

Investors Do Not Mind Their Utility Functions

The concave utility function that is both necessary and sufficient for the risk premium has implications for more than mere asset returns. Most fundamentally, it has implications about behavior and happiness. For example, people should be minimizing risk when they can, and they should be happier when they have more money. But even in these dimensions, the theory has been known to fail miserably, yet this failure is often seen as immaterial. Yet, given the failure of the pricing implication, it is not as if we can ignore this failure, because it gives us further evidence to suggest the foundation is seriously flawed, as opposed to imperfect.

Economists are used to dealing with what appears to many as absurd assumptions, because they often don't matter. In Milton Friedman's essay, *The Methodology of Positive Economics* he explains that assumptions should be judged by their consequences alone.¹ He discusses the example of how the expert billiard player makes shots as if "he knew the complicated mathematical formulas that would give the optimum directions of travel, could estimate accurately by eye the angles, and so forth, describing the location of the balls, could make lightning calculations from the formulas, and could then make the balls travel in the direction indicated by the formulas. Our confidence in this hypothesis is not based on the belief that billiard players, even expert ones, can or do go through the process described." In sum, if the assumptions work in predicting things, good enough for scientists.

As the implications, the predictions, of the risk-return theory do not work in general, the assumptions are fair game for examination. That is, perhaps we have had a bad century, and should believe in the theory because, fundamentally, the assumption and logic are correct, only the empirical implications counterfactual. Surely someone intelligent, someone thinking ahead, should be biased toward the theory, not data, because we know theories only work on average, and perhaps our sample is not sufficiently

large, or is biased. But, in face of bad evidence, an examination of the assumptions is appropriate. The common assumption that people are rational is often criticized because one can always find many examples of irrational behavior.

The assumption that people act as if they were rational generates lots of testable assumptions that seem to work: people buy more when prices are relatively low, and less when prices are relatively high, there are very few examples of arbitrage opportunities because others exploit them to make money, and so on. The alternative to rationality cannot be unspecified irrationality, but rather, a specific type of irrationality and here's where irrationality falters. Generally, irrationality deviates from rationality in an unbiased way, with just as many overconfident as timid, or those who underweight base-rate information versus those who overweight it. The burden is on those in favor of irrationality to show in what circumstances it is too much or too little.

Furthermore, heroic assumptions about zero transaction costs, no taxes, limitless borrowing, and perfect information, allow economists to isolate the main drivers of an idea by a model showing the bare logic of the theory presented, hopefully allowing one to test it, to identify how the violations of the theory cause deviations. A classic example is the Miller-Modigliani Theorem, which says that the debt-equity ratio does not affect the firm's value. In practice, the debt-equity ratio is very important for a company and everyone knew that, and so the Miller-Modigliani Theorem highlights the specific assumptions needed to generate this effect (often, asymmetric information, taxes, and so forth). Ideally, a model becomes not just a metaphor, but when parameterized, a realistic description of some interesting phenomena.

For these reasons, we can understand why the unrealistic assumptions of the CAPM, and the utility function that underlay it, were not seen as necessarily fatal. But the assumptions of the CAPM are truly in a class of their own. Sharpe's seminal paper on the CAPM that was to win him the Nobel Prize was initially rejected because of this very assumption. Dudley Lockett, the editor at the *Journal of Finance*, informed Sharpe that his assumption that all investors had the same expectations was so preposterous as to make his conclusions uninteresting. A new editor came in, and Sharpe's paper was published in 1964.²

Jason Zweig wrote a book about investing called *Your Money & Your Brain*, and he recounts a story about Harry Markowitz, who was then working at the RAND Corporation and trying to figure out how to allocate his retirement account. He knew what he should do: "I should have computed the historical covariances of the asset classes and drawn an efficient frontier." But, he said, "I visualized my grief if the stock market went way up and I wasn't in it—or if it went way down and I was completely in it. So I split my

contributions fifty-fifty between stocks and bonds.” So even the originator of the MPT did not follow its implications back in the day when he created the efficient frontier, and was thinking instead about his relative performance.

