BASEBALL ANALYST

JOURNAL OF SABERMETRICS

December, 1987 Vol. #33

This edition of the Baseball Analyst (what? This can't be DECEMBER already? You mean this thing is ON TIME?) contains four articles. Two of those articles are mine, "Alternatives for Career Projection,", which starts on page 2, and "Research in Progress", which is on a related subject and begins on page 10. The other two articles are "The 1988 Hall of Fame Ballot", by Robert O. Wood, which begins on page 5, and "Sending the Runner", by Doug Bennion which beginneth on Page 16.

Are y'all running out of ideas or energy? I haven't been getting much stuff in. There's a Christmas break coming up, so what better way to spend it than to turn those analytical arguments you have thrashing around in your brain into articles for the Baseball Analyst? If that's not the Christmas spirit, I don't know what is.

Actually, I have quite a few articles here--I just can't use them. The reason I just can't use them is ACGS--Academic Charts and Graphs Syndrome. You know what Academic Charts and Graphs Syndrome is: that's when you've got a three-and-a-half page article to turn in for Professor Sternlook, and you feel more comfortable turning in a 15-page package then a miserable little four-page flier, so you add on a cover letter, a bibliography, a page of footnotes, four pages of charts--12 lines each, but center them in the middle of the page and they look great--and four pages of graphs derived from the charts, and you've got a nice, weighty 15-page DOCUMENT to deliver on due date. I remember this, because I used to do it, too. Hell, I still do--I can turn a 400-page manuscript into a 1,000 page monster before you can bat out a bar graph.

The only problem is that if you do that to me, folks, I can't use it. We've got 20 pages an issue. Those 20 pages are precious. You people are paying--what, $16 a year for 120 pages of material? The other readers do not want to receive 35 pages of articles and 85 pages of charts and graphs. They want to receive 120 pages of articles, and By God, they shall, or as close to it as I can come. There's no way I can use an 11-page article is this publication. The absolute maximum is eight. My preferred length is four. So please, cut the charts and graphs, scrunch them onto one page if you have to have them, put the bibliography on the bottom of the last page and feed the footnotes to your pet muskrat. He'll appreciate them a lot more than I will. Thanks.

Editor and Publisher: Bill James Address: Box 171, Winchester, Ks 66097

ALTERNATIVES FOR CAREER PROJECTION

--Bill James

I have been trying to improve the career projection system. This is a sickness with me. At this time of year I really shouldn't be working on methods, and if I did choose to do so it would be something else, but I can't help myself; I love typing player records into the form and seeing what comes out and how far it misses the player's actual career, figuring out what I can do to reduce the error, etc. I figured since I am doing it anyway I would share with you all what little I have learned, and maybe profit by your feedback. I would be particularly glad to hear from anybody who knows how to break an addiction to running career projections.

There are about six ways to create a career projection system which might work better than the system I've been working with:

1. REFINEMENT OF 1BROCK.WK1.

1Brock.wk1 is my current projection method. An article on how it works and what is being done to refine it appears later in this issue.

2. CREATION OF AN INDEPENDENT CAREER PROJECTION SYSTEM ON A SPREADHSEET.

If I can do it, you could do it better if you wanted to.

3. TRANSFERRING 1BROCK.WK1 TO BASIC OR SOME OTHER COMPUTER LANGUAGE.

There would be obvious advantages. If the necessary info were plugged into a basic program, then to make adjustments and find out what the effect was on the total package would be a matter of perhaps a few minutes, rather than maybe ten hours. There would be other advantages, but also disadvantages. Sometimes I wish I had written it in basic.

4. DEVELOPING A "TARGET-ORIENTED" SYSTEM.

The Brock system is a series of sustained projections; age 24 if projected from 21-23, 25 if projected from 22-24, etc. The projected career total is the sum of the seasons. One could, instead, project the career totals, and then (if you wanted season totals) create season totals by dividing the career totals into season units.

An easy way to do this would be to use the player's record in seasonal notation. If you compare any player's final record with his record in seasonal notation in mid-career, you'll find profound similarities. For example:

NORM CASH	Years	AB	H	2B	3B	HR	Run	RBI	BB	SO	SB	Avg	Slug
AS OF 1963	4.0	502	145	18	4	32	91	95	100	78	6	.288	.532
AS OF 1967	7.7	526	145	19	4	31	87	92	90	82	5	.275	.500
FINAL	12.9	520	141	19	3	29	81	86	81	85	3	.271	.488

Cash's career had its ups and downs--but even from an early point, had you been able to guess accurately how long he would play, you could have made good estimates of his career totals by simply multiplying his performance per 162 games by the number of expected games.

DOC CRAMER	Years	AB	H	2B	3B	HR	Run	RBI	BB	SO	SB	Avg	Slug
AS OF 1936	4.9	669	204	33	8	4	107	62	39	34	4	.305	.396
AS OF 1940	8.4	684	209	33	9	3	111	63	42	29	5	.305	.390
FINAL	13.8	661	196	29	8	3	98	61	41	25	4	.296	.375

ED YOST	Years	AB	H	2B	3B	HR	Run	RBI	BB	SO	SB	Avg	Slug
AS OF 1950	3.4	594	153	28	7	6	88	49	107	64	5	.258	.360
AS OF 1954	7.2	595	154	30	6	9	98	52	121	66	6	.259	.375
FINAL	13.0	564	143	26	4	11	93	52	124	71	6	.254	.371

Those were chosen at random; I doubt that they are anything more than of normal similarity. There is an inherent similarity there that you can capitalize on to project career stats directly if you can just find a way to estimate how much longer the player will play.

