An Asset–Liability Version of the Capital Asset Pricing Model with a Multi-Period Two-Fund Theorem

M. Barton Waring and Duane Whitney

M. Barton Waring is a retired financial economist. He has written many articles addressing strategic asset allocation and investment policy.
barton.waring@aya.yale.edu

Duane Whitney is a principal at Barclays Global Investors in San Francisco, CA.
duane.whitney@barclaysglobal.com

"Everything should be made as simple as possible, but not one bit simpler."

—Attributed to Albert Einstein[1]

One of the most powerful, elegant, and delightfully simple constructs in modern portfolio theory is the Two-Fund Theorem—sometimes called the Two-Fund Separation Theorem or the Mutual Fund Theorem—first articulated by Tobin [1958]. A few years later, the capital asset pricing model (CAPM) of Sharpe [1964] created a new theory of market behavior. One of the key results in the process was a similar, although differently motivated, Two-Fund Theorem. Therefore, we will refer to the Two-Fund Theorem from time to time as the Tobin-Sharpe construct. As one of the key foundation stones of modern portfolio theory, the theorem can be summed up by saying that all desirable, or "efficient," portfolios can be constructed out of long and short positions in just two funds: one delivering the riskless rate of return, the other consisting of the market portfolio of all risky assets. This construct has become all but synonymous with the conventional CAPM and is as well accepted as any financial model—at an intellectual level. The caveat is that few if any practicing investment policy strategists use the strict two-fund form of the CAPM as a basis for developing actual investment policies. It seems that no strategist wants his clients' portfolios to have a substantial exposure to cash. Is the Two-Fund Theorem too simple, too elegant, to be useful in practice?[2]

Our intent in this article is to go beyond the asset-only Tobin-Sharpe construct in order to develop a capital asset pricing model parallel to Sharpe's CAPM, but incorporating the aggregate plans that investors have for using their assets, that is, the present value of future consumption. This present value is a key portion of the "economic liability" on the right-hand side of investors' individual economic balance sheets.

We use the terms "liability" and "economic liability" interchangeably, with the generality of meaning allowed to economists, but denied perhaps to lawyers. Our liability isn't necessarily debt enforceable in court, but the present value of expected future consumption, or spending. In the context of an individual investor, the liability might be the investor's retirement spending needs through his maximum possible age, including even his plans for bequests (neither of which is likely to be a legal obligation). For a defined benefit pension, the liability is the present value of probabilistically weighted expected future benefit payments, likely to be somewhat higher than the mere legal obligation (where the latter is usually measured with the limiting assumption that the plan is being immediately terminated with no further new accruals).

Our approach is that of proof by contradiction. First, we'll incorporate the liability using standard single-period forms of the CAPM and such standard forms of mean-variance surplus utility functions as have been in common use since Leibowitz [1987] and Sharpe [1990]. From these, we will derive our first effort, a *three-fund theorem* that is appropriate for investors with a liability and is intended to be consistent with the conventional CAPM. In this theory, an investor always holds three funds: a liability-matching asset portfolio (*LMAP*), the risk-free asset, and the risky asset portfolio of the conventional CAPM.

In the second section of the article, taking a closer look at the three-fund theorem, we will show that it is insufficient to describe a world in which investors hold assets for multi-horizon future use. If all investors held an *LMAP*, then all investors could not also hold a complete version of either the risk-free asset portfolio or the risky asset portfolio. This lack of macro-consistency is unacceptable, revealing an inadequate specification of both the risk-free asset and the risky asset portfolio in conventional approaches.

Accepting the premise that holding investment assets represents a choice by the investor to consume tomorrow rather than today, we will resolve this conflict by placing special weight on our earlier observation that *all* assets in the market are held to fund consumption deferred over multiple future periods. In a macro-consistent world, the *aggregate* risk-free asset must be equal in value and character to the aggregate of all of these deferred consumption ladders—the aggregate economic liability—and must match the liability in its key financial characteristics, including its term structure. This aggregate risk-free asset is not the single-period risk-free asset of the conventional CAPM; it is necessarily multi-period in nature. It might remove some of the mystery from the discussion to say that the new risk-free asset is a schedule of payments—a bond—having the same aggregate financial characteristics as the aggregate liability (sometimes called a superbond to reflect the fact that the payments prior to maturity are not all the same).

This logic leads us to recast the three-fund theorem as a new two-fund theorem in which the multi-period *LMAP* is seen as a more fully general kind of risk-free asset. Investors thus hold the following two funds: 1) a liability-matching asset sized to the full value of their liability and thus risk free, and 2) a fund held in addition to (not instead of) the *LMAP*, consisting of the world market portfolio of risky assets, where the size of the investor's exposure to the risky asset portfolio is determined by her risk tolerance. Note that the *LMAP* consists of a full term structure weighted by each investor's unique future spending plan and is inclusive of traditional "cash" versions of the risk-free asset.

An intriguing corollary affects the second fund: If the risk-free rate has a term structure, then the risk premium must reflect it; the risky asset portfolio cannot be adequately specified in terms of a single-period risk-free rate of return. This second fund, the risky asset portfolio (*RAP*), is a *revised* version of the world market portfolio in which the risk premia earned by its holdings are measured in excess of the return on a multi-horizon risk-free asset, in this case that of the aggregate market. This risk premium is quite different from that of the conventional asset-only two-fund theorem and of the capital asset pricing model, and it repairs the inadequacy of the single-period model revealed by our parsing of the three-fund theorem.

Our ambition isn't so much to develop a new equilibrium theory of the CAPM, however, but rather to give practitioners a familiar and tractable intertemporal model for developing investment policy. Although our views are presented within the framework of the conventional CAPM, using mean-variance utility, we believe them to be consistent with those found in more sophisticated asset pricing theories (including the arbitrage pricing theory of Ross [1976] and the large family of models inspired by Merton's intertemporal model [1973]). This article is thus intended to present a completely general approach to developing strategic asset allocation policy, one that is applicable to all investors—individuals, defined benefit (DB) retirement plans, foundations, and endowments.

The Original Two-Fund Theorem

As noted earlier, the original Two-Fund Theorem states *all* optimal portfolios consist simply of some combination of 1) an asset that is risk free over a single period (unspecified in theory, but in practice often one year and occasionally one month or one quarter), and 2) the world market portfolio of risky assets.[3] In effect, on a graph drawn in risk–return space, such as in Exhibit 1, the theorem states that the ideal and "most efficient" efficient frontier is on the straight line starting at the single-period risk-free rate (on the vertical, or expected return, axis), extending through and beyond the data point indicating the expected

EXHIBIT 1
The Two-Fund Theorem

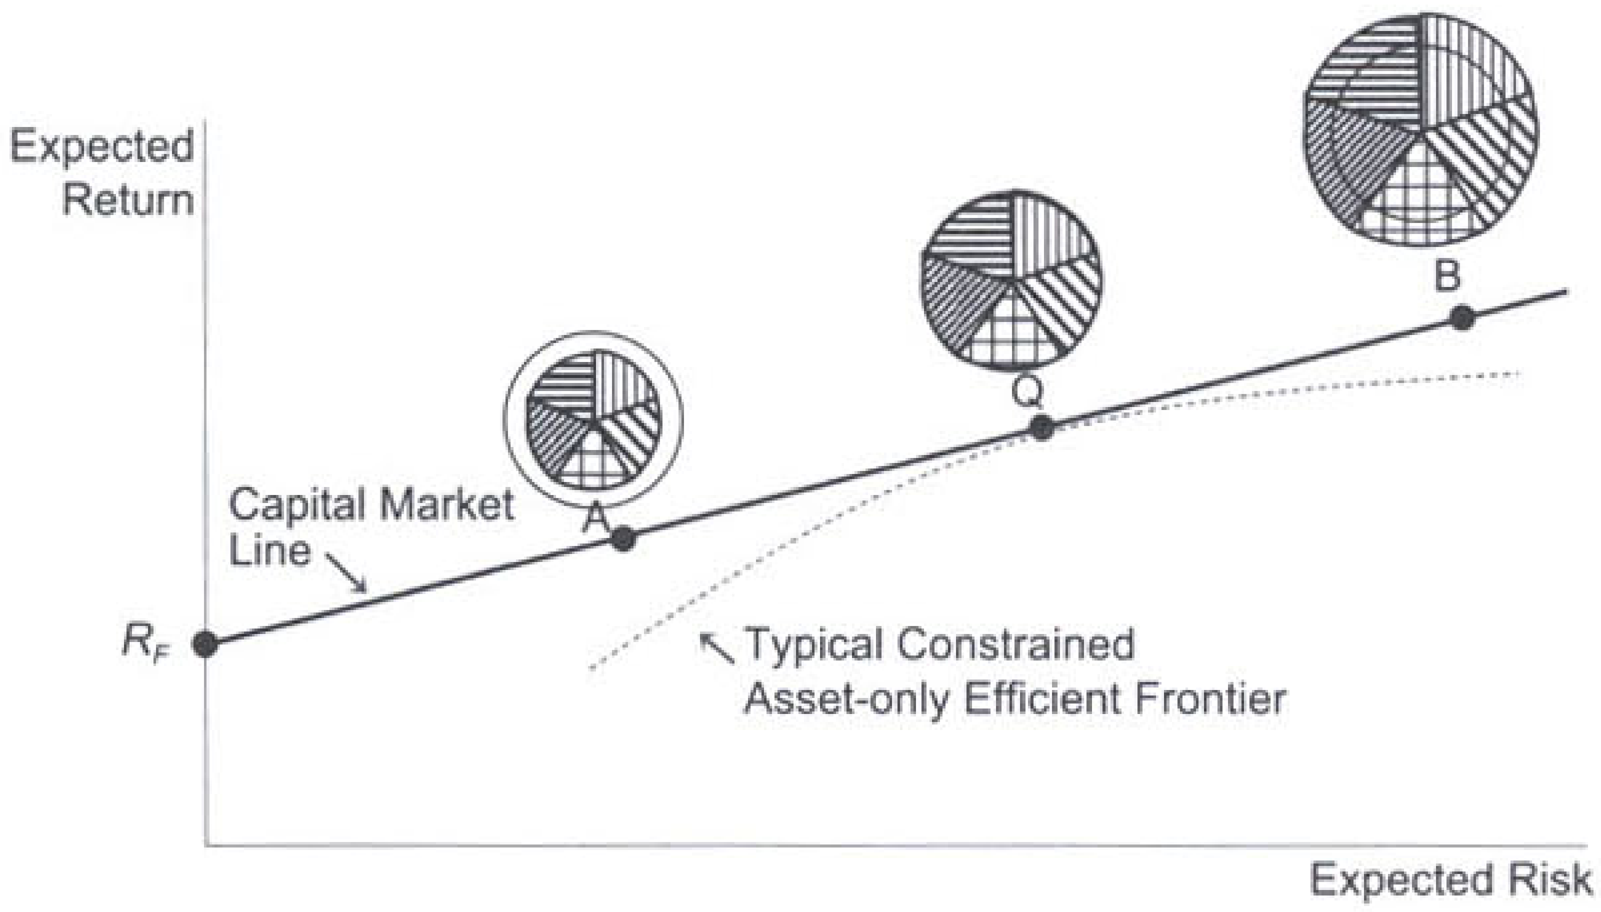

risk and return of the risky asset portfolio; each point on this line represents a particular weighting of the two funds.[4] This line is known in the literature as the capital market line, or CML. Although it required explanation originally, it is now widely understood and appreciated as a valid intellectual construct by all who study financial economics and parallels to it are found in nearly all models.

The mechanics of the Two-Fund Theorem are usually articulated in the form of a budget constraint that always has the total of cash and the market summing to 100%, as follows. Starting with 100% in the risky asset portfolio, if the investor desires a lower-risk position to the left of the world market portfolio (at point A in Exhibit 1), the investor delevers by lending cash (buying default-free bills or bonds) and holds only the remainder of the portfolio in risky assets. Alternatively, if the investor desires to *increase* leverage to achieve a higher-risk position to the right of the world portfolio (at point B in Exhibit 1), the investor borrows cash and uses it to purchase more than 100% in risky assets.

A more general articulation would be in risk-premium form—an investor always holds 100% of the portfolio in cash plus any desired exposure to the risk premia of the risky asset portfolio (no cash needed for risk premia).

The elegance of this theorem is that *the investor needs to make only one decision*: the level of aversion to investment risk. Once decided, risk aversion directly maps to a particular point on the ideal efficient frontier, and thus to an optimal mix of cash and risky assets—a policy portfolio. Strategic asset allocation policy? One decision, and *done*. What could be simpler?