And how was it even possible for investors to know about all those betas (implicitly, the covariances)? Arbitrage was from buying 1.5 beta stocks and shorting 0.5 beta stocks, creating a zero-cost 1.0 beta portfolio. How could they do this if, in the days before computers, they didn't have the ability to calculate betas? Presumably they intuited the betas, on average, and the wisdom of crowds generates such insights. The concept of *emergence* is the paradoxical ability of a swarm to exhibit greater intelligence than the individuals that make it up. Emergent behavior is evident in several natural realms, including ant colonies, brain cells, and city neighborhoods. All these systems solve problems by drawing on masses of individually stupid elements, rather than a single, intelligent executive branch. Thus, the market is, in theory, a prominent emergent thing, created by a bunch of semi-smart, semi-stupid individuals, to create what is commonly referred to as a hypersmart individual, a representative, rational, agent creating arbitrage portfolios.

But in economics you find that when someone proves that something can exist, others take this as a license for assuming they do. Assumptions are not expected to be true, so as long as there exist some assumptions that work, that is good enough. Tractability of the modeling, the usefulness of the assumptions, is in practice more important than the realism of the assumptions.

Those are assumptions whose violations are curious, but hardly fatal, because it is plausible these deviations from theory do cancel out. However, these are not merely pricing implications for the MPT, but behavioral implications, and these fail massively. Most theories are designed to explain some particular thing, but if they are describing reality, there will be other implications. False theories are actually best figured out through these incidental implications, because usually the theory was created post hoc, so it will explain the data very well on one issue. The MPT as originally constructed was not a post hoc theory explaining the data, it predicted a relationship that was not obvious. Only after its failure was documented, did it morph into a series of parochial, ephemeral, risk factors under the rationalization of the general equilibrium approach.

Let us consider the main behavior failures one by one.

BEHAVIORAL VIOLATION 1: INVESTORS TRADE TOO MUCH

In theory, investors buy efficient portfolios, and adjust their weightings based on different preferences, wealth, and so on. They do not buy and sell stock

based on bullish and bearish views on individual securities, because they all agree on the expected returns, and covariances, of those securities.

Thus, even those models that allow for asymmetries in information across traders, the volume of trade is mainly affected by unanticipated liquidity and portfolio rebalancing needs of investors. However, these motives would seem to be far too small to account for the tens of trillions of dollars of trade observed in the real world.³ The turnover of a passive fund is around 10 percent a year—and passive funds are what investors should buy, according to standard theory. In contrast, the average annual turnover rate on the New York Stock Exchange is currently around 100 percent, which implies that people are not buying factor proxies and rebalancing as their risk preferences change, but rather, reading the news and trying to get in front, or out of the way, of the next big move. This dissonance has led even the most ardent defenders of the traditional pricing models to acknowledge that the bulk of volume must come from something else—for example, differences in prior beliefs that lead traders to disagree about the value of a stock even when they have access to the same information sets. Nonetheless, being off by a factor of 10 on trading volume suggests the canonical model is missing something fundamental.

BEHAVIORAL VIOLATION 2: TOO MANY FUNDS

In the CAPM, there is the One-Fund Separation Theorem, which states that everyone invests in the same risky portfolio, which is the market portfolio. From this fact, the linearity of return in beta is a function of arbitrage. For the APT, this argument is generalized to a handful of factors that are unspecified—maybe oil, the dollar, and so on. The driver is the same idea, that investors merely want access to a factor proxy. In contrast, there are *thousands* of funds available, and new ETFs, or stocks representing passive indexes, are created every day, 1,000 times as many as implied by theory. One could argue that there should be more than one fund for each factor on institutional grounds, that some marketing issues are important, but when theory says, “five” and reality says, “13,000,” this is a material miss.