What you would need to do then, rather than concentrating on writing formulas to carry the player from year to year, would be to write formulas to evaluate the factors determining his longevity as a player, and to project changes in his productivity such as are possible. To the extent that changes in the player's record in seasonal notation occur over time, they are somewhat predictable. Stolen bases, triples and batting average probably almost always decline. At bats per game almost always increase in mid-career, and decrease toward the end of the career.

You can use other information to make the projection--the player's batting record in his most recent season, for example, or his defensive records. Yost' increase in power was probably predictable based on his playing his early years in an awful home run park, hence (in his early years) hitting 80% of his home runs on the road. The same effect would hold for Bob Elliott, discussed in the other study.

Among all the approaches discussed here, this is probably the most viable, the one which all things considered holds the most promise for accurate projection and creates the fewest obstacles. I use the Brock system because, among other things, its fun. I like to play around with it. But if I were making a strictly logical decision of how to approach the projection of a player's career stats, this is the one I would pursue.

5. PROJECTION BY SKILLS ASSESSMENT

Remember the system in the 1987 Abstract which proposed a way of evaluating the speed of each player, based on his performance in six areas? Well, suppose that you extended that approach to where you were assessing the player's abilities in a wide variety of areas. George Brett, Jim Rice and Barry Larkin might look something like this:

	Brett	Rice	Larkin
Speed:	4.1	3.8	7.1
Plate Discipline:	7.4	6.0	4.1
Plate Mechanics:	9.8	8.8	3.5
Reflexes:	6.0	5.1	6.6
Throwing arm:	5.0	4.1	7.0
Durability:	3.5	4.5	6.5
Quickness:	4.1	2.8	7.1
Balance:	9.5	8.0	6.0
Flexibility:	4.4	4.5	7.1

Such a system would need three parts. The first part would take the player's known production, and break it down into these elements. The second group would project changes in those elements--speed declining, plate mechanics improving rapidly at first, less rapidly later, reflexes declining in the mid-thirties, durability declining, plate discipline improving over the years, etc. The third group of formulas would translate these abilities back into production, with batting average being the outcome of all of the things which change batting average.

This approach has a couple of attractive advantages over the Brock systems. It has a theoretically underpinning: one knows WHY one is making the changes that one projects. The method would force its creator to learn about specific attributes, each interesting. How rapidly does speed decline? How much does plate discipline increase? Can one predict with any accuracy whether or not it will? At what age do the reflexes go?

In the Brock systems, I am attempting to project the end results of these processes, with no data about the processes themselves. This calls for basic research; the Brock system is applied research. It is applied research with intuition substituting for the basic knowledge.

That being said, it is questionable whether this approach would ever yield projections which were more accurate than what we can get with Brock. The complexities involved in ACCURATELY breaking a player's record into component abilities and then re-combining them into the same record are enormous. How well the changes in those component abilities could be projected is problematical. It's a nice idea but I don't know that it's ever going to work.

6. PROJECTION BY SIMILAR PLAYERS

Another way to project a player's expectation would be to identify past players who were similar at the same age. You would need an enormous data bank (or perhaps a series of smaller data banks) which would contain possible comparisons. You would feed into the system the information about, let's say, Mike Marshall, and it would identify the most comparable players. You might tell the system that Marshall was 27 years old, hit .290 with 15 homers in 100 games, had career totals of 105 home runs, a .271 batting average with 629 hits. You might add that he has back trouble, is 6'5", hits right-handed, can't run and plays in the worst hitter's park in baseball. The system would then search the data for the most similar players, and report that the most similar players were. . .I don't know, Mike Epstein, 1970, Willie Stargell, 1967, and Ron Blomberg, 1975. I'm sure you would be able to find a better group than that. The projected career for Marshall would be represented by the output of the most-similar players, with him who are #1 most similar being weighted more heavily than him who am #10.

An advantage of this system is that it yields an objective fact: this group of players, who averaged almost exactly the same as Marshall through age 27, went on to post career totals of exactly this. That is meaningful to some people who would dismiss our other projections.

The biggest problem with this system is that most players are unique in a critical respect. It is their uniqueness which makes them intriguing. Take the Cardinals. The group selected as similar to Vince Coleman would probably include the young Omar Moreno, the young Willie Wilson, maybe Tim Raines or Rickey Henderson. They are certainly similar to Coleman, and the averages of those selected would parallel Coleman in almost every category--except that they would have 70 stolen bases a season, not 110. The group for Ozzie Smith would be terrific shortstops who hit about like Ozzie, but they would be not quite as terrific at short. The group representing Jack Clark would be a group of awesome hitters who get hurt a lot, but would they be AS awesome with the bat, and would they get hurt AS often as Jack? Probably not. The group representing Terry Pendleton would probably be good defensively, but would they be AS good as Pendleton?

