Sharpe's State-Preference Approach and Beyond: *A Practitioner Overview*

H. RUSSELL FOGLER

H. RUSSELL FOGLER is the principal of Fogler Research and Management in Gainesville, FL.

When heterogeneous investors and tail risks are considered, the capital asset pricing model's (CAPM) linear discounting seems too simple. Envision a "typical" investor retiring in either 1929 or 1973 or, to a lesser extent, possibly 2000. If this investor spends 4% to 5% of his or her initial wealth annually over the next 10 years, purchasing power (after inflation) diminishes by 40% to 50% by the end of the decade. The retiree's loss of purchasing power is very great under this scenario (state) with a beta in the range of 0.8 doing little to ameliorate the pain.

In a new book, Nobel Laureate William Sharpe presents 23 case studies in a more robust economic setting than CAPM.[1] Risk is defined by economic states. Heterogeneous investors are involved and sometimes choose portfolios other than the market portfolio.

Sharpe's presentation is grounded in individual consumption and *state claims*. State claims, sometimes called *contingent claims*, *pure securities*, or *Arrow/Debreu securities*, pay one consumable unit when a future state occurs, but zero if the state does not occur. Sharpe sees these state claims as "the simplest possible type of security. In a sense, they are the atoms out of which security matter is constructed" [2007, p. 39].

In Case 1, a stock has four consumable payoffs (5 units, 3 units, 8 units, or 4 units) depending on which of four future states occurs. Rocket science isn't required to see that the stock can be equivalently decomposed into 20 atoms or state claims (5 claims for the first state, 3 claims for the second state, 8 for the third state, and 4 for the last state).

By explaining how investors value such atoms, Sharpe introduces the foundations of the latest academic research on asset pricing. His discussion clarifies concepts such as *complete/incomplete markets, spanning (covering) a risk space, pricing kernels, stochastic discount factors, and risk-neutral probabilities*. The presentation drops the CAPM linearity assumption and incorporates investors' greater aversion to downside risk (i.e., paying more for downside protection).

What's the practical impact? How does this approach modify investment guidance based on the CAPM? Has the role of the CAPM beta changed? Should investors expect asset returns to be linear in beta as in the CAPM? What are the implications for other asset pricing models such as arbitrage pricing theory (APT)? Could these new considerations clarify anomalies found by CAPM researchers, such as the *peso problem*, *nonstationary volatility*, *tail risks*, *correlations going to one* in crises, *smiles* in option pricing, and the *equity premium puzzle*?

This article provides a practitioner's glimpse into Sharpe's framework, and addresses the investment questions just enumerated. In the course of the discussion, we will see

- It's All about the "Total Pie"
- How Corporations Slice the Pie
- How Financial Engineers Dice the Pie
- How APSIM Prices Atoms of the Pie
- How "Kernels" Price the Slices of the Pie
- How to Personalize Portfolios with Sharpe
- Beyond Sharpe's Approach

IT'S ALL ABOUT THE "TOTAL PIE" (CASE 1)

Sharpe's initial case involves two commercial fisherpersons, Mario and Hue. These two fisherpersons operate on the California coast and face two risks: the size of their catch and where the most fish will be found.

On a bad day, their total catch, or "size of the pie," is 80 fish; on a good day, it is 120. But fish are fickle and on any given day the majority may swim either north or south.

Exhibit 1 shows the probabilities and outcomes for both Mario and Hue. For example, Mario's Monterey Fishing Company is located in Southern California and Mario faces a bad-day catch of 50 fish when most of the fish swim south (BadS), which has a 15% probability. On such a BadS day, it is even worse for Hue because she will catch only 30 fish with the same probability; her fishery is located in Northern California at Half Moon Bay. Regardless of where the fish swim, the size of the pie is only 80 on a bad day or 120 on a good day.

Having read about Sharpe's *free* APSIM software, Mario and Hue decide to see if they can diversify their risks.[2] After a few minutes learning the program, they input that data in Exhibit 1. Because Hue is more risk averse than Mario, she input 2.5 as her measure of risk aversion, while Mario used 1.5. Additionally, each specified a time preference of 0.96, which is *approximately* the same as asking for a periodic 4% return on a riskless asset. (The APSIM software requires the additional input of an initial wealth position as well as how much will be consumed in the initial time period, but for simplicity this input and some corresponding output is being omitted in these examples.)

Exhibit 2 shows the benefits of diversification. For example, Column 1 shows that Hue can consume 42.2 fish in the state of BadS even though only 30 fish will be in the North where she is located. Mario will be consuming 37.8 fish, although he was able to catch the 50 that were in the South; this is his price for being less risk averse.

In a recent discussion, Sharpe indicated that he focuses on the size of the total pie (i.e., 80 or 120) which is unaffected by how diversification slices the pie. His analogy helps in the more complex cases to follow, as well as clarifies the diversification decision of Mario and Hue. The fishers attain only systematic (market) risk, adjusted to reflect their risk preferences, while assuming no diversifiable (non-market) risk. For example, Hue does not care whether the fish swim north or south. She gets the same amount for either direction (42.2 for a bad catch or 57.3 for a good catch) with the actual amount depending only on the systematic (market or size-of-the-pie) risk of 80 or 120. She has diversified away all unsystematic (non-market or diversifiable) risk.

The results in Exhibit 2 illustrate in consumption what Sharpe labels the market risk/reward theorem (MRRT), and it applies equally to the CAPM beta risk:

(MRRT) Only size-of-the-pie (market) risk is rewarded.

The rationale is also clear for why a CAPM market risk/reward corollary (MRRC) works:

(MRRC) Do not take non-market risk.

EXHIBIT 1
Probability/Payoff Table

	BadS	BadN	GoodS	GoodN
Probability of State	0.15	0.25	0.25	0.35
Payoffs:				
Mario	50	30	80	40
Hue	30	50	40	80
Size of Market Pie	80	80	120	120

EXHIBIT 2
Consumption after APSIM

	BadS	BadN	GoodS	GoodN
Probability of State	0.15	0.25	0.25	0.35
Payoffs:				
Mario	37.8	37.8	62.7	62.7
Hue	42.2	42.2	57.3	57.3
Size of Market Pie	80	80	120	120

HOW CORPORATIONS SLICE THE PIE

How is it possible to so neatly divide the pie and adjust for its size-of-the-pie risk? Easy, if you are creative with two elements—creating securities and specifying a trading market.

