



Canadian Labour Market and Skills Researcher Network

Working Paper No. 75

Heterogeneous Worker Ability and Team-Based Production: Evidence from Major League Baseball, 1920-2009

Kerry L. Papps
University of Oxford
IZA

Alex Bryson
National Institute of Economics and Social Research
Centre for Economic Performance

Rafael Gomez
University of Toronto
London School of Economics

March 2011

CLSRN is supported by Human Resources and Skills Development Canada (HRSDC) and the Social Sciences and Humanities Research Council of Canada (SSHRC). All opinions are those of the authors and do not reflect the views of HRSDC or the SSHRC.

Heterogeneous Worker Ability and Team– Based Production: Evidence from Major League Baseball, 1920-2009

Kerry L. Papps^a, Alex Bryson^b and Rafael Gomez^c

March 2011[†]

^a Corresponding author; University of Oxford and IZA, Bonn. Address: Nuffield College, New Road, Oxford, OX1 1NF, United Kingdom; Email: kerry.papps@economics.ox.ac.uk; Tel: +44-1865-278 994; Fax: +44-1865-278 621.

^b National Institute of Economic and Social Research and Centre for Economic Performance. Email: a.bryson@niesr.ac.uk; Address: 2 Dean Trench Street, Smith Square London, SW1P 3HE United Kingdom.

^c University of Toronto and London School of Economics. Email: ralph.gomez@utoronto.ca; Address: 1101 Pharmacy Ave, #301, Toronto, ON M1R 2H2, Canada.

[†] The authors would like to thank Robert Simmons for helpful comments and suggestions and Jim Davis for supplying us with metropolitan area population figures.

ABSTRACT

A detailed longitudinal dataset is assembled containing annual performance and biographical data for every player over the entire history of professional major league baseball. The data are then aggregated to the team level for the period 1920-2009 in order to test whether teams built on a more intermediate distribution of observed talent perform better than those teams with either too high or too low a mixture of highly able and less able players. The key dependent variable used in the regressions is the percentage of games a team wins each season. Our finding is that conditioning on average player ability, dispersion in team pitching and hitting talent prior to the start of a season is related in a non-linear way to subsequent team performance. This suggests that there is an optimum heterogeneity of ability at the team level that maximises joint output. This result is robust to the inclusion of team fixed effects as well as year dummies and after controlling for the potential endogeneity of skill dispersion. These findings have potentially important applications both inside and outside of the sporting world.

JEL Code: J24, J21

Keywords: Baseball, Inequality, Team-based Performance

Executive Summary

Is skill dispersion a source of competitive advantage for teams? Or does a more compressed distribution of skills foster greater collective output? While it is established that the aggregate endowment of skill is an important determinant of team performance, Alex Bryson, Kerry Papps and Rafael Gomez investigate whether the distribution of skills inside a team can play a significant role in the determination of overall team performance.

Most managers would attest that mobilising collective effort is markedly different than managing a collection of individuals that require no collaborative interactions. All else equal, at the individual level, the most skilled worker would always be preferred over the least skilled. In a team-based environment, however, it is not necessarily the case that output is maximised by attracting individuals with the largest aggregate endowment of skills, without regard to what effect those individuals have on the distribution of skills within the team.

Recently, economists have become interested in the importance of skill dispersion as a determinant of outcomes as varied as aggregate trade flows (Ohnsorge and Trefler, 2007), educational achievement (Lavy et al., 2009), and firm productivity (Hamilton et al., 2003). The consensus is that the distribution of skills, whether at the level of a nation, region, classroom or organisation matters. Where disagreement still exists is over the effects produced by greater or lesser skill dispersion. In some studies, greater spreads of talent seem to be beneficial to aggregate output while in others the opposite appears to hold true.

Our research asks whether or not there is an optimal spread of talent that maximises performance in settings where team-based production is operative. Specifically, we consider whether it is optimal for managers to assemble teams solely on the basis of average ability (irrespective of the effect this may have on the distribution of skills) or whether organisations should manage the selection of workers so as to prevent too wide a gap opening up between the best and poorest performers. Put another way, if a manager is forced to choose two workers (e.g., baseball players in our case) whose average ability is the same (e.g., a combined historical batting average of 0.275), is it better to approximate the average more closely (0.270 and 0.280 respectively) or should one star (0.325) and one less able player (0.225) be hired? And at what point would too large or too narrow a spread in ability be damaging to team chemistry and performance?

Our analysis is based on annual performance and biographical data for every player over the entire history of major league baseball, which for the purposes of this study, are aggregated to the team level for the period 1920-2009. As individual performance measures, we use earned run average (ERA) for pitchers and on-base plus slugging percentage (OPS) for hitters. A low ERA or a high OPS indicates a good player. Our key finding is that heterogeneous ability measured prior to the start of a season is related in a non-linear way to team success, in that teams with a 'balanced' distribution of

on-base plus slugging percentage among hitters and earned run average among pitchers win more games than teams with either too low or too high a skill spread.

One of the major implications of this research is that in settings where output is jointly determined, such as in any team-based enterprise, considerations of skill distribution are likely to emerge. We find evidence of an optimal level of performance dispersion within baseball, even though both offensive and defensive roles require less on-the-field coordination and interaction than other team sports such as basketball, football or hockey. This is because baseball output is still subject to intermediation, which much like assembly line work, depends in large measure on who precedes or follows a player in a pitching rotation or batting order.

The question of the optimal spread of individual ability within any given team, however, will always be decidedly hard to answer and future work would be productively directed at specifying what organisational attributes foster a more optimal spread of talent.

1. Introduction

“The real reason the [Chicago] Bulls won six NBA championships in nine years is that we plugged into the power of *oneness* instead of the power of one man. Sure, we had Michael Jordan, and you have to credit his talent. But at the other end of the spectrum, if players 9, 10, 11, and 12 are unhappy because Michael takes twenty-five shots a game, their negativity is going to undermine everything. It doesn’t matter how good individual players are – they can’t compete with a team that is awake and aware and trusts each other. People don’t understand that.”

Phil Jackson (2004), *The Soul of Teamwork*.

Does the superstar model of team building actually work or does having a more even distribution of talent improve team performance?¹ As reflected in the quote above, one of the key determinants of joint production is individual ability: “without the right horse (i.e., a Michael Jordan) you won’t win the race”. However, the other essential ingredient is team chemistry: how employees, irrespective of their individual abilities, work together. One way of ensuring team chemistry is by being as attentive to the distribution of team skills as to the overall average. Rather than simply hiring the best *individuals* that money can buy, the best managers attempt to maximise joint output by recruiting, assembling and motivating the best *group* of workers possible. This is true whether the manager oversees the sales force of a New Jersey real estate office, runs the shop floor of a manufacturing plant in Shanghai or coaches from the sidelines of a football pitch in Leeds.

In this paper we consider whether it is optimal for firms to hire employees solely on the basis of their average ability (irrespective of the effect this may have on the distribution of skills) or whether firms should manage the selection of workers so as to prevent too wide a gap opening up between the best and poorest performers. Put

¹ Throughout the paper we use ‘human capital’, ‘talent’, ‘ability’ and ‘skills’ interchangeably. Though we are aware of the subtle-to-large differences in interpretation these terms entail, the fact that we are drawing from a number of differing literature streams requires us to be a bit more flexible in our use of terms representing player quality.

another way, if a manager is forced to choose two workers (e.g., baseball players) whose average ability is the same (e.g., a combined historical batting average of 0.275), is it better to approximate the average more closely (0.270 and 0.280 respectively) or should one star (0.325) and one less able player (0.225) be hired? And at what point would too large (or too narrow) a spread in ability be damaging to team chemistry and performance?