That is true of every team: the best players are all unique in one way or another. The other systems proposed here would have no trouble with that. If Coleman stole 140 bases in 1988, we'd just enter 140 in the stolen base column. The skills assessment system can make Ozzie as good defensively as he needs to be. The Brock system, although it does not, still could. A system of identifying similar players simply couldn't.

Building a workable career career projection system is not a weekend project. It takes a lot of time, a lot of hit-and-miss. I'd like to set up some sort of a system in which the rest of you could participate in the process of refining and improving a shared method, but at this point that just isn't practical for me to undertake. Unless somebody else organizes it, We're going to have to work on individual initiative.

THE 1988 HALL OF FAME BALLOT

Robert O. Wood

As the baseball Hall of Fame (HOF) balloting is currently taking place, I thought it would be interesting to both anticipate the results as well as determine who ought to be deemed worthy of the honor. This article consists of three tables, one for those players for which 1988 is their first year of HOF eligibility, and two others for those players who were on last year's ballot and failed to be elected.

Robert McCleery and I have developed a career valuation formula both for hitters, and separately, for pitchers. The hitters formulation was described in volume 28 of the Analyst. Suffice it to say here that the result is a numerical formula which sensibly combines a wide variety of a hitters' career statistics, taking into account such factors as defense, era illusions, and park illusions. We undertook the same type analysis on behalf of pitchers -- unfortunately, it is not applicable to relievers.

We scaled the formulas so that the minimum requirement for a first-ballot Hall of Famer is a career value of 300. Players congregating around this milestone include hitters Reggie Jackson, Ernie Banks, and Paul Waner, and pitchers Robin Roberts, Jim Palmer, and Red Ruffing.

As a second point of reference, we feel that a player with a career value of 230 deserves virtually no chance of making the HOF. Notables with values near this mark include hitters Bobby Murcer, Ken Keltner, and Dick Groat, along with pitchers Bobo Newsom and Steve Rogers. Between deserving (300 value) and undeserving (230 value) lies a discretionary "gray area" in which a player can be said to have a chance of entering the HOF. We assign probabilities in this 70 point range in a linear fashion, ranging from 100% to 0%. For example, a player with a value of 265 deserves a 50% chance of HOF enshrinement.

In the 1986 Baseball Abstract, Bill James sets forth the de facto HOF voting criteria. He stresses that he does not necessarily agree with the numerical point system which he discovers, only that it seems to reflect past HOF balloting. Among other failings, it blissfully ignores era effects, so that the superficially gaudy hitting stats of the lively-ball era of 1929-1939 are considered as evidence of the superiority of those hitters as compared to hitters of other eras. The point system (herein referred to as "BJHOF") is designed so that a player with BJHOF of 130 or more is certain to be voted in, with BJHOF of 70 or less a player has no chance of entering HOF. In the gray area between 70 and 130, the probability is linear so that a BJHOF of 100 implies a 50% chance.

I should remark here that the probabilistic nature of the two scales need not reflect stochastic (or "random") voting behavior of the individual voters. Rather it is best thought of as reflecting our incomplete knowledge regarding the over 400 different HOF thresholds of the voters.

Table I consists of those players appearing this year for the first time on a HOF ballot. The columns give our value, the associated deserved probability, BJHOF, and its associated actual probability. [The reader is encouraged to consult the original sources for a thorough understanding of both the Wood-McCleery value scale, and the Bill James HOF point scale, and why we refer to the former as the "deserved" probabilities and the latter as the "actual" probabilities.]

Table II consists of those players who have appeared on past HOF ballots. It contains the same categories as did the first table. Table III presents the past "voting record" of these holdovers. It gives the entire history of each player's HOF votes received, expressed as a percentage of the total ballots cast in that year. Of course, one needs to be named on at least 75% of the ballots to be elected. In the tables, the entry 'n.r.' denotes a player we have not ranked, and would have a value substantially below 200.

Note that the probabilities appearing in the tables do not necessarily correspond to the probability of being elected this year, but attempt to capture the probability of ever being elected to the HOF (either via the 15 years of BBWAA balloting, or via the Veterans Committee).

Several interesting conclusions can be gleaned from the tables. In Table I, we see that the three legitimate just-eligible players -- Stargell, Smith, and Tiant -- will likely be undervalued by HOF voters. Nevertheless, it is probable that Stargell will be elected this year, since he is similar (in both BJHOF and via our methods) to McCovey, and he was a first ballot electee. It is unlikely that Tiant will ever be elected, whereas Reggie Smith has absolutely no chance. Indeed, Sparky Lyle could be a more likely candidate than either of these two.

Holdovers which we think merit election include Santo, Boyer, Allen and Cepeda. Note that all of these had poor early returns, and are hoping to get up a full head of steam before their time runs out. It is likely that only one of this group will ever be elected by the BBWAA, probably Cepeda. The others will have to await the Veterans Committee.

On the positive side, it is probable that Bunning will be elected this year, a decision we can live with. On the other hand, it is unnerving to note that Oliva is making substantial progress in his quest. We have nothing against Tony, only that he is likely taking votes away from more deserving players.

Ours is not the last word, nor should it be. This short article attempted merely to present the "facts" to the readers, and stay clear of the intellectual morass that most Hall of Fame discussions have sunk to.

