3

THE MATHEMATICS OF CHANGE

A New World Model

What do Aesop's fable about a race between a tortoise and a hare, the Dow Jones industrial average, and a Formula 1 car have in common? They all deal with change—in the distance run by two animals in a storybook race, in a stock market index (a measure of the performance of a widely held group of stocks[1]), and in the speed of a car. Fortunately, to analyze the data in all three cases, the mathematics is the same.

3.1. The Measure and the Change in the Measure

Perhaps the most fabled race in all folklore is that of the tortoise and the hare. Aesop gives little information about the race, but knowing only that the tortoise traveled at a constant speed (as tortoises do) and that the hare raced ahead and fell asleep (as hares seemingly do), we can plot their data in a distance versus time graph to see what conclusions can be reached (see Figure 3.1). This is not unlike how one might think through the various routes to or from work, factoring in traffic bottlenecks and speed traps along the way.

Since Aesop provides no length or time for the race, I have had to make some assumptions. First, I have assumed a 2-meter/minute tortoise, which is reasonable for a real-life tortoise and is equivalent to about 0.1 miles per hour.[2] I have also assumed a 200-meter course, which thus takes 100 minutes for our steady, 2-m/min tortoise to complete and also seems reasonable given how long an average hare might nap. Last, I have assumed that the hare runs as fast as the fastest human,[3] at about 600 m/min, and that he fell asleep at the midpoint, or 100 m.

Figure 3.1 Distance versus time for the tortoise and the hare

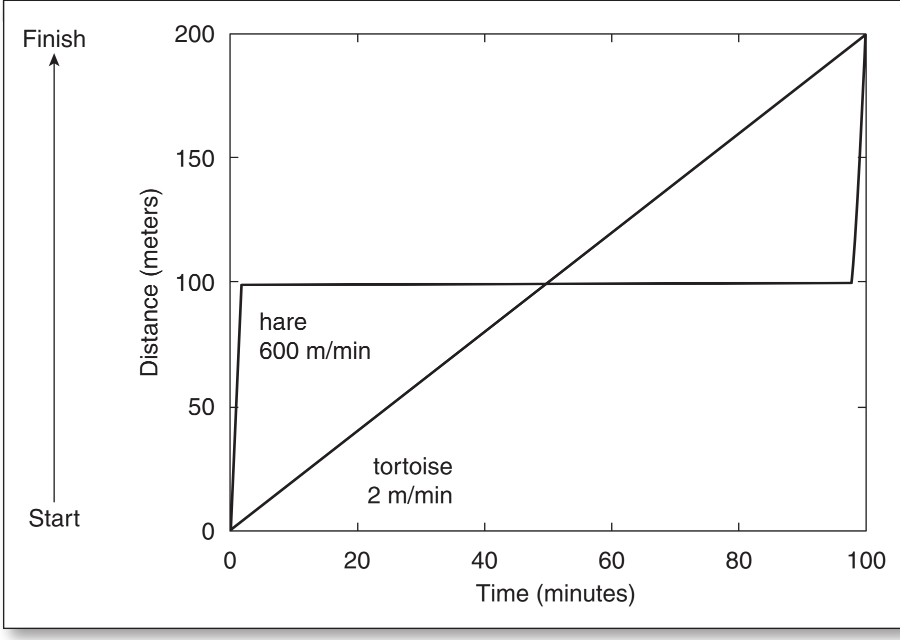

With the data at hand, we can now make a few important observations. I had always thought the race was a longer distance over a shorter time. I don't know if this was because of Disney's cartoon version, but if the race was, say, of marathon length (26 miles, 385 yards), a 2-m/min tortoise would take more than 2 weeks (without rest) to finish, and no hare would likely nap for that long (well, maybe a talking Disney cartoon hare would)—not to mention the patience the referee fox would need. Perhaps the hare hibernated during the race (Disney again?), happily dozing for 2 weeks, but that would only complicate the fable. As well, a faster race (say, the length of a cartoon) would require a faster tortoise, which is not only unrealistic but defeats the idea of the fable. Suffice it to say, the race was over a reasonable distance such that a steady tortoise could finish in a realistic time and a haughty hare could fall asleep and awake in time (well, almost in time).

It would also seem from the numbers that our hare was haughtier than originally supposed—he was so self-sure and cocky that he couldn't keep it together for 200 meters, or about 20 seconds. Perhaps our hare had a short attention span, one possible interpretation of the fable. Or perhaps he was more stupid than haughty, foolishly tiring himself out in a sprint when a slow but steady jog would have sufficed, given the unlikely competition. Perhaps he just overestimated his resources. Of course, we don't need exact numbers to understand a fable, but it does seem that retelling the story with reasonable numbers makes it even more of a cautionary tale about being haughty, distracted, and stupid.

But with such real-world numbers, we can now understand more about how data changes in time—for example, in an airplane flight data recorder, where the distance and speed data can reveal the cause of a crash; in inflation indices, where price instead of distance is measured; or in the stock market, where extreme changes in the price of a stock can result in huge windfalls or losses.