We address this question with annual performance and biographical data for every player over the entire history of major league baseball, which for the purposes of this study, are aggregated to the team level for the period 1920-2009. Our key finding is that heterogeneous ability measured prior to the start of a season is related in a non-linear way to team success, in that teams with a ‘middling’ distribution of on-base plus slugging percentage among hitters and earned run average among pitchers win more games than teams with either too low or too high a skill spread. Though some heterogeneity is observed along the dimensions of pitching and non-pitching talent, our results are robust to the inclusion of team fixed effects and strengthen when we account for the endogeneity of talent dispersion. Possible rationales for this pattern of findings are discussed in the text and rely on mechanisms such as greater levels of teammate cooperation, inter-player learning and emulation, and ease of coaching associated with an optimally balanced line-up.

This paper proceeds as follows: Section 2 describes the conceptual framework employed in the paper. Section 3 describes data and measures used in our analysis. Section 4 details our results. Section 5 demonstrates the potentially important applications beyond baseball. A summary and discussion appear in Section 6.

2. Conceptual Framework

a. Links with Relevant Literature

This paper is situated between the growing literature measuring the effects of skill inequality and ethnic diversity on outcomes such as cross-country trade flows (Bombardini et al. 2009), urban productivity (Ottaviano and Peri 2006) and team performance (Kendall 2003; Hamilton et al. 2003; Frick 2003; Berri and Jewell 2004; Simmons et al. 2009) and the well-established literature on the importance of peers and local environment in determining outcomes of individuals, groups and regions (Angrist and Lang 2004; Benabou 1994, 1996; Glaeser et al., 2008; Gould et al. 2004, 2009; Hoxby 2000; Lavy et al. 2009; Oreopoulos 2003). These studies have established that across many dimensions (e.g., income, skill, ethnicity) too much inequality or diversity, however measured, acts as a drag on performance.²

Several papers confirm that the distribution of a country's stock of human capital can be the source of comparative advantage and can trigger specialisation in sectors characterised by higher substitutability among workers' skills (Ohnsorge and Trefler 2007; Bougheas and Riezman 2007; Bombardini et al. 2009). This is quite a novel finding given that the comparative advantage literature up to now has tended to show how the abundance (*i.e.* the average level) of factors of production acts as the driving force in determining international specialisation and development. In this new research stream the whole distribution of factor endowments, rather than just their average, can help explain observed trade flows. Industry-level bilateral trade data show that human capital dispersion, as measured by standardised literacy scores, has a significant effect on trade and that the effect is of a magnitude comparable to that of

² Boisjoly et al. (2006: 1902) suggest that programmes that explicitly try to counteract differences across groups can help improve along dimensions such as attitudes, but "have little or no effect on harder-to-change behaviour (such as befriending or socialising with someone from another racial/ethnic group) and long-term goals".

average (or aggregate) endowments (Bombardini et al. 2009).

This paper builds upon the work cited above but differs by looking squarely at the effect of heterogeneous ability within and across organisations. Here the literature is not as extensive and is focused mostly on how individual (not team) output responds to co-worker behaviour or ability (Winter 2004; Ichino and Maggi 2000; Mas and Moretti 2009). The most closely related papers to ours are Kandel and Lazear (1992), Hamilton et al. (2003), and Gould and Winter (2009). Kandel and Lazear's (1992) model of team production within the firm shows how social pressure can solve the free-riding problem by making more productive agents limit the shirking of less hard-working ones. Hamilton et al. (2003) examine detailed company records in a garment plant and find that with average ability held constant, the coefficient on the ratio of the maximum to the minimum individual productivity of team members is positive and significant, meaning that more heterogeneous teams are more productive. They show that more productive workers appear to improve the performance of the less able through a combination of worker helping effects, emulation and peer pressure. Gould and Winter (2009) take baseball teams as their unit of analysis and find that teammates have a positive effect on individual performance where individual performance is a complement in production (i.e., amongst batters or amongst pitchers) but that better teammate ability affects output negatively when it is a substitute in production (i.e., when there is a talented pitching staff batters perform less well and vice versa).

Like Hamilton et al. (2003), we examine whether heterogeneous ability amongst workers improves or dampens team performance, and like Gould and Winter (2009) we treat baseball pitching and batting talent separately and use conventional measures of batting and pitching efficiency as proxies for worker ability. Our paper differs by

examining the reasons why within groups with strong output interdependencies – where the structure of output depends on collaborative effort or intermediation – too unequal a distribution of talent, conditional on average talent being the same, may be detrimental to organisational performance.³ We also explore the potential for an optimal degree of inequality with respect to performance, in that there are some strong reasons to believe that heterogeneous ability could be good for performance but only up to a point.⁴ This logic is depicted in Figure 1, where skill dispersion and team output are positively related up to region A-B, after which, in a region of high skill dispersion, B-C, the relationship turns negative. The implication is that team winning percentage will not be highest where skill dispersion is highest, but rather where heterogeneous ability is moderate, such as point B. This set-up provides an extension of the predictions found in Hamilton et al. (2003) for garment factory workers.

[Figure 1]

The remainder of this section describes how the basic structure of output in baseball relates to the literature on diversity and team performance surveyed above. It then outlines the reasons why there might be an optimal level of heterogeneous ability in the baseball setting.

It should be stressed that we are not attempting to replicate the rich literature examining the effects that team wage inequality has on performance (see Simmons and Frick, 2008a, 2008b and references contained therein). Rather, we are interested in examining the relationship between heterogeneous *ability*, as distinct from income, and team outcomes.

³ In baseball, the production of a run is much closer to the production of output on an assembly line and hence closer to intermediation in the sense employed by Lazear (1995).

⁴ See Freeman and Gelber (2006) for more on the idea of optimal inequality and incentives.

b. Structure of Production in Baseball

Baseball teams attempt to win games by scoring and preventing as many runs as possible. Team wins offer a clear measure of team performance in baseball: they are the objective being maximised by most baseball managers/coaches.⁵ This may not be the only measure being maximised by the ownership of the team, however, as the quest for the most wins may dampen profitability if the marginal return from every win exceeds the cost of acquiring that win. To the extent that bringing on board better talent does not translate into more revenue for the club, the goal of maximising the win ratio (or ‘winning percentage’ as is commonly used by sports writers) is not perfectly aligned between principal and agent (ownership and manager/coach). Nevertheless, the win ratio should translate into profitability over the long run and hence is of worth to a majority of owners.

Given that a baseball team endeavours to win and that winning is a matter of scoring and preventing as many runs as possible, there are easily-observable offensive and defensive determinants of the overall winning percentage measure in the form of on-base plus slugging (OPS) and earned run average (ERA) respectively. The two objectives (scoring and preventing runs) within a given baseball team entail varied levels of cooperation and interdependence. Because one player principally controls the number of opposing team runs (i.e., the pitcher through a low ERA) and because a pitcher essentially plays on his own, the effect of heterogeneous ability among pitchers might be expected to exert a weaker effect on overall team performance than

⁵ Depending on a team’s standing and point reached in the season, the objective of maximising wins may change. Once a team has secured a playoff spot, they may wish to rest players or, conversely, once a team is clearly out of the pennant race, a manager may be inclined to give younger players more playing time. Both strategies are arguably undertaken to maximise long-run winning percentage, but in the short run clearly have the effect of reducing the likelihood of beating an opponent.

offensive players.⁶ Much greater will be the impact of first moments, or overall team ERA, which should be unambiguously inversely related to winning percentage. Having said this, pitchers spend a considerable amount of time training (if not playing) together and communicating in the bullpen, their success is also determined to a large extent by who precedes and succeeds them in the pitching order. For these reasons, pitchers (like non-pitching talent) would be expected to interact and be affected by the ability differentials of fellow pitchers.

c. Heterogeneous Ability and Team Performance: The Positive Range

In addition to pure hitting and pitching ability, constructing a team naturally requires attention to other dimensions of observable (e.g., speed) and non-observable (e.g., leadership) player quality, presenting teams with certain trade-offs. For a uniformly top-hitting club, acquiring a player in the off-season who is a weak hitter involves a sacrifice in terms of anticipated run production, but the team may be willing to make this sacrifice if the player is noted for his leadership abilities. The team increases the heterogeneity of its offensive production with such a change, but this is likely an optimal move away from equality of hitting talent as it could bring the right ‘mix’ to the team.

