SED's New Asset Allocation Logic
The Three Principal Ideas Explained via a Socratic Dialogue

November 2005

Horace W. Brock, PhD
Strategic Economic Decisions, Inc.
www.SEDinc.com

SED'S NEW ASSET ALLOCATION LOGIC

– THE THREE PRINCIPAL IDEAS EXPLAINED VIA A SOCRATIC DIALOGUE –

(November 2005)

Author's Note: In Chapter II of our September 2004 report, we spelled out a new logic for investment managers—with particular attention to portfolio theory and asset allocation. The theory involved was new in the sense that it was based upon a much more general concept of "equilibrium" that is embedded in standard Modern Portfolio Theory (MPT) —Markowitz optimization in particular. Even so, our theory did not reject classical MPT, but rather generalized it. Indeed, the classical theory remains within the new theory as a limiting special case (both conceptually and mathematically), as should be the case according to the Correspondence Principle of the philosophy of science (i.e., the "nothing is really true" axiom of good science).

The three most fundamental conclusions of this ambitious chapter were that:

• The concept of the efficient frontier is invaluable. However, Markowitz identified the wrong frontier in his theory. For there exists a *superior* frontier whereby one can achieve greater return without greater risk than is possible for every point along the classical Markowitz frontier. This is achieved by replacing the concept of optimal portfolios with that of optimal strategies. The distinction is absolutely fundamental.

• The traditional tension between so-called strategic and tactical asset allocation ceases to exist once it is understood that truly strategic management *must* also be tactically optimal. This is because the new theory focuses on optimal strategies rather than optimal portfolios, and an optimal strategy is "contingency-dependent" or tactical by construction. As a result, tactical and strategic management are fused at the hip.

• The concept of passive versus active management needs to be completely rethought, as does the concept of benchmarking. Traditional thinking here is wrong, and has led to an endless proliferation of inconsistent concepts of benchmarking, as well as to invalid extensions of the classical concepts of "alpha" and "beta".

Response: The *good* news was that many (most) readers sympathized with my assertion that the classical MPT paradigm is in meltdown mode, and welcomed an alternative perspective. This was especially true since the new theory justifies at a deep theoretical level what many eminent investment managers have been doing and saying and feeling deep down for years. The *bad* news is that the subject is inherently difficult. Thus, despite my efforts to write a precise, clear and relatively non-mathematical exposition of the new theory, many readers found some of the arguments hard to follow. Some kindly communicated questions that I promised to address in future reports.

This is the first such essay, and is written as a Q&A Socratic Dialogue. Its goal is simply to restate the principal insights in a *very* user-friendly manner. Having been on the road talking with clients since last September, I believe I have now learned how to express the main ideas in a more intuitively appealing manner, and am thus better able to answer some of the questions put to me than was true a few months ago. Questions not addressed in this chapter will be dealt with in future essays. This topic will *not* go away given the advent of paradigm meltdown. – *H. W. Brock*

QUESTION: When all is said and done, what are the most important and relevant concepts in the repertoire of investment theory? That is, what should *every* investor be most aware of?

ANSWER: There are two, and they are fundamentally linked for reasons unearthed by Harry Markowitz in 1952. *First,* there is the crucial importance of diversification—diversification best achieved by holding a "portfolio" of different assets that are somewhat uncorrelated. *Second,* there is the concept of "portfolio efficiency": A portfolio is said to be efficient if its expected return is as high as it can be *given* its riskiness. That is, if a portfolio is efficient, than there can be no other portfolio which offers more expected return for the given level of risk involved.

Q: What you endorse is pure Markowitz. This being so, why was a thoroughgoing reworking of classical theory necessary?

A: Because Markowitz introduced these concepts in a very specialized context: one where the stochastic process governing the returns of the various asset classes is i.i.d. (returns are "identically and independently distributed"). If this is the case, then wealth is a random walk, markets have no patterns or "memory", and there is no mean reversion. In this highly restrictive context—*and only in this context*—Markowitz identified the efficient frontier and demonstrated how to construct a portfolio lying on it. Yet "i.i.d." is a very strong assumption. Among other problems it trivializes the *dynamic* aspects of portfolio theory and asset allocation. Unfortunately, this is one reason why most real-world so-called dynamic investment strategies are fundamentally ill conceived.

Q: What's the most intuitive way to understand this special i.i.d. case? What is the essential idea here?

A: The coin flip process is the classical example here. On any given flip of the coin, the probability of heads/tails is always 50/50 for an unbiased coin *regardless* of the proportion of heads/tails observed up until that flip. This is a way of saying that there is no discernable pattern in the outcomes of the process over time. To put it in terms of probability theory, the *unconditional* probability of the next outcome is always the same as the *conditional* probability conditioned on the past **k** outcomes, *whatever these were.* Historical observations thus provide "no forecasting information". Equivalently, we say that such processes possess "no memory". Equivalently, we say that such processes possess "no stochastic volatility" in that the structure of underlying risks never changes.

Note that most processes familiar from daily life—like rainfall, temperature, stock market returns, aging, lovemaking, and business cycles—evince notable patterns and thus stochastic volatility. The probability of "what's next" (e.g., a cold day) *does* indeed depend upon "where you come from" (the past week's temperature). That is, most processes do possess memory.

Why Classical Portfolio Theory Was Necessarily "Static"

Q: Why do you say that the random walk assumption trivializes the dynamic aspect of portfolio theory?

