A New Psychological Perspective on Investor Decision Making

Harvey A. Pines, Ph.D.

Psychology is enjoying a bull market. It is difficult these days to find an analysis of Wall Street happenings that does not at some point refer to "market psychology," "investor sentiment," "contrary opinion" or some other term related to psychology. In this article I want to bring to your attention a new psychological perspective, one that until now has received relatively little notice outside of the academic milieu. Drawing on recent research findings defining this perspective, I will identify several important, but non-obvious, ways in which the individual investor's judgment and choices may be biased or erroneous. Then, several suggestions will be made for tactics to counter these sources of bias and error.

The psychological perspective I shall present arises from the relatively new field of *cognitive psychology*. In contrast to earlier work, cognitive psychology does not speak to the social influences on investor behavior emphasized by David Dreman in his important work on contrarian investing, or to the effects of motivation and personality eccentricities discussed by other writers. The field of cognitive psychology is based on an analogy between *information processing* in computers and the same processes in human beings. How we acquire and remember information, and how we process that information to arrive at judgments and decisions, are the focal point for cognitive psychologists. This concern with information processing activity is especially appropriate in an investing environment in which the dominant academic hypothesis about the dynamics of the stock market—the efficient market hypothesis—is couched in terms of the dissemination of and responsiveness to information.

Among the most fundamental information processing activities for investors are assessing probabilities, judging values and combining this information into an overall evaluation leading to a choice from among different courses of action. These behaviors ultimately rely on subjective judgments, even when performed in the context of formal models of security evaluation and portfolio management.

What does cognitive psychology tell us about our ability to carry out these judgments, to process the vast amount of information that daily impacts investors? It is becoming abundantly clear that there are several facts about us as human decision makers.

- We have a limited capacity to process large amounts of complex information.
- We try to cope with that limited capacity by using a variety of simplifying rules-of-thumb or *heuristics* which, while effective in economizing on the information processing activity, are fallible guides to action and, under certain circumstances, become a liability rather than an asset.
- In assessing probabilities of events, we tend to overweigh information that is either minimally relevant or completely irrelevant, and to underweigh information which should form the basis of our assessments.
- We regularly underestimate important statistical influences on events.
- Our judgments and choices exhibit important inconsistencies.
- Our assessment heuristics lead us to beliefs about the predictability and control of events which may not be true in the uncertain world of investing.
- We are measurably *overconfident* in our judg-

ments, in spite of all of the above.

From what evidence does the cognitive psychologist draw such pessimistic conclusions? Being committed to empirical investigation, these psychologists often construct decision environments in the laboratory to test their hypotheses. Efforts are then made to validate these observations by studying decision making behaviors in more natural settings, with experts in the field as subjects of study. I shall present a few examples of the judgmental and decision tasks that have been used by cognitive researchers. While these tasks often do not involve investment judgments directly, their relevance to such judgments will be made clear.

Heuristics: Simplifying a Complex World

Consider the following two questions, both of which involve an assessment of frequency or likelihood:

Question 1. John and Mary are part of a family with six children, three boys and three girls. The following are three of several possible orders in which they might have had their children: a) GBBGBG b) BBBGGG c) BGGBGB. Is one of these sequences more likely to have occurred than the other two? Why?

Question 2. Does the letter k appear more often as the first letter of a word or as the third letter?

The well-known cognitive researchers, Professors Daniel Kahneman and Amos Tversky, have proposed that in an effort to reduce the complexity and demands of exhaustively calculating the likelihood of an event, we seek to use heuristics to arrive at our probability estimates. Two important and commonly used heuristics are the *representativeness heuristic*, according to which estimated probability is a function of the similarity of two events, and the *availability heuristic*, according to which estimated probability is a function of the availability of information in our memory.

To illustrate the *representativeness heuristic,* consider the first question. Most people believe that either GBBGBG or BGGBGB is more likely to occur than BBBGGG. This judgment is made, according to Kahneman and Tversky, because we employ the following heuristic reasoning: since a child's sex is randomly determined, we should observe a sequence that is representative of (typical of or similar to) the occurrence of a random process. The irregularly mixed sequences appear to be more representative of a random process than the regular sequence, so we judge the regular sequence as least likely. In fact, each of the above sequences is equally likely to occur. Thus, our seemingly plausible rule-of-thumb has, in this case, led us to an erroneous conclusion.

