J Prod Anal
DOI 10.1007/s11123-009-0152-8

Pay, productivity and aging in Major League Baseball

Jahn K. Hakes · Chad Turner

Abstract Using panels of player pay and performance from Major League Baseball (MLB), we examine trends in player productivity and salaries as players age. Pooling players of all ability levels leads to a systematic bias in regression coefficients. After addressing this problem by dividing players into talent quintiles, we find that the best players peak about 2 years later than marginal players, and development and depreciation of performance appear to be more pronounced for players with the highest ability levels. Within-career variation, however, is less pronounced than between-player variation, and the performance level of players within a given quintile will typically remain lower than the talent level for rookies in the next higher quintile. We also find preliminary evidence that free agents are paid proportionately to their production at *all* ability levels, whereas young players' salaries are suppressed by similar amounts.

Keywords Major League Baseball (MLB) · Career dynamics · Player salaries and performance · Quintile analysis

JEL Classification J3 · J24 · J42 · L83

J. K. Hakes
Department of Economics and Management, Albion College, Albion, MI 49224, USA
e-mail: jhakes@albion.edu

C. Turner (✉)
Department of Finance and Economics, Nicholls State University, Thibodaux, LA 70310, USA
e-mail: chad.turner@nicholls.edu

1 Introduction

The correspondence between salaries and worker productivity is a central issue in labor economics. Due to the abundance of performance and salary data for athletes in professional team sports, a great deal of empirical work has studied salary patterns in baseball and other sports, with sub-fields examining salary discrimination and the salary effects of arbitration and free agency, among other topics.[1] Although this path has been trod frequently, we believe there is still room in the literature for exploration yielding new insights.

Our contribution is to use quintile analysis to examine how productivity patterns vary between cohorts of players with similar ability and see how well the labor market accommodates ability variation in setting salaries. To illustrate the extent of the bias produced by traditional ordinary least squares (OLS), we run regressions estimating separate salary and performance paths for each talent quintile. This technique will reduce the tendency of pooled regressions of salary or productivity on experience to yield a "flatter" time profile than is actually the case.[2]

[1] In baseball, the majority of this research stems from the seminal work of Scully (1974). Kahn (1991) summarized the early racial discrimination literature, which has since expanded to include analyses of coaching discrimination in MLB (Singell 1991), the NFL (Madden 2004), and the NBA (Humphries 2000; Kahn 2006). Marburger (1994, 2004) looks at final offer arbitration in MLB, while Zimbalist (1992), Kahn (1993), Vrooman (1996) and Miller (2000), among many others, look at the salary effects of free agency in baseball.

[2] We note this technique is not the same as quantile regression, although both terms have a common element of placing greater analytical weight on particular centiles of a distribution. While it would be natural to consider using quantile regression, utilizing this technique would result in an observational bias in the same manner as

Published online: 16 September 2009

The primary advantage to studying productivity in Major League Baseball (MLB) is the abundance of quantitative data on multiple components of individual performances. Sabermetricians, statistically minded members of the Society for American Baseball Research (SABR), have produced a great number of performance metrics derived from these components in an attempt to quantify production with a single metric. The performance measure we utilize is on-base percentage plus slugging average (OPS). While some of the sabermetric alternatives are preferable to OPS in specialized scenarios, the more widespread acceptance of OPS is worth the potential very small cost in efficiency for our purposes.

In Sect. 2, we discuss the econometric problems involved in career path estimations, and how quintile analysis mitigates some of those problems. Section 3 discusses the data set used and our empirical method. Section 4 looks at player productivity over the span of a career and Sect. 5 presents career salary paths. Section 6 offers preliminary findings on the crude relationship between pay and performance in each league, while Sect. 7 summarizes our findings and details our plans for further research.

2 Variation of individual characteristics over time

2.1 Career paths in productivity

OLS regression is a tool so familiar to empirical economists that its use to study relationships such as that between player salaries and years of experience is virtually reflexive. We do not wish to combat the intuition that regression analysis is an appropriate tool here—especially as we will use that technique for our own modeling—but rather to note an observational deficiency in the data, that in the absence of correction, would lead to biased estimators.

Before we proceed, it is important to be clear about the terminology we utilize, as it is common to use the terms ability, talent, performance, and productivity interchangeably. We wish to determine how performance and pay across players of different talent levels, and therefore, some care is required in distinguishing between talent or ability on the one hand and performance and productivity on the other. We use the terms ability or talent to refer to the player's innate ability level, which we will later measure, with slight modification, as the player's peak performance measure. We model ability (talent) as being time invariant for each player, but of course varying across players. We use the term production or performance to refer to the observed measure of productivity—in economic terms the marginal physical product. For a player of a given ability level, production will vary in a predictable fashion as the player ages.

We can then model players in a top-tier professional sports league as individuals in the extreme right-hand tail of the ability distribution for playing that sport. To obtain a roster spot, a player must have a productivity level above "replacement level." A replacement level player has ability only marginally above that of the top player not contracted to play in the league, and typically is a benchwarmer or part-time fill-in at the top level, or a top player or prospect in a minor or secondary league.[3] Typically, pay and productivity are only recorded for those players who have been selected into the top tier league.

In the absence of entry barriers, a young player would be able to hold a spot on a professional roster once his performance level reached replacement level, and as he further developed he would continue to remain employed until diminishing skills (due to injuries or age) rendered him once again below replacement level. It has been solidly established that baseball hitting performance climbs until peaking at about age 27, and slowly trails off thereafter.[4] Adapting that methodology to other sports has led to similar findings, with minor variations of peak age, but in each case describing an age function that resembles an inverted-"U" and which is commonly estimated as a quadratic function. In Eq. 1, this would be noted as $\beta_1 > 0$, and $\beta_2 < 0$, where α is the vertical intercept, and ε_i is a stochastic error term for player i.

Footnote 2 continued
would occur from using OLS. At very young (and very old ages), only the highest ability players will be in the sample, and thus, say the 50th percentile of the conditional distribution of performance at, say age 19 will be estimated using information from only these high ability players (whose median would likely be a q5 player). However, the 50th percentile at, say age 28, will be estimated utilizing information from players of different abilities (whose median would likely be a q3 player). The result is a dampening of the distinctions between groups of players with different abilities, which is precisely what we attempt to avoid by dividing players into quintiles. For an early example that uses quantile regression to examine salary discrimination in the NBA, see Hamilton (1997). More recently, Berri and Simmons (2009) and Vincent and Eastman (2009) utilize quantile regression to examine discrimination in the NFL and NHL, respectively. To the extent that entry barriers in college football make entry age into the NFL more homogeneous, the compositional effect will be reduced.

[3] Many players in baseball make frequent journeys back and forth between the top minor league level (AAA) and "The Show," earning them the label of "4-A" players.

[4] See James (1982), Krohn (1983), Albert (1999), Schulz et al. (1994), and Fair (2008) for estimation of production peaks and Fort (1992) and Horowitz and Zappe (1998) for salary estimation models.

$$\text{Productivity}_i = \alpha + \beta_1 \text{age}_i + \beta_2 \text{age}_i^2 + \varepsilon_i \tag{1}$$

We can speak of a player's career path in productivity, or "productivity path" as the locus of Productivity_i level over the relevant range of ages, and if the correct functional form for Eq. 1 is indeed quadratic, we can then discuss the shape of productivity paths in terms of the parameters α, β_1, and β_2.