TABLE 1 -- HALL OF FAME BALLOT FIRST TIMERS

	VALUE	PROB	BJHOF	PROB
Stan Bahnsen	n.r.	0.00	14	0.00
Mark Belanger	143	0.00	52	0.00
Ross Grimsley	n.r.	0.00	24	0.00
Larry Hisle	178	0.00	22	0.00
Al Hrabosky	n.r.	0.00	16	0.00
Grant Jackson	n.r.	0.00	12	0.00
Randy Jones	n.r.	0.00	34	0.00
Bill Lee	n.r.	0.00	15	0.00
Sparky Lyle	n.r.	0.00	89	0.32
Lee May	214	0.00	35	0.00
John Mayberry	180	0.00	21	0.00
Lynn McGlothen	n.r.	0.00	5	0.00
Doc Medich	n.r.	0.00	9	0.00
John Milner	128	0.00	3	0.00
Willie Montanez	163	0.00	21	0.00
Manny Mota	155	0.00	20	0.00
Ken Reitz	130	0.00	6	0.00
Joe Rudi	197	0.00	20	0.00
Reggie Smith	262	0.46	59	0.00
Jim Spencer	145	0.00	3	0.00
Willie Stargell	305	1.00	106	0.60
Luis Tiant	268	0.55	94	0.40
Del Unser	145	0.00	1	0.00
Rick Wise	211	0.00	32	0.00

TABLE II — HALL OF FAME BALLOT HOLDOVERS

	VALUE	PROB	BJHOF	PROB
Dick Allen	276	0.66	96	0.43
Bobby Bonds	263	0.47	62	0.00
Ken Boyer	278	0.69	74	0.07
Jim Bunning	269	0.56	97	0.45
Orlando Cepeda	276	0.66	108	0.63
ElRoy Face	n.r.	0.00	45	0.00
Curt Flood	225	0.00	60	0.00
Elston Howard	233	0.05	60	0.00
Harvey Kuenn	232	0.03	72	0.03
Don Larsen	n.r.	0.00	20	0.00
Mickey Lolich	251	0.30	87	0.28
Roger Maris	218	0.00	77	0.12
Bill Mazeroski	231	0.01	47	0.00
Minnie Minoso	255	0.36	69	0.00
Thurman Munson	261	0.44	84	0.23
Tony Oliva	245	0.21	107	0.62
Vada Pinson	260	0.43	89	0.32
Ron Santo	286	0.80	85	0.25
Joe Torre	273	0.61	79	0.15
Maury Wills	210	0.00	87	0.28
Wilbur Wood	218	0.00	63	0.00

TABLE III — HALL OF FAME VOTING RECORD

[Percent of votes received out of total ballots cast]

ELIGIBILITY YEAR	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Dick Allen	4	0	7	10	13									
Bobby Bonds	6													
Ken Boyer	2	4	4	5	5	0	0	0	0	0	17	22	23	
Jim Bunning	38	48	34	46	41	33	37	50	54	66	70			
Orlando Cepeda	12	19	10	16	31	29	36	43						
ElRoy Face	6	9	7	8	5	6	5	9	16	16	17	19		
Curt Flood	4	2	3	0	0	0	0	0	7	11	12			
Elston Howard	5	6	14	11	11	7	8	21	10	9	11	14	12	11
Harvey Kuenn	15	15	15	22	23	15	21	26	32	34	35			
Don Larsen	8	6	12	10	8	12	8	8	8	6	7	8	8	7
Mickey Lolich	20	20	20											
Roger Maris	21	19	22	19	22	29	29	23	17	18	27	32	42	43
Bill Mazeroski	6	8	9	9	7	13	18	22	24	30				
Minnie Minoso	21	20												
Thurman Munson	15	6	5	7	8	8	7							
Tony Oliva	15	20	31	29	36	39								
Vada Pinson	4	1	3	0	.5	10	12							
Ron Santo	13	15	19											
Joe Torre	5	11	11	14	11									
Maury Wills	30	38	38	41	22	21	26	24	29	27				
Wilbur Wood	3	4	5	6										

RESEARCH IN PROGRESS

--Bill James

My current system for projecting a player's career totals is called "1Brock.wk1". I have written about similar systems before. I wanted to report here on the progress that I have made in improving it, correcting its errors, etc. I wanted to put on record a few statements about where it is in December, 1987--how accurate, how inaccurate, what is being done to improve it. Later on I'll offer you a copy of it if you want one--but I'm warning you, it's an addictive little bugger.

Those of you familiar with Lotus 1-2-3 will recognize the extent; the system, like the earlier career projection systems which I have written about, is set up on a spreadsheet. Although I used a similar name, and although it does incorporate many of the features of its fore-runner and possibly even (I'm not sure) a few of the same formulas, the new system is not essentially a refinement of the old one. It is a new method, just created a few months ago,

1Brock is not a complex theoretical method; indeed, what bothers me most about it is that it has no "original idea", no underlying premise or theory to support it. What you do is, you type into the spreadsheet the player's record through a given point of his career. The system, by a series of formulas, takes the last few years and projects the next one to be pretty much the same, with a few changes. Home runs per year increase if the player is young, decrease later. Speed elements--stolen bases, triples--decline earlier and more rapidly than other things. Walks never really decline. As the player ages, his productivity shrinks.