Exhibit 3 shows how securities can be created in units of fish. First, to allow a fuller market for saving and leveraging, a bond has been created that is riskless (i.e., it pays one fish regardless of the state). Then, both Hue and Mario's companies issue 10 shares to divide their catch. Thus, if it is a bad catch because most of the fish go south, Hue's company will catch 30 fish and each share of stock will give 3 fish to the owner. (APSIM requires both a riskless security and wealth positions for immediate consumption and/or trading, but for simplicity these features are not considered here).

The trading mechanism uses a simulated market-maker who iteratively discovers Hue and Mario's supply and demand curves for the securities. Equilibrium is reached when no further trading occurs. The program then outputs ex ante consumption, prices, returns, and portfolios.

Exhibit 4 shows the stock returns that can be expected relative to the market return; the lines result from connecting the extreme points of the portfolio returns. Exhibit 4 indicates that Mario and Hue's expected returns vary in a linear fashion with the market, although this is not always the case, as will be discussed later. Finally, it is important to note that Exhibit 4 is not the typical security market line with beta on the horizontal axis; rather, it merely illustrates how individual portfolio returns are related to market returns and that the market return must be linear by definition of the horizontal axis. Later in the article, however, portfolio returns will be shown that are nonlinear to the market return.

EXHIBIT 3
Payoffs of Securities

Securities:	BadS	BadN	GoodS	GoodN
Bond	1	1	1	1
MFC(10 shs)	5	3	8	4
HFC (10 shs)	3	5	4	8
Size of Market Pie	80	80	120	120

HOW FINANCIAL ENGINEERS DICE THE PIE (CASES 6 AND 7)

Exhibit 5 shows a more realistic and larger investment situation with 10 future return states and 5 consumption levels (i.e., consumption under Boom1 equals consumption under Boom2, consumption under Depression1 equals consumption under Depression2, and so on). By contrast, Mario and Hue faced 4 future return states and 2 future total consumption levels.

As more realistic situations are considered, it is helpful to distinguish between *complete* and *incomplete markets*. In simple terms, if there are n states of market risk ($n = 10$ in this case), investors can fully diversify if there are at least n securities with linearly independent returns. (Exhibit 5 has only six securities, so such diversification does not apply here, as will be shown.) The n linearly independent assets are said to *span* or *cover* the risk space, and the market is said to be *complete*. The simplest way to complete a market is adding state claim securities (atoms) whose returns are linearly independent for different states, by definition.

An *incomplete* market is one that has insufficient assets with independent returns to provide investors with a combination that achieves diversification over all economic states. As will be shown shortly, financial engineers create securities when markets are so incomplete that fees can be obtained by creating new securities. Before examining these concepts, the rest of this section shows how Exhibit 5 relates to complete and incomplete markets.

To consider market trading, two investors Quade and Dagmar (Sharpe [2007, pp. 64–70]) are given more complex risk-preference functions than in the preceding Mario-Hue case. Quade, as the name implies, has a risk-averse quadratic function. In contrast, Dagmar's relative risk aversion decreases as consumption increases.

Exhibit 6 shows a very different situation than the equilibrium in the Mario-Hue case. In this case, the investor (Quade and Dagmar) portfolios exhibit nonlinear and kinked returns relative to the market. The nonlinearity is easy to grasp intuitively because it reflects Quade's and Dagmar's respective utility functions. The Quade portfolio holds securities that provide more safety at lower levels of market return and correspondingly gives up higher returns at higher levels of the market return. In contrast, Dagmar's expected returns reflect her lower degree of risk aversion.

EXHIBIT 4
Returns for Mario-Hue (Case 1)

