The Journal of Behavioral Finance
2004, Vol. 5, No. 1, 2–7

THE PRACTITIONER'S PERSPECTIVE

Bleed or Blowup? Why Do We Prefer Asymmetric Payoffs ?

Nassim Nicholas Taleb

In some strategies and life situations, it is said, one gambles dollars to win a succession of pennies. In others one risks a succession of pennies to win dollars. While one would think that the second category would be more appealing to investors and economic agents, we have an overwhelming evidence of the popularity of the first. A popular illustration of such asymmetry in returns is evident in the story of the Long Term Capital Management hedge fund. The fund derived steady returns over a dozen quarters then lost all of them in addition to almost all its capital in a single observation (see Lowenstein, 2000)–only for the main principals to restart a new, albeit milder, version of the strategy. Is there a systematic bias in favor of such return profiles?[1]

Negative (or "left") skewness can be presented by considering a stream of gambles that differ from most symmetric lotteries generally presented in the literature (where the agent usually has a 50% probability of realizing a given gain, G, and a 50% probability of realizing a loss, L). The asymmetric case we consider has, for a given expectation, both probabilities markedly diverging from 50%. A considerably negatively skewed bet can present more than 99% probability of making G and less 1% probability of losing L. While such skewness may sound extreme we will see that there is an abundance of instruments in the financial markets that actually deliver such payoffs (one may even say that *almost all* derivative products offer asymmetric properties). More technically, the mathematical representation of negative skewness defines it as a negative third central moment, the product of the probabilities by the cube of the payoffs deducted from the mean.

Would an economic agent facing a stream of stochastic monetary payoffs prefer negative skewness? Given a profile of monetary gambles, would he prefer to "bleed" (i.e. undergo small but frequent losses) or "blowup,"[2] i.e., take severe hits concentrated in small periods of time? Statistical properties of popular classes of investments, earnings management on the part of corporations (where corporations manage their earnings to moderately beat estimates most of the time and take hits on occasion[3]), and mechanisms like covered call writing (where investors clip their upside gains for a small fee) shows a strong evidence for the predilection for negative skewness on the part of investors. Indeed such preference is mostly expressed in the growth of classes of financial securities like hedge funds that, according to the empirical literature (see Fung and Hsieh, 1997, Kat, 2002), seem to be severely plagued by such properties, even possibly designed to cater to the investors' biases.

We divide biases into two categories, namely, a) cognitive, as agents may not understand the true implications of skewness, or why the expectation of a payoff is not necessarily better even if it generates steady returns; b) behavioral, as they may prefer a set of payoffs to another. The aim of this short discussion is to connect this preference for skewness with research that has been done in the behavioral literature—and describe further experiments that would be needed to confirm it. It is organized as follows. Skewness is discussed, along with its prevalence in the growing new investment classes. We examine three major angles in behavioral research: a) the belief in the "law of small numbers" and aspects of the overconfidence literature, b) prospect theory, and c) the promising field of hedonic psychology, an offshoot of prospect theory. The aim of this note is not the display of the evidence but hints and direction for confirmatory research.

Skewed Payoffs and Financial Instruments

How are these payoffs constructed? Instruments abound. Consider the following examples, which we divide into direct (i.e. strategy analyzed on their own merit) and comparative (strategies analyzed in comparison with a benchmark).

Examples of Directly Negatively Skewed Bets

Loans and credit-related instruments. You lend to an entity at a rate higher than the risk-free one prevail-

Nassim Nicholas Taleb is founder and chairman of Empirica LLC, and a fellow and adjunct professor at the Courant Institute of New York University.

Requests for reprints should be sent to: Nassim Nicholas Taleb, Empirica LLC, 110 East 59th Street, New York, NY 10022. Email: nnt@fooledbyrandomness.org

ing in the economy. You have a high probability to earn the entire interest amount, except, of course in the event of default where you may lose (depending on the recovery rate) approximately half of your investment. The lower the risk of default, the more asymmetric the payoff. The same applies to investments in high yielding currencies that are pegged to a more stable rate (say the Argentine peso to the dollar) but occasionally experience a sharp devaluation.