In a set of formulas off-screen to the right, the system is calculating the player's offensive ability for each projected season, and comparing that with a minimum standard, calculated from the players age, the defensive position he plays and the league average of runs per game. When the point is reached at which the player's offensive productivity slips below the established levels, the system begins to reduce his playing time in the subsequent projections. Usually three or four years after his playing time begins to slip, the player will be retired. The system never holds a Day in his honor.

In the August, 1985, edition of the Baseball Analyst, Dallas Adams reported on a study of the accuracy of one of the earlier Brock systems, Brock4. He found that that system tended to be seriously over-optimistic in projecting a player's career, that in particular it tended to project players to hit 43 percent more home runs than they actually did, a tremendous error, and to score 45% more runs then they actually did. It was my intention, when I asked Dallas to undertake that research, to use his findings to modify and improve the existing Brock system. For several reasons, this didn't work out. That system was set up to run on a KayPro II, a rather primitive home computer, and as a result the system

a) was extremely difficult to work with. The system took several minutes to read when called up and about two minutes to recalculate a single run, so that it to revise the system and do a study of its new accuracy, a study necessarily involving at least a hundred players, was an extremely time-consuming process, and

b) was taxing to the very limits the active memory of the machine, so that at any moment, doing anything at all to try to make the system 1% more accurate, you were likely to encounter sudden messages like "Out of virtual memory!" and "Memory Full--Please Delete Something to Continue!", accompanied most often by an unplanned secession of the immediate activity.

These problems left me feeling trapped with Brock7.pc, blocked at the limits of my knowledge of how to use the machine. Although I continued to tinker with it, I hadn't done anything substantive to improve the system in a year or more.

With a new computer, an AT Clone, I have overcome these problems. 1Brock.wk1, though larger and more complex than Brock7.pc, reads in a matter of seconds and recalculates almost instantly once new values are added. When I built the new projection system, using the principles of the old one but not the exact formulas largely because it was easier to make up new formulas than to copy the old ones over, my first goal was to start incorporating into the system various types of information which there simply was not room for before. I provided a place to tell the machine if the player was a catcher, so it could shorten his career. I used a three-level cutoff for minimal performance levels, rather than a two-level cutoff as before. I used more "cross-fertilization" of columns, doing such things as basing future home runs to a tiny degree on past doubles as well as past home runs, reaching back a little further in the past to evaluate the player's talents, giving a small edge in future batting average to a player who steals bases, etc. A side effect of these changes was a dramatic increase in the stability of the projections. With Brock7, one could sometimes add 200 hits to a player's projected future by just adding 3 hits and a home run to one year or something. Those sorts of things can't really happen with the new method.

Then I began checking the accuracy of the system. Building on the advantages outlined before, I developed a pattern of evaluating the system by doing runs of 25 players at a time, but with the players projected from every (eligible) year of their careers. This is a check which is considerably less thorough and less reliable than what Dallas did, which was to check 200 individual players, balanced for characteristics, but it is considerably more reliable than just checking 25 players from one point each, which would require most of the same time. Twenty-Five players from each eligible year yields somewhere between 150 and 300 entry points--not as good as 200 distinct entry points, but better than 25. To do a "cycle" like this takes an average of about 90 minutes a day for a little more than a week.

At first, I was getting the same kind of errors I had before--home runs projected 22% high, triples 13% low, etc. I was, however, generally surprised by the progress I was making, and after several cycles I would now like to report that I have reached a level of accuracy that I frankly would not have thought possible a couple of years ago.

The basic measure of accuracy that I use in studying the system is similarity scores, comparing the similarity of the actual career totals to those projected by 1Brock. I use the similarity system exactly as it is laid out in the 1986 Abstract. The comparison is based on career totals, not just on the projections; however, since the past totals are the same for both the projection and the actual career, and since the similarity scores method is based on absolute differences, not relative differences, it makes virtually no difference whether you use all career totals or just the projections. The only areas of the comparison which would be altered by using just the projections would be the percentages, batting average and slugging percentage, and the only case in which there would be a big difference there would be the case where a player might be projected to hit .262 in the rest of his career but actually hits .133 and wipes out in 21 games or something. That would create an artificial 129-point penalty for the .129 difference in batting average, which obviously would be more misleading than substantive, and since this is the only case in which it makes any difference whether we compare career totals or projected totals obviously it makes more sense to compare career totals.

When I started running studies, the average similarity score for a projection was in the neighborhood of 820. At the moment it is around 880--actually 883.5919 in the last study, which involved 261 projections. Let's see if I can find one here to illustrate that. . .here's one--Bob Elliott, projected from 1945.