Let's start with the tortoise, since she is the easier to analyze. As seen in Figure 3.1, the tortoise's distance is as straight as an arrow: 20 meters after 10 minutes, 100 meters after 50 minutes, and 200 meters

Figure 3.2 Derivatives explained for the tortoise and the hare

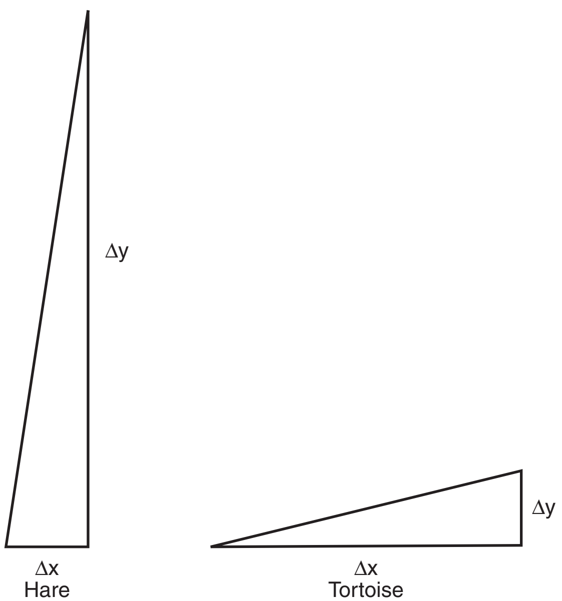

$$\frac{\Delta y}{\Delta x} = \frac{(100 - 0)\text{ m}}{(1/6 - 0)\text{ min}} = 600\text{ m/min}$$

$$\frac{\Delta y}{\Delta x} = \frac{(200 - 0)\text{ m}}{(100 - 0)\text{ min}} = 2\text{ m/min}$$

after 100 minutes, which is represented as $d = 2t$ (d is distance, and t is time). The hare's speed was also assumed constant, though in spurts, and is seen as two spikes in his d-t graph. The hare raced ahead (greater slope), slept (no slope), and madly dashed at the end to catch up (greater slope again). Since he raced his first 100 meters in 10 seconds (600 m/min), this is represented as $d = 600t$, i.e., 300 times faster than the tortoise.

Most of us are familiar with slopes (e.g., the change in height over the change in distance equals the slope of a hill, stairs, or a ladder, or distance versus time equals the speed of a car or plane). A slope shows the change of one variable with respect to another, as in a physical slope with respect to distance (the "rise over the run," as seen in Figure 3.2 on the previous page) or, more abstractly, versus time. Mathematically, a change over a change is called a derivative.[4]

As seen in Figure 3.3 in a plot of the Dow Jones index daily high over a monthly 50-year period (top graph) and the change in the index over the same period (bottom graph), slopes or derivatives can reveal seemingly hidden information. For example, spikes appear in the index data but are not nearly as pronounced as those in the derivative data. One can also see that the spikes are more frequent than those in the tortoise and the hare data, indicating a higher volatility. Furthermore, since 2000, the Dow Jones shows an increased spikiness, i.e., a much greater volatility. Note that the data is no different; only the presentation is—the change in the measure instead of the measure.

In auto racing, the same analysis applies, where speed versus time data is sampled using multiple high-frequency, onboard car sensors that record data in real time or at numerous sector points around the course. Figure 3.4 plots a Formula 1 car's speed versus time around the track at Monaco,[5] where the slopes show acceleration (positive slope) and deceleration (negative slope), marking where better pull-away speed and braking translate to better performance. Again, the spikiness is a direct measure of the change in speed as a car speeds up and slows down around the track.

In physics and economics, predicting a future trend (e.g., a ballistic trajectory or future price) is possible by quantifying the change in the "spikiness." Given the distance or price versus time, one calculates the car speed or price change by differentiation (calculating the derivatives

Figure 3.3 Dow Jones index and change in index (monthly from 1959–2009)