Derivative instruments. A trader sells a contingent claim. If the option is out-of-the-money the payoff stream for such strategy is frequent profits, infrequent large losses, in proportion to how far out-of-the-money the option is. It is easy to see in the volumes that most traded options are out-of-the-money.[4] Note that a "market neutral" or "hedged" (i.e., made insensitive to the direction of the market)[5] such strategy does not significantly mitigate such asymmetry, since the mitigation of such risk of large losses implies continuous adjustment of the position, a strategy that fails with discontinuous jumps in the price of the underlying security. A seller of an out-of-the money option can make a profit as frequently as he wishes, possibly 99% of the time by, say selling on a monthly basis options estimated by the market to expire worthless 99% of the time.

Arbitrage. There are classes of arbitrage operations such as: 1) "merger arbitrage" in which the operator engages in betting that the merger will take place at a given probability and loses if the merger is cancelled (the opposite is called a "Chinese"). These trades generally have the long odds against the merger. 2) "Convergence trading" where a high yielding security is owned and an equivalent one is shorted thinking that they converge to each other, which tends to happen except in rare circumstances. The hedge fund boom has resulted in a proliferation of packaged instruments of some opacity that engage in a variety of the above strategies—ones that are revealed through naive statistical observation.

Example of Comparatively Skewed Bets

Covered calls writing. Investors have long engaged in the "covered write" strategies by selling an option against their portfolio, thereby increasing the probability of a profit in return for a reduction of the upside potential. There is an abundant empirical literature on covered writes (see Board, Sutcliffe and Patrinos, 2000, for a review, and Whaley, 2002 for a recent utility-based explanation) where investors find gains in utility from capping payoffs as the marginal utility of gains decreases at a higher asset price. Indeed the fact that individual investors sell options at cheaper than their actuarial value can only be explained by the utility effect. For a mutual fund manager, doing such "covered writing" against his portfolio increases the probability of beating the index in the short run, but subjects him to long term underperformance as he will give back such outperformance during large rallies.

Properties of a Left-skewed Payoff

In brief, a negatively skewed stream offer the following attributes:

Property 1: Camouflage of the Mean and Variance

The true mean of the payoff is different from the median, in proportion to the skewness of the bet. A typical return will, say, be higher than the expected return. It is consequently easier for the observer of the process to be fooled by the true mean particularly if he observes the returns without a clear idea about the nature of the underlying (probability) generator. But things are worse for the variance as most of the time it will be lower than the true one (intuitively if a shock happens 1% of the time then the observed variance over a time window will decrease between realizations then sharply jump after the shock).

Property 2: Sufficiency of Sample Size

It takes a considerably longer sample to observe the properties under a skewed process than otherwise. For example, consider a bet with 99% probability of making G and 1% probability of losing L. In this example, the properties will not reveal themselves 99% of the time–and when they do, it is always a little late as the decision has already been made. Contrast that with a symmetric bet where the properties converge rather rapidly at the square root of the number of observations.

Property 3: The Smooth Ride Effect.

As mentioned above, the observed variance of the process is generally lower than the true variance most of the time. This means, simply, that the more skewness, the more the process will generate steady returns with smooth ride attributes, concentrating the variance in a brief period, the brevity of which is proportional to the variance. In other words, an investor has, without a decrease in risk, a more comfortable ride most of the time, with an occasional crash.

Overconfidence and Belief in the Law of Small Numbers

The first hint of an explanation for the neglect of the small risks of large losses comes from the early literature on behavior under uncertainty. Tversky and Kahneman (1971) wrote "We submit that people view a sample randomly drawn from a population as highly representative, that is, similar to a population in all essential characteristics." The consequence is the inductive fallacy: overconfidence in the ability to infer general properties from observed facts, "undue confidence in early trends" and the stability of observed patterns and deriving conclusions with more confidence attached to them than can be warranted by the data. Worst, the agent finds causal explanations or perhaps distributional attributes that confirm his undue generalization.[6]

It is easy to see that the "small numbers" get exacerbated with skewness since the observed mean will usually be different from the true mean and the observed variance will usually be lower than the true one. Now consider that it is a fact that in life, unlike a laboratory or a casino, we do not observe the probability distribution from which random variables are drawn. We only see the realizations of these random processes. It would be nice if we could, but it remains that we do not measure probabilities as we would measure the temperature or the height of a person. This means that when we compute probabilities from past data we are making assumptions about the skewness of the generator of the random series–all data is conditional upon a generator. In short, with skewed packages, Property 1 comes into play *and* we tend to believe what we see.