	G	AB	H	2B	3B	HR	Run	RBI	BB	SO	SB	Avg	Slg
Projected	1842	6736	1935	360	110	90	979	919	787	479	67	.287	.413
Actual	1978	7141	2061	383	94	170	1064	1195	967	604	60	.289	.440

WITH STATS THROUGH 1945 REMOVED:

	G	AB	H	2B	3B	HR	Run	RBI	BB	SO	SB	Avg	Slg
Projection	935	3309	921	172	45	45	477	354	422	256	28	.278	.398
Actual	1071	3714	1047	195	29	125	562	630	602	381	21	.282	.451

In that case the system didn't anticipate Elliott's power surge when he moved out of massive Forbes Field in 1947, and thus was completely off on the home run total, leading to large "penalties" for home runs, RBI and slugging percentage although the system was pretty accurate at projecting how long the player would continue to play, which is ordinarily the crucial thing. In the usual case the system will be off a little on everything. Rick Monday's first projection, from the 1968 season, is perhaps more typical in this respect:

	G	AB	H	2B	3B	HR	Run	RBI	BB	SO	SB	Avg	Slg
Projected	2133	6945	1804	301	75	183	850	842	967	1861	155	.260	.404
Actual	1986	6136	1619	248	64	241	950	775	924	1513	98	.264	.443

WITH STATS THROUGH 1968 REMOVED:

	G	AB	H	2B	3B	HR	Run	RBI	BB	SO	SB	Avg	Slg
Projection	1844	6016	1566	262	61	161	738	753	847	1595	137	.260	.404
Actual	1697	5207	1381	209	50	219	838	666	804	1247	70	.265	.451

At the risk of being obvious, if you get an 880 projection early in a player's career, as we did with Rick Monday, that's pretty satisfactory; if you get an 880 in mid-career, as with Bob Elliott, that's decent. But if you get an 880 toward the end of a player's career, as for example with Willie Mays (also from 1968) then you feel like the system really ought to be able to do better.

The other basic measure of the performance of the system is category error, meaning "How do the projected totals for all players compare to the actual aggregate totals." This is what Dallas Adams reported on in his study a couple of years ago. These are the category errors for the system in the latest run:

Games played:	.96 (Meaning 4% too low)
At Bats:	.99
Hits:	.98
Doubles:	.98
Triples:	1.01
Home Runs:	.97
Runs Scored:	.99
RBI:	.97
Walks:	.98
Strikeouts:	.97
Stolen Bases:	.88
Batting Average	.998
Slugging:	.994

These figures obviously would be different, and moved away from the 1.00 goal, if we compared only the projected parts of the careers--although again, this is irrelevant to attaining the goals of the method, since we're trying to refine the system to reach 1.000, and 1.000 would be 1.000 anyway.

The error on stolen bases here is somewhat misleading, since one of the 25 distinct players in the study was Lou Brock. Brock had 17 projection points (each year 1963-1979) and 938 actual stolen bases in each projection, so that he alone accounts for about 40% of the stolen bases in the study--a unique situation which doesn't, and really couldn't, occur in any other category. The projections for Lou Brock run generally low, and thus the projections for stolen bases are well off the mark. The system may run a little low on stolen bases, but not (in general) 12% low. Before doing another cycle of refinements and evaluations, I would select another 25 players (meaning another 150-300 projections), find the category errors in them, and then try to design ways to eliminate the error in each category.

In this particular case, one of the largest errors is in the games played category. Assuming that the second group of players also shows this error--I would certainly think it likely that they would--what this would mean is fairly obvious. Almost the right number of at bats but in 4% too few games means not enough games as a part-time player at the end of the career. Most players at the end of their careers have a couple of years with low ratios of at bats per game. We can move toward eliminating this problem by increasing the gap between the level of offensive productivity needed to continue playing regularly and the level of productivity needed to continue playing part-time, thus creating more part-time seasons.

Another problem with the projections is that home runs run, at the moment, 3% too low. In this case (again assuming this holds up on the second study) what I would do is try to find the age at which the problem occurs, and make an adjustment in the most effected period. My process for evaluating the system includes the player's age, which makes it easy to find the category error for each age. The category errors for home runs projected from each age (21-35) are given at left at the top of the next page:

Age	HR%
21	.90
22	.81
23	.86
24	.99
25	1.05
26	.94
27	.98
28	.94
29	.89
30	.96
31	.97
32	1.01
33	.99
34	1.00
35	.99

The projection from age 22 (.81, 19% too low) is based on just three players, so that really isn't terribly meaningful (unless a similar gap shows up in the second group). The projection from age 25 (5% too high) is based on 17 players, so that one I would probably make a small adjustment for. When I check the HR formula for age 26, I find that it calls for home runs at age 26 to be increased by 3% from the home run rates of the previous four years. Since it projects too high, I'll cut that to maybe 2%. Remember, the projections for the years after age 26 will be based on what the player does at 26, so the exact effect of each change is somewhat unpredictable. Cutting home runs at age 26 down by the full five percent might cut home runs over a period of years by seven percent, and might cause the group of players to retire a year too early, and might cause a category error for age 25 of .88 (12% too low) or something, as well as throwing off the errors for all of the other categories. At this point of the process, with a system that already has a considerable degree of accuracy, I'd rather take a small gain than risk a setback (or, actually, risk doing a complete cycle of the study for no gain, since obviously if the new projections turned out to be less accurate than the old ones I would simply start over on the next cycle with the previous version of the system.)

Anyway, the projections from age 29 are 11% too low. The formula for home runs at age 30 is odd in that it is based only on the previous two seasons of home run production, while others are based on three or four years and often draw on the doubles category as well as home runs. That change, then, seems pretty obvious: pull it in line with the other formulas to project home runs. But that isn't the main reason this figure is too low. The .89 figure here is just part of a pattern; EVERYTHING projected from age 29 is too low. The biggest problem is that batting averages are cut 18 points for the age 30 projection, which (it is obvious now) is too much, forces the players out before their time.

