A Mean-Variance Benchmark for Intertemporal Portfolio Theory

John H. Cochrane*

April 6, 2008

Abstract

By reinterpreting the symbols, one-period mean-variance portfolio theory can apply to dynamic intertemporal problems in incomplete markets, with non-marketed income. Investors first hedge non-traded income and preference shocks. Then, their optimal payoffs are split between an indexed perpetuity and a "long-run mean-variance efficient" payoff, which avoids variation over time as well as variation across states of nature. In equilibrium, the market payoff and the average outside-income hedge payoff span the long-run mean-variance frontier, and long-run expected returns are linear functions of long-run market and outside-income-hedge betas. State variables for investment opportunities and outside income are conveniently absent in these characterizations.

*Graduate School of Business, 5807 S. Woodlawn, Chicago IL 60637, 773 702 3059, john.cochrane@chicagogsb.edu, http://faculty.chicagogsb.edu/john.cochrane/research/Papers/. I acknowledge research support from CRSP and from a NSF grant administered by the NBER. I thank Lars Hansen, John Heaton and many seminar participants for helfpul comments.

1 Introduction

This paper studies long-horizon portfolio problems. I allow asset return dynamics, dynamic trading, non-market income, and preference shocks. Markets are incomplete, meaning that investors may not be able completely to hedge outside income or state-variable shocks. I mix three ingredients: First, I focus on the optimal stream of final payoffs or dividends, rather than focusing on the dynamic trading strategy that delivers those payoffs. Second, I characterize available asset payoffs by first constructing a discount factor or contingent claims price vector, following Cox and Huang (1989), or, ultimately, Arrow and Debreu (1954). Most importantly, I use the mean-variance approach that is so useful in one-period problems to characterize the dynamic problem with incomplete markets.

The basic idea is simply to treat time and probability symmetrically. I define an "expectation" that sums over time as well as states of nature, for example

$$\mathcal{E}(x) \equiv \frac{1-\beta}{\beta} E \sum_{t=1}^{\infty} \beta^t x_t. \tag{1}$$

Thus, I write the price of a security given a scaled discount factor m_t and a payoff (dividend and purchase or sale) stream x_t

$$p(x) = E \sum_{t=1}^{\infty} \beta^t m_t x_t = \frac{\beta}{1-\beta} \mathcal{E}(mx).$$

With this notation, we can transparently apply all of asset pricing and portfolio theory that naturally flows from $p = E(mx)$ in two-period models to a multi-period environment, by simply replacing E with $\mathcal{E}$ and reinterpreting the symbols.

Applying this idea to portfolio theory, the familiar ideas from two-period mean-variance analysis apply to the optimal payoff stream in a dynamic intertemporal setting. With no outside income, the investor obtains a payoff on the "long-run mean-variance frontier." The long-run frontier is defined using $\mathcal{E}$ and payoffs to a one dollar investment in the place of E and one-period returns. Long-run variance prizes stability over time as well as stability across states of nature.

If all investors are of this type, the market payoff, which is a claim to the aggregate consumption stream, is also on the long-run mean-variance frontier. Then, each investor's optimal payoff is a linear combination of an indexed perpetuity and the market payoff, weighted more to the former if the individual's risk aversion is small relative to average risk aversion and vice-versa. A long-run CAPM holds: each assets' long-run expected return is proportional to its long-run beta.

This characterization is particularly simple because state variables do not appear, even though there are time-varying investment opportunities and unspanned shocks. The dynamic *portfolio* that achieves a long-run mean-variance payoff may have such loadings. It may load more heavily on securities at *times* when they have high average returns, just as the one period portfolio loads more heavily on *securities* with higher average returns, and it may load on securities that can hedge such shifts in investment opportunities. But this variation in portfolio weights is just a means to the end, the long-run mean-variance efficient payoff, not the end itself.

When investors have outside income, we can define an outside-income hedge payoff by a long-run regression, i.e. using moments $\mathcal{E}$, of the outside-income stream on the set of available payoffs. Then, the investor wants a long-run mean-variance efficient "total payoff," composed of his asset market payoff and the hedge payoff he holds implicitly by his ownership of the outside income stream. Equivalently, he holds a short position in the hedge payoff, plus a long-run mean-variance efficient asset payoff, or he distorts his asset payoffs away from the long-run mean-variance efficient payoff in a way that that recognizes the riskfree, priced, and idiosyncratic components of his outside-income hedge payoff.

In a market of such investors, each individual holds a payoff that combines the market payoff, and a payoff consisting of the difference between his and the average outside-income hedge payoff. Investors with average outside income and risk aversion just hold the market, despite dynamics in returns and the presence of outside income. Investors with no outside income sell the average outside income hedge payoff. A long-run multifactor model emerges: long-run expected returns are proportional to long-run market betas and long-run betas of the asset payoff with respect to the average outside-income hedge payoff.

Again, this characterization is simple because of all the things that are absent. State variables for individual or aggregate outside income streams are absent, for the same reasons that state variables for investment opportunities are absent. Such state variables can matter in constructing the *portfolios* that support the hedge payoffs, but they do not matter in this characterization of the actual optimal payoffs. Outside income itself, of course, remains in the description of the optimal payoff, but in a simple way. We just add or subtract its hedge payoff, constructed by long-run regression. Background risk effects are absent. Agents who have to hold an unspanned component of outside income risk may, in general, act in a more risk-averse manner. This effect is absent in this long-run mean-variance setting as it is absent in one-period mean-variance settings.

Why is this interesting?

Dynamic portfolio theory is important. The fact that returns are not i.i.d. means that "long-run" and "short-run" investors may opt for different strategies, that asset demands may reflect hedging motives, and that investors may follow dynamic strategies to exploit time-variation as well as cross-sectional variation in asset return moments. These effects are potentially large.

Dynamic portfolio theory is hard. To cite one example, consider the question of how one should optimally invest in a stock index vs. a riskfree rate, given that index returns are predictable from variables such as dividend-price ratios. This classic simple Merton (1969, 1971a) problem has been attacked by Kim and Omberg (1996), Brennan, Schwartz and Lagnado (1997), Campbell and Viceira (1999), Barberis (2000), Brennan and Xia (2000, 2002), Wachter (2002), Sangvinatsos and Wachter (2005), Liu (2007), and many others. These papers are technically complex, and yet there is no closed-form solution for the simple benchmark case of power utility, an infinitely lived investor and imperfect hedging of the forecasting variable shock.

In this context, we can still very easily characterize optimal payoffs, as above, without deriving the *portfolio strategy* that supports those payoffs. The standard dynamic-programming approach must solve both problems at once, and the strategy problem is typically much harder. In addition, the dynamic-programming approach requires us to re-solve the whole problem every time a different asset is added to the mix.

Separating the *economic* characterization of the final payoffs from the *technical* or financial-engineering problem of constructing a portfolio that achieves those payoffs is really a central success of classical portfolio analysis. Though Markowitz (1952) derived the mean-variance frontier more than 50 years ago, we still have no settled way to actually *compute* that frontier. It is a difficult problem, resulting from the difficulty of estimating large covariance matrices, saying anything at all about mean returns, and dealing with parameter uncertainty and drift. Computation is approached differently for different asset classes, trading restrictions, data sets, time horizons, conditioning information sets, parameter knowledge, and all the other peculiarities of a given application. Much of the alpha-touting high-fee money management industry amounts to selling one or another solution to this problem. Classical mean-variance analysis brilliantly declares victory and goes home just before this hard part begins. My aim is simply to extend the same useful conceptual framework to intertemporal problems.

The final-payoff view of portfolio theory can also offer clearer intuition. For example, when one examines payoffs it is immediately obvious that a 10 year zero-coupon real bond is the riskless asset for an investor with a 10 year horizon, or that an indexed perpetuity is the riskless asset for an infinitely-lived investor. In the standard Merton (1969, 1971a) intertemporal approach we instead think of these as assets that just happen to hedge changes in their investment opportunity sets: long bond prices go up when interest rates go down. It takes work for Campbell and Viceira (2001) and Wachter (2003) to prove what seems obvious by inspection of the payoffs.

The final-payoff view of dynamic intertemporal portfolio theory, the analogy between static and dynamic optimization emphasized by my notation $\mathcal{E}(x)$ in (1), and the potential to characterize final payoffs quickly without solving for the dynamic strategy that supports them, are of course well known in finance since the work of Cox and Huang (1989), He and Pearson (1991), He and Pagès (1993), Schroder and Skiadas (1999), and many others. The theoretical advantages of this approach have not as yet been of widespread practical use in incomplete-market setups, however, because of the difficulty of finding which of the infinite number of potential discount factors results in a tradeable payoff. (Sangvinatsos and Wachter (2005) is a notable exception.) Restricting attention to linear marginal utility underlying mean-variance analysis avoids this central problem. Portfolio formation is a linear operation, so once we find a traded discount factor, which is on the long-run mean-variance frontier, we know it is the correct discount factor for the portfolio problem. Classical mean-variance analysis has always handled market incompleteness transparently, and my extension handles it transparently in the same way.

Expressing portfolio theory in terms of the market payoff is an especially useful discipline. Most portfolio theory seems to apply to everyone. But the average investor must hold the market, so any advice given to A must be mirrored in the opposite advice given to B. If not, the market equilibrium underlying the advice must change as soon as any measure of A investors take the advice. Investors can't all be smarter than average either. Phrasing portfolio advice in terms of *deviations* from the market portfolio, driven by *deviations* of the investor from average, along a described dimension of heterogeneity, helps to preserve portfolio advice from this Catch-22. It also may be of practical use. Investors may not be able to answer well "what's your risk aversion?" but they may be able to answer "are you more or less risk averse than the average investor?"

Incorporating outside income into portfolio theory is important. Real investors have houses, jobs, businesses, or other non-market income and assets. These investors should start by hedging non-marketed risks. They should short the zero-price, zero-premium payoff closest to non-marketed income, reducing overall risk *for free*. Then, they should adjust their asset payoffs to reflect the priced components of non-marketed income. The classic case that labor income looks like a bond moves the optimal asset portfolio to stocks, at least until retirement. Investors whose human capital is correlated with value returns should shun value stocks, despite their premium. Fama and French (1996) attribute the value premium to this effect.

Strangely, though it has been included almost from the start in portfolio theory, hedging non-market income is as rare as hedging state variables in practice. Steel workers, and their pension funds, do not short the Steel industry portfolio, or even the Auto industry. One would expect a class of money managers to have emerged that developed expertise in this hedging, understanding portfolios of traded assets that can hedge common sources of outside income, and selling these individual-specific tailored-portfolio services for fees. This has not happened. Academic research in asset pricing has focused almost entirely on finding "priced" factors, alphas for the one investor who has no outside income, and has ignored finding and characterizing *non-priced* factors that, by providing free insurance, are potentially the most important for typical investors.

Incorporating outside income is also hard, especially when markets are incomplete, and perhaps this is why it is overlooked. Outside income not likely to be i.i.d., so state variables need to be hedged. The optimal portfolio is peculiar to each investor, and we have to find asset portfolios that reliably hedge a useful fraction of outside-income volatility. Most of all, we can observe the *stream* of outside income, but we cannot easily observe its *value*. To apply standard return-based portfolio theory, one has to turn the income stream into a price. For example, Jagannathan and Wang (1996) assume an AR(1) labor income process and a constant discount rate, so that labor income growth measures human capital returns. However, price changes dominate high-frequency asset returns, and most asset price changes result from discount rate changes or changes in otherwise unobserved earnings forecasts. This procedure would not do a very good job of estimating stock returns from the stock dividend stream, so one may question whether it does a good job for labor income. (The AR(1) is not the only procedure of course. Campbell (1996) uses multivariate forecasts and assumes that the labor income stream is discounted at the stock expected return. Heaton and Lucas (2000b) and Davis, Kubler and Willen (2006) include labor income and variables that forecast labor income as Mertonian state variables in a dynamic portfolio theory, solved by dynamic programming. Still, fitting a labor income stream into a portfolio theory based on high-frequency returns remains difficult.)

