

TOPIC 13: THE CURSE OF THE COVER OF SPORTS ILLUSTRATED

Before we study this topic, you should read the following article from the archives of Sports Illustrated itself from January 2012: "[That Old Black Magic.](#)" *Sports Illustrated*. January 21, 2002, 5061. (The printed version of the article is attached).

This entertaining article gives a number of anecdotal examples of the supposed curse and goes on to give the results of the investigation into the curse carried out by the authors. The authors give an account of a study they conducted by looking at the covers of 2,456 issues of Sports Illustrated, finding that 913 of the subjects had suffered some verifiable misfortune compatible with their definition of misfortune. They express surprise at the percentage of those featured who met with misfortune but do acknowledge that cause and effect cannot be determined from the data. The authors also discuss a widespread preference in society for superstitious and mystical explanations over statistical arguments and logical explanation.

A thorough scientific investigation of the topic would require a lot of time spent collecting data on all of the athletes featured on the cover of Sports Illustrated. Rather than pursue this line of argument, we will take the opportunity to explore some of the subtle points of probability that often lead us to misjudge chance. We will discuss some common misconceptions about coincidence and ranking along with regression to the mean, which is the most commonly cited explanation for the large proportion of magazine coverees who fall to the hand of the curse. Hopefully this will serve to enable you to make your own judgements as to whether 37.2% is an unexpectedly large number of cover subjects to meet with "misfortune".

As a first step let us use our expertise to translate our question about the jinx into a question about conditional probability, namely **"is the probability of "misfortune" for an athlete given that he/she has "just" appeared on the cover of sports illustrated greater than the probability of "misfortune" for any athlete of the type that appears on the cover of Sports Illustrated?"** At this point, you will realize the answer to this question depends greatly on your definition of "misfortune" and requires some investigation into what qualifies an athlete, coach or member of the general public to appear on the cover of Sports Illustrated.

1. COINCIDENCE

As humans, our attention is often drawn to some events, while others go unnoticed. Certainly remarkable [coincidences](#) have the power to captivate us and distort our perception of relative frequency and chance, often leading us to attribute their occurrence to cause and effect or supernatural forces. In fact many athletes link an unusually good performance with something unusual that happened at the time of the performance or before, perhaps an item of clothing they were wearing during the exceptional performance or what they had for breakfast beforehand. Often success is attributed to the unusual circumstances and from then on the lucky socks are worn during each game or the lucky breakfast precedes it. Of course these rituals may serve to boost the athlete's confidence and improve performance, however, I suspect that their performance statistics would be the same on average even without the aid of the lucky socks.

In addition to the danger of mistakenly attributing cause and effect in this way, we have a tendency to underestimate the likelihood of some coincidences, causing us to attribute them to supernatural forces. The most commonly quoted example of this is [the birthday problem or birthday paradox](#). As it turns out, if you have as few as 23 (randomly chosen) people in a room, the chance that at least two of them share a birthday is slightly greater than 50%. With 50 randomly chosen people in a room, you are almost certain to find two with the same birthday with the probability at 97%. We can arrive at this conclusion easily using our calculation techniques from basic probability (see the link to wikipedia above).

Another apparently remarkable set of coincidences are those [between the lives of US presidents Abraham Lincoln and John F. Kennedy](#). Again, the number of these coincidences is easily explained by the fact that people searching for any coincidence have tens of thousands of well documented life events

for each president to choose from. For example, the political careers of both presidents shared a span of about 16 years which were a century apart; Lincoln from 1834 to 1865 and Kennedy from 1947 to 1963. If we considered 100 political or personal events that might have occurred (relatively randomly) over the span of 16 years for both presidents, for example; years when elected to some office, years when they became involved in a war, years when they moved from one state to another, years when children were born, years when family members passed away etc.. . For each of these events we get a pair of years, one for the occurrence of the event for each president (for example if the event is when both presidents ended their term in the House of Representatives, then we get a pair of years 1849 for Lincoln and 1953 for Kennedy). In a conservative estimate of the probability that the years would coincide for any particular event, one might imagine that the years for these events were randomly determined from the 16 year span and finding the pair of years is akin to rolling a pair of (imaginary) 16 sided dice. Using the same reasoning as we did with the experiment of rolling a pair of six sided dice, we get that the probability that the uppermost faces would coincide each time the pair of 16 sided dice are rolled is $1/16$. Thus we would expect about 6 or 7 coincidences from this exercise alone. Now throw in coinciding days of the year, names, birthplaces and life events of family members, successors and assassins and it would be surprising if we did not get at least 20 coincidences.

