Public Choice
https://doi.org/10.1007/s11127-021-00884-z

Informed voters and electoral outcomes: a natural experiment stemming from a fundamental information-technological shift

Shane Sanders[1] · **Joel Potter**[2] · **Justin Ehrlich**[3] · **Justin Perline**[4] · **Christopher Boudreaux**[5]

Received: 6 February 2020 / Accepted: 27 January 2021

Abstract
Do informed electorates choose better candidates? While that question is straightforward, its answer often is elusive. Typically, candidate-quality information is neither salient nor subject to exogenous change. We identify a natural experiment within a non-political election setting that is transparent and features exogenous change in the candidate-quality information frontier. The setting is *Major League Baseball's* (*MLB*) annual selection of two most valuable players, a challenging environment with an innately heterogeneous candidate set, and the exogenous change is the development of the pathbreaking, comprehensive player-value measure *Wins Above Replacement* (*WAR*) in 2004 and its subsequent calculation for all retrospective *MLB* player-seasons. *WAR*'s development and rapid popularization informed voting from 2004 onward. Retrospective calculation allows us to draw back the curtain and evaluate how pre-2004 voters behaved with respect to revealed candidate quality. From negative binomial, fixed-effect regression models, we find robust evidence of significant, substantial, pivotal behavioral change on the part of voters since 2004.

Keywords Voter information · Voting behavior · Candidate quality · Information frontier shift · Wins above replacement · Major league baseball

JEL Classification D72 · D80 · Z29

✉ Shane Sanders
sdsander@syr.edu

[1] Syracuse University, 301 MacNaughton Hall, #316, Syracuse, NY 13244, USA

[2] University of North Georgia, 105 Barnes Hall, Dahlonega, GA 30597, USA

[3] Syracuse University, 300 MacNaughton Hall, #320, Syracuse, NY 13244, USA

[4] Pittsburgh Pirates, Pittsburgh, PA, USA

[5] Florida Atlantic University, 777 Glades Road, Kaye Hall 145, Boca Raton, FL 33431, USA

Published online: 13 February 2021

1 Introduction

Do informed electorates choose better candidates? While that question is straightforward, its answer typically is elusive in large political elections. Under normal circumstances, the level and quality of information possessed by voters is neither salient nor subject to exogenous change. Herein, we identify a natural experiment within a non-political election setting that is transparent and features an exogenous change in the technology or frontier of candidate quality information. The setting is voting in the selection of *Major League Baseball's American League* and *National League Most Valuable Players* (*MVPs*), and the exogenous change is the development of the comprehensive player value measure *Wins Above Replacement* (*WAR*). The WAR metric represented an exogenous change because the measure was not created for the purpose of electing *MLB*'s two *MVPs*. Rather, it was created during the 2004 offseason for the de facto purpose of informing popular "sportometric" (Goff and Tollison 1990) research on baseball. However, it had a tertiary effect of making candidate information salient.

Despite its non-political setting, the present study focuses on a transparent election of substantial financial significance.[1] In that environment, we consider the role of information in addressing the question of whether informed electorates choose better candidates (Banerjee et al. 2011; Pande 2011). The study's setting contains several generalizable elements. For example, we are able to consider voters' response to changes in overall candidate quality and also their response to changes in constituents of candidate quality. We also consider the extent to which voters respond to non-meritorious candidate factors. For example, it may be that voters respond to information about candidates' races, ages, or teammate quality, even conditional on the knowledge of specific candidate quality. In the present setting, *MVP* candidates produce baseball value in a team setting. Therefore, we also consider whether voters respond to teammate quality, conditional on the candidate's own quality. That possibility has analogs to political elections in that political candidates typically campaign as members of a political party. Each of those considerations is examined empirically both in general and also conditional on changes in the candidate quality informational frontier that forms the paper's natural experimental setting. Through its focal points, the study seeks to understand generalizable aspects as to the behavior of newly informed voters.

The present study represents a coordination problem and not so much a problem of convergence to true player value.[2] In particular, we consider whether voters coordinated on different salient proxies for player value before the creation of *Wins Above Replacement* (*WAR*) than afterwards.[3,4] *WAR* provides a refined estimate to answer the following question (Slowinski 2010): "If this player got injured and [his] team had to replace [him] with

[1] Some teams adopt *MVP* awards as an incentive mechanism. Some players are offered large contractual bonuses for winning being selected as an *MVP*. For example, star *MLB* player Mike Trout earns an additional $500,000 by his team, the Los Angeles Angels, in the event of winning an *MVP* award.

[2] We thank an anonymous editor for pointing this out.

[3] The formula for position players is WAR = (Batting Runs + Base Running Runs + Fielding Runs + Defensive Runs Saved + Position Adjustment + Replacement Runs)/(Average Runs Needed to Obtain an Additional Win).

[4] The formula for pitchers is WAR = [(RA9avg—RA9)*(IP/9) + Rlr]/(Rpw) where RA9avg is the number of runs an average pitcher is expected to give up in 9 innings, RA9 is the average number of runs given up per nine innings, IP is the number of innings pitched, Rlr is replacement level runs, and Rpw is the average number of runs needed to obtain an additional win.

a freely available minor league player, how much value would the team be losing?". In the case of a player with a *WAR* value of 5, replacing that player with a minor league player would result in an estimated five games lost to the team, *ceteris paribus*. As such, *WAR* is designed to measure a player's performance independent of his teammates.[5] Following the development of *WAR*, the measure may represent only one of multiple inputs that voters consider in selecting a most valuable player. For example, some voters may consider both *WAR* and other performance measures, such as a player's win pivotality in *close* games or a player's win pivotality in *important* games, when estimating player value.

1.1 MLB's election(s): a brief institutional background

Major League Baseball was formed as a confederacy of two leagues—the National League (*NL*; 1876-) and the American League (*AL*; 1901-)—in 1903. Since 1911, the *NL* and *AL* each have chosen distinct *Most Valuable Players* following every regular season. Beginning in 1931, the MVP awards have been selected by the *Baseball Writers' Association of America* (*BBWAA*). The *Baseball Almanac* (2019) summarizes the award's early history:

> There have been three different official *Most Valuable Player* awards in Major League Baseball history, since 1911; the Chalmers Award (1911–1914), the League Award (1922–1929), and the Major League Baseball *Most Valuable Player* Award [1931-]. The … *MVP* is presented annually by the *BBWAA*. It is considered by *MLB* as the only official *Most Valuable Player* Award and symbolizes the pinnacle of a player's personal achievement during any single season of play.

In 1938, the *BBWAA* began electing *MVP*s on the basis of voting by its members. Initially, three *NL* (*AL*) award voters were designated for each *NL* (*AL*) team city. That number was reduced to two in 1961. For several years, then, selections have been made by 60 *MLB* voters, 30 participants for each league's election. The voting rule employed is a weighted, truncated Borda method of marks that has also been called a corner-weighted version of Borda's rule (see, e.g., Greenwich et al. 2019). Studies have examined the weighted, truncated Borda method with applications to Formula One racing and ski jumping competitions (Kaiser 2019). When the Borda count is applied to elections, all voters rank candidates or policy alternatives from most preferred to least preferred. Typically, the lowest ranked alternative is assigned 1 point, the next lowest 2 points, and so on until the highest ranked alternative receives n points (Kaiser 2019). Here, n is the number of candidates involved in an MVP election. Points are tallied for each candidate, and the candidate with the largest total wins.

In *MLB*, players receive *MVP* votes much like candidates in a political election. However, the voting procedure adopted for selecting a winner is a weighted truncated Borda count rather than the classic Borda method of marks described above.[6] It is weighted because the point intervals between different consecutive pairs of ranked options are not the same; it is truncated because only a subset of all alternatives is evaluated (i.e., other, lower ranked candidates receive zero marks). Each voter ranks 10 players who earn points

[5] A teammate effect is, however, possible. For instance, if a good player has an even better player hitting behind him in the batting order, it is plausible the first batter will receive better pitches to hit and would therefore produce more runs and wins (which is not captured by *WAR*).

