Diary of a Mad Sports Bettor

by James Holzhauer

In 2006, I took a hiatus from my job search to try my hand at betting on baseball for a living. Five years later, I retired at age 26. These are my stories.

I graduated from college in 2005. At that time I was dabbling in sports betting to pay the bills while I interviewed to become an actuary. Those interviews were full of banalities, topped by the world's most asinine question: "Where do you see yourself in five years?" I don't remember my exact response, but I certainly didn't explain that Tampa Bay would go from worst to first, shocking the world and setting me up financially for life. (At that time, I suspect that discussing the Devil Rays as a potential contender might not have been the best way to make myself look hireable.) I received zero job offers, but I did appreciate the irony of insurance companies not hiring someone because he seemed uncertain about the future.

Where do I see myself in 2016, five years from now? I can only guess, just as I can only guess who will win this year's World Series, AL Cy Young, or Ford Frick Award. Forecasts—for life or for baseball—are just an educated stab in the dark, and we should stop pretending this isn't the case.

I first started betting baseball seriously during the World Baseball Classic in March 2006. Many people would opt not to gamble on games featuring players they've never heard of; as a cocky 21-year old, I had no such misgivings.

Placing winning sports bets often hinges on an information advantage. The problem is that most publicly available information is already factored into the betting odds.

Consider the example of home-field advantage. The home team in baseball wins about 54 percent of the time; if bookies were not aware of this fact, you or I could make a fortune just blindly betting on every home team. Because the sports betting market adjusts the odds to reflect the home-field advantage, this information alone will not give anyone an edge. (If, however, I have reason to suspect that a specific team has a greater or lesser home-field advantage than 54 percent, and the betting odds don't reflect this outlier, that could be used as part of a profitable approach to gambling.)

If you walk into a sports book and eavesdrop on bettors explaining the logic behind their picks, they often say things like "Lincecum is on a roll" or

"The A's have no offense." Everyone knows these things already! Widely known information is of little or no value in sports betting.

Bookies offer an option called futures wagering, where the bettor backs a certain team to win its division, the League Championship Series or the World Series. We can't predict the future, but we can guess more accurately than the bookie. High-stakes gamblers tend to stay away from futures because they have low maximum bets and require complex math, but those who are savvy with statistics and simulations can make a killing on them.

Inside information is a different animal. There's a classic episode of *The Simpsons* where the degenerate Moe comes to the Simpsons' house before young Bart and Lisa face off in a pee-wee hockey game:

> *"How are the little kids doing? I mean, really, how are they doing? Any disabling injuries, something, say, that the gambling community might not yet know about? (grabs Bart's leg) Come here, let me see those knees."*

Moe didn't get what he was looking for, but the trainer for a major league team knows when his staff ace is more fatigued than usual or when his star first baseman is about to get a day off. This information, unknown to the public, would be of tremendous value to someone looking to place a large bet against that team. One hopes that no current trainers are committing the baseball equivalent of insider trading.

I felt I could have a legitimate information advantage in betting the World Baseball Classic. The bookies don't pay as much attention to lightly bet events because it's more important for them to focus on accurately handicapping NFL or NBA games, where many more bets are coming in. Hours of thought go into setting the point spread for an NFL game, but a bookie isn't going to waste time looking up Michihiro Ogasawara's wOBA; it's far more efficient for him to just make a guess at the correct game odds and limit his risk by not accepting large bets. (This is why small bettors can get their biggest edges on futures bets, which high rollers typically ignore.) I studied the WBC teams intently, determined to know far more about them than the bookies did.

Clay Davenport translated each team's statistics so they could be compared against each other, and his methodology seemed mathematically sound. Comparing the translations to the bookies' odds, it seemed clear that the smart bets were on the underdogs to topple the favored Americans and Dominicans. Canada was a 9-1 underdog in its game against the USA: my $1,000 bet collected $9,000 in profit when the Canadians pulled off the upset. Other bets I cashed included Mexico to win its initial pool (12-1) and Japan to win the tournament (30-1). I ended up making 50 cents of profit for every dollar I bet on the WBC, an encouraging start to my baseball betting career.

The World Baseball Classic gave me both the money and the confidence I needed to approach gambling seriously for the 2006 major league season. I programmed a spreadsheet to simulate the season thousands of times, and loaded it with projections from PECOTA, ZiPS and CHONE. The results recommended a number of wagers, but three stood out above the rest.