Benefits in team performance emanating from an expansion of heterogeneity could occur for several reasons. First, there is often a gap between the capability of a top player to perform some task during a game and his willingness to do so. A team full of high-performing batters might not be the easiest group to convince to ‘take one for the team’. The sacrifice play – whether it is designed to move a batter closer to

⁶ There are other factors limiting an opponent’s ability to score that are not directly related to pitching talent/ability such as quality of the catcher as well as infield and outfield defensive ability. But arguably, good pitching talent aids these team-defensive abilities, which are themselves less interdependent than other elements of production.

home plate or to generate a run – is at times integral to the manager’s ability to ‘manufacture’ runs. Assembling a team that increases the variance of OPS, over some range, could therefore be related to increases in subsequent team performance (controlling for defensive record) if this increased widening of non-pitching skills brings a new and complimentary dimension to a team.

Second, there may be other less obvious ways in which assembling an unequal batting roster at the start of a season may benefit a team’s subsequent performance. For instance, some widening of the skill distribution would allow the better offensive players to transfer their skills or training techniques to the lower-performing (and presumably less-experienced) hitters on the team.⁷ These and other instances of positive co-worker emulation could yield the patterns found in the experimental peer group literature (Falk and Ichino 2003) and studies of college-roommates (Sacerdote 2001), whereby better performers enhance the performance of those with lower incoming academic ability.

Due to the nature of defensive production in baseball (i.e., the fact that pitchers are segregated from the rest of the team in a ‘bullpen’ during a game and that a lower individual ERA will unambiguously lead to lower opposition scoring regardless of the distribution of pitching talent) one might expect to see weaker associations between ERA inequality and team performance. Pitchers, however, should also benefit, over some range, from differences in teammate ability. Firstly, the learning effects noted for non-pitchers are likely to be operative for pitchers as well. Secondly, teams that concentrate on acquiring starting pitchers with the ability to last many innings during a game (“to go deep in the pitch count”) would, in theory, enhance the effectiveness

⁷ Unfortunately, this might even extend to the use of performance-enhancing drugs, as evidenced by the revelations of New York Yankees pitcher Andy Pettitte and former star slugger Jose Canseco, who claim to have learned about the use of performance enhancing drugs from fellow teammates. Gould and Kaplan’s (2010) article in this issue suggests that this does indeed happen.

of relief pitchers as well by lowering the latter's pitch counts. This would in all probability increase the ERA of the starter, given the strong positive correlation between cumulative and immediate innings pitched and ERA (Bradbury 2010), but a game is more likely to be won in this instance from the effective combination of the starting pitcher and the relief pitcher.⁸ The widening of assembled pitching ability prior to a season could therefore be a predictor of team effectiveness during the season.

d. Heterogeneous Ability and Team Performance: The Negative Range

When assembled skill dispersion is at somewhat higher levels, a negative relation (such as the B-C range in Figure 1) may be expected for several reasons. First, for pitching staff in particular, wildly differing ERAs would allow opposing teams to exploit weaker pitchers by matching opposing hitting line-ups more easily or by forcing good pitchers to waste pitches and leave a game earlier than anticipated.⁹ When there is a more even distribution of pitching talent, opposing offenses lose this potential strategy and so, other things equal, we would expect that teams with too great an ERA dispersion at the start of season will have a lower winning percentage.

Scoring runs is somewhat different to preventing them in that offensive success depends crucially on who precedes and follows a batter in the order. Although an individual batter can score a run with no teammates on base by hitting a home run, only 18% of total runs were scored in this fashion over the period 1920-2009. Furthermore, a team with poor batters preceding and following a few otherwise very talented hitters will engender a pitching response from the opposing team in which the

⁸ Each additional pitch over a 10-game average increases ERA by 0.022 (Bradbury 2010).

⁹ Hitters are often instructed by managers to 'keep an at-bat alive' by swinging at pitches in a defensive manner, never intending to get a hit but to foul as many balls as possible and thereby wear down an opposing pitcher's arm.

least favourable pitches will be thrown to the best hitters, thus limiting the opportunity for good players to get on base or score runs.

From a manager's perspective, therefore, two teams with otherwise identical team OPSs may perform quite differently in terms of overall run production because of this interdependence in hitting and opponent strategy. Moreover, as noted by Gould and Winter (2009), batters behave more as complements than substitutes in production in that they are differentiated by fielding position: 'skilled positions' such as short-stop or 'power positions' such as first base and catcher tend to have different offensive expectations. Differences in the distribution of hitting ability across teams may therefore play a role in accounting for winning percentage differences.

Based on the above reasoning, conditioning on average offensive talent, the more unequal the OPS for a team at the start of a season, the lower will be overall team performance.

e. Is There an Optimal Distribution of Ability?

Given that there are two directions in which the distribution of skills could affect team performance, it may be that there is an intermediate range of skill inequality that maximises team output. Below a certain level of heterogeneity, players do not benefit from the assistance and motivation resulting from playing alongside teammates with complementary skills and greater talent. Beyond a certain level of heterogeneity, further increases in the variance of batting talent allows opposing teams' pitchers to exploit the weaknesses of lower-performing hitters and prevents good pitches from ever reaching the best hitters on the team. This could be the case if better hitters are either unwilling to help their teammates or if they are simply unable to do so because the gap between observed talent is too large to bridge. The best hitters would still

remain so in a rank order sense, but the effectiveness of a team would fall. Finally, one should not overlook the benefits that a more balanced distribution of skills might bring to a manager by reducing coaching difficulty.¹⁰ This is reinforced once again by Phil Jackson, who in an interview on the subject of “balance and basketball” stated that:¹¹

“... it’s the unselfish players – players who are more interested in reading what’s happening and keeping the flow going on the floor – who are the most valuable players that you have. They may only be averaging seven points a game... but their ability to play in a selfless manner gives the team its real opportunities. In those individuals, the power of *we instead of me* is more advanced. They feel more responsibility to the group, and that’s why you’re better off with maybe two very, very talented but perhaps selfish people on the team than five or six or seven. That’s why teams that are less talented but more selfless and group-oriented can have more success.”

To summarise, we expect skill distributions of intermediate orders for batters and pitchers to be the most likely to result in team success, while particularly high or low levels of skill inequality should be associated with the least successful team outcomes.

3. Data and Measures

A detailed longitudinal dataset was assembled from version 5.7 of Sean Lahman’s Baseball Archive (available from www.baseball1.com). This contains annual performance and biographical data for every player over the entire history of the major leagues; however, for the purposes of this study, the data were aggregated to the team level for the period 1920-2009. As well as providing an adequate sample period, 1920 was a landmark season in the major leagues as it was the first following both the Black Sox scandal and Babe Ruth’s now legendary move to New York. Full

¹⁰ The idea of a ‘balance is best’ approach to assembling an effective team is anticipated in the pluralist employment relations literature (Gomez et al. 2004).

¹¹ Interview at: http://findarticles.com/p/articles/mi_m1058/is_n35_v113/ai_18965404, accessed on March 12, 2010. Full text of article from *Christian Century* Magazine, December 4, 1996.

details of the construction of each of our variables can be found in the appendix. For the purposes of this analysis, a team that shifts cities is considered to be a continuation of the previous team, since typically most players and managers move with their teams. Details of the 30 teams thus defined are given in Table A1. The dependent variable in the regression analysis is the percentage of games a team wins each season.

Players are classified as either pitchers or non-pitchers. Although there are a number of possible performance statistics, this paper will focus on only two: earned run average for pitchers and on-base plus slugging for non-pitchers.¹² To form the best estimate of a player's ability in the face of year-to-year fluctuations in performance, these were calculated over each player's entire career as of the end of the previous season.

Career ERA was averaged over all pitchers within a team, using the number of innings pitched during the previous season as weights. Similarly, career OPS was averaged over all non-pitchers, weighted by the number of at-bats in the previous season multiplied by the sum of at-bats, walks, sacrifice flies and hit-by-pitches. This weight is appropriate since it is used as the denominator of OPS, just as ERA is defined as the number of earned runs conceded per nine innings pitched. In addition to weighting, a minimum playing time cut-off was also used to exclude players who made only a few appearances in the previous season.¹³ As measures of dispersion in performance, the weighted coefficients of variation in career OPS and ERA were constructed.