[6] Young (1995) discusses Borda voting and its implications for the likelihood of identifying the "correct" (Condorcet) top candidate should one exist.

according to the following allocation: (14, 9, 8, 7, 6, 5, 4, 3, 2, 1). The points for all players are summed to determine the final score. Under that system, the player-candidate with the highest vote-point total for a given league-year is crowned league-year *MVP*. Under MLB's procedure, a total of $(14 + \sum_{i=1}^{9} i) \times 30 = 1770$ points are available for a given league's race and one player can receive as many as $(30 \times 14) = 420$ points.[7]

1.2 Informed voting and the development of WAR: a natural experiment?

Baseball is a game of specialization such that cross-positional value comparisons have an apples-to-oranges quality, especially in the absence of a comprehensive, cross-positional measure of contributions to winning games. Consider a comparison between a starting pitcher and a shortstop. The pitcher attempts to throw the ball past the batter and either strike the batter out or force an out with the help of his defending teammates. Hence, the pitcher's position is almost purely defensive. However, the shortstop (i.e., the player between 2nd and 3rd base) plays substantial roles on defense (i.e., catching the ball and stopping the opposing player from advancing) and offense (i.e., hitting the ball). The shortstop accumulates win value by mixing his offensive and defensive skills, but the pitcher accumulates value almost exclusively by throwing baseballs and does so by taking the pitching mound about once every six days. Without adopting some exchange rate mechanism, it is difficult to compare the two players' respective contributions. And yet, MVP voting demanded just such a comparison for decades. In 2004, *WAR* was created by *Baseball Prospectus* writer Jay Jaffe (2004) as a cross-positional measure of player win value (above league replacement level value at player position). From 2004 on, *WAR* has been calculated for all *MLB* player-seasons by *Baseball Prospectus* and other baseball publications. What is most important, it also has been calculated retrospectively for the universe of professional baseball. *Baseball Reference* (www.baseball-reference.com) features retrospective calculations beginning with the first season of the first professional baseball league (*National Association of Professional Baseball Players*, 1871–1875) and continuing through the respective histories of the *NL* (1876-) and *AL* (1901-). Those histories were archived using the massive, crowd-sourced data collection project initiated by retrosheet.org, which was created by University of Delaware biology professor David Smith in 1989 (Smith 2000).

The concurrent assembling of past and present *WAR* data since 2004 has created something of a natural experimental setting. Since 2004, *MLB's MVP* voters have cast ballots largely with knowledge of player *WAR* values. Before 2004, voters were comparatively uninformed. Despite that lack of information, retrospective *WAR* values subsequently became available for past player seasons such that we can evaluate both earlier and later *MVP* voting behavior with respect to *WAR* values. Given *WAR*'s prominence (on, e.g., baseballprospectus.com, fangraphs.com, and baseball-reference.com) and the institutional nature of the *BBWAA* as a defined group with a voluminous archived website, as well as regular chapter and national meetings, *Baseball Writers* who now vote for *MLB* Awards every year presumably are better informed voters than their pre-2004 counterparts.[8]

7 In our dataset, the *NL* and *AL* were imbalanced in terms of numbers of teams from 1998 (when the Milwaukee Brewers moved to the *NL*) through 2012 (after which the Houston Astros moved to the AL). That imbalance affected the number of ballots cast for each league during those years. We account for it empirically by re-scaling ballots to a maximum of 420 possible points for each player during those seasons.

8 Many of the counterparts are earlier versions of the same person.

Baseball Writers refer frequently to *WAR* (see, e.g., Madden 2017, *NY Daily News* and Slowinski 2010 for summaries of the measure's importance to *Baseball Writers*).

1.3 Literature and motivation

The present study takes advantage of the development and retrospective calculation of *WAR* as a natural experiment by which to ask a question that Banerjee et al. (2011) considered previously in a starkly different voting context. Specifically, we ask whether informed voters make better choices in the presence of objective information about candidate quality. Whereas fans, participants, and voters may not want the *MVP* race to boil down to a contest of highest *WAR*, neither are they likely to wish for the race to be too far removed from considerations of objective player values. While *WAR* is the central measure of players' contributions to winning baseball games, voters also may consider alternative dimensions of player win value, such as the player's specific contribution toward winning in close or important games. *WAR* considers how a player's productivity relates to team winning given average game settings and therefore does not weight player performances in winning pivotal game scenarios. Banerjee et al. (2011) conducted field voting experiments in India and found evidence that non-partisan, third-party public disclosure of incumbent legislator report cards led to larger vote shares for high-performing incumbents. Within the present voting context, the *WAR* measure can be thought of as similar to a non-partisan, third-party evaluation of candidates. It evaluates candidates in a way that imposes no a priori subjective criteria and is "computationally agnostic" to the contest's results. Pande (2011) reviews the literature and concludes the findings of Banerjee et al. (2011) to be robust across different voting environments. Baron (1994) constructs a model of electoral competition with informed and uninformed voters. He finds that informed voters have the effect of reducing the legislative influence of interest groups in equilibrium as well as the overall level of campaign financing by interest groups.

Pietryka and DeBats (2017) examine the relationship between voting behavior—including voter information—and proximity to elites in one's social network. Ashworth and De Mesquita (2014) demonstrate within a set of formal models that the relationship between voter information and democratic performance may be positive or negative depending on a set of parametric values. Thoth and Chytilek (2018) find that voters facing time pressure shift their information-gathering efforts from accuracy to efficiency when evaluating candidates. Specifically, voters restrict their attentions to a smaller set of policies in those evaluations. MLB's voters may face time pressures in that the performance differences of a season's top players often are slight and subtle and may require viewing hundreds of hours of game footage to perceive. In a given season, a player could be on the roster of any of the 30 *MLB* teams, each of which plays a 162-game schedule. MLB schedules 2,480 games in a season and the average game runs a little over three hours, according to *Baseball Reference*. As such, approximately 7,300 h—more than 304 24-h days—of regular season games are played every year. Faced with that massive quantity of game performances, time-friendly measures to evaluate player-candidates likely are necessary for voters. Several other contributions to the relevant literature study alternative dimensions of information and voting (see, e.g., McMurray 2015; Jensen et al. 2015; Garmann 2017; Ginzburg 2017).

MLB and sports leagues in general are well-established in the literature as laboratories for public choice, economics of regulation, and labor economics study (see, e.g., Shughart 1997, antitrust legislation related to baseball; Ross and Dunn 2007, income tax responsiveness of *MLB* All-Stars; Goff et al. 1998, moral hazard in baseball; Hill and Groothius 2001, redistribution of rents among teams; Ehrlich et al. 2020, labor policies during pandemics).[9] As both "America's Pastime" and a sport that lends itself to statistical analysis, *MLB* provides a long history of labor productivity data that is unparalleled in many respects within American industrial enterprises and even among other popular American sports.

The remainder of the paper is structured as follows. Section 2 presents the dataset used in our study. Controlling for a set of team and season fixed effects negative binomial vote-count regression models,[10] Sect. 3 specifies and tests a model of informed voting. Specifically, the model asks whether *MLB* elections from 2004 relate more strongly to an objective, comprehensive, cross-positional measure of on-field value (i.e., *WAR*) than do prior elections dating from 1980. That is, are informed voters making choices that allow them to identify more accurately the true "*Most Valuable Player*" (rank-ordering) for each league-season? In the absence of information on players' *WAR*, we consider the factors that may have been more salient for evaluating MVP candidates before 2004. Section 4 presents and discusses the study's central results. Section 5 concludes.

1.4 Data description, summary and visualization

We collected data on MLB's top 50 seasonal *WAR* leaders for every regular season from 1980 to 2017. The sample comprises 38 vote-years and 76 league-level elections. We rely on the *Baseball Reference* version of *WAR* for our study rather than the *Fangraphs* version, the *Baseball Prospectus* version, the *openWAR* version, or another implementation. While all *WAR* measures generate positively correlated sets of *WAR* values of moderate to high strength (see, e.g., Baumer et al. 2015), slight methodological differences exist across the implementations; those differences are beyond the scope of the present study.[11] We chose *Baseball Reference WAR* because it is popular, accessible, and was created with retrospective analysis in mind.[12] Allowing for possible ties in seasonal *WAR* value (at the fiftieth-highest value), the sample contains 1907 player-season observations rather than 1900. Each *MLB* voter is charged with identifying and ordering ten *Most Valuable Players* in a given league; *Baseball Writers* typically exhibit strong consensus as to whom should be on a given ballot. For example, the union of all AL ballots in the year 2000—approximately midway through our sample—contained 19 AL players and 22 *NL* players.

