Point Shaving: Corruption in NCAA Basketball

By JUSTIN WOLFERS*

A new field of "forensic economics" has begun to emerge, applying price-theoretic models to uncover evidence of corruption in domains previously outside the purview of economists. By emphasizing the incentives that yield corruption, these approaches also provide insight into how to reduce such behavior. This paper contributes to this agenda, highlighting how the structure of gambling on college basketball yields pay-offs to gamblers and players that are both asymmetric and nonlinear, thereby encouraging mutually beneficial effort manipulation through "point shaving." Initial evidence suggests that point shaving may be quite widespread.

The incentives for gambling-related corruption derive from the structure of basketball betting. To highlight a simple example, the University of Pennsylvania played Harvard on March 5, 2005, and was widely expected to win. Rather than offering short odds on Penn winning the game, bookmakers offered an almost even bet (bet $11 to win $10) on whether Penn would win relative to a "spread." In this example, the spread was −14.5, meaning that a bet on Penn would win only if Penn won the game by 15 or more points, while a bet on Harvard would be successful if Harvard either won, or lost by 14 or fewer points.

The incentive for corruption derives directly from the asymmetric incentives of players, who care about *winning* the game, and gamblers, who care about whether a team beats (or *covers*) the spread.

Indeed, the example above is ripe for corruption: the outcome that maximizes the joint surplus of the Penn players and the gambler occurs when Penn wins the game, but fails to cover the spread (and the gambler has bet on Harvard). The contract required to induce this outcome simply involves the gambler offering a contingent payment to the player, with the contingency being that he pays only if Penn fails to cover the spread. Given the player's (approximate) indifference over the size of the winning margin, even small bribes may dominate his desire to increase the winning margin above 14 points, and this, in turn, yields large profits for the gambler who has bet accordingly. The betting market offers a simple technology for the gambler to commit to paying this outcome-contingent bribe: he can simply give the player the ticket from a $1,000 bet on his opponent not covering the spread.

Such attempts to shave the winning margin below the point spread are colloquially referred to as "point shaving" and form the focus of my inquiry. I start by outlining the type of corruption that theory suggests will be most prevalent:

- Players will be bribed *not* to cover the point spread: It is easy for a player to reduce his effort in response to marginal financial incentives. By contrast, if he usually plays at close to maximal effort—and if the point spread is set on this assumption—then inducing even greater effort will be impossible, or expensive.
- Favorites are more likely to shave points than are underdogs: For an underdog to commit not to cover the spread implies committing to losing the game, while the favorite can both win and not cover. Thus, the payment required to motivate the underdog not to cover is typically larger than that required to induce the favorite to shave points.

† *Discussants:* Raymond Fisman, Columbia University; Bruce Sacerdote, Dartmouth College.

* The Wharton School, University of Pennsylvania, 1456 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104, and CEPR, IZA, and NBER (e-mail: jwolfers@nber.org). Thanks to Steven Levitt and Betsey Stevenson for helpful feedback, and to Ravi Pillai, Lisa Rutherford, and Bryan Elliott for exceptional research assistance.

- Stronger favorites are more likely to shave points than weak favorites: For example, when the spread is three points, an attempt to win by only one to two points may backfire, either leading the player to lose the game (raising his cost of point shaving) or the gambler to lose his bet (lowering the benefit). By contrast, a team that shaves an expected 14-point victory to a margin of nine points faces very little risk.
- While point shaving may affect whether a strong favorite covers the spread, it should have no effect (or minor effects) on the team's chances of winning the game. Thus, point shaving will lead us to observe "too many" teams failing to cover, yet still winning the game.
- Given the discontinuity in gamblers' pay-offs at the spread, there may be a sharp difference in the probability of strong favorites just failing to cover, compared to the proportion just covering the spread.

Further, point shaving is less likely when the probability of detection is high, and more likely when the reward is large. Unfortunately, in the sports betting context, these two factors are likely highly correlated, as televised games, high profile, and high-stakes games all offer larger potential profits due to thicker betting markets, but invite greater scrutiny and hence risk of detection.

I. A Prima Facie Case

Data on the outcomes of 36,003 NCAA Division I basketball games played between 1989–1990 and 1996–1997 were provided by VegasInsider.com, and data for an additional 37,775 games played between 1997–1998 and 2004–2005 were extracted from Covers.com by a Web crawler. Bookmakers tended to take bets only on more popular match-ups and, hence, only 60 percent of these games yielded useful betting data, for a final sample of 44,120 games.