To date, we know of no studies searching for or finding systematic differences in the productivity paths of star players (the highest peak ability), journeyman players (of middling ability), and marginal players (only slightly above replacement level). The typical practice is to pool all players when estimating Eq. 1, although Krautman (1993) did acknowledge the estimation problems posed by inherently different ability levels. The one notable exception is Schulz et al. (1994), who found the productivity of their group of elite players declined more slowly after the peak age than did those of non-elite players. Our paper extends the approach of Schulz et al. by using a more detailed segmentation of the player population, applying the technique to salaries, and considering the economic implications of the resulting estimates.

Pooled regressions implicitly assume that the only structural difference between a star player and a marginal player is a higher value of the intercept term, α, for the star. That assumption leads to the implications that (a) player productivity improves or declines by the same amount at a given age, regardless of ability cohort, (b) players of all ability levels hit their "peak" at the same age, (c) star players reach replacement level at an earlier age than marginal players, and (d) that star players will remain above replacement level longer than marginal players, and will be forced out at later ages.

Pooled empirical estimation not only leaves testable implications unexamined, but results in biased coefficients as well. The mechanism for selecting observations is based on whether performance is higher than the replacement level threshold. While all professional players will be in the sample at their peaks, yielding an average productivity level at the peak age that represents the central tendency of all players, only the stars' productivity will be sufficiently high for their statistics to be observed at younger and older ages, leading to a positive observational bias that increases in magnitude as the distance (in years) from the peak age increases. This theoretical effect is illustrated in Fig. 1, and shows that an estimated regression line from pooled player data (the dashed line) will tend to underestimate the magnitudes of the β_1 and β_2 terms in Eq. 1.

On the left-hand side of the relevant range of ages, this effect might be somewhat lessened in some leagues (particularly the NFL and NBA) by entry barriers ostensibly designed to keep players out of the league until they have attended college. Though the magnitude of the mismeasurement would be reduced as a result, the bias from aging players in the right tail would persist. If the bias is originating mostly from the right tail of the age distribution due to left-tail truncation, we would also expect estimated peak ages to be biased upward.

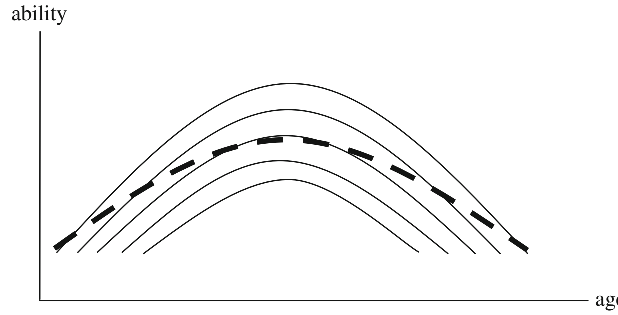

Fig. 1 Theoretical illustration of observational bias

The approach we use to reduce this bias is to divide our sample of players into ability cohorts that are more homogeneous than the pooled sample. As the variation in ability is reduced within each cohort, the resulting estimates are subject to less bias caused by differences in the age at which the players rose above or fell below replacement level. The estimated β_1 and β_2 terms should then be larger in magnitude, with estimates that consistently approach the true parameter value as the number of equally sized cohorts approaches infinity.[5]

2.2 Career paths in salary

Estimation of player pay has followed a substantially different methodology than that used for player performance, with most models estimating salaries directly as a function of productivity (or ability).[6] Our purpose here, however, is to establish time paths for individual player salaries as a function of experience ($\exp_i$). Therefore, we will formulate an empirical equation with a structure similar to the productivity equation above.

$$\ln(\text{salary})_i = \alpha + \beta_1 \exp_i + \beta_2 \exp_i^2 + \gamma' \mathbf{Pos_i} + \varepsilon_i \tag{2}$$

Equation 2 has a right hand side analogous to the productivity model, although it also includes a vector of control variables indicating player i's defensive position

[5] As we note above, segmenting the sample reduces but does not eliminate the bias associated with players of different ability levels. Further increasing the number of groups would reduce the bias further, but would reduce the sample size within each cohort and the number of observations of performance for each age level. We choose five groups somewhat arbitrarily bearing sample size considerations in mind.

[6] As we are focusing upon improved measurement of the individual components, the pay to performance ratios we construct will be somewhat crude. However, we will compare our broad-brush results to those from adaptations of the canonical Scully (1974) model in Sect. 6.

and it uses the natural logarithm of salary as the dependent variable to preserve normality of the residual terms. In comparison with traditional salary models following Scully's (1974) example, this model is extremely sparse. It lacks not only the normal control variables for negotiating freedom, market size, player awards, etc., but also the variable which typically is the primary regressor—productivity.[7]

The use of experience rather than age is necessitated by the nature of baseball's collective bargaining agreements. Aside from productivity, the most important predictor of player salary is freedom to contract, which has been defined in the league agreements as a function of playing experience. Consequently, our salary fits are far more efficient using experience.

If we were to estimate Eq. 2 using a pool of all players, it would be subject to the observational bias we described in Sect. 2.1. While we begin our analysis by estimating fits for productivity using age in order to illustrate our premise of replacement level observational bias, we subsequently estimate productivity equations using the experience measure in order to allow subsequent comparisons between pay and performance using a common time dimension.[8, 9]

Prior to estimating Eq. 2, however, we must ascertain whether the semilog quadratic model is appropriate. Two sources of potential difficulty are readily apparent. First, on the left side of the age range, reserve clauses and other restrictions on negotiating freedom for inexperienced players would be expected to lead to depressed salaries during the early stages of player careers. Indeed, measuring the extent of salary suppression is one of the main objectives of the pay-and-performance literature. Separate regressions for free agents and players under restricted bargaining may present significant improvements in the accuracy of estimates.[10]

Second, whereas player performance has been empirically determined to fade as a player ages past his peak, empirical evidence of player salaries falling as they age has not been as well-documented. One exception to this has been Horowitz and Zappe (1998), who concluded both that "once the average player has put in his 9 years… eroding skills result in lower pay," and that the effect was much smaller for former-star players.[11] Upon observing the data we will assess the appropriateness of the functional form of Eq. 2.

3 Data

The Baseball Archive database, version 5.3, edited by Sean Lahman contains season-by-season data on player performance, salaries, and many other variables that would serve as useful controls in a structural analysis.[12] As salary data is only available for the years 1985–2005, we select that as our sample period.

In order to compare productivity across a wide range of positions, we focus on hitting productivity, and have

[7] Our omission of the traditional set of controls from the Scully (1974) style models is intentional. We seek to look past the structural sources of pay variation and instead simply *describe* the pattern of pay variation for player groups similar to those in Eq. 1 and not, at this stage, to identify causative factors. Any resulting "omitted variable bias" is similarly intentional, as our objective is to compare and contrast the dynamics for performance and for pay, to confirm the general similarity, and to note whether there are systematic and large magnitude differences at particular ages.

The remaining distinctions we have added are to control for different player "markets". The division of players into ability quintiles allows a crude control for ability, as a star quality player and a journeyman are not close substitutes. We include indicator variables for position both because there are separate markets for shortstops and first basemen, for example, between which there is limited ability to substitute, and also as a concession to identifying player MPP. Players with similar offensive productivity at more challenging defensive positions are demonstrating ability along another dimension and should merit a salary premium.

[8] As could be predicted, the effect of the switch from age to experience in productivity Eq. 1 is that higher-ability players peak later in their careers, as they debut at younger ages, allowing more playing time prior to their peak age.