These difficulties notwithstanding, a large literature addresses the incorporation of outside income or assets such as housing into portfolio theory. Mirroring the difficulties of dynamic portfolio theory, we do not have an analytic solution to the most basic problem, power utility, lognormal i.i.d. returns, and a lognormal diffusion for outside income. (Koo (1998) and Duffie, Fleming, Soner, and Zariphopoulou, (1997) characterize this problem.) A bit more progress has been made with CARA utility (Svensson and Werner (1993), Teplá (2000), Henderson (2005)), which alas is not much more realistic than quadratic. Most of the applied literature studies numerical solutions to particular calibrations, or approximations that require numerical evaluation in a calibrated environment to see any basic patterns. Highlights include Heaton and Lucas (2000a), (2000b), Davis and Willen (2000), Munk (2000), Lynch (2001), Viceira (2001), Flavin and Yamashita (2002), Yao and Zhang (2005), Lynch and Tan (2008), Benzoni, Collin-Dufresne, and Goldstein (2007). Even here, most attention is still devoted to the simple stock/bond and long/short bond split; characterization of which risky portfolios hedge the risky components of outside income is still in its infancy.

In this context, the simple long-run mean-variance benchmark can allow one to include non-market income very easily. We can compute the required long-run covariance matrix with no information on the values of non-market income streams, and we can construct a hedge payoff for labor income by the long-run analogue to a simple regression. One may hope for better performance: asset market *payoffs* or dividend streams may be better correlated with outside income *streams*, and that correlation easier to measure, than income growth rates or imputed values are correlated with asset returns.

Characterizing optimal dividend streams may eventually suggest a reason why stocks pay dividends and bonds pay coupons. Consumers could in principle synthesize such securities from dynamic trading of stocks that repurchase shares instead of paying dividends and zero-coupon bonds. But if dividends and coupons reflect the optimal final payouts many consumers desire, it may be less surprising that these are the basic marketed securities.

Of course mean-variance analysis is only a first step. We eventually want to know how the nonlinearities of marginal utility and non-normality of available returns affect optimal payoffs. Nonetheless, classical one-period mean-variance analysis continues to dominate most academic application or extension of portfolio theory, and it dominates all practical and industry analysis, even when long-term investors are making dynamic investment decisions involving non-normal payoffs, such as in actively-managed mutual and hedge funds. Perhaps there is some wisdom in the old joke about the drunk who looks for his car keys under the streetlight. My hope is that the long-run mean-variance frontier can form an analogous simple conceptual benchmark for dynamic intertemporal portfolio theory with incomplete markets.

As in the one-period (discrete-time) case, however, mean-variance analysis is a "benchmark," not an "approximation." Long-horizon returns are potentially far from normally distributed, and even power utility with lognormal returns can lead to optimal payoffs substantially different from the long-run mean-variance frontier. More importantly, the long-run mean-variance characterization also does not help much to *compute* optimal portfolios in the standard challenging environments. Finding a traded discount factor sounds easy, but in fact it can be a problem of the same order of complexity as solving the dynamic programs or minimax problems. This is not its purpose. As in one-period analysis, the point of mean-variance analysis is that we are able quickly and intuitively to characterize portfolio choice, and think about the choice between large "funds," even though construction of those funds remains challenging.

Additional Literature

The basic idea of treating a discounted sum as an expectation in (1) comes from Hansen (1987), which includes an extensive analysis of asset pricing with quadratic utility. Magill and Quinzii (2000) is the most recent paper in this line, and the one most directly related to this paper. Magill and Quinzii also specify a quadratic utility investor, they characterize the optimal payoff in terms of a "least variable income stream," a concept similar to my "long-run mean-variance frontier," and they derive a similar long-run CAPM. Hansen and Sargent (2004) argue in a wide variety of circumstances for the approach of this paper: study exact solutions to problems with linear-quadratic preferences, rather than approximate or numerical solutions to problems with more realistic preferences. Hansen and Sargent also show how to extend quadratic specifications to include habits, recursive utility and so forth. My approach to the mean-variance frontier comes from Hansen and Richard (1987), also summarized in Cochrane (2004).

This work falls into two larger-scale trends in financial research. First, and at a most basic level, the empirical finding that returns are not i.i.d., and that discount rate news rather than cashflow news drives much price variation, requires a rewriting of most procedures in finance, including asset pricing, corporate finance such as cost-of-capital calculations, as well as portfolio theory. If returns and outside income were i.i.d. there would be no need for dynamic, or long-run, portfolio theory.

Second, the focus on the payoff stream in portfolio theory mirrors a renewed interest in payoff streams and long-run analysis in asset pricing more generally, for example Menzly, Santos and Veronesi (2004), Bansal and Yaron (2004), Hansen Heaton and Li (2005), Bansal, Dittmar and Lundblad (2005), Lettau and Wachter (2007), Gabaix (2007) and many others. Return betas are driven by the comovement of tomorrow's price with factor prices, and we now attribute much price movement to discount rates rather than cash flows. Beta is therefore largely endogenous, and so makes less sense as the central explanatory variable. These authors try instead to account for prices based on long-run cash flow streams. One can foresee a day when prices are our central endogenous variable, not an ad-hoc sorting characteristic (book/market ratio), in which the stream of cashflows is the central exogenous variable, and when one-period returns are barely mentioned; that we will treat stocks as we now do bonds. This approach is nothing new from a pure theory point of view; we have been able to write $p_t = E_t \sum_{j=1}^{\infty} m_{t+j} D_{t+j}$ as long as we have been able to write $1 = E_t(m_{t+1}R_{t+1})$. The challenge these authors are facing, and to which this paper makes a small contribution, lies in specifying workable applications.

2 Asset pricing environment

This section sets up notation to think about dynamic intertemporal portfolio problems in analogy to one-period problems, by treating date and state symmetrically, and interprets the resulting quantities.

2.1 Payoffs, prices, and discount factor

x denotes a *payoff*. In a one-period setting, the payoff is the amount x_1 that an investor receives at date 1, in each state of nature, for a time-zero price p_0. In an intertemporal setting, the payoffs are the streams of dividends $\{x_1, x_2, ...\}$, or $\{x_t dt\}$ in continuous time, resulting from an initial purchase. (The appendix summarizes notation.)

Returns are price-one payoffs. We can form returns by dividing a payoff by its initial price,

$$y_t = x_t / p(\{x_t\}),$$

where $p(\cdot)$ means "price of." In the intertemporal setting, the "return" to a particular date – the dividends or coupons accruing to a one-dollar purchase – has the units of a yield, or coupon rate; it is a number like 0.04 not 1.04. I use the notation y and the word "yield" rather than the word "return" to help keep the typical units in mind.

In a one period model, the risk free asset pays one unit in each state. The natural risk free payoff is thus one in all states and dates, a perpetuity,

$$x_t^f = 1.$$

The *risk free yield* is then naturally,

$$y_t^f = 1/p(\{1\}).$$

The riskfree yield is also a number like 0.01, not a number like the 1.01.

Excess yields are zero-price payoffs, which we can construct by differencing any two yields or returns,

$$z_t = y_t^1 - y_t^2;\ p\left(\{z_t\}\right) = 0.$$

When not required for clarity, I'll drop the time subscripts, e.g. $y \equiv y_t$ and sequence notation, e.g. $p\left(\{x_t\}\right) = p(x)$.

I use the notation $\mathcal{E}(x)$ to denote different operations, depending on context.

$$\begin{aligned}
\text{one period:} \qquad & \mathcal{E}(x) \equiv \frac{1}{\beta}E(\beta x_1) \\
\text{intertemporal, discrete:} \qquad & \mathcal{E}(x) \ \equiv \frac{1-\beta}{\beta}E\sum_{t=1}^{\infty}\beta^t x_t \\
\text{infinite period, continuous:} \qquad & \mathcal{E}(x) \ \equiv \rho E\int_0^{\infty} e^{-\rho t}x_t dt.
\end{aligned}$$

One can similarly write environments with a terminal date T and with a separate terminal payment.

The $\mathcal{E}$ operator takes a sum over time, weighted by β^t or $e^{-\rho t}$, as well as a sum over states, weighted by probability. Weighting is not essential, but it allows us to produce finite values for a larger set of payoff processes in an infinite-period environment. It will be useful, but not necessary, to pick β or ρ as an agent's subjective discount factor. One can substitute more general weighting functions. With the square-integrable assumption (2) below, $\mathcal{E}(xy)$ is an inner product, defining a Hilbert space as in Hansen and Richard (1987) and Magill and Quinzii (2000). Then, we can think of dividend streams x as vectors and $\mathcal{E}(xy)$ as an inner product.

I call $\mathcal{E}(x)$ the "long run mean" and I call $\tilde{\sigma}^2(x) \equiv \mathcal{E}(x^2) - \left[\mathcal{E}(x)\right]^2 = \mathcal{E}\left[(x - \mathcal{E}(x))^2\right]$ the "long run variance" of the payoff stream x_t. The "variance" concept prizes stability over time as well as stability across states of nature. A variable x_t that varies deterministically over time still has long-run variance.

With this notation, I write the fundamental pricing equation as $p = k\mathcal{E}(mx)$, meaning

$$\begin{aligned}
\text{one period:} \qquad & p_0 = \beta\mathcal{E}(m_1 x_1) \ = E(\beta m_1 x_1) \\
\text{intertemporal, discrete:} \qquad & p = \frac{\beta}{1-\beta}\mathcal{E}(mx) = E\sum_{t=1}^{\infty}\beta^t m_t x_t \\
\text{intertemporal, continuous} \quad : \quad & p = \frac{1}{\rho}\mathcal{E}(mx) = E\int_0^{\infty} e^{-\rho t}m(t)x(t)dt,
\end{aligned}$$

where m_t is a stochastic discount factor. (More precisely, it is the discount factor scaled by the weighting function; m_t is $u'(c_t)$ not $\beta^t u'(c_t)$.) The appearance of the constant k in the fundamental pricing equation is inelegant. However, we gain more in convenience by defining "long run mean" with weights that sum to one than we lose by introducing this constant in the pricing equation. Nothing essential is affected by the choice of notation.

$\underline{X}$ denotes the *payoff space* of all payoff streams that investors can buy, $\underline{Y}$ denotes the set of price-one returns or yields, and $\underline{Z}$ denotes the set of price-zero excess returns or yields,

$$\underline{Y} \equiv \{y \in \underline{X} : p(y) = 1\},$$

$$\underline{Z} \equiv \{z \in \underline{X} : p(z) = 0\}.$$

In many circumstances, I limit the payoff space $\underline{X}$ to include only square-integrable payoffs

$$\mathcal{E}(x^2) < \infty, \tag{2}$$

In an infinite-period model, this requirement limits us to payoffs that do not grow too fast, i.e. that do not vary too much over time, as well as limiting variance in the usual sense.

I let investors buy any portfolio of payoffs, which means that $\underline{X}$ and $\underline{Z}$ are closed under linear combinations,

$$x^1, x^2 \in \underline{X} \rightarrow ax^1 + bx^2 \in \underline{X}. \tag{3}$$

In an intertemporal context, we also want to allow *dynamic trading*, or equivalently we want to allow entrepreneurs to sell payoffs to *managed portfolios*. If $\{x_t\} \in \underline{X}$, we also allow

$$\{x_0, ..., x_{t-1}, x_t + p_t(x)\} \in \underline{X} \tag{4}$$

$$\{0, ... - p_t(x), x_{t+1}, x_{t+2}...\} \in \underline{X}. \tag{5}$$

In continuous time, with a set of assets with excess return process dr_t^e and riskfree rate r_t^f we allow payoffs x_t generated by

$$dV_t = \left(r_t^f V_t - x_t\right) + w_t' dr_t^e.$$

As usual, I limit dynamic trading so that the investor cannot generate arbitrage opportunities. The time-zero value of wealth must tend to zero, $\lim_{T\rightarrow\infty} p(W_T) = 0$, and the size of trading weights must be limited. (See for example Duffie (2001).) The restriction (2) can further limit dynamic trading. Since the limitations on dynamic trading do not appear in any of my analysis, this discussion intentionally omits the substantial technical details.

I do not assume that the payoff space $\underline{X}$ is complete, meaning that every random variable x, or even every variable with $\mathcal{E}(x^2) < \infty$, can be traded. I explicitly allow for two sources of incompleteness: the investor may have a labor or business income stream that cannot be completely hedged with traded assets, and there may be state variables for investment opportunities whose shocks cannot be spanned by those of traded assets.