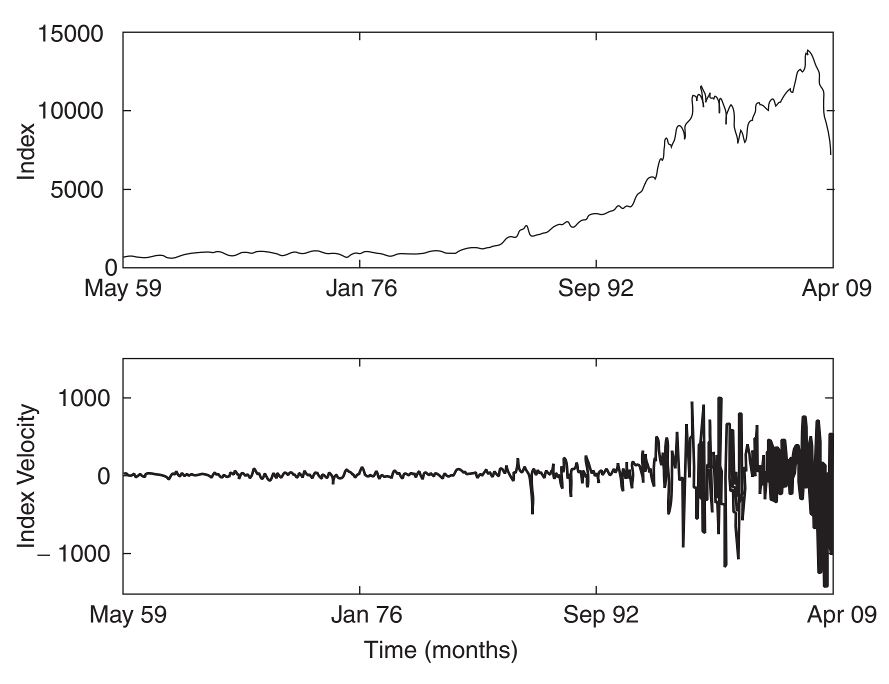

Figure 3.4 Monaco Formula 1 velocity versus sector-point data

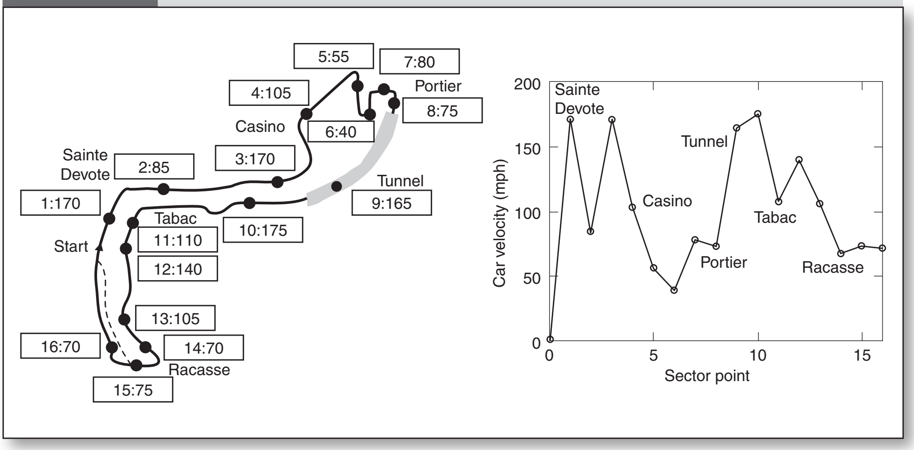

or difference between two points[6]) and the acceleration by taking the derivative again. In the financial markets, this is how a dealer estimates whether a stock or index will continue to peak or plunge and by how much, where a low rate of change indicates a more prolonged trend or high resistance to change (inertia). Growth or contraction in the economy is also measured by such change.

Calculating differences is not hard—we just take the differences between two points. If the tortoise achieved her top speed of 1/30 m/s (i.e., 2 m/min) and the hare his breezy 10 m/s (i.e., 600 m/min) in 10 seconds, then the tortoise's acceleration was 1/30 m/s^2 and the hare's 1 m/s^2.[7] The calculation is no different for the Formula 1 or Dow Jones data, although there are many more sample points—in this case, sector speeds or daily ticker prices.

What is most important with derivatives, however, is that one calculates by working backward, looking closely at how the data changes in time. In physics, the motion is exactly as prescribed by the applied force, as in gravity acting on an astronaut or an apple, which ultimately predicts a future trajectory. In economics, where the data doesn't necessarily follow a prescribed rule or law, the forces aren't as easily determined and future trends are harder to predict, but they can be found in the data.

As for the meaning of the change, one has to look at the cause. In physics, acceleration is associated with a force—recall James Bond's contorted face in *Moonraker* as the astronaut's centrifuge trainer was increased to more than 13 g[8] (without any force, there can be no *change* in the motion). The hare can thus be likened to a drag racer that burned out because of high stresses on the engine, related to a large change of speed in a short period of time. In hare terms, his body gave out. For the tortoise, although she was tired at the end of the race and "comfortably dozing after her fatigue," she didn't overly stress her body (no acceleration, no force) at her slow but steady tortoise pace of 2 m/min. As for racing, drivers accelerate and brake by applying pressure to the pedals while changing gears to improve performance. In the markets, however, the causes are much less clear and often colored by the politics of the day (which we'll look at later).

But change must always be considered—how fast, how slow, and at what rate—to understand the data as presented in any functional relationship, from sports standings to inflation, from everyday electrical usage to global temperatures, and from stock markets to races.

3.2. Distance Versus Speed, or Who to Bet on in the Crunch

The Greek letter for change is delta (Δ), and as we saw above, the change in a measure can be more important than the measure itself, more formally referred to in mathematics as a derivative ($\Delta y/\Delta x$).[9] Here, we look at an example from the world of sports that shows how discrete (or slowly changing) data—such as sports standings, rather than the more challenging continuous (or fast-changing) data of race cars and instant stock prices—is particularly dependent on change and the rate of change.

At the end of the 1992–1993 National Hockey League season, the New York Rangers led the New York Islanders by 3 points with 12 games to play as both teams battled for the coveted last playoff spot. The question is, could one tell from the data which team would go on to win? Or, if one was a bettor, which team would one bet on? The graph in Figure 3.5 shows both teams' performances through the year (top graph) and the

Figure 3.5 Who wins: the Islanders or the Rangers?