[9] As noted above, we would prefer to estimate career paths of performance relative to age, and ultimately wish to compare the evolution of pay and performance across a common metric. We illustrate below that either age or experience can explain essentially the same fraction of variation in performance paths; a switch from using age to experience reduces R^2 by an average of 0.0002 across the five models (where each quintile is modeled separately). However, when we later attempt to predict salary paths, a switch from using age to experience increases R^2 by an average of over 0.15 across the five models. In using experience, we accept the reduction in ease of interpretation in performance in exchange for the enhanced efficiency of the salary estimations.

[10] Fort (1992) estimated separate parabolic arcs for salary trends of players of below mean and above mean age. Although our differing arcs "hinge" at the point of free agent eligibility rather than mean age, the econometric objective is the same.

[11] This may be because star players refuse to accept salary cuts as their contracts come up for renewal after a performance decline, leading to observational selection with a short lag, or that perhaps the marginal revenue product (MRP) of star players—the theoretical source of employer willingness-to-pay for the player's services—does not decline as ability does. The latter possibility allows for veteran fan favorites and stars that either retain their popularity and drawing power through the end of their careers, or watch their fame decline more slowly than their fading skills, as was the case with the end-of-career versions of Cal Ripken, Jr. and Willie Mays. Although the discussion is beyond the scope of this paper, there is also a large literature in personnel economics that studies reasons why salaries may vary from MRP under competition. Common theories of this sort include efficiency wages to combat shirking and incentives for career contracting.

[12] The Baseball Archive database is available at http://www.baseball1.com.

eliminated pitchers from our sample.[13] Our preferred measure of batting performance is OPS (on base percentage plus slugging), which has been shown to be both simple to calculate and an accurate predictor of team output (wins). As OPS measures production per unit of playing time, whereas the chief sabermetric alternative, runs created, is a count statistic of production that increases with playing time, the use of OPS provides conservative estimates of within-career variation.[14]

Although performance as measured by OPS is not subject to general price inflation or reserve clauses as salaries are, league-wide OPS figures do vary over the period of our sample. The annual league-wide OPS ranges from 0.707 in 1988 to 0.796 in 2000, so that indexing by annual average is necessary. Furthermore, there are large differences in OPS across defensive positions, as the average OPS for first basemen is more than 125 points higher than that for shortstops (0.827 and 0.700). Failure to account for the positional differences would mean that shortstops and catchers would be significantly over-represented in the lowest ability quintile, while first basemen would be under-represented. We have therefore indexed OPS to correct for between-year and between-position variations.

We assign players to talent quintiles according to their peak ability. A player's peak ability is measured by the indexed OPS level of his third-best season of more than 130 at bats, including all seasons from 1985–2005, plus earlier seasons of players who were active prior to 1985.[15] For our empirical results, the sample consists of all player–seasons where we have both salary data and performance over at least 130 at bats for those position players (non-pitchers) with an established peak ability measure.[16] Using the reference levels for each player's peak ability, we established cut lines between players so equal numbers of player–seasons were represented in each of five quintiles.[17] Table 1 reports means of all measures used.

Table 1 Summary statistics for analysis, overall and by talent quintile, 1985–2005

Variable	q1	q2	q3	q4	q5	Pooled
OPS	0.666	0.708	0.737	0.763	0.846	0.744
Indexed OPS	0.890	0.948	0.990	1.027	1.123	0.996
Age	28.2	29.0	29.5	30.1	30.2	29.4
Experience	4.85	5.97	6.62	7.69	8.56	6.74
1B d.v.	0.104	0.097	0.134	0.088	0.170	0.119
2B d.v.	0.113	0.136	0.134	0.132	0.091	0.121
3B d.v.	0.110	0.117	0.127	0.149	0.113	0.123
Catcher d.v.	0.148	0.116	0.163	0.172	0.131	0.146
Outfield d.v.	0.400	0.401	0.313	0.388	0.401	0.381
Shortstop d.v.	0.125	0.133	0.130	0.071	0.093	0.110
Year	1995.4	1995.2	1994.8	1993.6	1994.6	1994.7
Plate appearances	325.1	400.5	434.2	477.7	534.8	434.5
At bats	292.4	359.2	387.3	423.8	464.8	385.5
Runs created	37.1	51.4	60.7	71.1	98.0	63.7
On-base pct.	0.306	0.322	0.332	0.337	0.365	0.332
Slugging pct.	0.359	0.386	0.406	0.426	0.481	0.412
Salary (millions)	0.57	1.11	1.57	1.93	3.80	1.79
ln(salary)	12.75	13.29	13.57	13.84	14.44	13.58
Adj. salary (lnnsal)	0.622	0.742	0.813	0.899	1.007	0.817
Ratio of lnnsal/nOPS	0.713	0.793	0.830	0.884	0.903	0.825
N (player–seasons)	1204	1222	1204	1211	1209	6050
N (players)	285	184	161	154	126	910

Note: Quintile sizes are not exactly equal because the career length of marginal players might force some player–seasons across percentile boundaries

[13] Defensive ability has proven both difficult to measure and, consequently, difficult to establish as a significant predictor of salaries except (weakly) through simple binary variables for defensive position.

[14] James (1988) concludes there is very little evidence of team complementarities in terms of batter "protection" and line-up effects. While batting order has a marginal effect on count statistics such as RBIs and runs scored (and this is why these statistics are poor choices to measure hitting performance), teammates and lineups do not statistically affect a player's underlying hitting productivity, which is what we use as a performance measure (OPS).

[15] Indexed OPS controls for systematic changes in the league-wide OPS level at each position across seasons. The use of the third-best season is somewhat arbitrary, but serves three purposes. First, it avoids the overestimation of ability of players who did extremely well in an injury-shortened season or a partial-season "cup of coffee" call up from the minors. Second, the cutoff removes very marginal players with fewer than three qualifying seasons from further consideration in the sample. Third, it reduces the likelihood that a player with a positive outlier season is assigned to the wrong ability quintile. To further limit the sample, players younger than age 29 in 2005 are removed due to the possibility that they might not have peaked yet, as their inclusion might taint the cohorts. Our results our not sensitive to this assignment mechanism, and using the 2nd or the 4th season to identify peak ability yields very similar results.

[16] The seemingly arbitrary threshold of 130 at bats is chosen based on the rules establishing "rookie status" by MLB (for the purpose of awarding the Rookie of the Year award). A player who has logged fewer than 130 at bats is still considered a rookie in the following season.

[17] The cutoff levels of indexed OPS are 0.958, 1.023, 1.073, and 1.153. The indexed values are relative to averages which condition for year and defensive position. The use of player-seasons as the unit of measure, coupled with the longer career lengths of high-ability players, means that there are fewer players in the top (5th) quintile than in lower ones. There are 126 players in 5th quintile and 285 in the 1st quintile. There are slightly more than 1200 player-seasons in each quintile.

Table 2 Estimated indexed OPS (iOPS), by age and ability quintile, 1985–2005

Panel A: Regression diagnostics

Coefficient	Model 1	Model 2	Model 3
Intercept	0.810** (0.079)	0.267** (0.062)	0.626** (0.229)
Age	0.015** (0.005)	0.047** (0.004)	0.022 (0.016)
Age^2	−0.0003** (0.0001)	−0.0009** (0.0001)	−0.0004 (0.0003)
Diff. intercepts	No	Yes	Yes
Diff. age coefficients	No	No	Yes
Diff. age^2 coefficients	No	No	Yes
Observations	6050	6050	6050
R^2	0.002	0.400	0.401
F-statistic (Chow)		999.9**	1.5

Panel B: Fitted equations and peaks from Model 3, by quintile

Quintile	Intercept	β_1 (age)	β_2 (age^2)	Peak age	Peak iOPS
1	0.626	0.022	−0.0004	25.6	0.906
2	0.437	0.040	−0.0007	26.8	0.969
3	0.512	0.039	−0.0007	26.1	1.017
4	0.456	0.044	−0.0008	27.5	1.056
5	0.311	0.060	−0.0011	28.2	1.163
Pooled	0.810	0.015	−0.0003	28.3	1.024

* Significant at 90% confidence level

** Significant at 99% confidence level

Standard errors in parentheses

4 Career productivity paths

4.1 Productivity paths with respect to age

We will estimate player development and deterioration alternately by age and by years of experience. Conditioning with respect to age is the intuitive way to analyze the validity of the four implicitly assumed characteristics of career performance paths as presented in Sect. 2.1 and to observe selection effects early and late in careers. However, changes with respect to experience are necessary to make comparisons to salary.

