Uncovering the Mysteries of the Knuckleball

by Alan M. Nathan

The knuckleball is perhaps the most mysterious of baseball pitches. It is thrown slower than normal pitches, it is barely spinning, and it appears to float to home plate. Its motion seems to be very unpredictable by anyone: the batter, the catcher, and even the pitcher himself. Moreover, there is a common perception that a knuckleball does not follow a smooth path to home plate, but instead "dances," "flutters" or "zigs and zags," with seemingly abrupt changes of direction. It's no wonder that it is so difficult both for the batter to hit and for the pitcher to command. Indeed, very few pitchers in the history of Major League Baseball have had much success mastering the knuckleball. That makes the recent success of R.A. Dickey, winner of the 2012 National League Cy Young Award, all the more impressive.

As a result of Dickey's success in the past few years, there has been renewed interest in knuckleballs in the baseball world. Indeed, many recent articles have been written in both the mainstream media and the blogosphere attempting to answer the question: What makes an effective knuckleball? Moreover, Dickey's success has also driven new efforts to train knuckleball pitchers early in their careers. One such effort is the International Knuckleball Academy, a training ground based in Vero Beach, Fla., where young pitchers can learn to master the craft of knuckleball pitching. The ultimate goal is to bring knuckleball pitching into the mainstream of baseball.

It is perhaps a very fortunate accident of nature that Dickey's success with the knuckleball has occurred during an era in which pitch-tracking technologies have revolutionized our ability to analyze pitching. These technologies, PITCHf/x and TrackMan, allow an analysis of the trajectories of all pitches, whether knuckleball or otherwise, with unprecedented precision. As a result, we are now able to quantify both the magnitude and the direction of knuckleball movement, to compare it to that of ordinary pitches, and to investigate its erratic behavior. These technologies also allow us to investigate more detailed aspects of the trajectory, allowing us to quantify the extent to which the common "zig-zag" perception is supported by the data. A parallel development has been wind-tunnel experiments to investigate the aerodynamic forces on a knuckleball and how they depend on the seam orientation. Taken together, the trajectories and the forces are leading to a better understanding of how a knuckleball behaves and why it does what it does.

In this article, I will discuss much of the recent developments, starting with a brief review of the new technologies. Then I will focus on how those technologies are used to gain new insights into the character of knuckleball trajectories, including its movement. I will then attempt to reconcile those trajectories with what we know about the forces on knuckleballs from wind tunnel experiments. I will conclude with a detailed case study that shows great promise in uncovering the mysteries of the knuckleball.

Review of Pitch-Tracking Technologies

What we know about knuckleball trajectories comes from two relatively new technologies that are utilized in major league ballparks these days. One of these is PITCHf/x, a video-based tracking system that is permanently installed in every major league ballpark and has been used since the start of the 2007 season to track every pitch in every MLB game. The system consists of two cameras mounted high above the playing field that determine the location of the pitched baseball over most of the region between release and home plate at 1/60-second intervals. Those positions constitute the "raw trajectory," which then is fitted to a smooth function using a constant-acceleration model, so that nine parameters (9P) determine the trajectory: an initial position, an initial velocity, and an average acceleration for each of three coordinates. From the 9P description, all of the relevant parameters for baseball analysis can be determined, including the release velocity, the location at home plate, and the movement. All the 9P information for each pitch is freely available to the public.

The other new technology is TrackMan, a phased-array Doppler radar system that is installed in several major league ballparks. It determines the full pitched-ball trajectory from release to home plate in 1/250-second intervals. Much like PITCHf/x, all of the relevant parameters for baseball analysis are determined by fitting the raw trajectory data to a smooth function, which is proprietary but probably not very dissimilar to the 9P fit used in PITCHf/x. TrackMan data are not publicly available, due to their business relationships with more than half of the major league clubs.

