SO, WHAT *IS* DRIVING THE HOME FIELD ADVANTAGE?

Hint: Vocal fans matter,
but not in the way you might think

It had the makings of a nearly perfect day. Jack Moore had just finished his sophomore year at the University of Wisconsin and was home for a few summer weeks, living with his folks. Marooned in the Mississippi River town of Trempealeau, Wisconsin, Jack was blissfully free of pressure, with generous rations of free time. He had a job coaching baseball, but the games didn't start until the evening. On this Friday of the 2009 Fourth of July weekend, Jack's beloved Milwaukee Brewers were playing an afternoon road game against their rivals the Chicago Cubs.

Air-conditioning blasting, Jack flicked on the cable to the regional sports network and sat down on the couch to watch. The Brewers were coming off a magical 2008 season in which they won 90 games and reached the playoffs. In the off-season, Milwaukee's ace, C. C. Sabathia, was poached by the Yankees. It was the numbingly familiar fate of a small-market team: The Brewers had been unable to match New York's $161 million contract offer. Jack was okay with that. A math major, he knew the economic realities and understood why Milwaukee could not afford to retain a star at those prices. Besides, the Brewers' 2009 incarnation was easy to root for, a fun team with a winning record, filled with young and energetic players.

The game was a rare Wrigley Field pitching duel pitting the Cubs' ace, Carlos Zambrano, then a Cy Young Award candidate, against Milwaukee's veteran Jeff Suppan. The game was tied 1–1 after nine innings, which was all good with Jack, a former high school baseball player who was thoroughly capable of appreciating a low-scoring affair. "It was one of those games," he recalls, "that remind you why you like baseball so much."

Then, in the bottom of the tenth inning, Jack's idyllic afternoon was ruined. The Brewers had summoned Mark DiFelice, a right-handed pitcher who had recently won his first Major League game at age 32. When the Cubs loaded the bases, DiFelice faced Chicago's third baseman, Jake Fox, a utility man who'd ricocheted between the majors and the minors. With a full count, two outs, and the decibel level soaring at Wrigley Field, DiFelice threw four consecutive pitches that Fox fouled off. On the next pitch of the at-bat, DiFelice reared back and fired a cutter that froze Fox and shot past him. After an awkward pause, home plate umpire Bill Welke popped up from his crouch and . . . stood idly. Ball four. The winning run had been walked home: Cubs 2, Brewers 1.

The crowd goes wild. Jack Moore of Trempealeau, Wisconsin, goes ballistic. "For five minutes, I just screamed words you can't print," he says. "Anyone who knows baseball knew that was a strike." For years, fans in Jack's position would bitch and moan and dispute balls and strikes until last call. But this was 2009, and Jack wasn't interested in an argument; he was interested in a straight, objective answer. He fired up his Internet browser, logged on to MLB.com, and clicked on Pitch f/x. Sure enough, DiFelice's pitch was gut-high and clearly within the upper-inside part of the strike zone. Minutes after the game had ended, right there in his parents' home in small-town Wisconsin, a 19-year-old was able to confirm his suspicions. The ump had blown the call, permitting the home team to win.

■ ■ ■

What sports fan doesn't harbor a belief that the officials are making bad calls against his or her team? It's a home crowd that voices this displeasure the loudest. The criticism ranges from passably clever ("Ref, if you had one more eye, you'd be a Cyclops!") to the crass ("Ref, you might as well get on your knees because you're blowing this game!") to the troglodytic ("You suck!"). Dissatisfaction is voiced individually and also collectively, often in a stereo chant of "Bullshit! Bullshit!" In Europe—quaint, civilized Europe—there are even various soccer websites that enable fans to download antireferee chants as ringtones.

What we've found is that officials *are* biased, confirming years of fans' conspiracy theories. But they're biased not against the louts screaming unprintable epithets at them. They're biased *for* them, and the bigger the crowd, the worse the bias. *In fact, "officials' bias" is the most significant contributor to home field advantage.* "Home cooking," as it's called, is very much on the menu at sporting events.

A statement like that had better have some backing, and we're prepared to provide it. Warning: An assault of numbers awaits. But stick with us and we'll walk you through it. We think the payoff is worth it.

■ ■ ■

Let's start by determining how to measure ref bias. You could examine the accuracy of calls made by the officials and whether that accuracy differs for calls favoring the home team versus the away team. But doing that is a challenge because it requires a great deal of subjectivity as well as a deep knowledge of the circumstances of the game. Was it really a foul? Was it really pass interference? What else was happening during the game at that time? In light of the speed of the game and the reactions of players within the game, it is nearly impossible to control for all the potential factors that could lead to differing calls for the home and away teams.

Suppose we find that more fouls are called against road teams than against home teams—which, by the way, is often the case. Does this indicate a referee bias in favor of the home team? Maybe, but not necessarily. What if teams play more aggressively on the road? After all, road teams know that statistically, they are already more likely to lose. Or what if the road team, exhausted from those back-to-back games, lacks the energy for proper defense and clutches and grabs instead? They might be inclined to commit more fouls regardless of any referee bias, and so it's difficult to identify the *causal* factor. Are referees *causing* more road team fouls because of bias against the road team? Or are players causing referees to call more fouls because of more sloppy or aggressive play? Or is there a third factor causing both?

We looked for a component of the game the refs control that isn't influenced or affected by players. We found it in a sport for which we have not had much success in explaining its sizable home advantage—soccer. It also turns out that had it not been for a diligent grandmother from Spain religiously watching and recording years' worth of Sunday evening matches, we might not have discovered this bias at all.

In soccer, the referee has discretion over the addition of extra time, referred to as "injury time," at the end of the game to make up for lost time resulting from unusual stoppages of play for injuries, penalties, substitutions, and the like. This extra time is rationed at the discretion of the head referee and is not recorded or monitored anywhere else in the stadium.

As best he can, the referee is supposed to determine the accumulated time from unusual stoppages—itself a subjective measure—and add that time at the end of regulation. So does the referee's discretion favor the home team? If so, he would lengthen this time when the home team is behind at the end of the game and reduce it when the home team is ahead, extending or shortening the game to increase the home team's chances of winning.

Using handwritten notes that his elderly mother had gathered logging matches she'd watched from her living room in Spain, Natxo Palacios-Huerta, a London School of Economics professor, joined with two colleagues from the University of Chicago, Luis

Garicano and Canice Prendergast—all soccer fanatics—to study the officials' conduct during games. The researchers were, quite justifiably, struck by what they found. Examining 750 matches from Spain's premier league, La Liga, they determined that in close matches with the home team ahead, the referees ritually shortened the game by reducing the extra time significantly. In close games in which the home team was behind, the referees lengthened the game with extra injury time. If the home team was ahead by a goal at the end of regulation, the average injury time given was barely two minutes, but if the home team was behind by a goal, the average injury time awarded was four minutes—twice as much time. Sure enough, when the score was tied and it wasn't clear whether to increase or decrease the time for the home team, the average injury time was right around three minutes.