A: Because an i.i.d. assumption mathematically implies that the risk/return structure of the asset markets (e.g., the means, variances, and covariances of the assets) *never changes.* In recent years, this restrictive assumption has become known as the "no stochastic volatility" assumption. Yet stochastic volatility of many kinds is known to exist, e.g., market cycles and GARCH phenomena (time-varying variances). Now from an investment strategy standpoint, if the structure of risks and returns never changes across time, then why would your optimal portfolio ever change—assuming that your risk attitude is constant? It will not. And that is why, in Chapter 11 of his 1959 monograph, Markowitz ends up recommending that the investor hold the same efficient portfolio *always.* [This is the chapter in which he extends his theory to the multi-period dynamic case, an opts for the i.i.d. assumption.]

Q: Was this the origin of the so-called "policy portfolio"?

A: Yes. And it has stuck and dominates investment theory to this day, even though adhering to a fixed portfolio can be shown to lead to very poor performance over long periods (e.g., long bear-market cycles).

The New Theory – Three Main Insights

Q: So your new theory accepts Markowitz's concept of the efficient frontier, but generalizes it to much more general non-random walk situations. Is that it?

A: Yes. We set out to show what it means to construct an "efficient portfolio" when the environment is not i.i.d. and stochastic volatility *is* present. And in doing so, we focused on the dynamic portfolio optimization process since this is what we all contend with in daily life.

Q: What were your most important discoveries?

A: Three insights that, if implemented, would fundamentally change the theory and practice of investment management.

(1) Efficient *Strategies* Replace Efficient Portfolios

First, when you drop i.i.d., everything changes much more than I had realized before. For example, the concept of the efficient frontier must be fundamentally modified. Once it is, then it turns out there is no such thing as an efficient portfolio. For efficiency now refers to the risk/return properties of strategies—not portfolios proper. A *strategy* is defined as a rule that specifies for an investor of a given risk attitude the best portfolio to hold in any particular "state". If there are 12 relevant states, then there will be 12 state-dependent optimal portfolios. These are denoted by the 12 dots appearing on the right of Figure 1 taken from the master essay.

Whenever an investor finds himself in a particular state, e.g., state 8, then he will always switch to the state-dependent portfolio appropriate to that state, e.g. portfolio 8. As the sequence of future states evolves over time, the investor will "loop" between these portfolios (dots) forever in the precise manner prescribed by the strategy. In doing so—*and only by doing so*—he will achieve a payoff on the true efficient frontier (this is shown below).

FIGURE 1: SOLUTIONS TO THE PORTFOLIO PROBLEM

The Optimal Strategy: Classical Versus General Case

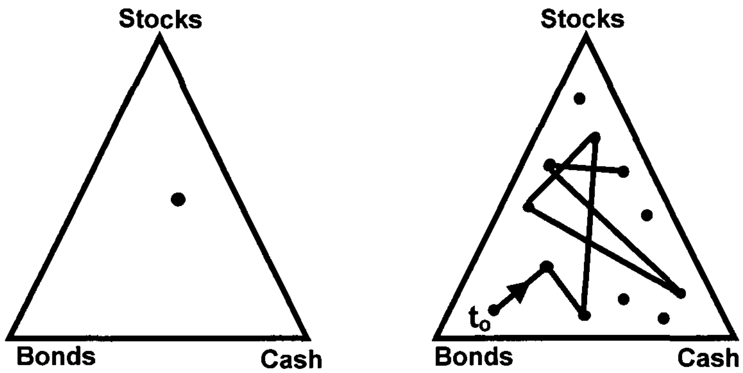

Classical Case: One Dot

This single dot represents the "Classical Policy Portfolio"

General Case: A Loop

The 12 dots taken together represent the "Optimal Policy Function"

Source: SED

Q: What exactly are these "states" you speak of? And exactly how is the investor's optimal strategy of portfolio switching determined?

A: These states refer to some set of "recent histories" that can be statistically shown to have "predictive power". For example, many scholars have shown that you can *improve* stock market forecasting by "conditioning on" the values of P/E ratios, interest rates, and earnings during the past five years. We reviewed such findings in the master essay. Any such history containing predictive value can be a "state".

Take a completely different example: In meteorology, the state could be the mean and the standard deviation of temperature during the past week. This particular history possesses great predictive power, confirming that temperature is not i.i.d. but rather exhibits seasonal patterns (cycles).

Q: And didn't you argue that it is an *empirical* matter to determine which particular states (recent histories) have maximal predictive power—and what it is they predict?

A: Yes, empirical and not ideological. Moreover, new "SNPS" econometric procedures make it possible to determine which states are most informative, as well as how many observations on them (e.g., the last five quarters—no more) are needed to arrive at an optimal state-dependent forecast. [This required time period is known as the "memory" of the stochastic process.] We reviewed this powerful new SNPS logic both in the text and in Appendix A.

As for how to determine the optimal strategy of portfolio switching for an investor with a particular risk attitude (i.e., how to determine the specific dots within the triangle), well-known techniques of dynamic programming can be utilized, although notable short-cuts now exist to facilitate this optimization.[1] We also discussed this matter at length.

Q: You have suggested that there is an analogy here to a farmer's decision as to how to optimally rotate crops (equivalently, switch portfolios) seasonally?

