Are Traded Players "Lemons?" Part II

Phil Birnbaum
September 23, 2010
Update September 26, 2010

DRAFT – comments appreciated

In economics, the "Market for Lemons" represents the idea that, in a situation where the seller of a good has better information on its quality than the buyer, the market will be inefficiently biased towards lower-quality goods.

The reasoning goes like this. Suppose the value of the average car of a certain year and model is $5,000. If I own an example that has never needed a repair and has always worked perfectly, then maybe it's worth $6,000. But I can never get $6,000 for it, because, even though I know how good it is, there's no way for me to convince the buyer than I'm not lying.

On the other hand, suppose my car is a lemon. It's been in the repair shop countless times, and, although it's running fine right now, I'm sure it'll break again sometime soon. I'd be thrilled to sell it for $5,000, because it's really only worth $4,000.

So, what happens? The owners of the good cars hang on to them, and the owners of the lemon cars are enthusiastic about selling. So the quality of the cars offered for sale is now below average.

Buyers eventually realize that the cars offered are worse than average, and so they're now willing to pay only $4,500. That means that, now, sellers of average cars, which are worth $5,000, also become unwilling to sell. This drops the quality even further, and the price drops even further, and so on, until, in theory, there are no cars for sale at any price.

In practice, of course, it doesn't go that far. There are indeed ways for sellers to convince buyers that their cars are of higher quality – warranties, service records, and so on. Also, some owners of $6,000 cars will sometimes be willing to sell their car even though they can only get $5,000 for them.

But, still, theory indicates that the quality of used cars on the market should be somewhat lower than the quality of used cars in general.

Is that also true in the market for ballplayers? The situation is similar. The team "owning" the ballplayer has more information than the other team thinking about acquiring him in a trade. The current team might know things that aren't public knowledge – if the player is gaining weight, or if he's not keeping himself in the best of shape, or if he's just gone off the steroids that made him such a great hitter recently.

If that's the case, then we should find that players who are traded turn out to be less successful than expected.

A few years ago, I did a study to test that. I looked at players who had accumulated 1,000 runs created in their careers so far, and then were traded. I found that untraded, presumably "non-lemon" players, were much more likely to reach difficult career goals than traded players, by a ratio of 3 to 1.

Back then, the only tool I had to predict future performance was Bill James' "Favorite Toy". In the years since, various other projection methods have been developed, including Tom Tango's "Marcel" system. A player's Marcel projection is his weighted average of his previous three seasons, regressed to the mean, and park and era adjusted. Despite its apparent simplicity, it works fairly well. A detailed analysis of Marcel's strengths and weaknesses is beyond the scope of this paper. Actually, it's really not, but I'll save it for a future study, because that way I get twice as many publications and citations, which makes me more likely to get tenure. (Ha, ha! Academia joke.)

Recently, researcher Jeff Sackmann was kind enough to publish a listing of every player's Marcel projection since 1901. Armed with the Sackmann database, I am now able to expand the original study to see if the results still hold.

We'll start with the most obvious test of the lemons hypothesis: whether traded players beat their Marcels. I took all hitters from 1901 to 1975, and divided them into two groups: those who changed teams in the off-season following that year – the "traded" players -- and those who did not – the "not traded" players. (I stopped at 1975 to ensure that I was looking mostly at trades, and not at free-agent signings.)

Then, I pulled the players' next-year Marcel projections from the Sackmann database. To make sure I was looking only at established players, I eliminated all players whose projections were for fewer than 400 plate appearances.

If the Lemons hypothesis holds, you'd expect the non-traded players to have beat their Marcel projections, and the traded players to have fallen short.