The Two-Fund portfolios that lie on the CML also share the property of being completely described by their (single-factor) CAPM beta relative to the risky asset portfolio. (An asset portfolio with a beta of 0.5, for example, has only half of its assets invested in the world market portfolio of all risky assets, which we designate as Q.) We make that point in Exhibit 1 by showing the "pie" of asset class weights as being the same, just larger or smaller than the investable cash value, as beta changes. The "slices," or asset classes (more formally, the asset classes can be generalized as factors for a multi-factor model representing portfolio Q), each have a constant (market-weighted) percentage of Q, regardless of the overall beta.

Most of the efficient frontiers used today by practitioners are still asset-only frontiers rather than surplus frontiers and, as such, do not include the economic liability or any other liability. Moreover, as illustrated by the dotted line in Exhibit 1, constraints—such as "no cash or small-cap" and "international no greater than 20%"—are nearly always placed on the optimization and thus prevent the frontier from taking the basic linear shape of the CML. This reflects the discomfort of strategists with the Two-Fund Theorem's recommendation that a portion of the portfolio be held in the form of cash; the result of these constraints is that the portfolio holds the market portfolio of bonds rather than cash, a more desirable result in their eyes.

In the constrained world of curved efficient frontiers, the relative proportions of asset classes *vary* with risk level from those in portfolio Q, with conservative portfolios holding a differently weighted mix than aggressive portfolios, a result at odds with the constant market-weighted percentages found in the Two-Fund Theorem. Given that many of the constraints used have little theoretical grounding, it is a surprise that such results are widely accepted as being "correct" by many or most strategists. This practice has done much to cloud the understanding of institutional and retail investors of investment policy, thus motivating us to find a "better" two-fund version of policy development for use by practitioners.

CURRENT METHODS FOR INCORPORATING THE LIABILITY: SURPLUS OPTIMIZATION

The most commonly accepted, academically supported, and practical framework for incorporating the liability is *surplus optimization*. This is a variant of the traditional form of asset-only utility used under the CAPM in which the liability is included as an asset held short. The optimal set of investment choices for any investor that is subject to a liability is determined by optimizing a mean-variance *surplus* utility function,[5]

$$\max_{\beta_A} \; U_S = R_{f(S)} + \beta_S \mu_Q - \lambda \beta_S^2 \sigma_Q^2 \qquad (1)$$

U_S is the utility of the surplus; $R_{f(S)}$ is the net or surplus risk-free rate over the period of the single-period optimization; β_S is the investor's surplus beta exposure; and μ_Q is the expected equilibrium risk premium of portfolio Q (i.e., $\mu_Q = R_Q - R_f$, the total return of Q less R_f). Q's variance is σ_Q^2. The optimal risk level is a function of the investor's aversion to risk expressed in terms of the variable lambda, λ.

Equation (1) is expressed in a succinct and compressed "surplus" form; we can unpack it a bit to aid the reader. We can expand surplus beta, $\beta_S = (\beta_A - \frac{L_0}{A_0}\beta_L)$, which is the net or difference in betas between the assets and the liabilities, with the liabilities being corrected for their size relative to the assets. And we can expand the surplus risk-free rate, likewise a net of the risk-free rates of the assets and the liabilities, $R_{f(S)} = (R_{f(A)} - \frac{L_0}{A_0} R_{f(L)})$. The variables A, L, S, and Q indicate the assets, liability, surplus, and world risky asset portfolio, respectively, whether used as variables, subscripts, or parenthetical notes to those variables. The subscript 0 indicates beginning-of-period values. Also, although we show the source of each risk-free rate element in a parenthetical with A, L, or S, as appropriate, the rate itself is, of course, the same regardless of source.

We can make the following substitutions so as to visualize surplus utility more completely:

$$\max_{\beta_A} \; U_S = \left(R_{f(A)} - \frac{L_0}{A_0} R_{f(L)} \right) + \left(\beta_A - \frac{L_0}{A_0}\beta_L \right)\mu_Q - \lambda \left(\beta_A - \frac{L_0}{A_0}\beta_L \right)^2 \sigma_Q^2 \qquad (2)$$

Equation (2) reflects an optimization problem having as its objective function the maximization of the rate of return, or growth, of the surplus (or, if L is greater than A, maximizing the rate of decrease of the deficit), while controlling the risk, or variance, of that rate of return.[6]

Why Should Investors Incorporate Their Economic Liability When Developing Investment Policy?

A plan to spend later rather than now is, after all, the reason for holding investments in the first place, and the first objective of those investments must be to cover those future spending plans and perhaps provide the means to increase them. This is simply deferred consumption, a classic Econ 101 trade-off: People are choosing to consume goods and services tomorrow instead of today. This trade-off is a key premise of this article. It is convenient to say that investors have an economic liability, whether explicit or implicit, representing their planned future consumption.

The value of the liability is the present value of deferred consumption, or, in more specific terms, the present value of expected future consumption; we will use these terms interchangeably. For pension plans, it is the present value of expected future benefit payments (a broader measure than just the accumulated benefit obligation, or ABO). For foundations and endowments, it is the present value of future spending. For individuals, it is the present value of expected future spending during the period of their maximum possible life, a fixed-term annuity (or, if available, the value of a risk-free life annuity).

In fact, it is fair to make a statement of this relationship as an important identity: Wealth *is* the present value of deferred future consumption, from the tautology that the value of an investor's economic assets necessarily equals the present value of future consumption. As Rubenstein [1976] so nicely put it, wealth is defined in terms of the claims to future consumption that it *endows*.

Don't be misdirected, however; liability-relative investing is still about how to invest the *assets*. The present value of expected future consumption is included in the optimization, but it comes in a fixed value and with fixed factor, or market risk, characteristics—having some correlations to the assets, to be sure. The size and character of the liability is not subject to change from the optimization. In fact, because the liability's characteristics are "givens," the liability becomes a special sort of benchmark that has to be matched or beaten by the assets, revealing the risk and return benefits of hedging the liability. The liability becomes a central part of the investment policy construct.

The Investor's Economic Balance Sheet: Can the Investor Be Over- or Underfunded?

In the classic economic balance sheet, assets equal liabilities plus owner's equity. For an individual investor, one can imagine a liability (the present value of deferred consumption) that is smaller than the assets, and thus one can imagine a concept of "owner's equity" and "surplus" to represent assets not planned for personal consumption (i.e., overfunded). We can move that "owner's equity" into the liability category, if we treat it as the present value of an investor's final deferred consumption expenditure, or gifts and bequests to heirs. In this general sense, the assets must equal the liabilities at all times.

After all, what does it really mean, economically, for an investor to be in deficit? Or in surplus? An investor cannot owe to herself more, or less, than she has. The concept of over- or underfunding requires us to accept that an investor is dealing with only some of the elements of that balance sheet and not others (as happens routinely for DB retirement plans).[7]

And we can also imagine a present value of planned future consumption, again for an individual person, that is greater than that person's current assets—an aspirational, but not currently funded, plan for future spending that, in turn, leaves room to imagine a "deficit." But if we were to provide room in the asset column for economic assets, however, such as the present value of expected future savings from labor income, then it is difficult to believe that this view would be very useful, as any deficit from spending plans that have a present value greater than available wealth is just so much wishful thinking. For a foundation or endowment, these additions would include the present value of expected future contributions from fund-raising activity; for individuals, the present value of expected future savings from labor income; and for a DB pension plan, the present value of expected future contributions.[8]

Further, imagine segregating the liability side of the balance sheet into different buckets representing, for example: 1) a survival level of consumption, 2) an additional amount needed to maintain the standard of living to which the investor has become accustomed during his working life, 3) an additional amount that would provide for a luxury level of consumption, and 4) another amount for desirable planned bequests. The investor could declare that he is in surplus or deficit relative to any or all of these amounts. But, regardless, any portion of this layer cake that, at the end of the day, exceeds the economic assets simply vanishes as so much wishful thinking. So it would add little to the analysis to consider an investor as being over- or underfunded. You've got what you've got! (An analysis of the investor's potential level of deferred consumption might, however, be relevant to the investor's risk aversion, according to Merton [2006a].)

Of the common types of investors, only DB pension plans appear at first blush to be capable of being under- or overfunded; actually this is only in the benefit-security sense of segregating specific physical assets held in trust to be set against a portion of the present value of expected future benefit payments, the accrued portion of that liability.[9] But, in truth, there is much more to the economic balance sheet of pension plans, and they also are always in balance (Waring [2009]).

For purposes of this article, it is necessary to limit the discussion, keeping this in its simplest form, and not attempting to optimize the full economic balance sheet. We will assume that individuals have no labor income, and thus no present value of future savings or additional present value of future consumption associated with such future savings, nor will we consider the house or the expected bequest from Uncle Dan.[10] Ditto for foundations and endowments—we'll assume no present value for future contributions or other revenue sources, nor spending plans conditional on them. We'll assume that the ABO is the benefit-security measure of the DB plan liability, avoiding discussion of the full economic liability or of the integrated balance sheet. The path to including these matters is straightforward, however, and may merit a follow-up paper.

Despite a bias that economic assets must equal economic liabilities, we have continued to include the customary inequality term in our mathematics, L_0/A_0, for the convenience of readers who desire to treat DB plans or other investors as if they are over- or underfunded, and we'll show how this affects optimal solutions. The effects are quite different and much simpler than they are often described.

FIRST, A *THREE*-FUND THEOREM FOR INVESTORS WITH A LIABILITY

Let's start by using a CAPM of conventional, asset-only, single-period derivation, optimizing our portfolio using the surplus utility variant shown in Equation (2). And because the world is messy, let's use the academics' technique of keeping the function artificially cleaner than it really is (for the moment) by assuming that the factor model of the liability and thus also of the liability-matching asset portfolio are both on the capital market line and there are no residuals (returns not linear with beta) from any source. This simplification lets us focus on easily understood linear surplus-optimal solutions, which have the convenient property that they are all completely describable with the CAPM beta, nicely consistent with the classic Two-Fund Theorem.

In Part I of Appendix A, we derive the first-order conditions for this single-factor surplus utility function. The optimal beta solution for an investor with a liability is described in Equation (A2). The equation is in linear form (intercept-slope), with risk aversion, λ, as the explanatory variable and the ideal or optimal asset beta, β^*_A, as the dependent variable,

$$\beta^*_A = \underbrace{\left(\frac{L_0}{A_0}\beta_L\right)}_{\text{Intercept}} + \underbrace{\left(\frac{\mu_Q}{2\sigma^2_Q}\right)}_{\text{Slope}}\frac{1}{\lambda} \tag{3}$$

The terms in parentheses are constants. The risk aversion term, λ, is a variable specific to the investor. The slope term, $\mu_Q / 2\sigma^2_Q$, defines the angle of the line relating λ to the optimal asset beta. An asterisk indicates that a given value is optimal (in this case for beta, but later in the article for returns and risks).

Liability-Matching Asset Portfolio

One of the first questions to answer is how to find the asset beta that provides the lowest surplus-risk level that can possibly be achieved. To do this assume that the investor's aversion to risk, λ, is so large that when multiplied by the slope term the product becomes effectively zero, in which case a portion of the equation "goes away" in a numerical sense. In this event, the optimal asset beta reduces to its minimum value, the beta of the liability. We'll call this the liability-matching asset portfolio, or *LMAP*, $\beta^*_A = \beta^*_{A(LMAP)} = \frac{L_0}{A_0}\beta_L$.

When holding this portfolio, the asset betas and the liability betas cancel out, and the optimal surplus beta is zero, $\beta^*_S = \beta^*_{A(LMAP)} - \frac{L_0}{A_0}\beta_L = 0$. Because both the risk premium and risk are direct functions of beta, this would mean that no surplus risk whatsoever would exist in this portfolio.[11] Therefore, we can also refer to the liability-matching asset portfolio as the minimum surplus variance portfolio (MSVP).[12]

Holding Equities: The Investor's "Risky Asset Portfolio"

Most investors will have some lesser degree of aversion to investment risk than that just described. These investors will want to hold some amount of exposure to equities and to other elements of the risky asset portfolio in the hope of trying to increase the amount of future consumption that can be supported by their assets. Note, however, that the liability-matching asset portfolio will still be present (and will still be, at least notionally, the same size as the liability) no matter what value is assigned to the risk aversion term, λ.

So, we can firmly conclude that the fund, $\beta^*_{A(LMAP)}$, must *always* be held by *every* investor with a liability and, in addition, the investor will hold whatever amount is desired of the risky asset portfolio.

Thus this investor will always hold three funds: 1) the liability-matching asset portfolio of real interest rate exposures (for most investors; and for DB pension plans, some exposure to the inflation rate), 2) the single-period risk-free rate portfolio, and 3) if the sponsor has a tolerance for the risk involved, an additional exposure to the risky asset portfolio of risk premia (i.e., portfolio Q, generating surplus beta, β_S). The second and third "funds" are identical to the two funds of the Tobin-Sharpe construct.