How might the representativeness heuristic lead our investing judgment astray? One closely related example would be predictions by technical analysts based on patterns of short-run price changes. Here the heuristic might be: Because there are apparent systematic or orderly patterns to short-run price changes, judge the underlying market process also to be orderly. In fact, the apparently regular sequence of price changes may be randomly generated. Use of the representativeness heuristic makes it very difficult to "see" or acknowledge this randomness, even if present. A second example of representativeness-induced bias might arise in evaluating a newly formed company's business prospects. In today's economic environment, we might intuitively, but erroneously, predict a high technology firm as more likely to succeed than, say, a company that makes baby products. This judgment would be arrived at, not by a careful analysis of the firm's business environment, but because the technology company is more representative of companies that are visibly successful in today's markets. Other applications of the representativeness heuristic will be identified below.

To see how the *availability heuristic* affects our likelihood judgments, examine question two. Although most people judge the letter k to occur more frequently as a first letter (know), it is actually three times as likely to appear as a third letter (acknowledge). This error in frequency estimation can be accorded to our tendency to search our memory for relevant instances of what we are trying to estimate and then judge their probability by the ease with which we find such instances. Normally, more frequent events are more available in memory. In applying the availability heuristic, we appear to reverse this association and assume that if an event is more easily recalled it must have occurred more frequently. However, availability of information in our memory is influenced by many factors unrelated to the frequency of an event, e.g., its recency, its interest, its importance or emotional significance to us, or simply the attention given to it. Words beginning in k may be more readily noticed and thus remembered, but they are not more frequent.

The impact of the availability heuristic on investors' judgments is widely visible. A somewhat unusual example comes from a study of the purchase of disaster insurance. It has been documented that right after an earthquake there is a sharp increase in the purchase of earthquake insurance, but that the purchase rate soon drops back to normal as the

earthquake ceases to be a recent event. Although one earthquake does not significantly alter the objective probabilities of earthquake occurrence, its increased availability in memory temporarily changes our subjective probability assessment. The availability heuristic also explains why our investment judgments are unduly affected by the most recent change in stock market prices; yesterday's price activity tends to stand out in our memory, and, as a result, it disproportionately influences our estimate of future price changes. More generally, our estimates of the likelihood that prices will appreciate or decline may be significantly influenced by the ease with which we can retrieve bullish or bearish information from our memory. Therefore, memorable events like a dramatic daily rise or fall in the market averages, pronouncements by well-known, but not necessarily expert public figures, or vivid coverage by the media, are all likely to unjustifiably alter our individual, subjective estimates of future price changes.

Even the wording of a question may affect our probability estimates because of the availability heuristic. Consider the following example:

a) How likely is it that there will be a resurgence of inflation?
b) How likely is it that, with the present expansion of the money supply, there will be a resurgence of inflation?

The second question may induce a narrow focus, making information about past instances of money supply growth especially available in memory. The first question may lead instead to the recall of available facts about the more general condition of inflation. The resulting probability estimates may vary considerably, depending upon which of these sets of information becomes available to the decision maker.

The availability heuristic also affects our probability estimates when we construct a *causal scenario*. Let us say that you are asked to indicate the likelihood that there will be increased earnings in the forest products industry next year. We may try to imagine a series of events which would lead to such an increase in earnings, a causal scenario. We then make a judgment about the likelihood of an increase in earnings, by the *ease* with which we were able to imagine the scenario. Here is one way this might occur. Let us say that you were watching the PBS television show, "Wall Street Week," and a reputable analyst explains why she believes the forest products industry will post record earnings next year. She constructs a cogent scenario leading to this event, a scenario including a period of declining interest rates. Several days later you read in the paper that there was a downtick in interest rates. The analyst's well-constructed scenario easily comes to mind and suddenly it seems as if the forest products industry is on the verge of an earnings explosion. In this example, the causal scenario was "pre-constructed" by the analyst, but if we were trying to construct the scenario ourselves, the same principle would apply: the easier it is to imagine a chain of events leading to an end state, the more likely that end state will appear to be.