Didn't happen. Here are the non-traded guys:

	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Actual	3354833	469674	935302	153295	38863	65090	424233	343366	298863	69926	23223	0.279	4.96
Projected	3597776	503430	1005885	163727	43668	68528	451972	354455	323786	77363	22165	0.280	4.95

Table 1: Players who weren't traded between seasons: actual performance vs. Marcel projection. (1901-1975, only players with minimum 400 PA projected)

Aside from the fact that the projections were a little too high in terms of playing time, the non-traded batters peformed almost exactly as expected. To show you that more easily, here's the same comparison, but normalized to 500 plate appearances:

Not Traded	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Actual	3354833	454	64	126	21	5	9	57	46	40	9	3	0.278	4.96
Projected	3597776	455	64	127	21	6	9	57	45	41	10	3	0.279	4.95

Table 2: Players who weren't traded between seasons: actual performance vs. Marcel projection. (1901-1975, only players with minimum 400 PA projected.) Normalized to 500 AB.

The first "AB" is the actual number of AB in the sample; the second AB is normalized to 500PA. The last column is "Runs Created per 27 Outs", a measure of what a team of nine identical batters would score per game.

So, as I was saying, the non-traded guys performed exactly as expected. Here are the traded guys:

Traded	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Actual	381026	455	61	123	20	5	8	55	45	41	10	3	0.270	4.57
Projected	465129	456	61	125	20	5	8	56	44	41	10	3	0.274	4.66

Table 3: Players who were traded (switched teams) between seasons: actual performance vs. Marcel projection. 1901-1975, only players with minimum 400 PA projected.

Now, we have something … at least a little bit of something. The biggest thing here is that the traded guys played significantly less than expected, by 18%. That doesn't necessarly mean there's a lemon effect – it could be that the reason these guys were traded is that they were mediocre in the first place, and the team was finally able to replace them with someone better. In that case, that their playing time dropped after the trade would be only natural.

I suspect that the difference in AB is real, and statistically significant (I suppose I should test that) – my point is that it might be something other than a lemon effect.

Other than playing time, there's not much difference in their performance. They gained a walk on their projection, but lost a couple of hits. Still, the difference in RC/G is almost certainly not statistically significant, and, regardless, we can't be sure Marcel is accurate enough that the difference is real and not just a shortcoming of the projection system. Overall, I'd say there's no evidence of a lemon effect so far.

How is it, then, that my previous study did find a lemon effct? Well, one thing is that the previous study used a different criterion for inclusion: it was restricted to batters who had 1,000 or more runs created in their careers before they got traded.

So let's try that. Here are the non-traded players:

Not Traded	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Actual	230758	443	72	132	22	5	13	68	57	36	9	3	0.298	6.34
Projected	252226	446	72	133	22	6	13	68	54	35	10	2	0.298	6.21

Table 4: Players who weren't traded between seasons: actual performance vs. Marcel projection. 1901-1975, only players with minimum 400 PA projected, only players with 1000 actual RC before that season

And the traded:

Traded	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Actual	37289	449	65	127	22	5	10	63	51	39	8	2	0.283	5.33
Projected	46839	451	66	128	22	5	10	64	49	37	9	2	0.284	5.39

Table 5: Players who were traded (switched teams) between seasons: actual performance vs. Marcel projection. 1901-1975, only players with minimum 400 PA projected, only players with 1000 actual RC before that season

A very, very slight bit of evidence for an effect, here: the non-traded guys beat their projections by 0.13, but the traded guys undershot by 0.06. The 0.19 difference is pretty small. However, the traded players did undershoot their playing time estimates more than the non-traded players did. So maybe the effect is mostly limited to playing time.

That gives us a hypothesis for why the previous study came up with such a large lemons effect. That study was based on long-term career Runs Created milestones. But achieving a difficult goal in a raw number requires significant amounts of playing time. Given that traded players lost 10 percent more plate appearances than non-traded players *in the very next season*, and given how the probability of hitting a far-away milestone is very sensitive to total playing time … well, it seems like we now have a possible explanation of how traded players' "Favorite Toy" chances dropped by two-thirds.