Table 2 shows the parameters that result from alternative specifications of Eq. 1. For all models reported, each player–season is weighted by plate appearances to avoid bias caused by part-time players, especially in the extreme tails of the age distribution. Panel A of Table 2 reports the results from three model specifications that allow testing of Implication (a). Model 1 naively pools all player seasons without any controls for ability quintile. Model 2 adds differential intercept terms to the Model 1 configuration. The associated Chow F-statistic shows that the improved explanatory power of the model, from 0.2 to 40.0 percent of the variation in ability, is statistically significant. Model 3 adds differential slope and quadratic terms to Model 2. The F-statistic of 1.5 confirms that the very small increase in goodness-of-fit is not statistically significant.

The rejection of the null hypothesis of identical intercept terms in Model 2 confirms that pooling of the ability quintiles is inappropriate, even though the evidence is not strong enough to disprove Implication (a)—that players develop at the same rate. Even so, the fitted equations reported in Panel B of Table 2 for each ability quintile show a clear trend where each successive quintile has a higher slope at younger ages, and a larger magnitude of the quadratic term. The additional curvature means that the best baseball players will exhibit more within-career variation than marginal players, who never rise far above replacement level and lose their jobs when they return to that level.

The Peak age column in Panel B of Table 2 shows a general trend whereby the relatively lower ability MLB players in the bottom three quintiles begin to fade sooner. Whereas players in quintiles 1 through 3 peak at ages 25.6 to 26.8, players in quintile 4 peak at 27.5 years of age, and players in quintile 5 peak at age 28.2. The fact that true superstars peak somewhat later than journeymen do is important to know for those attempting to forecast the value of mid-career players on the free-agent market, especially when careers last an average of less than 7 years.[18] The result represents a refutation of Implication (b).

The final result we show in Table 2 returns to the concept of observational bias. The Chow F-statistic in Model 2 confirmed that pooling of all five ability quintiles was inappropriate. In Sect. 2.1, we predicted that the effect of the pooling was to "flatten" the estimated productivity paths. The bottom row of Panel B in Table 2 shows the estimated pooled regression alongside the five fitted equations for each quintile. The parameter estimates on the Age and Age^2 variables in both leagues have lower magnitudes in the pooled regression than for *any* individual quintile in MLB. In addition to this flattening, we see that the effect of the biased slope coefficients is that the estimated peak age is biased upward by 1.5 years above the median peak age of the constituent quintiles, and above the point estimate for any one of them.

Figure 2 illustrates the fitted productivity paths for each quintile and the flatter regression line for the pooled sample, as well as the contingent mean values of indexed OPS by age. Each fitted regression line is discontinued at the point when there are fewer than 10 players remaining in the subsample at that age level. The entry ages for players in all quintiles are similar, but the high level of ability of upper quintile players as they debut lends support to Implication (c), that star players reach replacement level while younger.[19]

It can also be seen in Fig. 2 that the players with lower peak ability retire at much younger ages, so that the remaining players force the conditional mean OPS upward for the upper range of ages. That the upward observational bias is mostly occurring in the right tail of the age distribution is consistent with the overestimate of peak age in Table 2. Even though Fig. 2 shows that higher ability

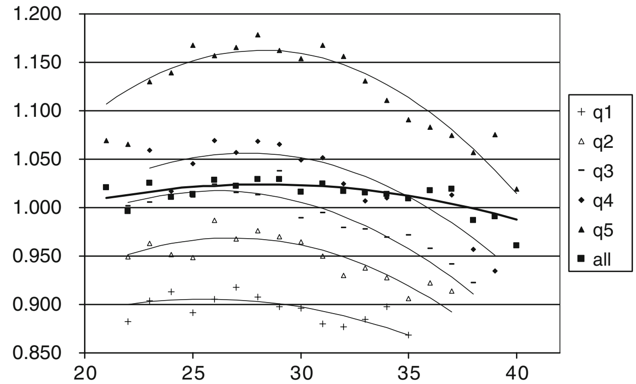

Fig. 2 Estimated indexed OPS (iOPS) by age and ability quintile, 1985–2005. Age (in years) is on horizontal axis; OPS indexed by year and position is on vertical axis. Estimated data points only shown for age-quintile combinations where there are at least ten observations. The corresponding estimated equations are shown in Table 2

players retire before they return to replacement level, the production paths extended enough years to verify Implication (d), that star players remain above replacement level longer than marginal players.

4.2 Productivity paths with respect to experience

To anticipate the comparison of productivity to salaries in Sect. 5, we fit Eq. 1 to the performance data using experience rather than age. Although we saw evidence of sample truncation in the left end of the age distribution, some highly talented players debuted while relatively young, so we would expect somewhat different results from this set of regressions. Expanding upon the notion in footnote 8, the switch from age to experience can be thought of as a relative rightward shift of the higher quintile productivity paths in Fig. 1.[20] In addition to delaying the peak productivity points for the highest ability quintiles in terms of experience, the shift would mean that we would expect rates of player development at a given age to differ between quintiles. The best players would be expected to develop faster at low levels of experience and to have productivity deteriorate more slowly at post-peak levels of experience *even when Implications (a) and (b) are true.* One testable implication of this is that we would expect the Chow F-statistics for differential slopes to be statistically significant, even when they were not when conditioning by age, as was the case in Table 2.

We again begin with a pooled model, allow differential intercepts in Model 2, and allow differential slope and quadratic terms in Model 3, still weighting each player–season

[18] So as not to overstate this finding, it should be noted that the estimated raw OPS levels for top players at ages 26 through 33 are all within 15 points of one another, and given the stochastic error in year-to-year performance these fits could quite easily be thought of as statistically equivalent (the peak resembling a plateau more than a point). Also, the Chow test in Model 3 fails to reject the null that players in different quintiles have statistically equivalent aging patterns aside from having differential quintile intercepts.

[19] The apparent similarity in age of MLB debut across quintiles may be slightly exaggerated in Fig. 2. While players in lower quintiles might be brought up as part-time players while young, and play their way into a regular job over two or three seasons, this is less common for top prospects who are projected to become stars. Because free agent eligibility is determined by time on a major league roster rather than playing time, having a young player sit on the major league bench has an opportunity cost of future productivity rents. This opportunity cost will be higher for players who are far above replacement level, so teams are reluctant to promote them until there is a starting position open for them. So while all quintiles are being called up at similar ages, the early observations for the top quintiles tend to represent much more playing time, and have more weight in the Table 2 regressions.

[20] The mean and standard deviation of age at debut for the highest ability players (q5) are 21.6 and 1.6 years of age, respectively, while the corresponding measures for the lowest ability players (q1) are 23.4 and 1.9 years.