Both of the above examples highlight the common mistake of confusing the probability of a particular coincidence with the probability of some coincidence. The probability that at least one person out of twenty three chosen randomly will have a birthday the same as mine is one minus the probability that all of their birthdays are different from mine, $1 - [(364)/(365)]^{23} \approx .06$. This is much smaller than the chances that some two among them will share a birthday (approx .5). Similarly looking for any coincidence when examining the lives of Abraham Lincoln and John F. Kennedy makes the discovery of such a coincidence much more likely than if we were looking for a specific coincidence.

Getting back to the evidence offered for the curse in the article from Sports Illustrated itself, let us ignore the anecdotal evidence lest we fall into the trap of having our minds swayed by the few but sometimes striking examples. Let us instead try take a look at how wide the net has been cast for qualifying coincidences by looking at the authors' definition of what should be counted as a misfortune. The authors enumerated six categories of misfortune:

- (1) an individual slump,
- (2) a team slump,
- (3) an individual blunder or bad play,
- (4) an individual injury or death,
- (5) a bad loss or lousy performance by a team or individual,
- (6) and a failure to win a title after having been featured during the postseason.
- (7) They added a seventh category to accommodate miscellaneous calamities, like Nike's stock plunge shortly after CEO Phil Knight appeared on the cover in 1993.

The time frame within which the misfortune had to hit is described in a sport specific manner but not exhaustively. Starting on the day the magazine hit the stands, the misfortune had to be relatively immediate:

- (a) In baseball or basketball a shooting slump or losing streak had to set in within two weeks of a cover appearance,
- (b) In football, loss or lousy performance had to take place the next weekend,
- (c) for Olympians, their showing at the games was compared to the medal each was forecast to win and injuries were examined for the month following the cover appearance.

Although these **categories** seem to narrow the search for misfortunes, they **cover a wide range of possible coincidences**. The category miscellaneous obviously covers any possible coincidence that might be considered a misfortune and allows just about anything to be counted. It also appears that a player featured on the cover of the magazine for any reason, breaking a record, or an exceptional personal performance, might add to the list of misfortunes with an individual slump, a personal injury, a bad individual performance, a loss by his/her team, a slump in his/her team's performance, a failure by the individual to meet the expectations of sports analysts or a failure by the individual's team to meet expectations of sports analysts.

Adding team losses, slumps and bad performance to the list of misfortunes opens the door to a much wider range of qualifying coincidences. Often the top players in any given sport belong to different teams and whenever two of the teams meet, one must lose. In addition, players who break records or have outstanding individual performances do not always belong to one of the top teams in their league.

More recently, Sports Illustrated has produced double and triple issues, such as the ones show below for the midseason NBA report, featuring three different players from three different teams. I'm not sure how the authors' dealt with these issues but it certainly increases the probability of a misfortune if multiple players from different teams for the same sport are featured on the cover.



The above framework may give the illusion of setting up a scientific experiment (consider a cover and search for a qualifying misfortune in the given time period following the date) and recording the outcome. However the categories of outcomes (along with “no misfortune”) given do not constitute a sample space since the **categories overlap**. This is evident from the fact that the percentage of misfortunes cited in the various categories add to more than 100%.

Of the 2,456 covers, 913 featured a person who, or team that had suffered some verifiable misfortune or loss, that is 37.2% of the covers looked at. We have :

- (i) No misfortune : 62.8%.
- (ii) Bad losses or lousy performance by a team [(5) and maybe (2)] : 19.6% (= 52.7% of 37.2%).
- (iii) Decline in individual performance (1) : 16.59% (= 44.6% of 37.2%).
- (iv) Bad loss or lousy performance by an individual (5) : 9.37% (25.2% of 37.2%).
- (v) Postseason Failure (6) : 4.08% (13.4% of 37.2%).
- (vi) Injury or death (4): 4.38% (11.8% of 37.2%).
- (vii) Blunder or bad play (3): 1.67% (4.5% of 37.2%).







