By the Numbers

Volume 18, Number 3 The Newsletter of the SABR Statistical Analysis Committee August, 2008

Review

Academic Research: Are Umpires' Ball-Strike Calls Influenced by Previous Pitches?

Charlie Pavitt

The author reviews a recent study that suggests that the more balls an umpires has recently seen, the more likely he is to subsequently call a strike on a borderline pitch.

Clare MacMahon and Janet L. Starkes, Contextual Influences on Baseball Ball-Strike Decisions in Umpires, Players, and Controls, Journal of Sports Sciences, May 2008, Volume 26, Number 7, pp. 751-760

This piece is a bit more in line with what the members of our sibling Science and Baseball Committee would be interested in, but I think it is worth our attention too. Based on the authors' literature review, there has apparently been a bit of experimental research on bias in umpire calls on borderline pitches. One by Rainey, Larson, and Stephenson published in the *Journal of Sport Behavior* back in 1989 found that umpires in a "wild pitcher" condition shown a sequence of ten pitches in which seven were clearly balls were more likely to call a subsequent ambiguous pitch a strike than umpires in a "control pitcher" condition shown a sequence in which only three were clearly balls. This surprises me, as I would have expected the opposite given at least the television announcers' claims that pitchers with reputations for good control get calls that pitchers with wild reputations don't receive. My expectation follows from the role of prior expectations in most of the judgments people make, but some sort of psychological contrasting effect seems to be happening here instead.

The current study, including experienced umpires, university-level players, and undergraduates claiming little baseball-relevant experience, is in the same vein. A series of pitches to batters was videotaped and, in a pre-test, called by a group of umpires; the pitches in which the umpires all agreed were considered unambiguous strike or ball cases, and the pitches in which they did not agree were considered ambiguous and used as the test cases in the main study. One set of trials was performed through showing the judges a sequence of unambiguous cases and then an ambiguous case to judge. After being shown a series of either one or two clear balls, all judges were a bit more likely to call an ambiguous pitch coming next as a strike (about 48 percent of the time) than after observing one or two clear strikes (about 45 percent of the time). These results were consistent with the Rainey et al. findings and again inconsistent with what I would have hypothesized. A different set of trials asked the judges to evaluate ambiguous pitches after being told what the count was before the pitch. There was a slightly lower proportion of strikes (34.2%) with 0-2 counts than with 3-2 (36.9%) or 3-0 (37.0%), and 0-0 winds up in the middle (35.5%). The 3-2 is the interesting one, in this case apparently along the lines with past qualitative observations the authors cite (Larsen and Rainey, *Journal of Sport and Exercise Psychology*, 1991) that umpires have a bias toward keeping the game moving along.

Charlie Pavitt, chazzq@udel.edu ♦

In this issue

Informal Peer Review

The following committee members have volunteered to be contacted by other members for informal peer review of articles.

Please contact any of our volunteers on an as-needed basis – that is, if you want someone to look over your manuscript in advance, these people are willing. Of course, I'll be doing a bit of that too, but, as much as I'd like to, I don't have time to contact every contributor with detailed comments on their work. (I will get back to you on more serious issues, like if I don't understand part of your method or results.)

If you'd like to be added to the list, send your name, e-mail address, and areas of expertise (don't worry if you don't have any – I certainly don't), and you'll see your name in print next issue.

Expertise in "Statistics" below means "real" statistics, as opposed to baseball statistics: confidence intervals, testing, sampling, and so on.