Table 3 Estimated indexed OPS (iOPS), by experience and ability quintile, 1985–2005

Panel A: Regression diagnostics

Coefficient	Model 1	Model 2	Model 3
Intercept	0.981** (0.005)	0.887** (0.005)	0.904** (0.010)
Experience	0.011** (0.001)	0.007** (0.0006)	0.001 (0.004)
Experience^2	−0.0006** (0.0001)	−0.0007** (0.0001)	−0.0004 (0.0003)
Diff. intercepts	No	Yes	Yes
Diff. exp coefficients	No	No	Yes
Diff. exp^2 coefficients	No	No	Yes
Observations	6050	6050	6050
R^2	0.011	0.396	0.401
F-statistic (Chow)		964.0**	6.4**

Panel B: Fitted equations and peaks from Model 3, by quintile

Quintile	Intercept	β_1 (exp)	β_2 (exp^2)	Peak exp	Peak iOPS
1	0.904	0.001	−0.0004	1.9	0.905
2	0.969	0.001	−0.0004	1.4	0.970
3	1.012	0.003	−0.0006	2.7	1.017
4	1.034	0.008	−0.0008	5.3	1.055
5	1.096	0.018	−0.0012	7.5	1.165
Pooled	0.981	0.011	−0.0006	8.9	1.028

* Significant at 90% confidence level

** Significant at 99% confidence level

Standard errors in parentheses

observation by the number of plate appearances and indexing to control for seasonal and positional variation in OPS. The coefficients are reported in Table 3. As predicted, the Chow test in Model 3 of Panel A rejects the null hypothesis of identical slopes at the 99 percent confidence level. Panel B of Table 3 presents the fitted equations, which clearly show that the higher ability quintiles are exhibiting more curvature from within-career variation and have a higher initial slope in the pre-peak years of experience, as we observed in the regressions using age reported in Table 2.

Unlike the models using age in Table 2 where the intercept had no economic interpretation, the intercepts in Panel B of Table 3 estimate the mean indexed OPS for players in that quintile at the time of their debut. The Chow test in Model 2 of Panel A serves to confirm what we saw in Figure 2, that higher quintile baseball players are not debuting as soon as they reach replacement level, but have higher productivity (that is, statistically significant differential intercepts) as they enter the league. Even though the sample doesn't observe the highest ability players at replacement level, Implication (c) would be supported by any pre-debut development function that does not impose a discrete "jump" that disproportionately benefits future stars.[21]

Panel B of Table 3 reports the experience level of peak performance for each ability quintile cohort and for the pooled sample using the estimated equations. While the lower three quintiles peak early, with slow deterioration beginning after only two or three years of experience, players in quintiles 4 and 5 continue developing until peaking after 5.3 and 7.5 years of experience, respectively. While we previously observed an overestimated peak with the pooled regressions using age, the estimated peak from the pooled regression is vastly overestimated, at nearly 9 years of experience, due to earlier retirements of lower ability players. The reason for the increased distortion here, of course, is that because all players are in the population at zero years of experience, the pooled sample regression in Table 3 perfectly exemplifies the extreme where all the observational bias effects are in the right hand tail.

[21] Implications (c) and (d) could both be formally tested using forecasting techniques that project the estimated regression equation

Footnote 21 continued
outside the data range. All of those techniques, in one way or another, factor the extent of data extrapolation into their forecasting errors and confidence intervals, and in this manner would be conducive to formal hypothesis testing even at age and experience levels where the are no observations. We defer from using this approach, however, as the nature of our project is inherently empirical and we wish to avoid the assumption, necessary for extrapolation or forecasting, that the structural form of productivity growth or decline is unchanged outside our data range.

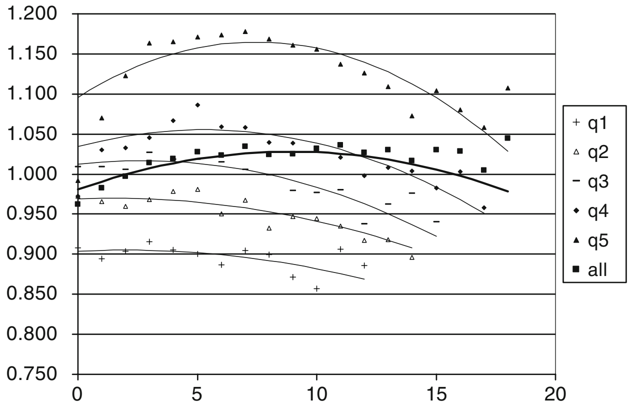

Fig. 3 Estimated indexed OPS (iOPS), by experience and ability quintile, 1985-2005. Experience (in seasons) is on horizontal axis; OPS indexed by year and position is on vertical axis. Estimated data points only shown for experience-quintile combinations with at least 10 observations. The corresponding estimated equations are shown in Table 3

Figure 3 suggests that high ability players in baseball could be identified quite early in their career, and shows that high ability players continue to enjoy high productivity until quite late in their careers.[22] While it is true that older players' abilities do fade, as a general rule they retire before they fall into mediocrity, and they are still having very productive years after nearly all the quintile 1 players in their "rookie class" have retired.

5 Career salary paths

Table 4 shows annual arithmetic mean and geometric mean salaries in MLB for each season from 1985 to 2005. Arithmetic means have increased from about $448,000 to over $3.3 million in that 21-year span, while geometric means have increased from about $267,000 to over $1.57 million. Use of the geometric mean as the measure of central tendency mimics the traditional preference of log-salaries to dollar salaries in labor market regression analysis to counteract the right-skewness of the salary distribution.

Table 4 Average Major League Baseball (MLB) salaries, 1985–2005

Year	Arithmetic mean for all players ($millions)	Geometric mean for all players ($millions)	Arithmetic mean for free agents* ($millions)	Geometric mean for free agents* ($millions)
1985	0.494	0.372	0.777	0.682
1986	0.448	0.280	0.837	0.715
1987	0.453	0.267	0.876	0.716
1988	0.474	0.288	0.978	0.804
1989	0.527	0.307	1.050	0.857
1990	0.572	0.341	1.171	1.001
1991	0.896	0.498	1.832	1.523
1992	1.075	0.553	2.179	1.760
1993	1.052	0.469	2.519	1.925
1994	1.143	0.517	2.473	1.805
1995	1.153	0.450	2.823	1.825
1996	1.228	0.511	2.908	2.028
1997	1.451	0.638	3.314	2.404
1998	1.505	0.666	3.627	2.573
1999	1.643	0.734	4.004	2.950
2000	2.198	1.021	4.511	3.278
2001	2.508	1.093	5.590	4.081
2002	2.741	1.213	5.754	4.008
2003	2.967	1.327	6.135	3.940
2004	2.892	1.282	6.004	3.743
2005	3.335	1.572	6.078	4.080

Mean salaries for free agents in the rightmost two columns have been weighted by plate appearances and do not include salaries of pitchers

To control for inflation and other institutional changes, we index a player's salary relative to the geometric mean salary of players who are free agent eligible in MLB that season. Players with 6 or fewer years of experience were omitted from the calculation of the geometric mean due to their limited ability to contract, which suppresses their salaries below competitive market levels.[23] Table 4 also shows the means for free agent position players, weighted by plate appearances. The arithmetic mean for this group increased over time from $777,000 to $6.08 million, while the geometric mean increased from $682,000 to $4.08

[22] As an anonymous referee has pointed out, the 3rd best season might be difficult to operationally utilize to forecast the ability level of a young player. At the very least, such an algorithm would require three qualifying (>130 AB) seasons for the player. While the focus of our paper is not to predict the ability level of a player based on playing statistics, we still contend that such players could be identified at an early age. For example, a slight modification of our methodology, using the 2nd best qualifying season up to and including the season where the player's (opening day) age was 25 would allow for fairly accurate prediction of a player's eventual talent quintile. While this prediction would be limited to players who have had two qualifying seasons by age 25, who are disproportionately high ability players, of those player assigned to q5 by the alternative metric, 78.6% are ultimately assigned to q5 by the method described in the paper (3rd best of all seasons).