BEHAVIORAL VIOLATION 3: UNDERDIVERSIFICATION

William Goetzmann and Alok Kumar (2005) document extreme underdiversification among investors using more than 40,000 equity investment

accounts.⁴ More than 25 percent of investor portfolios contained only one stock; more than 50 percent of them contained fewer than three stocks. This is massively irrational nondiversification in the CAPM context, as people appear to be assuming diversifiable risk when they do not have to. From a Sharpe ratio perspective, this can only be justified if these investors generate large returns above the market, when in fact evidence suggests traders are, on average, worse than average.

Alternatively, one potential explanation for underdiversification is that investors may consciously choose to remain underdiversified so they can increase the likelihood of extreme positive returns, or in other words, to capture higher levels of positive skewness in their portfolios. If investors have a preference for skewness, they prefer positive skewness in return distributions. Diversification then is a two-edged sword: it eliminates undesired variance in return distributions, and also eliminates desired skewness. Consequently, portfolios that are efficient in a mean-variance-skewness framework only appear to be inefficient when evaluated using a mean-variance framework.

BEHAVIORAL VIOLATION 4: NO FUNDAMENTAL ANALYSIS

Asset pricing theories generally assume investors are all agreeing about returns on stocks, and no one is doing any fundamental analysis on cash flows. Eugene Fama says fundamental analysis is a waste, in that markets are pretty efficient, at least sufficiently so that merely reading about a firm's financial statements is useless, and given the lame results of analyst recommendations or mutual fund returns, this seems rather obvious. There are not any obvious examples where a group of investors are outperforming passive investors: mutual funds, retail trader, equity analysts, so when one argues that some investors pay for more information and therefore garner higher returns because of it, no one has identified such information, other than material trade secrets that are illegal in most markets. When asked for evidence that such professionals have value, one is usually greeted with anecdotes: Warren Buffett, Peter Lynch, and so on. But these anecdotes prove nothing; they are, at best, exceptions to the rule.

Theory and evidence argue for the futility of fundamental analysis, yet it flourishes. There is a large economy built on feeding the demand for fundamental analysis, and people spending time, money, and exposing themselves to idiosyncratic variance based on the idea that the more research you do, the more you listen to analysts discussing companies on TV, the more money you

will make. Your average broker has a large set of tools and experts aimed at meeting this demand, which seems, in aggregate, irrational, because in aggregate there is no evidence that it delivers value.

BEHAVIORAL VIOLATION 5: BUY RECOMMENDATIONS EXCLUDE FIRMS WITH *MERELY* LOW RISK

In the CAPM world there should be many situations in which a lower-than-average-risk stock will be expected to outperform on a risk-adjusted basis, but not an absolute basis, yet you will never see a brokerage tout a buy or strong buy recommendation that has a lower-than-average market return.⁵ If people are simply maximizing their *risk-adjusted* returns, why are there *never* any recommended stocks that are expected to underperform the market, though at sufficiently low risk to be attractive? If we go outside of brokerages, the less-sophisticated Internet touts of bizarre investment schemes all have one common characteristic: They all promise greater-than-average returns. That lower-than-average return recommendations are absent suggests people only take risk if they expect an absolute greater-than-average return, which would seem to exclude half of all assets.

BEHAVIORAL VIOLATION 6: AGENTS DO NOT AGREE

The assumption that agents agree on asset return, variances, and covariances, is rather absurd, and indeed, was why Sharpe's initial submission of what was to be called the CAPM was rejected. But for some reason this did not seem a big deal after a while. Rubinstein's Aggregation Theorem, in 1974, outlined the assumptions necessary to represent the economy as a single representative agent. It is common to model the consumer for the United States as a single person, which makes models tractable. The assumptions necessary for this, however, are rather severe. In essence, Rubinstein says that if everyone has the same beliefs and preferences, wealth, patience, and there are complete markets, you can caricaturize them as a single person. The CAPM assumes investors have identical beliefs, too. While this can get very technical, at some level it is rather obvious: If everyone is the same, you can model an economy as a large, single person.