The present dataset was retrieved, with permission, from baseball-reference.com using the RVest package in the statistical software program *R*. The primary variables at the player-season level include[13]: *(election) vote point count, WAR value, age, whether on a team in a big (top-5) market, whether traded during season, (primary) league of play, (primary) playing position and (primary) team-of-play*. Summary statistics for

[9] More generally, sport often is used to measure the empirical incidence of aggregation paradoxes. For a recent contribution along those lines, see, e.g., Boudreau et al. (2018).

[10] Negative binomial models are chosen over Poisson models owing to strong evidence of overdispersion in vote-count data.

[11] All WAR measures use the same inputs but slightly different specifications.

[12] Specifically, the measure was designed for description rather than for forecasting. *Baseball Reference* Founder and CEO Sean Forman confirmed that point in a conversation with two of the present authors. Of course, retrospective analysis of the recent season is an input into *MVP* decision-making.

[13] We do not control for game characteristics because it is unlikely that they will affect an individual's WAR score and *MVP* votes. The reason is that because if a player is a season's WAR leader, he always has participated in a large number of games ($n > 100$). Over the course of a season, we therefore expect game characteristics to revert to the mean against an assortment of different opponents.

Table 1 Summary statistics of key variables

Variable	Obs	Mean	St dev (s)	Min	Max
Votepts	1907	54.33	92.68	0	420
WAR	1907	5.60	1.29	3	12.7
Age	1907	28.30	3.69	19	44
Big market	1907	0.30	0.46	0	1
Multiple teams	1907	0.02	0.15	0	1
American league (AL)	1907	0.50	0.50	0	1
Pitcher (P)	1907	0.31	0.46	0	1
Catcher (C)	1907	0.05	0.21	0	1
First base (1B)	1907	0.10	0.30	0	1
Second base (2B)	1907	0.07	0.26	0	1
Shortstop (SS)	1907	0.07	0.26	0	1
Third base (3B)	1907	0.11	0.31	0	1
Right fielder (RF)	1907	0.09	0.29	0	1
Center fielder (CF)	1907	0.11	0.31	0	1
Left fielder (LF)	1907	0.08	0.27	0	1
Designated hitter (DH)	1907	0.01	0.12	0	1

these variables are provided in Table 1. Note that *vote point count* is rescaled from 1998 through 2012 to account for league imbalance.

Player age in the sample ranges between 19 and 44, which is consistent with most of the overall age range for *Major League Baseball*. Incredibly, the average *WAR* value in the sample is 5.60. That value suggests that *MLB* stars generate exceptional average win value. A team of replacement players commonly is estimated to win about 48 games in a season, *ceteris paribus* (*Baseball Reference* 2012). Using that benchmark, we can consider the implied success of a (25-player) *MLB* team in which players average a *WAR* value of 5.60. In fact, such a high level of win production could not exist on one team (owing to crowding out effects). A single team would have to win $48+(5.60\times 25)=188$ games to support such a *WAR* average. However, only 162 MLB games are played in regular season. The sample is balanced in terms of league representation. In the dataset, playing *position* is represented by a set of *position* dummy variables, {P, C, 1B, …, DH}. We observe substantial variation in sample representation by playing *position*, an observation that stands to reason because (players at) one position innately are more valuable than (players at) others. We observe that more than 28% of sampled player-seasons were conducted in the service of a team in a city of top-5 market size (represented as the variable *big market*). Eight of 30 (26.7% of) *MLB* teams were located in a top-5 market throughout the sample; it therefore appears that top baseball talent does not gravitate disproportionately to MLB's biggest markets despite the absence of a salary cap in *MLB*. We base our market-size variable on the US metropolitan areas with the five largest populations according to the 1980, 1990, 2000, and 2010 US Censes. Surprisingly, the five metropolitan areas on that list did not change from 1980 to 1990, to 2000, or to 2010. In the estimation section, we will consider the distinct question as to whether being in a "big market" helps a player in terms of MVP voting points. As that variable is related to team fixed effects, the two will be treated as empirical substitutes in that section. Lastly, we observe that players represented *multiple teams* in 43 player-season pairs.

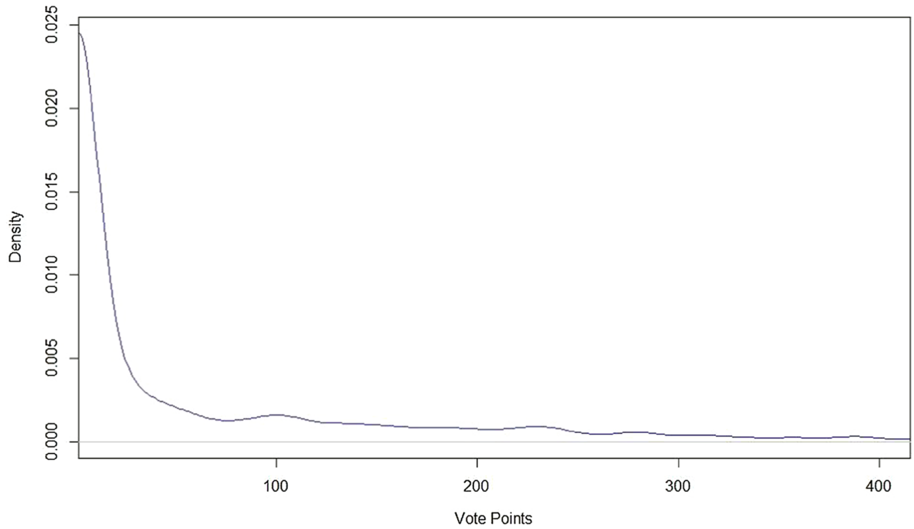

Fig. 1 Kernel density plot of vote points

The sample contains player-season level vote counts, but not voter-level (ballot-level) observations. Ballot-level data were not disclosed by the *BBWAA* until 2018. As such, we treat aggregate player *vote point counts* as our main left-hand side variable. While voter characteristics largely are unavailable to the public, we note that all active members of the *BBWAA* are eligible to vote. As such, the base of voters for the MVP award is quite persistent over time. To verify that conclusion, we examined a subset of individual ballots that have been made public since 2012 (but unfortunately not before the creation of *WAR*) at https://bbwaa.com/voter-database/. Using single letter search queries, we sampled the user database randomly to identify 40 *BBWAA* members who voted for the MVP awards at least once during the observation period. Of those 40 members, 35 had voted in multiple elections during the period of our sample. While we cannot control for individual voter characteristics, the observed persistence in the set of voters, which is an institutional characteristic, suggests that voter characteristics also are persistent across time. We take advantage of a long longitudinal dataset with substantial variation in player characteristics. We also benefit from the natural experimental context of *WAR*'s development (with subsequent retrospective calculations). From our 1907 player-seasons, 1125 candidates for MVP received at least one vote point. With complete consensus across ballots (as to whom should receive a vote), 760 players would have received points. With no consensus (with respect to players in the dataset), all 1907 sampled player-seasons would have received votes. Thus, we observe that *Baseball Writers* exhibit moderately strong consensus as to whom should be on the ballot(s). The present study considers whether such a level of consensus was formed among informed voters or formed among uninformed voters (e.g., an echo chamber) over time. Figure 1 displays a kernel density plot of the *vote point count* within our sample.