The literature on "small numbers" implies that agents have a compressed, narrower distribution in their minds than warranted from the data. The literature on overconfidence studies the bias from another angle by examining the wedge between the perception of unlikely events and their actual occurrence as well as the failure to calibrate from past errors. Since Alpert and Raiffa (1982), studies have documented how agents underestimate the extreme values of a distribution in a surprising manner; violations are far more excessive than one would expect: events that are estimated to occur less than 2% of the time will take place up to 49%. There has since been an abundance of literature on overconfidence (in the sense of agents discounting the probability of adverse events while engaging in a variety of projects), see Kahneman and Lovallo (1993), Hilton (2003).

"Every Day is a New Day": The Implications of Prospect Theory

Prospect theory derives its name from the way agents face prospects or lotteries (Kahneman and Tversky, 1979). Its central idea is that economic agents reset their "utility" function to ignore, to some extent, accumulated performance and focus on the changes in wealth in their decision making under uncertainty. One may accumulate large quantities of wealth, but habituation makes him reset to the old Wall Street adage "every day is a new trading day," which means that he will look at gains and losses from the particular strategy, not the absolute levels of wealth and make decisions accordingly. The reference point is the individual's point of comparison, the "status quo" against which alternative scenarios are compared. Moreover prospect theory differs from "utility theory" *per se* in the separation of decision probability from the "value function." Decision probability, or weighting function, has the property of exaggerating small probabilities and underestimating large ones.

It is noteworthy that prospect theory was empirically derived from one-shot experiments with agents subjected to questions in which the odds were supplied.[7] It is the value function of the prospect theory that we examine next, rather than probabilities used in the decision-making. The normative neoclassical utility theory stipulates an increased sensitivity to losses and a decreased one to gains (investors would prefer negative skewness only for their increase but not decrease, in wealth). On the other hand, the value function of prospect theory documents a decreased sensitivity to both gains and losses, hence a marked overall preference for negative skewness. At the core, the difference is simply related to the fact that operators are more concerned with the utility of changes in wealth rather than those of the accumulated wealth itself, creating a preference for a given path dependence in the sequences of payoffs.

To see how the empirically derived version of utility theory presents asymmetric higher order properties, consider the following proposed representation of Tversky and Kahneman (1992). One has a value function, $V^+(x)$, for x positive or 0, and $V^-(x)$ for x strictly negative.

$$V^+(x)= x^\alpha$$

$$V^-(x)= (-\lambda)(-x)^\alpha$$

From this it is easy to see that, with $\alpha<1$, that V is concave in the profit domain and convex to in the loss domain. Looking at the second derivatives, one observes:

$V^{+''}(x)= \alpha (\alpha-1) x^{\alpha-2}$; it is negative: a large profit has an incrementally smaller impact on the "utility" of the individual.

$V^{-''}(x)= -\lambda \alpha (\alpha-1) (-x)^{\alpha-2}$; it is positive, larger losses have a numbing effect.

Hence, the value of a large loss is higher than the sum of the value of losses. In other words the agent's utility resides in incurring a sharp hit than the same amount in piecemeal tranches. A loss of 100 (blowup) is better (from the value function standpoint) than 100 times a loss of 1 (bleed).

FIGURE 1
The Kahneman-Tversky Value Function

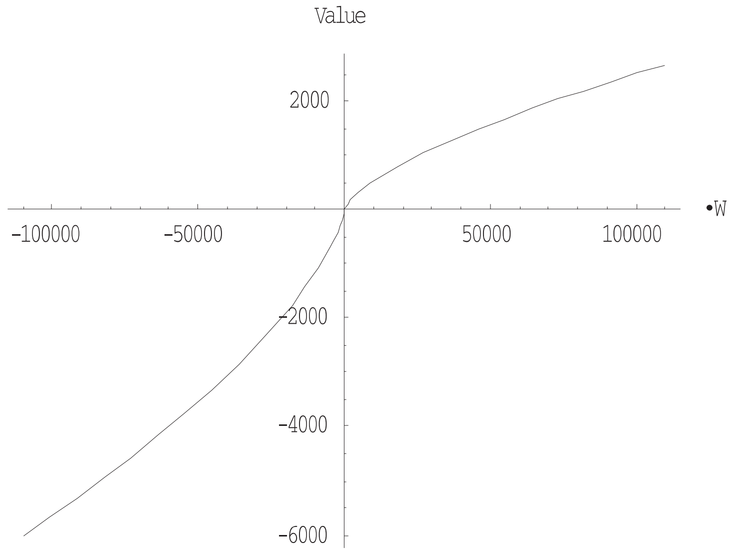

Note: The estimation in Tversky and Kahneman (1992) for λ =2.25, α = .65: losses are 2.25 times worse than profits. Note that a loss of \$100,000 has for value –56,517, a loss of \$1000 has for value –982, 1/57th of the pain. The relationship reverses for the gain domain.