The Character of Knuckleball Trajectories

To study the character of knuckleball trajectories, it is necessary to have access to the raw trajectory data rather than the 9P or similar smooth-function fit to the data. Such a fit is entirely appropriate for studying ordinary pitches based on what we know about the aerodynamic forces that governs the motion of those pitches. What we know is that these forces are nearly constant over a typical trajectory, so the 9P model is an excellent approximation of the actual trajectory. An even better approximation would be to let the aerodynamic forces be proportional to the square of the velocity. I call this the 9P* model and will discuss it further below. We know far less about the forces responsible for the knuckleball, so we really do not know how well it is described by the 9P or 9P* models. Indeed, those are precisely the issues we want to investigate, and that requires access to the raw tracking data. Accordingly, I obtained raw PITCHf/x data from four high-quality starts by either Tim Wakefield or R. A. Dickey during the 2011 season:

- May 1, 2011, 266 total pitches (70 knuckleballs by Wakefield)
- June 19, 2011, 278 total pitches (96 knuckleballs by Wakefield)
- Aug. 29, 2011, 201 total pitches (77 knuckleballs by Dickey)
- Sept. 17, 2011, 235 total pitches (87 knuckleballs by Dickey)

Each pitch was tracked over the approximate region of seven-to-45 feet from home plate, resulting in about 20 sequential points on the trajectory, a few more or less depending on the release speed and precise tracking region. The 9P* model was used to fit each trajectory to a smooth curve.

One way to characterize the smoothness of the trajectory is to compare the fitted smooth curve to the actual data. The distribution of root-mean-square (rms) deviations of the data from the smooth function, in units of inches, is shown in the box plots for each of the four games, where the white and grey boxes are for ordinary and knuckleball pitches, respectively.

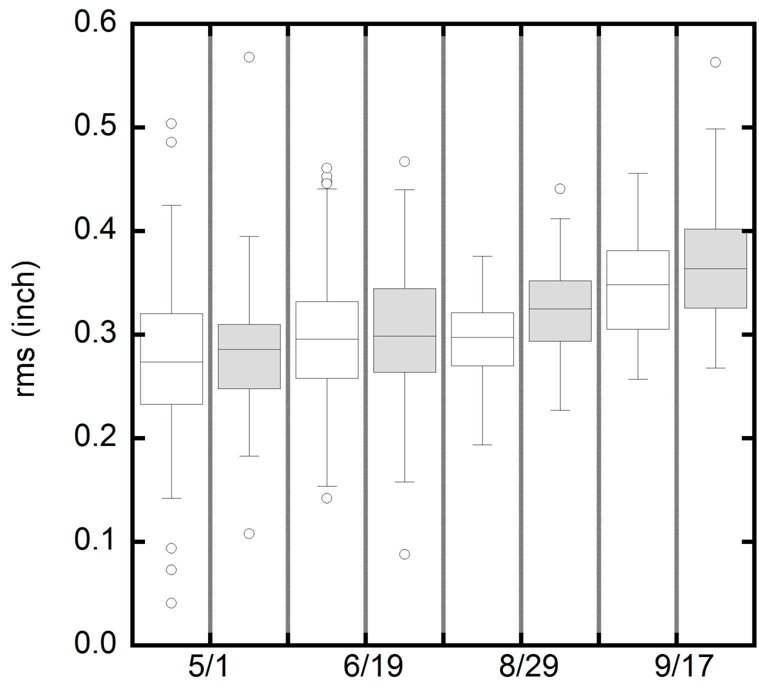

For ordinary pitches, which are expected to be very accurately described by the 9P* curve, the non-zero rms values are almost certainly due to random noise that is inherent in any measurement system. Taking median values as our estimate, that noise is approximately 0.3 inches per point on the trajectory, a very small value. Indeed, one would be very hard pressed to obtain tracking data more precise than these. There are small differences from one game to another that probably can be attributed to differences in camera placement and/or calibration.

But the truly remarkable thing about these plots is that they show quite clearly that the rms distributions for knuckleballs are not very different from those of ordinary pitches. Relative to ordinary pitches, the knuckleballs have a slightly higher median value (the horizontal line in each box) and a very similar interquartile distance (shown by the height of the box). We end up with the unavoidable conclusion that within the overall precision of the tracking system (~0.3 inches), knuckleball and non-knuckleball trajectories follow a 9P* curve nearly equally well. A similar conclusion is reached if a constant acceleration (9P) model is used.

To demonstrate this point more vividly, I now show a bird's-eye view of four different trajectories from the Aug. 29 game.