Figure 1 depicts a smoothed representation of a discrete variable (*vote point count*). The figure suggests that, even among seasonal *WAR* leaders, most player-seasons result in a small number of vote points. The median *vote point count* in the sample is 4, and the 75th percentile is 64 points. Given our model's specification (i.e., of a count data model), we account for the extreme right skewness of the dependent variable, its discrete

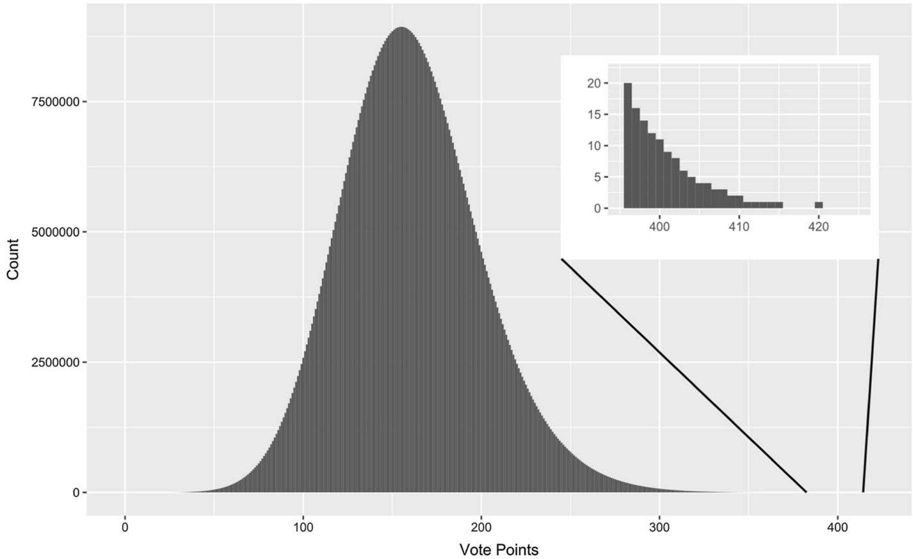

Fig. 2 Distribution of possible player vote point counts

distribution, and its left truncation (at zero). *Vote point count* is a count variable in that (a) it is non-negative integer valued, (b) it in fact tabulates how many points voters allocate to a given player (from their bundle of points), and (c) the possible points allocated to a player represent a set of consecutive, non-negative integers along with one possible non-consecutive value obtained by a player sweeping all first place votes. That is, $Vote\ Point\ Count_i = \left\{ v_i \in N : 0 \le v_i \le 415 \cup v = 420 \right\}$. To demonstrate that *vote point count* occupies that co-domain, Fig. 2 displays a histogram of the distribution of possible *vote point counts* for an individual voter (i.e., from the set of all possible 30×1 voting vectors for a given voting individual), where the frequency count on the *y*-axis represents the number of possible vectors yielding the corresponding number of vote points.

Before formally modeling the relationship of interest, let us consider a scatter plot with player-season *vote point count* on the vertical axis and player-season *WAR* value on the horizontal axis. In Fig. 3a, we adopt color-coded data points and corresponding color-coded trend lines to depict the uncontrolled relationship between the two variables both before and after the creation of *WAR*, respectively.

The uncontrolled relationship between *vote point count* and *WAR* was positive both from 1980 through 2003 and from 2004 through 2017. However, simple regression trend lines suggest that the responsiveness of *vote point count* to changes in *WAR* noticeably has been stronger in the latter period. Multivariate regression analysis will inform us as to whether this apparent difference is significant and substantial conditional on a set of control variables. When considering Fig. 3a, one might wonder whether the apparent slope difference is indeed explained by the creation of *WAR* or if it might suggest a more general (gradual) improvement in voters' knowledge and ability over the course of the dataset. We are dealing herein with panel data having multiple observations per period. Therefore, an endogenous structural break estimation is not supported. However, we do in fact know when the information technology breakpoint of interest occurred (a few months after the 2003 vote), such that an exogenous structural break estimation is appropriate. Moreover, we know that

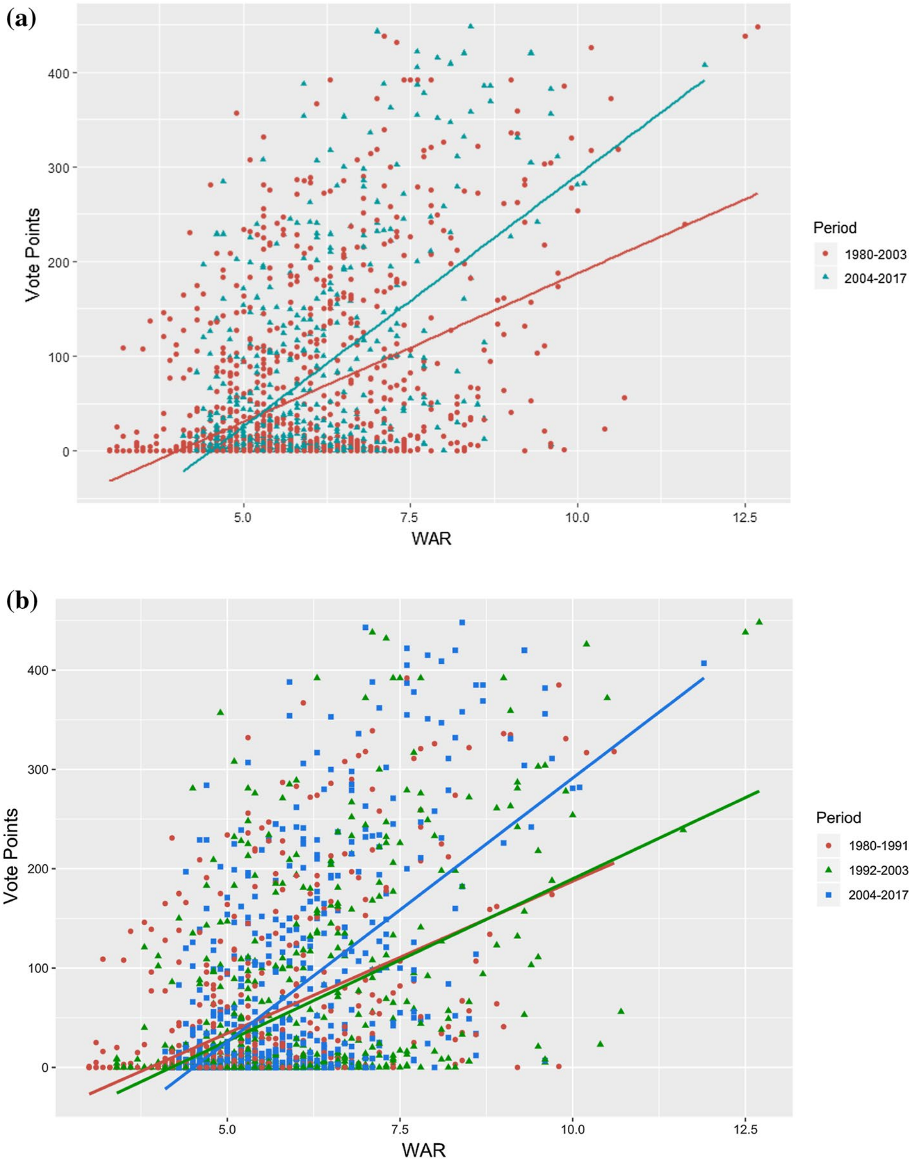

Fig. 3 **a** Plot of MVP vote points against WAR for 2 time periods (before and after information technological shift). **b** Plot of MVP vote points against WAR for 3 time periods (before information technological shift in 2 periods and then after)

WAR is the first and only virally adopted cross-positional player value measure (i.e., the first such measure to be featured on leading baseball statistics sites Baseball Prospectus, Baseball Reference, ESPN, and Fan Graphs). As such, *WAR* was not established in a sea of comparable performance measures. Rather, it was a groundbreaking measure in the area of comprehensive sabermetric player analysis. Baumer and Matthews (2014) state in an article subtitled *There is No Avoiding WAR*:

> While there have been many important contributions to the field, arguably the most prominent success story in recent years has been wins above replacement (*WAR*). *WAR* is an all-encompassing assessment of a baseball player's contribution to his team, measured in wins added relative to a hypothetical (and often vaguely defined) "replacement player." ….While predecessors of *WAR* (like Bill James's Win Shares) have been around for some time, the modern incarnation of *WAR* may now reach a broader audience than any sabermetric stat since OBP. Perhaps the most telling indication of *WAR*'s permeation of the baseball landscape was the announcement that *WAR* will appear on the back of Topps baseball cards in 2013.