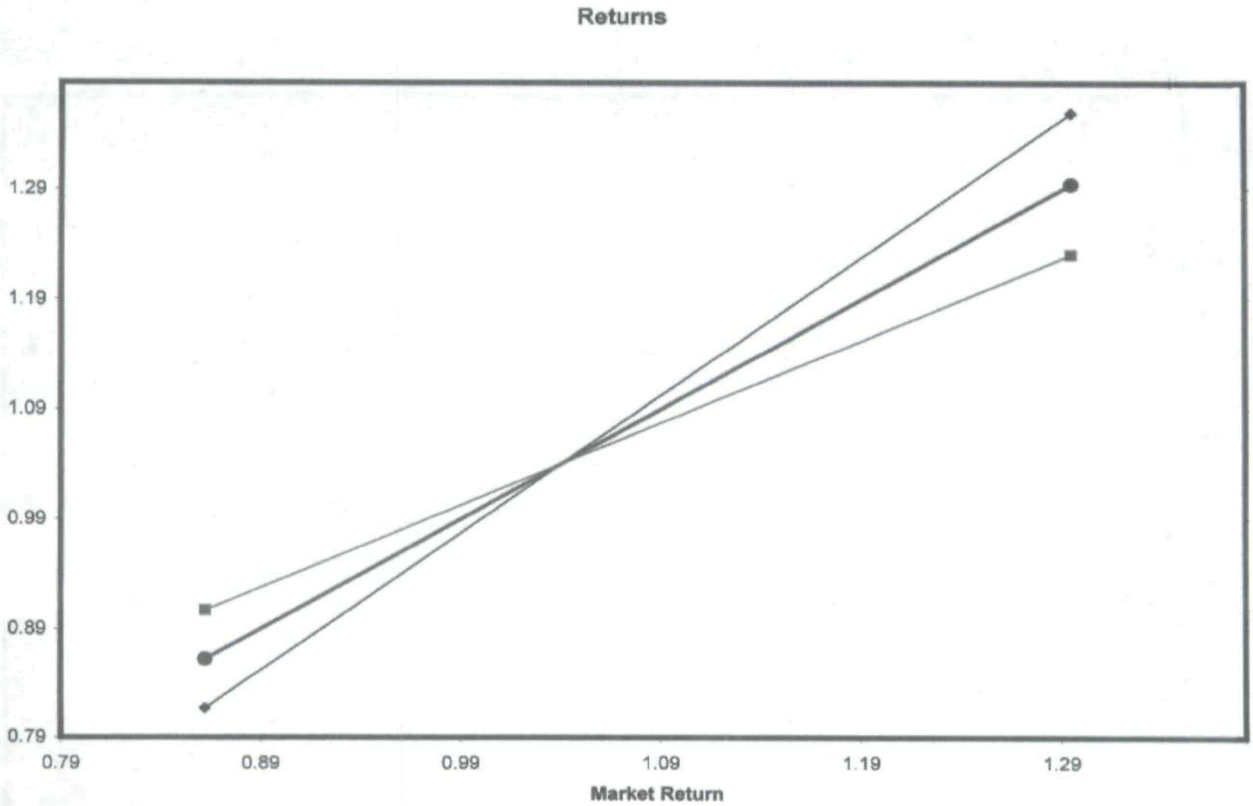

Source: APSIM (◆ Mario • Market ■ Hue).

EXHIBIT 5
Quade-Dagmar (Cases 6 and 7) Securities and States

			SECURITIES					
	Probability		STBond	GovBds	NonGBd	ValueSt	GthStx	SmStx
	0.05	Depression1	1.000	0.959	0.890	0.877	0.779	0.856
	0.05	Depression2	1.000	1.067	1.047	0.744	0.716	0.806
	0.10	Recession1	1.000	0.897	0.904	0.923	0.961	0.899
STATES	0.10	Recession2	1.000	0.914	0.929	0.930	0.958	0.821
	0.20	Normality1	1.000	1.069	1.054	1.151	0.981	1.114
	0.20	Normality2	1.000	0.964	0.994	1.185	1.058	1.167
	0.10	Prosperity1	1.000	1.193	1.313	1.109	1.119	1.119
	0.10	Prosperity2	1.000	1.043	1.038	1.240	1.305	1.192
	0.05	Boom1	1.000	1.106	1.122	1.179	1.337	1.419
	0.05	Boom2	1.000	1.109	1.114	1.295	1.306	1.261

Source: APSIM.

EXHIBIT 6
Quade-Dagmar (Case 6) Kinked Returns from Incomplete Market

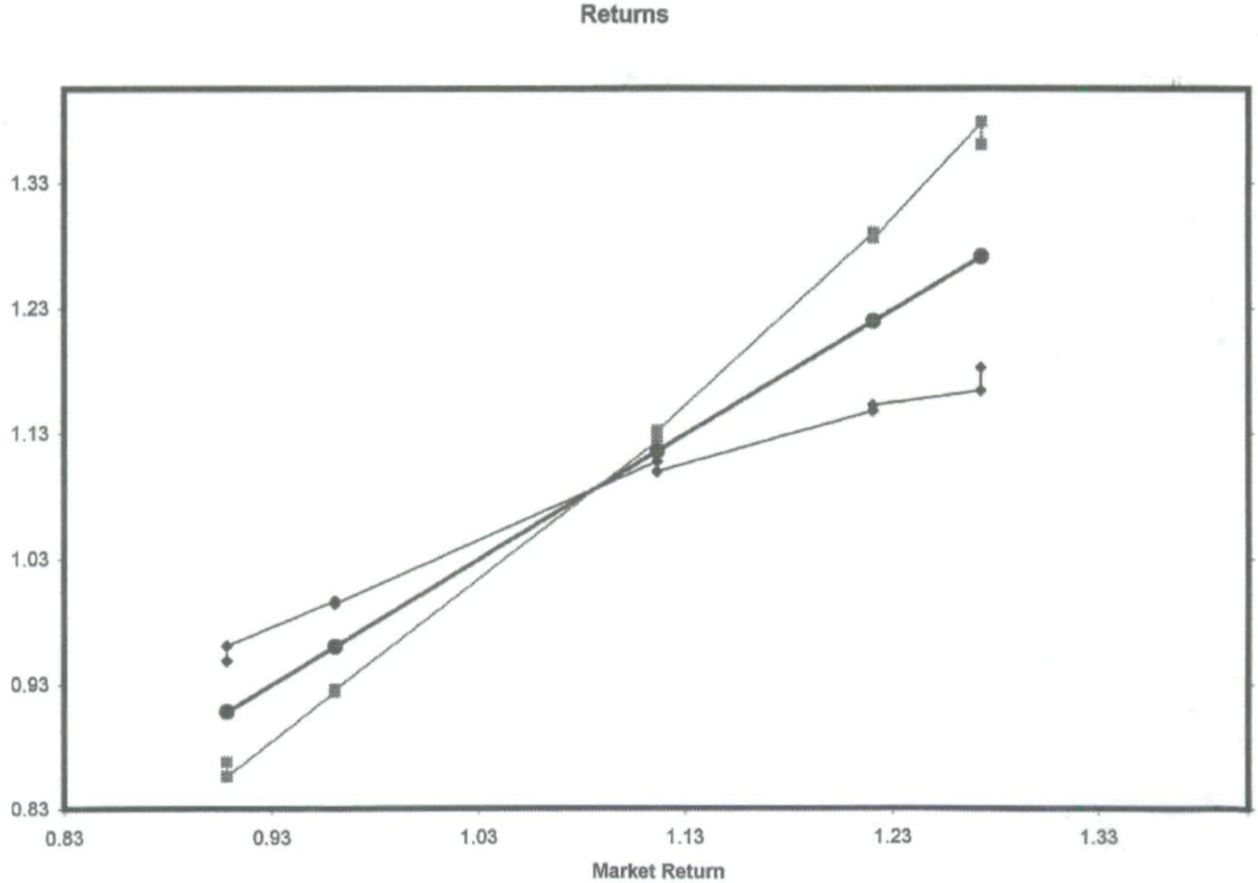

Source: APSIM (◆ Quade • Market ■ Dagmar).