Finally, to avoid the mechanical correlation between team performance and

¹² OPS is preferable to the more common batting average as it captures the overall contribution of a player towards scoring runs, rather than just his ability to reach base.

¹³ A consequence of this is that players who are in their first season in the major leagues will be excluded from our OPS and ERA variables for that season.

individual performance that exists within a year, all the independent variables measure the career performance of players at the end of the *previous* season, thereby capturing the ability of players at the point a team chooses its roster (which is just before the start of a given season).

a. Potential Endogeneity and Choice of Instruments

One potential criticism of the use of the performance dispersion measures in the regression equation is that they may be endogenous. Teams choose which players to sign each season and less successful teams may have more incentive than others to modify their roster to achieve a more desirable level of performance dispersion, primarily because they face greater constraints in attracting and retaining top-priced talent. In this case, OLS estimates will underestimate the effect that a given shift towards optimal inequality has on winning percentage. There is also the possibility that managers may simply misjudge what the optimal level of inequality is for their team, but the direction of this bias would be indeterminate *a priori*.

To control for the potential endogeneity of skill dispersion with respect to win rate we construct an instrumental variable based on exogenous changes in team skill dispersion. This is defined as the difference between a team's coefficient of variation in either career OPS or career ERA and the coefficient of variation excluding the players who left a team at the end of the previous season for reasons beyond pure management control. This could occur if players retire, are chosen in the special drafts used to select players for the expansion teams or if they leave the team as free agents.

Additional instrumental variables that are used include a team's total home attendance during the previous season and the population in a team's metropolitan area. These are used as proxies for a team's likely income from ticket sales and

broadcasting revenue, respectively, and hence capture a team's ability to afford star players. Finally, a set of year dummies are used as instruments since distinct trends in performance dispersion are observed over the sample period, but not in team win rate, which is always 0.5 by construction.¹⁴

b. Descriptive Statistics

Table 1 presents the descriptive statistics for all the variables used in the analysis. Since some lagged variables are used in the analysis, the first season for each new team is excluded from the analysis. This results in the loss of 14 observations. The final sample used for the analysis comprises 1908 observations.

[Table 1]

Averaged over all teams and years, the coefficients of variation in OPS and ERA in our sample are 0.108 and 0.168, respectively. However, the two measures of team inequality vary widely both across teams and over time. Figure 2 plots the average team coefficients of variation in each season between 1920 and 2009. ERA dispersion fell in the 1920s, likely because of the disappearance of 'spitball' pitchers, but has become more dispersed in recent years, perhaps due to the 1980 and 1990 collusion eras and the expansion of the leagues (Schmidt and Berri 2003; Bradbury 2007). In comparison, dispersion in OPS has been consistently lower and more stable but has declined somewhat since the early 1970s.

[Figure 2]

¹⁴ Among other factors, we expect the year dummies to capture the constraints placed on teams' purchasing decisions by the prevailing wage distribution. This has become much more dispersed since the introduction of free agency in 1977, making it relatively more expensive to afford highly-skilled or experienced players. Indeed, our regression results are qualitatively similar if league-wide average salary and minimum salary are used as instruments in place of year dummies.

4. Results

a. Baseline Regression Results

Column 1 of Table 2 presents estimates using the following equation for team i in year t :

$$WPCT_{it} = \beta_0 + \beta_1 OPS_{it} + \beta_2 ERA_{it} + \beta_3 OPSCOV_{it} + \beta_4 ERACOV_{it} + \varepsilon_{it}, \quad (1)$$

where $WPCT$ is winning percentage, OPS is team OPS, ERA is team ERA, $OPSCOV$ is OPS coefficient of variation and $ERACOV$ is ERA coefficient of variation. As a baseline, the model is first estimated using OLS with no team fixed effects. In column 1 of Table 2 we see that after controlling for average batting and pitching talent of the team, inequality in both ERA and OPS, as measured by the coefficients of variation, is negatively related to team performance. The higher/lower the average OPS/ERA, the greater the winning percentage of the team, indicating that assembling a team with higher average talent positively affects subsequent team performance, as expected.

[Table 2]

b. Is There an Optimal Degree of Heterogeneous Ability?

Table 2 column 2 presents the results when OLS is used to estimate the following equation, where we add the squared terms to Equation 1 to capture any non-linear effects arising from the inequality of team batting and pitching ability:

$$WPCT_{it} = \beta_0 + \beta_1 OPS_{it} + \beta_2 ERA_{it} + \beta_3 OPSCOV_{it} + \beta_4 OPSCOV_{it}^2 + \beta_5 ERACOV_{it} + \beta_6 ERACOV_{it}^2 + \varepsilon_{it}. \quad (2)$$

As before, team OPS has a positive effect on winning percentage and team ERA has a negative effect. Dispersion in ERA among pitchers no longer has an effect on a team's success; however, there appears to be a weak inverse U-shaped relationship between winning percentage and dispersion in OPS. The coefficients indicate that, *ceteris paribus*, a team will win most games when its coefficient of variation in OPS

is 0.062; that is, when the standard deviation of its OPS is 6.2% of the mean OPS. The estimates indicate that raising the OPS coefficient of variation from the mean to one standard deviation above the mean (*i.e.* from 0.108 to 0.135) will result in a team losing roughly 1.3 more games in a 162-game season. Most teams had levels of inequality greater than the estimated optimum, implying that many teams would benefit from a reduction in skill dispersion. This also explains why the linear effect of team inequality is negative in column 1.

c. Persistent Team Winning Legacies

There are a number of reasons to imagine that teams might have persistently high or persistently low winning percentages. Teams have stadia of different dimensions, located at different altitudes, with different climates and featuring different playing surfaces (Schell 2005; Bradbury 2007). Moreover, teams from larger markets are likely to be able to purchase better players and therefore are likely to have higher winning percentages. Teams may also foster cultures of success, which are transmitted to players well after the winning legacies were first established.

To examine this, a full set of team dummies is added to equations (1) and (2) in columns 3 and 4 of Table 2. Comparing columns 2 and 4 we see a marked improvement in the precision of our model. The addition of team fixed effects increases the size of the estimated coefficients and the optimal coefficient of variation in OPS increases to 0.079 from 0.062 in the non-fixed effect estimates. These results are robust to controls for autocorrelation using the method derived by Baltagi and Wu (1999).¹⁵

¹⁵ The Baltagi and Wu (1999) method fits cross-sectional time-series regression models when the disturbance term is assumed to be first-order autoregressive.

d. IV Results

As discussed in the previous section, our performance dispersion measures may be endogenous with respect to winning percentage. To correct for this problem, we present the results of instrumental variable estimation in Table 3. In these specifications, the instruments (discussed in section 3a) are a set of variables that are thought to influence performance dispersion but do not have any direct effect on winning percentage, other than through their effects on performance dispersion.

[Table 3]

In column 2 of Table 3, instrumental variables estimation is used without team effects. Compared with the OLS estimates in column 2 of Table 2, OPS dispersion is now more highly curved with respect to winning percentage and ERA dispersion is now found to be significant and have a similar non-linear relationship. A one standard deviation increase in the OPS coefficient of variation starting from the mean is now found to result in a team losing 1.7 extra games a season. These results are consistent with the proposition that teams pick players at the start of the season to optimise skill dispersion, leading OLS to underestimate the effects of talent inequality.

The regression coefficients indicate that a team will win most games when their OPS coefficient of variation is 0.107 and when the ERA coefficient of variation is 0.192. When team dummies are added to our quadratic IV estimates in column 4 of Table 3, the effects of OPS and ERA dispersion are little affected and the optimal levels of dispersion are the same as in the model without team effects.