According to Baumer and Matthews (2014), the only sabermetric measure comparable to *WAR* in impact is *On Base Percentage*, which measures a player's value in one aspect of the game. For the present setting, then, no other technological breakpoint candidates exist other than 2004. Therefore, we will conduct a *Chow* (exogenous structural break) *Test*—through our empirical specification itself—in the estimation section to follow. The structural break testing initially will center on the year 2004 before considering the possibility of an earlier breakpoint. Despite not having a tractable endogenous structural break test at our disposal, we will rely on one means of testing for a more gradual change in voter behavior. In the visualization to follow, as well as in the estimation section, we divide our pre-2004 sample into two equal sub-samples: 1980–1991 and 1992–2003. If voters were simply improving in their ability to measure player value gradually, we might expect the slope of the trend line between *vote point count* and *WAR* to increase incrementally between both the first (1980–1991) and second (1992–2003) sub-samples and between the second and third (2004–2017) sub-samples. The trinary color-coded plot in Fig. 3b helps us consider that relationship across time.

In terms of the (slope of) relationship between vote points and *WAR*, no substantial difference is evident between the first two periods. As such, Fig. 3b is quite similar to Fig. 3a as far as identifying a potential change in the relationship of interest beginning in 2004 goes. To examine further if 2004 represents a structural breakpoint, we plot the mean season-league Spearman rank correlations between player *WAR* and player MVP vote points for each of the three periods in Fig. 4.

The time series plot of Spearman rank correlation coefficients shows little difference in means between the first two periods. However, the sample mean increased markedly during the period from 2004 onward. In our estimation models, we will consider both binary and trinary partitions of the dataset within a multivariate regression setting.

2 Model specification

Given that the dependent variable, *vote point count*, is a non-negative integer-valued variable, we specify a count model herein. For all but one model, we forego the specification of a traditional linear model in favor of a count model because count data are zero-censored and typically right-skewed. The observations likewise are sparse because the variable takes on only integer values, and that sparseness contributes to non-normally distributed residuals if one explains variation in a count dependent variable by estimating a traditional linear model (Cameron and Trivedi 2013). A chi-squared test reveals over-dispersion in our dependent variable such that we adopt a *negative binomial* count model rather than a *Poisson* count model. The baseline model contains *team*, *season*, and playing *position* as fixed effect variables; a Hausman specification test supports selection of fixed over random effects modeling for each

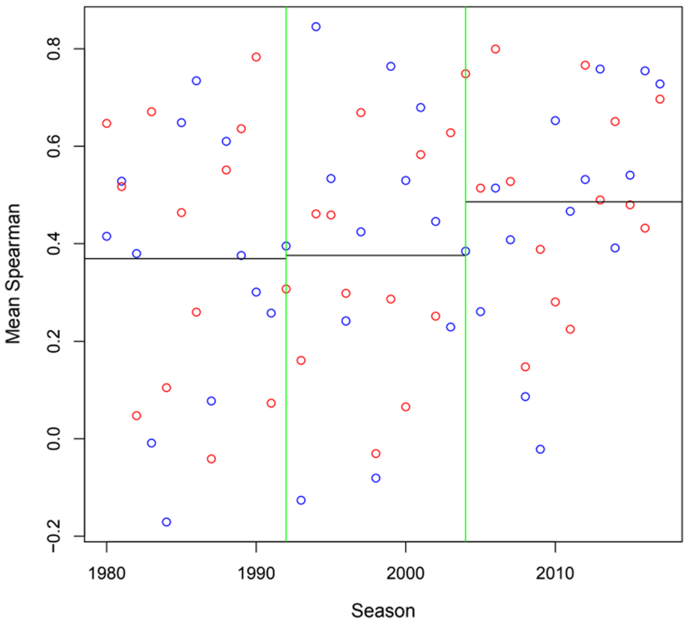

Fig. 4 Mean spearman rank correlation between WAR and MVP votes by time period

independent variable. Given the institutional and empirical setting of the dataset, we forego entering *player* fixed effects. Many sampled players (283) appear in the dataset in one season only. Several other players (70) appear in the dataset in more than one season but do not earn an MVP vote in any season. Moreover, *team* and *position* fixed effects are mutually compatible within one specification, whereas the specification of *player* fixed effects precludes entering both *position* and *team*, as explanatory variables because players tend to remain on the same team's roster and play the same *position*. Thus, specification of *player* fixed effects would lead to (a) loss of substantial data representing specific types of player-seasons (sub-sample selection bias) and (b) preclusion of our other fixed effect variables. As such, we decide that these data do not have classic micro-panel characteristics and specify a panel structure at the levels of *team*, *position*, and *season*.

Our set of fixed effect, negative binomial regressions appear as follows. The baseline model tests for a general relationship between *vote point count* and *WAR* from 1980 to 2017. We specify one model with *team* fixed effects and one model with *big market* in place of *team* fixed effects. We do so because *big market* is collinear with *team* fixed effects and thus cannot enter the same model. The baseline models are specified as follows.

$$votepts_{i,t} = \beta_0 + \beta_1 WAR_{i,t} + \beta_2 age_{i,t} + \beta_3 age^2_{i,t} + \beta_4 multiple_t eams_{i,t} + \beta_5 \mathbf{position}_i + \beta_6 \mathbf{team}_i + \beta_7 \mathbf{season}_t + \epsilon_{i,t} \tag{1}$$

$$votepts_{i,t} = \beta_0 + \beta_1 WAR_{i,t} + \beta_2 age_{i,t} + \beta_3 age^2_{i,t} + \beta_4 multiple\_teams_{i,t} + \beta_5 bigmarket_i + \beta_6 \mathbf{position}_i + \beta_7 \mathbf{season}_t + \epsilon_{i,t} \tag{2}$$

Nothing ensures a priori that *vote point count* and *WAR* bear a significant relationship over the full data sample. The joint purpose of the baseline models is to test whether any such relationship exists. As *WAR* was not created until 48 of the 76 sampled votes had been cast, voters made decisions in the absence of *WAR* for almost two-thirds of the sample period. Over that timeframe, voters cast ballots presumably based on their observations of the game and on available (decentralized) player statistics. Therefore, no a priori assurance exists that voters adopted *WAR* in their voting behavior following its creation. While *Baseball Writers* certainly have featured *WAR* prominently in (public) baseball articles since 2004, it is an empirical question as to whether that adoption carried over to their (individually private or anonymous) voting behavior.

Following the baseline treatment, we adopt a *Chow Test* (for difference in slopes) to test for differences in the relationship of interest during the 2004–2017 sub-period. Like the scatter plots presented in the previous section, the same test is conducted with 1980–2003 as the (single) comparison period and (in subsequent specifications) with 1980–1991 and 1992–2003 as the (dual) comparison periods. The intercept in our models represents the expected *vote point count* of a replacement player-season (i.e., a player with a *WAR* value of 0). Such player-seasons are well below average by definition and thus not observed in the data. Our two-period *Chow Test* specifications will be conducted as follows.

$$\begin{aligned} votepts_{i,t} = {} & \beta_0 + \beta_1 WAR_{i,t} + \beta_2 \left(WAR_{i,t} x after\_2003_{i,t} \right) + \beta_3 age_{i,t} + \beta_4 age^2_{i,t} \\ & + \beta_5 multiple_t eams_{i,t} + \beta_6 \mathbf{position}_i + \beta_7 \mathbf{team}_i + \beta_8 \mathbf{season}_t + \epsilon_{i,t} \end{aligned} \tag{3}$$

$$\begin{aligned} votepts_{i,t} = {} & \beta_0 + \beta_1 WAR_{i,t} + \beta_2 \left(WAR_{i,t} x after\_2003_{i,t} \right) + \beta_3 age_{i,t} + \beta_4 age^2_{i,t} + \beta_5 multiple\_teams_{i,t} \\ & + \beta_6 big\_market_{i,t} + \beta_7 \mathbf{position}_i + \beta_8 \mathbf{season}_t + \epsilon_{i,t} \end{aligned} \tag{4}$$