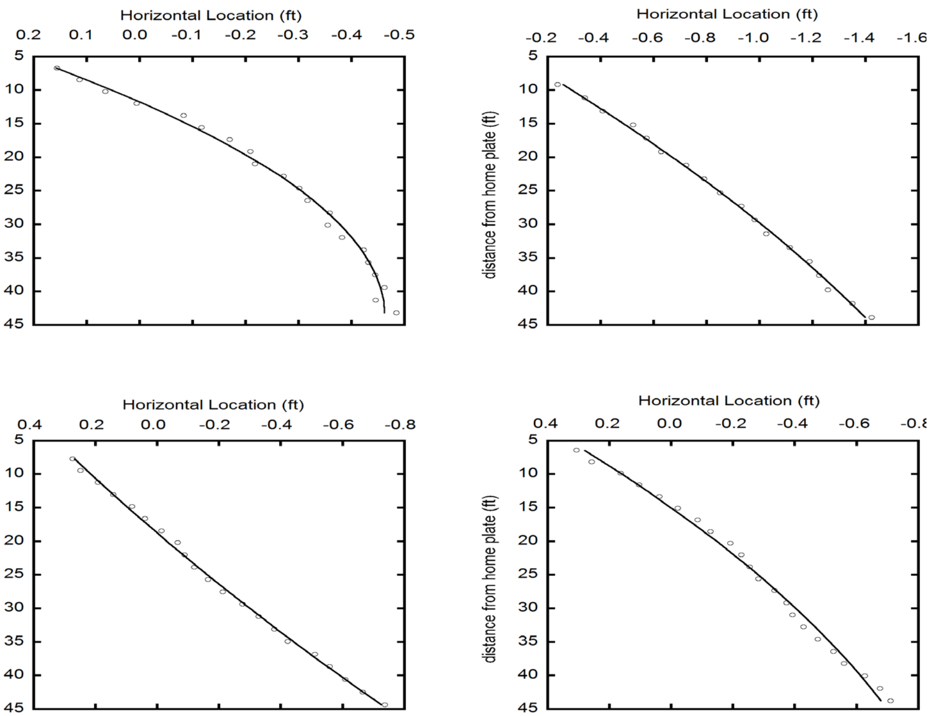

The upper two are ordinary pitches, the lower two are knuckleballs. The left side shows a trajectory with rms close to the median; the right side shows a trajectory with the largest rms value for that game. Note that the release point is near the bottom of each plot, home plate near the top. The first thing to note is that, on average, the curve follows the general trend of the data, with maximum deviations no more than half an inch. This goes a long way toward justifying the use of the 9P or 9P* fit for baseball analysis. More on this below. However, while the curve describes the average trend of the data, a closer inspection shows an interesting feature for the lower-right plot, the knuckleball with the largest rms value. Despite not deviating from the curve by more than half an inch, the data do not randomly fluctuate but rather appear to oscillate about the curve, making one full oscillation over a distance of about 25 ft. This behavior suggests that the pitch undergoes several distinct changes of direction, none of them very large. Despite ending up in the heart of the strike zone, this particular pitch was taken for a called strike, apparently fooling the batter completely. I will return to an analysis of a similar pitch below. Despite the unusual behavior of this pitch, I emphasize that it is definitely the exception rather than the rule, with the vast majority of pitches from these games following the smooth curve very closely.

The Movement of a Knuckleball

Next, I want to investigate the movement of a knuckleball, where "movement" is defined as the deviation of a pitch from a straight line with the effect of gravity removed. To appreciate why a knuckleball is so different, it is helpful to place it in context by discussing what is normal. All pitches, whether normal or otherwise, slow down due to air resistance, losing about eight percent of their initial speed by the time they cross home plate. They also deviate from a straight-line path due to the combined effects of gravity, which pulls everything down, and the other "aerodynamic forces." For normal pitches, which are spinning rapidly, the aerodynamic force causing the movement is called the Magnus force. The strength of the Magnus force increases as the spin rate increases. The direction of the Magnus force is such as to deflect the ball in the direction that the front edge of the ball is turning, as seen by the batter. Thus, balls thrown with backspin deflect upwards, those thrown with topspin deflect downward, and those thrown with sidespin deflect left or right. And, of course, anything in between is possible, depending on how the spin axis is oriented. For example, a two-seam fastball thrown by a right-handed pitcher will deflect up and to the catcher's left.

Now let's take a look at the movement of a subset of Dickey's pitches thrown during the 2012 season. They are shown on a polar plot, where the distance from the center is the release speed at 55 ft., and the angle is the direction of movement, as seen by the catcher.