A: Yes. Seasons are indeed analogous to market patterns and cycles. Depending upon a particular farmer's risk attitude and other circumstances, dynamic programming will solve the problem of optimal seasonal crop rotation given the various risks involved. You will of course protest: "But seasons unlike markets have an unambiguous pattern—there is no so-called stochastic volatility."

But you are wrong, and should ask yourself: "If there is no unpredictability to the cycle of the seasons, why then does the *Farmer's Almanac* always outsell any other annually updated book published?" Farmers bet on its predictions of the first snowfall, and of many other seasonal "firsts" and "lasts" as well! And come to think of it, hurricane season is very unpredictable too, with only one storm this year, four last year, two the year before, etc. Then, there is global warming and cooling—a notoriously unstable process. The point is that stochastic volatility is everywhere around us. Markets may seem very different, but they are not.

Q: How controversial are the statements you have just made?

A: On the one hand, many investors find these assertions a bit hard to apprehend since they have been inculcated in the Markowitz paradigm of invariant risks and returns, and of the primacy of portfolios as opposed to strategies. On the other hand, investors intuitively appreciate the need for "flexibility" and are thus sympathetic to the new ideas of variable risk/return structures and of strategies.

Yet regardless of how people think, the propositions stated thus far are *theorems* that will are necessarily true except in the special i.i.d. case. And only in this limiting case will holding one fixed "policy portfolio" be optimal, with the payoff from doing so lying on the true efficient frontier.

---

[1] If an investor's utility function is approximated by a logarithmic function, then the optimization problem can be simplified and solved "myopically" without dynamic programming at all, as Robert Merton and Paul Samuelson first demonstrated back in 1969.

A careful reading of Markowitz's highly technical 1959 monograph makes clear that he appreciated some of these points long ago. He talks about the need for a dynamic programming solution to the portfolio problem in the "general" case where the risk/return structure changes over time (stochastic volatility). But in opting for an i.i.d. framework, he simply sided-stepped the matter.

(2) Unification of Tactical and Strategic Allocation

A second insight was the realization that the great divide between tactical and strategic asset allocation was in fact a red herring. For when judged from the standpoint of a generalized theory, no such divide can exist at all. This is because the "strategy" of portfolio switching required to generate a payoff on the true efficient frontier will always require optimal tactical adjustments by virtue of being "state dependent". [Recall the strategy "loop" of Figure 1.] In short, strategic and tactical decisions are joined at the hip and are unified by the concept of a strategy.

Q: Could you explain this a bit more? How, for example, does the ability to "tactically" manage market beta fit in here? Or to take another example, how does the related strategy of classical market timing fit here? What's the exact relationship between tactical and strategic in these contexts?

A: The blunt answer: Strategies like "managing beta" and "market-timing" do not fit in very comfortably. This is a very subtle issue, so let's proceed carefully. First, recall that "active" investment managers attempting to liberate themselves from the tyranny of the one-policy-portfolio-forever paradigm introduced these tactical strategies. They were utilized to capitalize on "deviations" from equilibrium market returns, and the optimal strategic policy portfolio that was predicated on these equilibrium returns. However, given the dominance of the policy portfolio as the centerpiece of "responsible" investment management, all strategies for deviating from it were defensively dubbed "tactical". The very terms suggests that one should never stray very far from home base.

Q: What's wrong with this view? How does it conflict with the new paradigm?

A: It conflicts fundamentally with it, and moreover it prevents investment managers from doing what they should be doing. Here is why—at two different levels.

Level 1: At the deepest level, the concept of "equilibrium returns" lying at the heart of MPT plays no role in the theory of general equilibrium that underlies our theory, namely the Theory of Rational Beliefs.[2] This new theory was developed at Stanford University in the 1990s and demonstrates that it is *incorrect* to identify long-run average returns (e.g., 7.3% on stocks) with the concept of "market equilibrium". For the market is in fact always at equilibrium—*regardless* of whether values are much higher than/lower than their statistical average. That is, all equilibria enter equally into the theory.

[2]A Rational Beliefs Equilibrium conceptually and mathematically generalizes the traditional concept of a Rational Expectations Equilibrium, the bedrock of classical MPT, CAPM, and Efficient Markets Theory. The REE remains as a special case of an RBE. We argued in the master essay why the concept of a RBE was the appropriate basis for a theory of the portfolio.

This being true, the concepts of a "deviation from equilibrium" and of "tactically exploiting such deviations" have no legitimate meaning. This realization changes everything, as we now see at Level 2.

Level 2: The theory now says: OK, market returns and valuations range all over the map—and each one is a temporary equilibrium. This being true, let us look for *patterns* in market returns and valuations generated by the sequence of equilibria over time. In particular, let us identify the *statistically average dynamical pattern of returns* (e.g., statistically average bull/bear market cycles). Finally, let us adopt an investment strategy that *exploits* these patterns, and in so doing add value over the static strategy of sticking to the fixed policy portfolio. (Remember that this policy portfolio will only be optimal in an i.i.d. world devoid of patterns.)

Q: And the "loop" strategy of Figure 1 determined by dynamic programming is just what the doctor ordered here, correct?

A: Yes. But now stop and ask yourself a critical question: Does the decision to play this optimal pattern-exploiting strategy represent an active or a passive management stance?

Q: Good question. I don't know. It smells active .....