The variable, *after_2003*, is an indicator that is set equal to 1 for MVP voting years 2004–2017 and 0 otherwise. We interact that variable with *WAR* to compare the relationship between *vote point count* and *WAR* before and after the measure's creation. The variable *after_2003* does not enter the right-hand side of the model alone for reasons discussed previously, namely, that no reason exists for expecting a difference in intercept values across time periods. Moreover, the intercept is extrapolative within the model, as replacement players are not expected to receive votes. We next consider a model that is consistent with a three-period *Chow Test* for changes in slope.

$$\begin{aligned} votepts_{i,t} = {} & \beta_0 + \beta_1 WAR_{i,t} + \beta_2 \left(WAR_{i,t} x after_{2003_{i,t}} \right) + \beta_3 \left(WAR_{i,t} x between\_92\&03_{it} \right) + \beta_4 age_{i,t} + \beta_5 age^2_{i,t} \\ & + \beta_6 multiple\_teams_{i,t} + \beta_7 \mathbf{position}_i + \beta_8 \mathbf{team}_i + \beta_9 \mathbf{season}_t + \epsilon_{i,t} \end{aligned} \tag{5}$$

$$\begin{aligned} votepts_{i,t} = {} & \beta_0 + \beta_1 WAR_{i,t} + \beta_2 \left(WAR_{i,t} x after\_2003_{i,t} \right) \_ \beta_3 \left(WAR_{i,t} x between\_92\&03_{it} \right) \\ & + \beta_4 age_{i,t} + \beta_5 age^2_{i,t} + \beta_6 multiple\_teams_{i,t} + \beta_7 big\_market_{i,t} + \beta_8 \mathbf{position}_i + \beta_9 \mathbf{season}_t + \epsilon_{i,t} \end{aligned} \tag{6}$$

The variable *between_92&03* allows us to trisect the data so as to test for possible evidence that voting behavior actually began to change before 2004 (e.g., as some more gradual result of the sabermetric movement in general). In essence, models (5) and (6) allow us to conduct a *Chow Test* for changes in the relationship between *vote points* and *WAR* for the periods 1980–1991, 1992–2003 and 2004–2017. If the creation of *WAR* were integral to changing that relationship, we might expect it to remain fairly stable before 2004 and to change significantly (and perhaps substantially) thereafter.

3 Estimation and results

In this section, we report the estimation results for the six fixed effect, negative binomial models specified in the previous section. These results are presented in Table 2.

Over the full sample, Table 2 demonstrates that the relationship between *WAR* and MVP *vote points* is positive and significant. Models 3 and 4 show that the estimated slope of that relationship increased significantly (became significantly more positive) following the creation and publication of *WAR*. Models 5 and 6 provide evidence that the increase did not arise as a gradual and vague response to the sabermetric era in general (e.g., not as a process that began in years prior to 2004 and built from there). Rather, the models demonstrate that voters behaved in a manner from 1992 to 2003 statistically equivalent to their behavior from 1980 to 1991. It is only in the 2004–2017 sub-sample that we observe a change in voter behavior. As such, we have evidence that (informed) post-2003 MVP voters allocated vote points in ways that are more consistent with actual player value than did their relatively uninformed counterparts of earlier periods. Furthermore, the explanatory power of the model (i.e., in terms of improved R^2) rises following the creation of *WAR*.

For the purpose of assessing explanatory power before and after the development of *WAR*, we divide our sample into two sub-periods (1980–2003 and 2004–2017). We then estimate models (1) for each period separately as ordinary least squares, fixed effects regressions. We do so because fixed effects, negative binomial estimation does not yield R^2 coefficients or any other intuitive measure of explanatory power. We find that overall R^2 rises from 0.366 in the former sample to 0.464 in the latter. That is, the percentage of variation in vote points explained by variations in player characteristics and performances rose by almost 10 percentage points in the latter period. It is not simply that voting was significantly and substantially more responsive to estimated player contributions to winning games beginning in 2004. From 2004 forward, voting has been explained better by measures of player value (i.e., has been less subject to noise).

Negative binomial regression coefficient estimates can be interpreted as log-linear semi-elasticities. For example, model (3) suggests that each additional unit of *WAR* increased players' expected *vote point count* by 52.8% prior to 2004 and by 52.8 + 7.8 or 60.6% after 2004. Models (4)–(6) provide similar semi-elasticity results. If pre-2004 voting had been as responsive to *WAR* as subsequent voting, *ceteris paribus*, several pre-2004 races would have pivoted (e.g., the 1999 and 2001 AL races). Furthermore, we replace *WAR* with *offensive WAR* and *defensive WAR* for position players and with *pitching WAR* or *pWAR* for pitchers. We estimate a modified version of model (3) so as to determine what specific aspects of performance were treated differently by voters after 2004. We run the modified version as a negative binomial model for both (1) position players and for (2) pitchers. Although we estimate the full specification of Table 2's Model (3), we report only coefficient estimates for the variables of interest in Table 3 (for brevity).

The results of Table 3 suggest that voters responded to *all forms of* change in candidate-quality: *offensive WAR*, *defensive WAR*, and *pitching WAR*. For each constituent value measure, a significant increase in the corresponding coefficient from 2004 is observed. Before the advent of *WAR*, no statistical evidence is found that *ceteris paribus* improvements in defensive performance (*defensive WAR* or *dWAR*) raised one's *MVP vote point*

Table 2 Main estimation results

	(1) VotePts	(2) VotePts	(3) VotePts	(4) VotePts	(5) VotePts	(6) VotePts
WAR	0.545****	0.529****	0.528****	0.514****	0.544****	0.532****
	(0.018)	(0.017)	(0.018)	(0.017)	(0.020)	(0.020)
Age	−0.004	0.052	0.027	0.113	0.034	0.116
	(0.089)	(0.086)	(0.088)	(0.085)	(0.087)	(0.084)
Age^2	0.000	−0.001	−0.000	−0.002	−0.001	−0.002
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Multiple teams		−0.791***		−0.849****		−0.849****
		(0.257)		(0.258)		(0.258)
Bigmarket		0.039		0.045		0.053
		(0.061)		(0.060)		(0.060)
War ×after03			0.078****	0.077****	0.065****	0.063****
			(0.012)	(0.012)	(0.015)	(0.014)
War × between _92&03					−0.024	−0.027*
					(0.016)	(0.015)
First base (1B)	0.795****	0.817****	0.758****	0.780****	0.776****	0.808****
	(0.120)	(0.116)	(0.119)	(0.116)	(0.120)	(0.117)
Second base (2B)	−0.081	−0.082	−0.150	−0.180	−0.146	−0.168
	(0.145)	(0.141)	(0.144)	(0.142)	(0.145)	(0.142)
Third base (3B)	0.010	0.048	−0.007	0.029	−0.000	0.043
	(0.126)	(0.122)	(0.125)	(0.122)	(0.126)	(0.122)
Catcher (C)	0.038	0.142	0.040	0.135	0.051	0.149
	(0.161)	(0.156)	(0.161)	(0.156)	(0.162)	(0.156)
Center fielder (CF)	−0.141	−0.105	−0.176	−0.161	−0.165	−0.141
	(0.127)	(0.123)	(0.127)	(0.123)	(0.127)	(0.124)
Designated hitter (DH)	0.983****	1.009****	0.987****	1.016****	1.011****	1.047****
	(0.225)	(0.211)	(0.222)	(0.208)	(0.222)	(0.209)
Left fielder (LF)	0.076	0.046	0.061	0.023	0.068	0.049
	(0.134)	(0.126)	(0.135)	(0.128)	(0.135)	(0.129)
Pitcher (P)	−1.291****	−1.197****	−1.317****	−1.238****	−1.304****	−1.217****
	(0.124)	(0.120)	(0.124)	(0.120)	(0.124)	(0.121)
Right fielder (RF)	0.264**	0.322***	0.219*	0.273**	0.237*	0.299**
	(0.129)	(0.125)	(0.128)	(0.124)	(0.129)	(0.126)
Constant	−4.483****	−5.053****	−5.054****	−5.988****	−5.212****	−6.093****
	(1.344)	(1.279)	(1.330)	(1.263)	(1.331)	(1.259)
Observations	1907	1907	1907	1907	1907	1907