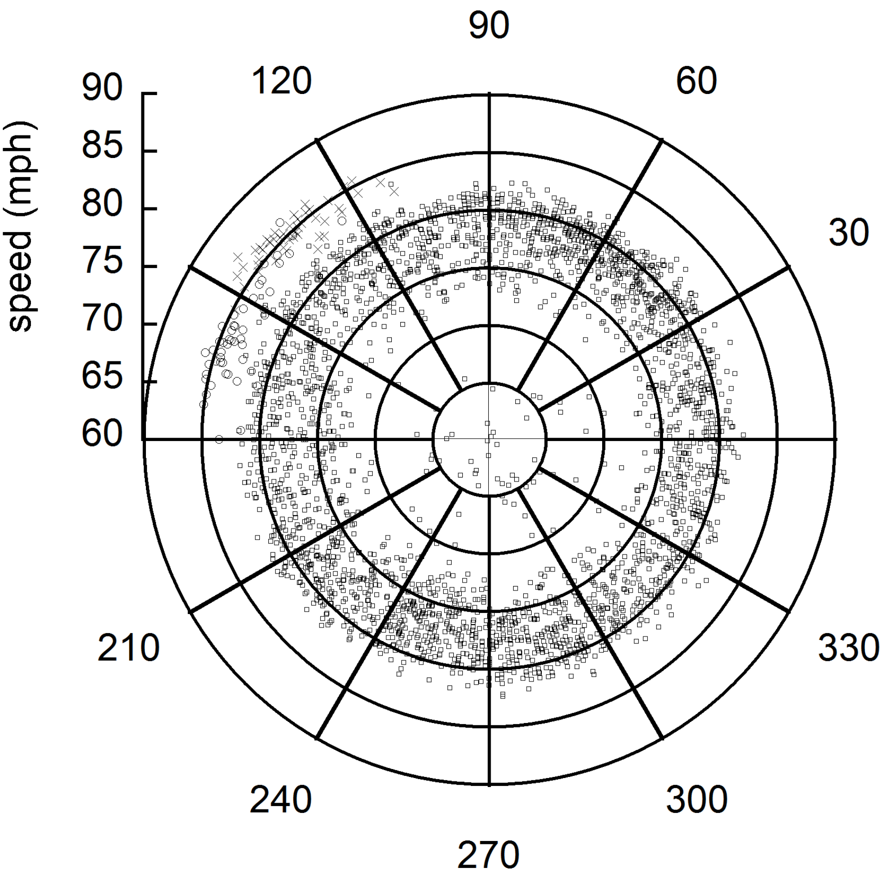

The ordinary pitches fall into two categories: four-seam fastballs denoted by the x's and two-seam fastballs (or sinkers) denoted by the o's, both of which are thrown at mid-80s speed. Relative to the four-seam, which has mainly upward and a little arm-side (or tailing) movement, the two-seam has less upward and more tailing movement. These features are all rather standard fare for anyone who has ever done this type of analysis. The important feature of these pitches is the clustering behavior of the movement. That is, the angular range for the movement of a given pitch type is limited in extent and, therefore, rather predictable. This predictability conforms with what we know about how pitchers pitch. Namely, for a given pitch type, a pitcher tries to release the ball with a consistent axis of rotation. Our understanding of the Magnus force tells us that if the axis of rotation is consistent, then the direction of movement also will be consistent.

Now take a look at the knuckleballs, denoted by the squares. In contrast to the ordinary pitches, there is no clustering of the knuckleball movement. The knuckleball seems to be erratic in the sense that the direction of movement is entirely random from one pitch to another and, therefore, unpredictable. Another way to say the same thing is that whatever aerodynamic force is acting on the knuckleball to cause the movement seems to have a random direction from one pitch to another. This is one feature of knuckleballs that we will try to understand in the context of the wind tunnel experiments.

Let me now summarize what we have learned from analysis of trajectories:

- To within our ability to measure with the tracking systems, knuckleball trajectories mainly follow a trajectory as smooth as that of ordinary pitches. There are exceptions, but they constitute a relatively small percentage of knuckleball pitches.
- The movement of knuckleballs is erratic and unpredictable.

I next turn to a discussion of wind tunnel experiments to see the extent to which they can explain what we observe from trajectory analysis.