The fact that categories (1), (2) and (5) together account for 122.5% of covers with misfortunes recorded implies that at least 22.5% of the covers associated with a misfortune have more than one associated misfortune, which in turned means that categories (1), (2) and (5) overlap. Although this is not a violation of truth and overlapping categories are common in media representation of statistics, they often lead to a distortion of the perception of the outcomes on the part of the reader and make it impossible to determine the exact percentage of outcomes in the intersection of the categories. The fact that a significant number of the 52.7% of the bad losses or lousy performances by teams may have followed a cover that featured one of the team members for an outstanding individual accomplishment makes the statistic much more believable.

2. PREDICTIONS

A number of the misfortunes cited involve a failure of the players or teams to meet with the expectations of the editors and sports analysts for the magazine. A good deal of the anecdotal evidence given concerns teams or players touted as the “best” in their sport. Bad or unexpected losses by teams or individuals and postseason failure are categories of misfortune, and the performance of Olympians is compared to their predicted medal showing. This leads us to examine two commonly discussed questions in sports; “Which team/individual is the best in their sport?” and “Which team/individual is going to win a match?”. When

debated among fans, the first question often has as many answers as it does debaters. What is often surprising is that sports analysts also often disagree on the answers to both questions. The book “Who’s # 1?” [1] outlines a number of commonly used algorithms used to determine rankings in sport. Quite often, the resulting rankings disagree on the order in which teams should be ranked. Even when rankings roughly agree, the predictive accuracy of the algorithms vary from sport to sport and they are far from 100% accurate. In fact assuming that either of these two questions has a definite answer is a mistake and makes us wonder **how many of the above listed “misfortunes” stem from random error in prediction?**

2.1. Does a “Best” team or player exist in every sport and if so how do we decide who is best? In reality, we often find that there are a number of good teams and athletes with differing strengths. We see below that each statistic from the NBA page for ESPN has a different player as the leader for that statistic.

NBA Statistics - 2015			
Season: 2015 Regular Season			
Offensive Leaders		Defensive Leaders	
POINTS	PPG	REBOUNDS	RPG
	1. James Harden, HOU 27.0		1. DeAndre Jordan, LAC 14.2
	2. Russell Westbrook, OKC 26.5		2. Andre Drummond, DET 13.0
	3. LeBron James, CLE 26.0		3. DeMarcus Cousins, SAC 12.2
	4. Carmelo Anthony, NY 24.2		4. Pau Gasol, CHI 12.1
	5. Anthony Davis, NO 23.9		5. Tyson Chandler, DAL 11.7
Complete Leaders		Complete Leaders	
ASSISTS	APG	BLOCKS	BLKPG
	1. John Wall, WSH 10.2		1. Anthony Davis, NO 2.67
	2. Chris Paul, LAC 9.9		2. Serge Ibaka, OKC 2.37
	3. Ty Lawson, DEN 9.8		3. DeAndre Jordan, LAC 2.32
	4. Rajon Rondo, BOS/DAL 8.3		4. Rudy Gobert, UTAH 2.32
	5. Russell Westbrook, OKC 8.1		5. Pau Gasol, CHI 2.13
Complete Leaders		Complete Leaders	
FIELD GOAL %	FG%	STEALS	STPG
	1. DeAndre Jordan, LAC .720		1. Stephen Curry, GS 2.16
	2. Tyson Chandler, DAL .677		2. Russell Westbrook, OKC 2.04
	3. Mason Plumlee, BKN .591		3. Tony Allen, MEM 1.94
	4. Amir Johnson, TOR .588		4. James Harden, HOU 1.93
	5. Jonas Valanciunas, TOR .556		5. Monta Ellis, DAL 1.92
Complete Leaders		Complete Leaders	

We also find that the team or athlete deemed to be the best by one algorithm is often different from the team or athlete deemed to be the best by a different algorithm. The following image of rankings for 10 teams in college baseball taken from the website masseyratings.com maintained by K. Massey, shows much variation in rankings.