Biases in the Selection and Effect of Information

Modern cognitive psychology suggests that our probability estimates tend to be more influenced by some kinds of information than by others. To experience this process, try answering the next two questions:

Question 3. A recent survey of occupations in a large midwestern city found that among its inhabitants were approximately 1250 long-haul truck drivers and about 30 professors of classical literature. Consider Bill, a resident of that city. Bill is short, slim, and likes to read poetry. Is Bill more likely a professor of classical literature or a long-haul truck driver?

Question 4. A panel of psychologists interviewed a sample of 30 engineers and 70 lawyers, and summarized their impressions in thumbnail descriptions of these individuals. The following description has been drawn at random from the sample of 30 engineers and 70 lawyers.

> "Jack is 39 years old. He is married with no children. A man of high ability and motivation, he promises to be quite successful in his field. He is well liked by his colleagues." Is Jack more likely to be an engineer or a lawyer?

In both questions there are two kinds of information available: singular or *case data*, consisting of evidence about the unique case under consideration (Bill, Jack), and distributional information or *base-rate data*, consisting of knowledge about the distribution of relevant events (numbers of truck drivers, lawyers, etc.). Considerable research has shown that in making these judgments we tend to focus on the unique qualities of the case data and fail to give sufficient weight to the base-rate data. In answering question 3, we tend to focus on Bill's personal characteristics, and then employ the *representativeness heuristic* to arrive at a decision. Since most people believe that "slim, short, and reads poetry" are characteristics representative of classics professors, they assign Bill to that category. The fundamental error here, however, is the inappropriate neglect of the fact that there are far more

truck drivers than classics professors. Indeed, when people know only the base rate and nothing about Bill they most often, and correctly, conclude that Bill is more likely to be a truck driver. The addition of case data, however, causes them to almost completely ignore the base rate.

It is not inappropriate to factor case data into our assessments. However, in comparison to normative statistical models, models that may sometimes be used to estimate how much we should adjust for the case data, we appear to weigh this kind of specific information far too heavily. The same bias affects answers to question four. Here the case data is intentionally neutral in its implications for classifying Jack as a lawyer or engineer. As a result, most people neglect the base rate showing more than twice as many lawyers in the sample from which Jack was randomly drawn, and erroneously conclude that Jack is equally likely to be a lawyer or engineer. Kahneman and Tversky believe that this prevalent tendency to insufficiently consider information about the distribution of similar events is perhaps the major error of intuitive predictions. They also suggest that this error is likely to be magnified the better one is acquainted with the specific case or the more intensely involved one is with it.

Consider an example of how this bias might operate in an investment decision context. Widget-Tech is going public. Let us imagine a base-rate distribution such that 70 percent of new issues are lower in price 12 months after they are on the market and 30 percent are higher in price. What is the likelihood that Widget-Tech will end the next 12 months higher in price? Using the base-rate information alone you might be inclined to let other investors subscribe to this new issue. Now, consider the unique case data your broker provides you about Widget-Tech. It makes rapid pulse lasers for storing data on optical disks. Lasers. Optical disks. Data storage. There goes the base-rate; here comes the representativeness heuristic. Furthermore, if your source of this information is someone you "know" inside the company, you are even more likely to neglect the base-rate data.

Why do we neglect base-rate data and overemphasize case data? While research is still continuing into this question, there are some useful clues. To begin with, the delivery system for these two kinds of information is different: base-rate data is typically obtained from printed tables or figures while case data is often received in the context of a "personal" encounter, e.g., from a friend, broker, TV commentator, etc. Cognitive research suggests that our information processing system is more readily attuned to vivid, concrete, often emotionally laden case information than it is to pallid, dry, statistical data. In other words, vivid case information may be more available in memory and, in accordance with the *availability heuristic*, will be more likely to affect our probability judgments. To illustrate, consider:

Question 5. Which of the following causes a greater percentage of deaths in the United States each year?

a) accidents *or* cardiovascular disease (e.g., strokes, heart attacks)
b) suicide *or* homicide

Most people estimate that accidents kill as many or more persons each year as cardiovascular disease and that homicide is a more frequent cause of death than suicide. In fact, there are 10 times as many deaths due to cardiovascular disease than accidents and 30 percent more deaths are due to suicide than homicide. Moreover, contrary to most people's estimates, asthma causes more deaths than tornadoes, and diabetes more deaths than breast cancer. Professors Slovic, Fischoff and Lichtenstein, who have gathered these observations, believe that the vivid portrayal in the media of homicide, accidents, tornadoes, etc., make these causes of death more available in memory. When we make use of the availability heuristic to estimate these events' probabilities, the more vivid ones are judged as more likely to happen.