The information provided in Exhibits 2 and 3 illustrate the geometry of the optimal solution. In Exhibit 2, the liability (L) is placed at a random but relatively low risk position situated directly on a *negative* capital market line, which represents the other side of the balance sheet. The liability-matching asset portfolio, *LMAP*, is placed on the positive, or asset, CML, located a bit further to the right than the liability itself in recognition that many pension plan sponsors anticipate treating the liability as larger than the assets (the *L/A* ratio is greater than unity); otherwise, the *LMAP* would be directly above the liability.

To generate the surplus expected return equation for the optimal portfolio in three funds, we substitute our solution back into the return components of utility from Equation (2),

$$R^*_S = \left(R_{f(A)} - \frac{L_0}{A_0}R_{f(L)}\right) + \left(\underbrace{\left(\frac{L_0}{A_0}\beta_L + \frac{\mu_Q}{2\lambda\sigma^2_Q}\right)}_{\beta^*_A = \beta^*_{A(LMAP)} + \beta^*_{A(RAP)}} - \frac{L_0}{A_0}\beta_L\right)\mu_Q \tag{4}$$

Exhibit 2
Include the Liability

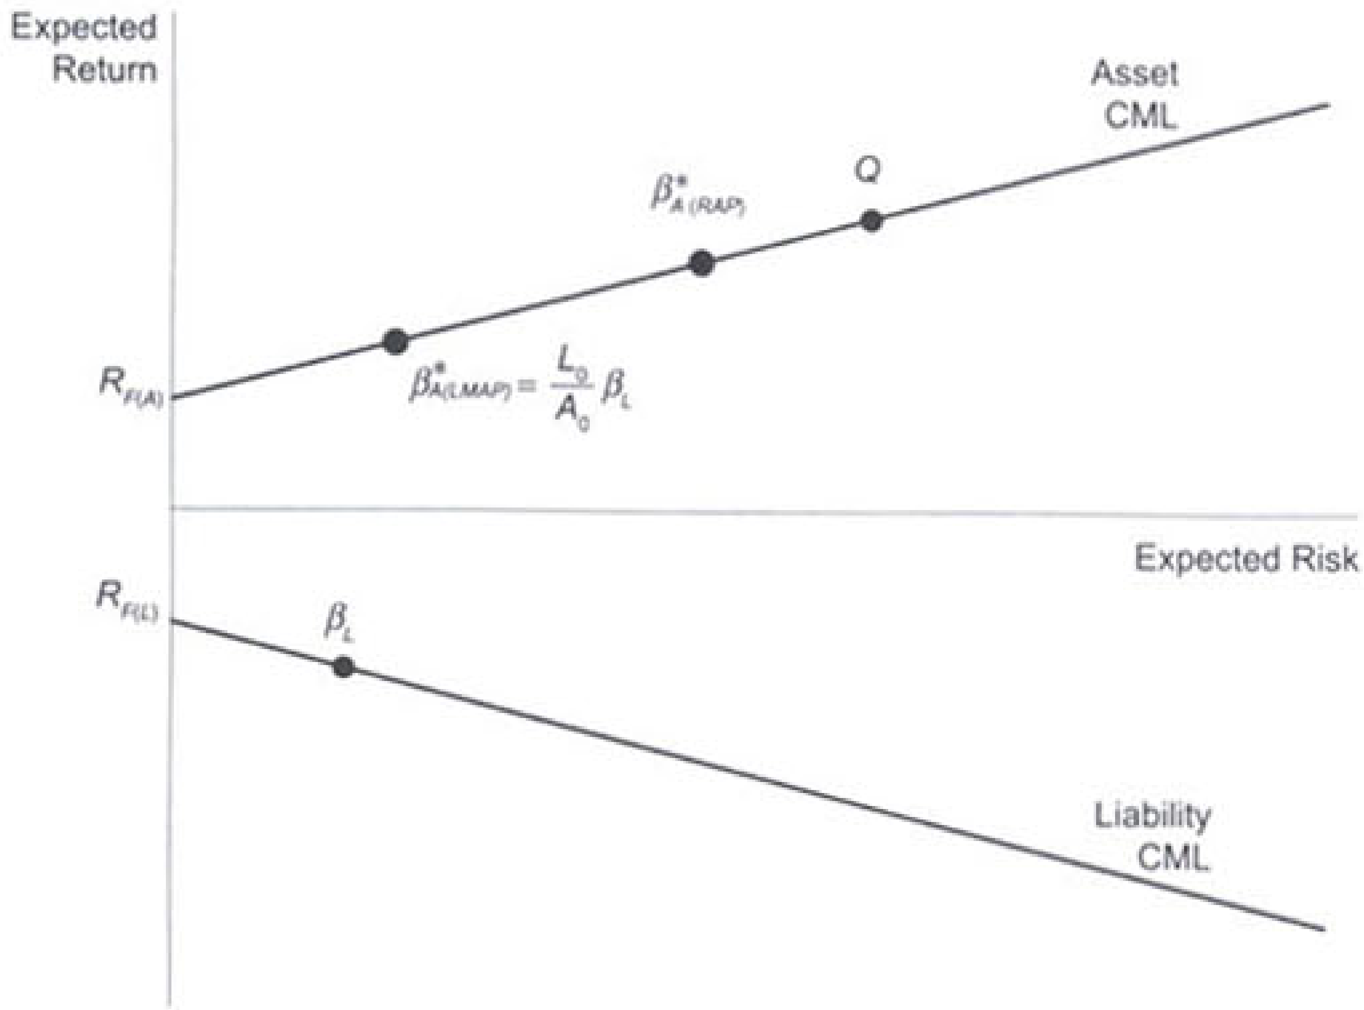

Exhibit 3
Three Fund Theorem

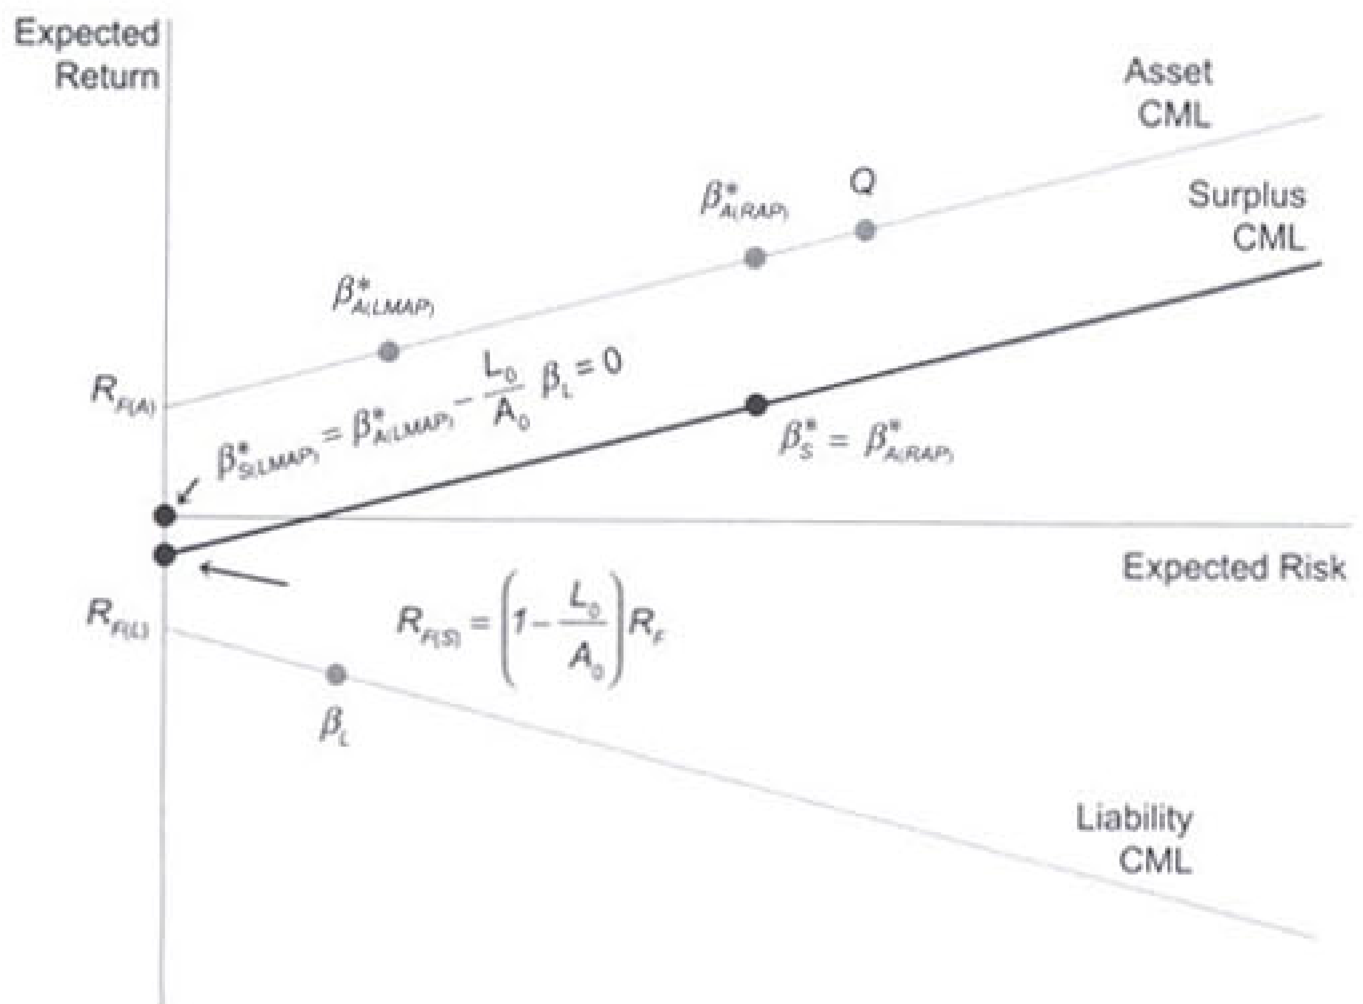

then decompose it into its *LMAP* and risky asset portfolio (*RAP*) components, the latter being its beta exposure to Q,[13]

$$R_S^* = \underbrace{R_{f(A)} - \frac{L_0}{A_0}R_{f(L)}}_{\substack{\text{The Surplus Risk-Free Rate}\\ R_{f(S)}}} + \underbrace{\left(\beta^*_{A(LMAP)} - \frac{L_0}{A_0}\beta_L\right)}_{\substack{\text{The Surplus } LMAP\\ \beta^*_{S(LMAP)}=0}}\mu_Q + \underbrace{(\beta^*_{A(RAP)})\mu_Q}_{\text{The } RAP} \tag{5}$$

In Exhibit 3 we complete the picture. The first point to identify for a surplus investment policy is the single-period net risk-free rate for the surplus, $R_{f(S)} = (R_{f(A)} - \frac{L_0}{A_0}R_{f(L)})$. This rate is slightly negative as is expected for a slightly underfunded DB pension plan. Any underfunding means that the sponsor is paying implied interest on the shortfall (which would be a zero term under our general assumption that $A = L$).

Next, the beta of the liability-matching asset portfolio and the liability beta cancel each other out, $\beta^*_{S(LMAP)} = (\beta^*_{A(LMAP)} - \frac{L_0}{A_0}\beta_L) = 0$, so that there is no surplus beta exposure—neither return nor risk—from this source. This "surplus *LMAP*" can then be plotted at the origin.

The span of choices for the remainder, the nonzero portion of optimal surplus beta representing the risky asset portfolio, $\beta^*_{A(RAP)}$, forms a capital market line as expected. But because the return to beta is only a risk premium rather than a total return, the total surplus return and risk position is in the new position anchored to the net risk-free rate and is shown as the "Surplus CML" described by Equation (5). This is the surplus efficient frontier that sets out the true strategic asset allocation policy opportunity set for an investor (one having a liability on the CML), and it is on this line that the total return surplus optimal solution, β^*_S, will sit.

Once the liability has been matched—the first fund is key—surplus optimization reduces to be just like an asset-only problem in the sense that the decision about how much of the risky asset portfolio and how much of the risk-free asset to hold is made without further regard to liability considerations. In other words, with the liability matched, the balance of the solution is just as in the Two-Fund Theorem.

Making This More General: A Liability Not on the CML

At the beginning of this section we made the unrealistic assumption that the liability portfolio was *on* the CML—which is only possible if all of its factor-beta exposures, or asset class factors, are proportional to those of the risky asset portfolio Q (i.e., if we can describe it accurately with a single factor-beta). It would in fact be rare for this to happen. Liability portfolios are random relative to efficient portfolios and are not expected to be on the CML.