So to increase the home run projection, we have to tinker with the batting average--but when we change the batting average at age 30, we change all of the projected figures for everything before the age of 30. It is, in this complex, just so, so easy to over-react to a problem and wind up with a system that suddenly projects 7% too many games and a similarity average of 855. So in the next cycle, I'll change the formula for batting average at age 30--but only by a couple of points. A small gain, rather than losing ground.

There are times when this type of age analysis points out the problem glaringly--for example, in the last cycle triples were projected 7% too high, but triples from the age of 23 (seven players) projected 53% too high. Obviously, then, I cut the growth of triples at the early ages, and most of that problem went away.

I also, in each cycle, run category errors and average similarities for the players at each position. This is less meaningful because whereas each player is 23 one year, 24 the next, etc., a player who is a second baseman one year tends to be a second baseman the next, so that the second base group doesn't contain information from enough distinct individuals to be very reliable. Still, if we find that games played are running low, it is useful to know for which group of players they are running particularly low. If you find that the projections for second basemen are running especially low, you can make an adjustment so that the retirement levels for second basemen are a just a little bit more liberal, thus making the projections for second basemen a little bit more accurate.

The average similarities, for obvious reasons, improve with age--indeed, for many players the projections become more accurate almost with each passing year. The average similarities for groups in the last cycle were:

Catchers	877.1	Age 24	790.9
First Basemen	881.5	Age 25	813.6
Second Basemen	905.6	Age 26	812.3
Third Basemen	924.4	Age 27	850.5
Shortstops	934.7	Age 28	872.2
Left Fielders	884.3	Age 29	876.5
Center Fielders	841.4	Age 30	872.0
Right Fielders	880.2	Age 31	899.4
		Age 32	914.4
		Age 33	933.8
		Age 34	938.4
		Age 35	946.2
		Age 36	956.1

Players projected to steal more than 200 bases:	844.3
Players projected to hit more than 200 HR:	872.1
Players who walk more than once per 8 at bats:	902.0
Players who walk little:	866.4

The projections for center fielders are low because one of the center fielders in the group was Mickey Mantle, whose projections are the most inaccurate of any player (projects for most of his career as a player who will get 3,000 hits.) The projections for catchers are low because catchers' careers are shaped by injuries, and thus damned difficult to project accurately. The fact that the projections for all of the "hitting" positions (outfield and first base) are lower than for the infield is intriguing, and may suggest some systematic adjustment.

Well, enough about the system that exists. The main reason for writing was to invite you all to get involved, if you can find a way, in creating a better career projection system. To begin with, I'll be happy to send any of you a copy of 1Brock.wk1, if you'll send me a blank diskette (360K or 1.2 Meg, IBM compatible), the mailing expenses and a couple of sheets of cardboard to mail it in. If you want to take 1Brock as a basis and improve it, you're welcome to try.

It is my intention to run career projections for a good many players and include them in the 1988 Abstract. I don't know right now whether I will run the projections with the current system, or go through another cycle to improve the system before running the projections. I know what I SHOULD do, but the problem is that one always knows what needs to be done to make the system work better, and it's hard to walk away from it when you know you could make it better. Once you implement the changes, you have to run the projections to make sure they work. Once you run the projections and make sure they work, you have new ideas about how to make it better. It's a vicious cycle.

How much room for improvement is there? As I said, a year ago, working on the old KayPro, I didn't think I'd ever get where I am. But obviously, this system could work a lot better than it does. I think it definitely is possible to reach 900 (average similarity), and that it may well be possible to push it to 920. There is no reason to accept category error at all, even for a sub-group. Above that, I'll wait and see.

SENDING THE RUNNER

I was somewhat surprised with the outcome of this exercise; perhaps you will be too.

The issue is this: There is a runner on first, with none or one out, and a full count on the batter. Under what circumstances is it right to send the runner, and what is gained?

The Blue Jays did it often last year. I must admit I thought they used the strategy too frequently, but I hadn't thought it through and no doubt the Jays had.

Jimy Williams sent anybody! I remember Heathcliffe and Ernie Whitt being thrown out ignominiously at 2nd, on the tail end of strikeout-throwout double plays. And Jimy didn't seem to mind who was batting. He sent the runner even when the risk of strikeout was high, say when Barfield or Moseby were at bat.

I wondered if there was some way to quantify the pros and cons of this strategy --- if we could develop guidelines as to its optimum use. Let's try. I'm going to deal with the none-out position; an analysis of the one-out position results in the same approximate conclusions.

One of three relevant events may occur at the plate; (1) the hitter may hit safely, (2) he may hit into a contact out (grounder or fly-out or pop-up), or (3) he may strike out. Two other possible events are not relevant. If he fouls off the ball, he merely defers the issue to the next pitch, and if he walks, the running start is not a factor.(There are also remote chances of passed balls or wild pitches or hit-by-pitcher, but I will ignore these events.)

I intend to measure the impact of a running start on run production for each of the three relevant events. First I will gauge the runs gained or lost by each event, then I will determine the net overall impact on runs by examining a range of probable frequencies for each event.

In order to quantify runs gained or loss, I will use the concept of the "run value" of an inning, with which most readers of this article will be familiar. Pete Palmer in "The Hidden Game of Baseball" showed that it is possible to assign a historically derived run potential to an inning based on the position of any runner(s) and the number of outs. Throughout this article I have used Palmer's values.