More confusing are the kinks at points in the return curves. The kinks indicate that, although trading has stopped, both investors have assumed non-market risk. As each investor tries to attain his or her goals for safety and higher return, the only security choices lead to multiple outcomes depending on the economic scenario. For example, even though total consumption and market returns are equal under either Boom1 or Boom2, Dagmar will have a higher return under Boom1 if her portfolio invests in the Sm/Stx Index Fund to achieve her desired higher marginal utility for high-consumption states.

One way to eliminate multiple outcomes (kinks) is to create a *complete market* by introducing tradable state claims for each state. As mentioned previously, such securities pay off in one unit of consumption if the corresponding state occurs and zero otherwise. Exhibit 7 shows the return curves when the market is complete; by trading both state claims and index funds, Quade and Dagmar attain their desired consumption without the quandary of multiple outcomes. In this case, Quade and Dagmar do not have to assume non-market risk—each follows a *non-linear market-based strategy* (discussed more extensively later in the article). In the real world, however, financial engineers are unlikely to introduce state claims, not because of the cost, but because of insufficient demand.

A second way to potentially eliminate multiple outcomes is for financial engineers to create derivative products that appeal to investors' unique needs and contribute to a completion of the market. The situation facing Quade and Dagmar is such an opportunity. In fact, in Case 20, puts and calls are used with index funds to create portfolio returns *nearly* identical to the complete market result in Exhibit 7 (Sharpe [2007, pp.153–158]). Such financial

EXHIBIT 7

Quade-Dagmar (Case 7) Smooth Returns from Complete Market Equilibrium

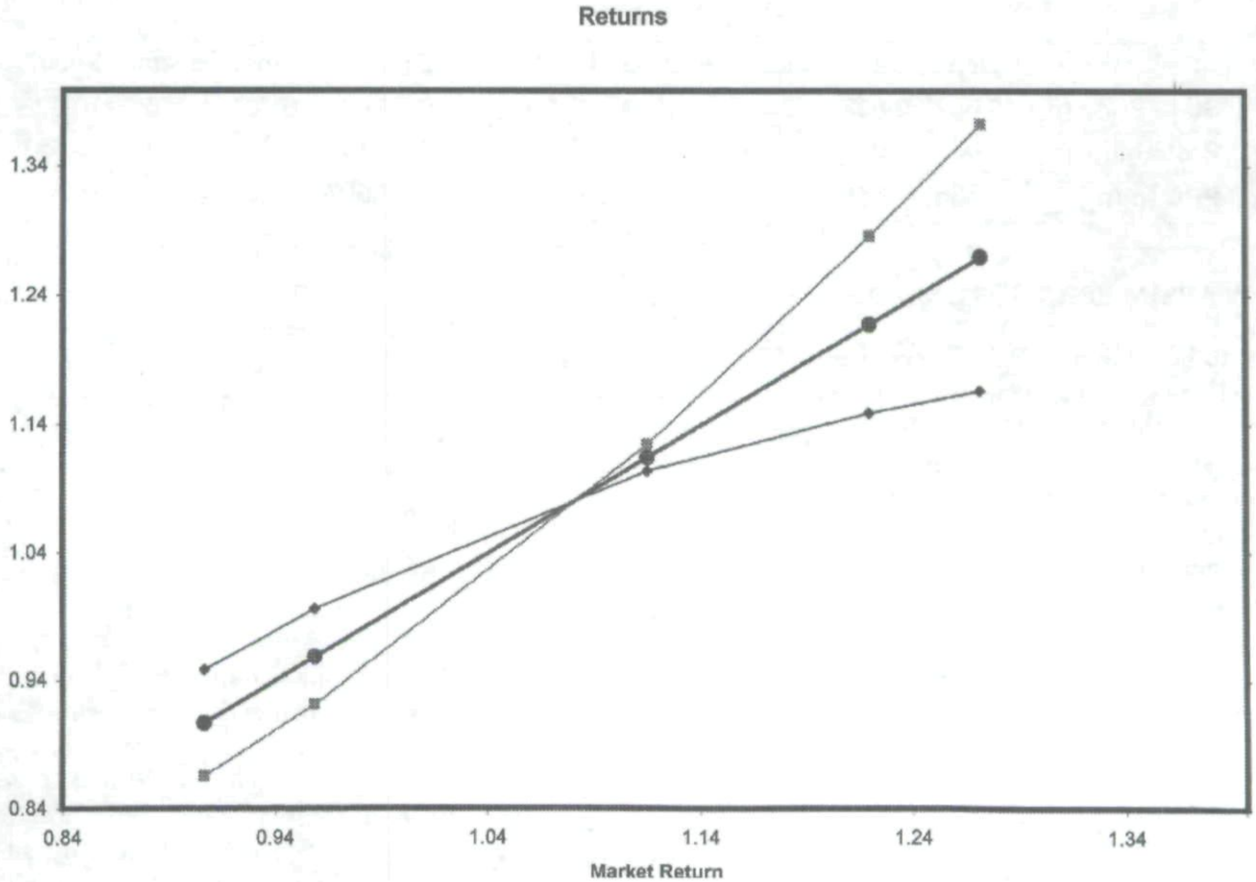

Source: APSIM (◆ Quade • Market ■ Dagmar).

engineering illustrates how an incomplete market can be transformed into a "sufficiently complete" market, as will be defined below after more explanation about APSIM's pricing. Indeed, when fees are sufficiently large, financial engineers may sell a single contract, using dynamic rebalancing and/or options to create a *replicating portfolio* of potential state claims to hedge an investor's risk of a particular state occurring (Sharpe [2007, pp. 84, 158–161]). Additionally, a specific state-risk slice of the pie may be cut into smaller pieces (tranches) to provide derivatives with different payoffs for the same state. Such slicing and dicing allows financial engineers to create virtually any shape of return distribution—it is only a question of sufficient investor demand to cover associated costs.

Once a market is completed, prices are available for each state claim. In the previous example in which the market was completed by trading state claims, the equilibrium price for $1 payoff in each state is shown in Exhibit 8 (for now, ignore the last row labeled "PPC," although it will be important later).