For our first approach, it was sufficient to describe the risky asset portfolio with the traditional single factor-beta. But to adequately allow inefficient *interior* portfolios, we need to describe the portfolio and the liability in a more granular, multiple-asset-class or multi-factor model. **F** will be a unique vector, or list, of asset class weights or factors (more precisely, this approach will work with a set of multiple factor-betas from any market risk model, of which asset class weights are just one very basic example).

We also need to establish notation for describing the time dimension of the liability, and later of the assets, to incorporate the spot curve of real yields and future cash flows across the total period of planned consumption. We could set up mathematical machinery to represent an overlay of these longitudinal periodic descriptors of the liability, on top of the cross-sectional asset class weights, but this notational Rubik's Cube is unnecessary because we can accomplish much the same purpose more simply. Rather than detailing the cash flows period by period, they can instead be satisfactorily described in terms of certain summary characteristics, their real interest rate duration (and for DB plans, their inflation duration), a form that compresses a great deal of information from the underlying cash flows into just those two values. The use of durations is an especially convenient shortcut for our purpose because durations can equivalently and more generally be thought of as real interest-rate and inflation factor-betas. Durations have associated returns that can stand in for yields and, in that sense, are not unlike our asset-class-weight vector and may, in fact, be included as elements within it.[14]

For individuals, foundations, and endowments, the liability is most simply modeled solely in real interest rate terms—a complex inflation-protected coupon super-bond—because the key goal of these funds is to protect consumption power (similar to Rubenstein [1976]). Again, this is a pattern of deferred consumption that is also describable in summary form in terms of factor-beta exposures—in this case a $\mathbf{F}_{\mathrm{L}}$ vector consisting of only one element, the real interest rate duration of the deferred consumption ladder. (An equivalently simple model of a DB pension liability would include both the real interest rate and inflation durations of the liability as the only nonzero elements of $\mathbf{F}_{\mathrm{L}}$.)[15]

More detailed models of pension liabilities than this are possible, of course, and might pick up useful factor-betas for equity risks, yield curve twist, convexity, and other exposures, but real interest rates and occasionally inflation rates appear to have most of the explanatory power; further, real and nominal risk-free bonds have the valuable property that prices move inversely to yield, thus protecting consumption power when held.

With this setup, we can show how a flexible multi-asset-class or multi-factor optimization of a liability that is not on the CML varies in its optimal solution from the outcome of Equation (3). The optimal solution vector (exclusive of the single-period risk-free asset) is described in Part II of Appendix A, Equation (A5), repeated here,

$$\mathbf{F}_{\mathrm{A}}^{*} = \underbrace{\frac{L_0}{A_0}\mathbf{F}_{\mathrm{L}}}_{\substack{\text{The } \textit{Liability-Matching} \\ \textit{Asset Portfolio} \\ \mathbf{F}^{*}_{\mathrm{A}(LMAP)}}} + \underbrace{\left(\frac{\mathbf{V}_{\mathrm{Q}}^{-1}\mathbf{r}_{\mathrm{Q}}}{2}\right)^{\mathrm{T}}\frac{1}{\lambda}}_{\substack{\mathbf{F}^{*}_{\mathrm{A}(RAP)}\text{:} \\ \text{the } \textit{Risky Asset} \\ \textit{Portfolio}}} \tag{6}$$

$\mathbf{F}_{\mathrm{A}}^{*}$ is the optimal vector of asset class exposures, $\mathbf{V}_{\mathrm{Q}}^{-1}$ is the inverse of the variance-covariance matrix of the asset classes, $\mathbf{r}_{\mathrm{Q}}$ is the vector describing the expected returns of each of these factors; and $\mathbf{T}$ is the transpose operator.

Notice that this solution is completely parallel in form to Equation (3) other than being expressed in a multi-factor vector-matrix form in place of the single factor-beta form. This means that the interpretation is also parallel to that of the more constrained version; the intercept term is still readily interpretable as the liability-matching asset portfolio (and as the minimum surplus variance portfolio) because it is the only term that remains when the investor's risk aversion is very high.

And when risk aversion is not so high that the first term becomes a zero term, the investor holds (in addition to the liability-matching asset portfolio) an amount of exposure to a term (in parentheses) that is clearly the vector-matrix analog to $\mu_{\mathrm{Q}}/2\lambda\sigma_{\mathrm{Q}}^{2}$ term of Equation (3), the single factor-beta exposure of the investor's risky asset portfolio, $\beta^{*}_{A(RAP)}\mathbf{F}_{\mathrm{Q}}$.

To generate the optimal surplus return for an investor with this unconstrained liability, we again substitute our solution back into the return components of utility from Equation (A3) and again decompose it into its *LMAP* and risky asset portfolio (*RAP*) components,

$$R_{\mathrm{S}}^{*} = \left(R_{f(A)} - \frac{L_0}{A_0}R_{f(L)}\right) + \left(\underbrace{\frac{L_0}{A_0}\mathbf{F}_{\mathrm{L}} + \left(\frac{\mathbf{V}_{\mathrm{Q}}^{-1}\mathbf{r}_{\mathrm{Q}}}{2\lambda}\right)^{\mathrm{T}}}_{\mathbf{F}^{*}_{\mathrm{A}}=\mathbf{F}^{*}_{\mathrm{A}(LMAP)}+\mathbf{F}^{*}_{\mathrm{A}(RAP)}} - \frac{L_0}{A_0}\mathbf{F}_{\mathrm{L}}\right)\mathbf{r}_{\mathrm{Q}} \tag{7}$$

$$R_S^* = \underbrace{R_{f(A)} - \frac{L_0}{A_0} R_{f(L)}}_{\substack{\text{The Surplus Risk-Free Rate} \\ R_{f(S)}}} + \underbrace{\left(\mathbf{F}^*_{\mathbf{A}(LMAP)} - \frac{L_0}{A_0} \mathbf{F_L} \right)}_{\substack{\text{The surplus } LMAP \\ \mathbf{F}^*_{\mathbf{S}(LMAP)}=0}} \mathbf{r_Q} + \underbrace{\mathbf{F}^*_{\mathbf{S}(RAP)}}_{\substack{\text{The surplus } RAP, \\ =\mathbf{F}^*_{\mathbf{A}(RAP)} \\ =\beta^*_{A(RAP)}\mathbf{F_Q}}} \mathbf{r_Q} \tag{8}$$

Equation (8) shows the solution in three funds, parallel to Equation (5): In addition to holding the single-period risk-free surplus asset, *it is always optimal to hold the liability-matching asset portfolio, regardless of whether the liability is on the* CML. And it is always *optional* to hold an additional risky asset portfolio exposed to the rest of the markets. So an investor holding a liability needs, at minimum, a liability-matching asset portfolio, but may also want a portfolio containing risky assets.

Exhibit 4 illustrates the nature of this more realistic liability-matching asset portfolio. Because the liability model for an investor, $\mathbf{F_L}$, is an interior portfolio consisting of, say, only TIPS of a 15-year duration (representing the full vector of spending plan cashflows), the optimal liability-matching asset portfolio is also an interior portfolio, $(L_0/A_0)\mathbf{F_L}$, consisting simply of the same exposure, perhaps levered up or down in duration to reflect a belief that the assets aren't as large or as small as the liability.

In fact, because no decision really needs to be made about holding the liability-matching asset portfolio (the *LMAP* is independent of risk aversion or other preferences[16]), the only strategic asset allocation policy decision that the investor needs to consider is how aggressive or conservative to make surplus beta. In other words, how averse is the investor to risk—what is his or her lambda?—and thus how large or small is the risky asset portfolio?

Thus, we return to the very same interesting generality derived from the Tobin-Sharpe construct: Once a liability-matching asset portfolio is held, the rest of the strategic asset allocation policy decision is made just as it is for an asset-only policy under the classic Two-Fund Theorem.

THE THREE-FUND THEOREM LEADS TO A NEW VERSION OF THE TWO-FUND THEOREM AND THE CAPM

Our three-fund theorem is as good as we can make it using a conventional CAPM as a base, and at first blush, seems eminently workable and suitable. In fact, it would not detract from other common, current practices if investors were to use it as a primary tool of policy development: It would substantially improve policy solutions simply by clarifying the purpose, nature, size, and role of the *LMAP* in hedging consumption and spending power; for this reason, we are willing to encourage its use by practitioners.

But let's take a closer look and decide whether we can improve it further. There is a problem—if all investors were to hold an *LMAP*, then all investors could not also simultaneously hold the conventional risky asset portfolio (in the aggregate), because, as the risky asset portfolio is conventionally defined, a portion of the latter portfolio (real bonds at a minimum) would already be held in the investors' *LMAP*s. The sum of all investments can be no greater than the market or we face a double-counting inconsistency.

This problem constitutes a critique, but not of any attempt to incorporate the liability into the CAPM.