Note, by comparison, the attributes of the conventional Von Neuman-Morgenstern utility of wealth (instead of payoffs). This results in asymmetry in preferences: U(W) is concave for all levels of wealth W which makes the investor prefer the frequent small losses to the large occasional one, as well as the frequent small profits to the large infrequent ones. In the domain of gains or increase in wealth there is a convergence between the two methods of viewing utility.

So far the value function seems to confirm the preference for skewness hypothesis. Prospect theory has been subjected to all manner of experiments and the concavity in the domain of losses has shown to be robust. However we are confronted here with a modicum of ambiguity–the overestimation of the odds in the probability weighting function seems incompatible with the statement in the previous section concerning the underestimation of the outliers.

We can safely ignore the probability weighting function, as we are looking at the results of risk taking in a framework of purely inductive inference, where the probabilities and the risks are not supplied, only discovered by agents and therefore subjected to cognitive biases. Recent research (Barron and Erev, 2003) shows experimental evidence that agents underweight small probabilities when they engage in sequential experiments in which *they derive the probabilities themselves*. Whether this comes from biases in our inductive inference machinery or the fact that we do not handle abstract probabilities properly (the "risk as feeling" theories) is to be ascertained. Note that the intuition of the problem is presented in an early paper by Slovic, Fischhoff, Lichtenstein, Corrigan, and Combs (1977), with the explanatory title of "preference for insuring against probable small losses"; they attributed their results consistent with the neglect of large infrequent losses to the "the sequential nature of the problem."[8]

Hedonic Adaptation and Quality of Life Perspective

The central idea behind recent research on well-being in hedonic psychology is the existence of a set-point of happiness, to which the agent tends to revert after some circumstantial departure. A positive or negative change in material conditions brings some changes in the individual well-being, but soon the process of habituation causes the reversion to the old level of life satisfaction. It is the equivalent of the utility curve resetting at the origin in the prospect theory case and the new changes in conditions mattering more–as if the accumulated changes did not bring a permanent change in one's utility.

The idea has been called the "hedonic treadmill" after the seminal paper, Brickman and Campbell (1971). Studies document that paraplegics after suffering from the onset of the impairment converge soon after to the general level of happiness of the population. Lottery winners also do not seem to hold on to significant permanent gains in their happiness and well-being. Academics granted tenure are no happier a few months later than they were before. The same applies to general societies experiencing abundance. Such mechanisms of adaptation are the backbone of the research on happiness and economics–an emerging branch in research, see Layard (2003), Frey and Stutzer (2002). Indeed utility was associated by Bentham to a measure of individual and communal happiness; it seems that economics has made a return to it.[9]

Seen in the context of skewness, the notion of habituation implies the following: *concavity* of good events; *convexity* of bad ones. Indeed the value function of "hedonic enjoyment" is deemed to have the same higher order properties as prospect theory (Kahneman, 1999). Again consider if the value function in the positive domain, $V^+[x]$ is concave, then the implication is that it is better to receive a steady flow of "good" events than the same quantity in one block–whether they are monetary gains or flows of enjoyment.

Consider two economic agents operating with mirror portfolios and strategies. The first, whom we call Nero, loses $1000 for 99 consecutive weeks (he "bleeds"), then makes $99,000 on the 100th. The second, whom we call Carlos, has the exact opposite payoff (he "blows up" on the 100th week). According to such aspect of the hedonic literature, Carlos' well-being and quality of life should be superior to those of Nero. The arithmetic sum of the pleasure/pain should swing squarely in his favor: consider that Carlos will experience 99 pleasurable weeks, will go to work every Monday with the expectation of more good news, and that the pain experienced from the loss will be short lived since he will recover from it soon after. As to Nero, the exhilaration of the gain will not compensate the lengthy bleed period. As such this theory provides the explanation that, everything else being held equal, for a given mathematical expectation of a payoff, a negatively skewed one provides higher quality of life. A few questions remain, however, to answer before the above argument can be accepted.