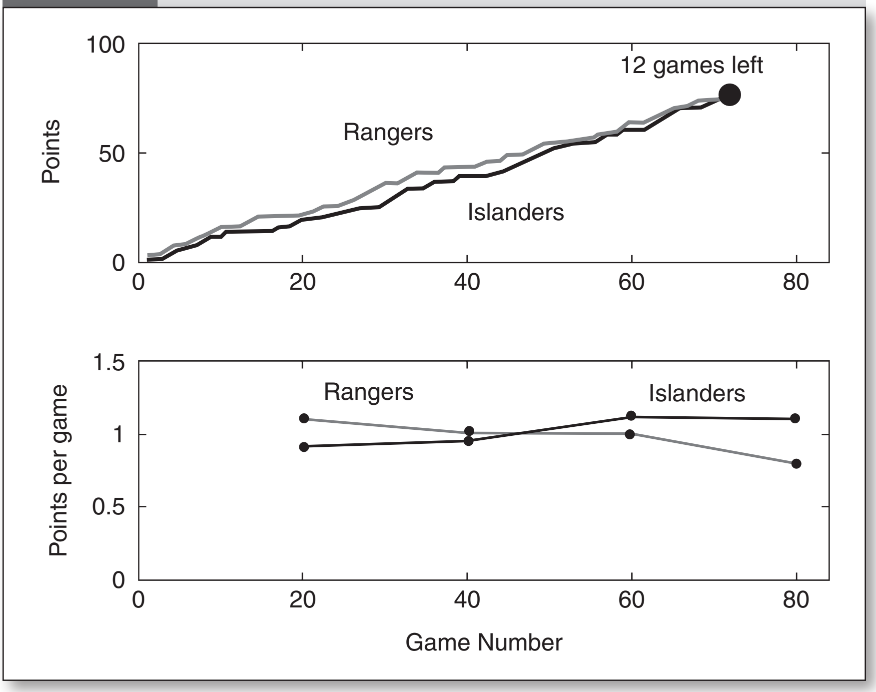

change in their performances over four 20-game periods (bottom graph). With the change in the data now plotted as shown, the question is asked again: Who would you bet on to advance—the Rangers or the Islanders?

The data is an almost static representation compared with the real-time race data or daily Dow Jones data, but the same analysis applies, that of measuring the change in the measure (points per game, velocity, inflation) rather than the measure (points, distance, price). As can be seen in the bottom graph in Figure 3.5, the Islanders steadily improved as the season advanced, whereas the Rangers got worse and, with only four games to play, were already in a tailspin despite having been ahead of their crosstown rivals for most of the season. As it turned out, the Rangers lost 11 of their last 12 games and the Islanders won by 7 points, a trend readily seen in the derivative data (points *per* game)—although one could just as easily say the Rangers lost it, as the derivative data prior to their meltdown shows.

Other teams have also managed to turn newfound "form" into success. The 1951 New York Giants beat the odds in the best-ever baseball pennant race by coming back to win after falling 13 1/2 games adrift of the Brooklyn Dodgers with 7 weeks to play. Over the final weeks, the Giants, with Willie Mays and Bobby Thomson, went 37 and 7 while the Dodgers won only 19 games, miraculously ending the season tied. The pennant winner was decided in the playoff by Bobby Thomson's famous "shot heard round the world." In 2011, the St. Louis Cardinals fashioned a similarly miraculous comeback, trailing by almost 10 games with 5 weeks to play. They won the National League wild card on the last day of the season and went on to win the World Series.

Analyzing sports scores over a limited number of games is a way of quantifying such form and can be a much better indicator of future success than league standings, as understood by the former head of President Obama's National Economic Council, Larry Summers, who as a schoolboy would compare baseball teams' midseason positions to their final positions to determine any telling correlation (Cohan, 2009). In the same way, a snooker player who is behind by 70 points can still win if the table is favorably positioned for a game-winning clearance, or a last-place race car with fresh tires can pass everyone and win. In each case, the change in the data is more revealing than the data itself.

Changing economic data can also show the strength of the economy, as reflected in the gross national product from one quarter to the next and as indicated by growth after a reduction rather than in annual figures (even better if presented monthly or weekly, as we will see in Chapter 5). From such simple examples, one easily sees how static data can be highly suspect and doesn't always represent the true state of affairs, as in the government statistic "economy grows by 3.4% in first quarter" or "unemployment is 2.5% less than last year," which may not represent current conditions, especially during times of economic turmoil.

Furthermore, if inflation measures or consumer price index data lag the real world, such data may be no better than common intuition gleaned from experience, such as department stores or restaurants offering more-than-usual reductions. Admittedly, change is hard to gauge and depends on the sample period (daily, weekly, monthly),[10] but we should at least know that the *change* in a measure holds more information than we think.

3.3. Inflation Decreases, So Why Don't Prices Go Down? It's All in the Delta

Another example from the world of economics highlights how change isn't properly understood and can be confusing even to the professionals. Prior to the 2009 economic downturn, business was booming in Ireland for more than a decade, with 7% average annual growth recorded since 1995. In times of great economic prosperity, housing prices and tax revenue can increase dramatically, as can inflation. In Ireland, particularly Dublin, inflation was on the rise.

One night during the peak, however, the news reported that inflation had decreased over a particular period, and a news reporter asked a representative of a national union to comment on the change. The union representative said the usual good things about the economy and that he welcomed the decrease in inflation, although he hadn't yet seen a corresponding decrease in prices.

