenue. Clearly we expect such an alteration to redistribute wealth from the large to the small market teams. Revenue sharing should cause the profits of small market teams to rise and those of large market teams to fall. Also, the market values of all teams would be expected to move toward greater equality after the redefinition of these property rights.

In contrast, no reallocation of players would be expected. As long as teams can exchange players (exchanges occur via trades and free agency in the current system), players will be employed by the team where their economic value is greatest. This is true regardless of the sharing of local media revenue.

The economic value of a player obviously depends upon his ability to generate revenue for the owner-- mostly via ticket sales. Suppose Barry Bonds is most valuable to the San Francisco Giants in comparison to all other teams. This greater value may be due in part to the fact that his father once played for the Giants and is currently a team coach. In large part, it is due to the fact that San Francisco is the center of a large metropolitan area with a well-developed local media. The reassignment of team rights to local media revenues will not change Barry Bonds' value to the Giants versus a team like the Pittsburgh Pirates. He will still sell more tickets in San Francisco than in Pittsburgh, and he will remain more valuable to the Giants. Therefore, they would be expected to make the highest bid for Bonds, and he is expected to continue playing for them as long as his economic value is greatest to that team.

---

5. R.H. Coase, "The Problem of Social Cost," *Journal of Law and Economics*, Vol. 3 (October 1960), pp. 1-44. Coase recently won the Nobel Prize in Economics for this work.

The implication of the Coase theorem is that a reassignment of property rights to local media revenues will not impact on the allocation of players between teams and therefore should not impact on the competitive balance between the teams. In general, the same types of players will have their greatest economic value with the same types of teams, and this redistribution of revenue will not cause any shift in the number of star players between large and small market teams.

Our analysis must be qualified with two assumptions. First, we assume that local media revenue depends primarily on the size of the media markets in which teams play their home games rather than on team performance. Second, we assume that team owners behave as profit maximizers. If instead, owners are willing to sacrifice profits in order to maximize the chance of winning championships, access to more revenue may cause small market teams to buy more stars which in turn may cause them to become more competitive.

II. THE EVIDENCE

We define the small media-market teams in Major League Baseball to include the Cincinnati Reds, Cleveland Indians, Kansas City Royals, Milwaukee Brewers, Minnesota Twins, Pittsburgh Pirates, St. Louis Cardinals, and Seattle Mariners. We define the large media-market teams to include the California Angels, Chicago Cubs, Chicago White Sox, Los Angeles Dodgers, New York Mets, New York Yankees, Oakland A's, and San Francisco Giants. We also define a third category of teams that play in near-large media markets. These teams include the Atlanta Braves, Boston Red Sox, Detroit Tigers, Houston Astros, Philadelphia Phillies, and Toronto Blue Jays. The Baltimore Orioles, Montreal Expos, San Diego Padres, and Texas Rangers are difficult to classify for various reasons, so these teams are omitted from our analysis. Clearly, minor variations on this classification scheme are plausible, but we believe our scheme generally captures the concept of large versus small media markets.

The data in our analysis include the 1982 through 1992 baseball seasons. These eleven seasons clearly fall within the free agent era. There are two types of data typically used to assess competitive balance. The first relates to turnover in the winners of division championships, league championships, and the World Series. The eight small market teams have won 12, 8, and 5 championships respectively. The eight large market teams have won 16, 6, and 3 championships respectively, and the six near-large market teams have won 14, 6, and 2 championships. On a per team basis, the small-market teams lead in World Series competition with 0.625 per team and are tied for first with the near-large teams in league championships at 1 per team. In division championships, the small-market teams have 1.5 per team, the large market teams have 2 per team, and the near-large teams have 2.33 per team.

The more comprehensive indicator of competitive balance is regular-season winning percent. The winning percent for each team from 1982-92 is presented in Table 1, and the percents are cumulated for each of the three groups of teams using total wins and games played by these groups. The small-market teams have a cumulative winning percent of .489 while the large-market teams and near-large teams have cumulative winning percents of .509 and .508 respectively.

Finally, data on 1993 team payrolls are presented in Table 2. The means are also presented for the three groups of teams. These data indicate that the large market teams and the near-large market teams spend approximately $35.5 million per team, while the small market teams spend approximately $28.4 per team. The difference on average is just over $7 million per team which is slightly more than the current cost of one elite free agent player.

III. CONCLUSIONS

A .509 season versus a .489 season is a winning percentage differential of .020. This is our estimate of the inherent long run advantage of large media-market teams over small market teams. A .509 winning percent translates into a 82-79 won/loss record (one rain-out), while a .489 percent translates approximately into a 79-83 record.