But in practice this is quite obviously violated, as disagreements on stock prospects are a staple of financial journalism. Wealth and preferences, also, are quite different from person to person. Ed Miller postulated a simple idea, that if we prohibit short selling, then the greater the disagreement

for a particular stock, the greater the price. This is because if you assume expectations are normally distributed around a mean, the greater dispersion means the most optimistic will be very optimistic, and these optimists will set the price. That is, say there is an asset with a mean price of 100, but people have a distribution of standard deviation of 10 points around that. If the top 15 percent set the price, its price is 115. But if the standard deviation is 20, the top 15 percent would set a price of 130. This higher price means a lower future return, so several researchers have looked at opinion dispersion to predict returns using that as a model, but its success has been disappointing.

Several researchers have recently used the Miller model to motivate using opinion dispersion, or even raw uncertainty, as explaining returns, especially in empirical examinations of the effect of analyst disagreement and returns. The key is that models can be very elegant and consistent, but inevitably they have extremely unrealistic assumptions—like the CAPM or general equilibrium models with a representative agent, or they can be rather ad hoc, such as Miller's, which are more intuitive but less rigorous.⁶ The problem is that economics has big problems modeling different beliefs, because under standard assumptions, it is irrational to agree to disagree. Of course, if the Miller model worked, it would bring forth many troublesome implications because disagreement measures are generally correlated with volatility, so it would imply that higher volatility is correlated with lower returns. But, leading researchers invoke the Miller hypothesis because it seems to explain a rather obvious pattern, that highly volatile assets have lower returns than average. Economics likes rigor, but over time, a simple story like Miller's, inconsistency with standard theory notwithstanding, will be accepted if it works.

BEHAVIORAL VIOLATION 7: THE HOME BIAS

One of the easier ways to reduce portfolio volatility is to invest in more than one's own home country. Further, if covariance between the market portfolio and aggregate consumption growth are correlated in a particular country, the move to other countries would reduce this correlation as well. It has long been known that, despite the gains from cross-border diversification and the increased integration of financial markets worldwide, the strong investor preference for domestic firms is pervasive in international financial markets. This home bias phenomenon is ubiquitous across developed and developing markets, that is, worldwide. In a more recent study based on worldwide equity fund holdings data in 1999 and 2000, Chan, Covrig, and Ng (2005)⁷ document the existence of home bias in every single country in their sample of 48 countries around the world. Plausible explanations for the investor

strong preference for domestic equities include the existence of cross-border boundaries that give rise to exchange rate risk, variation in regulation, taxation, accounting standards, corporate governance, and transaction costs, among others, information asymmetries, and more recently, investor behavioral factors (familiarity, culture, language, and geographic proximity).

THE ROTTEN CORE: THE UTILITY FUNCTION

The genesis of the risk premium is based on the standard utility function, which is concave in wealth, broadly defined. It generates some reasonable implications, in that faced with a random payout worth \$50 on average, most people prefer the average value, the \$50. But clearly, this should imply volatility of some nature would be positively correlated with returns, and here the theory fails. Economists postulate that people maximize utility and that utility is whatever people like “Consumer’s market behavior is explained in terms of preferences, which in turn are defined only by behavior.”⁸ Thus, considering the behavior and pricing anomalies discussed earlier, this suggests something is wrong. Looking further at utility, we see this might be the basis of our problem.

ABSURD EXTRAPOLATIONS

Matthew Rabin of the University of California at Berkeley notes that a consumer who from any initial wealth level turns down gambles where she loses \$100 or gains \$110, each with 50 percent probability will turn down 50-50 bets of losing \$1,000 or gaining any sum of money. Mathematically, you can extrapolate this using the pure logic from the concavity of a utility function that generates this property, to find that this implies you would reject an offer to gain \$1 billion dollars and lose \$1,000. Any utility function that is strictly concave, that would imply you do not like modest gambles, implies that you are extremely averse to really favorable gambles with 10 times the exposure—in other words, an absurdity.⁹

The major solution to the problem is the one proposed by Prospect Theory, where outcomes are considered relative to a reference point (usually the status quo), rather than to consider only the final wealth. This theory was championed by Nobel Prize winners Kahneman and Tversky in 1979.¹⁰ The essence of the approach is to put a strange wiggle in utility curves right around the current wealth level. Please see Figure 5.1.