To confirm that, let's try to look at the entire career, instead of just the next season. That's not possible to do with just the Sackmann database, because Marcels provide projections only for a single year out. To look at how the rest of a career plays out, we'll need a different strategy.

Let's try a controlled experiment. For every traded player, we'll find a similar non-traded player, and compare how their careers played out.

How do we find similar players? By using "Similarity Scores," a method developed by Bill James and explained in his 1986 *Baseball Abstract*. It's an algorithm that takes two batting lines and assigns a number that represents how similar they are.

So here's what I did: for every traded player from 1901 to 1975 (minimum 400AB projection), I went through every non-traded player, and found the one who was the same age, played the same position, and whose Marcel projection was most similar.[1]

For instance, after the 1965 season, Ken Boyer was traded from the Cardinals to the Mets. I searched all non-traded 3B seasons by 35-year-olds in the Sackmann database to find the one whose season projection was most similar to Boyer's 1966 projection. That turned out to be Frank Malzone in 1965. Here are the two projections:

Projections	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
1966 Ken Boyer	508	71	137	21	4	16	81	55	73	3	4	0.270	5.02
1965 Frank Malzone	513	58	139	19	2	14	61	32	46	1	1	0.271	4.36

These are fairly similar, but not that great; the similarity score comes out to 965, which is fairly typical. But some of the pairings are much better. This one is 974:

Projections	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
1974 Willie Davis	547	73	157	24	7	14	69	28	55	16	5	0.287	4.99
1969 Felix Alou	575	72	168	27	4	14	50	35	54	8	7	0.292	5.11

And here's a mediocre one, a 915. There are only a few this bad in the sample.

Projections	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
1928 Rogers Hornsby	511	108	176	32	7	20	106	69	36	6	0	0.344	9.07
1935 Ch. Gehringer	547	103	178	39	6	12	100	71	28	7	5	0.325	7.38

While not every projection is perfect, the differences should all cancel out over the entire sample. And they do. Here are the aggregate totals for the traded player and the control group:

Projections	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Traded	465129	456	61	125	20	5	8	56	44	41	10	3	0.274	4.66
Control	459717	457	60	125	20	5	8	56	43	41	9	3	0.274	4.65

Table 6: Players who switched teams between seasons, and their control group: Marcel projections for how they would do in the subsequent single season. 1901-1975, only players with minimum 400 PA projected.

The traded group was projected to have about 2% more AB than the control group … but everything else is almost identical.

So, the next step was to look at the remainder of the careers for the two groups. If there is indeed a lemon effect, the traded players should have finished up significantly worse than the control group players.

[1] Actually, I tweaked Bill's algorithm a bit. Because the Sackmann database didn't project games played, I didn't include that. Instead, I tripled the adjustment for AB. In addition, I deducted one point for every two calendar years' difference between the traded player and the similar player, under the theory that the closer the players were in baseball history, the more similar they truly were.

They kind of did:

	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Traded	1573625	454	60	123	20	5	8	56	46	40	9	0	0.271	4.59
Control	1782315	454	61	125	21	5	8	58	46	40	8	0	0.275	4.79

Table 7: Players who switched teams between seasons, and their control group: Actual performance of both groups in the the subsequent single season. 1901-1975, only players with minimum 400 PA projected.

Again, there's a large effect in playing time, and a small effect in performance.

The traded players had 12 percent fewer at-bats; there were 1,040 players in each group, so the difference is almost exactly 200 AB per player per career. The traded players also created 0.12 fewer runs per game – that means that if the traded group had played as well as the control group, the difference would be about 9 career runs per player. Since there were about three full-time seasons per player, we're talking about three runs a year of "lemon" effect.

So, overall, the career lemon effect of being traded is two hundred fewer at bats, and nine runs.

Anyway, that's every traded player. Let's now cut to the important case, the one where we got the most serious results above: players who had at least 1,000 Runs Created before being traded.