What happened when the home team was *significantly* ahead or behind? In games that were not close, there was no bias at all. The extra time added was roughly the same whether the home team was ahead by two goals or more or behind by two goals or more. This makes sense. A referee has to balance the benefit of any favoritism he might apply with the costs of favoritism—harm to his reputation, media scrutiny, and potential reprimands. Adding additional injury time when the score was so lopsided was unlikely to change the outcome and therefore accrue much benefit, so why do it and risk the potential cost?

The study also looked at what happened when, in 1998, the league altered its point structure from awarding teams two points in the standings for a win (and one for a draw and zero for a loss) to three points for a win. That change meant that a win was suddenly worth a lot more than it had been before and the difference between winning and tying doubled. What did this do to the referee injury time bias? It increased it significantly. In particular, preserving a win against the possibility of a tie now meant a lot more to the home team, and so the referees adjusted the extra time accordingly to reflect those greater benefits.

This wasn't unique to Spain. Researchers began looking for the same referee biases in other leagues—not hard given the global popularity of soccer. They found that the exact same injury time bias in favor of the home team exists in the English Premier League, the Italian Serie A league, the German Bundesliga, the Scottish league, and even MLS in the United States.

If referees are willing to alter the injury time in favor of the home team, what else might they be doing to help ensure that the home crowd leaves happy? We found that referees also award more penalties in favor of the home team. Disputed penalty shots and goals tend disproportionally to go the home team's way as well. Looking at more than 15,000 European soccer matches in the English Premier League, Spanish La Liga, and Italian Serie A, we found that home teams receive many fewer red and yellow cards even after controlling for the number of penalties or fouls on both teams. The dispensing of red and yellow cards has a large impact on a game's outcome. A red card, which sends the offending player off the field, reduces a team's chances of winning by more than 7 percent. A yellow card, which precedes a red card as a stern warning for a foul and may therefore cause its recipient to play more cautiously, reduces the chances of winning by more than 2 percent. These are large effects. When a single yellow card, followed by a red card, is given to a visiting player, it means the home team's chance of winning, absent any other effects, jumps to 59 percent. Add the injury time, fouls, free kicks . . . and it suddenly isn't so surprising that the home team in soccer wins nearly 63 percent of its games.

But could this be limited to the idiosyncratic world of European soccer? Surely, American sports wouldn't be subject to the same referee bias . . . would they?

Remember how, despite a significant home team advantage, athletes do not hit or pitch, shoot free throws, slap goals, or pass the football appreciably better at home than they do on the road? This prompts the question: What *do* home teams do better that allows them to achieve a higher winning percentage?

In baseball, it turns out that the most significant difference between home and away teams is that home teams strike out less and walk more—a lot more—per plate appearance than do away

teams. This could be for lots of reasons. One interpretation: Home team batters see the ball better or away team pitchers exhibit less control. But this contradicts our earlier results for batters and pitchers—in controlled, isolated environments, they hit and pitch the same at home as they do on the road. And as we've seen, road players in MLB aren't performing worse because they're exhausted from the travel.

Balls and strikes are the domain of the head umpire. Could the umpire be biased toward the home team? This would explain the differences in strikeouts and walks despite the lack of any difference in hitting and pitching.

But strikeouts and walks are not the right statistic to measure, because many strikes occur when a batter swings and misses or fouls off a ball. In such cases, there is no umpire discretion. A better metric to look at is *called* balls and strikes.* In other words, look only at pitches that do not involve swinging by the batter. It turns out that home batters receive far fewer called strikes per called pitch than away batters do.

It's even more apparent when we look at called strikes and balls at different points in the game. Certain situations have a much bigger impact on the game's outcome than others. Fortunately for us, Sabermetrics, an analysis of baseball through objective evidence, provides another useful tool to gauge the importance of a particular situation. A stats wizard, Tom Tango, devised a metric called the Leverage Index to measure the relative importance of any game situation. The idea is to take every game situation and consider every possible scenario that could occur in that situation, the likelihood of each scenario playing out from that point, and what effect each of those scenarios would have on the ultimate outcome of the game. Add up all these possibilities, their likelihood of occurring, and their potential impact on the game and you have a measure of how crucial the current situation is. A Leverage Index of 1 is the average situation; an index of 2 means the situation is twice as crucial. Here are two extreme examples: Down by four runs with two outs and nobody on base in the bottom of the ninth, where the game isn't in much doubt, translates into a Leverage Index of 0.1—the situation is one-tenth as crucial as the average situation. Down by one run in the bottom of the ninth with two outs and the bases loaded, where the game is on the line, gives a Leverage Index of 10.9. It is almost 11 times more crucial than the average situation.

Using the Leverage Index to examine called strike and ball counts in different situations, we found, just as with the soccer referees, that in low-leverage situations, when the game is not in much doubt, the home team advantage in receiving fewer called strikes and more balls goes away. But as the following chart shows, the called-strike advantage for home teams grows considerably as the game situation gets more and more important. In noncrucial and average situations, the home team receives about the same strike calls as, or even a few more strike calls than, the away team per called pitch, but that changes dramatically when the game is on the line. In crucial situations, the home team receives far fewer called strikes per called pitch than does the away team.

DIFFERENCE IN PERCENTAGE OF CALLED PITCHES THAT ARE CALLED STRIKES ON HOME VS. AWAY BATTERS

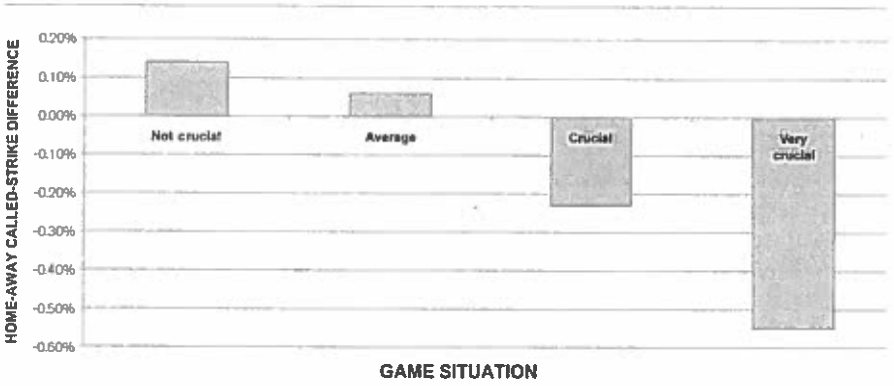

* Eliminating intentional walks.

This makes sense. If the umpire is going to show favoritism to the home team, he or she will do it when it is most valuable—when the outcome of the game is affected the most. You might even contend that in noncrucial situations the umpire may be biased against the home team to maintain an overall appearance of fairness.