For example, Exhibit 8 shows that $1 in a depression is worth over three times as much as in a boom (i.e., 0.082 vs. 0.026). Certainly, in a depression, money for food and shelter is valued more highly than an extra restaurant dinner. Also, $1 in stable economic times is worth more than the other dollars because the investor has a greater chance (40%) of receiving it than the dollars in other state claims. Most importantly, the sum of the prices of the states provides the current one-period discount rate of approximately 0.95. This is equivalent to the present value of a riskless $1 because one of the states must occur, and therefore the investor is guaranteed to receive a $1 payoff.

EXHIBIT 8

Quade-Dagmar (Case 7) Complete Market State Prices and PPCs

	STATE									
	Dep1	Dep2	Reces1	Reces2	Norm1	Norm2	Pros1	Pros2	Boom1	Boom2
State Price (p_s)	0.082	0.082	0.139	0.139	0.168	0.168	0.061	0.061	0.026	0.026
Probability (π_s)	0.05	0.05	0.10	0.10	0.20	0.20	0.10	0.10	0.05	0.05
PPC (p_s/π_s)	1.635	1.635	1.389	1.389	0.840	0.840	0.607	0.607	0.520	0.520

HOW APSIM PRICES ATOMS OF THE PIE

To better understand the state prices in Exhibit 8, it is necessary to understand how investors set reservation prices to trade in APSIM. Each heterogeneous investor is assumed to maximize the combination of the current period's utility ($U(c_t)$) and the expected value of next period's utility ($U(c_{t+1})$), where $U(\cdot)$ represents the investor's utility function and c_t represents consumption in period t. Giving the investor a time discount factor of d for a one-period delay of gratification, it can be shown (Cochrane [2005, p. 5]) that the investor will be willing to pay p_t for a consumable payoff of x_{t+1}, or

$$p_t = E_t\,[d(U'(c_{t+1})/U'(c_t))\,x_{t+1}] \qquad (1)$$

where $U'(\cdot)$ is the first derivative of $U(\cdot)$ and $E[\cdot]$ is the expectation operator. Simply stated, investors maximize their expected utility by bidding for securities based on the marginal rate of substitution (MRS) between today's consumption and future consumption. Sharpe [2007, pp. 35–45] skips the mathematical derivations and simply explains how investors work toward maximization by considering four elements: the time discount for state j, consumption in each state, the marginal utility of such consumption, and the probability of the state.

APSIM's pricing iterations begin with all investors submitting their reservation prices. The reservation price indicates the investor will sell (buy) the security if the market-maker's trade price is higher (lower) than the investor's reservation price. Once all investors have submitted their reservation prices, the market-maker announces a trade price that is an average of the individual prices. Then, each investor determines the quantity of shares he or she would trade at the market-maker's price, again, maximizing expected MRS. The market is cleared by matching buyers and sellers for the smaller of either the shares offered or bid for. This process is repeated until the market-maker sees that prices are changing less than some preset limit. At that point, the program reports the last reservation prices for both securities and state claims. If the markets did not trade in state claims, APSIM reports state prices as a simple average of the averages of buyers' final state reservation prices and sellers' reservation prices (Sharpe [2007, pp. 87–88]).

When APSIM's market-maker concludes trading, the final reservation prices indicate the degree of market completeness. If state claims were traded, all investors' reservation prices for the states will be the same, at least within the preset limit that ends trading. This is a complete market. It is also possible to have what Sharpe calls a *sufficiently complete market*, which is defined as a market in which either all investors' state claim prices are the same or in which no trading would occur if trading in state claims were introduced (Sharpe [2007, pp. 88–89]). As discussed in the next section, sufficiently complete markets are similar to complete markets with one difference. An incomplete market, by contrast, does not qualify as either complete or sufficiently complete.

HOW "KERNELS" PRICE THE SLICES OF THE PIE

Two methods can be used to price assets—state prices or pricing kernels. Although pricing kernels are, by definition, for continuous probability distributions, Sharpe clarifies kernels by using analogous discrete state probabilities. Both methods produce the same asset price.

In complete markets, asset pricing is straightforward with state prices—sum state prices times state payoffs for all respective states:

$$P_i = \Sigma_s(p_s)\,X_{is} \qquad (2)$$

where P_i represents the price of asset i, p_s is the price of a state claim in state s as calculated in the preceding section, and X_{is} is the ith security's payoff in state s (Sharpe [2007, p. 84]).

In the example of financial atomization given in the introduction to the article, a stock is created with four consumable payoffs (5 units, 3 units, 8 units, or 4 units) depending on which of four future states occurs. As stated earlier, the stock can be equivalently decomposed into 20 atoms or state claims (5 claims for the first state, 3 claims for the second state, 8 for the third state, and 4 for the last state).

This example, using Equation (2), indicates that the price of this stock would be the sum of each of the claims multiplied times the respective state claim price (i.e., $P = (p_1 \times 5) + (p_2 \times 3) + (p_3 \times 8) + (p_4 \times 4)$).

In complete markets with the no-arbitrage condition, the Law of One Price (LOP) prevails. The LOP follows directly from Equation (2) because every asset would be priced by the same set of state claim prices. If any asset was not priced by Equation (2), an arbitrageur would buy or sell the mispriced asset with the proceeds being invested in, or taken from, the correctly priced assets. Such arbitrage, of course, requires sufficient liquidity and no other transaction costs.

In contrast to the use of state prices, which merely sums the prices times the payoff, *pricing kernels* use an expected-value specification. Sharpe designates the pricing kernel as the set of PPC_s, or price per chance for state s (Sharpe [2007, pp.74 and 90]). The pricing equation can be derived directly by multiplying Equation (2) by (π_s/π_s) where π_s is the probability of state s:

$$P_i = \Sigma_s\, \pi_s(p_s/\pi_s)X_{is} \qquad (3)$$

Or,

$$\textit{letting } PPC_s = (p_s/\pi_s)\text{: } P_I = \Sigma_s \pi_s(PPC_s)X_{is} \qquad (4)$$

As the previous section discussed, state claim prices have four elements, of which one is the probability that a specific state occurs. The PPC subsumes this probability, thus aligning the prices of every state claim on a comparable basis. For example, the *PPCs* in Exhibit 8 show that risk-averse investors pay more on a relative basis for protection (a "Mafiaesque" concept) against the worst case ($PPC = 1.635$) compared to normal states ($PPC = 0.84$).