Member	E-mail	Expertise
Shelly Appleton	slappleton@sbcglobal.net	Statistics
Ben Baumer	bbaumer@nymets.com	Statistics
Chris Beauchamp	cbeaucha@asinc.ca	Statistics
Jim Box	jim.box@duke.edu	Statistics
Keith Carlson	kcsqrd@charter.net	General
Dan Evans	devans@seattlemariners.com	General
Rob Fabrizzio	rfabrizzio@bigfoot.com	Statistics
Larry Grasso	l.grasso@juno.com	Statistics
Tom Hanrahan	Han60Man@aol.com	Statistics
John Heer	jheer@walterhav.com	Proofreading
Dan Heisman	danheisman@comcast.net	General
Bill Johnson	firebee02@hotmail.com	Statistics
Mark E. Johnson	maejohns@yahoo.com	General
David Kaplan	dkaplan@education.wisc.edu	Statistics (regression)
Keith Karcher	karcherk@earthlink.net	Statistics
Chris Leach	chrisleach@yahoo.com	General
Chris Long	clong@padres.com	Statistics
John Matthew IV	john.matthew@rogers.com	Apostrophes
Nicholas Miceli	nsmiceli@yahoo.com	Statistics
John Stryker	john.stryker@gmail.com	General
Tom Thress	TomThress@aol.com	Statistics (regression)
Joel Tscherne	Joel@tscherne.org	General
Dick Unruh	runruhjr@iw.net	Proofreading
Steve Wang	scwang@fas.harvard.edu	Statistics

Models of Consistency

Abbott Katz

Some players have consistent careers, in the sense that their batting averages don't vary much from season to season. Other players are inconsistent, with their season averages all over the place. Is consistency an actual attribute of the hitter, or is it just a manifestation of normal, statistical variation? Here, the author investigates.

Immortality has at last redounded to Alex Kampouris. After decades of crouching beneath the radar, his day has come, and his story can finally be told: Kampouris may very well be the most consistent hitter in major league history.

Check it out. Of all the hitters since 1900 who've mustered at least five 200+ at-bat seasons[1], it was Kampouris -- an infielder who played for 4 teams in the 30s -- who realized a standard deviation across his yearly batting averages of .0034, the most infinitesimal such figure in the majors in the past 107 years.

Of course, that superlative may qualify as faint praise for Kampouris, as his lifetime average of .243 hardly positions him as a must-have for fantasy-game players or card collectors. Nevertheless, there is something curious, and stirring, about his 5 qualifying seasons:

```
1935    .246
1936    .239
1937    .249
1938    .249
1939    .249
```

Would that my watch ran so dependably.

Among current players, it's Kevin Mench whose yearly output runs straightest and narrowest, registering an steadfast .0093 SD to date. And stranded in bewildering isolation at the other end of the curve is Don Padgett, a first-baseman/catcher for the Cards from 1937-1941 and 1946-48, whose SD of .0582 makes him a statistical outcast - demonstrably baseball's all-time most inconsistent hitter, even in view of his not-at-all-bad lifetime average of .288.

Padgett, who almost surely played against Kampouris, inaugurated his career with a promising .314 in 446 at-bats, but then sophomore-slumped to a .271 in 1938. Following that offensive retreat, he was granted 233 ABs in 1939, and hit…. .399. Padgett then succeeded that mighty and inexplicable season by crash-landing to .242 and .247 the next two years, thus securing his place in the archives (Padgett in fact played 3 additional seasons after World War II, none of which met the 200-at-bat criterion).

These weird and wonderful outlier careers touch off an intriguing question: Pure hitting talent aside, do some players possess a special talent for consistency, for managing to calibrate their performances to within a few points of one another year after year - or have players such as Kampouris merely been flung to the statistical margins by the reliable vagaries of normal distributions?

The answer, one not likely to surprise BTN readers, appears to be the latter. SDs correlate very weakly with lifetime average[2] (.177), and number of seasons (.125), and don't appear to be associated with any other special predilection for constancy -- meaning that better hitters are, if anything, a wee bit more "inconsistent" than inferior ones, albeit at a higher talent point. The average hitter career SD: .0235. Even Old Reliable himself, Tommy Henrich, registers but a bit-better-than-average SD of .0202.