Think back to that Jake Fox pitch in the Cubs-Brewers game, on a 3–2 count with the bases loaded and a tie game on the line in the bottom of the tenth inning. It was an astronomically high-leverage situation. Knowing the statistics, you would have bet the house that the pitch wouldn't have been called a strike. And it wasn't.

Let's look at other calls that fall under the domain of the umpires, in particular, close calls that typically elicit a home crowd reaction. Two good examples would be stolen bases and double plays. We found that home teams are more likely to be successful when stealing a base and when turning a double play, yet the distance between the bases is identical in every stadium—stolen base success can't be driven by home field idiosyncrasies. In addition, the success rates of home teams in scoring from second base on a single or scoring from third base on an out—typically close plays at the plate—are much higher than they are for their visitors in high-leverage/crucial situations. Yet they are no different or even slightly less successful in noncrucial situations. (Third-base coaches, take note: If it's a close game and you're playing at home, windmill your arms and send the runner!)

But the most damning evidence of umpire bias might be a function of a tool that was employed for the specific purpose of eliminating umpire bias. Remember the Pitch f/x system that tracks the characteristics of each pitch, including location? Well, its predecessor—a digital technology called Umpire Information System (UIS) from QuesTec—was installed five years earlier by Major League Baseball for the specific purpose of monitoring the accuracy of umpires. According to Major League Baseball, QuesTec was implemented in six ballparks in the first year; by the time it was discontinued in 2008, 11 ballparks had the technology.* With two cameras positioned at field level and two in the upper deck, QuesTec combined the four images to track where the ball crosses the plate, and it was used by baseball executives to determine how closely an umpire's perception of the strike zone mirrored reality.

We also used the presence of QuesTec to evaluate umpire accuracy, but in a different way. We asked whether the same umpire behaved differently when he *knew* the cameras were monitoring him. If the home field advantage in called strikes disappears when the umpires know they're being watched—while everything else stays constant—it's pretty clear that official bias underlies it. Imagine you own a coffee shop and put out a jar in which patrons can donate or take loose change. You notice at the end of each day that the jar is empty. You deduce that either some customers are taking advantage by depleting the jar or your employees are stealing the coins. You tell only your employees that you are installing a hidden video camera. If the change jar is full at the end of each day, you're pretty darn sure it was your employees, not customers, who were to blame.

To test our theory, we first compared all pitches, about 5.5 million of them, from 2002 to 2008 made in stadiums using QuesTec versus those without it. For example, we looked at all called pitches when the Astros visited the Cardinals (at their non-QuesTec stadium) and when the Cardinals visited the Astros (at their QuesTec-equipped stadium).

What did we find? Called strikes and balls went the home team's way, *but only* in stadiums without QuesTec, that is, ballparks where umpires were not being monitored. This is consistent with an umpire bias toward the home team causing the strike-ball discrepancy. We also found something surprising. Not only did umpires

* According to Major League Baseball, the 11 franchises whose ballparks were equipped at various times with QuesTec were the Arizona Diamondbacks, Boston Red Sox, Chicago White Sox, Cleveland Indians, Houston Astros, Los Angeles Angels of Anaheim, Milwaukee Brewers, New York Mets, New York Yankees, Oakland A's, and Tampa Bay Rays.

not favor the home team on strike and ball calls when QuesTec was watching them, they actually gave *more* strikes and *fewer* balls to the home team. In short, when umpires *knew* they were being monitored, home field advantage on balls and strikes didn't simply vanish; the advantage swung all the way over to the visiting team.

We then looked at the same pitch counts in low-leverage (not crucial) and high-leverage (crucial) points in the game. Again, when a plate appearance is expected to have little effect on the outcome of the game, there is no bias for or against the home team. Umpires call things evenly whether QuesTec is present or not. But when the at-bat can have an impact on the game, we found both biases to be even more extreme. That is, when the game is on the line, home teams in non-QuesTec stadiums get a big strike-ball call advantage and those in QuesTec stadiums get a huge strike-ball call *disadvantage.*

In practical terms, when the umpire is *not* being monitored by QuesTec, a home batter in crucial game situations will get a called strike only 32 percent of the time if he doesn't swing. In the same situation, a batter from a visiting team gets a called strike 39 percent of the time. That's a big difference. Now consider the same two situations when the umpire *is* being monitored by QuesTec. Here the home batter gets a called strike 43 percent of the time, and the away batter only 35 percent of the time.

If we were consultants to a team equipped with umpire-monitoring technology, our first piece of advice would be: Get rid of QuesTec; it's wrecking your home field advantage. (How many teams would have agreed so readily to QuesTec if they knew these numbers?) Of course, if we were consulting for MLB, we might have encouraged them to install the technology in *all* the ballparks or at least tell the umpires that was the case. (Today, that's essentially what MLB has done.)

We also found the same results for the QuesTec stadiums before and after the system was installed. The called strike-ball differences between home and away teams declined sharply after QuesTec installation, and the decline was particularly pronounced in crucial situations. Even the *same* umpire behaved differently depending on whether QuesTec was present, calling more strikes and fewer balls on home batters when he was being monitored and doing the opposite when he wasn't.

Why would the home team advantage for strike-ball calls, particularly in crucial situations, switch completely in the other direction when QuesTec is present? You'd think the advantage would just disappear, creating no bias, but in fact the bias goes in the opposite direction. We suspect that as with the referees in soccer, umpires have to balance the costs and benefits of any bias (conscious or not) they might exhibit. If you know you are being monitored, you want to eliminate any perception of bias. And when the game is on the line, you know that any perceived bias will be scrutinized even more closely. With the speed of the game and the uncertainty of whether a 95-mph fastball hit or missed the outside corner of the plate, umpires may become overly cautious. Worried about accusations of home-team favoritism, the umpire seems to err in the other direction, particularly in situations that will be monitored and analyzed heavily afterward.

What about the other potential umpire biases we found that might benefit the home team, such as stolen bases, double plays, and other close plays? They remain the same whether QuesTec is present or not, which makes sense. After all, QuesTec monitors only the strike zone. It affects no other part of the game. Calls remain in favor of the home team because there's no "surveillance video" on those calls.

If QuesTec is our smoking gun in the case to prove umpire home team bias in Major League Baseball, Pitch f/x provides the ballistic support.* Using the Pitch f/x location data of the millions of pitches we examined earlier, we asked a series of questions: How likely is it that when the pitch is actually out of the strike zone, an umpire will call a strike on the home team versus the away team? How often is a ball called on a pitch actually within the

* Pitch f/x is now in every ballpark, and thus, one could argue, umpires are now monitored everywhere. However, Pitch f/x—unlike QuesTec—is not being used to evaluate performance. There's a big difference between casually monitoring umpires and bosses formally monitoring umpires.

strike zone when the home team is batting versus the away team, the situation that enraged the die-hard Brewers fan Jack Moore? What about pitches just in or just out of the strike zone? How do these calls change in critical situations?