D1A	MAS	MOR	NCB	CBN	PFG	D1B	BA	Rank,	Team,	Conf,	Record	Mean	Median	St.Dev
1	2	3	2	1	5	7	7	1	TCU	B12	6-1	3.50	2.5	2.51
3	10	2	3	5	3	2	6	2	Florida	SEC	6-1	4.25	3.0	2.71
2	1	20	1	2	2	3	4	3	Virginia	ACC	7-0	4.37	2.0	6.39
8	13	1	7	8	4	4	8	4	UCLA	P12	6-1	6.63	7.5	3.62
4	6	25	4	3	6	8	2	5	LSU	SEC	6-1	7.25	5.0	7.42
6	3	9	8	11	8	9	9	6	Texas	B12	7-2	7.88	8.5	2.42
7	9	23	6	6	7	6	3	7	Houston	AAC	6-1	8.37	6.5	6.14
5	11	11	9	7	9	13	5	8	Texas Tech	B12	8-1	8.75	9.0	2.92
12	12	14	10	15	10	5	12	9	South Carolina	SEC	6-1	11.25	12.0	3.06
9	26	54	5	4	1	1	1	10	Vanderbilt	SEC	4-2	12.63	4.5	18.68

In fact a theorem due to [Kenneth Arrow](#) shows that there is no perfect way to put all of these rankings together to get a “best” team.

Thus the idea that there is always one athlete or team unanimously considered the “best” in any sport is a flawed one and it should not surprise us that some teams, players or Olympians featured as the “best” in their sport on the cover of Sports Illustrated fail to live up to the expectations due to errors in judgement and prediction on behalf of the sports analysts. Even if everyone agrees on who should be considered to be

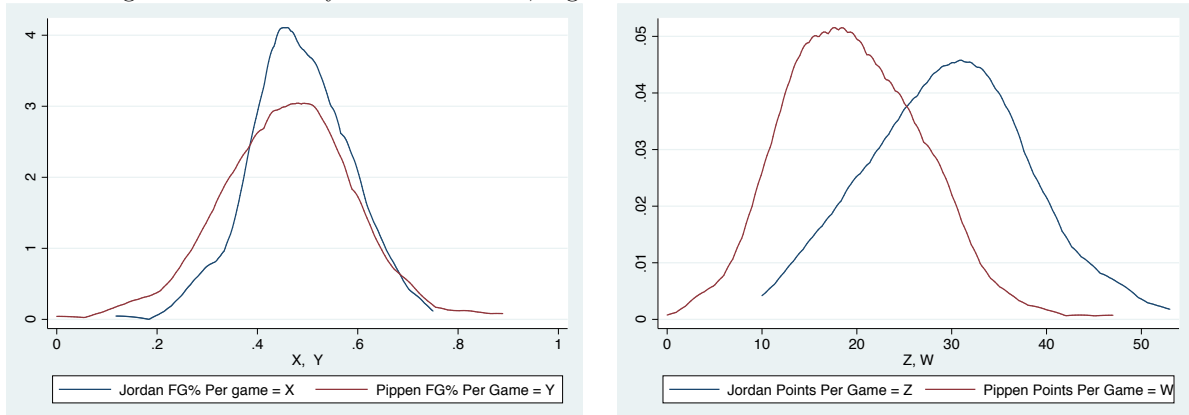
currently the best in a sport, this does not guarantee that future performance will stay at the same level or that losses are inconceivable. We will consider this in more depth in the next section on regression to the mean.

2.2. How accurate are the “Best” prediction systems? We already have enough experience with probability to know that prediction based on recent performance in sports such as American Football where the number of games per season is relatively small will undoubtedly be prone to error. The level of accuracy of predictions from the highly regarded ranking systems varies from system to system and from sport to sport. There are several different measures of accuracy in the literature. A straightforward measure of accuracy for a given ranking system in any given week of a tournament is given by the percentage of winners to date which are ranked higher than their opponents by the most recent version of this ranking system. A number of websites track the accuracy of ranking systems throughout a season. Among them [Wobus Sports](#) and [The Prediction Tracker](#). Surprisingly we find that 70-75% is the highest level of accuracy for the best (current) algorithms using the above measure of accuracy. The computer rankings tend to have a greater level of accuracy than polls such as the coaches poll which uses the judgement of insiders in the sport to determine rankings. The algorithm which ranks highest in accuracy will also vary from week to week. There are of course more sophisticated ways of measuring accuracy of rankings, but they are more difficult to explain and I’m sure there is no “best” method of measuring accuracy of ranking systems :)