From a baseball viewpoint, this .020 differential translates into 3.5 games in the standings over a 162 game season. This is a sufficiently small advantage so that well-managed small market teams may expect to be in pennant races with some regularity. In the long run, the large market teams are expected to win more games and therefore win more championships. But this long run differential due to city size is sufficiently small that the small market teams can compete.

TABLE 1

MAJOR LEAGUE BASEBALL TEAM RECORDS: 1982-1992

Team	Wins	Losses	Total Games	Winning Percent
Small Market Teams				
Cincinnati	881	899	1780	.495
Cleveland	789	993	1782	.443
Kansas City	908	872	1780	.510
Milwaukee	905	874	1779	.509
Minnesota	874	908	1782	.490
Pittsburgh	892	887	1779	.501
St. Louis	929	852	1781	.522
Seattle	794	987	1781	.446
Total	6,972	7,272	14,244	.489
Large Market Teams				
California	900	882	1782	.505
Chicago Cubs	865	910	1775	.487
Chicago White Sox	901	879	1780	.506
Los Angeles	912	867	1779	.513
New York Mets	948	831	1779	.533
New York Yankees	906	873	1779	.509
Oakland	939	843	1782	.527
San Francisco	874	908	1782	.490
Total	7,245	6,993	14,238	.509
Near-large Market Teams				
Atlanta	838	938	1776	.472
Boston	924	857	1781	.519
Detroit	933	848	1781	.524
Houston	886	896	1782	.497
Philadelphia	858	922	1780	.482
Toronto	986	795	1781	.554
Total	5,425	5,256	10,681	.508

TABLE 2

1993 MAJOR LEAGUE BASEBALL TEAM PAYROLLS

Team	Team Payroll[a]
Small Market Teams	
Cincinnati	$42,851,167
Cleveland	$15,717,667
Kansas City	$39,802,666
Milwaukee	$22,948,834
Minnesota	$27,284,933
Pittsburgh	$24,240,670
St. Louis	$22,615,334
Seattle	$31,461,333
Mean	$28,365,326
Large Market Teams	
California	$27,230,334
Chicago Cubs	$38,303,166
Chicago White Sox	$34,568,166
Los Angeles	$37,833,000
New York Mets	$38,350,167
New York Yankees	$40,405,000
Oakland	$35,565,834
San Francisco	$34,567,500
Mean	$35,852,896
Near-large Market Teams	
Atlanta	$38,131,000
Boston	$36,608,583
Detroit	$36,548,166
Houston	$28,854,500
Philadelphia	$26,812,334
Toronto	$45,747,666
Mean	$35,450,375

[a] These payrolls were reported in USA Today, April 5, 1993, page 4F.

From a statistical viewpoint, we can consider the 1982-92 seasons as a sample of all possible seasons and test whether the observed difference in winning percents is so large that it is unlikely to be due to random variation. First, we performed a test on the difference in cumulative winning percents of the eight large market teams versus the eight small market teams. Second we performed a test on the difference in means of the single-season winning percents of the teams in the two groups. Both tests show that the large market teams have significantly higher winning percents at the .05 level of significance. This indicates that, although the .020 differential is small, it is likely to be systematic rather than random.

Some large market teams will continue to attempt the purchase of a team championship in the free agent market. The data in Table 2 make this point clear. However, the recent experiences of the Angels, Dodgers, Mets, and Yankees demonstrate that this strategy is not always successful. Also, there is a special satisfaction for the fans in small cities when David slays Goliath as in the 1990 World Series.

It has been suggested that from a profit viewpoint, this differential is healthy for baseball. National media revenues may be greater in the long run if the large market teams win more often because post-season TV ratings are higher when they compete in post-season play[6].

Our point is that the sharing of local media revenues will not eliminate the advantage of the large media-market teams unless a centralized system for the assignment of players is implemented. For example, the commissioner could pick new teams every season, and player exchanges during the season could be made illegal. Obviously it is better for Major League baseball to live with the .020 advantage of large market teams that is inherent in a free market system.

Finally, the only result of revenue sharing would be the transfer of wealth as the market values of teams in large media markets decline and the market values of teams in small media markets increase. Nobody feels sorry for the wealth position of baseball owners in general regardless of the size of their team's market. But it is difficult to think of any good reason that owners of large market teams should be forced to transfer a part of their wealth to the owners of small market teams. How would this serve the "best interests of baseball"? Jerry Reinsdorf has been quoted: "I bought the Chicago White Sox--not the Seattle Mariners." His point should be well taken!