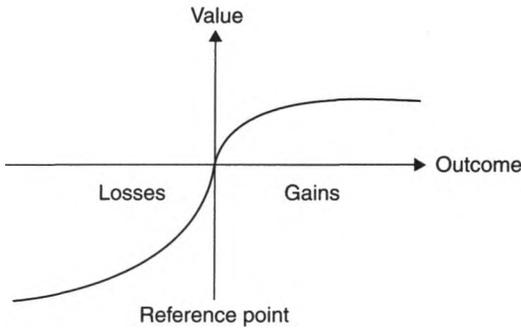


FIGURE 5.1 The Utility Function in Prospect Theory

The utility curve tries to capture three stylized facts:

1. Risk-averse behavior in gains (concave for right-hand side)
2. Risk-loving toward losses (convex for left-hand side)
3. Enjoy gains less than dislike of losses

This is an explanation for the fact that people simultaneously buy lottery tickets and insurance, but still their invest money conservatively. Lottery tickets are small average losses, and people are risk seeking here. Insurance is against big losses, and people are very risk averse here. Gains are also treated with moderate risk aversion. Indeed, Milton Friedman and Leonard Savage wrote about the first two points back in 1948, trying to explain why people buy lottery tickets, and also insurance, but noting the utility curve had an inflection point that allowed gambling, basically, the utility curve discussed earlier.

The main observation of Prospect Theory is that people tend to think of possible outcomes relative to a reference point rather than absolute wealth, a phenomenon which is called a framing effect. What affects the reference point? Researchers have suggested many things related to a person's recent history or the cross-sectional averages. When losses are looked at differently from gains, but only for certain-sized bets, theory begins to look more like an explanation of various choices—sometimes risk seeking, sometimes risk averse, depending on the data.

In 1952, Harry Markowitz tried to build on the Friedman-Savage model with a thrice-inflected utility of wealth function and argued that whenever an individual's wealth was at its customary level, it would be at the second

inflection point.¹¹ The problem with the Markowitz model was well explained by Alchian in 1953:

Markowitz recognized that until an unambiguous procedure is discovered for determining when and to what extent current income deviated from customary income, the hypothesis will remain non-verifiable because it is not capable of denying any observed behavior.

It is not clear at all how to make a general equilibrium model, where everyone has a quirky utility functions, some risk averse, some risk seeking, depending on whether they are buying or selling the bet in question. Since Kahneman received his Nobel Prize in 2002, many have noted the potential for behavioral finance to explain anomalies because it encompasses real people, who are often irrational, as opposed to *homo economicus*. The problem, however, is an embarrassment of riches. By allowing people to selectively apply risk aversion or risk seeking, by way of prospect theory, and various reference points therein, the “risk must be compensated by return” story clearly is qualified to a degree that makes it unrecognizable.

Way back in 1951, Nobel laureate George Stigler noted “each decade, for the past nine or ten decades, economists have read widely in the then-current psychological literature. These explorers have published their findings, and others in the field have found them wanting—wanting in useful hypotheses about economic behavior.”¹² And so it is with behavioral finance, which attempts to apply, selectively, risk loving and risk aversion, or anchoring and extrapolation, at various problems. The problem with this general approach is like the problems with specific non-normal distributions: they explain too much, because they predict biases exist in opposite directions, depending on the data. As Eugene Fama notes in his essay *Market Efficiency, Long-Term Returns, and Market Efficiency*, in an efficient market, deviations from efficiency will be half overreaction, half underreaction, which is about what the behavioralists have applied to finance. While I agree that the current paradigm is flawed, an alternative must be somewhat specific.

EASTERLIN'S PARADOX

The fundamental assumption of risk aversion is a concave utility function over more wealth. Increases in wealth increase our utility (aka happiness, or satisfaction) though at lower rates the higher it goes. Concavity of utility is a necessary and sufficient condition for risk aversion, and thus, risk premiums: risk should beget an additional return. This basic assumption is challenged

by the well documented finding that while in any given society the rich are, on average, happier than the poor, the trend in individual happiness is more or less flat once society passes a threshold of perhaps \$20,000.