Reconciling Wind Tunnel and Trajectory Data

Wind tunnel experiments are used to measure the aerodynamic forces on a baseball (or any other sports ball) as a function of wind speed, spin, etc. Much of what we know about air resistance and the Magnus force on sports balls comes from such experiments. For the knuckleball, the classic experiment was done nearly 40 years ago by Robert Watts and Eric Sawyer (*http://baseball.physics.illinois.edu/WattsSawyerAJP.pdf*). They investigated how the aerodynamic forces on a non-spinning baseball depend on seam orientation. A simplified version of their findings is shown on the plot, where positive/negative forces corresponds to up/down on the diagram of the ball, which also shows how the angle is defined.

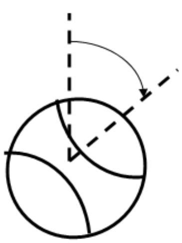

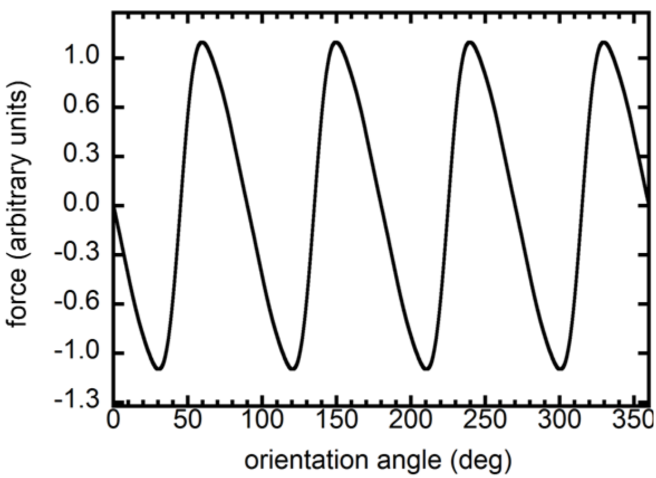

The wind blows from right to left, and as the orientation of the ball was changed, the seam pattern maps out the four-seam configuration. What we learn from this is that the size and direction of the force is strongly dependent on the seam orientation. Indeed, aerodynamicists tell us that when the on-coming air encounters a seam, the character of the air flow is altered in such a way as to deflect the ball one way or another, depending on the orientation of the seam relative to the air flow. The result is a sine-wave-like pattern that is repeated four times in one revolution, exactly what one would expect for the four-seam configuration. More recent data taken in the past year confirm this finding. Moreover, the more recent data also investigated the two-seam configuration. Strangely, that pattern is much more complicated and not yet fully understood. Therefore, pending a better understanding of that pattern, I will confine my investigations to the four-seam configuration. Note that the force scale is shown in arbitrary units. However, recent data confirm that the size of the knuckleball force at its peak is comparable to the Magnus force on a four-seam fastball.

Armed with a force-vs-orientation profile, it now is possible to do mathematical simulations of knuckleball trajectories to see how they compare with actual data. Some results are presented in the plots, which show a bird's-eye view similar to that shown of the actual data but with the horizontal scale in inches.