Figure 1 shows that the spread provides an unbiased forecast of game outcomes and explains much of the variation. Note that the efficient-markets hypothesis makes no predictions about whether the spread, *on average,* predicts the winning margin, but, rather, suggests the probability that a team beats the spread is unpredictable. Thus, the lower panel shows a more relevant metric: the frequency with which the home team covers, as a function of the spread. This panel yields little evidence of unexploited profit opportunities. Overall, the favorite beat the spread in 50.01 percent of games. Strong favorites (those favored to win by more than 12 points; $n = 9{,}244$) covered only 48.37 percent of the time, a statistically significant deviation from semi-strong form efficiency. While this evidence might be considered suggestive of point shaving, this implication follows only if the betting market systematically underestimates the incidence of point shaving.

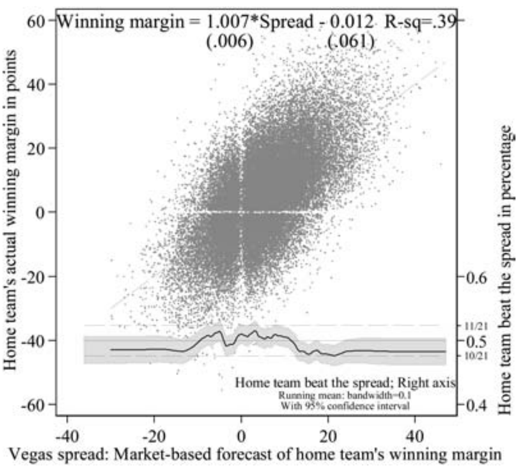

FIGURE 1. NCAA BASKETBALL: GAME OUTCOMES AND THE SPREAD

Note: Sample includes 44,120 NCAA Division I games from 1989 to 2005.

The more direct implications of the theory concern the full distribution of outcomes. The top panel of Figure 2 shows kernel density estimates of the distribution of winning margins—relative to the spread—for all games in the sample (excluding strong favorites). If the spread is a prediction market-generated median forecast (Wolfers and Eric Zitzewitz, 2004), then the deviation of game outcomes from the spread is a forecast error. These forecast errors are roughly normally distributed with zero mean and a standard deviation of 10.9 points.

More intriguingly, the bottom panel shows the comparable distribution for strong favorites. Compared to the normal distribution, this figure suggests that "too few" strong favorites beat the

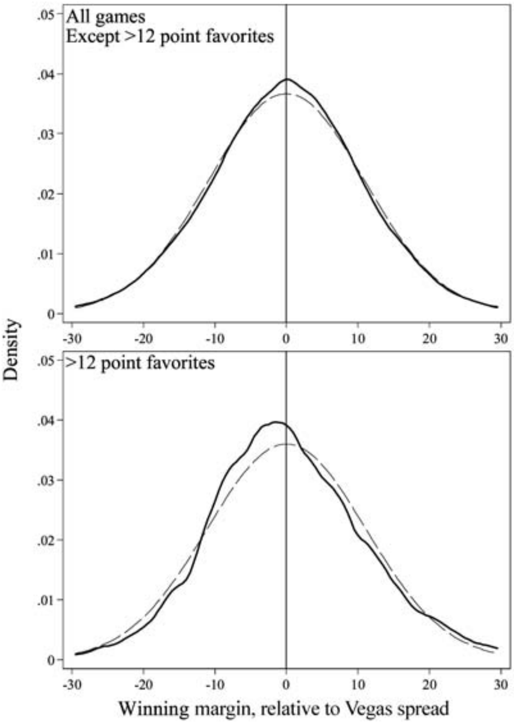

FIGURE 2. PROBABILITY DISTRIBUTION: GAME OUTCOMES RELATIVE TO THE SPREAD

Note: Dashed line shows normal distribution, with mean zero. Solid line shows nonparametric estimated distribution.

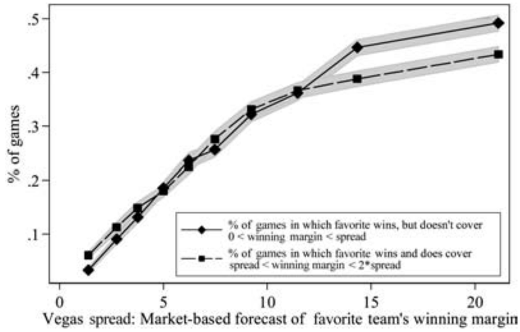

FIGURE 3. NCAA BASKETBALL: BETTING LINE AND BET OUTCOMES

Note: Shaded area shows 95-percent confidence interval.

spread and that this missing probability mass is largely displaced to outcomes in which the team wins, but fails to cover. The left tail of this distribution largely follows the estimated normal distribution, suggesting that while strong favorites differ in their behavior with respect to covering the spread, their likelihood of losing the game remains largely unchanged. These observations are consistent with some strong favorites point shaving.