Richard Easterlin's original work in this area, published in 1974, discovered what is now known as the Easterlin Paradox: as economic wealth doubled after World War II in the United States, surveyed happiness was the same.¹³ Initially, this reality was thought to apply only to developed countries, but recent research suggests that the paradox holds true for a large sample of much poorer countries as well. In both cases, wealthier people are, on average, happier than poor people are *within* the same country. *Across* countries or over time, however, there is a very weak correlation—if any—between average income levels and happiness.

Happiness data consists primarily of responses to a survey in which people are asked "In general, how happy would you say that you are—very happy, fairly happy, or not very happy?" Using self-reports to measure happiness immediately raises the question of comparability. But generally most people's concerns are about making a living and matters of family life, and these concerns seem about as pressing as they have always been. In a detailed analysis of data from 45 surveys covering data from 1950 through 1977 in the United States, happiness is basically unchanged over time. Local area surveys yield similar results; for example, a study of the Detroit area reports that there was no change in satisfaction among Detroit-area wives between 1955 and 1971 even though median family income rose 40 percent.¹⁴

Trends in life satisfaction in nine European countries from 1973 to 1989 are much like that for happiness in the United States. Satisfaction drifts upward in some countries and downward in others. The overall pattern shows little or no trend in a period when real GDP per capita rose between 25 and 50 percent.¹⁵

Japan provides the most stark data point. Between 1958 and 1987 real per capita income in Japan multiplied fivefold. Washing machines, refrigerators, and television sets went from being extremely rare to universal, and car ownership soared from 1 to 60 percent of the people. And just like in the West, there was no improvement in average subjective well-being. The utility function implied by this fact is clearly not increasing in wealth.¹⁶

But that is merely the past 50 years. Utility functions are even more absurd in light of what has happened since Adam Smith created economics as a discipline.

The increase in wealth in the past century raises some very fundamental question about assuming that utility is a function of absolute wealth. In the nineteenth century, life was comparatively nasty, brutish, and short. People existed on about 1,000 fewer calories per day than today in the West, and

so your average person was a runt: about 10 centimeters shorter than today in Western countries, and about 50 percent lighter in weight. Malnutrition was a fact of life, as deficiencies in niacin, thiamine, vitamin D, which were common and not understood, led to higher rates of chronic conditions that not only increased mortality but affected the quality of life. People had higher rates of diarrhea, heart disease, and circulatory problems, among other inconveniences.¹⁷

In 1880, the average worker worked 64 hours a week, with seven holidays. In 1995, the average workweek was about 37 hours a week, with 28 vacation days. Real GDP per capita was \$5,000 in the United States in 1900, and about seven times higher, \$37,000, in 2007. But this actually understates our wealth increase, because things like computers, the Internet, videos, and e-mail are all available to us now, and would have been simply inconceivable to our great-great grandparents. Clean water, available publicly everywhere, and ubiquitous air conditioning, are things we take for granted. We are all very rich in the context of human history, yet while reading a nineteenth-century novel, I feel like people then had the same concerns, anxieties, and happiness that I feel now. I do not feel like a very rich man reading about impoverished people, which, in a pure logic of my material comfort and options available, I objectively am.

Furthermore, it should be noted that while inequality persists, today's poor have benefited far more than the rich. The typical home owned by the poor is a three-bedroom house with one-and-a-half baths. Some 73 percent of poor households own a car or truck; nearly a third own two or more cars or trucks. Eighty percent have air conditioning; by contrast, in 1970, only 36 percent of the general U.S. population had air conditioning. Nearly 9 in 10 poor households own microwave ovens; more than a third have automatic dishwashers. Back in Napoleonic times, the poor were about five inches shorter than those in the leisure class—today the heights are almost equal.¹⁸

Poor households are well equipped with modern entertainment technology. Nearly all poor households have color TVs, but more than half actually own two or more color televisions. One-quarter own large-screen televisions, 78 percent have a VCR or DVD player, and almost two-thirds have cable or satellite TV reception. Needless to say, even the rich of 100 years ago did not have these conveniences. The fact that obesity is more common among the poor basically tells the story: Poverty is not what it used to be.