A: You've fallen for my trap. It is completely passive. Think back: No "market calls" or "timing judgments" arise in implementing the optimal strategy. A robot can compute it from the historical data on average return patterns via dynamic programming.

Q: I can see that it is passive, and yet it achieves what only active strategies could achieve classically. For example, aren't you saying that the robot-determined optimal strategy will incorporate the equivalent of a market timing strategy, provided that econometric analysis of historical data reveals the existence of long-cycles in the data?

A: Bulls-eye. If the data reveal "memory" (e.g., if bull/bear market cycles and/or other patterns exist in the data), then the optimal policy will automatically incorporate within itself the <u>right kind</u> and <u>right amount</u> of "beta management" and "market-timing". These terms themselves need never be introduced. They become redundant and confusing. They cannot be disentangled. Everything is strictly "passive". No one need "call" anything. Finally, strategic and tactical are fundamentally fused. The optimal strategy requires optimal tactical adjustments as the return-states evolve (the loop in Figure 1).[3]

Q: Where in the world does all this leave "active" management?

A: It is radically changed. This will be the subject of the section just below.

---

[3] To recall exactly how the optimal strategy can automatically incorporate optimal market timing in a strictly passive manner, refer to the detailed explanation of market timing in Section G of the master essay. The logic here is frankly *exceedingly* subtle, and the power of dynamic programming is seen to be no less than astonishing.

Q: So to conclude, you have established that thinking schizophrenically has sidetracked investment managers. First, they are told that a sacred equilibrium exists—a duo of equilibrium returns, and an associated policy portfolio. Second, they continuously debate: "Hadn't I better stick to the policy portfolio?" versus "Isn't this the perfect time to add value by deviating from the policy portfolio? But when and by how much?" And in thinking and acting this way, investors get confused, and inevitably invest suboptimally.

A: Yes, and today's confused debate about pension fund strategies (e.g., the debate precipitated by Peter Bernstein's call for flexibility and even market-timing) is proof positive of the point we are making here. Jargon is proliferating in this debate (seven confusing definitions of "beta" now exist, four of "equilibrium returns", five of "alpha", eight of "market timing", and twenty-two of "benchmark"), and confusion abounds. This is evidenced by the fact that most of those writing neither define their terms carefully nor produce a single general theorem about the nature of a truly optimal strategy in a non-i.i.d. world. We are left with a pastiche of "approaches", and many clients have expressed to us their dismay at the status quo.

Q: So this is a classic example of paradigm breakdown?

A: Yes. Classic. We need to go back to absolute basics. A message needs to go out to the profession that the problem of selecting an optimal sequence of portfolios (a strategy) *is* fundamentally a decision tree problem—no more or less. And as many MBAs learn in school, the correct way to "solve" such a tree for the desired strategy is to "roll the tree backwards from right-to-left" by utilizing the principle of Backwards Induction that lies at the heart of decision theory and in turn of dynamic programming. This truly should be the starting point for teaching dynamic portfolio theory.

Q: If much of this is not new, and stems from the work of such scholars as Samuelson and Merton, then why hasn't the CFA incorporated it in their curriculum?

A: Ah! Some five people have asked that very question. Perhaps there is a bureaucratic vested interest in clinging to the classical paradigm. But I am frankly at a loss to understand their inertia, especially given the salience of *Strategic Asset Allocation,* the much acclaimed 2002 book by John Campbell and Luis Viceira that gets this part of the story dead right, and thus exposes many short-comings of the CFA curriculum. But hopefully things will soon change for the better.

(3) Complete Reformulation of the Concept of "Active" versus "Passive" Management – and Benchmarking

The foregoing results led to a realization that the "active" versus "passive" distinction must be fundamentally rethought, along with the all-important concept of "benchmarking" performance. This part of our argument was novel since it drew on very recent work in general equilibrium theory that is not yet known within the finance community. It is also the most significant part of our theory since it turns most conventional wisdom about performance and how to achieve it on its head. The message to the pension consulting community is: "Go back to the drawing board, rethink your measures of performance, and work with clients to develop wholly new products that properly exploit patterns of returns. The opportunity to do so is great, and clients need help."

Q: In explaining your new theory of active/passive management, could you do so diagrammatically to make things a bit easier?

A Diagrammatic Explanation of these Basic Ideas

A: Yes. Consider Figure 2—a unifying diagram we should have introduced in the master essay, but didn't. The vertical axis represents "payoff" and the horizontal axis "risk".

FIGURE 2: CLASSICAL vs. TRUE EFFICIENT FRONTIER

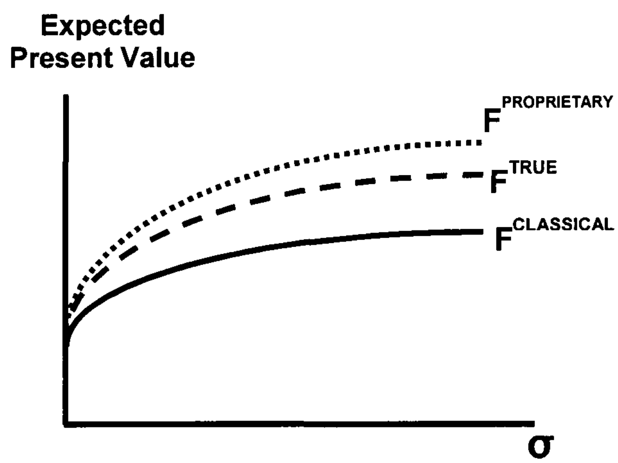

Source: SED

Q: Why is the vertical axis labeled "Expected Present Value"?