[23] In the MLB collective bargaining agreement, experience is measured in days of service on a team's roster, but data on service time is not publicly available. We estimate experience crudely by subtracting the player's debut year from the season year. We set the cutoff at 7 years rather than 6 both due to the frequent practice of short end-of-season callups for young players, which will lead to an early debut but little service time, and as this cutoff maximizes goodness of fit in our later regression models.

Table 5 Estimated adjusted MLB salaries (lnnsal), by experience and talent quintile, 1985–2005

Panel A: Regression diagnostics

Coefficient	Model 1 years 1–6	Model 2 years 1–6	Model 3 years 1–6	Model 4 years 7+	Model 5 years 7+	Model 6 years 7+
Intercept	0.458**	0.322**	0.418**	0.839**	0.682**	0.644**
	(0.014)	(0.019)	(0.024)	(0.041)	(0.042)	(0.190)
Experience	0.072**	0.072**	0.045**	0.043**	0.029**	0.027
	(0.007)	(0.006)	(0.012)	(0.007)	(0.006)	(0.038)
Experience^2	0.0038**	0.0032**	0.0028	−0.0018**	−0.0017**	−0.0009
	(0.0010)	(0.0009)	(0.0018)	(0.0003)	(0.0002)	(0.0019)
Diff. intercepts	No	Yes	Yes	No	Yes	Yes
Diff. exp coefficients	No	No	Yes	No	No	Yes
Diff. exp^2 coefficients	No	No	Yes	No	No	Yes
N	3211	3211	3211	2839	2839	2839
R^2	0.511	0.631	0.653	0.067	0.442	0.447
F-stat (Chow)		43.0**	25.5**		78.5**	3.0**

Panel B: Fitted equations for years 1–6 from Model 3, by quintile

Quintile	Intercept	β_1 (exp)	β_2 (exp^2)	lnnsal in year 5
1	0.418	0.045**	0.0028	0.712
2	0.461	0.053**	0.0048	0.844
3	0.432	0.075**	0.0028	0.879
4	0.399	0.115**	−0.0002	0.968
5	0.463	0.140**	−0.0025	1.101
Pooled	0.458	0.072**	0.0038**	0.912

Panel C: Fitted equations for years 7 and beyond from Model 6, by quintile

Quintile	Intercept	β_1 (exp)	β_2 (exp^2)	Peak exp	Peak lnnsal
1	0.644	0.027	−0.0009	14.3	0.834
2	0.902	0.019	−0.0013	7.7	0.976
3	1.035	0.012	−0.0013	4.7	1.064
4	1.060	0.019	−0.0013	7.1	1.126
5	0.961	0.045	−0.0023	9.9	1.186
Pooled	0.839	0.043**	−0.0018**	11.7	1.089

* Significant at 90% confidence level

** Significant at 99% confidence level

Regressions include d.v. to control for defensive position, and for defensive position at each quintile in Models 2 and 3

Standard errors in parentheses

million.[24] As applying natural logarithms to a simple ratio index presents mathematical problems, we have adjusted the formula using a scaling technique to assist computation of a statistic that is log-normally distributed, as are the salaries themselves.[25,26]

[24] The overall upward trend in salaries is not monotonic, interrupted briefly by collusion, expansion, and other short-term influences, and is faster than the rate of inflation. There were also particularly sharp rises in 1990 to 1992 and from 1998 to 2001. The annualized average growth rates of the arithmetic mean, geometric mean, and CPI-U are 10.8, 9.4, and 3.0%, respectively. As the correlation between the growth rates of geometric mean salary and arithmetic mean salary is over 0.90, normalizing relative to the algebraic mean does not meaningfully alter our results.

[25] As our goal is to relate a player's salary relative to the geometric mean salary of free agents, a natural starting point would be to simply calculate a simple logarithm of the ratio, ln(ratio) = ln(salary/gmsal), where gmsal is the league-wide geometric mean salary for free agents. To solve the mechanical problem of below-average salaries yielding negative ln(ratio) results, we simply add one and divide by a number, ln(100), which is sufficiently large such that transformed ratio is non-negative. Thus, lnnsal = 1 + ln(salary/gmsal)/ln(100). The function calibrates the statistic so that a player with the (geometric) mean salary has an lnnsal of 1.00, a player with double

To estimate the best-fitting regression line to these trends, we tried several structural break points and functional forms. The intuitive structural break between 6 and 7 years of experience is used to indicate the approximate date of passage into free agency. Before and after the onset of free agency, separate quadratic regressions estimate the parameters of the quadratic function, with each player–season in the sample weighted by the number of plate appearances. Although there are some "rough seams" between years 6 and 7, the overall fit is good.

The regression results in Panel A of Table 5 show that experience alone explains over half of the indexed log-salary ratio (hereafter, "adjusted salary") for players in their first six seasons, as their salaries increase with ability to contract. The Chow test F-statistic in Model 2, which adds differential intercepts for each quintile, refutes the null that players in all quintiles debut at the same salary level. A look at the intercepts reported in Panel B of Table 5, however, does not show a clear ordinal pattern between the quintiles. The low intercepts for the estimated adjusted salary paths for the quintiles of players as shown in Figure 4 are tightly packed, with rookie players receive only about 8 percent of the geometric mean for free agents.[27]

The Chow test in Model 3, which adds differential slopes and quadratic terms, indicates that the quintiles have dissimilar slopes, and the slope coefficients reported in Panel B show that salary increases are directly related to ability and tend to be log-linear across time, as the quadratic terms are of low magnitude and statistically insignificant. Figure 4 shows that the salaries of players with different ability levels are clearly separate by seasons 4 through 6, as negotiating power increases through final-offer arbitration and the desire of some teams to negotiate multi-year contracts for emerging stars prior to the onset of the player's free agent eligibility.[28] The effect of the salary

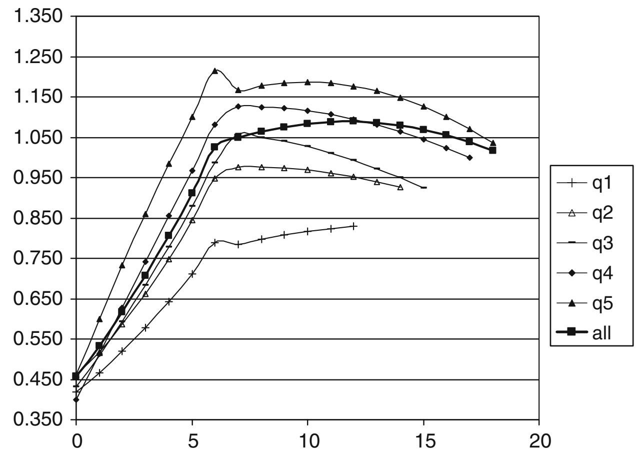

Fig. 4 Estimated adjusted MLB salaries (lnnsal), by experience and ability quintile, 1985–2005. Experience (in seasons) is on horizontal axis; adjusted salary (lnnsal) by year and position is on vertical axis. Estimated data points only shown for experience–quintile combinations with at least 10 observations. The corresponding estimated equations are shown in Table 5

increases can be seen in the rightmost column of Panel B, which show that by the end of the fifth year of experience, player salaries are correctly rank ordered by MPP, although only the very best players are expected to have salaries as high as the free agent geometric mean.