The following chart graphs the difference in the percentage of called pitches that are called strikes on home versus away batters for pitches within three inches of the dead center of the strike zone, pitches way out of the strike zone (at least three inches), and pitches within 1.5 inches of the strike zone, for example, just on or off the corners. We report the numbers separately for two-strike, three-ball, and full counts.

DIFFERENCE IN PERCENTAGE OF CALLED PITCHES THAT ARE CALLED STRIKES ON HOME VS. AWAY BATTERS

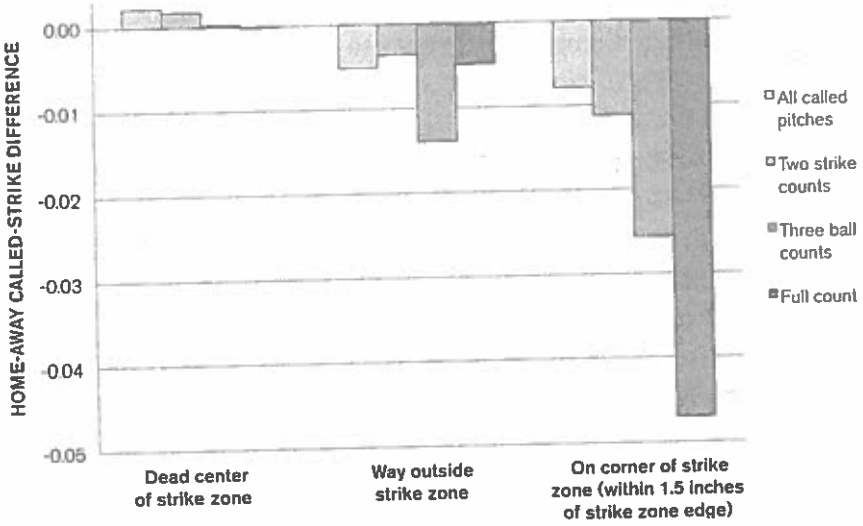

Note two points: (1) The home-away differences are largest for two-strike and three-ball counts and especially for full counts. (2) For the most ambiguous pitches—the ones on the corners—the home-away called-strike discrepancy is largest, which makes sense. The umpire has less discretion over pitches that are less ambiguous. Umpires will be reluctant to make a biased call if the pitch is obviously a strike. In fact, for pitches in the dead center of the strike zone, there is no bias at all. Umpires call these pitches correctly 99 percent of the time whether a home or a visiting batter is standing in front of them. For pitches way outside the strike zone, the umpire has a little more leeway and shows a slight bias in favor of home batters. The umpire has the most discretion for pitches on the corners, and there the home batter bias is largest.

Over the course of a season, all of this adds up to 516 more strikeouts called on away teams and 195 more walks awarded to home teams than there otherwise should be, thanks to the home plate umpire's bias. And this includes only terminal pitches—where the next called pitch will result in either a strikeout or a walk. Errant calls given earlier in the pitch count could confer an even greater advantage for the home team.

How much do these differences contribute to the home field advantage in baseball? Well, we need to know the value of receiving an extra pitch instead of striking out and the value of being awarded first base instead of facing another pitch, but here's a rough estimate. Taking the value of a walk and a strikeout in various game situations, this adds up to an extra 7.3 runs per season given to each home team by the plate umpire alone. That might not sound significant but cumulatively, home teams outscore their visitors by only 10.5 runs in a season. Thus, more than two-thirds of the home field advantage in MLB comes by virtue of the home plate umpire's bad calls.

We can't expect umpires to be perfect, and in fact, they call strikes and balls correctly 85.6 percent of the time. But the errors they do make don't seem to be random. They favor the home team.

Now that we understand that there is a bias in called balls and strikes, we get a different understanding of why the home team has better hitting and pitching stats. As we've seen, players aren't hitting or throwing any better at home versus on the road. But when you receive more favorable calls at the plate, this directly improves your hitting numbers. There is also an indirect effect. If home batters are benefiting from more favorable pitch calls, they face more favorable pitch counts and are in a better position to

swing at pitches to hit. And when home players are put in these situations, it is more likely that their teammates will be on base when they are at the plate, which gives them more opportunities to produce runs. In short, the direct effect from giving home batters fewer strikes and more balls alone seems to account for a sizable fraction of the home team's success in MLB. Add to this the indirect benefits and it could well account for just about all of the home team's advantage.

■ ■ ■

For evidence of official bias in the NFL, it makes sense to start by considering one obvious component in the control of the men in the striped uniforms: penalties. Home teams receive fewer penalties per game than away teams—about half a penalty less per game—and are charged with fewer yards per penalty. Of course, this does not necessarily mean officials are biased. Away teams might commit more violations and play more sloppily or more aggressively. But when we looked at more crucial situations in the NFL—much as with the Leverage Index or the pitch count in baseball—we found that the penalty bias is exaggerated. It turns out that more valuable penalties, those that result in first downs, also favor the home team.

The most compelling evidence of referee influence in the NFL comes from the introduction of instant replay, which gave coaches—and fans, players, and the media—a chance to review and potentially challenge the call on the field. The inauguration of instant-replay challenge came in 1999, and as with the QuesTec results in baseball, it coincided with a decline in the home team success rate in the NFL, from 58.5 percent (from 1985 to 1998) to 56 percent (from 1999 to 2008), a 29.4 percent drop in the home field advantage. Remember, the home advantage starts only when we get above 50 percent.

Coincidence? We can start by looking at turnovers. First, officials wield considerable influence here because they first determine whether there *was* a fumble. Second, they determine which team assumes possession of the football. Before instant replay, home teams enjoyed more than an 8 percent edge in turnovers, losing the ball far less often than road teams. When instant replay came along to challenge wrong calls, the turnover advantage was cut in half.

We can also distinguish between fumbles lost (possession changes hands) and fumbles retained (the team with the ball keeps possession). The home team does not actually fumble or drop the ball less often than the away team—in other words, they aren't "taking care of the ball" any better or worse than the away team. They simply lose fewer fumbles than away teams. After instant replay was installed, however, the home team advantage of *losing* fewer fumbles miraculously disappeared, whereas the frequency of fumbles remained the same. Home teams are as likely as ever to *drop* the ball, but now that visiting teams have the ability to challenge the call, home teams aren't nearly as likely to retain possession.

In close games, when referees' decisions may *really* matter—and when the crowd is really involved—home teams enjoyed a healthy 12 percent advantage in recovering fumbles. After instant replay was installed, that advantage simply vanished.

What about penalties? Instant replay is of limited use to us here because teams can't challenge a penalty call or a noncall. But if we examine the change in penalty discrepancy between home and away teams before and after instant replay, we have a placebo test of sorts. That is, we should not expect to see any changes in penalties. Sure enough, we don't. The discrepancy in number of penalties and yards per penalty given to home versus away teams hardly changed after instant replay. This helps confirm that it is instant replay, not something else, that has driven the recent changes in turnovers and winning percentage of home teams in the NFL.