1. It seems that to analyze the summation of utilities through time might not be straightforward as a measure of total utility–and it would be normative to assume that the agent *should* be subjected to one instead of another. Indeed research (see Kahneman, 1999) from such experiments as those of subjects undergoing colonoscopy, show that those do not base their decisions on past linear summation of utility, but to more complicated rules that tend to favor the peak and ending part of the sequence ("peak-end rule"). This gave rise to four possible utilities:

   a. experienced utility–the summation of the value function over the periods considered.
   b. remembered utility–the agent's recollection of the total experienced utility, often at odds with the previous one.
   c. predicted utility–the utility that the agent believes will result from the action.
   d. decision utility–the utility used in the decision process.

We assumed that the experienced utility (here the sum of the value function over time) was the one that mattered. Further experiments are needed to confirm such a point. Would it be the case that Nero, in spite of his negative experienced utility, would have a higher remembered utility from the episode?

2. It seems that such treadmill effects are selective and domain specific: there are things that lead to permanent happiness, or to an injection of utility that carries permanent effects. In all of these situations we do not revert to the origin or the set point. In some cases, repetition or duration of a constant stimulus even results in an increasing hedonic response–a process the literature calls sensitization, the exact opposite of the treadmill effect. The well-being literature (Frederick and Loewenstein (1998)) shows evidence that there are:

   a. Some things to which we adapt rapidly: (imprisonment, increases in wealth, and disabilities like paralysis),
   b. Condition to which we adapt slowly (the death of a loved one), and
   c. Things to which we do not seem to adapt (noise, debilitating diseases, foods, or an annoying roommate).

Now the question: do people adapt to "bleeding"? In other words do people increase in sensitivity to the pain of the "Chinese torture" treatment of slow losses? On this experiments should be done.

Conclusion

This discussion has explored skewness from the utility standpoint in addition to the perception of the probability of large adverse shocks. Prospect theory provides hints on this, but further research is needed to examine how agents react in a multi-period framework. This discussion also found evidence in the literature for the undervaluation of the probability of large adverse shocks when risks are neither salient nor directly observable. This may explain the appetite for negatively skewed payoffs. Finally, more research is needed for determining the relationship between utility of *streams* of payoffs and decision making.

Notes

1. We ignore in this discussion economic arguments justifying skewness, the most significant of which is the "moral hazard" argument. This argument stipulates that agents risking other people's capital would have the incentive to camouflage the properties by showing a steady income. Intuitively, hedge funds are paid on an annual basis while disasters happen every

four or five years, for example. The fund manager does not repay his incentive fee.

2. See Gladwell (2002) for a popular presentation of the difference between the two classes of strategies.
3. DeGeorge and Zeckhauser (1999) show the skewness in the distribution of the difference between announced and expected corporate earnings. For an illustration of the custom by a master of the practice see the memoirs of Jack Welsh (Welsh, 2001) who explains explicitly (and candidly) how he managed to use accounting methods to the smooth the earnings of the conglomerate GE.
4. See Wilmott(1998) and Taleb (1997) for a discussion of dynamic hedging properties for an option seller.
5. This is called "delta hedging" where the operator buys and sells the underlying security to respond to the changes in sensitivity of the option to the underlying security. As the underlying price rises the option trader may be insufficiently covered and would need to buy more of the asset. Likewise the operator needs to sell in response to the fall in the asset price. It is key here that volatility causes losses for such an agent –particularly discontinuities and jumps in the underlying security.
6. See Rabin(2000) for a modern treatment of the "small number" problem.
7. There have been few studies of sequential behavior –see Thaler and Johnson (1990) study of the sequential behavior of agents to see how they are affected by gains and losses.
8. There is a relevant recent piece of research in the recent literature on the affect heuristic (the tendency to determine the probability of an event by the emotional response that it causes). Hsee and Rottenstreich (2004) show that agents, when subjected to valuation "by feeling" (as opposed to valuation by calculating – a process that is not subjected to the affect heuristic) tend to be sensitive to the presence or absence of a given stimulus rather than its magnitude. This implies that a loss is a loss first, with further implications later. The same with profits. This explains the concavity/convexity of the value function. The agent would prefer the number of losses to be low and the number of gains to be high, rather than optimizing the total performance
9. Bentham's definition (Behnthm, 1789): "By utility is meant that property in any object, whereby it tends to produce benefit, advantage, pleasure, good, or happiness, (all this in the present case comes to the same thing) or (what comes again to the same thing) to prevent the happening of mischief, pain, evil, or unhappiness to the party whose interest is considered: if that party be the community in general, then the happiness of the community: if a particular individual, then the happiness of that individual."