EXHIBIT 4

Three-Fund Theorem with a Liability That Has Residual Beta Exposures

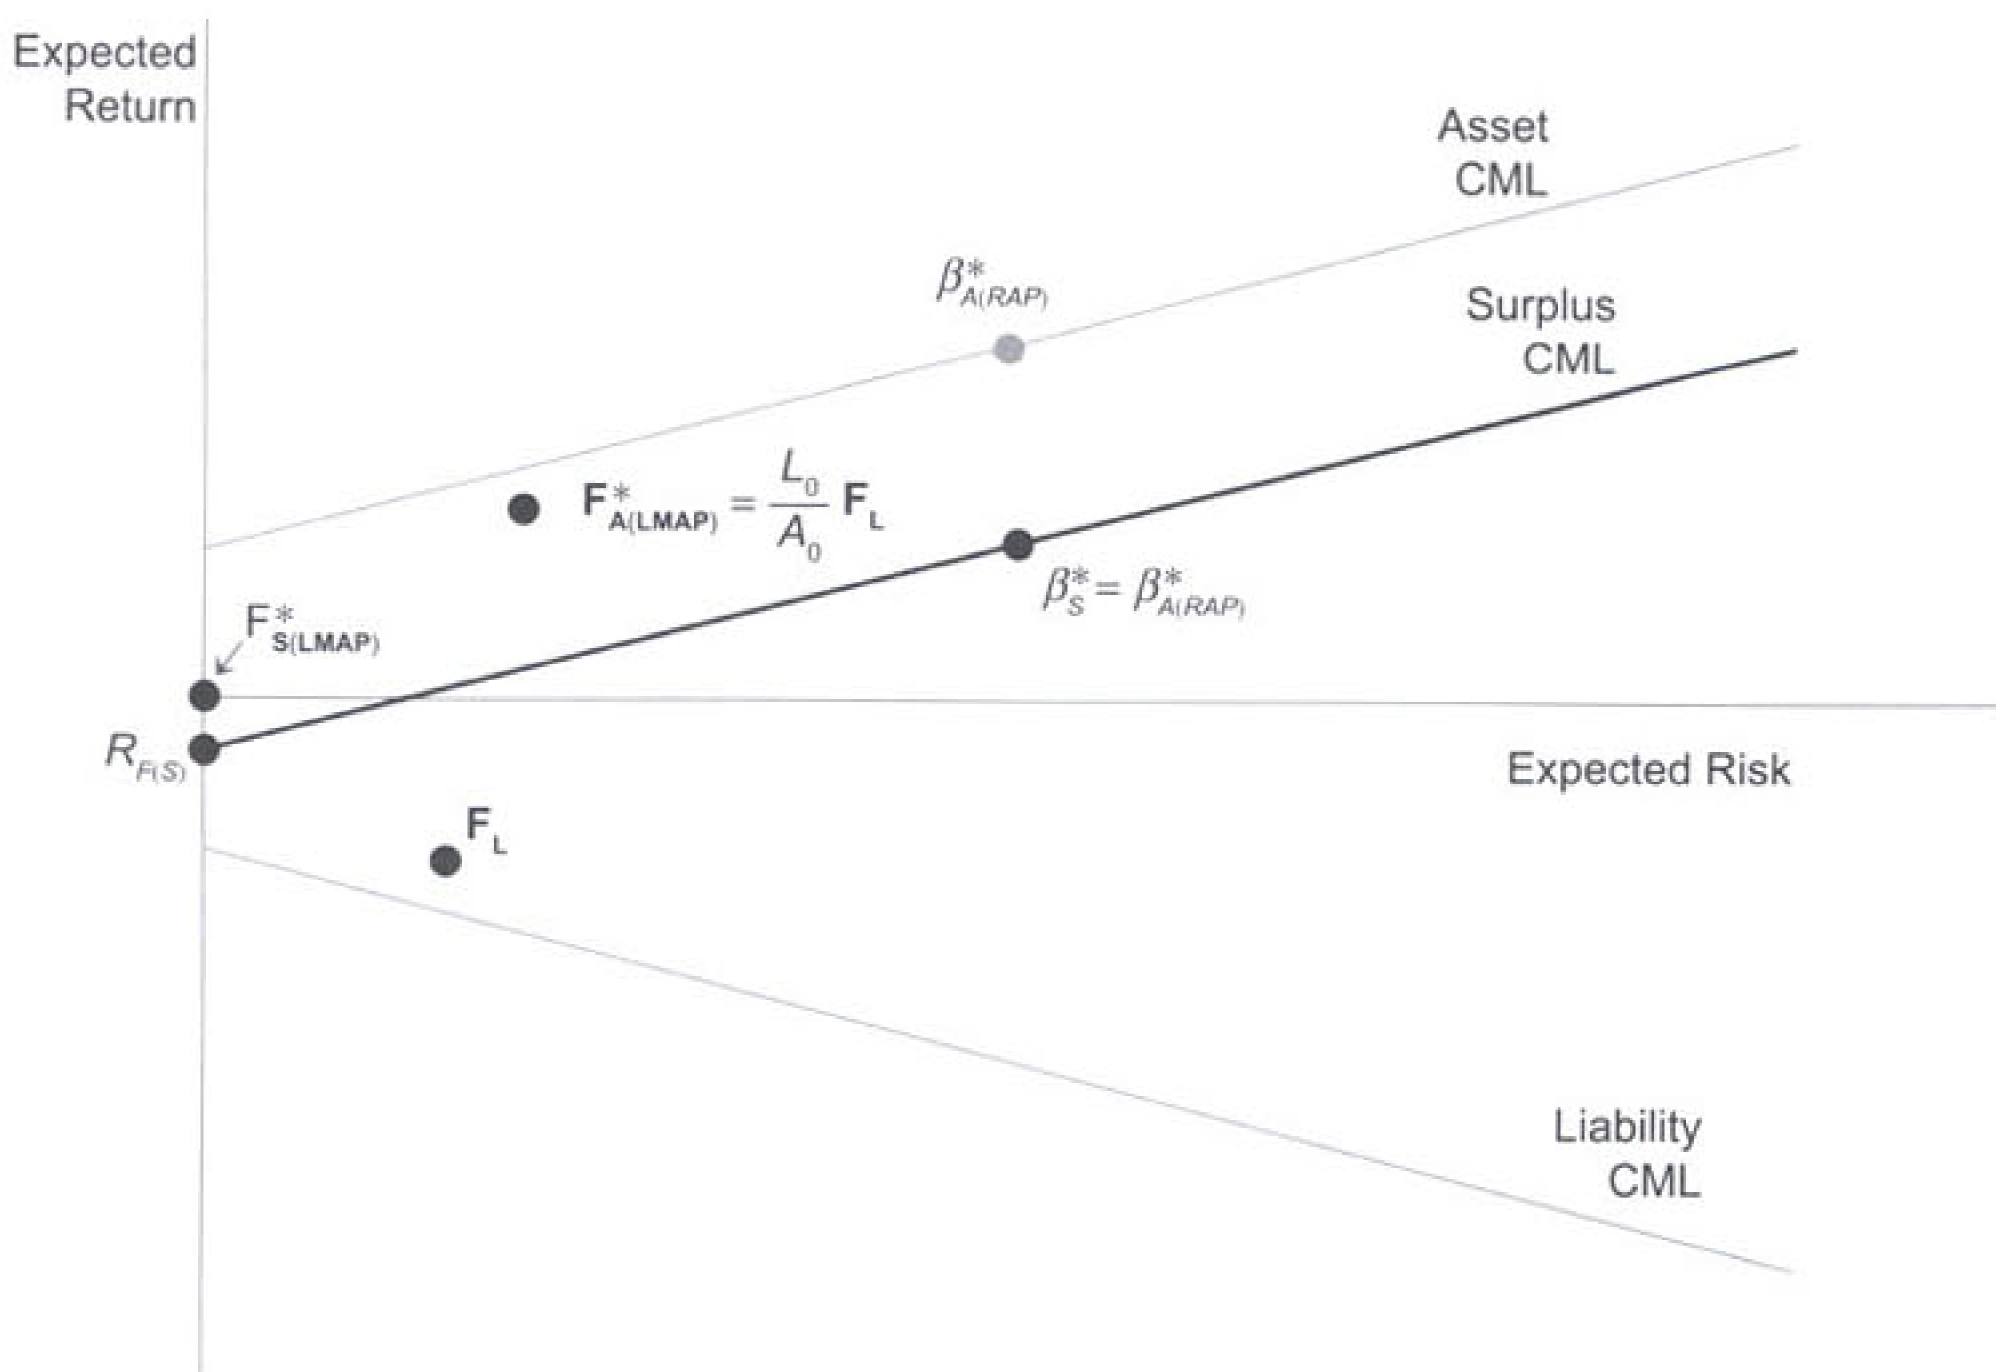

Rather, it is a critique of the underlying specification of the risk-free asset and of the risky asset portfolio. To resolve this conflict we have to move from a single-period focus to a multi-period focus and reconsider what we mean when we refer to risk-free assets.

The Case for Revisiting the Nature of the Risk-Free Asset

The academy has put much effort over many years into establishing a richly varied body of equilibrium theory for risky portfolio and asset prices, but relatively little effort into establishing an equilibrium theory for riskless assets and integrating such theory into more general pricing theory. This lapse of attention isn't trivial because it is difficult—if not impossible—to correctly develop the theory for risky asset pricing if the theory of riskless assets that underlies it is weak.

In his famous CAPM paper, Sharpe [1964] simply referred to the "pure rate of interest," without further explanation or in-depth discussion, as the rate of return on the risk-free asset. But it isn't really fair to single out Sharpe because his treatment of the risk-free asset is common to *all* the writers in that formative era, including those who are often credited, along with Sharpe, as having co-equally developed the capital asset pricing model. For example, Treynor [1962, 1999], Lintner [1965], and Mossin [1966] each did something similarly perfunctory; Mossin's quaint turn-of-phrase was that the pure rate of interest is the "price for waiting" as opposed to the "price of risk." A decade later, Ross' [1976] arbitrage pricing theory used a single-period risk-free rate, and Merton's [1973] intertemporal model incorporated an instantaneous risk-free asset for determining risk premia. With few exceptions, virtually all papers that discussed variations of asset pricing models included some variation of the same thing, *to this date*.

Yet as early as 1966, Modigliani and Sutch were articulating a view that has been widely accepted as far as it goes, but has yet to be fully reconciled with the CAPM—namely, the *preferred habitat theory* of interest rates, which is in fact the seed for a more complete theory of risk-free assets. This theory suggests that investors choose to own bonds with cash flows that have the same timing as their deferred spending needs in order to reduce or eliminate risk. To paraphrase Modigliani and Sutch [1966], a 10-year bond is riskless to an investor with an obligation 10 years out, but holding cash is very risky. Since then, this sentiment has been frequently echoed. Campbell and Viceira [2002] emphasized the same theory in their compact and wonderfully written book, asserting that "[a]t a theoretical level, these points have been understood for many years" (p. 6). Yet, it remains *nearly* universal to use some arbitrarily short "cash" rate as the risk-free rate when discussing various forms of the conventional single-period CAPM, and it seems to be *completely* universal when specifying risk premia. The situation is somewhat improved in the intertemporal CAPM literature, as Merton [1975, 1977] described his hedging portfolio to include a multi-horizon Modigliani-Sutch–inspired component, but did not treat it as part of the risk-free rate. Ditto for Ross' arbitrage pricing theory—it can include a multi-period portfolio of bonds, but doesn't re-examine the risk-free rate. The views of Modigliani and Sutch [1966] and the *LMAP*'s time dimension suggest a path for developing a better theory of risk-free assets and risk premia, and for resolving the problem shown in the three-fund theorem.

From the perspective of each asset owner, their investments exist simply for the primary, and perhaps sole, purpose of providing for, or "funding," their deferred future consumption—or the investor's "liabilities." And if this is true for each investor, it is also true for the overall market, which is simply the aggregate across all investors. In Appendix B, we use such an aggregate approach to derive a simple model for the pricing of risk relative to this aggregate liability, an equilibrium basis for a variant of the CAPM.

Pricing Risk Considering Investors' Aggregate Balance Sheet

Equation (B5) in Appendix B describes optimal asset-class or beta-factor holdings for the *i*th individual investor's portfolio (subscript *i* indexes unique investors, and *m* represents the aggregate market),

$$\mathbf{F}^*_{\mathbf{A},i} = \underbrace{\mathbf{F}^*_{\mathbf{A}(LMAP),i}}_{LMAP} + \beta^*_{A(RAP),i}\underbrace{(\mathbf{F}_{\mathbf{A},m} - \mathbf{F}_{\mathbf{L},m})}_{RAP} \tag{9}$$

Simplified, the investor should optimally hold a portfolio of factor-betas (asset classes) that matches the unique liability of the investor, the *LMAP*, plus a risky asset portfolio of returns in excess of the aggregate liability portfolio. In parallel, for the market as a whole,

$$\mathbf{F}_{\mathbf{A},m} = \mathbf{F}_{\mathbf{A}(LMAP),m} + (\mathbf{F}_{\mathbf{A},m} - \mathbf{F}_{\mathbf{L},m}) \tag{10}$$

The single-factor risky asset beta of the market, $\beta_{A(RA),m}$, is generally taken as unity by construction, and so is omitted.

These portfolios are similar to, yet importantly different from, those implied by traditional CAPM models. The parenthetical term redefines the risky asset portfolio from its classic form, Q, which is the portfolio of all assets less R_f. In contrast, the new risky asset portfolio consists of all assets excluding the aggregate liability-matching asset portfolio or, in other words, excluding both R_f and the longer-duration liability-matching assets taken in the aggregate—the latter having traditionally been considered part of the *RAP*.

To avoid confusion with Q or $\mathbf{F_Q}$, we will refer to this new market portfolio of risky assets as Q', or more formally in terms of its asset class risk-premium exposures, $\mathbf{F}'_{\mathbf{Q}} = \mathbf{F}_{\mathbf{A},m} - \mathbf{F}_{\mathbf{L},m}$.[17] In this way, we consider risk-free real bonds of all horizons (and of appropriate weights for each investor, asset, or the market) to be the building blocks of a multi-horizon risk-free asset for each investor and for the market. Our redefined risky asset portfolio is still priced on the basis of its market risk as in conventional theory, but only on that portion of those risks that isn't hedging the liability side of investors' aggregate balance sheet. This statement constitutes a new interpretation of the risky asset portfolio and the risk premium.

Expected Returns

We can again substitute these optimal holdings into the return components of our surplus utility function, Equation (A3), to get the optimal surplus return of this new form of liability-sensitive investor,

$$R^*_{S,i} = \left(R_{f(A)} - \frac{L_{0,i}}{A_{0,i}} R_{f(L)} \right) + \left(\underbrace{\mathbf{F}^*_{\mathbf{A}(\boldsymbol{LMAP}),i} + \beta^*_{A(RAP),i}\mathbf{F}'_{\mathbf{Q}}}_{\mathbf{F}^*_{\mathbf{A}}} - \frac{L_{0,i}}{A_{0,i}}\mathbf{F}_{\mathbf{L},i} \right)\mathbf{r}_{\mathbf{Q}} \tag{11}$$

$$R^*_{S,i} = \underbrace{R_{f(A)} - \frac{L_{0,i}}{A_{0,i}} R_{f(L)}}_{\substack{\text{The Surplus Risk-Free Rate} \\ R_{f(S),i}}} + \underbrace{\left(\mathbf{F}^*_{\mathbf{A}(\boldsymbol{LMAP}),i} - \frac{L_{0,i}}{A_{0,i}}\mathbf{F}_{\mathbf{L},i} \right)}_{\substack{\text{The Surplus } LMAP \\ \mathbf{F}^*_{\mathbf{S}(\boldsymbol{LMAP})}=0}}\mathbf{r}_{\mathbf{Q}} + \underbrace{\beta^*_{A(RAP),i}\mathbf{F}'_{\mathbf{Q}}}_{\text{The } RAP}\mathbf{r}_{\mathbf{Q}} \tag{12}$$

We multiply through by $A_{0,i}$, changing the expression from surplus return to the change in surplus wealth, and then group and identify the optimal asset-return and liability-return components, as follows:

$$\begin{aligned} A_{0,i}R^*_{S,i} &= A_0 \underbrace{\left[R_{f(A)} + \mathbf{F}^*_{\mathbf{A}(\boldsymbol{LMAP}),i}\mathbf{r}_{\mathbf{Q}} + \beta^*_{A(RAP),i}\mathbf{F}'_{\mathbf{Q}}\mathbf{r}_{\mathbf{Q}} \right]}_{R^*_{A,i}} \\ &\quad - L_{0,i}\underbrace{\left(R_{f(L)} - \mathbf{F}_{\mathbf{L},i}\mathbf{r}_{\mathbf{Q}} \right)}_{R_{L,i}} \end{aligned} \tag{13}$$

where $\mathbf{r}_{\mathbf{Q}}$ is the vector of expected returns in excess of a single factor R_f from the original CAPM across all asset class or factor-betas (i.e., inclusive of both risk and horizon premia).