If referee bias is driving these patterns and instant replay mitigated these biases, we should see that visiting teams are more successful when they challenge a referee's call using instant replay. In other words, if away teams are indeed getting more bad calls than home teams, more of those calls will be overturned on instant replay. We looked at the results of nearly 1,300 instant-replay challenges from 2005 to 2009 to examine the success rate of home team challenges versus away team challenges.

The results? It turns out that away teams are indeed more successful in overturning a call than home teams are, but only modestly so (37 percent versus 35 percent). Both are slightly more successful than official challenges (33 percent), which are challenges initiated by an official in the last two minutes of each half on close plays. These statistics are misleading, though, because as we saw in baseball and soccer, referees are less likely to make biased judgments when the game is no longer in doubt. So what happens if the home team is behind? When the home team is losing, a challenge made by the home team is successful 28.4 percent of the time. But a challenge made by the away team is successful 40.0 percent of the time. Thus, away teams seem to be getting more than their fair share of bad calls when they are winning, which is when bad calls would be most valuable to the home team.

Could referee bias explain a large part of the home field advantage in football? Absolutely. Again we see a dramatic reduction in the home team's edge when instant replay is introduced. Yet instant replay affords each team only a maximum of three incorrect challenges per game and is limited to certain circumstances. Clearly there are other calls not eligible for challenge that could favor the home team, such as penalties. The fact that home teams in football have better offensive stats—such as rushing more successfully and having longer time of possession—could be the result of getting more favorable calls, fewer penalties, and fewer turnovers. If you play at home and sense that you're less likely to get called for a penalty, you may be more inclined to block much more aggressively or challenge a receiver.

■ ■ ■

Recall that in the NBA home and away teams shoot identically from the free throw line. But home teams shoot more free throws than away teams—between 1 and 1.5 more per game. Why? Because away teams are called for more fouls, particularly shooting fouls. Away teams also are called for more turnovers and more violations. These differences could be caused by more aggressive or sloppy play on the part of road teams, which could be more tired because of the lopsided NBA schedule. But they are also consistent with referee bias.

To help distinguish sloppy play on the road from referee home bias, let's take a closer look at the *types* of fouls, turnovers, and violations that are committed by home and away teams. Certain fouls, turnovers, and violations require more referee discretion and judgment than others. For example, highly uncertain situations and close calls, where a judgment must be made, allow for greater referee influence, as opposed to something less ambiguous such as a shot clock violation that everyone can easily monitor because the 24-second shot clock is posted above the two baskets and a red light illuminates the glass backboard when the clock expires.

If sloppy or aggressive play by the away team is causing these differences, we should not expect to see the number of violations vary with how ambiguous or uncertain the fouls, turnovers, or violations are regardless of how much referee judgment is required. If you're playing badly, you're probably playing badly across many dimensions of the game.

We looked at calls requiring more or less referee judgment to see whether the home advantage was the same. Loose ball and offensive fouls seem to be the most ambiguous and contentious. Ted Bernhardt, a longtime NBA official, now retired, helped us with our analysis. "Blocking fouls versus charging fouls are by far the hardest calls to make," he says. It turns out that offensive and loose ball fouls go the home team's way at twice the rate of other personal fouls. We can also look at fouls that are more valuable, such as those that cause a change of possession. These fouls are almost *four* times more likely to go the home team's way than fouls that don't cause a change of possession.

What about turnovers and violations? Turnovers from shot clock violations, which aren't particularly ambiguous or controversial, are no different for home or away teams. Turnovers from five-second violations on inbounds plays, which are also fairly unambiguous because everyone can count (though referees may count a little slower or faster than everyone else and there is no clock indicating when five seconds has elapsed), are also not very

different for home and away teams (in fact, home teams receive slightly more five-second violations).

If, however, we look at the most ambiguous turnover calls requiring the most judgment, such as palming and traveling, we see huge differences in home and away numbers. The chance of a visiting player getting called for traveling is 15 percent higher than it is for a home team player. The fact that ambiguous fouls and turnovers tend to go the home team's way and unambiguous ones don't is hard to reconcile with sloppy play on the part of visiting teams. But it's exactly what you would expect from referee bias.

Identifying refereeing bias in the NBA is especially hard because context is so important, and some of the most controversial "calls" in basketball are in fact "no calls"—when a call is not made. But the evidence seems to suggest ref bias toward the home team. If bias clearly exists in soccer, baseball, and football, isn't it reasonable to suspect that NBA referees are vulnerable to the same influences?

Remember the Portland Trail Blazers playing so sluggishly in that dreary midweek road loss to the San Antonio Spurs? On the road in San Antonio, the Trail Blazers committed 13 fouls, the Spurs 14; each team had six turnovers. But if we look at the *types* of fouls and turnovers over which referees have more influence, we see that the Blazers were whistled for twice as many loose ball fouls as the Spurs and that five of the Blazers' six turnovers were on judgment calls made by the referee (one traveling, two ambiguous lost balls out of bounds, one offensive goaltending, and one questionable kicked ball). By contrast, all six Spurs turnovers were unambiguous (five bad passes/steals and one shot clock violation). In addition, more of the calls against the Blazers resulted in a change of possession favoring the Spurs. Perhaps it's not so surprising that the Spurs won.

Recall that only a few nights later the teams met again, this time in Portland. The Blazers won by 18 points. The same advantages conferred on the home team were present, though this time it was the Trail Blazers who were the beneficiaries. The Spurs were whistled for 25 fouls and 16 turnovers, compared with Portland, which had 18 fouls and 13 turnovers. The types of foul calls and turnovers tell an even stronger story. Among the visiting Spurs' 16 turnovers, 11 were of the more ambiguous variety, including a couple of debatable lost balls out of bounds, and two Spurs players were even called for palming. (To give you an idea of how rarely palming is called in the NBA, on average there is one palming call every five or six games.) Of the Blazers' 13 turnovers, 10 were unambiguous, consisting of two shot clock violations and bad passes that were stolen or thrown out of bounds. There were ten situations in which the ball was tipped out of bounds—eight went the Trail Blazers' way. More fouls resulting in a change of possession went Portland's way as well. If we tally the numbers across the two games, ambiguous turnovers went the home team's way 85 percent of the time, ambiguous fouls were charged to the visiting team 72 percent of the time, and the home team won by a collective 33-point margin.

So how much of the home court advantage in the NBA is due to referee bias? If we attribute the differences in free throw attempts to referee bias, this would account for 0.8 points per game. That alone accounts for almost one-fourth of the NBA home court advantage of 3.4 points per game. If we gave credit to the referees for the more ambiguous turnover differences and computed the value of those turnovers, this would also capture another quarter of the home team's advantage. Attributing some of the other foul differences to the referees and adding the effects of those fouls (other than free throws) on the game, this brings the total to about three-quarters of the home team's advantage. And remember, scheduling in the NBA explained about 21 percent of the home team's success, as well. That adds up to nearly all of the NBA home court advantage.