References

Alpert, M. and H. Raiffa, H. "A Progress Report on the Training of Probability Assessors" in D. Kahneman, P. Slovic & A. Tversky, eds., *Judgment Under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press, 1982, pp.294–305.

Barron, G. and I Erev. "Small Feedback-based Decisions and Their Limited Correspondence to Description-based Decisions." *Journal of Behavioral Decision Making,* 16, (2003), pp. 215–233.

Bentham, J. *An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press, 1789,1996.

Board, J., C.Sutcliffe and E. Patrinos. "Performance of Covered Calls." *European Journal of Finance*, 6, (2002), pp. 1–17.

Brickman, P. and D. T. Campbell. "Hedonic Relativism And Planning The Good Society" in M.H. Apley, ed., *Adaptation-Level Theory: A Symposium*. New York: Academic Press, 1971, pp. 287–302.

DeGeorge, J.P. and R. Zeckhauser. "Earnings Management to Exceed Thresholds." *Journal of Business*, 72, (1999).

Frederick, S., and G. Loewenstein. "Hedonic Adaptation" in *Well Being: Foundations of Hedonic Psychology*, ed., D. Kahneman, E. Diener and N. Schwartz. New York: Russell Sage Foundation, 1998.

Frey, B.S. and A. Stutzer. *Happiness and Economics: How the Economy and Institutions Affect Human Well-bein*g. New Jersey: Princeton University Press, 2002.

Fung, W. and D. Hsieh. "Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds." *The Review of Financial Studies,* 10, (1997), pp. 275–302.

Gladwell, M. "Blowing Up." *New Yorker,* April 22, 2002.

Hilton, D. "Psychology and the Financial Markets: Applications To Understanding and Remedying Irrational Decision-Making" in I. Brocas and J.D. Carillo eds., *The Psychology of Economic Decisions: Vol 1: Rationality and Well-Being*. Oxford: Oxford University Press, 2003.

Hsee, C. K. and Y. R. Rottenstreich. "Music, Pandas and Muggers: On the Affective Psychology of Value." Forthcoming, *Journal of Experimental Psychology.*

Kahneman, D. and D. Lovallo. "Timid Choices and Bold Forecasts: A Cognitive Perspective On Risk-Taking," *Management Science*, 39, (1993), pp. 17–31.

Kahneman, D. "Objective Happiness." In *Well Being: Foundations of Hedonic Psychology*, ed., D. Kahneman, E. Diener and N. Schwartz. New York: Russell Sage Foundation, 1999.

Kahneman, D. and A. Tversky. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica*, (1979).

Kat, H. "CTAs and Hedge Funds: A Marriage Made in Heaven." Cass Business School, London City University, Manuscript, 2002.

Layard, R. *Happiness: Has Social Science a Clue?* Robins Memorial Lecture, London School of Economics, 2003.

Lowenstein, R. W*hen Genius Failed: The Rise and Fall of Long-Term Capital Management*. New York: Random House, 2000.

Rabin, M. "Inference by Believers in the Law of Small Numbers." Economics Department, University of California, Berkeley, manuscript, 2000.

Slovic, P., B. Fischhoff, S. Lichtenstein, B. Corrigan and B. Combs. "Preference for Insuring Against Probable Small Losses: Implications For The Theory And Practice Of Insurance." *Journal of Risk and Insurance*, 44, (1977), pp. 237–258. Reprinted in P. Slovic, ed., *The Perception of Risk.* London: Earthscan.

Taleb, N. *Dynamic Hedging: Managing Vanilla and Exotic Options*. New York: Wiley, 1997.

Thaler, R., and E. Johnson. "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes in Risky Choice." *Management Science*, June (1990).

Tversky, A. and D. Kahneman. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*, 5, (1990), pp. 297!323.

Tversky, A. and D. Kahneman. "Belief in the Law of Small Numbers." *Psychology Bulletin,* 76(2), (1971), pp. 105–10.

Welsh, J. *Jack*: *Straight From the Gut*. New York: Warner Books, 2001.

Wilmott, P. *Derivatives*. Chichester: Wiley, 1998.