The instruments are always jointly significant in the first-stage regressions and each instrument except metropolitan population is generally statistically significant, although the exogenous OPS dispersion change variable has an insignificant effect on ERA dispersion and the exogenous ERA dispersion change variable and lag

attendance have insignificant effects on OPS dispersion. The null hypothesis of the Sargan-Hansen test cannot be rejected, suggesting that the instruments are valid.

In summary, if we assume that OLS underestimates the effect of talent inequality, then our preferred specifications reside in the IV estimates found in Table 3, which indicate that dispersion in OPS and ERA influences a team's level of success, having conditioned on the level of team ability, and that there is an 'optimal' degree of inequality among observed hitting and pitching talent.

e. Robustness Checks

To confirm the robustness of our results we undertake a series of sensitivity tests.¹⁶ By using the coefficient of variation rather than standard deviation, our measure of skill dispersion is already invariant to teams' overall ability levels. However, it is still possible that winning percentage is related to a team's skill dispersion *relative* to other teams during a given season. To examine this possibility, we included year dummies as added controls in our preferred specification (Table 3 column 4) rather than as instruments. The optimal levels of OPS and ERA dispersion were not found to change appreciably, although they indicated a greater penalty from being farther away from these optima.

The way that teams respond to inequality among players may have changed over the nine decades in our sample. Therefore, as a further test of how the relationship between team success and dispersion of team ability varies over time, we interacted the regressors and instruments in our preferred specification with dummies for which half of the sample they came from (*i.e.* 1920-1964 or 1965-2009). As reported in Table 4, the results become more strongly pronounced and more significant for both

¹⁶ Complete results for the robustness and sensitivity tests described in this section are available from the authors on request.

OPS and ERA dispersion in the later sample years (1965-2009), although the coefficients on the quadratic terms are not significantly different between the periods. One possible explanation for the stronger ERA relationships in the modern era could be the effects that changes to pitching styles and the role of relievers have had on pitching performance. The near complete absence of the relief pitcher in the early sample periods may have arbitrarily widened the dispersion of observed talent without really affecting underlying differences in ability. For OPS dispersion, the stronger results in the later period may be a function of the move to more specialised batting roles (especially the introduction of the designated hitter in the American League in the 1970s) and the fact that with more equality in batter talent over recent times (as noted by Bradbury, 2007) those franchises able to grab the high performing outliers (and thus increase dispersion of talent in their team) are increasingly likely to be successful.

[Table 4]

Our results are also found to be robust to the use of different measures of player inequality. If we replace the coefficient of variation of our offensive and defensive ability measures with the standard deviation or Gini coefficient, the results remain qualitatively similar. Using the 90/10 ratio yields an inverse U-shaped effect of OPS dispersion but an insignificant result for ERA dispersion. The results also remained significant when we replaced the career performance measures with ERA and OPS in the *previous* season only.

As noted in a large body of work, the role of salary inequality on team performance has been found to be negative (see Simmons et al., 2009 for a recent summary). To ensure that our baseball findings are robust to measures of payroll inequality, we added average team salary and a quadratic in coefficient of variation in

salary to the specification in column 4 of Table 3. The sample is restricted to the period for which salary data for individual players are available (i.e., 1985-2009). Our results, which are reported in column 1 of Table 5, show that the non-linear effects of ERA and OPS dispersion, though slightly less significant, are very similar when salary dispersion is added as a covariate.

[Table 5]

We replaced our dependent measure of team performance (the win rate) with separate equations for runs scored and runs allowed in columns 2 and 3 of Table 5, respectively. Only the OPS variables were included in the runs scored regression and only the ERA variables were included in the runs allowed regression, with year dummies included in both to capture trends in scoring. Estimating our models this way allows dispersion among hitters to have an effect on team output even if the team's pitchers are poor (meaning many games are lost regardless of offensive performance) and vice versa. OPS dispersion has an inverse U-shaped effect on runs scored and ERA dispersion has a U-shaped effect on runs allowed, indicating that there is an optimal level of dispersion in both offensive and defensive production.

Finally, we examine what happens if we use different measures of player ability for non-pitchers, other than OPS. Table 6 reports the results when batting average (BA), on-base percentage (OBP) and slugging percentage (SLG) are used in turn. The specifications with batting average and on-base percentage feature a significant inverse U-shaped relationship between win percentage and inequality in performance among non-pitchers, however there is no significant relationship when slugging percentage is used.

[Table 6]

5. Extensions and Applications Beyond Baseball

We have found that baseball teams assembled at the start of a season with either too large or too small a degree of inequality in OPS or ERA underperform relative to those teams with more intermediate skill distributions. The distribution of skill appears to matter for both offensive and defensive output. These findings suggest that teams with a healthy balance of stars and players on their way to becoming stars (and perhaps even players entering years of declining productivity but who provide experience) outperform teams at the extreme ends of the skill distribution.

Although our empirical analysis is performed using data on baseball players and their teams, these results are also likely to apply to many work environments where there is an element of joint production. Wherever workers have to perform their tasks in a team setting and there is a single product or ultimate output measure (e.g., consultancy work, academic departments, complex legal cases, film sets, NASA missions, restaurants, etc.), the idea of an ‘optimal’ distribution of worker ability is likely to be relevant. The task of a manager in these settings is not simply to hire individual workers with the best talent money can buy, or to hire a star and allow the rest of a team to catch up. Rather, it is as important to look at the effect that hiring someone will have on the dispersion of ability. In cases where work is highly interdependent and resources for the firm are constrained, it may even be best to look at distributional concerns first and absolute ability second before hiring.

Examples from the entertainment industry and personal computing illustrate this intuition. The Beatles, possibly the most successful pop band in history, benefitted for many years from: having a senior producer (George Martin) overseeing the song writing of two young music talents (John Lennon and Paul McCartney), who were each trying to top one another, a slightly younger and budding talent (George

Harrison) learning from these other two and one drummer (Ringo Starr) who, to quote from a *Saturday Night Live* sketch, ‘was just happy to be there’.

The Beatles’ experience contrasts with the experience of their contemporaries, the Beach Boys, who some have argued produced an album in 1966 (*Pet Sounds*) that was superior to anything the Beatles had released up to that point. However, relying as they did on the musical genius of one band member (Brian Wilson) the Beach Boys were more susceptible to burnout and decline. Indeed, the heavy reliance on one member’s output proved disastrous once he became severely depressed.

Though it cannot be quantified in quite the same the way as baseball performance can, the average musical proficiency of the Beach Boys and the Beatles was arguably the same up to 1966 but the distribution of band ability was decidedly not. The Beach Boys were too reliant on one star whereas the Beatles had a more balanced distribution of musical ability that was spread across the band. In the end, between 1966 and 1970 the Beatles recorded and released some of their best-selling albums whereas the Beach Boys managed only one major hit.

In a similar fashion, the world of personal computing would not have been possible without the collaboration of a large group of untested engineers working together with more seasoned talent at Xerox’s Palo Alto Research Centre throughout the 1970s and 1980s. In his book *The Soul of a New Machine*, Tracy Kidder explains how an upstart group of engineers, made up of a few senior designers working mostly with college graduates, created the new generation of 32-bit computers in the early 1980s by outperforming the original team composed of only ‘elite’ senior designers they were meant to support.

These examples and our findings reinforce what good organisations seemingly do every day: select the best group of workers possible and harness their collective

potential by being as attentive to the distribution of skills as on the average ability. An obvious question is why do so many organisations fail to attain the ideal distribution of talent, like many of the baseball teams in our sample? This may be because organisations face binding budget constraints. Managers of small organisations or teams may simply be unable to acquire the best talent. They may have a workforce with a very similar range of ability, but not enough star talent to pull up overall production. Conversely, in large firms with very few limitations on developing, finding or prying away the best human capital, the surfeit of star talent can often prevent the formation of a well-functioning team.

The other issue facing firms is the scarcity of talent. Assuming there is a distribution of ability (i.e., agents are not perfectly homogenous) a limit exists on the number of talented individuals that can be assembled on any one team at any given time. This forces firms to look at other margins for improvement beyond hiring the superstar talent, such as mentoring workers with less demonstrated ability but with huge potential, employing better human resource techniques and/or acquiring better coaching talent.

