#1 - 2012

NBIM DISCUSSION NOTE

This paper was part of the NBIM memo "On rebalancing" (February 2012)

Time-varying expected returns and investor heterogeneity: foundations for rebalancing

30/03/2012

What is the optimal rebalancing policy for a portfolio's equity and bond holdings? The classical answer, building on the seminal contributions by Mossin (1968), Merton (1969, 1971), Samuelson (1969) and others, is that investors should hold a constant proportion of bonds and equities in their portfolios. During the 1980s and 1990s, empirical and theoretical research began to challenge the fundamental premises for this result.

It has, for example, been documented that risk premia and expected returns vary over time. An implication of this is that the Mossin-Merton-Samuelson result is incompatible with market clearing. The risks associated with a portfolio with a constant proportion of bonds and equities might be very different from what was previously understood. A framework to analyse the expected returns and risks associated with a dynamic rebalancing regime must be based on models that account for investors' risk preferences and time-varying risk premia.

Analytical frameworks based on recursive risk preferences, long-run risk and time-varying volatility give a rationale for a dynamic rebalancing regime where different investors put different weights on different risk factors, such as short-run risk, long-run risk and volatility. Investor "preference heterogeneity" should not be interpreted literally as heterogeneity in "preferences", but rather as reduced-form representations of other forms of heterogeneity associated with, for example, market frictions or market incompleteness.

Introduction

Financial economics is about the valuation of cash flows that extend over time and are usually uncertain. Portfolio strategies are based on theories for valuing financial assets. The basic intuition that underlies valuation is the absence of arbitrage. In simple terms, an arbitrage opportunity is a money pump; an opportunity to consistently buy "cheap" and sell "expensive". Clearly, almost all investors would want to take advantage of such an opportunity. Equally clearly, in competitive markets such an opportunity cannot persist: when investors try to take advantage of a known arbitrage opportunity it will cease to exist.

In a separate note, we documented that, in a sample from 1970 to 2010, simple portfolio-rebalancing rules lead to higher expected returns and lower period-by-period standard deviations. From these findings, two questions might naturally arise. First, is a simple rebalancing regime an arbitrage opportunity that nobody has ever thought of before – or does rebalancing expose the owner of a portfolio to risks not accounted for by period-by-period standard deviations in a 40-year sample? And second, if an investor finds it optimal to rebalance, why would any other portfolio manager find it optimal to buy the assets the investor is selling and sell the assets the investor is buying? In this enclosure we will try to answer these and related questions.

All strategic investment decisions are, either explicitly or implicitly, based on a stylised representation of reality – an "analytical framework" or a "model". One challenge is to be aware of the modelling frameworks that are employed and have a thorough understanding of their strengths and weaknesses. Recent research in financial economics provides some analytical frameworks to answer the questions posed above.

For a large, long-term investor, an analytical framework for how to optimally time the market and adjust to time-varying expected returns needs to meet the following requirements. First, it ought to have some mechanism for accounting for key asset-markets features, including time-varying risk premia, and hence time-varying expected returns. Second, it ought to be based on a realistic approach to investment strategies and an implicit acknowledgement that other investors are doing as best they can, given the constraints they are subject to. And third, it ought to explicitly take market-clearing in to account, which in simple terms means that the market portfolio has to be owned by somebody: if one investor has sold an asset, there must have been somebody else who has bought that asset.

Mossin (1968), Merton (1969, 1971), Samuelson (1969) and others showed that, under a set of general, but restrictive assumptions, investors should optimally hold constant bond and equity portfolio weights, irrespective of time horizon and wealth level. Among the critical assumptions for this result were that expected returns are constant (expected returns are independently and identically distributed) and investors' risk preferences can be summarised by standard constant relative risk aversion. The future equity risk premium was proxied by the historical average equity excess return. Applied to the question of portfolio rebalancing, these results imply that investors should hold the same, constant proportion of bonds and equities in their portfolios.

During the 1980s and 1990s, empirical and theoretical research began to challenge whether risk premia in stock and bond markets are constant and could be proxied by historical average excess returns (see, for example, Ilmanen, 2011, pp. 57-58, for an overview and discussion). Shiller (1981) and LeRoy and Porter (1981) challenged the assumption of constant expected returns by documenting that, if risk premia are constant, financial markets seemed to be excessively volatile. Fama and French (1988) documented that the price-dividend ratio predicts realised returns, and Fama and French (1989) documented countercyclical returns for both stocks and bonds.

These findings undermine the classical justification for constant portfolio weights. First, they falsify the crucial assumption of constant expected returns. Second, ignoring any issues about how markets clear, a partial analysis of optimal portfolio weights for an investor with traditional constant relative risk aversion implies time-varying portfolio weights. As we will show later, in simple terms, this means that if one individual investor is oblivious to risks other investors take into account, she or he will naturally hold a different portfolio from other investors. Third, and most importantly, it is incompatible with market clearing. If risk premia are time-varying and equity and bond prices have different sensitivities to changes in risk premia, the ratio of the market capitalisations of bonds and equities will be time-varying. The implication of any analytical framework that takes market clearing into account will be that the individual investors' bond and equity portfolio weights will move with the relative sizes of the bond and equity market capitalisations.