I assume that prices and payoffs follow the law of one price, or linearity,

$$p(ax^1 + bx^2) = ap(x^1) + bp(x^2). \tag{6}$$

It is useful to construct a stochastic discount factor that is also a traded payoff. The standard conditions (Hansen and Richard 1987) on the payoff space $\underline{X}$ that *guarantee* the existence of such a discount factor can apply in this context as well. First, payoffs must have finite long-run variance as in (2). Second, we need to assume that investors can form arbitrary payoffs (3) and that the law of one price or linearity of the pricing function (6) holds. Third, the payoff space must be "complete" in the sense that if a sequence of payoffs is in $\underline{X}$ and converges, its limit point must also be in $\underline{X}$. ("Convergence" here uses $\mathcal{E}(x^2)$ as a norm.) With these assumptions, we can guarantee that there is a unique discount factor x^* in the payoff space, i.e.

$$\exists x^* \in \underline{X} : p = k\mathcal{E}\left(x^* x\right).$$

x^* is a *dividend stream* that acts as a discount factor *process,* i.e.

$$p(x) = E\sum_{t=1}^{\infty} \beta^t x_t^* x_t \text{ or } p(x) = E\int_0^{\infty} e^{-\rho t} x_t^* x_t dt.$$

The standard proof[1] by Riesz representation theorem then applies.

The conditions for this theorem are actually somewhat restrictive, and the conclusion and web-appendix explore some common setups in which they are violated. (Most of the trouble involves growth in infinite-period settings.) However, these are only sufficient, not necessary conditions for a traded discount factor. For example, short sale constraints imply a violation of the portfolio assumption (3) and the completeness assumption. A discount factor that is a linear function of asset payoffs may require negative weights and thus may not be tradeable. But then again, it might not require negative weights; we might get lucky and we might be able to construct a traded discount factor anyway, perhaps for a restricted range of parameters. Similar luck can happen in the infinite-period settings that cause problems here.

When there is a finite vector of basis payoffs $\mathbf{x}$ with prices $\mathbf{p}$, and ignoring dynamic trading beyond what is included in the basis assets (the basis payoffs may themselves be payoffs from managed payoffs), so the payoff space is $\underline{X} = \{\mathbf{c}'\mathbf{x}\}$, the usual discount factor construction applies. The payoff

$$x^* = \frac{1}{k}\mathbf{x}'\mathcal{E}(\mathbf{x}\mathbf{x}')^{-1}\mathbf{p} \tag{7}$$

is a discount factor, i.e. it satisfies $p = k\mathcal{E}(x^* x)$.

As the notation suggests, we can use the $\mathcal{E}$ operator to define long-run projections $proj(\cdot|\underline{X})$ and long-run regressions. A long run regression $y = xb + \varepsilon$; $b = \mathcal{E}(xx')^{-1}\mathcal{E}(yx)$ is a regression of one *dividend stream* against another, and prizes a fit over time as well as across states of nature.

2.2 Mean-Variance Frontier

The *long-run mean-variance frontier* consists of payoffs that solve

$$\min_{\{y \in \underline{Y}\}} \ \mathcal{E}(y^2) \ \text{ s.t. } \mathcal{E}(y) = \mu.$$

I follow the Hansen-Richard (1987) approach (see also Cochrane (2004)), which makes clear the intimate link between marginal utility, mean-variance frontiers, and discount factors, but the familiar Lagrangian minimization works just as well when the payoffs are created from a finite set of basis assets. The frontier is generated as

$$y^{mv} = y^* + wz^*. \tag{8}$$

[1] Ross (1976) first proved the basic theorem, and Cochrane (2004) presents a simple textbook discussion. See Hansen and Richard (1987) for the completeness assumption, which turns out to matter in these applications.

Here, y^* is defined by

$$y^* = \frac{x^*}{p(x^*)} = \frac{x^*}{k\mathcal{E}(x^{*2})}, \tag{9}$$

and z^* is defined by

$$z^* = proj(1|\underline{Z}). \tag{10}$$

z^* is the excess return "closest" to the perpetuity payoff.

y^* and z^* have the usual properties from the one period case, suitably reinterpreted. y^* is the discount-factor mimicking yield. For any discount factor m, we have $x^* = proj(m|\underline{X})$ and then (9). y^* is the minimum long-run second moment yield,

$$y^* = \arg \min_{\{y \in \underline{Y}\}} \mathcal{E}(y^2).$$

Since y^* is proportional to x^* it can be used to price other payoffs–any mean-variance efficient return carries pricing information:

$$\mathcal{E}(y^* y) = \mathcal{E}(y^{*2}) \quad \forall y \in \underline{Y}.$$

An explicit formula for the finite-basis case follows from (7) and mimics the standard formulas for one-period mean-variance frontier returns:

$$y^* = \frac{\mathbf{1}'\mathcal{E}(\mathbf{y}\mathbf{y}')^{-1}\mathbf{y}}{\mathbf{1}'\mathcal{E}(\mathbf{y}\mathbf{y}')^{-1}\mathbf{1}}. \tag{11}$$

z^* generates *long-run means* in the same way that x^* generates prices,

$$\mathcal{E}(z) = \mathcal{E}(z^* z) \;\; \forall \; z \in \underline{Z}. \tag{12}$$

Since z^* is a price-zero excess yield, y^* and z^* are orthogonal, $\mathcal{E}(y^* z^*) = 0$. If a riskfree yield is traded ($1 \in \underline{X}$) then z^* is simply

$$z^* = \frac{y^f - y^*}{y^f}. \tag{13}$$

One can avoid risk-free rate and inflation issues by focusing on the mean-variance frontier of excess yields,

$$\min_{\{z \in \underline{Z}\}} \; \mathcal{E}(z^2) \;\; \text{s.t. } \mathcal{E}(z) = \mu.$$

This frontier is generated simply by

$$z^{mv} = wz^*, \;\; w \in \Re.$$

When there is a riskfree rate, we can compute z^* by (13). With a finite set of basis assets, we can also calculate z^* analogously to the calculation (7) of y^*:

$$z^* = \mathcal{E}(\mathbf{z})'\mathcal{E}\left(\mathbf{z}\mathbf{z}\right)^{-1}\mathbf{z} \tag{14}$$

where $\mathbf{z}$ is a vector of excess yields.

All of these results can be derived by following exactly the Hansen-Richard (1987) approach in two-period models, but using $\mathcal{E}$ in the place of E. We show that any yield (return) can be written as $y^i = y^* + w^i z^* + \eta^i$, with $\mathcal{E}(\eta^i) = 0$, $\mathcal{E}(\eta^i y^*) = 0, \mathcal{E}(\eta^i z^*) = 0$. The mean-variance frontier is then the set of yields with $\eta^i = 0$.

Of course one can span the frontier with any other two returns as well, and one can extend any standard characterization of the mean-variance frontier to the long-run mean variance frontier by using moments $\mathcal{E}$ in the place of moments E.

2.3 Expected returns and betas

As in one-period asset pricing, we can connect discount factors, asset pricing, mean-variance frontiers, and expected return-beta models. I write here the case with a traded risk-free yield. First, we can write the fact that x^* is a discount factor as a single-beta representation for long-run expected yields. For any yield i,

$$\mathcal{E}(y^i) - y^f = \tilde{\beta}_{i,x^*}\lambda_{x^*};\text{ or } \mathcal{E}(z^i) = \tilde{\beta}_{i,x^*}\lambda_{x^*}$$

Second, we can use any long-run mean-variance efficient yield y^{mv} in such a representation, i.e.,

$$\mathcal{E}(y^i) - y^f = \tilde{\beta}_{i,mv}\left[\mathcal{E}(y^{mv}) - y^f\right]. \tag{15}$$

$\tilde{\beta}$ denotes a long-run regression coefficient, i.e. defined using long-run moments $\mathcal{E}$. Of course, the numbers, units and definition of moments are entirely different in the long run case.

Derivation. The derivation simply parallels standard derivations in one-period models. Denote $z^i = y^i - y^f$. Then,

$$\begin{aligned} 0 &= \mathcal{E}(x^* z^i) = \mathcal{E}(x^*)\mathcal{E}(z^i) + \widetilde{cov}(x^*, z^i) \\ \mathcal{E}(z^i) &= \frac{\widetilde{cov}(x^*, z^i)}{\tilde{\sigma}^2(x^*)}\left(-\frac{1}{\mathcal{E}(x^*)}\tilde{\sigma}^2(x^*)\right) = \tilde{\beta}_{i,x^*}\lambda_{x^*}. \end{aligned}$$

Second, decompose any yield into three orthogonal components,

$$y^i = y^f + w^i\left(y^f - y^*\right) + \eta^i;\ \mathcal{E}(\eta^i) = 0;\ p(\eta^i) = 0;\ \mathcal{E}(\eta^i y^*) = 0 \tag{16}$$

Choose any mean-variance efficient yield except y^f as reference,

$$y^{mv} = y^f + w^{mv}\left(y^f - y^*\right).$$

Now we can can rewrite (16) as

$$y^i = y^f + \frac{w^i}{w^{mv}}\left(y^{mv} - y^f\right) + \eta^i$$

The residual is orthogonal to the right hand variable, so this is a long-run regression with

$$\tilde{\beta}_{i,mv} = w^i/w^{mv}.$$

Equation (15) follows.

3 Portfolio problems

An investor has initial wealth W, a stream of labor or business income $e = \{e_t\}$, and he can buy payoffs $x = \{x_t\} \in \underline{X}$ at prices p. I assume no arbitrage in the available prices and payoffs, so there is a discount factor m that satisfies $p = k\mathcal{E}(mx)$. The investor's problem is then

$$\max_{\{x\in\underline{X}\}} \mathcal{E}\left[u(c)\right]\ \ s.t.\ W = k\mathcal{E}(mx),\ c = e + x. \tag{17}$$

As a reminder, we interpret these familiar-looking symbols as long-run portfolio problems, for example

$$\max_{\{x_t \in \underline{X}\}} E\sum_{t=1}^{\infty} \beta^t u(c_t) \;\; s.t.\; W = p(\{x_t\}) = E\sum_{t=1}^{\infty} \beta^t m_t x_t;\; c_t = e_t + x_t,$$

or

$$\max_{\{x_t \in \underline{X}\}} E\int_{t=0}^{\infty} e^{-\rho t} u(c_t)dt \;\; s.t.\; W_0 = p(\{x_t\}) = E\int_{t=0}^{\infty} e^{-\rho t} m_t x_t dt;\; c_t = e_t + x_t.$$

I have simplified the notation by using the subjective discount factor as the time-weighting function in $\mathcal{E}$, but clearly nothing essential hinges on this choice.

The first-order conditions state that at an optimum $\hat{x}$,

$$u'(\hat{x} + e) = \lambda m. \tag{18}$$

Marginal utility is proportional to a discount factor, i.e. for all $x \in \underline{X}$,

$$p = k\mathcal{E}\left[\frac{u'(\hat{x} + e)}{\lambda} x\right].$$

Inverting (18), the solution to the portfolio problem is characterized by

$$\hat{x} = u'^{-1}(\lambda m) - e. \tag{19}$$

The payoff (19) has a simple intuition: The investor consumes more $c = \hat{x} + e$ in "cheap" (low m) states and dates, and less in "expensive" (high m) states and dates, with u'^{-1} dictating how much or little to respond to this relative date- and state-price. The traded payoff $\hat{x}$ simply offsets the effects of outside income e. The Lagrange multiplier scales consumption up and down to match initial wealth W.

If markets are complete, the discount factor $m = x^* \in \underline{X}$ is unique, traded, and therefore often easy to find. Every payoff is traded, so the construction (19) satisfies the constraint $\hat{x} \in \underline{X}$. Hence, all we have to do is find the Lagrange multiplier λ to satisfy the initial wealth constraint.

I focus on the more interesting case that markets are not complete. Now, condition (19) is necessary, but not sufficient. There are many discount factors that price assets, and for only one of them is the inverse marginal utility in the space of traded payoffs $\underline{X}$. To solve this problem, I specialize to quadratic utility, so that marginal utility is linear. The payoff space $\underline{X}$ is closed under linear transformations (portfolio formation, equation (3)), so once we construct the traded discount factor x^*, we know that the inverse image of $x^* \in \underline{X}$ is guaranteed also to be in the space of payoffs $\underline{X}$, and this is the optimal payoff. Though stated with no outside income, this logic extends to outside income, as we'll see in a moment.

Analytically, I specialize to

$$u(c_t) = -\frac{1}{2}(c_t^b - c_t)^2. \tag{20}$$

c_t^b is a potentially time-varying stochastic bliss point. A time-varying or stochastic preference shift, represented here by a bliss point, may help to provide more realistic answers, for example by accommodating growth, life-cycle and household-composition effects, or to give a better approximation to nonlinear utility functions.

