CHAPTER 14

Conclusion

Science is a great many things, . . . but in the end they all return to this: science is the acceptance of what works and the rejection of what does not. That needs more courage than we might think.

Jacob Bronkowski

There have been many cases where success seems just around the corner, but *almost* is just as close in science as in a mathematical proof. Any field is populated by thousands of really smart individuals who work very hard, and would love to solve an outstanding puzzle. The easiest way to get published, get a job, is to empirically validate a popular theory—everyone in the field welcomes it. Thus, any 40-year-old unsolved puzzle is not the result of insufficient time like how long it took to find the first million digits of Pi; it is the result of an insoluble problem. Think about early thoughts that we were close to finding a cure for war, cancer, or poverty.

In 1942, Paul A. Samuelson laid down the gauntlet and showed economists the way it was going to be in his seminal *Foundations of Economic Analysis*. Unlike psychology or sociology, economists would use a common mathematical structure underlying multiple branches of economics from two basic principles: optimizing behavior of agents and stability of equilibrium as to economic systems. The emphasis on stability was quickly forgotten, but the optimizing and equilibrium is exactly what the founders of Modern Portfolio Theory were doing, following Samuelson's playbook precisely in terms of how to do good economic science. The problem is, their utility functions were wrong in a very fundamental way, emphasizing absolute wealth over relative wealth, or, more prosaically, ignoring that all professionals benchmark their returns, but by the time this became apparent it was too late for economics or finance to merely adjust theory to the data because there was too much path dependence. Thus, we have risk aversion, risk neutrality, and risk loving, all highly mathematically elaborate as having all sorts of consistent insights, but when used

selectively, depending on the data presented, all the theory on which it is built is rather pointless. The theory is a framework, and frameworks are lousy theories. This is why economists are not becoming like dentists fixing cavities as Keynes envisaged, but rather, like someone studying postmodern literary theory, explaining everything, predicting nothing. There are very few issues in economics that have been accepted as fact merely by theory, as opposed to their obvious empirical reality, suggesting that the wonderful, powerful, maximization principles expounded in Samuelson's *Foundations* invariably are merely rationalizations. The original expositions of risk are now seen as incredibly naïve, because volatility of stock wealth is clearly not helpful in explaining risk premiums. The resulting extension is useful within academic finance, and helpful at keeping those who aren't economists out of the conversation, but hardly useful at what theory should do.

Again and again, the bright thought has occurred,

> *If we can only define our terms, if we can only find the basic unit, if we can spot the right "indicators," we can then measure and reason flawlessly, we shall have created one more science.*
>
> Jacques Barzun

It is tempting to think that the essence of risk is an objective number that, combined with our risk aversion preferences, tells us what to do. That these are both so mathematical and precise gives us the illusion of science and progress, like finding the next element on the periodic table. Yet the malleability of both risk factors and preferences means that the conventional framework is something ignored at no cost. It is customary to note that risk is extremely difficult to quantify in practice, but rarely is it acknowledged to be a highly misleading way to approach a successful career in finance. It is clearly self-serving for economists to presume that their comparative advantage at looking at investments—portfolio mathematics, covariance factors—is the essence of finance, yet one must remember that most investors, even Warren Buffett, do not use equations to make investment decisions and use simple unvarying discount rates. The standard theory has some useful insights and tools for finding alpha, but it is neither sufficient nor necessary, which is why most successful finance professionals do not spend much time calculating covariances.

No group of full-time hairsplitters appreciates being told they are splitting hairs. They will roll their eyes, say, "You just don't get it," and take comfort in continued conversation among themselves. But the facts don't

change, so if you care about the future of your ideas it's good to have them on your side. No one wants to be like a Marxist professor, retiring in 1989 after decades extolling the superiority of East Germany, finding the data now too overwhelming to support his primary intellectual argument. Intellectually, that has to be the most devastating realization possible, which is why people like Noam Chomsky simply reinterpret their earlier position so that they were right all along. Science is done by actual people, after all.

The explosion of complexity in asset pricing in the face of more data is symptomatic of a bad theory. More and better data generally increase confidence in good theories as opposed to making them incredibly more convoluted, and refinements to good theories are merely made to capture second and third order discrepancies, such as in physics when things approach the speed of light. The Arbitrage Pricing Theory and the Stochastic Discount Factor approaches are not extending a theory to fit selective anomalies in the data; they are trying to generate a model that looks good from 30,000 feet. Currently, only someone on the moon thinks that the standard theories work well as an approximation, mainly because the general idea that risk, properly defined, is unpleasant, and so should require compensation, just seems to make sense.

People do not approach finance assuming assets have merely some objective mean and covariance numbers that ignores their search for alpha. To the extent assets are consistent with alpha, traditional risk metrics are generally negatively correlated with expected returns, as investors try to find their comparative advantage in the quickest and most productive way, by taking a big swing. Assets and ideas have alpha, but it takes negotiating skills to capture them, and highly parochial knowledge of a product space to separate the illusory alpha from that which is feasible.