Indeed, by plotting the averages of all 1,901 players between 1900 and 2007 who accumulated those five minimum 200+ at-bat seasons, we find that, among other things, their career SDs are *themselves* more-or-less normally curved, suggesting, of course, that SDs themselves are no less beholden to normal-distributional forces than other metrics. (See Figure 1.)

[1] Data source: Sean Lahman's player database.
[2] Defined here as a player's simple average of all qualifying seasons.

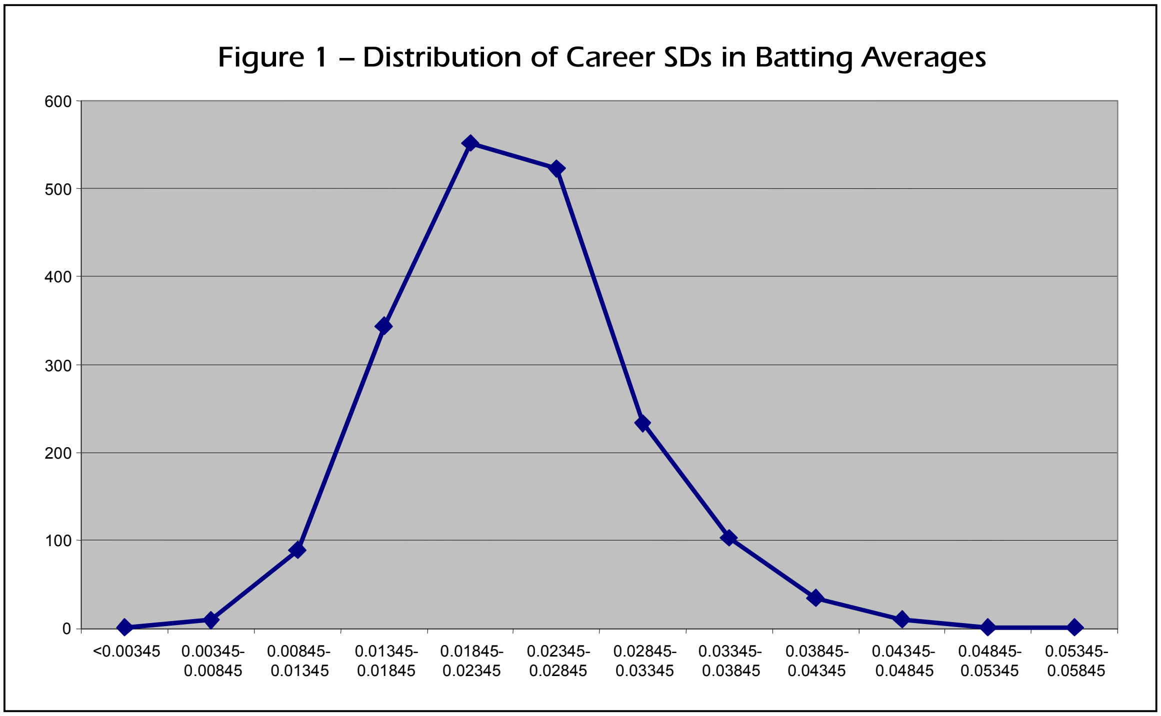

Moreover, players' individual career curves at least suggest this sort of normality as well. One hears about "career years", but every player has a career worst season as well, and the two tend to be positioned at similar removes from the players' lifetime BAs. If you extract both the best and worst seasons from each of the 1,901 hitters, as they do with Olympic ice skating scoring, average those two seasons, and compare those results with the average of the remaining middle years, the mean differential between the two sets is about only 7 points (see Table 1).

Table 1 – Player distribution of best-worst vs. intermediate years differentials

Interval	Freq
0 -0.003	499
0.003-0.006	477
0.006-0.009	295
0.009-0.012	214
0.012-0.015	157
0.015-0.018	117
0.018-0.021	65
0.021-0.024	29
0.024-0.027	27
0.027-0.030	9
0.030-0.033	5
0.033-0.036	2
0.036-0.039	1
0.039-0.042	2
0.042-0.045	2

While few of these disclosures may serve to astonish, one additional finding may prove just a touch more provocative. A small but notable longitudinal contraction in average career SD seems to have come about, if one organizes the data by the players' final years (I've omitted 2007 data, as most of these hitters will continue to play in 2008). See Figure 2.