The optimal portfolio is then characterized by

Proposition 1. *The optimal payoff for the investor (20) is given by*

$$\hat{x} = \left(\hat{c}^b - \hat{e}\right) - \left[p(\hat{c}^b - \hat{e}) - W\right] y^*, \tag{21}$$

where the hedge payoffs $\hat{c}^b$, $\hat{e}$ *are the projections of the preference shock and outside income on the set of traded payoffs, e.g.*

$$\hat{c}^b \equiv proj\left(c^b|\underline{X}\right), \; \hat{e} \equiv proj\left(e|\underline{X}\right) \tag{22}$$

W *is initial financial wealth,* y^* *is the discount-factor mimicking and minimum second-moment yield (9).*

Derivation. With the quadratic utility function (20), the first order condition (18) reads

$$c^b - \hat{x} - e = \lambda m.$$

Solving for $\hat{x}$ and projecting both sides on the set of traded assets $\underline{X}$ yields

$$\hat{x} = -\lambda x^* + \hat{c}^b - \hat{e}. \tag{23}$$

The wealth constraint states

$$W = p(\hat{x}) = -\lambda p(x^*) + p(\hat{c}^b - \hat{e}).$$

Solving for λ, substituting λ in (23), and using the definition $y^* = x^*/p(x^*)$ we obtain the optimal payoff (21).

The expression (21) offers a simple interpretation. (Many more will follow). The investor starts by buying a payoff $\hat{c}^b - \hat{e}$ that gets him as close to the bliss point as traded assets allow. We can also think of the payoff $\hat{c}^b - \hat{e}$ as the optimal hedge for preference shocks and labor income risk. It is formed by a "long-run" regression of the streams c^b and e on the yields or dividend streams of the traded assets.

Typically, complete hedging is not possible. Wealth W is lower than the cost $p(\hat{c}^b - \hat{e})$ of the hedge payoff. Thus, the investor shorts an optimal risky payoff y^* in order to buy the individual hedge payoff. y^* is proportional to contingent claims prices, so by shorting y^* the investor is shorting the "most expensive" payoff, in order to generate the largest funds possible at minimum risk. y^* is of course on the mean-variance frontier.

In sum, *each investor's optimal portfolio is a combination of a labor income and preference shock hedge payoff, plus an investment in a long-run mean-variance efficient yield.*

4 Yields and the market

Equation (21) is an unusual statement of a familiar result, and intuition is most of the point of the paper, so I focus instead on the long-run version of standard statements of mean-variance analysis. I cast results in terms of the yield (return) of the optimal portfolio; I characterize preferences by risk aversion rather than by a bliss point; I express the optimal payoff in reference to a mean-variance efficient payoff on the upper part of the frontier, rather than the minimum second moment yield y^* which is on the lower part of the frontier; I express the optimal portfolio as a set of distortions to mean-variance efficiency induced by nontraded income and preference shocks; and I express the optimal portfolio relative to a market yield.

4.1 No outside income or preference shocks

I start with the following special case: The preference-shock hedge payoff $\hat{c}^b$ is constant, the outside-income hedge payoff $\hat{e}$ is zero, and a risk free yield y^f is traded. (This case is slightly more general than a constant bliss point and no outside income. The bliss point may vary and there may be outside income, so long as their hedge payoffs are constant and absent, respectively.) This classic special case simplifies the formulas a great deal, and shows the structure of the main ideas. The more interesting case with outside income and preference shocks follows.

4.1.1 Mean-variance frontier

Proposition 2 offers a more familiar statement of a mean-variance frontier:

Proposition 2. *The yield of the optimal payoff is on the long-run mean-variance frontier, with a greater investment in risky assets for investors with lower risk aversion*

$$\hat{y} = y^f + \frac{1}{\gamma}\left(y^f - y^*\right), \tag{24}$$

where γ is the investor's coefficient of risk aversion,

$$\frac{1}{\gamma} \equiv \frac{\hat{c}^b - y^f W}{y^f W}. \tag{25}$$

Derivation. Since $\hat{c}^b$ is constant and a riskfree rate is traded, $p(\hat{c}^b) = \hat{c}^b / y^f$. Then, from (21),

$$\begin{aligned}
\hat{x} &= \hat{c}^b - \left[\frac{\hat{c}^b}{y^f} - W\right] y^* \\
\hat{y} &= \frac{\hat{x}}{W} = \frac{\hat{c}^b}{y^f W} y^f - \left[\frac{\hat{c}^b}{y^f W} - 1\right] y^* \\
\hat{y} &= y^f + \left[\frac{\hat{c}^b}{y^f W} - 1\right]\left(y^f - y^*\right),
\end{aligned}$$

and (24) follows.

Since y^* as minimum long-run second-moment yield is on the lower portion of the mean-variance frontier, expression (24) places the reference mean-variance efficient payoff on the more familiar, upper, portion of the frontier. We see a greater investment in this risky payoff for investors with lower risk aversion.

For quadratic utility, the risk aversion coefficient is

$$\frac{1}{\gamma} \equiv \frac{u'(c)}{-cu''(c)} = \frac{c^b - c}{c}.$$

Thus, we interpret γ as defined by (25) as the local risk aversion coefficient, evaluated at a value of consumption in which the investor invests all wealth in the riskfree rate. In expression (24), risk aversion is constant through time. Only $\hat{y}$ and y^* (really, $\hat{y}_t$ and y_t^*) vary over time. I later define a "time-varying" risk aversion coefficient, depending analogously on wealth at time t, and for other purposes, but this one does not vary.

With a finite set of basis assets, (14) gives us a very traditional expression of the optimal payoff,

$$\hat{y} = y^f + \frac{y^f}{\gamma}\mathcal{E}(\mathbf{z})'\mathcal{E}\left(\mathbf{z}\mathbf{z}'\right)^{-1}\mathbf{z}.$$

In the one period case, $y^f = R^f \approx 1$, and $E(R^{e2}) = \sigma^2(R^e) + E(R^e)^2 \approx \sigma^2(R^e)$ so the numbers are similar to the continuous-time result that the risky share is $\Sigma^{-1}\mu/\gamma$. For long-horizons, y^f is a number near zero, and means are as or more important than variances in forming long-run second moments, so the numbers can be quite different.

This is a result about long-horizon portfolio theory, in an environment with time-varying investment opportunities and incomplete markets. Yet the intensive dynamic trading, the usual rules for allocating wealth to securities based on state variables, and "hedging demands" for shocks to those state variables are absent from this representation. These demands can appear, in the *construction* of the mean-variance efficient payoff. In a static context, we construct mean-variance efficient portfolios by optimally allocating wealth across securities, trading mean return against variance and covariance. In this dynamic context, we may also optimally allocate wealth to securities as their means and variances change over time, as allocating wealth over time is fundamentally just like allocating wealth across assets; we may optimally hedge such changes in the investment opportunity set, and we may correctly retrieve payoffs from accumulated wealth. All this construction and financial engineering is hidden in a long-run mean-variance characterization, just as it is hidden in the classical one-period mean-variance characterization. That's the point – the hard part is producing a mean-variance efficient yield y^*. Once that is done, we can simply characterize the investor's choice between funds, just as we do in one-period models.

For this reason, the investor may *not* have to do all this dynamic portfolio allocation and payout construction. It would be perfectly natural for firms and fund managers to do the "hard part," and market long-run mean-variance efficient securities. If the investors' optimization takes place over an asset space that for this reason already contains long-run mean-variance efficient funds, the investor's job is very simple.

4.1.2 Market payoff

As this discussion suggests, constructing long-run mean-variance efficient payoffs may be difficult. However, as in the one-period case, the average investor must hold the market portfolio. This deep insight can offer a great simplification in applying portfolio theory: rather than think about your absolute level of risk aversion and explicit models of asset payoff moments, think only about how you are *different* than everyone else, and adjust your position in simple indices accordingly. As I argued in the introduction, thinking in relative terms can also avoid the pitfall that portfolio advice can't apply to everyone.

Following this insight, we want to express the optimal payoff *relative* to the market payoff, in an equilibrium in which investors have limited and described forms of heterogeneity.

Proposition 3. *If all investors are of this type, the yield of the investor's optimal portfolio is split between the riskfree yield and the market yield, which is a claim to aggregate consumption,*

$$\hat{y}_i = y^f + \frac{\gamma_a}{\gamma_i}\left(\hat{y}_a - y^f\right), \tag{26}$$

where the yield on the market payoff is

$$\hat{y}_a \equiv \frac{\sum_j W_j \hat{y}_j}{\sum_j W_j} = \frac{\sum_j \hat{x}_j}{\sum_j W_j} = \frac{\hat{x}_a}{W_a},$$

and aggregate risk aversion is defined as a wealth-weighted average of individual risk aversions,

$$\frac{1}{\gamma_a} \equiv \frac{\sum_j W_j \frac{1}{\gamma_j}}{\sum_j W_j}.$$

Derivation. Start with (24), sum over investors, and divide by wealth,

$$\begin{aligned}
\hat{y}_i &= y^f + \frac{1}{\gamma_i}\left(y^f - y^*\right) \\
\frac{\sum_j W_j \hat{y}_j}{\sum_j W_j} &= y^f + \frac{\sum_j W_j \frac{1}{\gamma_j}}{\sum_j W_j}\left(y^f - y^*\right) \\
\hat{y}_a &= y^f + \frac{1}{\gamma_a}\left(y^f - y^*\right) \\
y^f - y^* &= \gamma_a\left(\hat{y}_a - y^f\right).
\end{aligned} \tag{27}$$

Substitute this result in the right hand side of (24).

This result can help a lot in thinking about portfolios. Keep in mind that we still have time-varying portfolio opportunities described by unspanned state variables. In this case, the infinitely risk averse investor holds the perpetuity y^f, as pointed out by Campbell and Viceira (2001) and Wachter (2003). An investor whose risk aversion is the same as that of the average investor just holds the market payoff, as he must, despite any time-varying investment opportunities, and that market payoff pays aggregate consumption as its dividend. The novelty is that investors with risk aversion (say) greater than the market now can simply purchase a payoff that is a linear combination of the market and the real perpetuity.

The *yield* that is halfway between those of the market and the index perpetuity is not generated by a *portfolio* that constantly rebalances to 50/50 weights. Constant rebalancing gives rise to payoffs that are nonlinear functions of underlying payoffs. Constructing the dynamic portfolio rule to achieve the 50/50 payoff is not that easy. In an i.i.d. environment, one can construct the 50/50 payoff either by shorting a portfolio that itself shorts a mean-variance efficient portfolio, or by a strategy with suitable time-varying portfolio and payout weights. These are characterized in the web-appendix. The problem is even harder in a non-i.i.d. environment. Once again, *characterizing* the optimal payoff can be a lot easier than *constructing* it, which is one of my main points.

4.1.3 A long-horizon CAPM

In an equilibrium of investors who are all of the same type, but vary by risk aversion, we naturally can write a CAPM-like result:

Proposition 4. *For each asset i, the long-run expected yield follows a long-run CAPM,*

$$\mathcal{E}(y^i) - y^f = \tilde{\beta}_{i,a}\left[\mathcal{E}(\hat{y}_a) - y^f\right]$$

where $\tilde{\beta}_{i,a}$ is a long-run regression coefficient of yield i on the market yield.

This proposition follows simply from (15) and the fact (27) that the market payoff is long-run mean-variance efficient. Thus, in pricing as in portfolio behavior, Mertonian state variables disappear from long-run expected returns. Even though there can be a complex ICAPM representation of one-period returns, long-run returns obey the long-run CAPM. Again, time-series variation in expected returns is, from this perspective, no different from cross-sectional variation in returns. (MacGill and Quinzii (2000) derive a similar representation).

4.2 Outside income and preference shocks

Next, I allow outside income and preference shocks, while retaining a traded risk-free yield. To incorporate preference shocks, I separate the preference-shock hedge payoff into a constant and a variable component,

$$\hat{c}^b = \bar{c}^b \times 1 + \tilde{c}^b; \mathcal{E}\left(\tilde{c}^b\right) = 0.$$

Equation (21) now reads

$$\hat{x} = \left(\bar{c}^b + \tilde{c}^b - \hat{e}\right) - \left[p(\bar{c}^b + \tilde{c}^b - \hat{e}) - W\right] y^* \tag{28}$$

The variable part $\tilde{c}^b$ of the bliss point hedge payoff and the outside income hedge payoff $\hat{e}$ enter together in everything that follows from this equation, so to save some space I combine them in what follows; I use the symbol $\hat{e}$ to denote the quantity $\hat{e} - \tilde{c}^b$ and I say only "outside income hedge payoff" to refer to both components.