At first thought, the union representative's remarks don't seem out of place, but if one pictures an analogous situation—say, that of a driver braking—the fallacy is apparent. Let's say Joe Economist is driving along, sees a cat crossing in front of his car, and slams on his brakes. He decelerates immediately, but of course the car still moves forward before eventually coming to a stop. How hard Joe Economist slams on the brakes (and how fast he was traveling) determines how quickly he will stop, but he will still move forward after braking (think of how long a jet or the shuttle takes to come to a stop after landing). In the same way, although inflation might decrease, *prices* will still rise. In fact, inflation can continue to decrease *forever* and prices will still *rise*, although by less and less.

A decreasing inflation means only that prices are increasing by a *lesser* amount, as shown in Figure 3.6, where the inflation decreases from 4% to 2% to 1% yet prices keep rising. A decrease in prices

Figure 3.6 Prices versus time in decreasing inflation

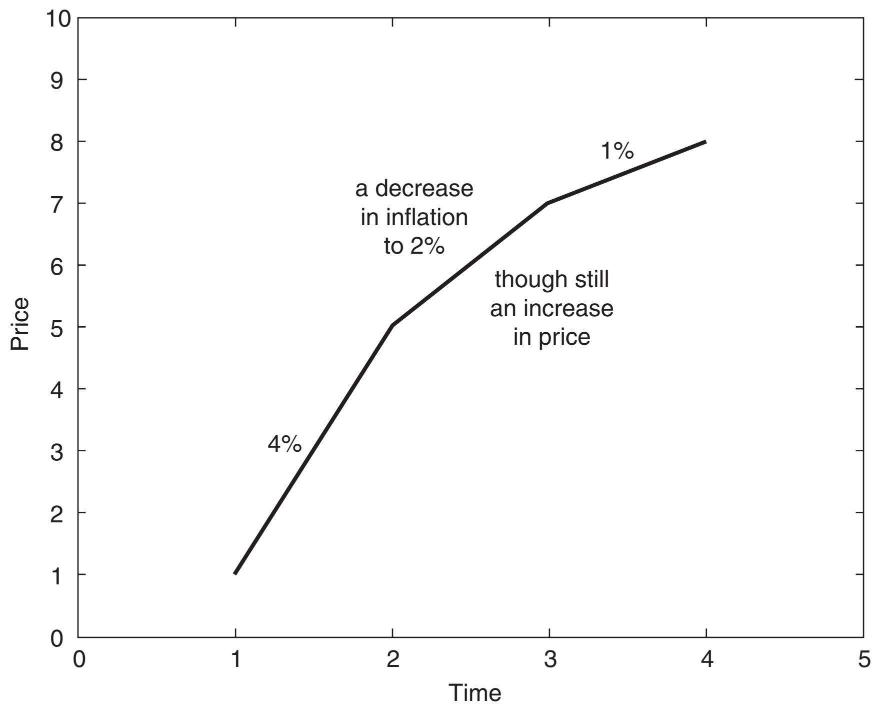

requires *negative* inflation (or deflation), not a decrease in inflation—analogous to a car reversing, something that will never happen when one slams on the brakes. In a car, as in the economy, another gear is needed: reverse.

As we saw above, a force is a change in the change of distance versus time, felt when one decelerates or accelerates—think of the seatbelt restraining your forward motion as you brake or of your back pushing against the seat as you accelerate to pass. In economic terms, the increase in prices is inflation, the *rate* of which can decrease even though prices increase.

Alas, the decreasing inflation in Ireland didn't last long and prices began increasing again in their usual way—until the 2009 economic melt-down, which saw a dramatic decrease in rising prices followed by an actual turnaround, i.e., deflation (or negative inflation). Death and taxes aren't the only certainties in life; inflation isn't far off, economic melt-downs notwithstanding.

3.4. The Measure and the Change in the Measure: More on Slopes

All sorts of examples help us see the difference between a measure and the change in the measure, such as inflation with price or speed with distance. A country's GDP may be increasing (positive slope) but its rate decreasing (negative *change* in the slope), which will eventually translate to a decrease. It seems, however, that the mind is drawn to the instantaneous measure and not to the more important change in the measure, as the following example illustrates (McDermott, Rosenquist, & van Zee, 1987):

> Figure [3.7] shows a position versus time graph for the motions of two objects A and B that are moving along the same meter stick.
>
> At the instant t = 2 s, is the speed of the object A greater than, less than, or equal to the speed of object B? Explain your reasoning.
>
> Do objects A and B ever have the same speed? If so, at what times? Explain your reasoning.