The first was a bet on who would lead the majors in home runs in 2006. Bookies typically set these odds by listing last year's top homer hitters as this year's favorites. With Barry Bonds and Sammy Sosa fading fast, there was no clear top slugger in the majors. The Phillies offered a good candidate: a new first baseman who had destroyed pitchers in Triple-A, slugged 22 dingers in half a season in the bigs in 2005, and was projected by PECOTA to lead the majors with 46 homers in '06—no one else was forecast for more than 41.

Because he was unproven and hit "only" 22 homers the year before, the sportsbooks didn't even list Ryan Howard as a candidate to win the home run crown until I wrote an email to the bookie requesting that he be added as an option. Howard, of course, smashed 58 homers en route to an MVP award. Unbelievably, a bet on Howard that year paid 40-1 odds: a $250 max bet cashed for a $10,000 profit.

The second and third recommended bets resided in the AL Central. In 2005 the White Sox had won 99 games and the World Series, while the Indians had looked like potentially the best team in baseball before choking away a playoff spot in the season's final week. Naturally they were listed as the two big favorites to win the division and advance in the playoffs, but the projected standings showed the Twins and Tigers only slightly behind Cleveland and actually *ahead of* Chicago. Joe Mauer, Justin Morneau and Francisco Liriano seemed poised to become household names for the Twins in 2006, while the Tigers had Curtis Granderson and Justin Verlander as their aces in the hole.

Fixated on the 2005 standings, the bookies listed the Twins as 7-1 underdogs to win the division, and the Tigers at an astonishing 30-1. After leading the Twins by 12 games at one point, Detroit lost the division crown to Minnesota in one of the worst late-season collapses ever, but both teams left the White Sox and Indians in the dust. The Tigers won the Wild Card and rebounded to take the AL pennant, which paid a whopping 60-1: a $1,000 max bet paid out sixty grand in profit. Had they won the World Series (Vegas made them a 2-1 favorite over the Cardinals before the first pitch) that would have paid 150-1.

The most important advice I can give regarding futures betting is that bookies tend to underestimate the frequency of low-probability events. Surprise playoff teams like the 2011 Diamondbacks aren't nearly as uncommon as we think: The 2002 Angels, 2003 Marlins, 2005 White Sox, 2006 Tigers, 2007 Rockies, 2008 Rays and 2010 Giants each entered those seasons regarded as 83-win or worse teams. All seven advanced to the World Series, and four won it all.

If you decide to bet futures, focus on a few undervalued longshot teams. You might catch this year's lightning in a bottle.

Before the start of every season, many pundits come out with their projected standings. If today's best prediction model forecasts 95 wins for the 2012 Yankees, how accurate will that forecast be?

The Yankees will deal with injuries, breakouts, collapses, trades, midge attacks and the vagaries of random chance. All of these cloud our forecast: My estimate is that they will win exactly 95 games just 4.4 percent of the time. That's a one-in-23 shot, which means that if we project the full major league standings, we will nail the W-L record of only one team in 23, or about one team per season.

What if we relax our standards of accuracy? I estimate that the Yankees have a 45.9 percent chance of winning 90-100 games. Our "95-win" Yankees will actually finish below .500 more frequently (one time in 18) than they will go 95-67. They also have a one-in-18 shot of winning 110 or more games. I believe that given the current limits on our ability to predict the future, this is the best we can do: miss the target by 15 or more wins in either direction "only" one time out of nine.

The White Sox entered 2007 with essentially the same personnel that had averaged 94.5 wins the previous two seasons. Many fans saw no reason that they wouldn't be competitive again, but Nate Silver's computer disagreed. Silver's PECOTA projections suggested a 72-90 finish for the South Siders, a sudden and steep dropoff. PECOTA saw an old lineup that would likely regress heavily; the 2006 team had benefited from unexpectedly big seasons from Paul Konerko, Jim Thome, Joe Crede and Jermaine Dye.

The *Chicago Tribune* featured an article, "Computer Crashes White Sox," which quoted GM Kenny Williams bashing the forecasting methods: "Don't you get tired of being wrong or is it you figure one of these days you're going to be right?" Konerko sarcastically laughed off the forecast: "Well, we're screwed now."