Exhibit 9 graphs the PPC_s from Exhibit 8 against consumption. The downward slope of the pricing kernel reflects investors' decreasing marginal utility for increasing consumption. This downward slope is necessary for achieving equilibrium.

Today, the starting point for most doctoral research on pricing is pricing kernels, for two reasons—simplicity and generalization capability. These characteristics are obvious in the formulation that Sharpe [2007, p. 89] denotes as Cochrane's Basic Pricing Equation (BPE),

$$\mathrm{P} = \mathrm{E}[\mathrm{mX}] \qquad (5)$$

where m is a continuous function equivalent to the PPC_s in the world of discrete states and $E[\cdot]$ is the expectation operator. The variable m is generally called a stochastic discount factor or price density function, and the designation m may be thought of as coming from investors' marginal rates of substitution (MRS) between the present ($t = 0$) and the future ($t + 1$), as discussed in the preceding section. A kernel is similar to the core or seed of a plant in that the kernel of a probability density function is the basic part from which everything else can be derived. Once $\exp(z)$ is known, the constant can be derived and the entire normal density can be determined by substituting the variables of interest into the relationship. Similarly, once the pricing-density kernel is specified, asset prices can be derived by substituting the corresponding state payoff function X.

What could be simpler than saying that all absolute asset pricing is simply the expected value of discounted cash flows, or $P = E[mX]$? This approach can also work for returns. If both sides of the equation are divided by P, the result is $1 = E(mR)$ where R is the return vector. The kernel function can then be specified in terms of relative pricing models such as CAPM and APT. Also, when the equation is put in return form, it is easy to see that m is a discounting function; for example, if the discount rate were 0.95, then \$1 invested today should return approximately \$1.053 at the end of the first period. Further, Sharpe [2007, pp. 90–98] shows that, with some algebraic manipulations, $1 = E[mR]$ directly implies that m covaries with both X and R when the BPE holds. And Cochrane [2005, pp. 7–8] showed that all asset pricing models can be viewed as special cases of BPE.

Having reduced all pricing to such simple terms, it is necessary to add an important caveat—the pricing statements for Equations (2) through (5) hold only when markets are either complete or sufficiently complete. When markets are only sufficiently complete, multiple solutions can provide different state prices, although the equations

EXHIBIT 9

Pricing Kernel Plotted against Consumption (Case 7)

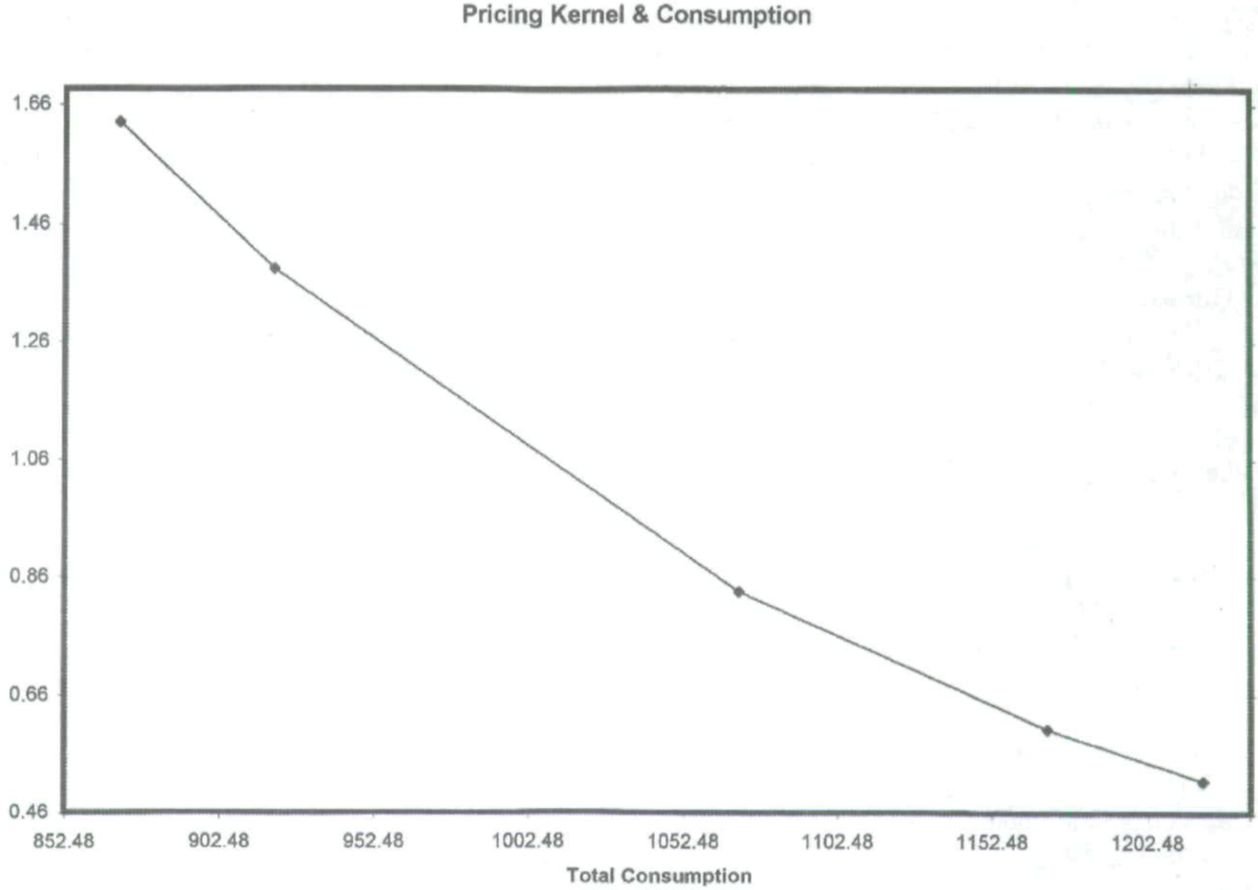

Source: APSIM.

and LOP hold for any set of such prices (Sharpe [2007, pp. 84–90]). Even if the market is only sufficiently complete, Sharpe makes the reasonable argument that APSIM's market-maker function provides a reasonable set of such prices that contain equilibrium information for the next trade if the market reopens. By contrast, if markets are incomplete, while the equations may still hold if no arbitrage is possible, the state prices are not unique and, further, contain no equilibrium information about the next trades if the market reopens.