To the extent that these inferences rely on comparisons of an empirical distribution with a specific parametric distribution, one might be concerned that this analysis rests heavily on assumed functional forms. Simply comparing the two empirical distributions to each other also yields similar conclusions, however.

To estimate convincingly the overall prevalence of point shaving and provide statistical tests would require explaining the empirical distribution of outcomes shown in Figure 2 as a mixture of the distribution of outcomes when point shaving and the distribution when not point shaving, weighted by their relative prevalence. While we are directly interested in these prevalence weights, the component distributions are not directly observable, presenting an identification problem. An explicit structural model of player behavior is required to recover the distribution of outcomes when point shaving and this is the focus of Wolfers (2006).

A simple shortcut yields some illustrative estimates. Under the null of no point shaving, and an assumption that the distribution of forecast errors is symmetric,

$$(1)\quad p(0 < \mathit{Winning\ margin} < \mathit{Spread}) = p(\mathit{Spread} < \mathit{Winning\ margin} < 2 * \mathit{Spread}).$$

Figure 3 shows these probabilities within each point-spread decile. Naturally the proportion of teams that win, but fail to cover, rises as the spread rises (as the range of outcomes this includes increases). Equally, as equation (1) suggests, the proportion of teams in the comparison set of outcomes is typically just as large, as it covers just as wide a range of outcomes. There is a clear difference among strong favorites, however, with significant evidence of "too many" teams winning but failing to cover, relative to the comparison set.

Among teams favored to win by 12 points or more, 46.2 percent of teams won but did not cover, while 40.7 percent were in the comparison range of outcomes. (The standard error on each of these estimates is 0.5 percentage points). If these proportions would have been equal in the absence of point shaving (at around 43 percent), then point shaving led roughly 3 percent of strong favorites who would have covered the spread not to cover (but still win). While this identifies the proportion of games in which the bribe changed whether the team covered, around half of the teams accepting bribes would have failed to cover regardless of the point shaver's behavior, suggesting that the proportion of strong favorites agreeing to point shave is twice as large, or 6 percent.

These calculations are only suggestive, as they derive from comparing observed outcomes with the no-point-shaving null, rather than a fully articulated model of point shaving. A more complete analysis would take account of the possibilities that many of those point shaving would have otherwise beaten the spread by a large margin; that some of those agreeing to point shave will accidentally cover the spread; and that yet others will shave too vigorously, losing the game.

II. Confounding Factors

A. *Effort*

A potentially confounding factor is that players may reduce effort or coaches may use second-tier players when the game outcome is no longer in doubt. To the extent that this incentive is symmetric, the losing team will also reduce its effort and the distribution of game outcomes will be largely unaffected. There is anecdotal evidence of an asymmetry in college sports, however, as "running up the score" is typically regarded as poor sportsmanship.

More important, the effort story yields implications for the distribution of winning margins, while point shaving changes the distribution of winning margins, *relative to the spread.* Hence, these competing explanations are separately identifiable. For instance, 12-point favorites influenced by bribes are sensitive to winning margins above 12 points, and 20-point favorites are sensitive to winning margins above 20 points. Figure 2 provided evidence on precisely this point, and the sharp contrast between the probability of barely failing to beat the spread and barely covering the spread is suggestive of point shaving. By contrast, under the effort interpretation, both 12- and 20-point favorites are sensitive to the *same* sportsmanship norm, and both will endeavor to avoid blowouts in roughly equal measure; this constraint presumably binds more often for the strong favorite, suggesting that the distribution of game outcomes will be more nonnormal (and specifically, left-skewed) as the spread increases. The circles in Figure 4 map the proportion of games ending in decisive victories at each level of the spread, comparing this with a simple baseline computed by assuming that the winning margin is drawn from a normal distribution with a median equal to the spread (as suggested under efficient markets), and the sample standard deviation of 10.9 points. The proportion of blowouts rises smoothly with the market forecast and, if anything, heavily favored teams appear to be involved in *too many* blowouts.