The expected optimal return of the asset portfolio, in isolation, in this liability-relative form of the CAPM is $R^*_{A,i} = R_{f(A)} + \mathbf{F}^*_{\mathbf{A}(\boldsymbol{LMAP}),i}\mathbf{r}_{\mathbf{Q}} + \beta^*_i\mathbf{F}'_{\mathbf{Q}}\mathbf{r}_{\mathbf{Q}}$. This is the sum of the returns for the three funds we have discussed—the traditional risk-free rate, the *LMAP* return, and the risky asset portfolio (*RAP*) return. But the risk premium of the risky asset portfolio is now explicitly redefined in terms of our multi-horizon risk premium, which is now defined as $\mu'_Q = \mathbf{F}'_{\mathbf{Q}}\mathbf{r}_{\mathbf{Q}}$.

So we have redefined the three-fund theorem in a manner that is macro-consistent and workable. Is there a way to make this more parallel to the original Two-Fund Theorem?

Is It Really Three "Funds" or Just Two? Combining R_f and the *LMAP* into a Single Fund

Because of the way we developed this model, starting with the notation and forms of the traditional single-period CAPM, the single-period risk-free rate was taken out of the *LMAP*, which is expressed as a risk premium (or horizon premium). But if we agree that the two together constitute the complete risk-free asset on the basis that the *LMAP*—at least the real interest rate portions—is required universally by all investors, then the total return for a particular investor is their sum,

$$R_{A(LMAP),i} = R_{f(A)} + \mathbf{F}_{\mathbf{A}(\boldsymbol{LMAP}),i}\mathbf{r}_{\mathbf{Q}} \tag{14}$$

To parallel traditional notation, we can define $R_{A(LAMP),i} = R'_{f,i}$ a version of the risk-free rate that incorporates the multi-horizon needs of the investor. Similarly, we can define the risk premium of the risky asset portfolio not as μ_Q, but as $\mu'_Q = \mathbf{F}'_\mathbf{Q}\mathbf{r}_\mathbf{Q}$. With this notation, we can express the optimal asset portfolio return extracted from Equation (13) in scalar algebra, as a CAPM variant, in clear parallel to the form of the classic CAPM,

$$R^*_{A,i} = R'_{f,i} + \beta^*_{A(RAP),i}\mu'_Q \tag{15}$$

Not only does this portfolio lie on the CML, but, by inspection of Equation (15), can be seen to lie also on the *security market line* (SML) discovered by Sharpe [1964] (if graphed in beta, β, versus expected return space, we would see a linear relationship between beta risk and expected return). We complete this as an asset pricing model by noting that if we have proven this latter point for the *optimal* asset portfolio, then we have proven it for all other portfolios and assets. For this conclusion we need only refer back to the same argument originally developed by Sharpe, the key to the CAPM: Like portfolios, assets must all lie on the SML and must also have expected returns proportional to their betas. Here, because there are as many risk-free rates as there are investors, there are as many CMLs and SMLs as there are investors and assets, but they all work with the same *RAP*.

While the asset portfolio must carry risk not only from the *RAP* but also from the *LMAP*, the *LMAP* risk disappears when matched to the liability. Thus, the total surplus risk if the investor holds this optimal asset portfolio would simply be

$$\sigma^{*2}_{S,i} = \beta^{*2}_{A(RAP),i}\mathbf{F}'_\mathbf{Q}\mathbf{V}_\mathbf{Q}\mathbf{F}'^{\mathrm{T}}_\mathbf{Q} = \beta^{*2}_{A(RAP),i}\sigma^{2'}_Q \tag{16}$$

What Will be the Experience of Our Investor?

If an investor with a long future of annual consumption flows followed this advice (an individual planning for retirement, for example), and actually established an *LMAP* that corresponded to the consumption plan "liability" (perhaps a ladder of zero-coupon TIPS), with no *RAP*, then the investor's consumption would remain absolutely level in real terms and at the same level as when the *LMAP* was established, a result of the unique property of these bonds that prices move perfectly inversely to returns. The derivative of both the rate of consumption and of the present value of expected future consumption, with respect to changes in the real interest rate curve, is of necessity zero; that is, if real interest rates decreased, increasing the liability, the *LMAP*'s value would also increase. The new higher value of the *LMAP* would support the identical periodic annuity payment for consumption as did the prior smaller *LMAP* with the higher real rate.

But if the investor also held a portfolio of risky assets, consumption would increase or decrease directly with the returns on that portfolio; the derivative of consumption (the *relative* proportion of wealth spent in each future period held constant) with respect to changes in portfolio value from risky asset returns is one.

Discussion and Implications

In this section, we discuss critical elements of the Two-Fund Theorem and noteworthy implications of its application.

Relationship to Full Intertemporal Asset Pricing Models

A large literature exists that describes other, more sophisticated intertemporal capital asset pricing models. In this literature it is common for the results to be expressed as two-fund, three-fund, or even n-fund theorems. The flagship of this literature is Merton's [1973] article, but there are many more articles of importance, by Merton and others.[18] The models that are presented in these articles have several strong virtues, which have made them the exclusive focus of a number of academics working today. These advantages include, among others, their ability to use native utility theory in place of the mean-variance approximation, their ability to explicitly provide for an investor's flow of future consumption needs, and their capability of dealing with an investor's need to hedge not only real interest rate risk over time, but other types of unique individual intertemporal risks as well. Unfortunately, these models have the disadvantages of being difficult, or in many cases impossible, to solve when applied to practical real-world problems and of being constructed in the somewhat-advanced mathematics of stochastic dynamic calculus—which is unfortunately impenetrable to most practitioners. Thus, these models have had little practical application despite their flexibility, elegance, and beauty.

The asset-liability version of the CAPM presented in this article is for all intents and purposes a simplified version of such an intertemporal model, one that combines many of the best features of the highly tractable, single-period, mean-variance, asset-only world of most practitioners today, and the more heady world of the Merton style of intertemporal model that uses utility theory and continuous-time mathematics. By remaining in a mean-variance framework, it is fully transparent and understandable by many, if not most, practitioners and is readily usable. If the single most important lesson of the full intertemporal model is that the investor can hedge future consumption over time, then this model also accomplishes that task. If additional important lessons address hedging other more unique risks faced by an investor (energy price risk for a utility company or for an airline, for example), then by extending this model to include the remainder of the organization's economic balance sheet beyond the investment assets already on hand, similar results can also be developed.

What this model *can't* do is to reflect the finer differences in the optimal portfolio that might come from use of the more sophisticated utility and from the smooth evaluation of continuous-time mathematics. We don't have any parallel to the Merton approach's ability to hold securities that hedge "wealth productivity shocks" or certain other risk exposures that may change the investment opportunity set; see Campbell and Viceira [2006] or consider the term structure of the risk–return trade-off (long-horizon serial correlation) shown by Campbell and Viceira [2005].

Are these shortfalls important? Yes, *but*—in either case, the primary *LMAP* portfolio will be essentially the same and the primary *RAP* portfolio will be the world market portfolio of risky assets (ex-*LMAP* assets). Other potential hedges will be missing, but in most cases they are relatively smaller in importance. It is hard not to conclude that the misnamed Pareto's Law, the 80/20 rule, is fully satisfied.

Guidance for a Sensible Spending Rule

The discussion in this article strongly implies that we determine our spending rule either with an explicit annuity (see Endnote 11) or by replicating an annuity-like payment stream. As a bonus, a replicated annuity gives the investor the option to hold a personal *LMAP* and to hold some risky assets. A replicated annuity assures that funds will last a lifetime by avoiding the risk of over-consumption during periods of low returns. The replicated annuity can be adjusted to front- or back-end loading on consumption, providing allowances for high late-life medical costs, a preference to spend more when relatively young than when relatively old, and other time preferences.

The key idea is that in each period the amount withdrawn from assets and consumed would not exceed the amount that would have been paid out and spent in the same period if, at the beginning of the period, the annuity holder had purchased a fairly priced level-payment (or any other desired shape) real annuity for her maximum possible remaining lifespan (for constant consumption, this provides about 3.0% of the portfolio value each year for a 40-year future when the real interest rate is 1%, and about 3.7% when the real rate is 2%). At the beginning of each successive period, the amounts of the payout and the consumption would be recomputed, given the then-available value of the asset portfolio after returns, whether good or bad.

If the investor holds as part of his portfolio an *LMAP* portfolio of real rate bonds that hedges the spending plan, the rate of consumption will be protected, as was just described. If the investor also holds a risky asset portfolio, then consumption can be ratcheted up (or ratcheted down) proportionally to the return of the risky asset portfolio each period.

Nonobvious Sources of Long-Horizon Risk-Free Returns

All asset returns, including risky assets, include an underlying term structure of multiple-horizon risk-free rates, R_f', embedded within their total return. Presumably these are weighted as a function of the size and timing of their expected future dividends (or other payouts to investors). This means that equities—and all other risky assets—have a multi-period horizon (that can also be summarized as duration) and that they are able to supply a portion of the *LMAP* portfolio $\mathbf{F}_{L,i}$ to the investor.[19] Because of the presence of these embedded long-horizon real riskless assets within all assets, including all equities and other risky assets, investors probably do have holdings that more closely approximate their *LMAP*s than we might otherwise think. To the extent that they do, this new two-fund theorem is a positive theory, not just a normative one.

An Indifference Proposition: The Present Value of Future Consumption Does not Change with Investment Policy

We have used a multi-period *risk-free* discount rate on the liability side of the balance sheet to determine the present value of the investor's expected future consumption. What if the investor is invested on the asset side in a portfolio including not just an *LMAP*, but also *risky* assets—should the discount rate for the liability also change?

Well, yes, and in fact it is axiomatic that both sides of the balance sheet, when properly displayed, should have the same discount rate. But changing the discount rate on the liability side would make no difference to either the present value of the liability or to the amount of the investor's wealth today: if the investor is invested in risky assets, for the same initial portfolio value, the expected level of future consumption will be higher than for an investor holding only the *LMAP*. But that higher expected level of future consumption on the liability side would be discounted back to present value using a discount rate including the risk premium from the risky assets on the other side, a higher discount rate—because it is, in fact, riskier. The bottom line is that the present value of future consumption would be the same regardless of investment policy.

What this means is that holdings of risky assets, $\mathbf{F}'_{A(RAP),t}$, *might* make an investor more wealthy and thus able to spend more, and will certainly do so if realized returns are close to, or above, their expected value. But the presence of risk in the risky asset portfolio means that returns of the assets may well be *less* than those of the *LMAP* alone and, in such a case, the realized level of future wealth, and thus of future consumption, will be *lower*, not higher, than if the investor had invested solely in the *LMAP*. Assuming that risk is fairly priced, the odds are exactly fair that the more aggressive policy will result in the investor having a *smaller* pool of assets from which to fund consumption as the investor having a *larger* pool. Taking on more risk does not create wealth in and of itself. Wealth creation (or destruction) depends on uncertain future *realizations*, not current *expectations*.

It is convenient to start investment planning by assuming that the investor wants his deferred consumption to be secure, which calls for use of the risk-free rate. But if the investor is willing to accept some level of uncertainty, then he can add risky assets to the portfolio, which increases the return and risk to both sides of the balance sheet. *Regardless, the present value of assets must equal the present value of liabilities and that value will not change with investment policy.*

So, contrary to common belief among investment professionals, holding a riskier portfolio with a greater expected return, rather than a more conservative portfolio, does not assure greater future consumption. Neither will a riskier portfolio assure the make-up of a shortfall between assets and a higher aspirational liability (for individuals or for endowments), nor a shortfall between assets and a DB retirement plan's liability.

Honoring the Traditional Budget Constraint: 100% Investment and No More

There is an apparent difference worth discussing between the two-fund theorem and the theorems of Tobin [1958] and Sharpe [1964]: the risk-free asset, or *LMAP*, is held in the full value of the assets, and the risky asset portfolio, which consists solely of risk premia, is held over and above the *LMAP*. This observation has caused some with whom we have discussed this article to be concerned that the optimal portfolio must sum to greater than 100%, a perhaps inappropriate degree of leverage.