We can understand now why a stock market "tip" from a friend, broker or the media may have a far greater impact on our choice of investments than duller, but presumably more reliable, investment data obtained from statistical tables or charts. A vivid, perhaps even dramatic, tip will be more available in our memory. Then, when we apply the availability heuristic, the events promised in the tip will seem most likely to occur.

Inconsistencies in Judgment

There are two kinds of judgmental inconsistency that an investor should be concerned about: *random inconsistency*, in which we make randomly different judgments of the same event, usually because of transitory conditions such as boredom, fatigue, distraction, etc., and *systematic inconsistency*, in which we make different judgments of the same event, because of some systematic bias in the way we think about the event. Random inconsistency was manifested when expert radiologists, observed as they attempted to detect the presence of lung disease on x-ray films, changed their minds in about 20 percent of the cases when reading the

same film on two separate occasions. In another study, this one of horse race predictions, expert handicappers changed their mind on between 22 percent and 39 percent of their judgments during repeated appraisals of several races. Surprisingly, the more information the handicappers had to work with, the greater their inconsistency! No wonder we adopt heuristics to simplify our decision making. If you have ever wondered "How could I have decided to buy that?" part of the reason may have to do with the random inconsistency of our judgments.

One factor that may produce random inconsistency seems especially worthy of note. Recent research conducted by Professor Gordon Bower of Stanford University suggests that our moods or feelings may influence our probability estimates. According to Dr. Bower's data, people who are happy or in a good mood tend to think it more likely that good things will happen to them and less likely that bad things will befall them. People in a depressed or unhappy mood view the likelihood of good and bad events befalling them in the reverse manner. Thus, depending on your emotional state of mind, it may seem that the security whose purchase you are considering is either about to double in price or to lose half of its market value. Professor Bower speculates that the basis of this mood-probability estimation effect may be the *availability heuristic*: When we are in a good mood we tend to retrieve positive or bullish facts from memory; when we are in a bad mood, we tend to recall negative facts. As we have already seen, the information available in our memory influences our probability judgments.

Random inconsistencies tend to cancel each other out over time. Of greater concern is investment judgment that displays systematic biases that lead to logically inconsistent choices. Consider the following question:

Question 6. Do you prefer a sure gain of $80 *or* a risky choice that offers an 85 percent of winning $100 and a 15 percent of winning nothing?

Most people prefer the sure $80, in spite of the fact that the uncertain choice has a higher "monetary expectation." (The monetary expectation or "expected value" of an alternative is the sum of its outcomes weighted by their probabilities. Hence, $100 x .85 plus 0 x .15 = $85 [versus $80 × certainty = $80]. The monetary value is what one would expect to get if the uncertain alternative were taken many times, i.e., the average gain.) Most investment theory assumes that investors are "risk averse," that they prefer a certain outcome to a risky alternative with an equal monetary expectation. Investors would be "risk seeking" if they reject the certain outcome in favor of the uncertain one. Economists generally assume that the rational investor is risk seeking only when the expected monetary gain of the risky alternative is large enough to compensate for taking a risk and foregoing the certain gain. Now consider the following choice:

Question 7. Do you prefer a sure loss of $80 *or* a risky choice that involves an 85 percent chance of losing $100 and a 15 percent chance of losing nothing?

Here, the risky alternative does *not* offer a greater monetary expectation (greater gain or lesser loss) to compensate for the risk. Rather, it offers the prospect of a greater loss. Thus, the rational investor should logically, and consistently, be risk averse and choose the certain loss. However, many studies indicate that a large majority of persons faced with this choice select the uncertain loss, i.e., they are risk seeking. In other words, our risk preferences are systematically inconsistent: we are risk averse in the face of gains, but risk seeking in the face of losses.