Long story short, referee bias could well be the *main reason* for home court advantage in basketball. And if the refs call turnovers and fouls in the home team's favor, we can assume they make other biased calls in favor of the home team that we cannot see or measure.

■ ■ ■

What about the NHL? By now you can probably guess what we found. Home teams in hockey get 20 percent fewer penalties called on them and receive fewer minutes in the box per penalty. (In other words, home teams are not just penalized less often but penalized for less severe violations.) The net result is that on average per game, home teams get two and a half more minutes of power play opportunities—a one-man advantage during which goals often are scored—than away teams. That is a *huge* advantage. To provide some perspective, the average NHL team succeeds in scoring a goal during a two-minute power play about 20 percent of the time. So if you take the power play advantage and multiply it by the 20 percent success rate (per two minutes), this gives the home team a 0.25-goal advantage per game. The average point differential between home and away teams in the NHL is 0.30 goals per game, so this alone accounts for more than 80 percent of the home ice advantage in hockey.

But is the penalty difference driven by refereeing bias? Repeating the same exercise we conducted for the NBA, we looked at more ambiguous calls—holding, hooking, cross-checking, boarding, tripping—and found that these penalties in particular went the home team's way. Less ambiguous calls such as too many men on the ice, illegal equipment, delay of game from sending the puck into the stands, and fighting had much less home team bias. Again, this is consistent with officiating bias—and not with tired or sloppy play from visiting teams.

Also, don't forget the shootout results we discussed earlier. Remember, in a shootout we found no home ice advantage. Not coincidentally, this is the only part of the game in which the referee essentially plays no role.

■ ■ ■

The fact that we can identify an officiating bias toward the home team is unsettling—that this may be the chief reason home field advantage exists in every sport is *very* unsettling. But *why* are officials biased toward the home team?

WHY DO OFFICIALS FAVOR THE HOME TEAM?

First let's be clear: Is there a conspiracy afoot in which officials are somehow *instructed* to rule in favor of the home team, especially since the league has an economic incentive to boost home team wins? Almost unquestionably no. We're convinced that the vast majority of, if not all, officials are upstanding professionals, uncorrupted and incorruptible, consciously doing their best to ensure fairness. All things considered, they do a remarkable job.

They are not, however, immune to human psychology, and that's where we think the explanation for home team bias resides. Despite fans' claims to the contrary, referees are, finally, human. Psychology finds that social influence is a powerful force that can affect human behavior and decisions *without the subjects even being aware of it.* Psychologists call this influence conformity because it causes the subject's opinion to conform to a group's opinion. This influence can come from social pressure or from an ambiguous situation in which someone seeks information from a group.

In 1935, the psychologist Muzafer Sherif conducted a study about conformity, using a small point of light in an otherwise dark or featureless environment. Because of the way the human eye works, the light appears to move, but the amplitude of the movements is undefined—individual observers set their own frames of reference to judge amplitude and direction. Therefore, each individual saw the "movement" differently and to differing degrees.

When participants were asked individually to estimate how far the light had moved, as one would expect, they gave widely varying answers. Then they were retested in groups of three. The composition of the group was manipulated; testers put together two

people whose estimate of the light movement when alone was very similar and one person whose estimate was very different. Each person in the group had to say aloud how far he or she thought the light had moved. Sherif found that over numerous trials, the group converged on a common estimate. The subject whose estimate of movement had been vastly different from that of the other two in the group came to conform to the majority view.

More important, when interviewed afterward, the subject whose initial estimate had been very different now *believed* his or her initial estimate was wrong. That is, that subject did not succumb to social pressure and state something he or she didn't believe; his or her actual perception of the light's movement had changed. The experiment demonstrated that when placed in an ambiguous situation, a person will look to others for guidance or additional information to help make the "right" decision.

After the Sherif study, Solomon Asch, a pioneer of social psychology, conducted an experiment in which he asked participants to look at two cards and decide which line (A, B, or C) on the card on the right in the following illustration was most like the line on the card on the left.

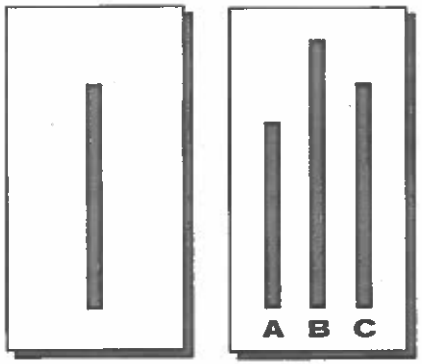

The answer, you probably guessed, is C. The participants, though, were asked to make this assessment in a group setting. Asch had put one unwitting subject in a room with seven confederates, or actors. The actors were told in advance how to respond. Each person in the room gave his or her answer, and the "real" participant offered his or her answer second to last. In most of the cases, the subject yielded to the majority at least once, even though he or she suspected it was wrong.

Asked why they readily conformed to the group even though they felt the answer was wrong, most participants said that they did not really believe their answer; rather, they went along with the others for fear of being ridiculed or thought "peculiar." A few, however, said that they really did believe the group's answers were correct. Asch also found that subjects felt enormous stress when making these decisions; giving a response that was at odds with the majority caused anxiety, even though they knew they were right.

The takeaway here is that human beings conform for two reasons: (1) because they want to fit in with the group and (2) because they believe the group is better informed than they are. Makes sense, right? If you are asked to make a decision and are unsure of your answer, wouldn't you look for other cues and signals to improve that answer? And don't you accord weight to people's answers by the confidence with which they provide them? After a difficult test in school, who hasn't polled other classmates for the answer to a question, paying particular attention to the responses of the known "A" students?

■■■

Now, back to referees. When humans are faced with enormous pressure—say, making a crucial call with a rabid crowd yelling, taunting, and chanting a few feet away—it is natural to want to alleviate that pressure. By making snap-judgment calls in favor of the home team, referees, whether they consciously appreciate it or not, are relieving some of that stress. They may also be taking a cue from the crowd when trying to make the right call, especially in an uncertain situation. They're not sure whether that tailing 95-mph fastball crossed the strike zone, but again, even if it's subconsciously, the crowd's reaction may provide a useful signal that changes their perception.

If beliefs are being changed by the environment, as psychology

shows, referees aren't necessarily consciously favoring the home team but are doing what they believe is right. It's just that their perceptions have been altered. In trying to make the right call, they are conforming to a larger group's opinion, swayed by tens of thousands of people witnessing the exact same play they did. As the saying goes in psychology, "I'll see it when I believe it." Referees, it's safe to assume, do not intend this favoritism. They're probably not even aware of it. But it is a natural human response.