I'll start by showing you the composite Marcel batting lines for the traded players and the controls, for the subsequent single season, just to show you how comparable they are:

Projections	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Traded	46839	451	66	128	22	5	10	64	49	37	9	2	0.284	5.39
Control	46303	453	66	129	22	5	10	61	47	36	9	3	0.285	5.30

Table 8: Players who switched teams between seasons, and their control group: Composite Marcel projections for both groups in the the subsequent single season. 1901-1975, only players with minimum 400 PA projected, only players with minimum 1000 RC before that season in the "traded" group.

Almost the same, right? So you'd expect that if you looked at the remainder of the careers of both groups, it would be fairly similar. That's especially true when you remember that the controls were also matched by age and position, and usually come from similar eras.

But, it turns out, there's a huge, huge difference in how their careers went from there:

Rest of career	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	AVG	RC/G
Traded	103242	450	63	125	22	4	10	64	50	39	8	0.278	5.12
Control	143367	449	66	131	23	5	10	64	51	33	8	0.292	5.62

Table 9: Players who switched teams between seasons, and their control group: Actual composite performance for both groups in the remainder of their careers. 1901-1975, only players with minimum 400 PA projected, only players with minimum 1000 RC before that season in the "traded" group.

Despite the fact that they were selected to be as close a match as possible to the traded group, the controls had almost 30 percent more bulk production than the traded group. In addition, they created an extra half run per game (25.5 hitless at-bats), which is about 7 runs per season, or almost 3/4 of a win per year.

There were 102 players in each group; dividing the above lines by 102 gives the average, rest-of-career production for a single member of each group:

Rest of career	AB	R	H	2B	3B	HR	RBI	BB	K	SB	AVG	RC/G
Traded	1012	143	282	48	10	23	143	112	87	18	0.278	5.12
Control	1406	206	410	71	15	31	201	159	103	24	0.292	5.62

Table 9: Players who switched teams between seasons, and their control group: Actual performance for average player in both groups for the remainder of his career. 1901-1975, only players with minimum 400 PA projected, only players with minimum 1000 RC before that season in the "traded" group.

So the traded group lost almost a third of their production over the remainder of their career. No wonder they failed to meet their "Favorite Toy" projections: if you need 500 runs and your expectation is 500 runs, you have a 50 percent chance of reaching your goal. But if you need 500 runs and your expectation is only 360 runs, the Favorite Toy gives you only a 22 percent chance.

These results are consistent with the original 2004 study. There, when non-traded players had a 51 percent chance of reaching their goal, traded players succeeded only 16 percent of the time (3 out of 19). That's in line with our rough 22 percent estimate here, which would be 4 out of 19.

When I did the original study, I was surprised that the effect was so large. I thought maybe it was just a fluke. This study seems to confirm that it's not. When a veteran player with over 1000 runs created is traded, it looks like you should immediately knock 30 percent off your estimate of his remaining career length, and you should expect his overall level of productivity during that time to be about 10 percent lower than you thought.

That's unexpected. But one aspect of the findings surprised me even more.

And that is: the effect seems to be delayed! For the first year after the player is traded, his first season with the new team, there's almost no effect. Here's the comparison:

1st year traded	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Traded	37289	449	65	127	22	5	10	63	51	39	8	2	0.283	5.33
Control	37022	451	65	130	22	5	10	62	49	33	8	3	0.288	5.46

Table 10: Players who switched teams between seasons, and their control group: Actual performance for average player in both groups in the subsequent single season. 1901-1975, only players with minimum 400 PA projected, only players with minimum 1000 RC before that season in the "traded" group.

In fact, the year after being traded, the players in the sample actually got a bit more playing time than the controls. Their performance was a little worse, but not much.