A: In the case of *dynamic* portfolio theory, the payoff from any strategy is represented as the (time-discounted) present value today of the sequence of payoffs resulting from pursuing that strategy. [Remember that a strategy is a rule specifying how to "loop" around the strategy simplex shown in Figure 1.] Moreover, just as any strategy will have an expected present value, so will it have an associated level of risk, and this is denoted on the horizontal axis.

Q: Exactly what does each of the three frontiers represent in the figure?

• **Bottom Arc:** The bottom frontier represents the entire set of payoffs achievable from investing in classical Markowitz policy portfolios. The point at the far left of this bottom arc corresponds to the "no risk/lowest expected return" payoff from holding a 100% cash portfolio over time. [Note that this is true for all three cases.] The payoff at the other extreme of the arc corresponds to holding 100% equities in perpetuity. And so forth. An investor's risk attitude will determine the right (utility maximizing) point on the frontier and hence the right portfolio for him.

Note that each point on the frontier is generated by an (efficient) *portfolio*, never by a *strategy*. Finally, observe that if markets in the real world were i.i.d. and devoid of any patterns, then this lowest arc *would* be the "true" frontier in the sense that an investor could not achieve a payoff anywhere above this arc.

• **Middle Arc:** The middle arc represents the non-i.i.d. real-world true frontier. Here it is possible to achieve a payoff above the classical frontier—*well above it in fact*—by pursuing an optimal strategy of portfolio switching as defined above. We have already noted why this is true:

> The strategy determined by solving the underlying dynamic programming problem will optimally exploit the dynamical patterns and "memory" of the risk/return structure unearthed by appropriate econometric analysis. This is true no matter what the nature of the pattern may be, e.g. long cycles, or intra-day trading patterns. Exploiting such patterns can increase expected value *without* increasing risk.

Conversely, this extra value (increase in efficiency) will be lost when such patterns are ignored as they are in a mean/variance framework that *assumes* no patterns even if they exist and when the investor sticks to his Markowitz policy portfolio forever.

Q: So to summarize, any point on the true frontier corresponds to the payoff of sticking to an optimal strategy forever. And given an investor's particular risk profile, there will always be an optimal strategy —one that puts him at the right point on the true frontier. Correct?

A: Yes. And the more risk-prone the investor is, the more his optimal strategy will place him in the upper right-hand quadrant (where there is more risk and more expected return). It is very similar to the Markowitz set-up.

Q: You state that this true frontier is "well" above the classical frontier. Why?

A: Because in many contexts—but not all—there is enough of a pattern in the data and thus enough "predictability" for a state-dependent strategy to yield a payoff dominating that of any fixed portfolio, as has been documented by Robert Shiller and many others in the empirical work cited in our master essay. It will be helpful to consider once again an agricultural analogy: How profitable is the farmer who acts as if the *seasonal* mean temperature is always the same as the *annual* mean temperature—and plants his crops accordingly? The answer: He will be broke *unless* he lives on the equator where the conditional mean is the unconditional mean (the process is i.i.d.) since there are no seasons. No seasonal crop rotation is needed.

Dynamic Passive versus Active Management – The Essential Difference

Q: In your remarks above, you suggested that a robot in principle could optimize for a client with a specific risk tolerance. The resulting optimal strategy will have a risk/return payoff that is on the true efficient frontier, correct? Is this robot equivalent to a passive manager in the new theory?

A: Yes. The situation is analogous to that of the classical paradigm. The passive manager starts off by crunching the historical data to learn the true risk/return structure describing the average dynamics of risk and return over the long run (e.g., the average behavior of all market cycles). Next, he explicitly accepts these "objective" econometric results as representing the best possible forecast of the future. That is, he believes that the average dynamical pattern of past returns will replicate itself in the future, or equivalently that he cannot arrive at a better subjective forecast on his own. *The decision to believe this is the hallmark of passive management.* Finally, he optimizes with respect to this econometric forecast and the client's risk posture. The result is the optimal passive strategy—the strategy whose payoff will be the right point for that client on the true efficient frontier.

You will better understand this characterization of dynamic passive management when it has been contrasted with the a concept of dynamic active management, and the associated concept of the "Proprietary Efficient Frontier"—the top-most arc in Figure 3 —to which we now turn.

Q: Stop. Before discussing the top-most arc, could you explain something? You keep referring to "the historical risk/return structure" as revealed by econometric analysis of the data. What exactly is this? Is this the matrix $\mathbf{R}^{\mathbf{M}}$ from your master essay?

A: Yes, it is, and it is so important to understand this that we have reproduced the matrix with a lengthy description of what it means in Figure 3. Recall that each row represents a so-called conditional joint probability distribution of returns. The decimals in each row give the probability of achieving the four different stock-and-bond return-states in the "next period" *conditional* on what has happened "in the recent past". These numbers are called transition probabilities, and should be read left-to-right. Intuitively, the matrix as a whole captures what history can tell us about the likelihood of "what's next" given "recent events".