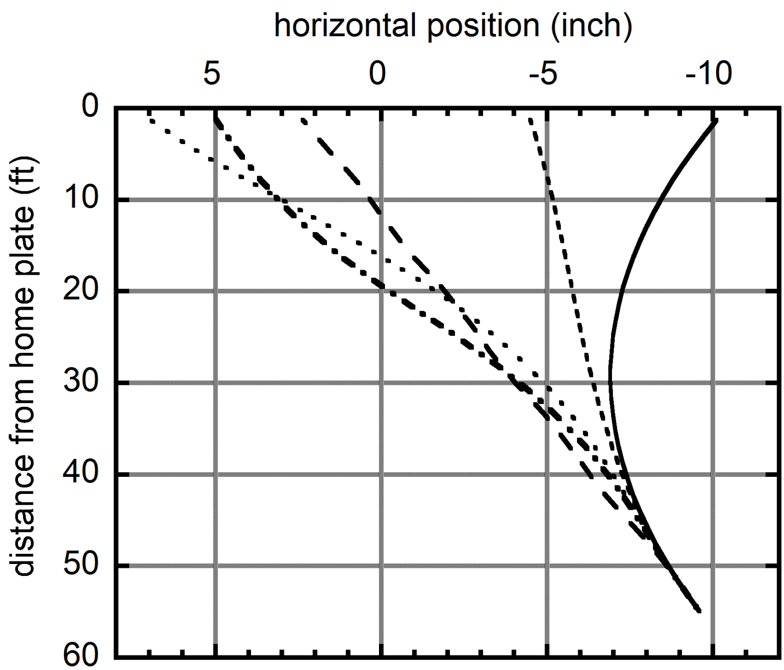

All pitches are thrown with the same release point, speed (75 mph), direction, and initial seam orientation (780). However, they are thrown with different rotation rates, given as the angle through which the ball turns between release and home plate: solid line 00; short dashes 180; dots 720; dash-dot 990; long dashes 3600. Therefore, all five trajectories start out at the same point and are heading in the same direction. However, their eventual fate differs greatly. At one extreme is the pitch corresponding to 00 (no rotation). In that case, the force is constant both in magnitude and direction, much like that of ordinary pitches, resulting in the largest amount of movement. The other extreme is the one corresponding to 3600 (one full revolution), so that the force profile undergoes four complete oscillations. In that case, the force is changing in direction so rapidly that the ball cannot keep up with it, so the ball ends up going more or less straight, i.e., no movement, just as shown in the plot. That confirms the well-known fact about knuckleballs that too high a rotation rate results in a considerable loss of movement. Between these two extremes, the simulations show that the character of the trajectories is a very strong function of the rotation rate, with small changes in the rate leading to large changes in where the ball crosses home plate. Although no plots are shown, the trajectory is also sensitive to small changes in the initial orientation of the seams. This sensitivity to small changes in the initial conditions is an example of chaotic dynamics and almost surely goes a long way toward explaining why the movement of the knuckleball is unpredictable.

Now take a closer look at the dash-dot trajectory, for which the force undergoes a bit more than one complete oscillation. The net effect is quite evident in the trajectory, as the ball initially breaks to the left, then breaks to the right about 20 ft. from home plate. The net movement is not great, only about 2.5 in. (compare with the long dashes), but this pitch is quite likely to fool both the batter and the catcher. Of course, it is only a simulation. Moreover, it utilizes only the four-seam data, whereas most knuckleballs these days are thrown in the two-seam orientation. Nevertheless, pending a more complete understanding of the two-seam force profile, it is encouraging that some of the same features observed in actual tracking data also show up in the simulation. To emphasize that point further, I will now turn to a detailed analysis of an actual pitch.

Anatomy of a Really Nasty Pitch

On June 13, 2012, Dickey pitched a one-hitter against the Tampa Bay Rays at Tropicana Field. In the bottom of the third, he struck out Will Rhymes with a truly nasty pitch. Modern technology has given us two different tools for analyzing this pitch. First, a high-speed video allows us to see the motion of the ball, including its rotation, in great detail. Second, the TrackMan system is set up in Tropicana Field, allowing a rather precise measurement of the entire trajectory of the pitch from release to glove. So let's take a look at how those two compare with each other and see if we can understand the trajectory based on what we know about the knuckleball forces.

First take a look at the animated gif of the high-speed video at *http://baseball.physics.illinois.edu/KBall/dickey-rhymes-hd.gif*, if you can. Being careful not to be fooled by the upward movement of the camera or the movement of the catcher, see if you can notice that the pitch undergoes two distinct changes of direction. Shortly after release, it appears to swerve to Dickey's right; later it swerves back to the left. Poor Will never had a chance. In fact, even the catcher has a lot of trouble with the pitch, as we can see from the movements of his head and glove. There seems to be no question that this pitch gets very high marks on the nasty scale.

The other interesting thing to note on the video is the spin, which is easy to follow by watching the progress of the MLB logo on the ball. Dickey has said that he tries to throw the ball in the two-seam orientation with about 0.5 revolutions of topspin. This pitch is nothing like that. Instead it undergoes about 1.5 revolutions between release and catch with nearly pure gyrospin; that is, the rotation axis is about the direction of motion, much like a spiral pass in football or a bullet fired from a rifle. If nothing else, this demonstrates the difficulty for even a skilled pitcher like Dickey in reproducing the initial conditions.

Now let's take a look at the TrackMan data, shown as the circles in the plot.