6. Conclusion

This paper analyses whether, after conditioning on average skill, team performance is significantly related to the distribution of skills within a team. Using data on professional baseball players, our results show that team winning percentage increases with the dispersion of batting and pitching talent up to a point and decreases thereafter. Hence, there exists an optimal level of heterogeneity in ability. This interpretation is robust to the inclusion of team fixed effects and to controls for potential endogeneity through IV estimation.

One of the major implications of this research is that in settings where output is jointly determined, such as in any team-based enterprise, considerations of skill distribution are likely to emerge. We find evidence of an optimal level of performance dispersion within baseball, even though both offensive and defensive roles require less on-the-field coordination and interaction than other team sports such as basketball, football or hockey. This is because baseball output is still subject to intermediation, which much like assembly line work, depends in large measure on who precedes or follows a player in a pitching rotation or batting order.

The question of the optimal spread of individual ability within any given team, however, will always be decidedly hard to answer and future work would be productively directed at specifying what organisational attributes foster a more optimal spread of talent.

Figure 1
Hypothetical Relationship between Skill Dispersion and Team Performance

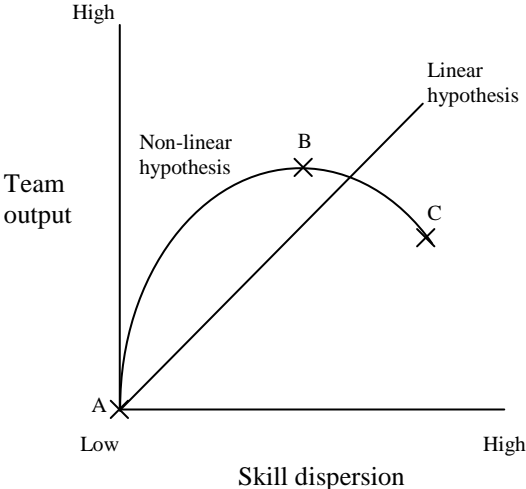
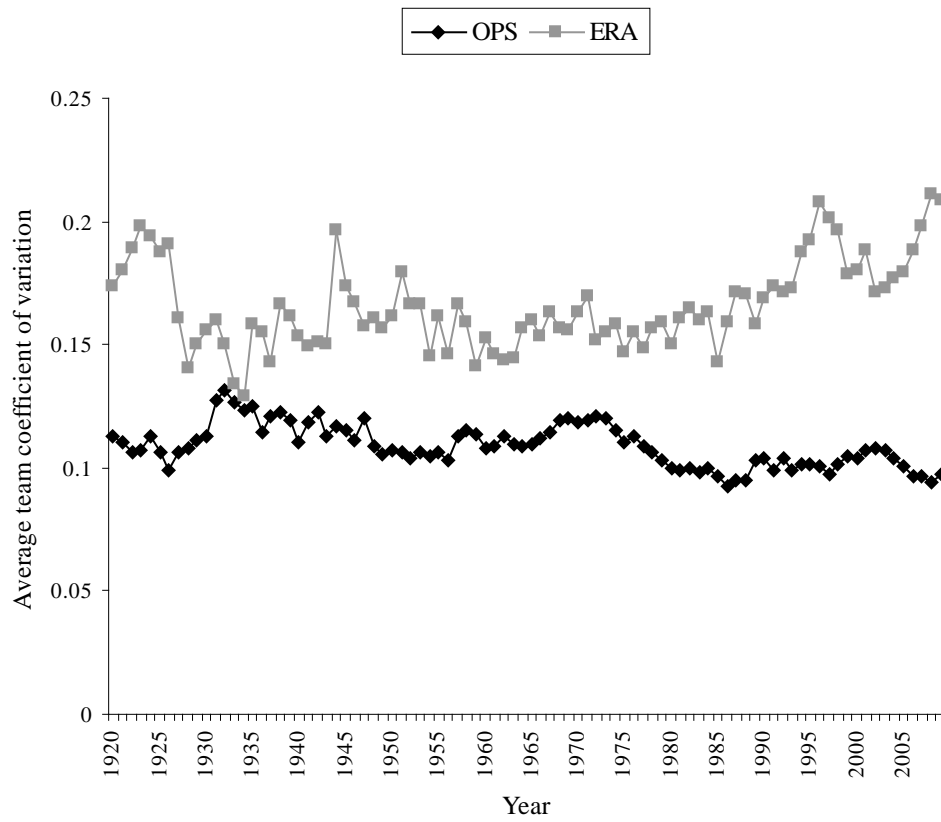


Figure 2
Performance Dispersion over Time



Note: The data points for each year reflect the coefficients of variation in OPS and ERA (as defined in the appendix) averaged across the teams in our sample.

Table 1
Descriptive Statistics for Major League Baseball Teams, 1920-2009

Variable	Mean	Standard deviation	Minimum	Maximum
1. Winning percentage	0.501	0.079	0.248	0.721
2. Team career OPS	0.756	0.042	0.626	0.933
3. Team career ERA	3.842	0.464	2.465	5.760
4. Team OPS coefficient of variation	0.108	0.028	0.034	0.231
5. Team ERA coefficient of variation	0.168	0.053	0.031	0.419
6. Exogenous change in OPS coefficient of variation	0.001	0.010	-0.038	0.091
7. Exogenous change in ERA coefficient of variation	0.012	0.033	-0.117	0.220
8. Lag home attendance	1,444,638	870,568	80,922	4,483,350
9. Metropolitan area population	2,630,713	1,458,968	633,887	10,109,895

Table 2
Distribution of Player Ability and Team Performance, 1920-2009
OLS Regression Results

Dependent Variable: <i>Winning Percentage</i>				
Independent Variables	(i) OLS	(ii) OLS	(iii) OLS	(iv) OLS
1. Team career OPS	0.924*** (0.038)	0.927*** (0.038)	0.853*** (0.040)	0.852*** (0.040)
2. Team career ERA	-0.082*** (0.003)	-0.082*** (0.003)	-0.077*** (0.004)	-0.077*** (0.004)
3. Inequality of team OPS	-0.259*** (0.054)	0.303 (0.295)	-0.271*** (0.056)	0.597** (0.298)
4. Inequality of team OPS squared	–	-2.439* (1.259)	–	-3.772*** (1.272)
5. Inequality of team ERA	-0.092*** (0.028)	-0.178 (0.127)	-0.095*** (0.029)	-0.174 (0.127)
6. Inequality of team ERA squared	–	-0.234 (0.333)	–	-0.216 (0.332)
7. Constant	0.161*** (0.027)	0.135*** (0.033)	0.197*** (0.028)	0.157*** (0.034)
Team effects	No	No	Yes	Yes
Adjusted R-squared	0.328	0.329	0.348	0.348
Number of observations	1908	1908	1908	1908

Note: Standard errors are presented in parentheses. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. Inequality of batting and pitching ability is estimated through the standardised coefficient of variation.

Table 3
Distribution of Player Ability and Team Performance, 1920-2009

IV Regression Results

Dependent Variable: <i>Winning Percentage</i>				
Independent Variables	(i)	(ii)	(iii)	(iv)
	IV	IV	IV	IV
1. Team career OPS	0.935*** (0.042)	0.927*** (0.047)	0.860*** (0.042)	0.846*** (0.046)
2. Team career ERA	-0.083*** (0.003)	-0.086*** (0.004)	-0.078*** (0.004)	-0.080*** (0.004)
3. Inequality of team OPS	-0.326*** (0.118)	2.699** (1.188)	-0.330*** (0.121)	3.069*** (1.143)
4. Inequality of team OPS squared	–	-12.660** (5.183)	–	-14.365*** (4.989)
5. Inequality of team ERA	-0.053 (0.060)	2.228*** (0.532)	-0.061 (0.061)	1.949*** (0.516)
6. Inequality of team ERA squared	–	-5.798*** (1.344)	–	-5.103*** (1.297)
7. Constant	0.157*** (0.028)	-0.198** (0.083)	0.194*** (0.029)	-0.152** (0.077)
Team effects	No	No	Yes	Yes
Number of observations	1908	1908	1908	1908

Note: Standard errors are presented in parentheses. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. Inequality of batting and pitching ability is estimated through the standardised coefficient of variation. The coefficient of variation terms are instrumented for by a set of year dummies, lag team attendance, metropolitan area population and quadratics in exogenous change in OPS coefficient of variation and exogenous change in ERA coefficient of variation.