That means the entire decline in the traded players' careers started in the second year after the trade. Here's the "rest of career" stats for both groups, but starting one year after the trade:

	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	AVG	RC/G
Traded	**65953**	**451**	**63**	**124**	**21**	**4**	**10**	**64**	**49**	**39**	**8**	**0.275**	**4.94**
Control	**106345**	**449**	**66**	**131**	**23**	**5**	**10**	**65**	**51**	**33**	**8**	**0.292**	**5.66**

Table 11: Players who switched teams between seasons, and their control group: Actual performance for average player in both groups for the rest of their career starting one full season after the switch. 1901-1975, only players with minimum 400 PA projected, only players with minimum 1000 RC before that season in the "traded" group.

Huge difference.

So in the first year, nothing happened. But in the years after that, the non-traded group played 60 percent longer than the traded group, and had a significantly better level of productivity during that time.

Just to make this absolutely clear: suppose you have two players, both the same age and playing the same position, with roughly equal Marcel projections. If one of them was traded last year (and so just finished his first year with his new team), and the other wasn't … and if the traded player also had 1,000 runs created to his credit before being traded … then, you should expect the remainder of that player's career to be *almost 40 percent shorter* than the other player's, even though they're the same age, play the same position, *and have almost exactly the same Marcel projection for next year*.

That's a bit of a shock. I'll have to do some significance testing to see if it could be random … but I doubt it.

Another interesting thing happens when you limit the sample to the more modern era –1949-1975 instead of 1901-1975. There, the "non-effect" period seems to be two years instead of one. The first year after the trade, the control group had only two percent more AB than the traded group (13457 to 13257). The second year, the effect was still about two percent (9125 to 8940). Finally, in the third year, the divergence starts, and the control group's playing time advantage grew to 40 percent (6508 to 4645).

Recently, over at "The Book" blog, there was some speculation on whether there might be a lemon effect in regard to free agents. Are teams good at predicting what their current players will do next year (over and above naive predictors like Marcel)? If so, they are likely to re-sign the ones with "good" outlooks, and let the ones with "bad" outlooks sign with other teams?

We can do a quick check on this by running the same study for the free-agent era. Among players who change teams in the off-season, I'm unable to distinguish free agents from trades. But, since we're limiting the sample to players with 1,000 runs created, the majority are probably indeed free agent signings.

I ran the test for the seasons 1985 to 2000, which resulted in 58 players from each group. Keep in mind that even with the sample ending in 2000, there may still be careers that haven't ended yet, so the results aren't perfect.

This time, the results go the other way: the free agents did better than the control group for the rest of their careers:

Career	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	AVG	RC/G
Traded	76317	440	65	120	21	2	15	63	60	72	10	0.273	5.45
Control	69913	447	62	120	22	2	17	66	53	74	5	0.268	5.24

Table 12: Players who switched teams between seasons and their control group: actual performance for both groups for the remainder of their careers. 1985-2000, only players with minimum 400 PA projected, only players with minimum 1000 RC before that season in the "traded" group.

The career length difference is about 10 percent.

The free agents also did better in their first year with their new team, even though they had a bit less playing time:

Next Season	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Traded	22297	440	65	121	21	2	15	65	60	73	9	3	0.275	5.51
Control	24630	449	60	120	21	2	17	67	51	72	6	3	0.267	5.17

Table 13: Players who switched teams between seasons and their control group: actual performance for both groups for the single subsequent season. 1985-2000, only players with minimum 400 PA projected, only players with minimum 1000 RC before that season in the "traded" group.

So it's probably fair to assume that the effect has something to do with being traded, and not with being let go as a free agent.

So what's going on? Are any of these results good evidence that there's a real lemon effect? Not necessarily.

Our control group was composed of players who had similar Marcels to the traded group. But a similar Marcel this year doesn't necessarily also mean a similar Marcel next year, or the year after that.

Suppose Marcel predicts that a player whose last three seasons are (120, 80, 80) has a projection next year of 80 – the same as a player who went (85, 90, 85). And that might be right. But what happens the season after that? It could be that the first player's prognosis should be worse – he's on the decline, while the second player is holding his ground.