Many economists, decision theorists and psychologists have grappled with the theoretical implications of this inconsistency, and it is not our intention here to delve into these theoretical issues. Rather, let us consider the implication of this inconsistency for helping us to understand what appears to be a widespread bias in investor decision making.

In a fascinating volume, titled *Winning: The Psychology of Successful Investing*, Srully Blotnick reported the results of a 10-year study of over 1,100 investors. This report, published by McGraw-Hill in 1979, constitutes one of the few publicly available surveys of actual, ongoing investor behavior. In his book, Blotnick notes that many investors at one time or another owned a stock that had either a spectacular rise or decline. What was especially notable, however, was that in these cases an investor was almost 21 times more likely to "ride a spectacular stock all the way down than to ride it all the way up." This pattern of behavior, points out Blotnick, is in direct contradiction to one of the most widely given pieces of advice on Wall Street: to wit, cut your losses and let your profits run. Why do so many investors appear to behave in a manner opposite to this dictum? Blotnick suggests that a rising stock creates a state of tension in the investor, tension that can be relieved by the sale of the security; a declining stock leads to a state of depression and helplessness that effectively paralyzes the in-

vestor into inaction.

From a cognitive perspective, however, something else may be going on as well. Consider the investor who owns a stock that has appreciated in price. That investor is now faced with the choice of a certain gain from the sale of the security or the prospect of holding the stock for a greater but uncertain gain. As we have just seen, for the risk averse investor it "makes sense" to take the certain gain. Now consider the investor who is holding a security that has declined in price. That investor is faced with the choice of a certain loss or the uncertain prospect of an even greater loss. Again, as we have just seen, when faced with losses investors appear to act as risk seekers. Under loss conditions, in other words, it appears to "make sense" to select the risky alternative and hold the security. The result of this systematic inconsistency is to suggest that we sell our winning stocks and hold our losers.

The bias in our decision making induced by this inconsistent risk preference is aggravated by the way we define or "frame" the consequences of choice. Consider two securities. Security A, let us say, has a 70 percent chance of bringing a profit, security B a 60 percent chance. We could just as well say that security A has a 30 percent chance of bringing a loss and security B a 40 percent chance. Whether we compare the securities in terms of potential gains or potential losses logically should have no effect on our decision, since the two frames are merely opposite sides of the same coin. Psychologically, however, the choice of frame can produce a considerable effect. Professors Tversky and Kahneman asked a large number of physicians to choose one of two programs of treatment for "an outbreak of an unusual Asian disease expected to kill 600 persons." In one frame, the choice was presented in terms of gains, i.e., in terms of lives potentially saved. Here the choice was given as between program A that would certainly save 200 lives *or* program B that offered a one-third probability that 600 lives would be saved coupled with a two-thirds probability that no lives would be saved. When the programs were presented in these terms, the majority of the physicians were risk averse; they preferred program A that offered a certain saving of 200 lives. Then the two *identical programs* were presented to another group of physicians in terms of lives lost, i.e., program A which would certainly result in 400 deaths *or* program B with a two-thirds probability that all 600 would die coupled with a one-third probability that nobody would die. When the choice of programs was presented in terms of lives to be lost, the majority of physicians chose the uncertain or risky program B. In other words, the physicians were risk averse in the face of gains, and risk seekers in the face of losses. But, the perception of gain and loss was induced entirely by the framing of the problem; the "real" choice was always the same.

Once again, the implications for investors are clear. Comparing choices in terms of profits or positive rates of return is likely to induce a different risk perspective, and consequently a different decision, than might occur if the same choice was framed in terms of losses. Whether we view the cup as half full or half empty is entirely in the eye of the beholder, but it appears to have an important effect on our decisions.

Summary Thoughts and Some Corrective Procedures

In this article, we have looked at investor decision making through the eyes of the cognitive psychologist. We have seen that man has limited capacities as a processor of investment information. Like all human decision makers, investors operate according to what Herbert Simon called "the principle of bounded rationality." Professor Simon, recipient of the Nobel Prize for his studies of organizational decision making, believes that we develop a much more simplified and personally biased view of decision problems than is typically assumed in the formalized, normative models employed by economists.