Finally, to the extent that markets are either complete or sufficiently complete, MRRT and MRRC hold. Thus, when markets are complete or sufficiently complete, investors should only follow market-based strategies because non-market risk has no benefit. Earlier Mario and Hue did this with linear market-based strategies, and Quade and Dagmar did this with nonlinear market-based strategies. Such strategies, however, do not say that asset returns are expected to be linear to the CAPM beta.

PERSONALIZING PORTFOLIOS WITH SHARPE

Sharpe's model has a one-period horizon and, thus far, three conditions have been invoked: "agreement among investors on the probabilities of future outcomes, the absence of outside sources of consumption, and the lack of state-dependent preferences among investors on the probabilities of future outcomes at the same time period" [2007, p. 108]. These simplifying assumptions need to be dropped because every practitioner recognizes that investors differ.

Sharpe's modeling recognizes this and stresses three personal differences, or what might be called the "three Ps":

preferences, positions and predictions. Only heterogeneous preferences have been discussed to this point, but Sharpe [2007] discusses other types of heterogeneity, such as differing positions with respect to taxes, job risks, nonsalaried income (Chapter 5), and differing predictions on future market returns (Chapter 6). In contrast to the homogeneous investors of the CAPM and the MRRC dictum, some investors are counseled not to invest the market weight in stocks that is correlated to their job risk. Such guidance is shown to produce complex pricing effects (e.g., Case 10 exhibits a smooth pricing kernel against consumption states, but a kinked or multiple outcome curve when plotted against market return states). Chapter 6 shows that, to the extent investors have differing predictions, they may affect prices as they choose to bet on various future outcomes.

BEYOND SHARPE'S APPROACH

Sharpe limits his analysis to one-period (two-date) settings. But as a practitioner I am concerned with multiperiod settings. To this end, the remainder of the article suggests useful extensions of Sharpe's book.

Personalizing Creates New Slices

Modeling personalized states can be difficult. For example, for a retiree who experiences a tail-risk situation of meager portfolio returns with negative cash flows (i.e., spending risk) over 10 periods, one difficulty is the specification problem of how to include a portfolio's cash inflows and outflows with their asymmetric compounding effects. Even with no outflows, a 20% loss requires a 25% gain to break even, and a greater gain if an outflow occurs during the losing period. Another difficulty is specifying the time horizon and associated probability. If, for example, there are three 10-year periods of low cumulative returns in the past 80 years, should the probability be 3/8 rather than 3/80, even if the 10-year periods were generated by five 1-year bear markets? And, if returns are in real terms, how is inflation to be considered? Sensitivity measures for portfolio returns in tail-risk states would certainly be useful. Maybe these sensitivities could be called *state betas* although they would not be related to the CAPM beta. Fortunately, in addition to multi-maturity bonds, financial engineers provide derivative products that can be used to hedge all but the most personal tail risks.

Ex Post Uncertainty Adds Factors

Equilibria are a sequence of equilibriums and, over multiple periods, risk can turn into uncertainty (defined by Knight [1921] as events for which ex ante probabilities cannot be defined). The best examples of this may well be found in multiperiod stochastic models that have heterogeneous investors with differing risk preference functions and differing forecasting models. Such models generate returns that mirror real markets with multiple ex post factor returns that reflect both permanent (ex ante) and transient (ex post) states. For further discussions of multiperiod simulations, see Farmer [2001]; Fogler [2001], and Kyle and Wei [2001], as well as Sharpe [2007, p. 209] on the complexity of modeling equilibrium outcomes in a multiperiod setting.

Applying a Personal Weighting Scheme

The preceding personalization and ex post considerations lead practitioners to develop their own state-space models, augmented by a widely used practitioner weighting scheme, which can be viewed as a transformation of Equations (3) and (4):

$$\textit{If financially engineered, } P_i = d\Sigma_s(\pi_s\ (\pi^*_s/\pi_s)\ X_{is}) \quad (6)$$

$$\textit{Or simply, } P_i = d\Sigma_s(\pi^*_s\ X_{is}) \quad (7)$$

where π^*_s is similar to $(\text{PPC}_s/d)\pi_s$ or $(m_s/d)\pi_s$ and is called the *risk-neutral probability* of state *s*, and *d* is a constant discount factor for the risk-free rate (Cochrane [2005, pp. 51–52; Sharpe [2007], pp. 84–85). The variable π^*_s reflects the financial engineer's subjective assessment of how to price risk. Although less than or equal to one and summing to one over the states of *s*, π^*_s potentially captures two subjective variables: a subjective probability of the state *s* and a reweighting of the probability to reflect a subjective estimate of the true *m*. Technically, π^*_s is a subjective probability for calculating expected value, but some relative pricing applications effectively adjust for implied utility. Risk-neutral pricing is simple and straightforward which has made it especially useful in pricing many derivative products.

The Main Factor is Still the Size of the Pie

Despite short-term price effects, markets ultimately price only the total pie. Historical charts are clear that prices fluctuate around total company cash flows, although the fluctuations are sometimes large.

Because of this total-pie effect, the CAPM beta is a great proxy for ex post portfolio volatility. When the no-arbitrage condition holds, Sharpe shows that "the expected return of a security or portfolio is related to its covariance with a function of the return on the market portfolio" [2007, p. 109]. Further, because a portfolio is a weighted average of individual securities and every asset's return covaries with the state's discount rates (i.e., the pricing kernel) it is not surprising that the CAPM beta would be a good proxy for volatility risk even if the linearity of the CAPM does not apply. Thus, even *power beta* (a concept defined by Sharpe to reflect risk preference functions with constant relative risk aversion) often exhibits very similar pricing relationships to the simple CAPM beta, as discussed in Sharpe [2007, pp. 99–101] and shown in APSIM's output of security market lines.