Without enrolling in the larger debate, these data tend to corroborate the well-known thesis of the late Stephen Jay Gould, who held that .400 hitters have been consigned to extinction owing to a general trimming of BAs on both edges of the curve[3]. If, for example, the effective band of BAs has been tightened from say, .210-.380 to .230-.350, one might then well expect a diminutive but real paring of SDs. But this point calls for additional scrutiny, as Gould's principal interest was in inter-player, as opposed to my intra-player comparisons.

Of course, other takes on "consistency" could be imagined. Bill James does just that[4], defining the parameter by assaying win shares across a player's career, his metric comprising far more moving parts than my lean measure. As might or might be expected, his rankings depart significantly, but not

[3] See, for example http://en.wikipedia.org/wiki/Batting_average#The_decline_of_the_.400_hitter, and http://www.pbs.org/newshour/gergen/november96/gould.htm. The thesis has been disputed, and a fuller exposition is for another time. See also http://danagonistes.blogspot.com/2004/08/where-have-400-hitters-gone.html, for a more extensive consideration.

[4] James' article is called "Measuring Consistency," at http://www.billjamesonline.net/ArticleContent.aspx?AID=158&Code=James01110 (subscription required)

completely, from mine. For example, James identifies Hank Aaron as his all-time most consistent player; I have Aaron slotted at 1628 out of my 1901, though to be fair, SDs separating contiguous players are often minute. Does Aaron's "inconsistency" make him any less great? Of course not; you'd rather have a player oscillating between .290 and .330 than one who beats out a yearly, bet-the-house-on-it .245. On the other hand, both James and I consign Rico Carty to the inconsistent sector; for me, he ranks 1842. Perhaps some bright-eyed researcher could correlate the two modes of consistency.

In any case, a statistical ingénue might be moved to wonder if superior hitters - who, after all, are better at what they do, and thus by definition exert greater control over their environment - could be expected to muster a greater consistency across their batting averages. But it doesn't seem to work that way, does it? The apparent limits of human precision seem to hold rather adamantly -- and that suggests in turn, that, had they been given a career do-over, Alex Kampouris and Don Padgett would have installed themselves closer to the overall SD average - even though, if statistical form holds, they'd have likely ended up near their respective .243 and .288, just the same.