The reconciliation lies in keeping the amount invested in the risk-free asset separate from the amount invested in risky assets. The first generates total returns, but the latter only generates risk premia. A portfolio of 100% cash is no more or less fully invested than a portfolio of 100% equities—but the equities implicitly represent 100% cash plus 100% in the equity risk premium. Viewed in this way, this solution is no different in its leverage characteristics than those of Tobin and Sharpe, and the portfolio is not more than 100% invested (see previous discussion).

CONCLUSION: IMPLICATIONS OF THE NEW TWO-FUND THEOREM

In this article, we have incorporated the liability side of the investor's balance sheet—the multi-period "economic liability" or present value of deferred future consumption—into our derivation of a new form for the CAPM. It turns out that a single-period form of the risk-free rate is simply not compatible with the multi-period form of most expected future consumption. These liabilities are the reason for which such investments are held, and therefore all investments represent future consumption possibilities, deferred in a multi-period context.

In a macro-consistent world, the *aggregate* risk-free asset, that is, the aggregate liability-matching asset portfolio, must be equal in value and character to the *aggregate* economic liability or deferred consumption plan, and must match the liability in its key financial (beta-factor) characteristics. Such a risk-free asset must of necessity have multiple horizons over time—a term structure—rather than a single-period short horizon of the risk-free asset of the conventional CAPM. And if the risk-free asset must have a term structure, then the risk premium must reflect it.

This gives us a useful iteration of the CAPM, incorporating an old theory of risk-free returns, that of Modigliani and Sutch [1966]. Our two funds are no longer the same as the two funds of Tobin and Sharpe, in that our specification of both the risk-free asset and of the risk premium is now multi-period to match our specification of investors' deferred consumption plans. And our *LMAP* portfolio of consumption-hedging, but otherwise risky, bonds is clearly priced on a different basis than are other "risky" assets.

This form of the two-fund theorem will be much more comfortable than the original Two-Fund Theorem to those strategists who never could bring themselves to divide the portfolio between single-period cash and risky assets, a solution that while elegant never seemed to fit the real world of investor needs. A risk-free asset consisting of an array of liability-cash-flow–matching risk-free real bonds will be much more intuitively acceptable. Strategists can comfortably return to a consistent usage of the world market portfolio of risky assets, using it as the base method of exposing portfolios to the market across asset classes, regardless of risk tolerance.

Tomorrow's strategic asset allocation policies need to be much better than today's, for all investors. Today, we spend a great deal of time on distractions, such as justifying policy allocations to hedge funds and "exotic asset classes," when there is an investment strategy elephant in the room that we are only slowly beginning to acknowledge. That elephant, the economic liability, needs to be attended to. If attention is not paid to it, investors of all types are focusing on the wrong issues and carrying excess and uncompensated risks. Surplus asset allocation in the context of a multi-period risk-free rate is the tool of choice for doing so.

Granted, this multi-period CAPM may be more difficult to parameterize than the original. A multi-period risk-free rate forces us to deal with the term structure, or at minimum, the duration, of our client investors and of the market as a whole, perhaps a more difficult task than in a single-period context. Planned future consumption for investors may also be difficult to estimate, but relatively level future spending plans will likely achieve acceptance as the default. In both cases, it may turn out that it is sufficient to simply estimate the real rate duration of future spending, rather than the exact term structure, as we do in this article.

A more complete solution will include more off-book economic assets, such as labor income, human capital, current consumption, and future savings, for individuals' retirement planning, and future gifts and contributions, as well as the economic value of assets, for pensions, foundations, and endowments. Of course, all the criticisms and weaknesses that have been asserted against the CAPM apply equally to this form of that approach. The fact is, however, that all capital asset pricing models, including the sophisticated intertemporal models, can benefit from including the present value of a liability expressed over time and tied in value to the available assets, as well as from writing budget constraints more carefully so as only to tie the cash position to investment values, not risk premia. Doing so will give these models an improved theory of the role of the risk-free asset and of the structure of the risk premium.

A practically implementable CAPM needs to be as simple as possible—*but not one bit simpler.* It is too simple when based on single-period cash without regard to the liabilities of investors. With the bit of complexity we have added, we hope that investors will conclude that this two-fund theorem is intuitive and practical, and therefore usable, improving strategic asset allocation practice.

APPENDIX A

Deriving the Three-Fund Theorem for Surplus Utility

Part I: Taking the utility function from Equation (2) in the text, we evaluate its first derivative with respect to the single-factor asset beta. This tells us how utility changes when beta changes,

$$\frac{\delta U_S}{\delta \beta_A} = \mu_Q - 2\lambda \beta_A \sigma_Q^2 + 2\lambda \frac{L_0}{A_0} \beta_L \sigma_Q^2 \tag{A1}$$

Setting this equal to zero, and manipulating, we learn that the solution for the utility-maximizing, or *optimal* beta, β_A^*, in the presence of a liability is

$$\beta_A^* = \left(\frac{L_0}{A_0} \beta_L \right) + \left(\frac{\mu_Q}{2\sigma_Q^2} \right) \frac{1}{\lambda} \tag{A2}$$

Part II: When discussing liability models that are *not* on the capital market line, we have to discuss asset class weights (factor-beta exposures) that are weighted differently than market weights (i.e., they are not proportional to portfolio Q).

Most liability models, being largely or solely combinations of long risk-free real and sometimes nominal bonds, have asset class weights (mixes of beta factors) that are not at all proportional to those in Q. In such a case, both the liability model and the *LMAP* will be off the CML and below it. In preparation, we can re-express Equation (2) in vector-matrix form, which allows us to move from the single-factor CAPM of Sharpe [1964] to a simple multi-asset-class (or multi-factor) model, so that the asset class weights can vary from those in Q,

$$\operatorname*{Max}_{\mathbf{F_A}} U_S = \underbrace{\left(R_{f(A)} - \frac{L_0}{A_0} R_{f(L)} \right) + \left(\mathbf{F_A} - \frac{L_0}{A_0} \mathbf{F_L} \right) \mathbf{r_Q}}_{\text{return components}}$$
$$\underbrace{- \lambda \left(\mathbf{F_A} - \frac{L_0}{A_0} \mathbf{F_L} \right) \mathbf{V_Q} \left(\mathbf{F_A} - \frac{L_0}{A_0} \mathbf{F_L} \right)^{\mathrm{T}}}_{\text{risk components}} \tag{A3}$$

Our asset class vectors, $\mathbf{F}$, are row vectors of weights for the various asset classes or other factor-beta exposures; all other vectors are columnar. Specifically, $\mathbf{F_Q}$ is the unique vector of asset class weights composing the market portfolio Q, and $\mathbf{F_A}$ and $\mathbf{F_L}$ are the weights of the asset portfolio and of the liability. $\mathbf{r_Q}$ is the set of equilibrium expected returns for these asset classes and $\mathbf{V_Q}$ is their variance-covariance matrix.

To get the optimal solution, we take the first derivative of surplus utility, Equation (A3), with respect to the vector of asset classes,

$$\frac{\delta U_S}{\delta \mathbf{F_A}} = \mathbf{r_Q} - 2\lambda \mathbf{V_Q} \mathbf{F_A^T} + 2\lambda \frac{L_0}{A_0} \mathbf{V_Q} \mathbf{F_L^T} \tag{A4}$$

And to solve for the optimal vector of asset class exposures, $\mathbf{F_A^*}$, we set the right side of the equation equal to zero, multiply through by the inverse of the covariance matrix, and rearrange to get a result that is very much parallel to Equation (A2),

$$\mathbf{F_A^*} = \underbrace{\frac{L_0}{A_0} \mathbf{F_L}}_{\substack{\text{The Liability-Matching} \\ \text{Asset Portfolio} \\ \mathbf{F^*_{A(\mathit{LMAP})}}}} + \underbrace{\left(\frac{\mathbf{V_Q^{-1} r_Q}}{2} \right)^{\mathrm{T}} \frac{1}{\lambda}}_{\substack{\mathbf{F^*_{A(\mathit{RAP})}}\text{:} \\ \text{the Risky Asset} \\ \text{Portfolio}}} \tag{A5}$$

Appendix B

A New Form of the Two-Fund Theorem

We know the optimal holdings for an individual investor from Equation (A5). But a single investor does not describe the whole market. The market, m, is the summation of all of these N optimal individual portfolios, including both the *LMAP* and risky asset portfolio portions (i.e., across Equation (A5)),

$$\mathbf{F^*_{A,\mathit{m}}} = \sum_{i=1}^{N} \left(\frac{L_{0,i}}{A_{0,i}} \mathbf{F}_{\mathbf{L},i} + \frac{(\mathbf{V_Q^{-1} r_Q})^{\mathrm{T}}}{2} \frac{1}{\lambda_i} \right) \tag{B1}$$

Because the market is simply the sum of all investors in it, we can move the first portion of the term in parentheses outside of the summation simply by replacing λ_i with the risk aversion of the market, λ_m, and we can drop the asterisk on $\mathbf{F^*_{A,\mathit{m}}}$ as the market portfolio does not need a reminder of its optimality. And the aggregate across all investors' *LMAPs*, the second term, is the *market* liability-matching asset portfolio, $\mathbf{F}_{\mathbf{L},m}$. The liability/asset ratio for the *aggregate* is unity, as every asset is someone else's liability. This gives us the market portfolio described in two components—a risky asset component and a liability-matching component,

$$\mathbf{F}_{\mathbf{A},m} = \mathbf{F}_{\mathbf{L},m} + \frac{(\mathbf{V_Q^{-1} r_Q})^{\mathrm{T}}}{2} \frac{1}{\lambda_m} \tag{B2}$$

The reason for aggregating optimal investment strategies across all investors is so that we can determine the market-clearing price of risk. We can solve for it by rearranging Equation (B2) as follows:

$$(\mathbf{V_Q^{-1} r_Q})^{\mathrm{T}} / 2 = (\mathbf{F}_{\mathbf{A},m} - \mathbf{F}_{\mathbf{L},m}) \lambda_m \tag{B3}$$

We substitute Equation (B3) into Equation (A5) to get a model of the optimal *individual* portfolio, but now using market equilibrium-risk pricing,

$$\mathbf{F}^*_{\mathbf{A},i} = \mathbf{F}^*_{\mathbf{A}(LMAP),i} + \frac{\lambda_m}{\lambda^*_i}(\mathbf{F}_{\mathbf{A},m} - \mathbf{F}_{\mathbf{L},m}) \tag{B4}$$

The ratio of the market and individual lambdas is the investor's single factor-beta for the risky asset portfolio. Therefore, we can make this look more familiar with a form parallel to the classic CAPM, but with a revised risk-free asset,

$$\mathbf{F}^*_{\mathbf{A},i} = \mathbf{F}^*_{\mathbf{A}(LMAP),i} + \beta^*_{A(RAP),i}(\mathbf{F}_{\mathbf{A},m} - \mathbf{F}_{\mathbf{L},m}) \tag{B5}$$

ENDNOTES

[1]See http://www.quotationspage.com/quote/2927.html, accessed on February 27, 2008.

[2]We would like to thank Laurence Siegel for his advice, comments, and suggestions, as well as Mary Brennan for her editing. And a special thanks to Paul Kaplan for not only suggesting that we challenge the notion that the liability-matching portfolio can be a third fund (after he reviewed an earlier version in which only the three-fund theorem was discussed), but for also suggesting later the important transformation in Equations (B3) and (B4), which is key to making this a fully general pricing model.

[3]In this reference to the "original" Two-Fund Theorem (which we'll capitalize), we're including the generally accepted body of knowledge that fleshed it out over the years. For example, Sharpe characterized the risky asset portfolio as consisting only of equities, but Roll [1977] showed why this portfolio should consist of *all* risky assets; we include such generally accepted incremental refinements.

[4]It relies on an assumption (among others) that the amount of cash in the portfolio is unconstrained; that is, cash can be borrowed or lent at its risk-free rate in any amount the investor chooses.

[5]For more detailed discussions of the derivation of surplus-optimization utility functions, and to see alternative formulations, see Waring [2004a, 2004b]. In general, these articles expanded on the basic premise of surplus optimization, integrating several ideas and derivations originally articulated separately by others (cited there, but particularly Leibowitz [1987] and Sharpe [1990]), and showed how to put the basic theoretical idea of surplus optimization into a practical and usable form. It described a total return (beta plus alpha) surplus efficient frontier using an objectively describable measure of the pension liability. For other approaches, see Campbell and Viceira [2002], who provided a summary of the literature of investment strategy technology, quite complete up to its date, with an emphasis on power utility of consumption (in both discrete and continuous time).

[6]By convention we generally refer to $A - L$ as the "surplus" even if it is a negative number (a deficit).