4.2.1 Mean-variance frontier and risk aversion

The mean-variance characterization now applies to the "total payoff," consisting of the asset payoff and the outside income hedge payoff:

Proposition 5. *The total yield, including the yield of the outside-income hedge payoff, is on the long-run mean-variance frontier,*

$$\hat{y}^T = y^f + \frac{1}{\gamma}\left(y^f - y^*\right). \tag{29}$$

The total yield and risk aversion are defined here by

$$\hat{y}^T \equiv \frac{\hat{x}+\hat{e}}{W+p(\hat{e})} \tag{30}$$

$$\frac{1}{\gamma} \equiv \frac{\bar{c}^b - y^f\left[W+p(\hat{e})\right]}{y^f\left[W+p(\hat{e})\right]}. \tag{31}$$

Derivation. Since $\bar{c}^b$ is constant and a riskfree rate is traded, $p(\bar{c}^b) = \bar{c}^b/y^f$. Then, from (28),

$$\begin{aligned}
\hat{x} &= \left(\bar{c}^b - \hat{e}\right) - \left[p(\bar{c}^b - \hat{e}) - W\right] y^* \\
\frac{\hat{x}+\hat{e}}{W+p(\hat{e})} &= \frac{\bar{c}^b}{y^f\left[W+p(\hat{e})\right]} y^f - \left[\frac{\bar{c}^b}{y^f\left[W+p(\hat{e})\right]} - 1\right] y^* \\
\hat{y}^T &= y^f + \left[\frac{\bar{c}^b}{y^f\left[W+p(\hat{e})\right]} - 1\right]\left(y^f - y^*\right),
\end{aligned}$$

and (29) follows.

The "total" payoff consists of the assets $\hat{x}$ actually held, and the outside-income hedge payoff that the investor holds implicitly. The investor calculates the "payoff equivalent" of his outside income and adjusts his asset payoff $\hat{x}$ accordingly.

For this representation, we interpret γ as defined by (31) as the local risk aversion coefficient, evaluated at a constant bliss point $\bar{c}$ and a value of consumption in which the investor invests all wealth, including the proceeds from selling the hedge payoff for outside income, in the riskfree rate. Again, this is just a sensible point on the utility function at which we take derivatives.

As in the simple case, time-varying investments and "hedging demands" for shocks to state variables are absent, now including state variables for time-varying and stochastic outside income as well as state variables for time-varying investment opportunities, though they appear in the dynamic strategies required to construct the mean-variance efficient and hedge payoffs. Outside income itself, represented in $\hat{e}$, does appear in the payoff, as it appears in one-period mean-variance problems.

4.2.2 Portfolio distortions

It is useful to characterize the distortions of the asset payoff, relative to the mean-variance frontier, induced by the presence of outside income and preference shocks.

We have the result already of course; we can just say "construct a long-run mean-variance efficient payoff, then subtract the hedge payoff for outside income." In equations, we can write from (29) and (30)

$$s_W \hat{y} = y^f + \frac{1}{\gamma}\left(y^f - y^*\right) - s_e y_e,$$

where y_e denotes the yield of the outside-income hedge payoff

$$y_e \equiv \frac{\hat{e}}{p(\hat{e})},$$

and the wealth shares s_W, s_e are defined as

$$s_W \equiv \frac{W}{W + p(\hat{e})};\ s_e \equiv \frac{p(\hat{e})}{W + p(\hat{e})}.$$

(We do not have $p(e) = p(\hat{e})$; we do not know how to assign prices for nontraded payoffs. Hence, though I call $W + p(\hat{e})$ "total wealth," it really is only "asset wealth plus the value of the outside-income hedge payoff.") Even more simply, we can multiply by total wealth, and express the asset payoff as an investment of total wealth in the mean-variance efficient payoff, plus sale of the outside-income hedge payoff.

$$\hat{x} = [W + p(\hat{e})]\left[y^f + \frac{1}{\gamma}\left(y^f - y^*\right)\right] - \hat{e} \tag{32}$$

However, the outside income hedge payoff $\hat{e}$ will usually contain some riskfree yield y^f and some of the mean-variance yield $\left(y^f - y^*\right)$. Buying something just to sell it again is a messy way to describe a portfolio, especially in a real world with short selling constraints and transactions costs. It is therefore interesting to describe the asset market yield $\hat{y}$ or payoff $\hat{x}$ directly, subtracting out the components of $\hat{e}$ that affect the riskfree and mean-variance investment. There are many ways to do this, depending on how one characterizes the frontier. I have characterized the frontier as an investment in the riskfree yield and a zero-cost payoff on the positive side of the mean-variance frontier, so I will characterize the distortions here in terms of the same two payoffs.

Proposition 6: *Break the outside-income hedge yield into three components, defined by long-run regression*

$$y_e = y^f + \tilde{\beta}_e\left(y^f - y^*\right) + \eta_e. \tag{33}$$

Then the asset yield can be expressed as

$$\hat{y} = y^f + \left[\frac{1}{s_W}\frac{1}{\gamma} - \frac{s_e}{s_W}\tilde{\beta}_e\right]\left(y^f - y^*\right) - \frac{s_e}{s_W}\eta_e \tag{34}$$

Derivation: From Proposition 5,

$$\begin{aligned}
\frac{\hat{x}+\hat{e}}{p(\hat{e})+W} &= \hat{y}s_W + y_e s_e = y^f + \frac{1}{\gamma}\left(y^f - y^*\right) \\
\hat{y} &= \frac{1}{s_W}\left[y^f + \frac{1}{\gamma}\left(y^f - y^*\right)\right] - y_e\frac{s_e}{s_W} \\
\hat{y} &= \frac{1}{s_W}\left[y^f + \frac{1}{\gamma}\left(y^f - y^*\right)\right] - \left[y^f + \tilde{\beta}_e\left(y^f - y^*\right) + \eta_e\right]\frac{s_e}{s_W}
\end{aligned}$$

And with $s_e + s_W = 1$, (34) follows.

To see what (34) means, suppose first that the outside-income hedge payoff is constant across time and states. Then $y_e = y^f$ so (34) becomes

$$\hat{y} = y^f + \frac{1}{s_W}\frac{1}{\gamma}\left(y^f - y^*\right).$$

We get the familiar result that an investor with a large "bond-like" outside income stream and hence low s_W should shift his portfolio more toward risky assets, or behave in a less risk averse manner. His effective risk aversion is reduced by the asset wealth share. (We can also express this result that the investor holds less of the riskfree asset and does not change the risky asset investment.)

In the more general case, the risky (priced) asset position is also modified. If the investor holds an outside income stream with a large beta on the mean-variance efficient yield, he sensibly reduces his risky asset investment. The s_W and s_e terms adjust a simple addition of total payoffs to the more familiar units of yield.

Finally, the investor sells the idiosyncratic component of the outside-income hedge payoff η_e. This payoff, as defined in (33), has zero price since $\mathcal{E}(y^*\eta_e) = 0$, and zero mean since $\mathcal{E}(y^f\eta_e) = y^f\mathcal{E}(1\times\eta_e) = 0$. It constitutes free insurance against outside income risks. Selling (or buying) this payoff ought to be the *first* thing every investor does, since it provides something for nothing. Characterizing such payoffs ought to be the first task of academic portfolio advice. As I speculated in the introduction, perhaps it is often ignored because *returns* on outside income are difficult to compute. The decomposition (33) may help in this regard, since this is a long-run, cashflow regression. We do not need a time series of values of the outside income stream to compute it.

The idiosyncratic component η_e can capture variation over time as well as states of nature. For example, if the investor has a certain wage stream and retires at a given date with certainty, then η_e goes from a positive to a negative loading on the indexed perpetuity on the retirement date, and "short η_e" generates the usual shift from stocks to bonds at that date.

As in equation (32), we get an even clearer result if we multiply (34) by W to express how the values of the various components add up to the optimal asset payoff,

$$\hat{x} = Wy^f + \left\{[W + p(\hat{e})]\frac{1}{\gamma} - p(\hat{e})\tilde{\beta}_e\right\}\left(y^f - y^*\right) - p(\hat{e})\eta_e.$$

Since we are normalizing our mean-variance representation to the riskfree yield as the only non-zero cost payoff, the investor starts by putting all market wealth in that payoff. Then, he invests in the zero-cost long-run mean-variance efficient payoff according to his risk aversion and *total* wealth, including the tradeable component of outside income wealth, but reduced in accordance to the mean-variance exposure he already implicitly owns by owning the outside income stream. Finally, (or, again, initially) he shorts the idiosyncratic component of his outside income stream.

The *nontradeable* component of outside income appears nowhere in these expressions. In general, "background risk" affects investment decisions, for example by making the investor behave in a more risk averse manner. This effect is avoided with quadratic utility. Background risk lowers utility but does not affect linear marginal utility, so does not affect decisions. Again, the point of a simple benchmark is to avoid complex, though potentially important, refinements.

4.2.3 Market payoff

Again, we want to express the optimal portfolio relative to the market payoff. We consider a market of investors who are all of the same type (quadratic utility), but with varying risk aversion (bliss point, initial wealth) and also varying outside income streams. Now investors will think about how their outside income stream differs from the market average, as well as how their risk aversion differs from the market average. Unsurprisingly, our basic result mirrors Proposition 3, but using the long-run mean-variance efficient total payoff yield.

Proposition 7. *The investor's total payoff is proportional to the aggregate total payoff, which is a claim to the traded component of aggregate consumption.*

$$\hat{y}_i^T = y^f + \frac{\gamma_a}{\gamma_i}\left(\hat{y}_a^T - y^f\right), \tag{35}$$

where the yield on the total aggregate payoff is

$$\hat{y}_a^T \equiv \frac{\sum_j \left[W_j + p(\hat{e}_j)\right] \hat{y}_j^T}{\sum_j \left[W_j + p(\hat{e}_j)\right]} = \frac{\sum_j \hat{x}_j + \hat{e}_j}{\sum_j \left[W_j + p(\hat{e}_j)\right]} = \frac{\hat{x}_a + \hat{e}_a}{W_a + p(\hat{e}_a)},$$

and aggregate risk aversion is defined as a wealth-weighted average of individual risk aversion,

$$\frac{1}{\gamma_a} \equiv \frac{\sum_j \left[W_j + p(\hat{e}_j)\right] \frac{1}{\gamma_j}}{\sum_j \left[W_j + p(\hat{e}_j)\right]}.$$

Derivation. Start with (29), sum and divide by wealth,

$$\begin{aligned}
\hat{y}_i^T &= y^f + \frac{1}{\gamma_i}\left(y^f - y^*\right) \qquad (36)\\
\frac{\sum_j \left[W_j + p(\hat{e}_j)\right] \hat{y}_j^T}{\sum_j \left[W_j + p(\hat{e}_j)\right]} &= y^f + \frac{\sum_j \left[W_j + p(\hat{e}_j)\right] \frac{1}{\gamma_j}}{\sum_j \left[W_j + p(\hat{e}_j)\right]}\left(y^f - y^*\right)\\
\hat{y}_a^T &= y^f + \frac{1}{\gamma_a}\left(y^f - y^*\right)\\
y^f - y^* &= \gamma_a\left(\hat{y}_a^T - y^f\right).
\end{aligned}$$

Substitute this result in the right hand side of (36). The interpretation as a claim to the traded component of aggregate consumption follows from the definition, $\hat{y}_a^T = \left(\hat{x}_a + \hat{e}_a\right) / \left(W_a + p(\hat{e}_a)\right) = proj(c_a|\underline{X})/p(proj(c_a|\underline{X}))$.

This advice is harder to implement when there is outside income. Now the investor's total portfolio, accounting for the outside income hedge portfolio he implicitly holds, is split between the index perpetuity (easy) and the market's *total* portfolio $\hat{y}_a^T$, accounting for the outside income hedge portfolio of the average investor. The latter quantity is hard to observe, of course, so it is convenient to reexpress the optimal portfolio with respect to the *traded* market portfolio, which is easy to observe.