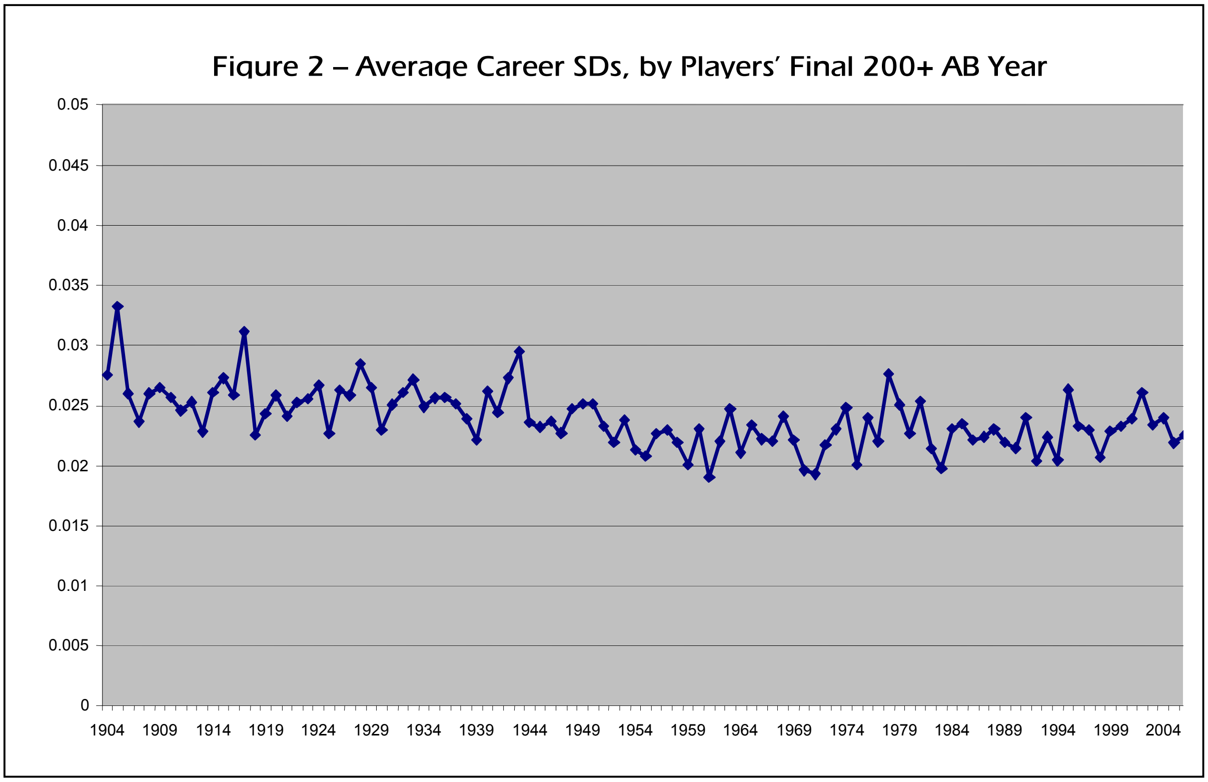

In light of these existential musings, some enterprising grant seeker might want to go on and study standard deviations in other sporting -- and daily -- activities. What, for example, is the characteristic SD among NBA foul shooters? Of minor league ballplayers? Of three-game series scores of PBA bowlers, or even Scrabble turns? Or college student GPAs? Again, the issue isn't talent as such - but rather consistency, at whatever competence point.

Hmmm…anyone have the address of the Guggenheim Foundation?

Abbott Katz, akatz@hotmail.com ♦

Submissions

Phil Birnbaum, Editor

Submissions to *By the Numbers* are, of course, encouraged. Articles should be concise (though not necessarily short), and pertain to statistical analysis of baseball. Letters to the Editor, original research, opinions, summaries of existing research, criticism, and reviews of other work are all welcome.

Articles should be submitted in electronic form, either by e-mail or on CD. I can read most word processor formats. If you send charts, please send them in word processor form rather than in spreadsheet. Unless you specify otherwise, I may send your work to others for comment (i.e., informal peer review).

If your submission discusses a previous BTN article, the author of that article may be asked to reply briefly in the same issue in which your letter or article appears.

I usually edit for spelling and grammar. If you can (and I understand it isn't always possible), try to format your article roughly the same way BTN does.

I will acknowledge all articles upon receipt, and will try, within a reasonable time, to let you know if your submission is accepted.

Send submissions to:
Phil Birnbaum
88 Westpointe Cres., Nepean, ON, Canada, K2G 5Y8
birnbaum@sympatico.ca

Get Your Own Copy

If you're not a member of the Statistical Analysis Committee, you're probably reading a friend's copy of this issue of BTN, or perhaps you paid for a copy through the SABR office.

If that's the case, you might want to consider joining the Committee, which will get you an automatic subscription to BTN. There are no extra charges (besides the regular SABR membership fee) or obligations – just an interest in the statistical analysis of baseball.

The easiest way to join the committee is to visit http://members.sabr.org, click on "my SABR," then "committees and regionals," then "add new" committee. Add the Statistical Analysis Committee, and you're done. You will be informed when new issues are available for downloading from the internet.