3. REGRESSION TO THE MEAN

A lack of understanding of regression to the mean is often cited as the main reason for a belief in the jinx of the cover of sports Illustrated. In the case of Bernoulli trials, we have seen that random fluctuation can lead to long runs of success’ in the data which does nothing to change the probability of failure in the next trial. For example, if I flip a coin many times, a streak of 10 heads in a row does not change the fact that the probability of a tail on the next coin flip is $1/2$. As we saw above performance for any athlete or team is measured by many performance statistics which vary from week to week for each player or team. For example the completion percentage will vary for any quarterback from game to game as will the number of passing yards. **Regression to the mean refers to the fact that extreme observations of these statistics for any given athlete or team are more likely to be followed by observations closer to the average for the athlete/team than by extreme observations.**

3.1. Distribution of Player Statistics. The graphs shown below show the distribution of the statistics field goal percentage (X, Y on the left) and the number of points scored (Z, W on right) per game for two NBA players, Michael Jordan and Scottie Pippen. The distributions were compiled using data from ESPN for over 200 games collected by Peter Ulrickson, a graduate student at Notre Dame.




From these pictures, which fit the profile of mound shaped distributions discussed in the last section, we see that the mean and standard deviation for any given statistic will vary from player to player. In the examples shown above, we have:

- For $X = \text{FG}\%$ per game for M. Jordan, $E(X) \approx 0.48$ and $\sigma(X) \approx 0.1$.
- $Y = \text{FG}\%$ per game for S. Pippen, $E(X) \approx 0.46$ and $\sigma(X) \approx 0.13$.
- $Z = \text{points per game}$ for M. Jordan, $E(Z) \approx 29.59$ and $\sigma(W) \approx 8.44$.
- $W = \text{points per game}$ for S. Pippen, $E(W) \approx 19.67$ and $\sigma(X) \approx 7.28$.

We see that some players may play at a consistently high level throughout their career with a high average and low standard deviation. On the other hand, some players may have relatively low career averages by comparison with their peers, but may have enough variation in performance to reach high levels of play 20% of the time or 10% of the time. These exceptional performances may be enough for them to surpass the performance of their peers in the short term, but (because of the nature of expected values) their average performance should regress to their overall average in the long term.


Lets suppose that X is some statistic which measures performance for a given athlete or team. In fact ranking experts often combine statistics such as those shown above to get a single overall measure of performance for the athlete. Lets assume that X is such a statistic for our given athlete or team with mean $\mu = E(X)$ and standard deviation $\sigma = \sigma(X)$. Lets also assume that for our given athlete or team, the values of the random variable X observed have a distribution similar to the mound shaped distribution shown below. Recall our empirical rule for mound shaped distributions in the last section, we know that the observed values of a variable X with a mound shaped distribution are distributed (roughly) in the following manner:



where $\mu = E(X)$ and $\sigma = \sigma(X)$. For this athlete/team we will observe values of X that are above $\mu + \sigma$ about 16% of the time and below that value 84% of the time. This means that for any observation of X above $\mu + \sigma$, the probability that it will be followed by something less than $\mu + \sigma$ is 0.84. An exceptional performance that results in an observation of X more than two standard deviations above the mean (above $\mu + 2\sigma$) has a 95% chance of being followed by a decrease in performance. In fact, the more outstanding the performance for the athlete, the greater the probability that it will be followed by a decrease in performance. Dramatic decreases in performance are more likely for athletes or teams with high standard deviations.

Note Reasoning is similar for distributions with other shapes. No matter what shape the underlying distribution of a statistic is the average over a sufficiently large number of independent trials has a mound shaped distribution according to the [Central Limit Theorem](#).

Example 3.1. The [total quarterback rating](#) is a statistic which gives a measure of the overall performance of the quarterback for each game. It ranges from 0-100 with an average quarterback scoring 50.