Not that Professor Simon, or other students of human cognition, view their research as deprecating the extraordinary cognitive skills possessed by man—quite the contrary. Expanding knowledge about how we learn and remember, about how we make inferences and judgments, makes clear that our abilities in these domains are marvelously complex. However, the same skills and heuristic strategies that allow our greatest achievements also can lead to errors in judgments with important consequences for investors. Baruch Fischoff expressed this point eloquently. Our judgment skills, he noted, "are good enough to get through life, poor enough to make predictable and consequential mistakes; they are clever enough to devise broadly and easily applicable heuristics..., unsophisticated enough not to realize the limitations of these heuristics."

Many readers were probably familiar with the concept of heuristics before reading this article, if not with the term itself. Consider some of the rules of thumb that abound on Wall Street: "concentrate on emerging growth stocks;" "buy low PE stocks;" "cut your losses and let your profits run;" "buy the

strongest stocks in the strongest groups." These "rules-for-investing" are heuristics, and one of their values is that they keep us from being overwhelmed with information. Cognitive psychology is now making us aware that we use similar rules of thumb at a much more fundamental level of investment decision making: in assessing the likelihood that an event will occur, in judging the risks and benefits of alternative courses of action, even in selecting information to attend to. At this more basic level, however, we may be unaware of our reliance on heuristics and unsuspecting of the biases and errors to which these useful strategies can lead.

Are there any tactics we can use to counter these biases? Recognizing that increased awareness is itself an important first step toward "debiasing" efforts, and with a caveat that the correctives offered here are only suggestive, what follows are a few possibly useful tactics.

1) **Altering availability** One technique for countering the biases induced by our use of the availability heuristic owes its inspiration to Louis Rukeyser, host of "Wall Street Week." As discussed earlier, guests on this program often construct persuasive and easily remembered causal scenarios to support their predictions about some future event, e.g., improved industry conditions, higher market averages, etc. After such a presentation, Rukeyser may ask his guest, "What could happen to change your mind?" From the perspective of cognitive psychology, this question forces the guest to provide the audience with an alternative causal scenario, one that makes an alternative forecast equally probable because it is now equally available in our memory. Try this tactic the next time a broker or financial adviser paints a particularly rosy picture for an investment. We can also apply this technique to our own forecasts. After constructing a scenario to explain why Widget-Tech's earnings are going to increase 20 percent a year for the next five years, construct another scenario that would lead to flat or even declining earnings. You will find that the 20 percent a year gain no longer seems quite so certain. Additionally, you will be in a more unbiased position to evaluate the causal chain that led to the initial forecast.

2) **Doing contrarian "exercises"** Being on guard against the availability heuristic is also important in combating inconsistency in our judgment, especially that induced by extremes of emotion. Recall Professor Bower's finding about emotion's effects on judgments of the likelihood that good or bad things will happen to us. It is easy to see how the euphoria occurring at market tops and the gloom and depression that are prevalent at market bottoms can bias our forecasts. These extreme emotional conditions would make bullish or bearish thoughts more readily available in our memory in a particularly undiluted form. At such times the ability to act contrary to the prevailing mood may be facilitated by contrarian "exercises;" development and rehearsal of causal scenarios with outcomes opposite to that which everyone else is expecting. The increased strengthening or availability of these "contrary cognitions" may serve as an antidote to the highly available news supporting the prevailing psychological environment. Beware if you find that you "can't think of anything that could turn this market around;" that "can't think of" means your availability heuristic is hard at work against you!