If you would like more information, send an e-mail (preferably with your snail mail address for our records) to Neal Traven, at beisbol@alumni.pitt.edu. If you don't have internet access, we will send you BTN by mail; write to Neal at
4317 Dayton Ave. N. #201, Seattle, WA, 98103-7154.

A Simple Method for Estimating Pitch Counts

Charlie Pavitt

Here, the author presents a method for estimating the number of pitches thrown in a game, based on a few simple numbers from the standard box score.

Pitchers aren't as tough as they used to be. Back a hundred years ago, they were real men: completing almost every game they started. But today's namby-pambies are given accolades if they complete seven innings. Wimps.

Well, I don't know about that. It could be that modern-day pitchers have to make significantly more pitches to get three outs in the typical half-inning than in much earlier times. Circumstantial evidence supporting this possibility is the marked increase in walks and, in particular, strikeouts over the interim. Using the 6th edition of *Total Baseball* as my readily-handy source, a game included an average of 6.70 walks, and 13.44 strikeouts a game in 1998, the last year included in the book; a hundred years previous, those respective numbers were 5.52 and a scant 4.62. Fewer strikeouts and fewer walks probably meant fewer pitches a hundred years ago. Unfortunately, we do not have and probably will never have much in the way of pitch count data before Project Scoresheet.

The only direct evidence of which I am aware was presented by Dan Levitt in the 2000 *Baseball Research Journal.* Apparently, the 1920 *Spalding's Guide* included pitch counts from the tainted 1919 World Series; Levitt compared this data with that of the 1997 Series and calculated a increase over time in average pitches per inning from 13.1 to 17.8, due both to more batters per inning (4.1 versus 4.6) and pitches per plate appearance (3.2 to 3.9). Consistently with my conjecture, the percentage of walks per plate appearance went from 6.9 to 13.0; for strikeouts the figures were 8.9 and 16.9. Dan ended his piece with the following: "Assuming the number of pitches as the correct measure of pitcher workload, on average, a 300-inning season in 1919 would be the equivalent of 221 innings today. Using the available World Series pitch count data, pitchers today work just as hard as their Deadball Era brethren; they just do it in fewer innings" (page 47).

Although we lack data, we ought to be able to estimate pitch count based on what information is available and a number of (perhaps questionable) assumptions. Tom Tango has worked on this problem (http://www.tangotiger.net/pitchCountEstimator.html -- reviewed by Phil Birnbaum in the May, 2003 *By the Numbers*) using 1999-2002 data and provided the following multiple regression predictor

```
Pitch Count = (3.3 * Balls in Play) + (4.8 * Strikeouts) + (5.5 * Walks)
```

which he felt was sufficiently accurate for most pitchers but not for those with extreme numbers of balls in play (Randy Johnson and Brad Radke his low and high examples). To increase accuracy, Tom computed specific formulas for estimating pitches per ball in play, walk, and strikeout that improves predictive accuracy; these are available from the online article.

I trust that what Tom did is good for its purposes, but I was after something different: a simple pitch count estimator using data that you can easily get from boxscores (which is not true for balls in play, which takes more work than I'd want). So, what makes a pitcher have to make more pitches in an inning? Giving up hits and walks. Strikeouts usually take up more pitches than balls in play, so I included that too. I decided not to worry about hit by pitches because they are relatively rare. I chatted about the issue with Editor Phil and he suggested double plays and caught stealing; I suppose I could have included the former but I didn't, and, as caught stealing data wasn't kept before 1950, I have to ignore that one. I did sum data across both teams, leading to the full game as constituting the unit of analysis; maybe I shouldn't have, but I did.