CONCLUSION

In his recent book, Sharpe's goal is to simplify asset pricing via state claims (atoms). In the introduction to his book, Sharpe notes that, in contrast to MBA curricula focusing on CAPM, "If you were to attend a Ph.D. finance class … you would learn about no-arbitrage pricing, state claim prices, complete markets, spanning, asset pricing kernels, stochastic discount factors, and risk-neutral probabilities" [2007, p. 4]. As such, the intended readership is "those who would like to understand more of the material now taught in the Ph.D. classroom but who lack some of the background to do so easily."

The concept of a pricing kernel (designated as m) is central to Sharpe's paradigm. This paradigm originates from academicians searching for a unified economic theory that begins with total economic output, encompasses utility maximization by heterogeneous consumers, and ends with, through market interacting, realistic equilibrium asset prices. As shown in the preceding sections, such a paradigm is likely to lead to higher discount rates for states with lower expected returns (*smiles* or convex functions) in asset pricing. As our retiree example illustrated, very sharp risk aversion seems reasonable for low-return states because multiperiod wealth expectations are diminished if spending is too high; such high risk aversion may partly explain the equity premium puzzle (i.e., excessively high equity risk premiums occurring in historical data). For other explanations and research on the equity premium puzzle, see Cochrane [2005, Chapter 21]). Interestingly, Sharpe's cases exhibit realistic pricing from his risk preference functions, but this may be due to the return estimates and risk preference parameters having been "calibrated" to approximate realistic results.

Four features of pricing kernels are especially appealing. The first is generality, because m does not require a specific $U(\cdot)$ and could be for one investor and *any* time period, although such a general use would not provide equilibrium pricing. Second, when the no-arbitrage condition holds, equilibrium pricing is clarified and specified in two ways—in absolute terms (by the Basic Pricing Equation, BPE) and in relative terms ($1 = E(mR)$). Third, the concept clarifies that systematic risk arises from m covarying with economic states (of course, this was true in the CAPM because m was specified as a linear function of the market return, which proxied for the economic state variable). Finally, fourth, m can also be generalized to any asset pricing model, including multifactor formulations such as the APT (e.g., $m = a + bR_M$ for CAPM and $m = a + b_1 f_1 + b_2 f_2 + \cdots + b_n f_n$ for APT, where R_M is the market portfolio's return and f_n represents the nth APT factor).

Using this framework combined with discrete economic states, Sharpe's cases and simulations summarize into three corresponding mega-messages:

- Discount rates co-vary with economic scenarios
- Pricing kernels are likely to be nonlinear
- Portfolio risk likely rests on more than the CAPM beta

Despite the generality of Sharpe's model with its inclusion of the three Ps of hetereogeneity (preferences, positions, and predictions), many CAPM insights remain valid or nearly valid. In particular, the CAPM dictum of MRRT often holds (even MRRC is seen to have wide validity), partly because financial engineers are creating securities leading to markets which approach the definition of "sufficiently complete." Also, linear and nonlinear market-based strategies are often appropriate as shown in the article's discussions of Sharpe's Cases 1, 6, 7, and 20.

Additionally, the state-preference framework clarifies CAPM issues such as nonstationary correlations going to 1 during crises, tail risk, anomalies, and the peso problem. As ex post realization of states occurs (nonstationarity), some low-probability states (tail risks) are likely to have all equity assets exhibiting negative (positive) returns while all the safest fixed-income assets are exhibiting positive (negative) returns; hence, the frequent observation that, in crises, correlations go to either +1 or −1. Further, because ex post occurrences of states are unlikely to equal ex ante probabilities, anomalies may be found in ex post data. One example of such anomalies is the peso problem, named for the high returns continuously achieved by shorting the dollar against the peso before a state of peso devaluation occurred.

Finally, because only approximately 20% of APSIM's cases have been discussed in the article, practitioners will be well served by learning Sharpe's other insights for heterogeneous investors. In the end, the markets may not fit models because of minor imperfections (or the models may not fit the markets!!!!), but we can expect good practice and good theory to align as financial engineers create new products to make the markets more complete. And, certainly, Sharpe's book and modeling provide practitioners with the foundation for understanding our next generation of investment practice and modeling.

ENDNOTES

[1]I want to especially thank Bill Sharpe for granting permission to paraphrase major sections of his text, as well as to utilize his cases 1, 6, 7, and 20. Also, Marty Gruber was especially helpful with observations on the current state of the art in academic asset pricing. Finally, both Ira Horowitz and Mary Ida Compton made substantial comments that have greatly improved the article.

[2]Sharpe provides his simulation at no cost and with his encouragement that we use it to develop and test our own ideas (APSIM@http://www.stanford.edu/~wfsharpe/apsim/index.html). APSIM has been designed to facilitate testing all sorts of individual preferences, even admitting kinked utility functions. Additionally, security market line (SML) graphs are provided for examining linearity in the CAPM beta, as well as Sharpe's power beta (Sharpe [2007], pp. 99–101). The software includes the inputs for all 23 cases in his text.

REFERENCES

Cochrane, J. *Asset Pricing.* Princeton, NJ: Princeton University Press, 2005.

Farmer, J.D. "Toward Agent-Based Models for Investment" in *Developments in Quantitative Investment Models*, AIMR Conference Proceedings. Charlottesville, VA: CFA Institute (formerly Association for Investment Management and Research), 2001, pp. 61–70.

Fogler, H.R. "Overview: Developments in Quantitative Investment Models" in *Developments in Quantitative Investment Models* AIMR Conference Proceedings. Charlottesville, VA: CFA Institute (formerly, Association for Investment Management and Research), 2001, p. 2–3.

Knight, F.H. *Risk, Uncertainty, and Profit.* Boston, MA: Houghton Mifflin Company, 1921.

Kyle, A., and X. Wei. "Contagion as a Wealth Effect." *Journal of Finance*, 56 (2001), pp. 1401–1440.

Sharpe, W.F. *Investors and Markets.* Princeton, NJ: Princeton University Press, 2007.