3) **Fixing the frame** This next suggestion is a bit more demanding. We have seen that how we frame an investment decision, e.g., whether in terms of future gains or future losses, may bias our risk preference. The dangers of this bias are probably less when applying a formal model of investment decision making, e.g., modern portfolio theory, where risk and reward are framed in more abstract, quantitative terms. More typically, however, we tend to think in terms of vivid, concrete, dollar denominated gains and losses, especially when confronting decisions about assets that have advanced or declined in price. How might we deal in a less biased manner with the decision over whether to hold these assets or sell them to realize a gain or loss? One useful tactic may be to "reframe" the sell alternative. Rather than thinking of the sale as resulting in a certain gain or loss, think of the proceeds of that sale as being invested in another risky asset, along with its attendant uncertain outcomes. Instead of being confronted with the alternatives of an uncertain hold versus a certain profit or loss, a condition that seems to produce significant biasing, the decision is "reframed" as a choice between two uncertain holds. In effect, this tactic frames the investment decision more clearly and appropriately as one of relative risk management, rather than as a choice guided by a dubiously applicable heuristic like "a bird in the hand is worth two in the bush."

4) **Basing forecasts on the base rate** If Professors Kahneman and Tversky are correct in their belief that overweighing the unique aspects of a case is the major error behind our intuitive predictions, then we need to be on guard not to neglect information about other cases of the same general class, i.e., not to ignore base-rate data. Analysts who prognosticate about future market directions pro-

vide examples of both the use and abuse of base-rate information. A positive example of using base-rate data is afforded by the commentator who evaluates the present market advance (or decline) by comparing it to a statistical summary of past market advances. For example, if the analyst provides data on the *average* duration or magnitude of previous bull markets, along with the *range* of those figures, i.e., the lengths of the shortest and longest advances, then she is providing useful base-rate information against which to judge present market conditions. However, the analyst who says "this advance reminds me of the 1951-53 market runup" is committing a double heuristic error: he is relying on *availability* of past instances in memory ("reminds me of") and is selecting some particularly salient, *representative* features of the present market to define a comparison. This error is the same as that committed in judging whether Bill was a classics professor or a truck driver by using the representative case information, "short-slim-reads poetry," and ignoring base-rate information about the number of professors and truckers. Of course, using base-rate information about past markets does not guarantee accurate predictions about future markets. It does, however, provide some protection against completely unrealistic expectations based on the current "case" and our involvement with it.

Indeed, none of the psychological insights or tactics presented in this article can give assurance of forecasting success. Furthermore, they certainly do not guarantee increased profits or rates of return. Their importance is best indicated by a line from *The Money Game*, Adams Smith's witty and literate survey of Wall Street and its psychology: "If you don't know who you are, the stock market is an expensive place to find out." Knowing who you are, whether in terms of your investment goals, your personality and motivations or the cognitive processes described in this article, may not result in larger profits, but such self-awareness may lead to smaller losses.

Finally, I would offer a caution to the reader who is provoked to conclude from this article that the biases and errors of individual investors "disprove" the thesis that the market acts as an efficient and rational price setting mechanism. Quite the opposite conclusion could be drawn. It can be shown on statistical grounds that when a reasonably large number of independent biases are operating together, a kind of "cancelling out" occurs. Consistent with the efficient market hypothesis, the collective, or *net*, effect of our individual cognitive biases and errors may be a market judgment that is more correct, more of the time, than any single one of us.

Cognitive Sources of Common Errors in Investor Judgment

Judgmental Bias or Error	Cognitive Source
1. Being bearish at market bottoms and bullish at market tops	Unavailability of a "contrarian," causal scenario
2. Selling winners and holding losers	Inconsistent risk preference and the biasing effects of decision frames
3. Over-optimistic appraisal of a security	Applying the representativeness heuristic to case data while neglecting base-rate data
4. Acting on unreliable "tips"	Vivid, "personal" information is more memorable than statistical data
5. Indecision and/or inconsistency	Effect of emotion on information availablity and probability judgments
6. "It's the 1960s again!": misleading market metaphors	Misapplying the availability and representativeness heuristics
7. The "sure thing" that wasn't	Over-reliance on one available, causal scenario

$ $ $

Dr. Pines was an undergraduate major in finance and investments at City College, CUNY before changing over to psychology. He received his Ph.D. in that field from SUNY at Buffalo in 1971 and is currently Associate Professor in the Psychology Department at Canisius College, Buffalo, New York. This article grew out of a course in "The Psychology of Investment Decision Making" that he recently developed for finance and psychology students at Canisius. Dr. Pines and his students are currently working on an "expert-based" computer system to mimic human stock market judgment.