Figure 3.7 Position versus time graph (McDermott et al., 1987)

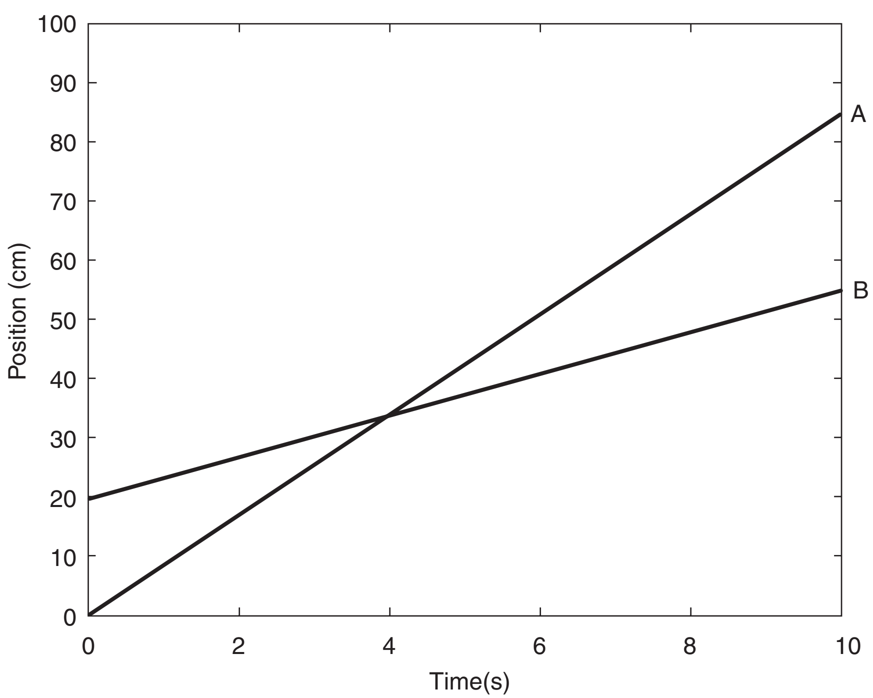

Many of us, however, confuse slope and height, thinking that at 2 s, B is going *faster* than A, when at 2 s, B is only *higher* than A. It is the slope (speed = distance/time) that matters and not the height (distance), yet most do not realize that B (although higher than A prior to 4 seconds) can *never* go faster than A. Recall the hare, who moved much faster than the tortoise but still lost the race (a distance, not a speed race).

As for the second part of the question, as convincing as it might seem that A and B are going the same speed at the point where they cross (4 s), they are not. Since both are straight lines, A is *always* moving at 8.5 cm/s and B is *always* moving at 3.5 cm/s (positively slug-like compared with the 600 m/s and 2 m/s of the hare and the tortoise). A and B may be temporarily at the same height (at 4 s) but will never be at the same speed once in motion.[11]

Similarly, as applied to markets, it does not necessarily follow that buying in a rising market is better than buying in a falling one. Here, the price matters more than the change in price. Thus, it is better to buy as close as possible to the bottom of the market, whether rising *or* falling, and sell as close as possible to the top (mathematically called a local minimum or local maximum). The old adage may have been "location, location, location," but the new adage is "timing, timing, timing." Timing is all about getting the change right.

The following teaser perhaps best highlights the difference between a measure and the change in a measure. If you can order the three ramps in Figure 3.8 from lowest to highest by height and by slope, then you are well on your way to mastering the mathematics of change.

Figure 3.8 Heights and slopes (Kampen, Wemyss, & Smith, 2009)[12]

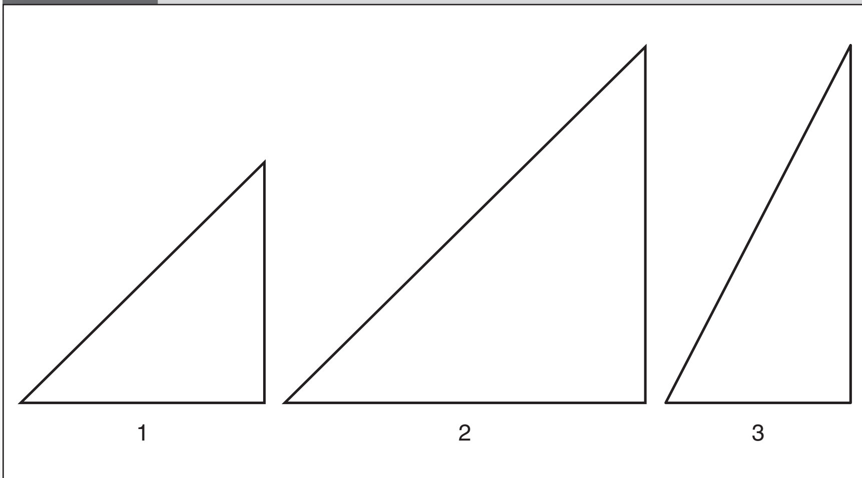

To be sure, numerical data is part of everyday life, from daily temperatures to hourly measures of our own health and well-being, which can be quantified by today's battery of health workers. But a *change* concerns us more often than the actual numbers. No one worries about a constant body temperature of 37° C (= 98.6° F), but when the thermometer starts to creep up or down at a perceptible rate, we become concerned, and all the more as the rate increases.

The same applies to trends in energy consumption or weather, as shown below in two sets of historical, or longitudinal, data (i.e., data versus time): electrical usage per hour over 7 days, as recorded by a utility provider (Figure 3.9), and global temperature per year from the mid-19th century to the present (Figure 3.10). Both sets of data show clear trends, which we can easily see and attempt to interpret now with a little more insight about change.