Remember, too, that on top of the anxiety caused by passionate and sometimes angry fans, the refs receive stress from their supervisors and superiors. In a variety of ways—some subtle, some not—officials must take in cues that the league has an economic incentive for home teams to do well. If your boss sent a subtle but unmistakable message that Outcome A was preferable to Outcome B, when you were forced to make a difficult, uncertain, and quick decision, how would you be inclined to act?

Let's look at our previous results on referees through the lens of psychology and our understanding of the human propensity to conform. The extra injury time in soccer? It is probably a response to social pressure, that is, the desire to please the crowd—and in some cases preserve personal safety. The strike-ball discrepancy in baseball and similar disparities in fouls and turnovers in basketball, along with penalties and turnovers in football and hockey, may also be the result of "informational conformity" in the face of social pressure—using the crowd as a cue to resolve an uncertain or ambiguous situation.

If this is true, psychology suggests that both the crowd size and the uncertainty or ambiguity of the situation should make a difference. Home team favoritism therefore should be greater the larger and more relevant the crowd and the more ambiguous the situation. We've already shown in a variety of ways how the more ambiguous the call—whether it is a 90-mph pitch on the corner of the strike zone in baseball, a fumbled football, a two- or three-step move without dribbling in basketball, or a questionable check in hockey—the more severe the home advantage.

What about the size of the crowd? Recall the original study of the Spanish La Liga. The authors found that the bias in regard to extra time was even more evident when the crowd was larger. Similarly, the studies in the English Premier League, Italian Serie A, German Bundesliga, and MLS also found that referee favoritism was more apparent when attendance was higher. Maybe most interesting was the study conducted in Germany, where many of the soccer stadiums also house a running track that acts as a moat, separating the stands from the field of play. In those stadiums, the referees are more removed from the fans. Guess what? The bias referees usually exhibit for the home team gets cut in half in those stadiums but is the same as it is in other leagues for German stadiums that do not contain a track. In the three European soccer leagues we examined, attendance also had a marked effect on the number of red and yellow cards the visiting team received relative to the home team. Other studies have also linked attendance to penalties and fouls, showing that the bias in favor of the home team grows with the crowd.

What about the extra walks awarded to home teams and the extra strikeouts imposed on away teams by the home plate umpire? These, too, occur predominantly in high-attendance games and are not present in the games with the lowest attendance. The chart below shows the net strikeout and walk advantage to home teams from bad umpire calls, reported separately for the games with the lowest and highest attendance (bottom and top fifth of attended games). Although there is virtually no home team strikeout or walk advantage in the least-attended games, the highest fifth of attended games account for more than half of the entire strikeout and walk advantage given to home teams each season. In the highest-attended games, home teams are given 263 fewer strikeouts than their opponents. In the lowest-attended games, that falls to 33 fewer strikeouts. Similarly, the home team receives 93 more walks than the visitors from bad umpire calls in the most-attended games relative to the least-attended ones.

HOME TEAM ADVANTAGE IN STRIKEOUTS AND WALKS FROM UMPIRE INCORRECT CALLS IN LOW AND HIGH ATTENDANCE GAMES

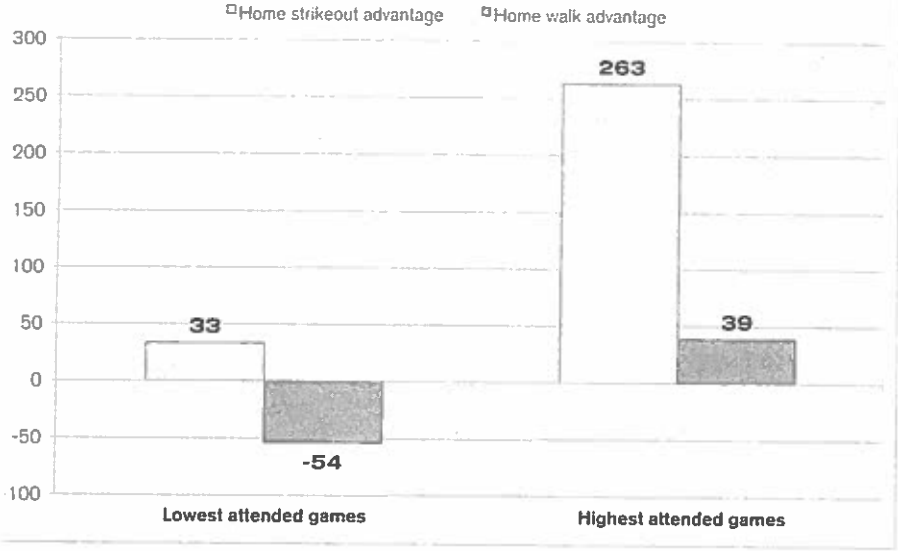

In the NBA, crowd size also affects the home-away differences, particularly for the more ambiguous calls. Recall how traveling is called 15 percent less often against home players. Looking at NBA games in the bottom fifth of attendance, this discrepancy goes down to 6 percent. But if we look at the most-attended games, the home team is 28 percent less likely to be called for traveling.

In the NHL, the bigger the crowd, the more penalties, fouls, and close calls that go against the visiting team, and once again, the effects are greatest for more ambiguous calls. Even in the NFL, in which most games are sold out, the home-away discrepancies in penalties and turnovers increase with crowd size. With virtually every discretionary official's call—in virtually every sport—the home advantage is significantly larger when the crowd is bigger.

In fact, in the least-attended games in each sport, the home field advantage all but vanishes. In MLB, if you look at the 20 percent least-attended games, the home field advantage is only 50.7 percent. In other words, home and away teams are about equally likely to win when the crowd is small. In the one-fifth of games with the highest attendance, however, home teams win 55 percent of the time in MLB. In the NBA, the least-attended games are won by the home team only 55 percent of the time, and the most attended games 69 percent of the time. In the NHL, the home team wins only 52 percent of the time in the lowest-attended games but 60 percent of the time in the highest-attended games. And in European soccer, the home team wins 57 percent of the time in the lowest-attended games and an astonishing 78 percent of the time in the highest-attended matches.

Wait a second, you might say. Doesn't this stand to reason? After all, crappy teams draw crappy crowds, so the games with the empty seats usually involve the worst teams. Never mind official bias; just look at the standings. You'd expect the Pittsburgh Pirates or the New Jersey Nets—lame teams, lame crowds—to win fewer home games than, say, the Boston Red Sox or the Los Angeles Lakers. True, but even after adjusting for the strength of the team we find similar effects. Also, it doesn't matter as much as you might think, because when a bad team hosts a good team, attendance often spikes. When LeBron James and the Miami Heat visit Memphis or Milwaukee, the crowds swell. The worst-attended games usually involve two terrible teams, and the most-attended games feature two great teams. So it turns out there isn't much of a difference in ability between the two teams in either case.