Note how the row probabilities change as we move from "history" to "history". The existence of such changes is what we have called "stochastic volatility". In a classic i.i.d. "no memory" Markowitz world, all row vectors will be identical, *and* will equal the unconditional probability distribution **R** seen up top in the graph. Classical theory assumes this very special case. And the static means, variances, and covariances of classical theory are deduced from the data in **R** alone.

From the investor's standpoint, what is important is that this matrix $\mathbf{R}^{\mathbf{M}}$ is that it contains *everything* that can be learned from historical data about the dynamics and the "memory" of asset returns. And it is these data that get crunched to construct the true efficient frontier.

Q: Didn't you stress that the best way to think of this matrix is to think of it as capturing the *average dynamics* of the asset return structure?

A: Yes. The transition probabilities capture the *average behavior of market cycles,* e.g., the average duration and amplitude of equity market cycles. We stress "average behavior" here because the behavior of such cycles varies over time, just as hurricane seasons do. All we can infer statistically is the average or mean behavior of such cycles, the variance of such behavior, and the covariance of such behavior across asset classes.

Q: It is far from obvious how a simple matrix of this kind can contain so much information.

A: Yes. Markov chains of this kind are incredibly powerful. In practice, information of the kind we are discussing might not appear in the discrete matrix form we have chosen (on the grounds of simplicity). An academic will simply say: "Let's estimate the joint conditional distribution *and* its memory using SNPS tests." And the variables involved will often be continuous, not discrete.

Now let's turn to dynamic *active* management, and to the meaning of the proprietary efficient frontier.

• **Top Arc:** This top-most entity is a wholly "subjective" version of the objectively true frontier in the middle. It will differ for every manager who possesses a proprietary forecast of the future risk/return structure that is *different* from the objective forecast $\mathbf{R}^{M}$ revealed by econometric analysis. In our original essay, this subjective forecast was denoted by the matrix $\mathbf{R}^{M^*}$, and it was interpreted as a "subjectively edited" version of $\mathbf{R}^{M}$. [The probabilities will differ from those in $\mathbf{R}^{M}$.] Any manager who is confident in his ability to improve the forecast embedded in $\mathbf{R}^{M}$ will be called an "active" manager, and the top arc will be the proprietary frontier that is derived from his subjective forecast $\mathbf{R}^{M^*}$.

The gap between the top two arcs represents the extra efficiency (value) that the active manager *believes* he can achieve via his proprietary forecast of the future risk/return structure $\mathbf{R}^{M^*}$.

Q: How does an investment manager earn this extra value? What kinds of strategies can he use, and how do these relate to traditional modes of adding value by "managing beta" and "adding alpha" via security selection? Do these concepts carry over to the new theory?

A: As I answer this question, you will see that the concepts of managing beta and adding alpha become quite muddled. Indeed, they are not really needed whereas they were indispensable in classical MPT. We stressed this point earlier in the passive context, and you will now appreciate it in an active context.

Assume that the active manager starts off with a complete history of returns for a particular set of asset classes or markets. These are the markets in which he is confident of being able to add value in one way or another. There are then four generic steps he should follow in attempting to do so:

First, he will utilize econometrics to construct the historically-based risk/return structure $\mathbf{R}^{M}$ needed for constructing the "true" frontier in Figure 2. This provides him with an objective starting point as to the average dynamics of the risks and returns in his chosen markets.

Second, he will edit this structure to reflect his own beliefs as to the risk/return structure that he believes best represents the future dynamics of asset risks and returns, and the result will be $\mathbf{R}^{M^*}$. This forecast will incorporate views of the following kind: "In our view, the duration of the bull market in the US will be longer than its historical average this time around, and will peak at a higher-than-average P/E ratio. Our reasons are........Conversely, we believe that the Russian cycle will be shorter and less pronounced than usual. With these data in hand, his proprietary frontier can be constructed from $\mathbf{R}^{M^*}$.

Third, he will optimize over $\mathbf{R}^{M^*}$ via dynamic programming, thereby determining the strategy whose payoff is the right point on the proprietary frontier given the risk attitude of the client involved. The result will be an optimal strategy dictating when and at what rate to exit/enter the various asset classes (recall Figure 1).

Fourth, as an extension of this third step, or an alternative to it, he can obtain subjective expertise on the future behavior of component securities *within* any of the different classes. He can then substitute these securities for the market indices utilized in steps two and three above. He will then reoptimize utilizing this security-specific data rather than market-specific data. Alternatively, he could use proprietary sector-specific information, and treat the set of sectors like the set of asset classes, and reoptimize. Finally, he could "overlay" his security-specific and sector-specific expertise on top of his market-specific expertise using a derivatives strategy, and then reoptimize. Note that for each of these efforts, optimizations will be required that *differ* from the bread-and-butter asset class optimization of step three just above.

Q: Let me be sure I understand the way in which beta management is fitting in here. You already explained that a *passive* manager will "automatically" manage beta risk to the extent that the optimal strategy given $\mathbf{R}^{M}$ will incorporate optimal market timing, etc. Indeed, you said that a robot could be programmed to optimize here since no proprietary forecast is involved.

But in the case of the *active* manager, beta is being managed in a way a robot could not mimic. For the strategy by which it is managed will derive from optimization with respect to the manager's proprietary forecast $\mathbf{R}^{M^*}$. In sum, the passive manager manages beta risk passively, whereas the active manager manages it actively. Correct?