Bettors, meanwhile, saw an opportunity. Each year, sportsbooks post Over/Under odds on season wins for every team. If a team has an Over/Under of 80.5 wins, those betting on the Over need the team to win 81 or more times that season to cash their bets. If they win 80 or fewer games, the Under bettors collect. When many more people bet the Over than bet the Under, the bookie will move the line upward to 81, 81.5, or maybe even 82 or 82.5.

This serves two purposes: it entices more people to bet the Under and balance the books; and it limits the amount that can be won by wiseguys betting on the Over. (Even if a professional likes the Over on 80.5 wins, he understands that the team will win exactly 81 or 82 games nearly 9 percent of the time—a huge difference—so he will stop betting once the line moves too far.)

Sportsbooks released an Over/Under of 89.5 wins for the 2007 White Sox, essentially flipping PECOTA's 72-90 forecast. This number was eventually bet down to 86 wins. That difference, 3.5 wins, is an unusually large move, indicating heavily unbalanced action on the Under—clearly some people believed in the computer forecast. In the end, the old offensive core tanked just as PECOTA predicted and the White Sox finished precisely 72-90. Score one for the computer.

The 2007 World Series pitted the Rockies, who had won 21 of their last 22 games, against the Red Sox, who were the far superior team on paper. Would Colorado's magical run continue or end in heartbreak?

The media often portray sabermetricians as heartless robots for not believing in the predictive power of clutch hitters or hot streaks. Clearly there were many believers in the Rockies: Based purely on the numbers, Boston should have been about a 75 percent favorite in the Series; but the

actual odds had them as only a 67.5 percent favorite. This level of disparity is rare in a heavily bet market, and it was a golden opportunity for me to speculate that the streak wouldn't continue. To me, betting on the Red Sox was like buying a stock for 10 percent less than its value.

Colorado looked outclassed right from the start, losing 13-1 in Game One. The Red Sox swept the Rockies, ignoring Alex Rodriguez's attempt to steal the spotlight by opting out of his contract during the clinching game.

When I was 13, I told my dad that there should be a stock exchange for sports teams, which would allow me to turn my obsession with statistics into a profitable venture. It turns out that the markets for sports betting and stocks are very similar: in each case, you have a large population of speculators gambling that they can outsmart a weakly efficient market.

If investors are rushing to buy Google stock, the price of a share will go up. If bettors are rushing to back the Yankees against the Red Sox in tonight's game, the betting odds will change, effectively increasing the price of a bet on the Yankees. Just as buyers and sellers determine the market price of a stock, bettors determine the market odds on a contest. Your bookie has no more control over the odds for a Yankees-Red Sox game than your broker has over the price of a share of Google; he's just there to execute your trade and take a commission.

The big difference—legality aside—is that a typical person can expect to make money investing in stocks and lose money betting on sports. But there are exceptions: the Tampa Bay Rays outperformed the Dow Jones pretty handily in 2008.

The 2008 Rays were quite the conundrum. The previous season, they had finished with the worst record in the league at 66-96 despite above-average production from their hitters and starting pitchers; all of this good work was undone by a historically bad bullpen and a historically bad defense. These are two aspects of team performance that tend to regress heavily toward the mean the next year, but the Rays took no chances: They underwent a complete bullpen makeover and made several moves to shore up the glovework.

The 2007 squad featured the iron gloves of Delmon Young, Brendan Harris, Ty Wigginton and B.J. Upton; the former three were shipped out of town and Upton was shipped out of the infield. On paper Tampa Bay actually came into 2008 with an above-average defense, as hard as that was to believe.

The World Series winner in any given year is usually not the best team in baseball. It's become popular to suggest that once the playoffs start, everyone has about a 1-in-8 shot of taking the trophy, and we might as well determine the winner via a series of coin flips.

This is a total oversimplification. In a typical year, the best team in baseball has a 25-30 percent chance of winning three playoff series. That's not as Calvinist as the NBA playoffs, but it's still a meritocracy of sorts. A truly great team will be a big favorite to prevail over a lesser squad; the "problem" is that major league baseball playoff teams are rarely so badly mismatched.

In Las Vegas, locals are inundated with ads from "touts" who claim they can win 75 percent of their sports bets. If this claim were accurate, a good tout could bet his own picks and easily turn a thousand bucks into millions within a year. It's worth asking why he would rather sell you his Five Star Lock of the Week for the bargain price of $19.95.