Still not convinced by the psychological explanation for referee bias? Consider a final study, this one performed in 2001. Researchers recorded videos of soccer matches, focusing on tackles during the game, and showed them to two groups of referees. The first group was shown the tackles with the crowd noise audible. The second group was shown the same tackles with the crowd noise muted. Both sets of referees were asked to make calls on the tackles they saw. The referees who watched the tackles with the crowd noise audible were much more likely to call the tackles *with* the crowd. That is, tackles made against the home team (where the crowd complained loudly) were more likely to be called fouls and

tackles made by the home team were less likely to be called fouls. The referees who viewed the tackles in silence showed no bias.

You probably guessed correctly which group of referees made calls consistent with the actual calls made on the field. Yes, the ones who could hear the crowd noise. Not only that, but the referees watching with sound also reported more anxiety and uncertainty regarding their calls, consistent with the stress they felt from the crowd. Imagine how much more intense that stress would have been if they had been on the actual field of play.

But perhaps the most convincing evidence for the effect of crowds on referees occurred when *no fans* were present. On February 2, 2007, supporters of two soccer clubs in Italy—Calcio Catania and Palermo Calcio—clashed with each other and police. It was a typical hooligan-induced riot, and following the episode the Italian government forced teams with deficient security standards at their stadiums to play their home games without *any* spectators present. Two economists (and soccer fanatics) from Sweden, Per Pettersson-Lidbom and Michael Priks, collected the data from the 21 soccer matches that were played before empty bleachers.

What they found was amazing. When home teams played without spectators, the normal foul rate, yellow card, and red card advantage afforded home teams disappeared entirely. Looking at the same team with the same crew of officials, the authors found that when spectators were no longer present, the home bias in favorable calls dropped by 23 to 70 percent, depending on the type of calls (a decline of 23 percent for fouls, 26 percent for yellow cards, and 70 percent for red cards). That is, the *same* referee overseeing the *same* two teams in the *same* stadium behaved dramatically differently when spectators were present versus when no one was watching.

When the economists also looked at player behavior, they found that, unlike the referees, the players did *not* seem to play any differently when the crowd was there yelling versus in an empty, silent stadium. Home and away players shot the same percentage of goals on target, passed with the same accuracy, and had the same number of tackles as they normally do. The absence of the crowd

did not seem to have any effect on their performance. This is in keeping with what we saw for NBA foul shooters, hockey penalty shots, and MLB batters and pitchers: Crowds don't appear to have much effect on athletes.

So it is that we assert that referee bias from social influence not only is present but is *the leading cause of the home field advantage.*

■ ■ ■

We started with three questions that any explanation of the home field advantage must address: (1) Why does it differ across sports? (2) Why is it the same for a particular sport no matter where the game is played? (3) Why hasn't it changed much over time?

To answer the first question, if the refs are responsible for the home advantage, it must be the case that referees are more important or have more influence in some sports (say, soccer, in which home teams have the greatest success) than in others (such as baseball, where the advantage is weakest). As it turns out, this is precisely the case. In soccer, the official has an enormous influence on the outcome of the game. One additional penalty, free kick, or foul can easily decide a game, in which one goal is often all that separates the two teams. In basketball, which has the second highest home team advantage, the official could call a foul on almost every play. By contrast, the umpire's role in baseball is limited relative to other sports. Most plays and most calls are fairly unambiguous; a home run is a home run—either it cleared the fence or it didn't. Most force-outs are not close. Sure, the umpire has discretion over called balls and strikes, but more than half the time the batter swings, eliminating umpire judgment.

In addition, crowd size, which we contend affects referee judgment, has more influence in the sports with the greatest home field advantage. Crowd size matters most in soccer (the sport with the highest home field advantage) and least in baseball (the sport with the lowest home team winning percentage) and is somewhere in between for the other sports. This is also consistent with referees mattering more in some sports (soccer) than others (baseball).

To answer the second question, referee bias also explains why

the home field advantage is the same for a particular sport no matter where it is played. Whether baseball is being played in the United States or Japan, whether it's basketball in the NBA, WNBA, or NCAA or soccer in France versus South Africa, the rules and, more important, the role of the referee are essentially the same, no matter where the game is played.

Finally, to answer the third question, referee bias also explains why the home team's success rate hasn't changed over a century. Although sports have altered their rules over time—raising and lowering the pitcher's mound, introducing a shot clock and the three-point line—the official's role in the game hasn't changed much. Umpires still call balls and strikes, referees still call fouls and penalties, and for over a century these calls have been made by human beings—none of them immune from human psychology.

Although we will never be able to measure or test all the decisions an official makes, if we can see that some biased judgments are being made, it is likely there are other biases going the home team's way that we don't see. Think of the father who comes home early from work and catches his teenage daughter kissing her boyfriend. He's upset about the kiss, but he's more upset about what else she might be doing when he doesn't happen to be looking.

■ ■ ■

Knowing what we now know, let's revisit that Cubs-Brewers game, the ten-inning affair that ruined a summer day for Jack Moore of Trempealeau, Wisconsin. You wouldn't deduce this by scanning a conventional box score or watching a *SportsCenter* highlight. But after revving up the Pitch f/x results, it becomes clear that when the umpire erroneously called a ball on a 3–2 pitch in the bottom of the tenth inning, enabling the winning run to score for the hosting Cubs, it marked the culmination of an afternoon filled with unfavorable decisions against the visiting Brewers.

According to Pitch f/x, Cubs hitters failed to swing at 25 pitches that were strikes. However, nearly a third of them were incorrectly called balls. As for the Brewers, they failed to swing at 32 pitches in the strike zone, only a quarter of which were called incorrectly as balls. Advantage, Cubs. In high-leverage situations, when batters had three balls, *not a single* strike was called on a Cubs hitter even when the ball was in the strike zone—including, of course, the final pitch of the game. But for Brewers hitters facing three-ball counts, *every* pitch in the strike zone was called a strike and *half* the pitches outside the strike zone were called strikes! Big advantage, Cubs.

Overall, the Brewers were deprived of three walks to which they were entitled and the Cubs were given two walks on strikes that were erroneously called balls, including the game-winner. That's a difference of five base runners in a game that ended with a final score of 2–1 in extra innings.

One last point: Recall how the closer officials are to the crowds, the more likely they are to favor the home team. Wrigley Field has about the smallest amount of foul territory in the Major Leagues, so the umpire is uncommonly close to the restless natives. And remember how attendance influences the home field advantage. Wrigley Field seats 41,118 fans and is generally nearly full. In fact, despite a long history of losing seasons, the Cubs have won 54 percent of their home games—above the league average. (It's just that they have been terrible on the road.) That particular afternoon drew a crowd of 41,204—more than 100 percent of capacity with standing room only.

When that mass of humanity on Chicago's North Side yelled at the players, they weren't affecting the outcome. When they yelled at the umpire, well . . . that's another story entirely.