[7]To make an analogy to a pension plan, an individual's liability is *always* "fully funded" (i.e., with an L/A ratio of unity). Unfortunately, this happy state is not all that it might seem, as the implied benefit level is all too often too low to be comfortable.

[8]See Waring [2004a] for a discussion of U.S. equity real interest rate and inflation durations and of how one can incorporate them into the surplus-optimization policy development effort.

[9]Today the benefit security measure of the liability is generally a conventionally measured ABO, a very unsatisfactory measure for that purpose. See Waring [2009].

[10]This doesn't merely restrict use to individuals who are retired or otherwise have no labor income, as it is equivalent to an assumption that the present value of future consumption associated with labor income is perfectly correlated with the present value of future savings from labor income (and definitionally equal). Thus, there is a hedge with respect to the surplus, but not with respect to retirement consumption; it carries whatever risks are associated with the investor's labor income and thus his or her savings. To the extent that these are market-like, the strategist may want to take those risks into account as adjustments to the risky asset portfolios we develop later in the article. While an optimization of the *total* economic balance sheet would improve on this "Kentucky windage" approach, our simplification isn't a bad working assumption even for a young person with much more human capital than investment capital. For a complete discussion, see Bodie, Merton, and Samuelson [1992] and Merton [1993].

[11]Merton, likewise, suggests that a life annuity is the ideal risk-free asset for individuals [2005], but here we use the *maximum possible* life, and a real *fixed* annuity, rather than life *expectancy* and a real *life* annuity, because at the current time the issuer default risk for available commercial annuities seems high enough to discourage their use. Also, investing in annuities requires funding in cash, making it difficult and impractical to layer any risky asset exposures on top of the "portfolio."

[12]The comment that all risk from the liability goes away is confined to market-related risk capable of being hedged. Not all liability characteristics can be hedged. As discussed in Waring [2004b], "liability alphas" are those parts of the liability return that have risks not hedgeable through any financial instrument. These include mortality risks and other demographic uncertainties unrelated to market movements. Further, one person's inflation rate might well be different from that of another given differences in their consumption baskets (see, for example, Jennings [2006]). And this, in turn, means that one person's risk-free real bond ladder is different from that of another, not only in duration, but in its consumption-hedging dimension (no instruments yet exist to make this observation actionable). Because there are no investment solutions to these risks, we

don't address them here. The only solutions come in the form of plan design alternatives, that is, building liability structures that don't create as many unhedgeable uncertainties. Fortunately, it appears that the contribution of risk from unhedgeable demographic uncertainties is much smaller than the risk that is contributed by holding assets not well matched to the liabilities, the current norm. So if we fix the latter problem, we will have fixed the bigger part of the problem.

[13]In Equation (5) and in Equation (8) that follows, we conform to usual practice and consider the optimal *RAP* to be the portfolio's exposure to the market risky asset portfolio, as conventionally defined, in our notation $\beta^*_{A(RAP),i}$. In the discussion following Equation (10), we will show that this presumption needs refinement.

[14]We are aware that this simplification, using a single duration rather than the full span of consumption-weighted cash flows with their individual zero-coupon durations and yields, gives a higher yield than would the true underlying superbond, but at the price of missing important convexity and nonparallel-shift hedging requirements. This may be a fair trade-off for small investors with high trading costs, but larger investors will want to model their liability and implement their *LMAP* in the more granular form. For an excellent approach to building more detailed and complete models of pension liabilities, see Fabozzi, Martellini, and Priaulet [2006]. Also, if future consumption is in more than one currency, that fact should be reflected by the presence of real-rate instruments matching those cash flows in those currencies (or effective equivalents using hedges) when building the liability model that will include similar positions in the *LMAP* portfolio and thus in the optimal solution. This is similar to Solnik [1974], but without requiring an additional fund.

[15]In a paper that is closely related in subject matter, but that comes at the topic from a completely different direction, Hoevenaars et al. [2007] built a sophisticated structural econometric model of the key asset classes and concluded, consistent with our observation, that long-term real bonds should be favored by liability-relative investors.

[16]This is consistent with the results of Rudolf and Ziemba [2004] in a full intertemporal model (p. 976).

[17]This discussion of $\mathbf{F}'_Q$ clarifies the caveat made in Endnote 13.

[18]See Campbell and Viceira [2002] for a general overview of this literature. Merton has been prolific; in addition to his other articles cited elsewhere herein see [2006b]. Others of note include Breeden [1979], and Rubenstein [1976], as well as authors cited elsewhere in the article. Of these models, all take future consumption into account, but Merton [1993], Rudolf and Ziemba [2004], and Martellini [2006] do so using an explicit liability. The last paper is the closest to ours in its results.

[19]The durations of risky assets have historically been difficult to estimate, but with a stream of future cash flows and a discounted price today, this suggests there must indeed be price sensitivity to changes in the underlying risk-free rate. The duration of a risky asset is usually debated as if it were in nominal form, but Waring [2004a] and Siegel and Waring [2004], following Goodman and Marshall [1988], separated nominal inflation into real interest rate duration and inflation duration. These authors suggested that the inflation duration of equities is roughly zero and that the real interest rate duration is in the high teens. Today, we would argue that the real interest rate duration is somewhat lower—in the high single digits, say 8, taking into better account the correlations between the real interest rate in the denominator and the dividend stream in the numerator of the pricing equation.

REFERENCES

Bodie, Zvi, Robert C. Merton, and William F. Samuelson. "Labor Supply Flexibility and Portfolio Choice in a Life-Cycle Model." *Journal of Economic Dynamics and Control*, 16 (1992), pp. 427–449.

Breeden, Douglas T. "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities." *Journal of Financial Economics*, Vol. 7, No. 3 (September 1979), pp. 265–296.

Campbell, John Y., and Luis M. Viceira. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford, U.K.: Oxford University Press, 2002.

——. "The Term Structure of the Risk–Return Tradeoff." *Financial Analysts Journal*, Vol. 61, No. 1 (January/February 2005), pp. 34–44.

——. "Strategic Asset Allocation for Pension Plans." In *Oxford Handbook of Pensions and Retirement Income*, edited by Gordon L. Clark, Alicia H. Munnell, and J. Michael Orszag. Oxford, U.K.: Oxford University Press, 2006.

Evans, John L., and Stephen H. Archer. "Diversification and the Reduction of Dispersion: An Empirical Analysis." *Journal of Finance*, Vol. 23, No. 5 (December 1971), pp. 761–767.

Fabozzi, Frank J., Lionel Martellini, and Philippe Priaulet. "Hedging Interest Rate Risk with Term Structure Factor Models." In *Advanced Bond Portfolio Management*, Frank J. Fabozzi Series. Hoboken, NJ: John Wiley & Sons, Inc., 2006.

Goodman, Laurie S., and William J. Marshall. "Inflation, Interest Rates, and Pension Liabilities." In *Asset Allocation: A Handbook of Portfolio Policies, Strategies, and Tactics,* edited by Robert Arnott and Frank J. Fabozzi. Chicago: Probus, 1988.

Hoevenaars, Roy P.M.M, Roderick D.J. Molenaar, Peter C. Schotman, and Tom B.M. Steenkamp. "Strategic Asset Allocation with Liabilities: Beyond Stocks and Bonds." White paper, 2007. Available at SSRN, http://SSRN.com/abstract=675681.

Jennings, William. "Disaggregated TIPS: The Case for Disaggregating Inflation-Linked Bonds into Bonds Linked to Narrower CPI Components." *Journal of Pension Economics and Finance*, Vol. 5, No. 3 (2006), pp. 325–343.

Leibowitz, Martin L. "Asset Allocation through Surplus Management." *Financial Analysts Journal*, Vol. 43, No. 2 (March/April 1987), pp. 29–40.

Lintner, John. "The Valuation of Risky Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets." *The Review of Economics and Statistics*, Vol. 47, No. 1 (February 1965), pp. 13–37.

Martellini, Lionel. "Managing Pension Assets: From Surplus Optimization to Liability-Driven Investment." White paper, 2006. Accessed at http://www.edhec-risk.com/ALM/managing_pension_assets/index_html/attachments/Managing_Pension_Assets.pdf on February 13, 2008.

Merton, Robert C. "An Intertemporal Capital Asset Pricing Model." *Econometrica*, Vol. 41, No. 5 (September 1973), pp. 867–887.

——. "Theory of Finance from the Perspective of Continuous Time." *Journal of Financial and Quantitative Analysis*, 10 (November 1975), pp. 373–413.

——. "A Reexamination of the Capital Asset Pricing Model." In *Risk and Return in Finance*, edited by J. Bicksler and I. Friend. Vol. 1. New York: Harper Collins, 1977.

——. "Optimal Investment Strategies for University Endowment Funds." In *Studies of Supply and Demand in Higher Education*, edited by C.T. Clotfelter and M. Rothschild. Chicago: University of Chicago Press, 1993.

——. "Allocating Shareholder Capital to Pension Plans." *Journal of Applied Corporate Finance*, Vol. 18, No. 1 (2006a), pp. 15–24.

——. "Observations on Innovation in Pension Fund Management in the Impending Future." *PREA Quarterly* (2006b), pp. 61–67. Available at http://www.people.hbs.edu/rmerton/PREA%20Journal%20Article%202%2006.pdf.

Modigliani, Franco, and Richard Sutch. "Innovations in Interest Rate Policy." *American Economic Review*, 56 (May 1966), pp. 178–197.

——. "Debt Management and the Term Structure of Interest Rates: An Empirical Analysis of Recent Experience." *Journal of Political Economy*, Vol. 75, No. 4 (1967), pp. 569–589.

Mossin, Jan. "Equilibrium in a Capital Asset Market." *Econometrica*, Vol. 34, No. 4 (October 1966), pp. 768–783.

Roll, Richard. "A Critique of the Asset Pricing Theory's Tests, Part I: On Past and Potential Testability of the Theory." *The Journal of Financial Economics*, Vol. 4, No. 2 (March 1977), pp. 129–176.

Ross, Stephen A. "The Arbitrage Pricing Theory of Capital Asset Pricing." *Journal of Economic Theory*, Vol. 13, No. 3 (December 1976), pp. 341–360.

Rubenstein, Mark. "The Valuation of Uncertain Income Streams and the Pricing of Options." *The Bell Journal of Economics*, Vol. 7, No. 2 (Autumn 1976), pp. 407–425.

Rudolf, Marcus, and William T. Ziemba. "Intertemporal Surplus Management." *Journal of Economic Dynamics and Control*, 28 (2004), pp. 975–990.

Sharpe, William F. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *Journal of Finance*, Vol. 19, No. 3 (September 1964), pp. 425–442.

——. "Asset Allocation." In *Managing Investment Portfolios: A Dynamic Process*, edited by John L. Maginn and Donald L. Tuttle. New York: Warren, Gorham & Lamont, 1990.

——. "Equilibrium." In *Macro-Investment Analysis* available at http://www.stanford.edu/~wfsharpe/mia/eq/mia_eq1.htm. Accessed on February 24, 2008, and viewed previously over a period of several years.

Siegel, Laurence B., and M. Barton Waring. "TIPS, the Dual Duration, and the Pension Plan." *Financial Analysts Journal*, Vol. 60, No. 5 (September/October 2004), pp. 52–64.

Solnik, Bruno H. "An Equilibrium Model of the International Capital Markets." *Journal of Economic Theory*, 8 (1974), pp. 500–524.

Tobin, James. "Liquidity Preference as Behavior towards Risk." *The Review of Economic Studies*, Vol. 25, No. 2 (1958), pp. 65–86.

Treynor, Jack L. "Toward a Theory of Market Value of Risky Assets." In *Asset Pricing and Portfolio Performance*, edited by Robert A. Korajczyk. London: Risk Books, 1999. This otherwise unpublished paper has been widely circulated since its original writing in 1962.

Wagner, W.H., and S.C. Lau. "The Effect of Diversification on Risk." *Financial Analysts Journal*, Vol. 27, No. 6 (November/December 1971), pp. 48–53.

Waring, M. Barton. "Liability-Relative Investing I: Be Dual Duration Matched and on the Surplus Efficient Frontier." *The Journal of Portfolio Management*, Vol. 30, No. 4 (2004a), pp. 8–20.

——. "Liability-Relative Investing II: Surplus Optimization with Beta, Alpha, and an Economic View of the Liability." *The Journal of Portfolio Management*, Vol. 31, No. 1 (2004b), pp. 40–53.

——. "A Pension Rosetta Stone: Reconciling Actuarial Science and Pension Accounting with Economic Values." White paper, Chicago: Society of Actuaries (2009), forthcoming.

Waring, M. Barton, and Laurence B. Siegel. "Don't Kill the Golden Goose! Saving Pension Plans." *Financial Analysts Journal*, Vol. 63, No. 1 (January/February 2007), pp. 31–45.