Proposition 8. *The investor's asset portfolio can be written in terms of the market asset yield, the market average outside-income hedge yield, and the individual outside-income hedge yield as*

$$\hat{y}_i = y^f + \frac{s_{Wa}\gamma_a}{s_{Wi}\gamma_i}\left(\hat{y}_a - y^f\right) + \frac{s_{Wa}\gamma_a}{s_{Wi}\gamma_i}\frac{s_{ea}}{s_{Wa}}\left(y_{ea} - y^f\right) - \frac{s_{ei}}{s_{Wi}}\left(y_{ei} - y^f\right) \tag{37}$$

where

$$\begin{aligned}
y_{ea} &\equiv \frac{\hat{e}_a}{p(\hat{e}_a)};\quad y_{ei} \equiv \frac{\hat{e}_i}{p(\hat{e}_i)} \\
s_{Wa} &= \frac{W_a}{W_a + p(\hat{e}_a)}; s_{Wi} = \frac{W_i}{W_i + p(\hat{e}_i)} \\
s_{ea} &= \frac{p(\hat{e}_a)}{W_a + p(\hat{e}_a)}; s_{Wi} = \frac{p(\hat{e}_i)}{W_i + p(\hat{e}_i)}
\end{aligned}$$

Derivation. From (35),

$$\hat{y}_i^T = y^f + \frac{\gamma_a}{\gamma_i}\left(\hat{y}_a^T - y^f\right)$$

$$\gamma_i\left(s_{Wi}\hat{y}_i + s_{ei}y_{ei} - y^f\right) = \gamma_a\left(s_{Wa}\hat{y}_a + s_{ea}y_{ea} - y^f\right)$$

$$\gamma_i\left[s_{Wi}\left(\hat{y}_i - y^f\right) + s_{ei}\left(y_{ei} - y^f\right)\right] = \gamma_a\left[s_{Wa}\left(\hat{y}_a - y^f\right) + s_{ea}\left(y_{ea} - y^f\right)\right].$$

Equation (37) follows.

To digest this result, start with the case that both individual and average outside income is a constant over time and states. Now, $y_e = y^f$, so (37) reduces to

$$\hat{y}_i = y^f + \frac{s_{Wa}\gamma_a}{s_{Wi}\gamma_i}\left(\hat{y}_a - y^f\right).$$

We see the same result as in Proposition 7 (35), except that "effective" risk aversion in asset markets is actual risk aversion times the share of wealth. Even an investor who is temperamentally of average risk aversion should invest his asset portfolio in a riskier manner, if he has a larger than average outside income.

The remaining terms of Proposition 8 direct the investor to think about the outside-income hedge portion of his portfolio in relative terms as well. If the investor has average shares and risk aversion $\gamma_i = \gamma_a$, $s_{Wi} = s_{Wa}$, and also the same outside income hedge payoff as the average investor, $\hat{e}_i = \hat{e}_a$, then the last two terms cancel and once again he holds the market payoff $\hat{y}_i = \hat{y}_a$ ignoring outside income. If only the individual and average hedge portfolios are different, (37) reduces to

$$\hat{y}_i = \hat{y}_a + \frac{s_{ea}}{s_{Wa}}\left(y_{ea} - y_{ei}\right).$$

The investor starts with the market portfolio, then holds only the *difference* between aggregate and individual outside income hedge portfolio. He sells to the average investor the payoff that that investor would like to short, and shorts the payoff that best hedges his own outside income.

In the fully general case the term $\frac{s_{Wa}\gamma_a}{s_{Wi}\gamma_i}$ in front of y_{ea} adjusts for larger or smaller effective risk aversion of individual and average. An individual who is very risk averse, either intrinsically or because he has a very large share of asset wealth in his portfolio, will provide less outside-income insurance to the market. The s_e/s_W terms adjust for the difference between each investor's desire to insure the actual flow of income, and our expression in terms of yields, which divide by prices.

As in Proposition 6, it is interesting to orthogonalize the payoffs. The interesting order of orthogonalization in this case is, first, the market asset portfolio, then the market average hedge portfolio, and finally the individual's hedge portfolio. (This is more interesting than orthogonalizing in the order y^*, z^* and η, or $y^f, (y^f - y^*)$ and η.) Thus, define components by successive long-run regressions

$$y_{ea} - y^f = \tilde{\beta}_{ea,a}\left(\hat{y}_a - y^f\right) + \varepsilon_a,$$

$$y_{ei} - y^f = \tilde{\beta}_{ei,a}\left(\hat{y}_a - y^f\right) + \tilde{\beta}_{ei,ea}\varepsilon_a + \eta_{ei}.$$

The three payoffs $\left(\hat{y}_a - y^f\right)$, ε_a, and η_{ei} are all zero cost. In addition, $\mathcal{E}(\eta_{ei}) = 0$. η_{ei} is again a zero-price, zero-mean, "idiosyncratic" component of the outside income hedge payoff. Substituting in (37), we obtain

$$\begin{aligned}\hat{y}_i &= y^f + \left[\frac{s_{Wa}\gamma_a}{s_{Wi}\gamma_i} + \frac{s_{Wa}\gamma_a}{s_{Wi}\gamma_i}\frac{s_{ea}}{s_{Wa}}\tilde{\beta}_{ea,a} - \frac{s_{ei}}{s_{Wi}}\tilde{\beta}_{ei,a}\right]\left(\hat{y}_a - y^f\right) \\ &\quad + \left[\frac{s_{Wa}\gamma_a}{s_{Wi}\gamma_i}\frac{s_{ea}}{s_{Wa}} - \frac{s_{ei}}{s_{Wi}}\tilde{\beta}_{ei,ea}\right]\varepsilon_a - \frac{s_{ei}}{s_{Wi}}\eta_{ei}.\end{aligned} \tag{38}$$

This expression appears formidable, because of the wide variety of differences between individual and average that it accommodates, so it is best examined with special cases. First, verify that an average investor, with $\tilde{\beta}_{ia} = \tilde{\beta}_a$, $\tilde{\beta}_{iea} = 1, \eta_{ei} = 0$ holds the market asset yield $\hat{y}_a$.

Second, consider an investor who is just like the average in risk aversion and shares, but has a different outside income process. He holds

$$\hat{y}_i = y^f + \left[1 + \frac{s_{ea}}{s_{Wa}}\left(\tilde{\beta}_{ea,a} - \tilde{\beta}_{ei,a}\right)\right]\left(\hat{y}_a - y^f\right) + \frac{s_{ea}}{s_{Wa}}\left[1 - \tilde{\beta}_{ei,ea}\right]\varepsilon_a - \frac{s_{ea}}{s_{Wa}}\eta_{ei}.$$

The first term directs him to hold more or less of the market portfolio depending on how *differently* correlated his income is to the market payoff from that of the average investor. If his outside income is uncorrelated with the market, $\tilde{\beta}_{ei,a} = 0$, for example, and the average outside income is so correlated, then he will hold more of the market payoff than average, selling outside income insurance to the average investor. If his outside income responds more to the market than that of the average investor, he will hold less of the market, as part of his outside-income hedge strategy. The second term directs the investor to hold more or less of the orthogonalized aggregate outside income payoff. If the non-market component of the investor's outside-income hedge payoff is uncorrelated with the non-market component of the average outside-income hedge payoff, the investor buys the average hedge payoff, again selling insurance to the market and earning the second factor risk premium. Vice versa, agents whose outside-income hedge payoff moves more than average will sell that payoff, despite losing its factor risk premium. Finally, as always, the investor sells the completely idiosyncratic, zero-price, zero-mean component of his outside-income hedge payoff.

4.2.4 A long-horizon multifactor model

In an equilibrium of investors, all of this type, but with varying outside income streams as well as varying risk aversion, we obtain a multifactor model:

Proposition 9. *The expected long-horizon yield of each asset i follows a multifactor model, with the market payoff and average outside income hedge payoff as factors,*

$$\mathcal{E}(y^i) - y^f = \tilde{\beta}_{i,a}\left[\mathcal{E}(\hat{y}_a) - y^f\right] + \tilde{\beta}_{i,e}\left[\mathcal{E}(y_{ea}) - y^f\right], \tag{39}$$

where $\tilde{\beta}_{i,a}$ and $\tilde{\beta}_{i,e}$ are long-run multiple regression coefficients.

Derivation. Since the total yield $\hat{y}_a^T$ is long-run mean-variance efficient, we have by (15) a single-factor model

$$\mathcal{E}(y^i) - y^f = \tilde{\beta}_{i,T}\left[\mathcal{E}(\hat{y}_a^T) - y^f\right]. \tag{40}$$

Since the total yield is composed of asset and outside income hedge portfolio yields, $\hat{y}_a^T = s_w \hat{y}_a + s_e y_{ea}$, I can show that this single-factor model is equivalent to the stated multifactor model. The multifactor long-run regression is

$$y^i - y^f = \alpha_i + \tilde{\beta}_{i,a}\left(\hat{y}_a - y^f\right) + \tilde{\beta}_{i,e}\left(y_{ea} - y^f\right) + \eta_i.$$

Our goal is to show $\alpha_i = 0$. The following long-run multiple regressions are equivalent, i.e. their intercepts and errors are equivalent, since the right hand variables span the same space. The $\tilde{\beta}$ coefficients are not the same across regressions, which is why I distinguish them with additional symbols.

$$y^i - y^f = \alpha_i + \tilde{\beta}^*_{i,T}\left(\hat{y}_a^T - y^f\right) + \tilde{\beta}^*_{i,e}\left(y_{ea} - y^f\right) + \eta_i$$

$$y^i - y^f = \alpha_i + \tilde{\beta}^+_{i,T}\left(\hat{y}_a^T - y^f\right) + \tilde{\beta}^+_{i,e}\left[\left(y_{ea} - y^f\right) - \alpha - \tilde{\beta}_{ea,T}\left(\hat{y}_a^T - y^f\right)\right] + \eta_i.$$

In the last regression, α and $\tilde{\beta}_{ea,T}$ are single regression coefficients of $y_{ea} - y^f$ on $\hat{y}^T - y^f$. By (40), this $\alpha = 0$. Since the right hand variables are now orthogonal, $\tilde{\beta}^+_{i,T}$ is also the single regression coefficient $\tilde{\beta}_{i,T}$ of $y^i - y^f$ on $\hat{y}^T - y^f$. Taking expectations, then,

$$\mathcal{E}\left(y^i - y^f\right) = \alpha_i + \tilde{\beta}_{i,T}\mathcal{E}\left(\hat{y}_a^T - y^f\right) + \tilde{\beta}^+_{i,e}\left[\mathcal{E}\left(y_{ea} - y^f\right) - \tilde{\beta}_{ea,T}\mathcal{E}\left(\hat{y}_a^T - y^f\right)\right]$$

By (40) again, though,

$$\begin{aligned}
\mathcal{E}\left(y^i - y^f\right) &= \tilde{\beta}_{i,T}\mathcal{E}\left(\hat{y}_a^T - y^f\right) \\
\mathcal{E}\left(y_{ea} - y^f\right) &= \tilde{\beta}_{ea,T}\mathcal{E}\left(\hat{y}_a^T - y^f\right),
\end{aligned}$$

so we are left with $\alpha_i = 0$.

In pricing as in portfolios, Mertonian state variables for outside income as well as investment opportunities disappear, but outside income itself remains. (State variables for outside income are present and potentially important in the traditional one-period ICAPM. Why they have been ignored for 40 years is an interesting question.) When the tradeable component of outside income does not average to zero, a second factor emerges. Assets have higher long-run expected returns if their cashflows have a higher long-run covariance with the dividend stream of the aggregate outside-income portfolio. For example, if as Fama and French (1996) speculate, the outside income of the average investor is correlated with the payouts of a class of "distressed" securities, then those securities will require higher long-run expected returns, and they will receive lower prices. We see a "value" effect in prices and a "value factor" in expected long-run returns.

We also can represent pricing with orthogonalized factors, which are possibly more interesting. Define again ε_a by a long-run regression.

$$y_{ea} - y^f = \tilde{\beta}_{ea,a}\left(\hat{y}_a - y^f\right) + \varepsilon_a$$

Then, our multifactor model becomes:

$$\mathcal{E}(y^i) - y^f = \tilde{\beta}_{i,a}\left[\mathcal{E}(\hat{y}_a) - y^f\right] + \tilde{\beta}_{i,\varepsilon}\left[\mathcal{E}(\varepsilon_a) - y^f\right] \tag{41}$$

Comparing the pricing results (39) and (41) to the portfolio expressions (37) and (38), we see the same right hand variables. The portfolio expressions (37) and (38) tell the investor how much to put in to the "priced assets" corresponding to the aggregate market portfolio and aggregate outside-income-hedge portfolio, and then to perfectly hedge residual, zero-price, mean-zero, risk. The portfolio shares advocated by (37) and (38) are given by "risk aversions" which combine true risk aversion, and aversions induced by the character of outside income.

5 Concluding comments

I do not present illustrative calculations. Interesting calculations require nontrivial modeling and data choices, whose solution rises far above the "illustrative" category. I can, however, discuss some of the challenges. The web-appendix includes details of some calculations I report.

Lognormal i.i.d.