(a) Let X be the quarterback rating for Quarterback A. Lets assume that X is a variable with a mound shaped distribution with mean $E(X) = 70$ and standard deviation $\sigma(X) = 9$. Suppose that Quarterback A has a quarterback rating above 85 for three games in a row, what is the probability that his quarterback rating for the next game will be at or below 61?

(b) Let Y be the quarterback rating for Quarterback B. Lets assume that Y is a variable with a mound shaped distribution with mean $E(Y) = 75$ and standard deviation $\sigma(X) = 5$. Suppose that Quarterback B has a quarterback rating above 85 for three games in a row, what is the probability that his quarterback rating for the next game will be at or below 60?

3.2. Fluctuation around the mean. As mentioned above, it is difficult for the observer to distinguish between short term runs of good performance due to randomness and a long term improvement in performance (which amounts to a shift in the density function for the player or team). For many athletes/teams exceptional performances due to random fluctuation in statistics have a very high probability of being followed by a large decrease in performance. Athletes and teams with consistently high performances are less likely to suffer a large decrease in performance after an outstanding performance. One would expect that those athletes/teams with high and consistent levels of performance would have repeated appearances on the cover of Sports Illustrated and their appearance would be less likely to be followed by a misfortune in the form of a dramatic decrease in performance. On the other hand one would expect an individual or team with a relatively high, but not exceptionally high average performance coupled with a large variance in performance to make it to the cover of Sports Illustrated perhaps once in his/her/their career but with a high probability of a dramatic decrease in performance shortly thereafter.

Example 3.2. Lets look at some data for NBA player Jeremy Lin who was featured on the cover of Sports Illustrated (and many other magazines) twice in February 2012 after a period of exceptional performance averaging over 25 points per game with the New York Knicks . The graph extends over a period from the beginning of the 2011-2012 season to the present (2015). It uses the results of 186 consecutive games in which he played collected from the ESPN website by Peter Ulrickson, a graduate student at Notre Dame. In the graph on the left, we show a moving average, the average number of points scored per game for the previous five games in which he played , starting with the fifth game in which he played in the 2011-2012 season on Feb. 14 2012. On the right, we also have a moving average. We first calculated the number of points per


minute for each game and then took the average for the previous 5 games starting on Feb. 14 2012. This takes into account the amount of actual playing time he has had in each game, which varies considerably.



The average for the data graphed on the left is approximately 13.43 and the standard deviation is approximately 4.07. The average for the data graphed on the right is approximately 0.44 and the standard deviation is approximately 0.103. We see that as time progresses the statistics fluctuate around the mean with sharp increases and drops in the graph reflecting periods of poor performance followed by dramatic increases in performance and vice versa. We see that most of the time the data is close to the mean (within 2 standard deviations) with a few exceptional periods performance followed by a regression to the mean.

If a player has a greater variance in performance, we can expect larger fluctuations around the mean as demonstrated in the results of the following simulation:

Example 3.3. We used excel to simulate the quarterback rating for two hypothetical quarterbacks over the course of 100 games. We assume that both quarterbacks, Quarterback 1 and Quarterback 2 had the same average quarterback rating at 75. Quarterback 1 had standard deviation of 5 and Quarterback 2 had a higher standard deviation of 10. We generated random values of the quarterback rating for each player for 100 games and then calculated the running average for the previous 5 games starting at the the fifth game as we did with the data for Jeremy Lin above. Below we see the averages for both players plotted on the same graph with the performance of Quarterback 1 shown in red and that of Quarterback 2 shown in blue.



We see that every now and then Quarterback 2 has an exceptional performance followed by a sharp drop in performance, all due to random variation. On the other hand the changes in performance in the data for Quarterback 1 are not so dramatic. Thus we see that large changes in performance are to be expected for players who exhibit more variation in play and should not be considered a result of bad luck.