We always think we are a 20 percent increase in wealth away from true happiness, but that is because we define happiness relative to a moving target. Like our definition of what is warm is a function of what we are used to—Minnesotans would find 40 degrees warm in January while in

Florida it would be considered cold—happiness is generally a function of what we are used to, though, at the extremes this clearly breaks down (just as -20 degrees is considered cold to both Minnesotans and Floridians). Thus the paradox of flat happiness in the face of rising wealth is explained as ever increasing aspirations. The idea that there is a paradox inherent in the drive for affluence is one of the key arguments for contemporary skeptics on economic growth. Many of the most influential books on the topic even have the word in their titles. In the past decade there has been Gregg Easterbrook's *The Progress Paradox*, David Myers's *The American Paradox*, and Barry Schwartz's *The Paradox of Choice*. In Richard Layard's *Happiness: Lessons from a New Science*, the opening paragraph states, "There is a paradox at the heart of our lives. Most people want more income and strive for it. Yet as Western societies have got richer, their people have become no happier."

A measurement issue is that, if you look cross-sectionally at status and happiness, those with higher-than-average status will be happier than those with lower-than-average status. But there is a bias to this reporting, because at any one time, those with higher-than-average status will have recently experienced positive shocks to their status, and so the short-run effect can be conflated with the relative position, so the empirical issue is whether happiness is more affected by recent gains in wealth (habituation) or the mere relative status.

Evidence in favor of the relative-to-consensus theory comes from several sources, and is considered to be greater than the habituation. A novel paper dealing with social comparisons is Knight, Song and Gunatilaka (2008).¹⁹ This paper appeals to cross-sectional information on 9,200 households in China, and thus refers to an economy that is very different from the Europe–North America nexus that has dominated the literature. The authors are not only able to identify which villages their respondents came from, but also confirm that 70 percent of individuals indeed saw their village as their reference group by simply asking them to whom they compare themselves, making their rural sample well-suited to the question of how important reference groups really are. Controlling for their own income and for village income, those respondents who said that their income was much higher than the village average reported much higher levels of happiness than those who said their income was much below the village average. Relative income was the most important right-hand-side variable.

Another result comes from surveys of hypothetical questions. *Homo economicus* should care about his absolute income independent of others. But given a choice between two hypothetical worlds, one in which he earns \$100,000 a year in perpetuity while others earned \$90,000, and another world in which he would earn \$110,000 while others earned \$200,000, almost everyone prefers the world which is, in aggregate, poorer, because

they would be relatively richer.²⁰ Or to take a more tangible example, as anyone with children can attest, one sure way to make Son One throw a fit, is to let Son Two watch TV and not let Son One. If you do not let either watch TV, there will be much less unhappiness, but kids are hypersensitive to their relative privileges.

In general, researchers find that relative income matters more than absolute income.²¹ That is, quantitatively, changes in relative income have much larger effects on happiness than do changes in absolute income.

SUMMARY

It is not merely the risk-return implications of the risk-return theory that are violated, but other implications of the utility function that generates this result. This implication that people should merely buy the market in various degrees of leverage has always been known to be wrong as a description of the typical investor, who often buys highly undiversified portfolios based on reading about company fundamentals and cash flows, a futile gesture that is costly, time consuming, and generates excessive volatility. The implication that people's happiness is a concave positive function of wealth is not merely a prediction, but an assumption of this work and it too is empirically violated. This is often waived away by the notice that even though there is a lot of irrational activity, the market is pretty efficient, and so, the behavioral anomalies cancel out with respect to pricing (and thus returns).

Regardless of the market anomalies presented in the data, it is also clear that it is not easy to become rich in the market, and so on that basis, the rational expectations assumption thrives. The failed assumptions in the risk premium framework have no compensating empirical successes. To fix the theory of risk, one needs to fix the utility function.