A: Yes. Both types of managers *must* manage what is called beta risk if they are to arrive at an optimal strategy for their clients, assuming a non-i.i.d. environment. This is because the failure to exploit stochastic volatility via appropriate market timing strategies will always generate an inefficient payoff *not* lying on the efficient frontier—the true efficient frontier in the passive case, and the proprietary frontier in the active case.

Q: But you are also permitting the active manager to supplement his active management of beta risk with alpha-like strategies of security selection and sector selection.

A: Yes. It is a different type of expertise, but is completely valid. But note that whereas there was a strict passive analogy to active beta management (the robot does it utilizing $\mathbf{R}^{M}$), there is no such passive analogy to active security selection. For there will be no historically-based risk/return structure relevant to the security selection that a robot could use. Security selection is inherently subjective from the start. Interestingly, sector selection strategies may represent a half-way house here: Risk/return structures of sectoral returns estimated from historical data could perhaps be exploited passively via a robot to arrive at an optimal rotation strategy.

FIGURE 3: THE RISK / RETURN STRUCTURE

DATA →

R =

$S^H B^H$	$S^H B^L$	$S^L B^H$	$S^L B^L$
.11	.41	.38	.10

R^M =

	$S^H B^H$	$S^H B^L$	$S^L B^H$	$S^L B^L$
History 1	.2	.3	.4	.1
History 2	.1	.6	.2	.1
History 3	.1	.2	.6	.1
History 4	.1	.5	.3	.1

For illustrative purposes, we have simplified matters by assuming that there are only four "return-states" (the 4 labels across the top), namely "S^H B^L" for "High Stock Returns, Low Bond Returns", etc., where High denotes above average returns, and Low denotes below average returns.

The graph shows that data on historical returns can be "mined" to yield two very different kinds of risk/return information: R, and R^M. R is the traditional "joint distribution of joint returns" which gives the relative frequency of all stock and bond returns jointly over time. Thus, the first element of R tells us that there is an 11% probability that stock and bonds returns will both be above average ("High") in any year. Importantly, the data in R can be crunched to yield the classic Markowitz Means, Variances, and Covariances of asset returns.

R^M is a much more general structure that captures the "average dynamics" or "patterns across time" of asset returns. It is called a "Markov Chain" and its rows should be read left-to-right. Each row gives the "transition probabilities" of returns "tomorrow" *given* the recent history of returns and of any other variables (like P/E ratios) whose recent values help investors to forecast "what's next". (Determining which "Histories" have the highest predictive power is a statistical matter.) We have illustrated this via four hypothetical histories.

Finally, note that R^M is a much more general structure than R. It can be crunched to yield R, but the converse is not true. R abstracts from all "dynamics" of markets, and merely contains information about *how often* each of the return states occurs. More formally, it is the static average (or "stationary measure") of R^M. For the technical reader, R^M is known as the "joint conditional distribution" of the historical data on returns.

Q: You are implicitly claiming that, in carrying out their respective missions, neither active nor passive managers need utilize the concepts of alpha and beta. Is this true?

A: Yes. These concepts arose within the linear CAPM framework where they were well-defined and had a clear meaning. That framework has collapsed, and these concepts are no longer necessary or even useful. In the new paradigm, all that matters is the investment manager's alleged expertise $\mathbf{R}^{M^*}$ *and his ability to utilize this by properly optimizing over it.* This expertise can take the form of proprietary beliefs about the dynamics of asset class returns or about individual securities or sectors.

Q: So the key to superior returns from active management lies in both superior expertise *and* superior optimization. Why do you stress both here?

A: Because the wrong optimization logic will generate an inferior and inefficient strategy *regardless* of the quality of the underlying forecast. And vice versa. This point is not at all understood in this profession, as evidenced by the plethora of static "optimizers" that have been utilized for 30 years. In real-world non-i.i.d. environments, this form of optimization yields very inefficient strategies, as everyone from Merton and Samuelson in the past to Cambell and Cochrane today recognize.

> Note how I keep returning to my initial point: Once it is realized that we must seek optimal *strategies* as opposed to optimal portfolios, everything changes. The data required are altogether different ($\mathbf{R}^M$ versus $\mathbf{R}$), *and* the logic of optimization required is far more sophisticated (dynamic versus static programming).

A Reconceptualization of "Benchmarking"

Q: My final question concerns the status of benchmarking in the new theory. What exactly is it?

A: Our new concept of benchmarking is a significant generalization of the original concept whose origins date back to James Tobin's celebrated "mutual fund separation theorem" of 1957. Tobin showed that all investors will hold only two assets, regardless of their risk attitude: **(1)** the riskless asset, and **(2)** the portfolio of all risky assets. The latter became known as the "market portfolio" and in practice has been proxied by broad market indices like the S&P 500 during the past three decades. Since "the market" was the *only* asset (other than cash) held by *all* investors of *any* degree of risk aversion, it was automatically the relevant performance benchmark for *every* active manager.

Q: Is there any problem with this classical concept?

A: On the surface, none. It's clean, and brilliantly insightful. Deeper down, yes. The assumptions generating the separation theorem imply a stationary world in which active money management could not exist in the first place. No benchmark would thus be needed, any more than it would be needed in a Lucasian world of Rational Expectations twenty years later. But let's not dwell on this fine point.

Q: How does the new theory extend this classical theory?