Table 4
IV Results Interacted by Sub-Period

Dependent Variable: <i>Winning Percentage</i> Independent Variables	Coefficients	
	1920-1964	1965-2009
Team career OPS	0.777** (0.071)	0.789** (0.060)
Team career ERA	-0.088** (0.006)	-0.067** (0.006)
Inequality of team OPS	1.323 (1.717)	2.887** (1.437)
Inequality of team OPS squared	-5.952 (7.345)	-13.618** (6.377)
Inequality of team ERA	0.232 (0.603)	1.499** (0.717)
Inequality of team ERA squared	-1.031 (1.501)	-4.061** (1.810)
Team-period effects	Yes	
Number of observations	1908	

Note: Standard errors are presented in parentheses. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. Inequality of batting and pitching ability is estimated through the standardised coefficient of variation. The coefficient of variation terms are instrumented for by a set of year dummies and the interaction of period with lag team attendance, metropolitan area population and quadratics in exogenous change in OPS coefficient of variation and exogenous change in ERA coefficient of variation.

Table 5
IV Results with Alternative Specifications

Dependent Variable:	Winning percentage	Runs scored	Runs allowed
Independent Variables	(i) IV	(ii) IV	(iii) IV
Team career OPS	0.511*** (0.139)	1253.644*** (75.374)	–
Team career ERA	-0.065*** (0.010)	–	118.506*** (6.488)
Team salary (in \$ million)	0.007* (0.004)	–	–
Inequality of team OPS	4.522* (2.526)	8404.215** (3962.307)	–
Inequality of team OPS squared	-22.329* (12.091)	-36851.480** (17614.390)	–
Inequality of team ERA	1.722 (1.095)	–	-4143.628*** (1341.472)
Inequality of team ERA squared	-4.578* (2.518)	–	10539.050*** (3234.990)
Inequality of team salary	-0.099 (0.126)	–	–
Inequality of team salary squared	0.049 (0.051)	–	–
Team effects	Yes	Yes	Yes
Year effects	No	Yes	Yes
Number of observations	704 [†]	1908	1908

Note: Standard errors are presented in parentheses. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. Inequality of batting and pitching ability and salary is estimated through the standardised coefficient of variation. The coefficient of variation terms are instrumented for by a set of year dummies, lag team attendance, metropolitan area population and quadratics in exogenous change in OPS coefficient of variation and exogenous change in ERA coefficient of variation. † Salary data is not uniformly and publically available until the 1985 season, hence the reduced number of observations associated with the 1985-2009 sample period.

Table 6
Results with Alternative Measures of Hitting Performance

Dependent Variable: <i>Winning Percentage</i>			
Independent Variables	(i) IV	(ii) IV	(iii) IV
Team career BA	1.544*** (0.146)	–	–
Team career OBP	–	1.875*** (0.134)	–
Team career SLG	–	–	1.269*** (0.064)
Team career ERA	-0.063*** (0.004)	-0.064*** (0.005)	-0.084*** (0.004)
Inequality of team BA	4.329** (2.157)	–	–
Inequality of team BA squared	-28.956*** (11.162)	–	–
Inequality of team OBP	–	7.883*** (2.361)	–
Inequality of team OBP squared	–	-43.603*** (12.098)	–
Inequality of team SLG	–	–	1.392 (1.063)
Inequality of team SLG squared	–	–	-4.342 (3.266)
Inequality of team ERA	2.630*** (0.571)	2.548*** (0.561)	1.775*** (0.477)
Inequality of team ERA squared	-6.854*** (1.447)	-6.548*** (0.561)	-4.582*** (1.207)
Constant	-0.046 (0.113)	-0.446*** (0.132)	0.045 (0.088)
Team effects	Yes	Yes	Yes
Number of observations	1908	1908	1908

Note: Standard errors are presented in parentheses. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. Inequality of batting and pitching ability is estimated through the standardised coefficient of variation. The coefficient of variation terms are instrumented for by a set of year dummies, lag team attendance, metropolitan area population and quadratics in exogenous change in hitting measure coefficient of variation and exogenous change in ERA coefficient of variation.

Table A1
Franchises in the analysis

Current team name	Metropolitan area(s)	Years in dataset
Arizona Diamondbacks	Phoenix	1998-2009
Atlanta Braves	Boston	1920-1952
	Milwaukee	1953-1965
	Atlanta	1966-2009
Baltimore Orioles	St Louis	1920-1953
	Baltimore	1954-2009
Boston Red Sox	Boston	1920-2009
Chicago Cubs	Chicago	1920-2009
Chicago White Sox	Chicago	1920-2009
Cincinnati Reds	Cincinnati	1920-2009
Cleveland Indians	Cleveland	1920-2009
Colorado Rockies	Denver	1993-2009
Detroit Tigers	Detroit	1920-2009
Florida Marlins	Miami	1993-2009
Houston Astros	Houston	1962-2009
Kansas City Royals	Kansas City	1969-2009
Los Angeles Angels	Anaheim	1961-2009
Los Angeles Dodgers	New York	1920-1957
	Los Angeles	1958-2009
Milwaukee Brewers	Seattle	1969-1969
	Milwaukee	1970-2009
Minnesota Twins	Washington	1920-1960
	Minneapolis	1961-2009
New York Mets	New York	1962-2009
New York Yankees	New York	1920-2009
Oakland Athletics	Philadelphia	1920-1954
	Kansas City	1955-1967
	Oakland	1968-2009
Philadelphia Phillies	Philadelphia	1920-2009
Pittsburgh Pirates	Pittsburgh	1920-2009
San Diego Padres	San Diego	1969-2009
San Francisco Giants	New York	1920-1957
	San Francisco	1958-2009
Seattle Mariners	Seattle	1977-2009
St Louis Cardinals	St Louis	1920-2009
Tampa Bay Rays	Tampa	1998-2009
Texas Rangers	Washington	1961-1971
	Dallas	1972-2009
Toronto Blue Jays	Toronto	1977-2009
Washington Nationals	Montreal	1969-2004
	Washington	2005-2009

References

- Angrist, J.D., Lang, K. 2004. Does school integration generate peer effects? Evidence from Boston's Metco Program. *American Economic Review* 94 (5), 1613-1634.
- Baltagi, B.H., Wu, P.X., 1999. Unequally spaced panel data regressions with AR(1) disturbances. *Econometric Theory* 15, 814-823.
- Benabou, R., 1994. Human capital, inequality, and growth: A local perspective. *European Economic Review* 38 (3-4), 817-826.
- Benabou, R., 1996. Equity and efficiency in human capital investment: The local connection. *Review of Economic Studies* 63 (2), 237-264.
- Berri, D., Jewell, R., 2004. Wage inequality and firm performance: Professional basketball's natural experiment. *Atlantic Economic Journal*, 32 (2), 130-139.
- Black, D., Henderson, V., 2003. Urban evolution in the USA. *Journal of Economic Geography* 3 (4), 343-372.
- Boisjoly, J, Duncan, G.J., Kremer, M., Levy, D.M., Eccles, J., 2006. Empathy or antipathy? The impact of diversity. *American Economic Review* 96 (5), 1890-1905.
- Bombardini, M., Gallipoli, G., Pupato, G. (2009). Skill dispersion and trade flows. NBER Working Paper No. 15097.
- Bougheas, S., Riezman, R., 2007. Trade and the distribution of human capital. *Journal of International Economics* 73 (2), 421-433.
- Bradbury, J.C, 2007. *The Baseball Economist: The Real Game Exposed*. New York: Dutton.
- Bradbury, J.C., 2010. *Hot Stove Economics: Understanding Baseball's Second*

Season. New York: Springer-Verlag.