In Figure 3.9, we see a trend in daily usage—increasing steeply in the morning, continuing through the day, and peaking in the evening. The

Figure 3.9 Electrical usage over 7 days (Ontario Hydro, 2009)

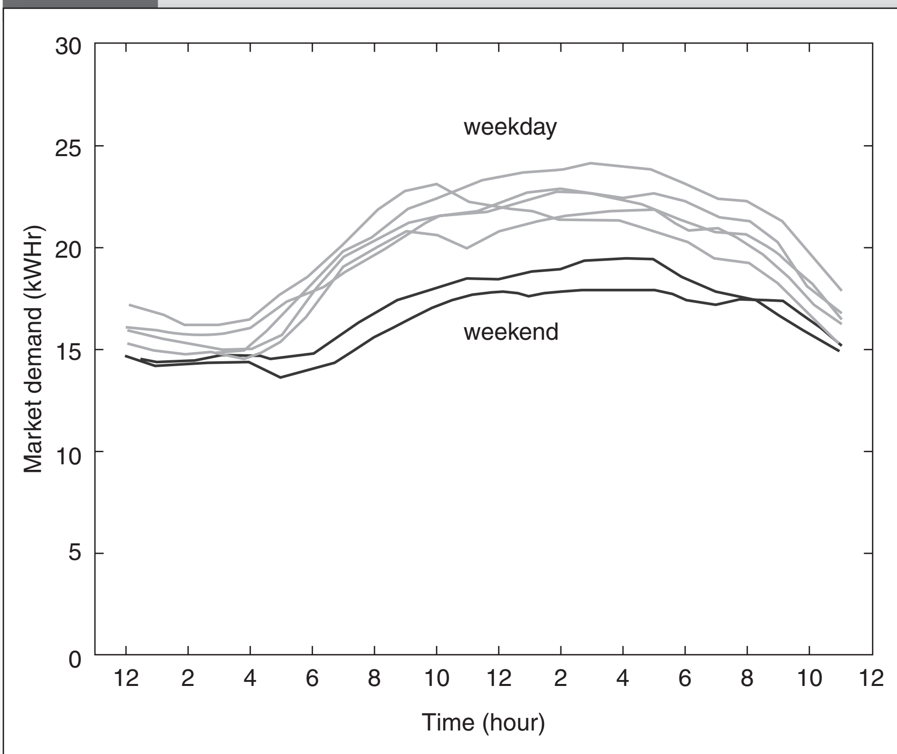

inference is clear: People get up, turn on appliances, work through the day, return home to cook, watch television, bathe children, etc., before turning off their lights to go to bed. There is also an increased load during the week (top five lines), suggesting that more electricity is consumed in the office than at home. Such data can inform power stations about how to balance electrical loads or give politicians ideas about encouraging peak-time reductions, staggered office hours, or public transit fare savings. Here, the plotted data helps us see a simple, obvious trend.

In Figure 3.10, the famous global warming "hockey stick" figure, the data shows a marked increase in temperature with time, plotted as relative "anomalies" to base years. The change in the data from one year to the next, however, is less predictable, as seen in the derivative data, but is certainly not as dramatic as the increase in the stock market data. And although cause and effect are not as easy to interpret—a subject we will look at later with the help of statistics and regression (see Chapter 6)—for

Figure 3.10 Global temperatures (1852–present)

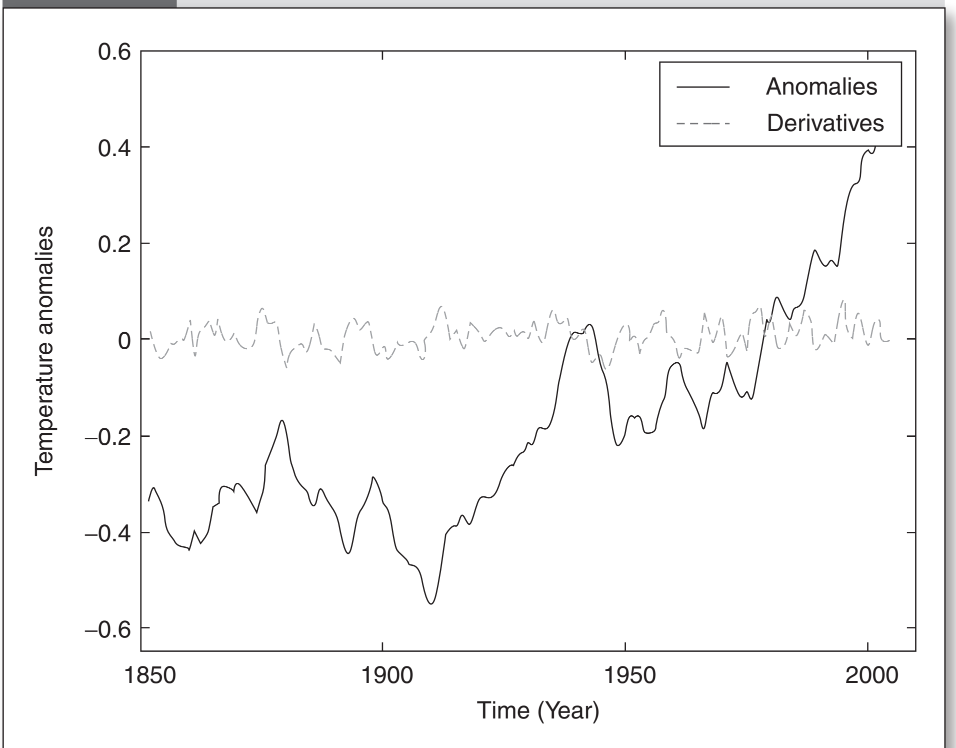

Figure 3.11 GNI versus time for United States and China (1970–2030; World Bank, 2011)