Standard errors in parentheses using default VCE

*$p<0.10$; **$p<0.05$; ***$p<0.01$; ****$p<0.001$

count by any amount! Rather, only *offensive WAR* (*oWAR*) and *pitching WAR* (*pWAR*) led to significant increases in *MVP* vote points before 2004. From 2004, improvements in *dWAR* were rewarded meaningfully in the *MVP* race. As defense is the least salient source

Table 3 Additional estimation results

	(1) VotePts	(2) VotePts
oWAR	0.569****	
	(0.02)	
oWAR_after	0.070****	
	(0.015)	
dWAR	0.002	
	(0.036)	
dWAR_after	0.104**	
	(0.049)	
pWAR		0.721****
		(0.058)
pWAR_after03		0.066**
		(0.032)
Observations	1324	583

Standard errors in parentheses using default VCE. Shortstop is the omitted position category

*$p < 0.10$; **$p < 0.05$; ***$p < 0.01$; ****$p < 0.001$

of value on the baseball field, its lack of (statistical) importance in *MVP* voting before 2004 stands to reason. With the publication of *dWAR* values, evidence exists that defensive value became a productive input in the *MVP* race for the first time. We also find evidence that voting from 2004 on rewarded marginal improvements in *oWAR* and *pWAR* more strongly. All three types of win value in baseball were compensated more strongly by *MVP* voters beginning in 2004.

3.1 Considering election "fairness" and relationships to political voting

The results reported in Table 3 suggest that voters are not only responsive to aggregate information on candidate quality. They are even responsive to information that allows them to grade candidates at the constituent source of value. That conclusion might be seen as analogous to the candidate-issue level in political voting. Table 3 also suggests that voters can be less responsive or even unresponsive to less salient candidate attributes. However, voters also are shown to respond to mechanisms that bring improved clarity to such attributes. In Table 4, we estimate other modified versions of model (3). The modified models consider the potential effects of player age, race, and teammate productivity on a player's *MVP* vote count. We consider those variables while controlling for the player, *WAR*, whether or not the vote took place before 2004, and other characteristics from model (3). In a meritorious MVP election, those additional variables would not influence the vote. That is, they are each irrelevant to the issue of a player's individual contribution to winning baseball games. Hence, many of the results of Table 4 may be seen as a test of the "fairness" or meritorious nature of *MVP* voting—before and after 2004—such that a fair vote would be taken as one that ignores player characteristics that are irrelevant to the issue of player on-field value.

Table 4 "Fairness" of voting behavior estimation results

	(1) VotePts	(2) VotePts
WAR	0.535****	0.521****
	(0.017)	(0.020)
war_after03	0.066****	0.101****
	(0.012)	(0.029)
Age	0.058	0.064
	(0.082)	(0.083)
Age^2	−0.001	−0.001
	(0.001)	(0.001)
teammate_war	0.027****	0.028****
	(0.003)	(0.003)
teammate_war_after		−0.003
		(0.002)
Asian	0.453*	0.584
	(0.275)	(0.477)
African American	0.081	0.102
	(0.079)	(0.094)
Hispanic	0.267****	0.257***
	(0.065)	(0.094)
Asian after03		−0.190
		(0.573)
African American after03		−0.075
		(0.149)
Hispanic After03		0.017
		(0.130)
Constant	−7.310****	−7.392****
	(1.225)	(1.244)
Observations	1907	1907

Standard errors in parentheses. Position variables specified but not listed for brevity

*$p<0.10$; **$p<0.05$; ***$p<0.01$; ****$p<0.001$

In the specifications of Table 4, we modify model (3) by replacing team fixed effects with cumulative *WAR* of the player's teammates in that season. Doing so, we find that a unit increase in cumulative teammate *WAR* increases player vote points by approximately 2.7–2.8% points both before and after 2004, while the main result of interest (i.e., the coefficient on *WAR x after_2003*) remains significant and positive in the modified model. That result presents evidence that *MVP* candidates receive credit for teammate contributions and that the credit was assigned both before and after 2004. That result is consistent with recent evidence from a study of the NBA that finds teammate spillover effects in productivity measures (Ghimire et al. 2020).

Table 4 likewise reveals no evidence that player age influences *MVP* voting. Some evidence that race can influence *MVP* vote points significantly exists even after controlling for on-field player value. Hispanic players are estimated to receive more *MVP* vote points than the reference group (i.e., non-Hispanic white players), *ceteris paribus*, at the $\alpha = 0.05$

significance level. A Hispanic player earns an estimated 25.7% to 26.7% more vote points than do other players, *ceteris paribus*. That bonus is equivalent to the estimated vote point effect of a player gaining half a unit bump in *WAR*. From the controls entered in Table 2′s specifications, we find other results that potentially are relevant from the perspective of political elections. Specifically, we find that democratic voting—in this case, equal voting representation by *MLB city*—can overcome certain exposure advantages that may otherwise accrue to "big city" or other high-profile candidates.

On the basis of those results, we conclude that *Baseball Writers* are not simply writing about advanced baseball statistics to add color to their articles. Since 2004, we find evidence that the creation of *WAR* has changed voting behavior in a high-stakes environment both significantly and substantially.

3.2 Robustness checks

We also assessed the robustness of our findings to alternative estimators, fixed effects, and standard error corrections. We outline the robustness checks briefly here, but the results are available by request. First, as discussed previously, we chose the negative binomial model over the Poisson model because of concerns about overdispersion. However, an alternative approach is to estimate a fixed effects Poisson model, which will not suffer from those problems as long as the standard errors are cluster-robust (Cameron and Trivedi 2009; 2013). We replicated our six models using that approach and found similar results. Next, instead of specifying season fixed effects, we instead specified player-level fixed effects in our six models. Again, we found very similar results. We also assessed whether the type of standard error adjustment influenced our findings. We replicated our estimates using bootstrapped standard errors, standard errors clustered at the season level, and standard errors clustered at the player level. The results are robust to all of those alternative estimation approaches.

4 Conclusion

This paper considers the effect of an exogenous change in the information frontier of voters in a unique election setting, namely, the selection of the most valuable players (MVPs) in the American and National leagues of Major League Baseball. We find that voters respond to an instrument providing new information about candidate quality. Voting behavior is more positively responsive to increases in candidate quality after the information frontier shift. As a result, voting became less noisy overall with respect to observable candidate characteristics. Neutral measures of candidate quality (i.e., measures that reflect only what a player does rather than who the player is) have ensured higher vote counts for more qualified candidates, and the evidence suggests that the response is of a magnitude sufficient to pivot elections. We observe that objective information about candidates is salient to voters' evaluations of candidates. In turn, voter behavior can change significantly and substantially, whereby voters respond positively to objective information on stronger candidates as in Banerjee et al. (2011). We find significant and substantial evidence of such responses not only with respect to the aggregate performance of candidates, but also with respect to individual performance measures that matter to the electorate. Voters respond to evidence of performance, where performance is found to be credited in each of its measured forms. One particularly non-salient candidate characteristic—on-field defensive performance—did not drive voting behavior whatsoever until voters were informed about such performance. Thereafter, voters responded positively to evidence of superior defensive skills. In that setting, voters do not exhibit a firmly fixed mindsets, whereby they might be found to begrudge evidence that calls into question their past voting behavior. Rather, members of the MVP voting body after 2004, many of whom also were members of the same voting body beforehand, exhibited an average tendency toward behavioral agility in processing new information about candidate quality. That is, they responded to the watershed introduction of the Wins Above Replacement (WAR) measure by which to compare player win value cross-positionally.