Falk, A., Ichino, A., 2003. Clean evidence on peer effects. *Journal of Labor Economics* 24 (1), 39-58.

Freeman, R., Gelber, A., 2006. Optimal inequality/optimal incentives: Evidence from a tournament. NBER Working Paper No. 12588.

Frick, B., 2003. Contest theory and sport. *Oxford Review of Economic Policy* 19 (4), 512-529.

Glaeser, E.L., Resseger, M.G., Tobio, K., 2008. Urban inequality. NBER Working Papers 14419.

Gomez, R., Budd, J., Meltz, N., 2004. Why a balance is best: the pluralist industrial relations paradigms of balancing competing interests. In Kaufman, B.E. (ed.), *Theoretical perspectives on work and the employment relationship*. Cornell University Press: Ithaca, 195-228.

Gould, E.D., Kaplan, T.R. (2010). Learning unethical practices from a co-worker: The peer effect of Jose Canseco. *Labour Economics*, forthcoming.

Gould, E.D., Lavy, V., Paserman, M.D., 2004. Immigrating to opportunity: Estimating the effects of school quality using a natural experiment on Ethiopians in Israel. *Quarterly Journal of Economics* 119 (2), 489-526.

Gould, E.D., Lavy, V., Paserman, M.D., 2009. Does immigration affect the long-term educational outcomes of natives? Quasi-experimental evidence. *Economic Journal* 119 (540), 1243-1269.

Gould, E.D., Winter, E., 2009. Interactions between workers and the technology of production: Evidence from professional baseball. *Review of Economics and Statistics* 91 (1), 188-200.

- Hamilton, B.H., Nickerson, J.A., Owan, H., 2003. Team incentives and worker heterogeneity: An empirical analysis of teams on productivity and participation. *Journal of Political Economy* 111 (3), 465-497.
- Hoxby, C.M., 2000. Peer effects in the classroom: Learning from gender and race variation. NBER Working Paper No. 7867.
- Ichino, A., Maggi, G., 2000. Work environment and individual background: Explaining regional shirking differentials in a large Italian firm. *Quarterly Journal of Economics* 115 (3), 1057-1090.
- Jackson, P., 2004. The Soul of Teamwork. *EnlightenNext Magazine*, May-July 2004. www.enlightennext.org/magazine/j25/teamwork.asp, accessed on March 10, 2010.
- Kandel, E., Lazear, E.P., 1992. Peer pressure and partnerships. *Journal of Political Economy* 100 (4), 801-817.
- Kendall, T.D., 2003. Spillovers, complementarities, and sorting in labor markets with an application to professional sports. *Southern Economic Journal* 70 (2), 389-402.
- Kidder, T., 2000. *The Soul of a New Machine*. San Francisco: Back Bay Books.
- Lavy, V., Silva, O., Weinhardt, F., 2009. The good, the bad and the average: Evidence on the scale and nature of ability peer effects in schools. NBER Working Paper 15600.
- Lazear, E., 1995. *Personnel Economics*. Cambridge, MA: MIT Press.
- Mas, A., Moretti, E., 2009. Peers at work. *American Economic Review* 99 (1), 112-45.
- Ohnsorge, F., Trefler, D., 2007. Sorting it out: International trade with heterogeneous workers. *Journal of Political Economy* 115 (5), 868-892.

- Oreopoulos, P., 2003. The long-run consequences of living in a poor neighborhood. *Quarterly Journal of Economics* 118 (4), 1533-1575.
- Ottaviano, G., Peri, G., 2006. The economic value of cultural diversity: Evidence from US cities, *Journal of Economic Geography* 6 (1), 9-44.
- Sacerdote, B., 2001. Peer effects with random assignment: results for Dartmouth roommates. *Quarterly Journal of Economics* 116 (2), 681-704.
- Schell, M.J., 2005. *Baseball's All-Time Best Sluggers: Adjusted Batting Performance from Strikeouts to Home Runs*. Princeton, NJ: Princeton University Press.
- Schmidt, M., Berri, D., 2003. On the evolution of competitive balance: The impact of an increasing global search. *Economic Inquiry* 41 (4), 692-704.
- Simmons, R., Kahane, L., Longley, N., 2009. The effects of coworker heterogeneity on firm-level output: Assessing the impacts of cultural and language diversity in the National Hockey League. Working Paper 005934, Economics Department, Lancaster University Management School.
- Simmons, R., Frick, B., 2008a. The allocation of rewards in athletic contests. In Howard, D.R., Humphreys, B.R. (eds), *The Business of Sports*, vol. 2. Westport, CT: Praeger Publishers.
- Simmons, R., Frick, B., 2008b. Pay and performance of players in sports leagues: international comparisons. In Howard, D.R., Humphreys, B.R. (eds), *The Business of Sports*, vol. 2. Westport, CT: Praeger Publishers.
- United Nations, 1998. *World Urbanization Prospects: The 1999 Revision*. New York: Population Division, Department of Economic and Social Affairs, United Nations.
- Winter, E., 2004. Incentives and discrimination. *American Economic Review* 94 (3), 764-773.

Appendix 1: Construction of variables

Earned run average is the number of earned runs a pitcher (as identified from the Baseball Archive dataset) concedes divided by the number of innings he pitches multiplied by nine. On-base plus slugging is the sum of a non-pitcher's on-base percentage (defined as times-on-base per plate appearance, excluding sacrifice hits, fielder's obstruction or catcher's interference) and his slugging percentage (defined as total bases scored per at-bat). We calculate these over a player's entire career, up to and including the previous season. OPS is then averaged over all the non-pitchers on a team, weighting each player by his number of at-bats in the previous season multiplied by the sum of his number of at-bats, walks, sacrifice flies and hit-by-pitches in the previous season. ERA is averaged over all the pitchers on a team, weighted by the number of innings pitched in the previous season.

The coefficients of variation in OPS and ERA are calculated similarly: the weighted standard deviation in each variable is calculated for each team in each season and this is divided by the weighted averages described above. Only pitchers who pitched more than 20 innings in the previous season and non-pitchers with more than 20 at-bats in the previous season are included in the calculation of the coefficients of variation.

The exogenous dispersion change variables are defined as the difference between a team's coefficient of variation in OPS or ERA across its players in the *previous* season (using their career performance values at the end of that season) and its coefficient of variation excluding those players who left the team for reasons out of a franchise's control before the beginning of the current season. Players may leave a team without the full consent of the team if they retire, become injured, are selected in an expansion draft or leave a team as a free agent. Retirement or permanent injury is defined as being not observed in any future season.

Expansion drafts were held at the end of the 1960, 1961, 1968, 1976, 1992 and 1997 seasons to stock the rosters of the teams joining the major leagues the following seasons. The exact rules varied from year to year, but in each case existing teams were able to protect a certain number of players from the draft; expansion teams could then select from the remaining players. In the drafts prior to 1992, the new teams were only allowed to select players from other teams in their own league. In the 1992 and 1997 drafts, teams could select players from any other major league team.

Free agency was introduced in 1977 and means that any player who is out of contract and has completed six full seasons of major league service is free to sign with any team he wishes. Since the dataset does not record exact service time (which includes time spent injured but still on a team's roster), a player is considered to be a free agent after playing in part of any six seasons.

Average attendance was taken from the Baseball Archive. The population data refer to the population of the primary metropolitan statistical area a team is located in, as listed in Table 1. Metropolitan population refers to a consistent set of primary metropolitan statistical areas constructed by Black and Henderson (2003). These are calculated from the decennial U.S. Censuses for 1920-2000, with comparable data for the same years for Toronto and Montreal taken from United Nations (1998). For non-Census years, the data were interpolated for interior years and extrapolated for 2001-2009, based on the 1990 and 2000 values. In cases where there is more than one team in a metropolitan area, such as New York and Chicago, the population is divided by the number of teams to give a crude measure of a team's effective fan base.