The sports betting marketplace is smart and when its odds are off, it's not by much. My lifetime return on investment betting individual baseball games was about 2 percent, and professional sports bettors are happy to do that well. Jonah Keri wasn't kidding: One has to rely on small edges adding up, whether it's on Wall Street, in a major league front office, or betting sports.

The Rays also featured an impressive young lineup with no real holes and three players (Carl Crawford, B.J. Upton, Carlos Pena) who had put up star-level performances in 2007. Top prospect Evan Longoria looked ready to star in The Show.

The rotation was anchored by the solid trio of Scott Kazmir, James Shields and Matt Garza. If the bullpen held together, there looked to be no reason the Rays couldn't win 90 games—after never winning more than 70 in their franchise history. In fact, PECOTA projected them to do just that, offering a 90-72 forecast. This drew the attention of a feature story, "Leap Year," in *Sports Illustrated*. Most readers laughed it off as another case of a computer ignoring the obvious.

This was a perfect storm for sports bettors. Here we had a perpetual last-place team that was bound for a huge improvement due to regression and smart player management. Unlike, say, the 2010 Miami Heat, the Rays made moves that flew under the radar and only a few sabermetrically inclined people took notice. The best betting values occur when a team has gotten significantly better or worse without attracting mainstream attention.

The 2008 Rays Over/Under opened at 68.5 wins. Betting opened before PECOTA and other forecasts became public, but the wiseguys were ready ahead of time. Within a few hours, enough money been bet on the Over to move the line to 71 wins. A day later it was 72.5, then 74 the next day. Eventually the line settled at 76.5 wins—an unprecedented eight-win move. It was as if the Rays had traded Cliff Floyd for Albert Pujols.

Of course, it all came together for Tampa Bay that year. The pitchers' BABIP went from a league-worst .331 in 2007 to a league-best .277 in 2008, a remarkable defensive turnaround. Despite an increase in xFIP from the year before, the starters posted a combined ERA of 3.95 versus a 5.20 mark in 2007. The Rays' bullpen ERA, an eye-popping 6.16 in 2007, came all the way down to 3.55, fifth best in the majors. Longoria burst out of the gate to anchor a lineup featuring nine regulars with 1.9 WAR or better.

Tampa Bay went 97-65, capturing the AL East crown and shattering even the loftiest of preseason expectations.

With so little mainstream media attention paid to the Rays in the preseason, it was no surprise that the bookies didn't give them much respect in the futures markets either. If the bookie feels a team has no chance whatsoever, he might simply make up odds that seem high enough to entice people to throw away some dollars on a prayer. The Rays were 50-1 longshots to win their division. After all, even if they overachieved,

how could lowly Tampa Bay ever hope to topple the mighty Red Sox and Yankees? A maximum bet of $1,000 cashed for a cool fifty grand. They were 125-1 underdogs to win the ALCS; the books took a maximum of only $500 on this, but that was enough to win $62,500. A bet on the Rays to win the World Series would have paid 300-1…and they opened the Series as 7-5 favorites over the Phillies.

It's fun to check your favorite team's live in-game win probability at FanGraphs, but its models assume each game begins as a 50-50 contest. How can we use the Vegas betting line to determine the *a priori* chances that a team will win tonight?

You may have looked at the odds in the sports section of your local newspaper and seen something that looks like this:

Phi	Halladay	-130
Atl	Hanson	+120

What these numbers mean is that bettors who want to take the Phillies must *risk* $1.30 for every dollar they hope to win, while Braves backers will *win* $1.20 for every dollar they risk. So, you could bet $13 on Philadelphia to win back $10, or $20 on Atlanta to win back $24, and so on.

The difference between the lines on each team is like the bid-ask spread on a stock: if we believe that the market is efficient, the odds must be set so that neither Atlanta bettors nor Philly bettors can make a quick profit. In gambling terms, this means that bets on both teams do not offer a favorable expected value (EV).

To calculate the EV of a bet, multiply the probability of winning by the amount that can potentially be won. For example, many states have a pick-3 lottery where you must correctly guess three ping-pong balls with values between 0 and 9 inclusive. If you match all three, you win $500. What is the EV of a lottery ticket?