Our study offers several implications. First, it speaks to the debate on whether informed voters choose higher quality governance or candidates (Banerjee et al. 2011; Pande 2011). Our finding—that an exogenous information shock affects voter behavior—suggests that they do. However, the results might also suggest that voters are conformists.[14] That is, once they accept that a computerized algorithm reflects the real qualities of candidate on the ballot, they vote accordingly. If the algorithm changes, the voters may change their votes in turn. Stated differently, voters can be manipulated easily. Future research might consider the conditions under which information leads to voter manipulation, which has come to light in recent years (Harvey 2016; Ziegler 2018). Second, it remains an open question as to which types of elections our findings extend. Do our findings extend to jurors and judges, whereby after receiving better information better decisions are made? Do our findings extend to political elections, whereby voters cast votes based on a basket of issues?[15] Ostensibly, our study extends more easily to scenarios like the latter, as WAR aggregates win value from different sources (i.e., defense, offense, and pitching). The latter involves ideology, multiple issues, and the risk of strategic voting and manipulation, which is more complex. Nevertheless, the questions remain open and ripe for future research.

Despite the availability of new information, some non-meritorious factors driving voter behavior persisted even after 2004. Most prominently, player race and teammate productivity are found to influence voting behavior throughout the sample period. That finding counters prior work on player race and elections to the *Baseball Hall of Fame*, which did not document evidence of racial discrimination (Jewell 2003). Of course, *MVP* voting and *Hall of Fame* voting are different. Future research might examine those and related issues, especially concerning whether information affects discrimination in voting.

Acknowledgments We would like to thank William Shughart (and other members of the 2019 Public Choice Society Meetings), Marek Kaminski, and anonymous referees for valuable feedback and suggestions. For contributions during the early manuscript phase, we would also like to thank directors and judges of the Carnegie Mellon Sport Analytics Research Competition, as well as Brian Taylor, Shana Gadarian, Simon Weschle (and other faculty members and graduate students at the SU Maxwell School Department of Political Science), Justin Ross, B. Andrew Chupp, and Seth Freedman (and other faculty members and graduate students at the IU School of Public and Environmental Affairs). We would also like to thank Lara J. Potter for her encouragement.

Funding This research was not funded.

Code availability Code is unpublished.

[14] We thank an anonymous reviewer for providing those thoughtful comments.

[15] Drazen and Eslava 2010 study voting scenarios based on a basket of issues.

Compliance with ethical standards

Conflict of interest There are no conflicts/competing interests with any of the authors.

Availability of data and materials Data is unpublished.

References

Ashworth, S., & De Mesquita, E. B. (2014). Is voter competence good for voters? Information, rationality, and democratic performance. *American Political Science Review, 108*(3), 565–587.

Banerjee, A., Kumar, S., Pande, R., & Su, F. (2011). Do informed voters make better choices? Experimental evidence from urban India. http://www.povertyactionlab.org/node/2764.

Baron, D. P. (1994). Electoral competition with informed and uninformed voters. *American Political Science Review, 88*(1), 33–47.

Baumer, B., & Matthews, G. (2014). A statistician reads the sports page: There is no avoiding WAR. *Chance, 27*(3), 41–44.

Baumer, B. S., Jensen, S. T., & Matthews, G. J. (2015). openWAR: An open source system for evaluating overall player performance in major league baseball. *Journal of Quantitative Analysis in Sports, 11*(2), 69–84.

Boudreau, J., Ehrlich, J., Raza, M. F., & Sanders, S. (2018). The likelihood of social choice violations in rank sum scoring: Algorithms and evidence from NCAA cross country running. *Public Choice, 174*(3–4), 219–238.

Cameron, A. C., & Trivedi, P. K. (2009). *Microeconometrics using stata*. College Station, TX: Stata Press.

Cameron, A. C., & Trivedi, P. K. (2013). *Regression analysis of count data*. Cambridge: Cambridge University Press.

Drazen, A., & Eslava, M. (2010). Electoral manipulation via voter-friendly spending: Theory and evidence. *Journal of Development Economics, 92*(1), 39–52.

Ehrlich, J. A., Ghimire, S., Khraiche, M., & Raza, M. F. (2020). COVID-19 countermeasures, sporting events, and the financial impacts to the North American leagues. *Managerial Finance* (Forthcoming).

Garmann, S. (2017). The effect of a reduction in the opening hours of polling stations on turnout. *Public Choice, 171*(1–2), 99–117.

Ghimire, S., Ehrlich, J. A., & Sanders, S. D. (2020). Measuring individual worker output in a complementary team setting: Does regularized adjusted plus minus isolate individual NBA player contributions? *PLoS ONE, 15*(8), e0237920.

Ginzburg, B. (2017). Sincere voting in an electorate with heterogeneous preferences. *Economics Letters, 154,* 120–123.

Goff, B. L., Shughart, W. F., & Tollison, R. D. (1998). Moral hazard and the effects of the designated hitter rule revisited. *Economic Inquiry, 36*(4), 688–692.

Goff, B. L., & Tollison, R. D. (1990). Sports as economics. In B. Goff & R. D. Tollison (Eds.), *Sportometrics* (pp. 3–11). College Station: Texas A&M University Press.

Greenwich, C. T., Gazianis, J. T., Khader, Z., Lytle, R., Owens, W., Keating, C., et al. (2019). A comprehensive review of voting systems and a proposal for connecticut. http://studentpolicyinitiative.org/projects/IRVPaper2.pdf

Harvey, C. J. (2016). Changes in the menu of manipulation: Electoral fraud, ballot stuffing, and voter pressure in the 2011 Russian election. *Electoral Studies, 41,* 105–117.

Hill, J. R., & Groothuis, P. A. (2001). The new NBA collective bargaining agreement, the median voter model, and a Robin Hood rent redistribution. *Journal of Sports Economics, 2*(2), 131–144.

Jaffe, J. (2004). The class of 2004: Analyzing the hitters. Baseball Prospectus. Retrieved 18 June 2019. https://www.baseballprospectus.com/news/article/2502/the-class-of-2004-analyzing-the-hitters/.

Jensen, N. M., Malesky, E. J., & Walsh, M. (2015). Competing for global capital or local voters? The politics of business location incentives. *Public Choice, 164*(3–4), 331–356.

Jewell, R. T. (2003). Voting for the baseball hall of fame: The effect of race on election date. *Industrial Relations, 42*(1), 87–100.

Kaiser, B. (2019). Strategy and paradoxes of borda count in Formula 1 racing. *Decyzje*, *31*, 115–132.

Madden, B. (2017). WAR shouldn't be so heavily emphasized when it comes to voting for the Baseball Hall of Fame. Retrieved 03 July 2019. https://www.nydailynews.com/sports/baseball/war-shouldn-heavily-emphasized-baseball-hof-voting-article-1.3684618.

McMurray, J. (2015). The paradox of information and voter turnout. *Public Choice, 165*(1–2), 13–23.
Pande, R. (2011). Can informed voters enforce better governance? Experiments in low-income democracies. *Annual Review of Economics, 3*(1), 215–237.
Pietryka, M. T., & DeBats, D. A. (2017). It's not just what you have, but who you know: Networks, social proximity to elites, and voting in state and local elections. *American Political Science Review, 111*(2), 360–378.
Ross, J. M., & Dunn, R. R. (2007). The income tax responsiveness of the rich: Evidence from free agent major league baseball all-stars. *Contemporary Economic Policy, 25*(4), 639–648.
Shughart, W. F. (1997). Preserve baseball's antitrust exemption, or, why the senators are out of their league. Steerike Four!: What's Wrong with the Business of Baseball? 143.
Slowinski, S. (2010). "What is WAR?" http://www.fangraphs.com/library/index.php/misc/war/.
Smith, D. W. (2000). The retro sheet. Retrosheet.org.
Thoth, M., & Chytilek, R. (2018). Fast, frugal and correct? An experimental study on the influence of time scarcity and quantity of information on the voter decision making process. *Public Choice, 177*(1–2), 67–86.
Young, P. (1995). Optimal voting rules. *Journal of Economic Perspectives, 9*(1), 51–64.
Ziegler, C. E. (2018). International dimensions of electoral processes: Russia, the USA, and the 2016 elections. *International Politics, 55*(5), 557–574.