A: It doesn't extend it as much as reconceptualize it. The logic here is subtle. In the new theory, each passive investor will possess a unique optimal strategy—the strategy whose payoff is that point on the true frontier that is consistent with his risk attitude.

> These strategies will always differ from client to client. As a result, clients will be "looping" throughout the space of all portfolios in different ways (recall Figure 1). *That is, each will exit and enter different markets at different rates and dates.* This follows from the nature of the solution to the specific dynamic programming problem each must solve.

This last point is critical to understand. Please read it carefully.

Now, suppose that an active manager approaches any one potential client and says: "Invest with me. I'll outperform your benchmark for a fee of 1.5%!"

Q: But exactly what is the right benchmark here? What does this term even mean?
A: Here is what it means. The benchmark for the money manager in the case of *this* client with *his* risk tolerance is the entity

$$\textbf{(1)} \qquad \mathbf{B = x + 1.5\%}$$

where **x** denotes the expected risk-adjusted payoff to the client from utilizing the *passive* strategy that is optimal for him given $\mathbf{R^M}$ and given his risk attitude. **1.5%** is the management fee. For it to be rational for the client to utilize the active manager, the expected risk-adjusted payoff **y** from utilizing the *active* strategy that is optimal given $\mathbf{R^{M^*}}$ and given the client's risk attitude must exceed **B**. That is, we require that

$$\textbf{(2)} \qquad \mathbf{y > B}$$

A tricky issue arises concerning the units of measurement of **x** and **y** here, and this is discussed in a footnote.[4]

Q: Since the benchmark will differ according to a client's risk aversion, you end up needing a *different* benchmark for each client with a different risk posture?

---

[4] The passive and active strategies that are optimal for this client will generate two risk/return payoffs **x** and **y** that differ both in expected return *and* risk. [Note that the two payoffs will lie on the "true" and "proprietary" efficient frontiers respectively in Figure 2.] Since these two payoffs must be compared for benchmarking, we render them *commensurate* by computing their so-called "certain equivalent" values. The certain equivalent value of a given lottery is defined in decision theory as that "certain prize" (measured in dollars or whatever) that makes you *indifferent* between receiving it or holding the risky lottery (here, a point on the efficient frontier). These values can be computed from the inverse of the utility function. It is this certain equivalent value that we refer to above when we speak of the "expected risk-adjusted payoff" for each strategy, namely **x** and **y**.

A: Yes, and that is what we stressed in the master essay. All the simplification permitted by the Tobin theorem is thus lost. The benchmark concept remains in the sense that, for any given client, there is a score **B** that must be exceeded. But this will differ from person to person, so the force of the separation theorem is lost.

Q: How important does the concept of "beating the market" remain in the new theory—a concept so critical in classical theory?

A: It becomes much less important. Two reasons will convince you of this. First, note an ambiguity in your question. You did not specify *which* market is to serve as benchmark: US equities, French bonds, gold? Never forget, in classical theory, only *one* market existed, namely the market of all risky assets everywhere. This has no relevance in a world where what matters is which market to move into/out of, when, and why.

Second, once you leave the i.i.d. world of the textbook, and think in terms of strategies rather than portfolios, the performance of any given market—or all markets combined—cannot serve as a benchmark for managing the client's portfolio. For what now matters is the performance generated by *shifting* from market to market (or sector to sector) as dictated by the optimal strategy. *It is thus the performance of strategies that matters.* The performance of individual markets is at best a halfway house to superior returns. This is a difficult concept for many to apprehend.

The Four Principal Hurdles to Understanding the New Theory

Q: To conclude this dialogue, what do you find to be the greatest obstacles to your clients confront in understanding all this? What are the stumbling blocks?

A: There are four. While we shall identify these here, but their discussion will be postponed to the next SED report.

• First is the widespread belief that it is "professionally correct" and "conservative" to adhere to a fixed policy portfolio through thick and thin. It most definitely is not.

• Second is the confusion centering around questions of the "adequacy" and "reliability" and "non-stationarity" of historical data. People will ask: "How much can we trust $\mathbf{R^M}$? Won't it change?" No, it won't, it can't, as it is constructed to be the stationary measure or "invariant" of the non-stationary stochastic process by which the empirical data of history are generated. It is the only thing that *doesn't* change.

• Third is the widespread misunderstanding of the concept of an "average" (e.g., the average returns or risks of an asset class). This misunderstanding is based upon the failure to distinguish between conditional and unconditional probabilities in data analysis. In this vein, the means, variances, and covariances utilized in the Markowitz theory are *not* in fact long-run statistical averages as is always asserted. Rather, they represent the average of an underlying family of statistical averages. Not understanding this difference leads to major problems.[5]

• Fourth is the issue of the "frequency domain" both of decision-making and of econometric analysis. The issue here is one of time frame consistency: If you are looking for "patterns" and "predictability" in the data, know whether you are engaged in hourly, weekly, monthly, annual, or decade-long bets, and then utilize a data generated on the same frequency basis. For data of different frequencies will possess altogether different patterns and thus risk/return structures.

In the next report, we shall investigate these four problems further.

■ ■ ■ ■ ■

[5] Mathematically, **R** can be shown to be the "static" average of the "dynamic" average represented by $\mathbf{R}^{\mathbf{M}}$ in Figure 3. Formally, **R** is the stationary measure of $\mathbf{R}^{\mathbf{M}}$, and is obtainable as its unique unit left eigenvector.