Suppose the next season, both players score 80. Now the first player is (80, 80, 80), while the second player is (90, 85, 80). Marcel correctly predicts that the second player will do better this year, and then next year, and so on.

Put more simply: suppose one player's recent performance is (9, 8, 7). And another player's is (6, 6, 6). Just looking at the trend, you'd expect both players to be a "6" next year. But they year after that, you'd expect the first player to be a "5", and the second player to still be a "6".

In a nutshell: just because two players are expected to have the same performance this year, that doesn't mean they're expected to have the same performance the year after that.

That could be what's happening here. The traded players could be on the decline, while the control group players aren't. That would explain the results perfectly: similar performance between the two groups the first year, with the traded group declining more after that.

And if you look at the trend, that's exactly what you see. Here are the RC27s for both groups for the four years before the trade, and the three years after:

```
Traded         6.72 6.23 6.13 5.34 trade 5.33 5.30 5.30
Non-Traded     6.03 5.75 5.63 5.69 trade 5.40 5.62 5.67
```

What we notice is that the traded players were probably significantly better than the non-traded players, but have declined faster, to the point where they're about the same by the time of the trade. Indeed, the traded players had an unusually large decline in the year before the trade – from 6.13 to 5.34. That's probably a good part of the reason their teams decided to trade them.

So there may not be a lemons effect here at all. It might just be that the Marcels don't capture something important about a player – the trajectory of his performance. It could be that everyone and his GM knew that the traded player is declining, and that's why he was traded. In that case, the trade doesn't actually provide new evidence that the player will decline: we already knew the player will decline, and the trade is simply because some other team values the declining player more.

How can we check that? We need to use something other than Marcels to form the control group.

So, for each player, I found the player who was the same age, playing the same position, who best matched his individual batting lines for the past two years. I calculated three similarity scores -- two years ago, one year ago, and the total – and averaged them out with weights of 3-2-1.

Actually, I made one other change – I relaxed the age requirement a bit in order to get a better match in the batting lines: I allowed the comparison player to be one year younger, and up to two years older, than the traded player. (The asymmetry is because there are more younger players than older players, and I wanted to keep the average age roughly the same.)

First, I'll show you that the tradeds are still roughly comparable to the controls. Then, I'll give you the results.

Here's how the two groups compared two years before the trade:

Two seasons ago	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Traded	49795	448	71	134	22	6	11	69	52	36	9	2	0.299	6.13
Control	50012	454	68	136	23	6	11	67	46	35	8	3	0.300	5.89

And here's one year before the trade:

One season ago	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Traded	47771	449	65	127	22	5	10	63	51	37	8	2	0.283	5.34
Control	48483	453	65	130	22	5	9	62	47	33	8	3	0.287	5.28

And here are the Marcel expectations for the first season after the trade:

Marcels	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Traded	46839	451	66	128	22	5	10	64	49	37	9	2	0.284	5.39
Control	47358	455	64	129	21	5	10	62	45	36	8	3	0.284	5.16

So, the tradeds performed a tiny bit better in the past, and are expected to perform a tiny bit better in the coming season. So how *did* they do compared to their Marcels? A little worse:

One year after	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Traded	37289	449	65	127	22	5	10	63	51	39	8	2	0.283	5.33
Control	39951	452	64	128	22	5	10	63	48	36	8	3	0.283	5.23

They also started to drop a bit in playing time. Last year, the traded players had only 700 fewer AB than the control group, but, this year, they had 2700 AB fewer. That suggests they're still declining faster than the control group. In terms of rate of performance, the drop is probably significantly worse than the 5.33 suggests -- you can't just go by the raw batting lines, because those results are affected by survivor bias. That is, the missing 2,000 AB probably came off the worst players. If you actually had those guys play their additional AB, the 5.33 would have dropped a fair bit.[2]

Now, the results we've been waiting for: the stats for the rest of their careers (including the year above):

Career post-trade	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	AVG	RC/G
Traded	103242	450	63	125	22	4	10	64	50	39	8	0.278	5.12
Control	118658	449	65	126	21	5	10	64	51	36	7	0.281	5.23

[2] This is the same logic that explains why 35-year-old hitters don't have stats that much worse than 26-year-old hitters – the older guys are only the ones good enough to still be in the league at 35, and have probably dropped substantially from the days when they were 26.