The lognormal i.i.d. case is a natural first environment for examining any portfolio theory, i.e. a constant riskfree rate r^f and risky assets that follow $dr = \left(\mu + r^f\right)dt + \sigma dB$. The web-appendix finds the long-run mean-variance frontier, and contrasts the portfolios that support optimal payoffs for power utility and quadratic utility in this environment. To form a long-run mean-variance efficient yield, The quadratic utility investor still holds a conditionally mean-variance efficient portfolio with weights $w = 1/\gamma_t \Sigma^{-1}\mu$, and consumes more as wealth rises. The difference is that local risk aversion γ_t varies over time in the quadratic case, declining as wealth rises, and the consumption-wealth relation has an intercept.

Alas, this calculation is not quantitatively realistic, because the long right tail of lognormal returns and mean-variance analysis do not mix well. Discrete-time returns in the lognormal i.i.d.

environment have arithmetic Sharpe ratios that go to zero as the horizon increases, while the maximum arithmetic Sharpe ratio available from dynamic trading goes to infinity. (It's tempting to take logs first, but mean-variance analysis applies to arithmetic, not log, returns.) As a result of this behavior, long-run mean-variance frontiers and quadratic utility are poorly behaved for the large risk premia we observe in the data. When $2r^f - \rho - \mu'\Sigma^{-1}\mu \leq 0$, finitely-lived quadratic utility investors finance early consumption by disastrous falls in consumption late in life or when markets fall; as the lifetime increases, this repayment is indefinitely extended, which means the investor exceeds any borrowing limit. The limit point seems to violate arbitrage. To solve the quadratic-utility infinite-horizon economy with large risk premia, we need to impose borrowing and arbitrage constraints, and they will bind, losing much of the simplicity which is the whole reason for examining the rather unrealistic utility function. At the same time x^* ceases to be tradeable, as its price goes to infinity, and $\mathcal{E}(x^{*2}) < \infty$ is violated.

This is not a "fatal flaw" for quantitative application (as opposed to merely useful conceptual benchmark) of these ideas. As I document in the web-appendix, actual index returns are far from lognormal. The long right tail predicted by the lognormal is missing, while the fat left tail of short horizon returns also disappears. Even at a 10 year horizon, index returns are better described by a normal rather than lognormal distribution. This finding is not that surprising: we know that there is some mean-reversion in returns, and that volatility decreases when the market rises. Both effects cut off the large troublesome right tail of the lognormal. However, it does mean that a quantitatively realistic calculation (one that violates $2r^f - \rho - \mu'\Sigma^{-1}\mu > 0$) must incorporate at least stochastic volatility and potentially mean-reversion, to say nothing of additional state variables, exceeding by far the back of any envelope.

Related, quadratic utility investors in equilibrium *choose* a market return that is not lognormal. When we specify $dr = \left(\mu + r^f\right) dt + \sigma dz$, we are specifying the underlying technologies. As wealth rises, investors individually and collectively rebalance away from risky technologies, and the market becomes less risky, so the market portfolio return endogenously loses the large right tail of a lognormal. If one takes the equilibrium point of view, or makes any calculations of portfolios relative to the market index, it does not make theoretical sense to specify a lognormal distribution for the market index.

Time-varying returns with unspanned state variables

Long-run mean-variance calculations with time-varying investment opportunities and incomplete markets are also harder than they may appear. For example, a typical environment might be a constant risk free rate r^f and risky returns that follow

$$\begin{aligned} dr_t &= \left[r^f + \mu(z_t)\right] dt + \sigma(z_t) dB_{1t} \\ dz_t &= \mu_z(z_t)dt + \sigma_{z1}(z_t)dB_{1t} + \sigma_{z2}(z_t)dB_{2t}. \end{aligned}$$

The discount factor in this situation must be of the form

$$\frac{d\Lambda_t}{\Lambda_t} = -r^f dt - \frac{\mu(z_t)}{\sigma(z_t)} dB_{1t} + \sigma_{\Lambda 2t} dB_{2t}$$

where $\sigma_{\Lambda 2t}$ is arbitrary. Choosing correctly $\sigma_{\Lambda 2t}$ so that $e^{-\rho t} u'(x_t) = \lambda \Lambda_t$ produces a tradeable x_t is precisely the central difficulty of applying the discount-factor approach in this situation. One might think that the *traded* discount factor is simply $\sigma_{\Lambda 2t} = 0$, cutting through the central difficulty for quadratic utility. After all, the resulting discount factor is the only one whose shocks are spanned by the shock dB_{1t} to the traded assets.

Alas, this is a subtle mistake. "Traded" means "an achievable *payoff* from a valid dynamic trading strategy." Having shocks spanned by the asset return shocks $\sigma_{\Lambda 2t} = 0$ is neither necessary or sufficient for Λ_t or $x_t^* = e^{\rho t}\Lambda_t$ to be achievable as the dividend process of a trading strategy. Typically, in fact, traded discount factors will have dB_2 loadings, $\sigma_{\Lambda 2t} \neq 0$.

As a very simple example in which to see this point, suppose the interest rate varies over time,

$$dr^f = \mu(r_t^f)dt + \sigma(r_t^f)dB_t,$$

and there are no other assets. The marginal utility of a quadratic utility investor moves immediately when there is news about investment opportunities dB_t. Loosely, if investment opportunities improve, the investor increases consumption immediately so that the first order condition and the wealth constraint both continue to be satisfied. Yet the discount factors are

$$\frac{d\Lambda_t}{\Lambda_t} = -r_t^f dt - \sigma_{\Lambda t} dB_t,$$

so the choice $\sigma_{\Lambda t} = 0$ means that marginal utility does not move when there is interest rate news. If marginal utility followed such a discount factor, either the first order condition or the wealth constraint would have to be violated – this is not a tradeable discount factor.

To see the point a little more formally in this example, consider a finite-horizon version with payoffs from 0 to T and specialize to $dr_t^f = \sigma dB_t$. A "traded payoff" is an $\{x_t\}$ that results from the trading strategy

$$dV = \left(r_t^f V_t - x_t\right) dt \tag{42}$$

with $V_T = 0$. Now you can see how a traded payoff x_t *can* load on unspanned shocks: you can choose to draw down x_t out of wealth in a way that responds to any variable. However, not every payoff is traded. Most obviously, if x_t draws down wealth based on some non-traded shock, some future x must pay back the debt. In this way, looking backwards, i.e. accumulating wealth by (42), an arbitrary payoff x_t will violate the terminal wealth constraint $V_T = 0$. Looking forwards, the time-t value of an arbitrary payoff,

$$V_t = E_t \int_{s=t}^{T} \frac{\Lambda_s}{\Lambda_t} x_s ds,$$

may result in dV_t that loads on dB_t, while a tradeable $\{x_t\}$ must result in a dV_t implementable by (42), with no loading on the unspanned shock dB_t.

We can see therefore in this simple example that a traded discount factor *must* load on dB_t in this way, by supposing the contrary and showing that the corresponding value process does load on dB. Try $\sigma_\Lambda = 0$. Then, the value of the portfolio that delivers $\{\Lambda_t\}$ as a payoff is, at any point in time,

$$V_t = \frac{1}{\Lambda_t} E_t \int_{s=t}^{T} \Lambda_s^2 ds = \Lambda_t E_t \int_{s=t}^{T} e^{-2\int_{\tau=t}^{s} r_\tau^f d\tau} ds.$$

Differentiating takes a little work, relegated to the appendix, but the result is

$$dV_t = (r_t^f V_t - \Lambda_t) - 2\sigma\Lambda_t \left[\int_{s=t}^{T} (s-t) e^{-2(s-t)r_t^f + \frac{2}{3}\sigma^2 (s-t)^3} ds\right] dB_t \tag{43}$$

This value loads on dB_t, so cannot be traded.

Finite basis

In conventional mean-variance analysis, as in one-period return based asset pricing more generally, we typically do not try directly to compute mean-variance frontiers of 8000 US stocks, let alone the plethora of other available assets. Instead, we typically form a much smaller number of portfolios first. Implicitly, we assume that the interesting variation in the larger set of securities is spanned by the much smaller number of portfolios.

The same approach may be valuable in addressing dynamic issues. Rather than try to find the exact optimum in an infinite-dimensional space of portfolio weights and payout rules, as I have done in the discussion so far, we can simply include a finite number of well-chosen dynamic trading strategies – a few well-chosen rules for portfolio weights w_t and payout rules x_t as a function of state variables – and then consider a static maximization over those options. (Brandt and Santa-Clara (2006) advocate and implement this strategy.) If we make this simplification, most of the technical difficulties vanish. With a finite vector of well-behaved basis payoffs $\mathbf{x}$, we can easily find a traded discount factor $x^* = \frac{1}{k}\mathbf{p}'\mathcal{E}(\mathbf{x}\mathbf{x}')^{-1}\mathbf{x}$ and the optimal payoffs, which will be of the form $\mathbf{c}'\mathbf{x}$. Since mean-variance analysis is really a benchmark rather than an approximation, this approach may also be useful when it is a *poor* approximation to the exact optimum, by forcing a choice among sensible options when that exact optimum exploits heavily the peculiarities of quadratic utility. However, as in conventional mean-variance analysis, the set of underlying portfolios, dynamic trading rules, and dynamic payout rules, must all be artfully chosen to capture the important variation in the assets at hand, rendering this approach as well one that takes more than back-of-the envelope work.

Outside income

Finally, the most interesting calculations in this setup may well be the calculations of hedge payoffs for outside income. A quick look at the fundamental expectation

$$\mathcal{E}(x) = \rho \int_0^\infty e^{-\rho t} x_t dt$$

reveals the obvious issues. Handling trends is particularly important, since time trends are a large part of "long-run variance." Estimating directly means and variances of this sort, like estimating spectral densities near zero, is obviously difficult with "short" datasets. It is attractive to estimate these quantities based on time-series models rather than simply substituting in data realizations. However, imputing long-run implications of short-run models has always been a dicey affair, and short-run model performance statistics can be misleading. At a minimum, specifying trends and cointegrating relationships is vital.

Bottom line.

Again, the difficulty of calculating long-run mean-variance frontiers in interesting applied situations is, in a sense, a feature not a bug, as is the difficulty of calculating one-period mean-variance frontiers. Both ideas are primarily useful as conceptual benchmarks, for the analysis that follows *given* a mean-variance efficient payoff or portfolio, for characterizing investor choices given a well-constructed set of funds, and for directing efforts at the hard, technical, financial-engineering task to the kinds of payoffs that investors desire.

6 References

Arrow, Kenneth J., and Gerard Debreu, 1954, Existence of an equilibrium for a competitive economy, *Econometrica* 22, 265-290.

Bansal, Ravi, Robert F. Dittmar, Christian Lundblad, 2005, Consumption, dividends, and the cross-section of equity returns, *Journal of Finance* 60, 1639–1672.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481-1509.

Barberis, Nicholas, 2000, Investing for the long run when returns are predictable, *Journal of Finance* 55, 225-264.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, 2007, Portfolio choice over the life cycle when the stock and labor markets are cointegrated, *Journal of Finance* 62, 2123-2167.

Brandt, Michael W., and Pedro Santa-Clara, 2006, Dynamic portfolio selection by augmenting the asset space, *Journal of Finance* 61, 2187-2217.

Brennan, Michael J., Eduardo S. Schwartz and Roland Lagnado, 1997, Strategic asset allocation, *Journal of Economic Dynamics and Control* 21, 1377-1403.

Brennan, Michael J. and Yihong Xia, 2000, Stochastic interest rates and bond-stock mix, *European Finance Review* 4, 197-210.

Brennan, Michael J. and Yihong Xia, 2002, Dynamic asset allocation under inflation, *Journal of Finance* 57, 1201-1238.

Campbell, John Y., 1996, Understanding risk and return, *Journal of Political Economy* 104, 298–345.

Campbell, John Y. and Luis M. Viceira, 2001, Who should buy long-term bonds?" *American Economic Review*, 91, 99-127.

Campbell, John Y. and Luis M. Viceira, 1999, Consumption and portfolio decisions when expected returns are time varying, *Quarterly Journal of Economics* 114, 433-95.

Cochrane, John H., 2004. *Asset Pricing, Revised Edition* (Princeton University Press).

Cox, John C. and Chi-fu Huang, 1989, "Optimal Consumption and Portfolio Policies When Asset Prices Follow a Diffusion Process," *Journal of Economic Theory,*49, 33-83.

Davis, Steven, Felix Kubler, and Paul Willen, 2006, Borrowing costs and the demand for equity over the life cycle, Review of Economics and Statistics 88, 348-362.

Davis, Steven, and Paul Willen, 2000, Occupation-level income shocks and asset returns: their covariance and implications for portfolio choice, NBER Working Paper 7905.

Davis, Steven J. and Paul Willen, 2001, Income shocks, asset returns, and portfolio choice, in Zvi Bodie, Brett Hammond and Olivia Mitchell, eds., *Innovations in Retirement Financing* (University of Pennsylvania Press).