Now, to complicate matters, all games are not the same length. A few more go 17 half-innings than 18 because of the 54 percent home field advantage. Then there are also games shortened by bad weather and lengthened by extra innings. Dividing the raw data by number of half innings solves that. The result is the following hypothesized multiple regression model:

$$\text{Pitch Count} = \alpha + (\beta_1 * \text{hits allowed/half inning}) + (\beta_2 * \text{walks/half inning}) + (\beta_3 * \text{strikeouts/half inning})$$

In addition, I examined to see if any of the relationships between pitch counts and predictors was nonlinear, as it apparently was for extremely high balls in play in Tango's data (see the diagram in his article), by squaring each predictor, and for two-and-three-way interactions among the predictors through multiplicative terms. Before the fact, I sure hoped none of these would have substantial impacts, because then the model would lose its intended simplicity.

I gathered one piece of additional data that I was interested in, game time, which Dave Smith told me were first appearing in boxscores during the 1920s. I wanted to see how strongly game time correlated with pitch count and also the validity of the same regression model and its possible complications with that rather than pitch count as dependent variable. For this reason, I wish that I had gathered one additional piece of data available to add to this regression model: within-inning pitching changes, which take up a considerable amount of time.

I thought it fitting to eschew existing data bases and input the data directly from boxscores (fortunately I had the time, and its mindlessness was a nice break from my usual work). My sample, all 2008 games, consisted of every boxscore published in the morning edition of the *Washington Post* from June 17th (right after I thought of doing this project) through August 7th and August 17th through the end of the season, with a vacation break interrupting; plus a few missing from the *Post* but appearing in *USA Today* when a houseguest left one. The reason I didn't add double plays or within-inning pitching changes to the data set was that I didn't talk to Phil until I had recycled a lot of the newspapers after data entry. Anyway, I hope the data set is approximately representative of 2008 baseball. It certainly is sufficiently large for my purposes: 1185 games. I estimated statistical power using *G*Power 3*, a wonderful, free on-line program from statisticians at Dusseldorf University [1] that I strongly recommend to all my friends and neighbors; just google its title, download it, and honor the authors' request to cite them if you use it in publicly-distributed research. With the basic three predictors, statistical power for the regression equation as a whole is a perfect 1.0 for what is conventionally considered to be both large and medium effect sizes and .97 for small.

As I intended to keep everything simple, SPSS was sufficient for the purpose. As for the results, we begin with descriptive indices in Table 1:

Table 1 – Descriptive Statistics

	Mean	Std. Deviation	N
PitchCount	293.58	39.804	1185
GameTime	174.58	25.925	1185
hitsperinn	1.0246	.28809	1185
walkperinn	.3690	.17196	1185
wiffperinn	.7714	.22313	1185

Keep in mind that the latter three rows are based on half-innings, which averaged 17.84 per game; multiplying gives us 18.20 hits, 6.60 walks, and 13.74 strikeouts per game, comparable to those for 1998 mentioned at the beginning of this report (there were 17.88 hits per game that year, which hasn't changed much from 1898's 18.42).

Table 2 shows the correlations of the various measures with each other.

Table 2 – Correlations

	PitchCount	GameTime	hitsperinn	walkperinn	wiffperinn
PitchCount		.896**	.390**	.566**	.176**
GameTime			.369**	.495**	.131**
hitsperinn				.153**	-.136**
walkperinn					.049
wiffperinn					

** significant at the 0.005 level (2-tailed). The remaining correlation of .049 is significant at .091.

[1] Described in F. Faul, E. Erdfelder, A.-G. Lang, and A. Buchner, G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods,* 2007, volume 39, pp. 175-191

The correlation between pitch count and game time is extremely high, to the extent that the two are almost interchangeable in practice. Those comparing pitch count and game time with both walks and hits per innings are, not surprisingly, reasonably substantial, as they both mean extra hitters for the pitcher to face. Those between the former two and strikeouts per innings, in contrast, are statistically significant but modest, implying that strikeouts don't take up that many more pitches and time than other ways of getting outs. Happily, those among the three predictors are all less extreme than .2, which means absence of multicollinearity such that the three all serve as independent factors in the following regression equations.