Run values are known for all combinations of position/outs. For example, the highest run value, that for an inning with the bases loaded and none out, is 2.25 runs. The lowest, that for an inning with none on and two out, is 0.10 runs. Run values for other combinations lie between these extremes. An inning with a man on first and none out is valued at 0.78 runs. Over the course of time, an average of 0.78 runs will score from that position.

If the Batter Gets a Hit

If the batter gets a hit, what is the impact of a running start? Well, it depends on the type of hit. If it's a triple or home run, the running start has no effect. But if it's a single or double, the running start aids the offence because it leads, in general, to greater base advancement.

Normally, on a single, a man on first will advance to third roughly 40% of the time. This success will vary, of course, with the speed of the runner. Fernandez might make it 50% of the time, and Whitt 30% of the time, but on balance an average expectation is about 40%.

Similarly, a double will score a man from 1st perhaps 50% of the time.

A running start will materially increase these successes. For this analysis I will estimate that on a running start, the runner will (1) advance to third on a single 90% of the time and (2) score on a double 90% of the time.

I won't bore you with the dirty details, but using these assumptions, and comparing run values of innings with and without a running start, I can demonstrate that the strategy gains an average of about 0.10 runs for each hit.

If the Batter Makes a Contact Out

What happens on a contact out? Well, about half of contact outs are fly-outs and pop-ups and these change the picture very little; the situation after the out -- the runner retreats -- is usually the same whether or not there is a running start. It is true that the incidence of line-out double plays would be slightly higher when the runner is started but these are rare events and the impact of this difference would be adverse, but very small.

The impact of the strategy on grounders is significant, however. For this exercise I've had to make a number of assumptions. I have guessed that in normal circumstances, 50% of ground balls would result in double plays, 25% would result only in the runner at second being thrown out, and 25% would result only in the runner at first being thrown out. These values clearly aren't precisely correct, but I don't think they're stupid and we're essentially interested in general impressions anyway.

I've further estimated that if the runner is started, the number of double plays is reduced to 20% from 50%, the number being thrown out at second is reduced to 0% from 25%, and the number being thrown out at first is increased to 80% from 50%. Again, crude but probably not dumb.

Given these assumptions, starting the runner gains 0.24 runs when the batter makes a contact out. <u>If we had used different sets of reasonable assumptions</u> the gain would be of the same general order.

If the Batter Strikes Out

The third and final relevant event is the strikeout. Here the strategy can either gain, when the runner steals second, or lose, when the runner is thrown out at second, the dreaded strikeout-throwout double play.

Again bypassing the arithmetic, the play gains 0.22 runs if the runner steals second, but loses 0.38 runs if the runner is thrown out.

So, we conclude that starting the runner (1) <u>gains</u> about 0.10 runs when the batter gets a hit, because the baserunner advances further more often, (2) <u>gains</u> roughly 0.24 runs on a contact out, because the number of double plays is reduced, (3) <u>gains</u> about 0.22 runs after a strikeout on a successful steal, and (4) <u>loses</u> 0.38 or so runs after a strikeout if the runner is thrown out.

Now, in order to do smart things with this information we must know the frequency of each of these events. Do you know them? I don't and I'm not sure anybody else does either. Well, I know the first -- on a 3-2 count the typical batting average is about .200

-- but I don't know how to apportion outs between contact outs and strikeouts.

However, not all is lost. We can work with general ranges. The following table presents an array of runs gained given different ranges of strikeout and stolen base probabilities. I will repeat that I don't suggest these values are precisely accurate, but their general range and sign would be roughly so.

Runs Gained by Running on a 3-2 Count

Strikeout %	Steal Success % 70	50	30	10
100	.05	-.05	-.14	-.24
75	.07	.00	-.08	-.15
50	.08	.04	-.01	-.06
25	.10	.07	.05	.03
0	.11	.11	.11	.11

This says, for example, that if 25% of a hitter's outs are strikeouts, and if the runner at first can steal second 50% of the time, the gain over the long run from starting the runner will be about 0.07 runs. A few general observations:

* With a decent baserunner, say one who will steal second at least 60% of the time, the table tells us to <u>always</u> send the runner, regardless of who is at bat.

* With a good contact hitter who will seldom strike out, <u>always</u> send the runner, no matter how slow he is. Heck, even if Pat Gillick is the runner, send him! There are few Ks to be doubled up on, and the other things that can happen are all good.

* Within normal ranges the general rule would be to send the runner. However, if the runner is very slow and the hitter tends to K a lot, apply some judgement. What, does Whitt

have a 10% chance of stealing a base? Does Barfield have a 75% chance of striking out if he doesn't hit safely? If so, you shouldn't send Ernie; on average you'll lose about 0.15 of a run doing so. Application of the strategy should be very liberal, but not automatic.

What does this mean in wins and losses? Check the table, what do you suppose the average gain per event would be ... 0.05 runs or so? How many times a season would this happen? Would it be as often as once a game? Well, 162 x 0.05 is about 8 runs. Over a season, an 8 run pickup might win you one additional game. Hardly spectacular, but it's there for the taking.

The Jays knew all about this. They'll take it. They may need it.