Duffie, Darrell, 2001 *Dynamic Asset Pricing Theory*, 3d Edition (Princeton University Press).

Duffie, Darrell, Wendell Fleming, Mete Soner, and Thaleia Zariphopoulou, 1997, Hedging in incomplete markets with HARA utility, *Journal of Economic Dynamics and Control* 21, 753–782.

Fama, Eugene F. and Kenneth R. French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55-84

Flavin, Marjorie, and Takashi Yamashita, 2002, Owner-occupied housing and the composition of the household portfolio, *American Economic Review* 92, 345–362.

Gabaix, Xavier, 2007, Linearity-generating processes: A modelling tool yielding closed forms for asset prices, Manuscript, New York University.

Hansen, Lars Peter, 1987, Calculating asset prices in three example economies, in Truman F. Bewley, ed.: *Advances in Econometrics, Fifth World Congress* (Cambridge University Press).

Hansen, Lars Peter, John C. Heaton and Nan Li, 2005, Consumption strikes back?, Manuscript, University of Chicago, forthcoming *Journal of Political Economy.*

Hansen, Lars Peter, and Scott F. Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587-613.

Hansen, Lars Peter and Thomas J. Sargent, 2004, *Recursive Models of Dynamic Linear Economies*, Manuscript, University of Chicago and New York University.

He, Hua, and Neil D. Pearson, 1991, Consumption and portfolio policies with incomplete markets and short-sale constraints: the infinite-dimensional case, *Journal of Economic Theory*, 54, 259-304.

He, Hua and Henri F. Pagès, 1993, Labor income, borrowing constraints, and equilibrium asset prices, *Economic Theory* 3, 663-696.

Heaton, John, and Deborah J. Lucas, 2000a, Portfolio choice in the presence of background risk, *Economic Journal*, 110, 1-26.

Heaton, John, and Deborah J. Lucas, 2000b, Portfolio choice and asset prices: The importance of entrepreneurial risk, *Journal of Finance* 55, 1163–1198.

Henderson, Vicky, 2005, Explicit solutions to an optimal portfolio choice problem with stochastic income *Journal of Economic Dynamics and Control*, 29, 1237-1266.

Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3-53.

Kim, Tong Suk and Edward Omberg, 1996, Dynamic nonmyopic portfolio behavior, *Review of Financial Studies* 9, 141-61.

Koo, Hyeng Keun, 1998, Consumption and portfolio selection with labor income: A continuous-time approach, *Mathematical Finance* 8, 49-65.

Lettau, Martin and Jessica Wachter, 2007, Why is long-horizon equity less risky? A duration-based explanation of the value premium, *Journal of Finance* 62, 55-92.

Liu, Jun, 2007, Portfolio selection in stochastic environments, *Review of Financial Studies* 20, 1–39.

Lynch, Anthony, 2001, Portfolio choice and equity characteristics: Characterizing the hedging demands induced by return predictability, *Journal of Financial Economics* 62, 67-130.

Lynch, Anthony, and Sinan Tan, 2008, Labor income dynamics at business-cycle frequencies: Implications for portfolio choice, NBER Working Paper 11010.

Markowitz, Harry, 1952, Portfolio selection, *Journal of Finance*, 7, 77-99.

Magill, Michael and Martine Quinzii, 2000, Infinite-horizon CAPM equilibrium, *Economic Theory* 15, 103-138.

Menzly, Lior, Tano Santos and Pietro Veronesi, 2004, Understanding predictability, *Journal of Political Economy* 112, 1-47

Merton, Robert C., 1969, Lifetime portfolio selection under uncertainty: The continuous time case," *Review of Economics and Statistics* 51, 247-257.

Merton, Robert C., 1971a, Optimum consumption and portfolio rules in a continuous time model, *Journal of Economic Theory* 3, 373-413.

Merton, Robert C., 1971b, An intertemporal capital asset pricing model, *Econometrica* 41, 867-87.

Munk, Claus, 2000, Optimal consumption/investment policies with undiversifiable income risk and liquidity constraints, *Journal of Economic Dynamics and Control* 24, 1315-1343.

Ross, Stephen A., 1978, A simple approach to the valuation of risky streams, Journal of Business 51, 453-475.

Sangvinatsos, Antonios and Jessica Wachter, 2005, Does the failure of the expectations hypothesis matter for long-term investors?, *Journal of Finance* 60, 179–230.

Schroder, Mark and Costis Skiadas, 1999, Optimal consumption and portfolio selection with stochastic differential utility, *Journal of Economic Theory*, 89, 68-126.

Svensson Lars E. O. and Ingrid M. Werner, 1993, Non traded assets in incomplete markets, *European Economic Review* 37, 1149-1168.

Teplá Lucie, 2000, Optimal hedging and valuation of nontraded assets, *European Finance Review* 4, 231-251.

Viceira, Luis, 2001, Optimal portfolio choice for long horizon investors with nontradeable labor income, *Journal of Finance* 56, 433-470.

Wachter, Jessica, 2002, Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets, *Journal of Financial and Quantitative Analysis* 37, 63–91.

Wachter, Jessica, 2003, Risk aversion and allocation to long-term bonds, *Journal of Economic Theory* 112, 325–333.

Yao, Rui, and Harold H. Zhang, 2005, Optimal consumption and portfolio choices with risky housing and borrowing constraints, *Review of Financial Studies* 18, 197–239.

7 Appendix

7.1 Notation

$\mathcal{E}(x)$: Long-term mean, e.g. $\mathcal{E}(x) = \rho E \int_0^\infty e^{-\rho t} x_t dt$.

$\tilde{\sigma}^2(x) \equiv \mathcal{E}(x^2) - [\mathcal{E}(x)]^2$: Long-term variance.

$\tilde{\beta}, \widetilde{cov}$: Long-term regression coefficient and covariance, $\widetilde{cov}(x,y) = \mathcal{E}(xy) - \mathcal{E}(x)\mathcal{E}(y)$; $\tilde{\beta}_{x,y} = \widetilde{cov}(x,y)/\tilde{\sigma}^2(x)$.

β, ρ: Discount rate and weighting function.

$x, x_t, \{x_t\}$: Stream of payoffs (dividends plus purchase-sales).

$p(x)$: Price of the stream $\{x_t\}$.

m_t: Scaled discount factor, e.g. $p(x) = E\sum_{t=1}^\infty \beta^t m_t x_t$.

y, y_t: Yield, payoff to a one dollar investment, $y_t = x_t/p(x)$.

x^f, y^f: Riskfree payoff and yield, $x^f = 1$, $y^f = 1/p(x^f)$.

z: Excess yield, payoff to costless investment, e.g. $z = y - y^f$.

$\underline{X}, \underline{Y}, \underline{Z}$: Sets of available payoffs, yields, excess yields.

$\mathbf{x}, \mathbf{y}, \mathbf{z}$**:** Vectors of N basis payoffs, yields, or excess yields.

x^*: Discount-factor mimicking payoff, i.e. $x^* \in \underline{X}$, $p(x) = k\mathcal{E}(x^* x)$.

$y^* = x^*/p(x^*)$: Minimum long-run second moment yield.

$z^* = \left(y^f - y^*\right)/y^f$: Mean-generating excess yield.

y^{mv}, z^{mv}: Mean-variance efficient yield, excess yield.

η: Idiosyncratic component of a yield, $y = y^* + wz^* + \eta$, and $p(\eta) = 0, \mathcal{E}(\eta) = 0$.

x^i, y^i, z^i: Generic ith asset.

$\hat{x}, \hat{y}$: Optimal payoff, yield of optimal payoff.

$\hat{y}_i, \hat{y}_a$: Yield on investor i's optimal payoff, and market average.

y_e, y_{ei}, y_{ea}: Yield on outside income hedge payoff, $y_e = \hat{e}/p(\hat{e})$, individual value, and market average.

$\hat{y}^T, \hat{y}_i^T, \hat{y}_a^T$: Optimal total payoff, including asset and hedge payoff $\hat{y}^T \equiv (\hat{x} + \hat{e})/(W + p(e))$, individual value, and market average.

W, W_t: Initial wealth and wealth at time t.

e, e_t: Stream of outside labor or business income.

$\hat{e}$: Hedge payoff for outside income , $\hat{e} = proj(e|\underline{X})$. (Includes preference shocks as well, $\hat{e} = \hat{e} - \tilde{c}^b$.)

η_e, η_{ei}: Idiosyncratic, zero-price, zero-mean component of outside-income hedge yield.

ε_a: Residual in long run regression, and orthogonalize outside-income hedge factor payoff, $y_{ea} - y^f = \tilde{\beta}_{ea,a}\left(\hat{y}_a - y^f\right) + \varepsilon_a$.

s_W, s_e: Share of asset and outside wealth, $s_W = W/\left[W + p(\hat{e})\right], s_e = p(\hat{e})/\left[W + p(\hat{e})\right]$.

$s_{Wa}, s_{ea}, s_{Wi}, s_{ei}$: Market average and individual shares.

c^b: Bliss point of quadratic utility, $u(c) = -1/2\left(c^b - c\right)^2$.

$\hat{c}^b$: Hedge payoff for the bliss point, $\hat{c}^b = proj(c^b|\underline{X})$.

$\bar{c}^b$, $\tilde{c}^b$: Constant and variable components of the preference shock hedge payoff, $\hat{c}^b = \bar{c}^b \times 1 + \tilde{c}^b$; $\mathcal{E}\left(\tilde{c}^b\right) = 0$.

γ: Relative risk aversion coefficient, defined locally at c, $1/\gamma = (c^b - c)/c$.

γ_i, γ_a: Investor i's risk aversion, and wealth-weighted market average risk aversion.

7.2 Algebra

This appendix derives equation (43).

$$
\begin{aligned}
V_t &= \Lambda_t E_t \int_{s=t}^{T} e^{-2\int_{\tau=t}^{s} r_\tau^f d\tau} ds. \\
&= \Lambda_t \int_{s=t}^{T} e^{-2(s-t)r_t^f + 2\sigma_t^2 \left(\int_{\tau=t}^{s} r_\tau^f d\tau\right)} ds \\
&= \Lambda_t \int_{s=t}^{T} e^{-2(s-t)r_t^f + 2\sigma_t^2 \left(\int_{\tau=t}^{s} d\tau \int_{u=t}^{\tau} dB_u\right)} ds \\
&= \Lambda_t \int_{s=t}^{T} e^{-2(s-t)r_t^f + 2\sigma_t^2 \left(\int_{\tau=t}^{s} (s-\tau) dB_\tau\right)} ds \\
&= \Lambda_t \int_{s=t}^{T} e^{-2(s-t)r_t^f + 2\sigma^2 \left(\int_{\tau=t}^{s} (s-\tau)^2 d\tau\right)} ds \\
&= \Lambda_t \int_{s=t}^{T} e^{-2(s-t)r_t^f + \frac{2}{3}\sigma^2 (s-t)^3} ds.
\end{aligned}
$$

Remembering $d\Lambda/\Lambda = -r^f dt$, we have

$$
dV_t = -r_t^f V_t dt + \Lambda_t d\left[\int_{s=t}^{T} e^{-2(s-t)r_t^f + \frac{2}{3}\sigma^2 (s-t)^3} ds\right]. \tag{44}
$$

To evaluate the second term,

$$
\begin{aligned}
& d\left[\int_{s=t}^{T} e^{-2(s-t)r_t^f + \frac{2}{3}\sigma^2 (s-t)^3} ds\right] \\
= \;& -1dt + \int_{s=t}^{T} \left(\left[2r_t^f - 2\sigma^2 (s-t)^2\right] dt - 2(s-t)dr^f + 2(s-t)^2 dr^{f2}\right) e^{-2(s-t)r_t^f + \frac{2}{3}\sigma^2 (s-t)^3} ds \\
= \;& -1dt + \int_{s=t}^{T} \left(\left[2r_t^f - 2\sigma^2 (s-t)^2\right] dt - 2(s-t)\sigma dB_t + 2(s-t)^2 \sigma^2 dt\right) e^{-2(s-t)r_t^f + \frac{2}{3}\sigma^2 (s-t)^3} ds \\
= \;& -1dt + \frac{V_t}{\Lambda_t} 2r_t^f dt - \left[\int_{s=t}^{T} (s-t) e^{-2(s-t)r_t^f + \frac{2}{3}\sigma^2 (s-t)^3} ds\right] 2\sigma dB_t. \qquad (45)
\end{aligned}
$$

Substituting, (45) in (44), we obtain (43).