We essentially need to match a three-digit combination from 000 to 999. There are one thousand such combinations, so the one you pick has a 1/1,000 chance of winning. Multiplying this probability by $500, we get an EV of 50 cents. The state will happily sell you a ticket for a dollar, knowing it is worth only half that much. If it instead sold tickets for 40 cents or raised the jackpot to more than $1,000 (thus giving each ticket an EV of more than $1) the lottery would effectively be giving money away.

Back to our baseball example: A Braves bettor can risk $100 and get back $220—$120 in profit plus his initial $100—if they win. The bookie

MLB is often praised for its competitive balance because nine different franchises won the World Series from 2001-2010. During that decade, the Spurs and Lakers combined for seven NBA titles, while the Patriots won three Super Bowls in a four-year span. Baseball looks pretty good in this comparison, but this isn't a good way to measure systemic imbalance.

Baseball produces fewer championship dynasties for two reasons. First, its playoff structure gives little advantage to the best teams: five or seven games aren't nearly enough to distinguish a 95-win team from a 90-win team. Furthermore, home-field advantage is practically meaningless in a best-of-seven series—the home team would win only 51 percent of playoff series if the teams were otherwise evenly matched.

wants to set the odds so the EV of this bet is less than $100; otherwise he would be selling you something for less than it's worth, like a foolish lottery operator offering a pick-3 ticket for 40 cents. The EV is calculated by multiplying Atlanta's chances of winning by the $220 we receive when the Braves do triumph:

EV = A%*220

If the EV of our ticket is exactly $100, we have a bet that will break even in the long run. This break-even case occurs when Atlanta's chances of winning are 45.5 percent:

A%*220 = 100
A% = 100/220 = 45.5%

If Atlanta's chances of winning are greater than 45.5 percent, they make for a profitable bet. So, assuming an efficient market, Atlanta's win probability must be less than this.

Similarly, a Phillies bettor can risk $130 and get back $230 if they win. The bookie wants the EV of a bet on the Phillies to be less than $130. Now a break-even winning percentage is 56.5 percent:

EV = P%*130
P% = 130/230 = 56.5%

If Philly's chances of winning are greater than 56.5 percent, this bet is profitable, so in an efficient market Philadelphia will win less frequently than this. Note that the sum of our two break-even win rates is 102 percent. This implies that with equal betting action on both teams, the bookie will pay out $100 for every $102 he takes in; he keeps the rest as his commission. The break-even percentages must always add up to more than 100 percent or the bookie would be giving money away.

If Philadelphia wins less than 56.5 percent of the time, it means Atlanta wins more than 43.5 percent of the time. We have now pegged Atlanta's chances of winning at somewhere between 43.5 percent and 45.5 percent, a pretty good estimate.

Today there are some online sportsbooks with minimal spreads between the odds on each team. For example, we might see Philadelphia -125 / Atlanta +124. Assuming an efficient market, we have pretty much

Second, baseball teams are less reliant on the impact of a single franchise player. The Cavaliers collapsed after losing LeBron James and the Colts collapsed without Peyton Manning; the Cardinals lost Adam Wainwright for the season and advanced to the World Series anyway.

Acquiring and developing Tom Brady or Kobe Bryant isn't a built-in advantage, but the Yankees' budget is. Baseball isn't any fairer; it's just a harder sport to dominate.

pinpointed the Braves' chances of winning at between 44.4 percent and 44.6 percent. Indeed, studies of gambling odds show that teams with these betting lines will win at almost exactly the frequency we calculate using this method.

Of course the sports betting market is not perfectly efficient—I couldn't have made a living if it was—but it's close enough that we can be reasonably confident the Phillies are about a 5-4 favorite to win this game.

In *Popular Crime*, Bill James describes how he used to anger his teachers by writing funny notes to his friends and tinkering with sports statistics during class. Later, he became the first accessible author of baseball analyses, precisely because he learned how to write from humorous notes instead of essay prompts. If Bill James didn't have such an endearing writing style, he would still be toiling in obscurity and sabermetrics may never have gotten off the ground.

I changed my focus of study five times in college. My degree is in mathematics, but if you ask me what I majored in, the most honest answer would be either fantasy baseball or Internet poker. Math was always my favorite subject in school, but my friends convinced me that a math degree wouldn't get me a desk job. This prophecy turned out to be both accurate and irrelevant.

The most important lesson the past five years have taught me: To be happy, make your own way in life. Study the things you're truly passionate about, in the classroom and in your free time. If you search hard enough, you'll find a way to make a living from it.