Panel C of Table 5 shows that baseball player salaries are relatively stable once a player reaches free agency. The slope coefficients yield downward trending estimated adjusted salaries as players reach their final seasons, but the standard errors are so large that the coefficients are not statistically significant. The 1.186 adjusted salary ratio for top players at their salary peak represents a salary 2.35 times the geometric mean salary for all free agents in the league.

The R^2 values in Models 4, 5, and 6 in Panel A show that it is the quintile differential intercepts that explain the lion's share of salary variation, so that the between-player variation is more significant than within-career variation in free-agent salaries. Even so, Fig. 4 shows that expected salaries fall relative to the geometric mean of salaries of all free agents. It is also that case that salaries fall in dollar terms for players as they near retirement, a pattern that is disguised by the observational bias in the right tail of experience.[29] An analyst studying the naive pooled sample model indicated by the thick trend line might conclude that

Footnote 25 continued
the mean salary will have lnnsal of ln(2)/ln(100) = 1.151, and a player with half the mean salary will have lnnsal of ln(0.5)/ ln(100) = 0.849.

26 The unadjusted salary data is quite skewed (skewness = 2.94, kurtosis 15.21). The skewness of the transformed salary measure (lnnsal) is −0.08, while the kurtosis is 1.82.

27 The mean adjusted salary (lnnsal) for players with 0 years of experience is about 0.45. Inverting the formula, the simple ratio of salary to the geometric mean salary is exp((lnnsal − 1) * ln(100)). For lnnsal of 0.45, this equates to about 0.0794, or about 8% of the geometric mean salary.

28 The mean adjusted salaries conditional upon years of experience confirm that the trend over seasons 1–6 is rising steadily, not in discontinuous jumps as would be true if freedom to contract increased discretely after the third and sixth season. One reason for this comes from possible measurement error in our estimation of arbitration and free agent eligibility. Another would be that teams will offer some players multi-year contract extensions before the player attains

Footnote 28 continued
eligibility status, and a portion of the player's expected salary increase at that future date can then be collected at the time of the contract extension, consistent with the notion of goodwill in contracting.

29 The unadjusted mean (dollar) salary level falls after 10 years of experience for all quintiles except the lowest ability players (q1).

end-of-career salaries decline far more slowly than they actually do.

6 Pay and productivity

Due to the multiple adjustments and instances of indexing we have applied to the raw data to assist our analysis of the age trends in performance and in pay, it is unclear what the "efficient" ratio of our constructed estimators would be in an ideally functioning labor market.[30] For this reason, our discussion of pay and performance will be of a relatively heuristic nature, and should be considered preliminary. That said, we wish to place our research in the context of the existing literature.

Previous attempts to compare player productivity to compensation have followed the seminal theoretical work of Scully (1974). Using a two-stage model, team revenue is shown to be primarily determined by team wins, establishing that the production of wins is the player's marginal physical product, and that the revenue accruing to the team from a player's performance statistics determine his marginal revenue product. In its final form, the second-stage equation usually takes the form of a semi-log model, as in Eq. 3.

$$\ln(\text{salary})_{\mathrm{i}} = \alpha + \beta \text{ performance}_{\mathrm{i}} + \delta' \mathbf{X}_{\mathrm{i}} + \varepsilon_{\mathrm{i}} \tag{3}$$

In Eq. 3, $\mathbf{X_i}$ represents a vector of exogenous control variables which influence salary, such as the (log) population of the player's host city, whether the player has been named the league MVP or an All-Star, and typically contains a pair of dummy variables to distinguish the levels of negotiating freedom. By changing the methodological approach to one that is comparing time paths of pay and performance with respect to changes in a mutual covariate, experience, the inclusion of control variables is not necessary.[31]

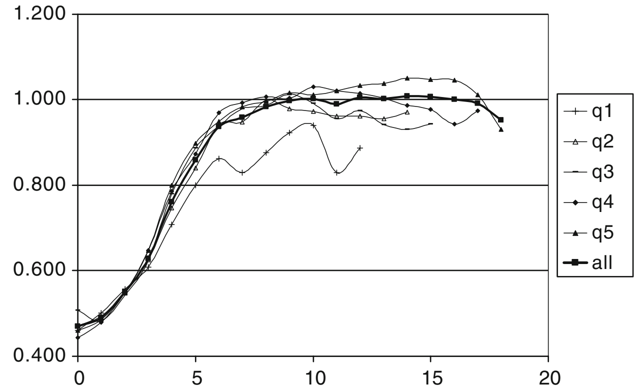

Fig. 5 Average ratios of adjusted salary over indexed OPS, by experience and ability quintile, 1985–2005. Experience (in seasons) is on horizontal axis; the ratio of adjusted salary over indexed OPS (lnnsal/iOPS), is on vertical axis Data points only shown for experience–quintile combinations with at least 10 observations

For the purpose of constructing a pay-to-productivity ratio for MLB, we use the adjusted salary formula discussed earlier in Sect. 5. The adjusted salaries are approximately normally distributed and exhibit symmetry between players receiving n times the average salary and players receiving $1/n$ times the average salary.[32] Dividing this measure by the indexed performance measure discussed in Sect. 4 results in the statistics displayed in Fig. 5. The key pattern here is that during the years of limited negotiating power, the ratio is quite low, and that beginning with the advent of free agent eligibility, the ratio becomes very close to 1.[33] Moreover, this is true for players in *all* quintiles, and is such that the ratio stays stable and does not experience a significant drop as the player ages.[34] This indicates that as productivity levels slowly decline, salaries decline at a proportionate rate. The increased variation in the data points in the right tail is likely due to reduced sample sizes in each conditional mean.

Although our methodology is very different from that in the traditional literature, the ratio analysis both confirms earlier findings regarding the underpayment of inexperienced players, and adds the novel finding that MLB free agents are paid approximately their MRP not just *on*

[30] If one were to assume that the efficient salary for a player of position-specific mean ability is the geometric mean salary of all free agents, and that deviations of ability from the position-specific mean should be rewarded with log-linear increases in salary, then a ratio of 1.00 could be considered efficient. It is not at all clear to us that either of those conditions should necessarily hold.

[31] See also footnote 7 where we contrast our approach to that of the traditional Scully approach. While it is true that the structural causes of changes in the ratio may go undetected with this method, the magnitude of the change in ratio will still be correctly measured. For example, suppose that the two main components of change in the ratio were increased freedom to negotiate contracts and migration of high ability players from small markets early in their careers to large markets mid-career. While we will remain unable to decompose our overall salary ratio increase into estimates for what proportion of the change in relative salaries is due to each effect, the fitted salary ratios should nonetheless correctly measure the sum total of all the partial effects.

[32] The only significant departure from normality is the spike that occurs at the truncation in the left tail caused by the league minimum salary.

[33] As mentioned previously, however, given the multiple manipulations of the two statistics being compared, we cannot claim that a ratio of 1.00 represents market efficiency, but only that relative player salary is increasing in the ratio.

[34] The contingent mean ratios for the lowest quintile of players vary more during years 6 through 11 than the ratios for the other quintiles, due to increased standard error from small sample sizes. But as the ratios are only varying between 0.83 and 0.93, this says more about the stability of the ratios for the other quintiles than about erratic pay ratios for lower ability players.

average, but *throughout* the talent distribution. Further, we find that the level of underpayment of inexperienced players is also consistent throughout the talent distribution. Our measures of underpayment for players are a bit higher than, but largely consistent with those reported in the previous literature.[35] Another source of possible insight into pay and performance issues is to inspect the peak levels of each reported in Tables 3 and 5. For each quintile, adjusted salary peaked at least 1.8 years after productivity peaked, with an average lag of 2.8 years. We repeat the caveat associated with multiple indexing, and we do not claim this is evidence of labor market inefficiency, but merely present this as an empirical finding so that future theories of baseball labor market efficiency might incorporate it.