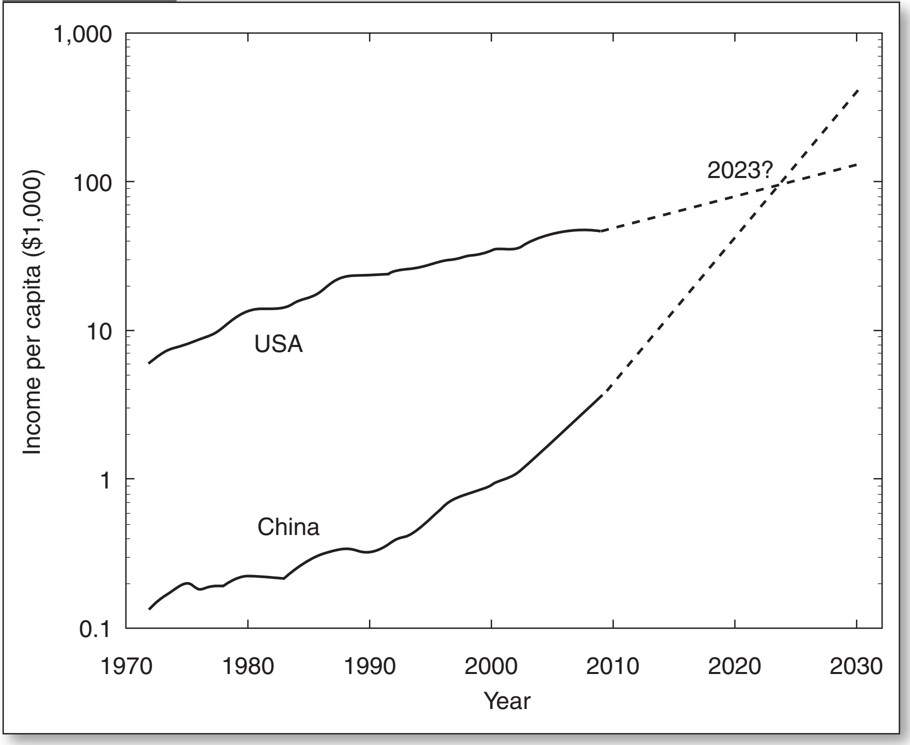

now, we can comment only on the appearance of a correlation that shows a steady increase.

Comparing increases in U.S. income per capita (gross national income, or GNI) to China's fast-growing economy can also show the extent to which China (and the rest of the world) is catching up. By plotting GNI versus time and extending the slope forward, one can even estimate a catch-up date, shown in Figure 3.11 to be in May 2023! But be careful—to achieve such a spectacular increase, China would have to continue growing as in the past 5 years—increasing annual income per person from $3,650 to almost $100,000—an unlikely if not impossible feat.[13] Nonetheless, if we could accurately model the decreasing increase in China's growth, we could come up with a more reasonable estimate, assuming everything stays the same.

So, can we predict the future from the past? That is the question one asks when showing data versus time. In some cases, the answer is yes, and in some cases no, but at least by plotting both the data *and* the change in the data, we can more easily see any changing trends and what might be the real cause if and when abrupt changes occur.

3.5. Turning Points: More Change and Illusive Growth

Formally, a derivative is the change of something with respect to something else, as we saw above in the stock market, Formula 1, and tortoise and the hare data. Speed is the change in distance versus time and acceleration is the change in speed versus time (first and second derivatives of distance), whereas inflation is the change in price (first derivative). We also saw that the spikiness is related to how much the data changes and that a slow change suggests a continuing trend and a fast change high volatility, as was especially seen in the stock market data after 2001, where the excessive spikiness in the index derivative gave a more explicit picture of the increased volatility. Although one can see that the index is changing, the effect is much clearer with the derivative.

Most of us know that the trajectory of a thrown ball follows a parabola (a quadratic function[14]), as seen when an outfielder throws out a runner at home plate or in canon fire on the high seas in an old war movie. The Dow Jones data looks similar, suggesting a similarly growing exponential relationship. But compared with the stock market data, one can see that the change in the Dow Jones has not been constant and has fluctuated dramatically in the past decade. The derivative data shows that the growth of the stock market has not continued unchecked.

Furthermore, the change was seen prior to the presumed global downturn of 2008—much sooner. John Casti (2009), whose research includes large-scale microsimulations of stock markets and road-traffic networks, stated that the financial market turnaround began in 2000, a result that can be readily verified by the derivative data. A snapshot of the Dow Jones data (see Figure 3.12) shows that the index leveled off well before the time generally assumed, which highlights the importance of the sampling period, where change appears less exaggerated over a shorter time.

The variation could be the result of noise (random fluctuations), and at first glance stock market data does include noise. For small time steps, the price variations (noise) follow a random walk, with varying step sizes corresponding to a power-law distribution (a so-called Levy flight, as noted by Benoît Mandelbrot). For large time steps, the noise is more Gaussian (think of the randomness of a series of coin flips).