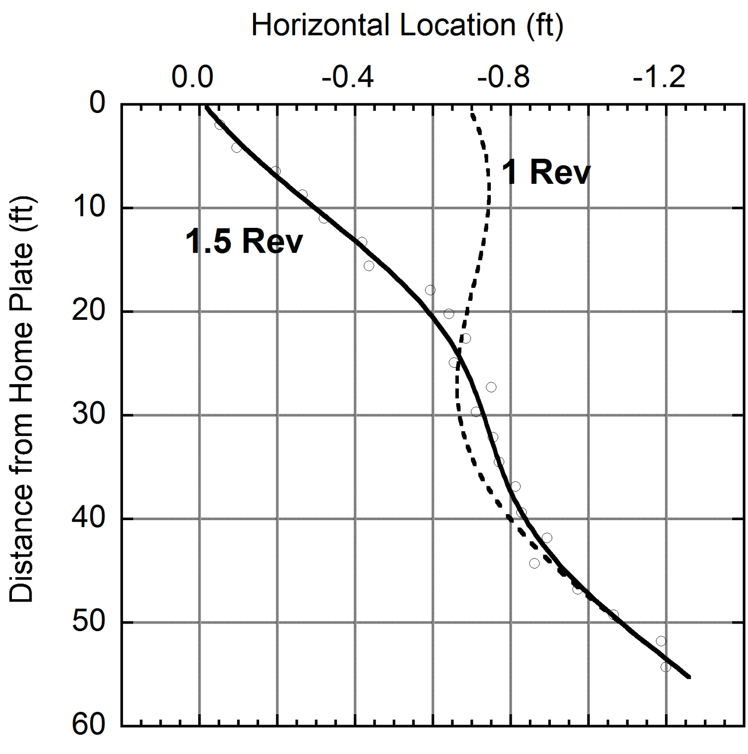

One can clearly see the same features that are observed in the video, namely, an initial break to the right followed by a second break to the left. Given that the ball is thrown with gyrospin and given the symmetry of the ball, the force that causes the movement must be constant in magnitude. Moreover, the direction of the force must rotate along with the ball about the direction of motion, initially pointing in the upward direction, then rotating clockwise as seen by the pitcher (counterclockwise as seen by the catcher). That makes it very easy to simulate this pitch with three unknown quantities that are adjusted to fit the data: the rotation rate, the initial direction, and the magnitude of the force. The result is shown as the solid curve, which does a remarkably good job at accounting for the tracking data. Now we come to the key point: Given a complete mathematical model that describes this pitch (the solid curve), we now can investigate what would happen with slightly different release conditions. So, for example, the dashed curve shows the trajectory expected if everything were exactly the same except the rotation rate, making 1.0 rather than 1.5 revolutions. The difference between the two curves is stark, with the actual trajectory going right over the middle of home plate, whereas the alternate trajectory hitting the outside corner, a difference of over eight inches. Indeed, judging from the movement of the catcher, the alternate trajectory was expected, whereas the actual trajectory seemed to break toward the middle of the plate at the last moment. By the way, there are comparable multiple breaks in the vertical direction that I am not showing, since they are masked by the much larger effect of gravity.

There is no question that this pitch is extremely effective at fooling the batter. Was it thrown this way intentionally by Dickey? Has he discovered a new "gyroball knuckleball?" Should he try to throw this pitch more often? Suffice it to say that I would love to know the answers to these questions.

Summary and Outlook

With the combination of precise tracking data, wind tunnel experiments, and high-speed video, we are slowly but surely uncovering the mysteries of the knuckleball. We now have a working hypothesis for the erratic nature of knuckleball movement, having to do with small changes in the release conditions. We understand that knuckleball pitches with multiple changes of direction are the exception rather than the rule, at least to within the precision of the tracking data. We further understand that the amount of movement is not necessarily indicative of the nastiness of the pitch, since a pitch thrown with multiple changes of direction—one guaranteed to fool the batter—necessarily has a net movement that is small. We understand why too high a rotation rate has the opposite effect, not fooling the batter at all, given that such a pitch travels pretty much in a straight line. And we have started to make progress reconciling these features of the actual tracking data with wind tunnel experiments.

There still are many things we would like to understand better, particularly the force profile in the two-seam orientation. Indeed, we really would like to understand the forces with the ball oriented in an arbitrary configuration. So there is still plenty of work to do. Nevertheless, the progress made thus far is very pleasing.

It is a pleasure to thank Sportvision and TrackMan for supplying me with raw tracking data. I also acknowledge many fruitful discussions with Professors Rod Cross, Robert Watts, and John Borg.