7 Summary of findings and future plans

Quintile data permit exploration for differing rates of skill development and deterioration for players of different levels of peak ability, and comparison of the variation in a player's productivity across his career to the variation of talent within the league. Even with star players retiring before they completely regress to replacement level and the partial left tail truncation caused by rent-maximization decisions of MLB owners, the relatively extended careers of players in higher quintiles result in biased slope coefficient estimates from pooled sample regressions on age or experience. As the observational bias is predominately in the right tail for age (and completely so for experience), estimated rates of development and deterioration are systematically underestimated and peak ages and experience levels are overestimated by between 2 and 5 years in pooled productivity regressions on age.

The data in MLB show some, but not all, of the characteristics that would be consistent of the theory of "nested" ability quintiles proposed in Sect. 2.1. Due to partial truncation of the left tail of the sample, young MLB players within a given ability quintile will debut at a productivity level similar to or above the peak ability level of players in the next lower quintile. In the right tail of the age distribution, star players tend to retire before their skills fall to the level of peaking journeymen, although mid-quintile players attempt to linger longer. In combination, these results strongly suggest, in accordance with the idea of nesting, that stars reach replacement level at younger ages than less-talented players, and remain above replacement level well after lower-peaking players have been forced into retirement.

It is not true, however, that the productivity paths of the various ability quintiles only differ through a fixed differential intercept. The data show evidence that higher ability players develop faster than lower peaking players. Higher quintile players show substantially more within-career variation. MLB players in the top two quintiles appear to peak about 2 years later than players in the lowest three quintiles. So as not to overstate the importance of these results, the variation in an individual player's ability due to development and deterioration at the MLB level (standard deviation of about 30 points) is of far smaller magnitude than between-player variation (standard deviation of over 112 points).

Use of the quintile methodology along with the calculation of adjusted salary to indexed productivity ratios permit a suggestive examination of underpayment of players with reduced negotiating power due to league collective bargaining agreements. As in previous literature, we find that salaries of young players are suppressed below those of similarly talented older peers. A novel result, however, is that we show that salaries of young players of high ability increase more rapidly than salaries of more marginal young players throughout their first years in the league, despite the theoretical ability of teams to exert monopsony power. The negotiating flexibility provided by multi-year contracts allows players to obtain a portion of the rents that would otherwise be extracted from them in their years prior to free agent eligibility, and this results in a smoothing of the career salary path. Interestingly, at all levels of ability, salary levels peak at least 1.8 years after hitting productivity peaks in baseball.

Being able to correctly estimate rates of productivity increase and decline allows competitively negotiated salaries to more closely approximate subsequent productivity, helps managers to make more cost-effective personnel decisions, and assists the calculations of rents in imperfectly competitive labor markets to inform industry-wide policy or regulatory decisions.

In our future research, we will conduct a similar analysis using player productivity and salary data from the NBA, in which imperfect competition in the labor market persists even for veteran players. In an extension of our MLB research, we will refine the methodology used to relate pay to performance, and additionally estimate separate peak ages for distinct skills within baseball and measure the improvements in productivity forecasts

[35] Burgess and Marburger (1992) find salaries of arbitration-eligibles to be 80 to 90% higher than for "ineligibles" (players in their first 3 years), and Kahn (1993) found that salaries for arbitration and free agent eligibles were 35 to 50% higher than ineligibles, but the comparisons in both papers were damaged by data from the collusion era. Marburger (1994) finds salaries for average players in their first 3 years to range from 18 to 41% of those of free agents. Vrooman (1996) estimates that Tier 1 (ineligible) players receive about 28% of the salary of free agents. Using data from 2000 to 2004, Hakes and Sauer (2006) produce a model that estimates Tier 1 player salaries at 17 to 21% of free agent salaries.

that result from replacing the current one-dimensional OPS model.

Acknowledgments The authors wish to thank Glenn Knowles and session attendees at the March 2007 Midwestern Economics Association meetings for their helpful comments on an earlier draft of this paper, as well as several anonymous referees.

References

Albert J (1999) Bridging different eras in sports: comment. J Am Stat Assoc 94:677–680

Berri DJ, Simmons R (2009) Race and the evaluation of signal callers in the National Football League. J Sports Econ 10:23–43

Burgess PL, Marburger DR (1992) Bargaining power and baseball. In: Sommers PM (ed) Diamonds are forever: the business of baseball. Brookings Institute Press, Washington DC

Fair RC (2008) Estimated age effects in baseball. J Quant Anal Sports 4:11

Fort R (1992) Pay and performance: is the field of dreams barren? In: Sommers PM (ed) Diamonds are forever: the business of baseball. Brookings Institute Press, Washington DC

Hakes JK, Sauer RD (2006) An economic evaluation of the Moneyball hypothesis. J Econ Perspect 20:173–185

Hamilton BH (1997) Racial discrimination and professional basketball salaries in the 1990s. Appl Econ 29:287–296

Horowitz I, Zappe C (1998) Thanks for the memories: baseball veterans' end-of-career salaries. Manag Decis Econ 19:377–382

Humphries BR (2000) Equal pay on the hardwood: the earnings gap between male and female NCAA division I basketball coaches. J Sports Econ 1:299–307

James B (1982) The Bill James baseball abstract. Ballantine Books, New York

James B (1988) The Bill James baseball abstract. Ballantine Books, New York

Kahn LM (1991) Discrimination in professional sports: a survey of the literature. Ind Labor Relat Rev 44:395–418

Kahn LM (1993) Free agency, long-term contracts and compensation in Major League Baseball: estimates from panel data. Rev Econ Stat 75:157–164

Kahn LM (2006) Race, performance, pay, and retention among National Basketball Association head coaches. J Sports Econ 7:119–149

Krautman AC (1993) Shirking of stochastic productivity in Major League Baseball: reply. South Econ J 60:241–243

Krohn GA (1983) Measuring the experience-productivity relationship: the case of Major League Baseball. J Bus Econ Stat 1:273–279

Madden JF (2004) Differences in the success of NFL coaches by race, 1990–2002: evidence of the last hire, first fire. J Sports Econ 5:6–19

Marburger DR (1994) Bargaining power and the structure of salaries in Major League Baseball. Manag Decis Econ 15:433–441

Marburger DR (2004) Arbitrator compromise in final offer arbitration: evidence from Major League Baseball. Econ Inq 42:60–68

Miller PA (2000) A comparison of free agent and arbitration-eligible salary structures in Major League Baseball. South Econ J 67:87–104

Schulz R, Musa D, Staszewski J, Siegler RS (1994) The relationship between age and Major League Baseball performance: implications for development. Psychol Aging 9:274–286

Scully G (1974) Pay and performance in major league baseball. Am Econ Rev 64:915–930

Singell LD Jr (1991) Baseball-specific human capital: why good but not great players coach in the major leagues. South Econ J 58:77–86

Vincent C, Eastman B (2009) Determinants of pay in the NHL: a quantile regression approach. J Sports Econ 10:256–277

Vrooman J (1996) The baseball players' labor market reconsidered. South Econ J 63:339–360

Zimbalist A (1992) Baseball and billions. Basic Books, New York