As you might expect, it appears the tradeds continued to drop faster. Overall, they wound up below the control group in performance, and got 13 percent fewer at-bats.

But I don't think there's a lemon effect here, because, even though we made efforts to create tight controls, the traded group was still declining faster than the non-traded group. And that was public information.

Indeed, the traded group was declining even faster than the above results suggest, as we'll see when we look at the relative performance three years before the trade. We haven't looked at that aspect yet. And, since our process controlled only for one year before and two years before, there is no reason that the three-years-before performances have to be close to each other. And they're not that close:

Three years ago	AB	AB	R	H	2B	3B	HR	RBI	BB	K	SB	CS	AVG	RC/G
Traded	51654	450	73	136	23	6	12	72	50	35	10	3	0.302	6.23
Control	47547	453	68	132	22	6	11	65	47	35	9	3	0.291	5.56

In the three seasons leading up to the trade, the traded guys were 67 points better than the controls, then 24 points better, then only 6 points better. It's safe to say, I think, that it was probably obvious that the traded players would continue to decline faster than their control group.

I think, after all this, we have to conclude that there's no real evidence of a lemons effect. The only time that we saw any evidence that *might* have been something was with this group of players, the 1000 RC guys. But, as it turned out, we found another reason those guys might have done worse than their controls – because they weren't all that comparable to the controls in the first place. The traded players were on the decline, while the controls weren't, and that seems to have explained the results just fine without the need to resort to other hypotheses.

We could probably try redoing the study with a better algorithm to match the traded players to controls. Perhaps we could use three or four years instead of two. However, the requirement that we come close in matching age and position means that there's a limited sample of players in the pool in the first place, and, the more seasons and characteristics we insist on matching, the worse the fit is going to be. It might still be worth a try, though.

In any case, we still learn a few things from this study. One important thing, which I haven't seen explicitly said before, is that there might such a thing as "momentum" in player declines. The traded batters dropped faster than their controls in the period leading up to the trade, and they continued to decline faster than their controls later. That means that estimators like Marcel might be too simple. They consider only the weighted average of the past three years, and not the pattern within those years. A player whose last three years are (6, 5, 4) might be projected by Marcel to be the same as a player who's gone (4.7, 4.7, 4.7) – but the guy on the decline should probably be projected lower.

If this is correct, it means that the "Favorite Toy", which, like Marcel is based on a weighted average performance of the last three years, is not at all accurate for these kinds of players. Recall that in the original study, traded players with 1,000 RC achieved a certain goal only 17 percent of the time,

whereas non-traded players with 1,000 RC had a 51 percent success rate. If that's not a lemons effect, and is simply a "declining players" effect, we have to concluded that the Toy is simply not accurate enough to be used late in a batter's career.

But … this "momentum of decline" hypothesis is not proven at all. It could be false, or it could be true but not account for 100% of the effect we saw. It could be that the decline in the traded players was only partly due to momentum, and partly due to the lemons effect. It could be that if we found non-traded players with simlar declines, they would have leveled off, thus proving that the traded players are indeed different in some way from non-traded players.

With such a small sample of players, it would be hard to tease the two effects apart. And it seems to me that the "decline" hypothesis is more plausible than the lemons hypothesis, especially in light of the fact that the effect seems occur only for older players.

The way forward, perhaps, is to work on a projection system that we can show unbiased for the kinds of older players we're interested in, which takes rate of decline into account. Marcels, which only project one year, and which may be biased too high for older, declining batters, aren't the tool we need.