Lawrence Hadley and Elizabeth Gustafson are Associate Professors of Economics at the University of Dayton, Dayton, OH 45469-2240, (513) 229-2403

6. Andrew Zimbalist, *Baseball and Billions*, Basic Books, New York, 1992, pp. 101-104.

****

1993: A TRIPLE MILESTONE YEAR

by Bill Gilbert

By any measure, 1993 will go down as a remarkable season for hitters. This was expected in an expansion year. However, a comparison with other notable hitting years suggests that a major reason for the offensive explosion this season is the arrival of a new wave of exceptional young players who can hit for both average and power. Entering the last month of the season, nine players have a good shot at hitting the "triple milestones" of a .300 batting average, 30 home runs and 100 runs batted in. If they all succeed, 1993 will rival the legendary hitting year of 1930 when ten players achieved the feat. The only other comparable years since World War II were 1953 and 1961 when eight players did it. The most recent notable hitting year, 1987, produced only four players with triple milestone years. In that year, home run rates were 20% higher than this year but batting averages were about the same.

In 1930, players achieving triple milestones included Hall-of-Fame sluggers Babe Ruth, Lou Gehrig, Jimmie Foxx, Chuck Klein, Hack Wilson and Al Simmons, all of whom had at least three triple milestone seasons. In addition, Wally Berger, Goose Goslin, Gabby Hartnett and Babe Herman had career years, achieving triple milestones for the only time.

In 1953, the first wave of post-World War II sluggers were in their prime years producing eight triple milestone seasons. Stan Musial did it for the fourth time and would repeat the next two years. Ted Kluszewski began a string of four straight triple milestone seasons. Duke Snider, Roy Campanella and Gil Hodges, all of the Brooklyn Dodgers, accomplished the feat in 1953 and they would combine for a total of nine triple milestone seasons before the end of their careers. Eddie Mathews, who hit .300 only three times in his career, did it in 1953 and repeated in 1959. Rounding out the list were AL MVP, Al Rosen and Gus Bell who achieved triple milestones for the only time in 1953.

By 1961, the first expansion year, a new wave of sluggers had reached stardom. Hall-of-Famers Hank Aaron, Willie Mays, Frank Robinson, and Mickey Mantle, a group comparable to the 1930 sluggers, all achieved triple milestones in 1961. Aaron and Mays each did it seven times, Robinson five and Mantle three. Orlando Cepeda reached triple milestones in 1961 and would do it two more times. Finally, Norm Cash, Jim Gentile and Dick Stuart had their career years in 1961, reaching triple milestones for the only time. Conspicuous by his absence from this group is AL MVP Roger Maris who hit 61 home runs in 1961, but batted only .269.

The 1993 group has the potential to match up with their earlier counterparts. Barry Bonds and Frank Thomas have achieved triple milestones in the past and appear capable of repeating for many years. Juan Gonzalez and Ken Griffey Jr. will both do it for the first time in 1993 at age 23 and may have the potential to challenge Ruth's record of 12 seasons with triple milestones. Mike Piazza is on the verge of becoming the first rookie to reach triple milestones since Walt Dropo in 1950 (who never did it again). Rafael Palmeiro and John Olerud have developed into complete hitters who could become regulars on the list. Albert Belle and Matt Williams are having career years with a shot at triple milestones in 1993 if they can add a few points to their batting averages in September.

Potential Triple Milestone Players - 1993
(Statistics include games of August 31)

Player	AGE	BA	HR	RBI	Comments
John Olerud	25	.382	23	97	Needs some homers.
Barry Bonds	29	.344	39	101	3rd time in 4 years.
Frank Thomas	25	.319	36	109	Also did it in 1991.
Ken Griffey, Jr.	23	.318	39	91	First of many.
Mike Piazza	25	.315	28	85	Needs some RBIS.
Juan Gonzalez	23	.313	40	102	Same as Griffey.
Rafael Palmeiro	29	.303	33	89	Could lose out on BA.
Albert Belle	27	.298	34	110	Batting average is key.
Matt Williams	27	.295	27	86	Needs a strong September.

While offensive numbers may be inflated in this expansion year, the more significant point is that a new wave of young sluggers has emerged on the scene that may rival the sluggers of the '30s, '60s and '70s.

Bill Gilbert, a member of the Society for American Baseball Research, is a free-lance baseball analyst and writer residing in Baytown, Texas.