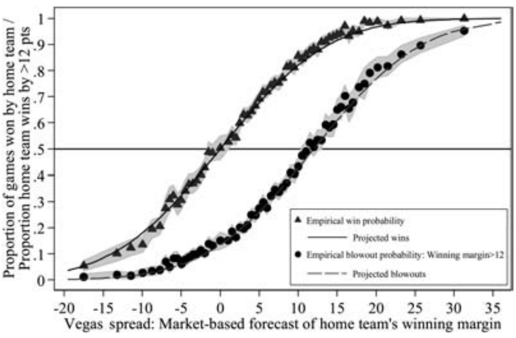

FIGURE 4. THE VEGAS SPREAD AND THE PROBABILITY OF VICTORY OR A BLOWOUT VICTORY

Notes: Projections based on winning margin ~Normal-(Spread, 10.9). Shading shows 95-percent confidence interval.

B. *Betting Market Inefficiency*

An alternative confounding factor derives from cognitive biases that might lead the spread to provide a biased forecast. For instance, overestimating the favorite's ability would lead the peak of the distribution of outcomes for strong favorites to occur at outcomes in which the favorite barely fails to cover, with fewer strong favorites barely covering. This bias affects the whole distribution of outcomes, yielding directly testable implications. For instance, this bias should lead the spread to be a poor predictor of who wins the game, suggesting too few victories by strong favorites. The triangles in Figure 4 show the proportion of victories at different levels of the spread, again mapped against a baseline that assumes that the spread is an unbiased forecast with normal errors. These data provide no evidence that the spread overrates the ability of strong favorites. If anything, strong favorites (particularly strong home favorites) tend to win *more often* than is suggested by the spread.

III. Discussion

These data suggest that point shaving *may* be quite widespread, with an indicative, albeit rough, estimate suggesting that around 6 percent of strong favorites have been willing to manipulate their performance. Given that around one-fifth of all games involve a team favored to win by at least 12 points, this suggests that around 1 percent of all games (or nearly 500 games through my 16-year sample) involve gambling-related corruption. This estimate derives from analyzing the extent to which observed patterns in the data are consistent with the incentives for corruption derived from spread betting; other forms of manipulation may not leave this particular set of footprints in the data, and so this is a lower bound estimate of the extent of corruption. Equally, the economic model suggests a range of other testable implications, which are the focus of ongoing research (Wolfers, 2006).

My estimate of rates of corruption receives some rough corroboration in anonymous self-reports. Eight of 388 Men's Division I basketball players surveyed by the NCAA (2004) reported either having taken money for playing poorly or having knowledge of teammates who had done so.

A shortcoming of the economic approach to identifying corruption is that it relies on recognizing systematic patterns emerging over large samples, making it difficult to pinpoint specific culprits. Indeed, while the discussion so far has proceeded as if point shaving reflected a conspiracy between players and gamblers, these results might equally reflect selective manipulation by coaches of playing time for star players.[1] Further, there need not be any shadowy gamblers offering bribes, as the players can presumably place bets themselves, rendering a coconspirator an unnecessary added expense.

The advantage of the economic approach is that it yields a clear understanding of the incentives driving corrupt behavior, allowing policy conclusions that extend beyond the usual platitudes that "increased education, prevention, and awareness programs" are required (NCAA, 2004, p. 5). The key incentive driving point shaving is that bet pay-offs are discontinuous at a point—the spread—that is (or should be) essentially irrelevant to the players. Were gamblers restricted to bets for which the pay-off was a linear function of the winning margin, their incentive to offer bribes would be sharply reduced. Similarly, restricting wagers to betting on which team wins the game sharply reduces the incentive of basketball players to accept any such bribes. This conclusion largely repeats a finding that is now quite well understood in the labor literature and extends across a range of contexts—that highly nonlinear pay-off structures can yield rather perverse incentives and, hence, undesirable behaviors.

REFERENCES

National Collegiate Athletic Association. *2003 NCAA national study of collegiate sports wagering and associated behavior*. Indianapolis, IN: National Collegiate Athletic Association, 2004.

Wolfers, Justin. "Point Shaving" University of Pennsylvania, Unpublished Paper.

Wolfers, Justin and Zitzewitz, Eric. "Prediction Markets." *Journal of Economic Perspectives*, 2004, *18*(2), pp. 107–26.

[1] While referees may also be corruptible, they are not vested in either the favorite or underdog winning and, thus, are unlikely to be the source of the systematic patterns previously identified.