4. CONCLUSION

Lets see how the above pieces might fit together to explain the levels of “misfortune” cited in the study for the cover subjects of Sports Illustrated. Lets consider the cover subject’s path to misfortune as a two step process. The first step is being chosen by the editors to appear on the cover and the second is what happens in the period following their appearance. Below I divide the subjects into distinct categories and then attempt to estimate the proportions of cover subjects in each category and the (conditional) probability that misfortune will befall them. It is not intended to be a scientific study, rather it is intended to demonstrate that the proportion of misfortunes cited in the article could have occurred in an entirely reasonable way if one bears in mind the points made above about regression to the mean, the inclusion of a wide range of coincidences and the unreliability of predictions. You are encouraged to substitute your own probabilities instead of mine if you think they are more fitting or if you wish to undertake a more scientific study of magazine covers and those athletes/teams featured.

First, lets categorize the cover subjects into the following non overlapping categories:

- P = predicted to win a in the near future or claimed to be the best in their sport with the expectation of winning in near future.
- RHPC = No predictions or claims of “best” involved, chosen for recent high performance in their sport and fit in the category of consistent high performers(small variance).
- RHPI = No predictions or claims of “best” involved, chosen for recent high performance in their sport and fit in the category of relatively large variance in performance.
- SE = something else

Next lets estimate the percentage of cover subjects that fall into each category. Unless one looks through all of the covers, exact percentages are impossible. I have outlined how I made my estimates below and you should feel free to substitute your own.

- P(10%): from browsing through covers for a few years, I estimate that 1 in 10 of the covers are devoted to prediction or have claims of a team or individual being the best in their sport.
- RHPC (20%) : from searching on wikipedia for the number of subjects who had repeat appearances (I assume some intersection between individual subjects and team subjects) I would suggest a ballpark figure of 500 cover subjects out of the 2,500 or so to date appeared on the cover more than once and most likely had a high level of consistency in performance. Again we are assuming that these are not among those already counted in P above.
- RHPI (60%) : Assuming that the category SE is small, say 10% we will put the remainder of the cover subjects in this category.
- SE(10%)

Now let’s estimate the probability of “misfortune” (M) for each category.

- P: Lets assume that the analysts at Sports Illustrated are among the best and that their predictions are correct 75% of the time, this means that the probability of “misfortune” of the failed prediction type for this category is 0.25. Of course this would increase if we added in team losses and injuries etc...
- RHPC: these are the consistent athletes with low probabilities of a dramatic decrease in performance. Lets assume a very conservative estimate of their probability of a noticeable decrease in performance at 0.05. Their probability of “misfortune” which would also include team performance would be somewhat higher than this.
- RHPI For these athletes and teams, they have a relatively high variance in performance. Being conservative, lets assume that most are operating at an average of one standard deviation above their own mean when appearing on the cover of Sports Illustrated. The (average) probability that they will perform below their own mean in the next game is about 50% (exactly 50 for our mound shaped symmetric distribution). Of course the probability that they will meet with misfortune will increase if we throw in team performance and injury, however lets remain conservative with a probability of misfortune at 0.5 for this category. (Note that using variations in team performance

to account for misfortunes also makes this category bigger and makes the probability of misfortune greater).

- SE: Here I use statistics from everyday life. For any given exam about one student in 300 will have some random injury, sickness or misfortune causing them to miss the exam. So I will estimate the probability of “misfortune” for this category at about $1/(300) \approx 0.003$.

The following table shows the percentage of all cover subjects that would experience “misfortune” from each category as a result of our somewhat conservative speculation:

Category	Percentage of covers in Category	Prob. Misfortune	% misfortunes from category
P	10%	0.25	2.5%
RHPC	20%	0.05	1%
RHPI	60%	0.5	30%
SE	10%	0.003	0.03%
			33.53%

The percentage of all cover subjects who are expected to meet with misfortune as a result of this speculation sums to 33.53%. I could revise it to get 37.2%, but given that I have been conservative with the estimates throughout, I do not find that 37.2% is an unnaturally high percentage of cover subjects to meet with misfortune. In a more scientific study one could take a random sample of say 200 covers and use the data to get better estimates for the percentages and probabilities above. One would also need to look at the statistics for the cover subjects to determine variability of play to find the probability estimates for those in the category RHPI and RHPC to complete the calculation. Hopefully this exercise has at least given you some tools with which you can think logically about the “curse”.

REFERENCES

1. Amy N. Langville & Carl D. Meyer, *Who's # 1*, Princeton University Press (2012).