If there were only one thing wrong with our current theory, it would probably have been fixed long ago, yet when there are multiple errors, any one solution is incomplete, and therefore just as wrong, so the existing paradigm maintains its position. Remember, the heliocentric alternative to the geocentric picture of the solar system only became better in an empirical sense when you combined heliocentrism with elliptical orbits and the conservation of angular momentum; just moving the sun to the center of the universe, but leaving the other laws the same, necessitated about the same number of epicycles. The transition needed three things to be an empirical improvement, which is why a long time between Ptolemy and Kepler was needed.

As the nature of argument in the academy is to present articles that address one thing at a time—one asset class, one general adjustment to a

utility function—the problems with the existing financial theory have been stuck in a local maximum through a local hill-climbing algorithm that takes too many small steps.

The evidence for the absence of a risk-return relation on average comes from looking at the scope of evidence, and seeing four basic facts:

1. For most assets, the rate of return is unrelated to its volatility, or correlation with various prominent time series.
2. Really safe assets such as short-term T-bills have lower than average return.
3. Really volatile assets such as out-of-the-money options, long shots at racetracks, or high beta stocks, have lower-than-average returns.
4. People trade too much and diversify too little.

The explanation of these disparate observations is not one uber-theory (a framework), but rather, the following somewhat disparate set of three ideas:

1. People are more envious than greedy.
2. People take some risks in all aspects of their life to find their alpha, based on a lot of unfounded hope.
3. People save some money in supersafe assets so that wealth cannot go to zero in financial panics when correlations and volatilities seem to go to extremes.

Now, a framework that fits these three ideas into one general utility function, is not helpful, but the standards of argumentation in academics rather demands that an innovation involving three steps be mushed into one, and a new empirical finding be presented one at a time. Thus, you often see papers like "A Unified Theory of Ten Financial Puzzles,"(answer: tail risk) as opposed to a paper that argued "Three new assumptions that explain four financial puzzles."[1] Isn't the former obviously better? Simplicity and elegance are good things, but we must remember Einstein's dictum to make things as simple as possible, not more so. This simplicity bias necessarily prevents one from seeing the problem, or a solution, because the empirical problem is created by patterns in vastly different areas that have different signs (volatility bad in T-bills, irrelevant in T-bonds, good in stocks); and the theory as to why people take risks based on hope, is quite different from why they benchmark with their peers and not zero or last year. As someone used to kludgy but useful models in default forecasting, I think an observation, and an explanation with a handful of parts is neither overfit nor in need of simplification.

The importance of a relative utility function leads to benchmarking against the consensus in investments, and leads to a symmetry in risk taking—too much or too little relative to the consensus—that leads to a zero risk-return trade-off. We are not so much greedy as power hungry, and power is about our relative wealth in our tribe. The first-order effect of any gain or loss I experience affects my absolute and relative position similarly, so standard implications are maintained: I dislike random volatility, risks from accidents, I like wealth, and so on. Yet in some areas, such as investments or policy preferences, the relativity of the decision is important in explaining behavior. People will never totally embrace selfish individualism because we have an instinct for egalitarianism in that we hate to be dominated, and the domination of one invites the scorn of many that one needs if one works in a large organization. Yet envy, like greed, has a societal silver lining, because if you are indifferent to the envy of others, you are indifferent to a lot of suffering that without context is mere whining.

The grand scope of the risk-return failure in all of the risk proxies, in spite of a massive search for 50 years, suggests a large amount of delusional hope among academics not unlike that of those silly investors in out-of-the-money options. The higher returns of minimum variance portfolios, portfolios with lower volatility and covariance with everything, can only be riskier in the Alice-in-Wonderland world where low volatility, lower tail risk, and low correlation with both the aggregate market and the business cycle, is risky.

The implications of the absolute risk aversion assumption embodied in the CAPM and its extensions are logical, powerful, and slightly surprising, things one wants in a theory, but it also empirically vacuous. The anomalies to the high risk–high return assumption are not exceptions to a general tendency. There is no general tendency within a variety of investments: equities, options, bonds, mutual funds, commodities, movies, lottery tickets, or horse races, among others. This irrelevance of risk to return is implied by a status-conscious investor benchmarking himself against others, and holds in both a utility and arbitrage argument. The standard assumption that relevant volatility is absolute, not relative, may be a good normative theory—that is, what you should do—but a preference for status generates a more accurate positive theory, and its assumption—caring only about relative status—is at least as plausible as assuming people care only about absolute wealth. In light of the severe underdiversification of most portfolios, the higher expected returns that drive them, and the common method of allocating assets to standard categories, this seems to be a more realistic description of the investing process in practice than the CAPM assumption that people invest on the basis of a target beta between their household wealth and the market.

In seeking alpha, we are looking for good ideas that play off our individual talents, so that when we succeed, it implies an ability for us to capture rents via our contacts or reputation. The really important part of any edge is our unique insight into an expected value, combined with efficient ability to implement it. Understanding risk, in terms of estimating the over-the-cycle losses of the venture, is essential. Understanding risk in terms of a particular discount rate is irrelevant.