(Unless you're majoring in English; then you're actually screwed.)

There's no magic formula to build a team for the playoffs, but one component that will improve a squad's postseason performance is rarely mentioned: a terrible fifth starting pitcher. You hear all about the increased reliance on closers and staff aces in the playoffs, but no player sees his role change more in October than the fifth starter, who almost never throws an important pitch in the postseason. The 2009 Yankees appreciated sending CC Sabathia to the mound, but not as much as they appreciated getting Sergio Mitre off of it.

If two playoff teams look roughly equal, the one with the worse fifth starter is likely the better bet. Call it the Zito Theory.

- James Holzhauer

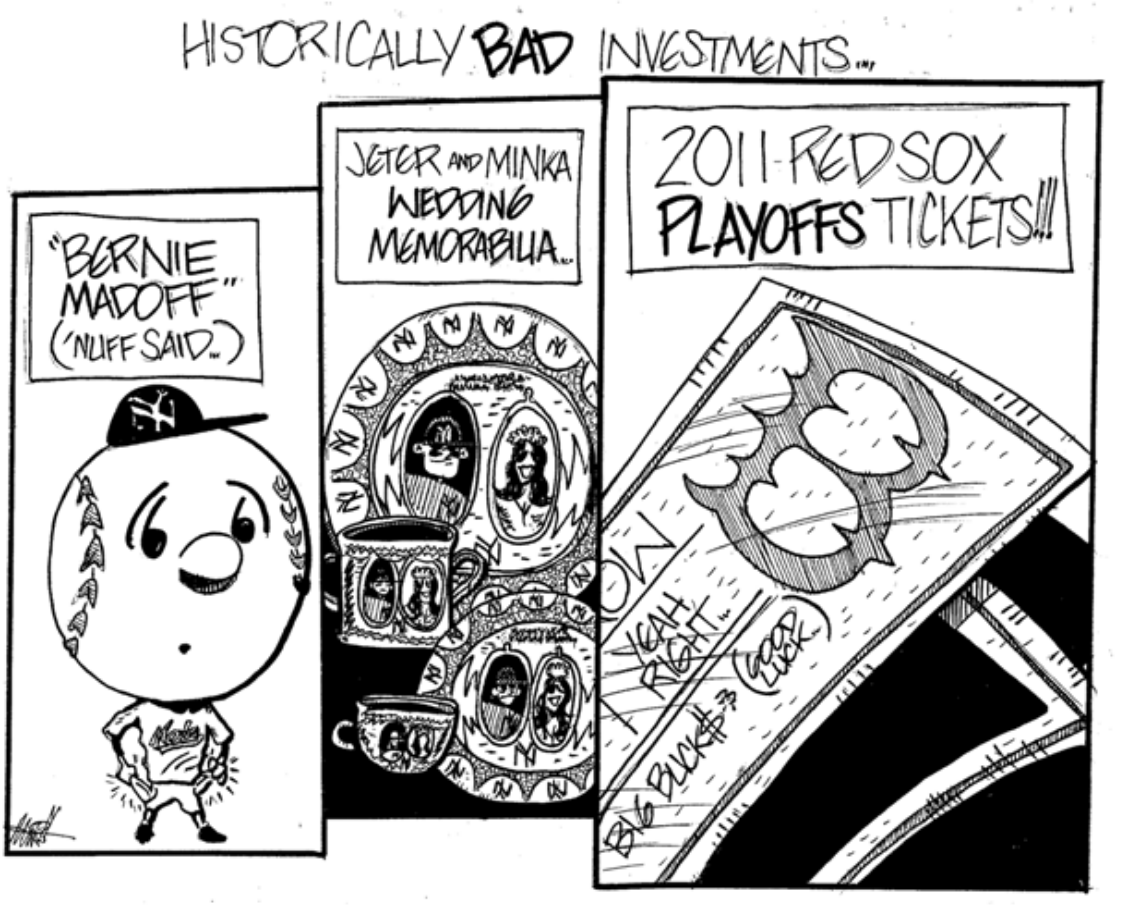
HISTORICALLY BAD INVESTMENTS...
"BERNIE MADOFF" ('NUFF SAID...)
JETER AND MINKA WEDDING MEMORABILIA...
2011 RED SOX PLAYOFFS TICKETS!!!
YEAH RIGHT...
BIG BUCKS (GOOD LUCK)