Table 3 presents the basic regression model for predicting pitch count. All three are significant predictors, with their relative importance (as indicated by beta) as expected given the strength of the correlations with pitch count examined just above. The model as a whole accounts for 45.3 percent of the variance in pitch counts. I was pleased to find that adding the squares of the predictors increased this figure by only 0.4, and including all interactions among them by a tiny 0.1, keeping things simple.

Table 3 – Coefficients for Predicting Pitch Count

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	175.151	4.791		36.556	.000
	hitsperinn	47.011	3.042	.340	15.455	.000
	walkperinn	116.672	5.055	.504	23.083	.000
	wiffperinn	35.280	3.886	.198	9.080	.000

So, we would predict pitch count with the following formula:

```
175.151 + (hits per inning * 47.011) + (walks per inning * 116.672) + (strikeouts per inning *
35.280)
```

We have no idea how many innings the average game was in 1898, so let us use the mean of 17.84 computed here. This means that for 1898, we would predict an average of

```
Pitches = 175.151 + ([18.42/17.84] * 47.011) + ([5.52/17.84] * 116.672) + ([4.62/17.84] *
35.28)
        = 175.151 + 48.54 + 36.10 + 9.14
        = 268.93
```

which was a surprise to me. The implication is that the average team's staff needs 147 pitches to complete a game now, and 134 a hundred ten years ago. Not much difference. The culprit seems to be the size of the regression constant, implying that getting 53.52 outs (17.84 multiplied by 3 outs per inning) one way or another should take about 175 pitches. There aren't that many more baserunners per game now than then (24.80 now versus 23.94 then). It all works out to the average number of batters per inning changing from 4.34 to 4.39 (all this assuming the 17.84 figure).

Given that Levitt's results made me think the over-time difference would be greater, let's look once more at his 1919/1997 Series comparisons. His contrast of 4.1 versus 4.6 batters per inning is much greater than what is found here, an increase of 12.2 percent as compared to 1.1. His difference in pitches per plate appearance (3.2 to 3.9) implies a rise of 21.8 percent. This inspired me to determine the predicted pitch count for 1919. The equation there was

```
Pitches = 175.151 + ([17.55/17.84] * 47.011) + ([5.35/17.84] * 116.672) + ([6.12/17.84] *
35.28)
        = 175.151 + 46.25 + 34.99 + 12.11
        = 268.51
```

which is almost the same as for 1898 and, as such, not much different than for 1998. Was offense unusually low for the 1919 World Series?

Just for the information, Table 4 presents the coefficients for predicting game time, accounting for 35.7 percent of the variance.

Table 4 – Coefficients for Predicting Game Time

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	106.628	3.384		31.507	.000
	hitsperinn	29.083	2.149	.323	13.536	.000
	walkperinn	66.069	3.570	.438	18.506	.000
	wiffperinn	17.857	2.745	.154	6.506	.000

All of this presumes that other factors impacting on pitch count and game time have undergone no basic changes over the interim. I doubt that this is the case for the latter. For example, I would guess that the time between innings used to be shorter before television required enough of an interval to fit in a sufficient number of commercials. I would also be confident that within-inning pitching changes, which as I noted earlier I wish I had counted, have increased across time. In addition, there is no good way of knowing whether there were as many Mike Hargroves and Al Hraboskys slowing up the proceedings a hundred years ago, as many visits to the mound by catchers and coaches without pitching changes, and as many on-the field arguments bringing the manager out onto the field (John McGraw versus Earl Weaver??). As for pitch count, all this assumes that games also consumed about 17.84 innings a hundred years ago. More to the point, it assumes that the 175.151 pitches just to get a game's worth of outs was valid back then. It may not have been, particularly if batters customarily did not let counts advance as much as now. Levitt's findings imply that this number may have been about 10 percent lower.

But unless a treasure trove of hundred-year-old detailed game accounts shows up in an attic someday, or someone invents a time machine and some of us volunteer to go back for a summer scoring games (an attractive opportunity), I fear we will never know for sure.

Charlie Pavitt, chazzq@udel.edu ♦