Two analytical frameworks have emerged as de facto standards for accounting for volatility of financial markets and predictability of excess returns: the habit model of Campbell and Cochrane (1999), and the long-run risk model of Bansal and Yaron (2004) drawing on a recursive specification of time and risk preferences in Epstein and Zin (1989). Assuming that all investors are identical, these frameworks are also an implicit acknowledgment that other investors are also doing as best they can. However, as we will get back to, most likely, only models with recursive preferences can take market-clearing into account and still provide a rationale for a rebalancing regime with weights that differ from market capitalization weights.

Recursive specifications of time and risk preferences account for how investors not only care about period-by-period return fluctuations, but also about long-term returns and the uncertainty associated with these. Within these analytical frameworks, decision problems become inherently dynamic and investors not only care about distributions over outcomes, but also, for example, when uncertainty is resolved. These preference specifications provide analytical frameworks within which it is possible to analyse dynamic asset-allocation questions and rebalancing. Differences in investors' preferences for the resolution of time-varying uncertainty might give rise to different rebalancing strategies.

For analytical tractability, when analysing rebalancing within models where investors have heterogeneous recursive preferences over time and risk, we assume that markets are frictionless and complete. There are, however, few reasons to believe that markets are either frictionless or complete. For example, frictions such as credit constraints may prevent some investors from borrowing precisely when they need the extra liquidity; some investors might also be regulated such that they reduce their risk exposure in times when expected returns are high. Finally, investors might be overconfident or in other ways boundedly rational, and their decision structures might be such that it is hard to actually implement long-term time-consistent investment strategies.

Investor "preference heterogeneity" should therefore not be interpreted literally as heterogeneity in preferences narrowly defined. Instead, it could be interpreted as reduced-form representations of other forms of heterogeneity associated with market frictions or market incompleteness.

In this note, we will

1. Review the evidence for time variation in risk premia and expected return variation unrelated to market betas

2. Describe frictionless theories with preference heterogeneity that might account for rebalancing

3. Give a brief overview of some economic theories that might account for sources of preference heterogeneity

Excess volatility and predictability

During the last 30 years, several key aspects of asset market data have posed serious challenges to economic models. For a long time, it proved difficult to account for the observed equity premium and the low risk-free rate (see Mehra and Prescott, 1985; Weil, 1989; Hansen and Jagannathan, 1991). The literature on variance bounds highlights the difficulty in accounting for the high market volatility (see Shiller, 1981; LeRoy and Porter, 1981). The conditional variance of the market return, as shown by Bollerslev et al. (1988), fluctuates across time and is very persistent. Price-dividend ratios seem to predict long horizon equity returns (see Campbell and Shiller, 1989).

Later research has shown that these "puzzles" and empirical regularities are closely connected. All the serious challenges to economic models listed above – often referred to as "puzzles" – are defined conditional on measures and definitions of the price of risk. Most analytical frameworks, including the canonical capital-asset-pricing model, assumed independent and identically distributed returns; in other words that expected returns don't change over time.

The early literature on excess equity and price-dividend ratio volatility following Shiller (1981) and LeRoy and Porter (1981) emphasised "efficiency" and the importance of changes in the information sets. Since then, researchers have found that the actual empirical phenomenon of excess volatility observed in the data might be a phenomenon of forecastable excess stock market returns. Early contributions, such as Campbell (1991) and Cochrane (1992), conclude that nearly all the variation in the price-dividend ratio is attributable not to variation in expected growth in future real cash flows accrued by the owners ("dividends"), but how the price-dividend ratio co-varies with future returns. This suggests that risk-premia move over time and account for the majority of variability in the price-dividend ratio.

Cochrane (2005, Table 20.3, p. 400) provides an estimate of the magnitude of excess volatility by providing empirical decompositions of the variation in the price-dividend ratio, using annual data for value-weighted NYSE stock prices:[1]

$$\frac{100 \ \times cov\,(p_t - d_t, \sum_{j=0}^{15} \rho^j \Delta d_{t+1+j})}{var(p_t - d_t)} = -34$$

$$\frac{100 \ \times cov\,(p_t - d_t, \sum_{j=0}^{15} \rho^j \Delta r_{t+1+j})}{var(p_t - d_t)} = 138$$

These estimates show that almost all the variation in the price-dividend ratio is attributable to changing return forecasts (the price-dividend ratio co-varies with future returns). This means that almost all of the variation in the price-dividend ratio is attributable to excess volatility.

Not shown in the above numbers – but what is now a well-known empirical phenomenon – is that almost all of this return forecastability is attributable to forecastable variation in the excess returns. In other words, the covariance of the price-dividend ratio with future excess returns – and not the covariances with future dividends or risk-free rate – accounts for almost all of the variability in the price-dividend ratio given by the figure of 138 per cent above.

Virtually none of the variability is attributable to the forecastability of the riskless real interest rate (see Campbell et al. 1997 and Cochrane 2005 for overviews). Thus, empirically, high prices, relative to dividends, reflect lower expected excess returns or lower risk-premia. Similarly, Campbell (1991) finds that almost all of the variability in return surprises is attributable to revisions in expectations of future excess returns, rather than in the riskless interest rate. All of this is evidence of long-horizon excess stock return predictability. For further details on asset-market volatility and time-varying excess returns, see Appendix.

1 The decomposition does not have to lie between 0 and 100 per cent: high prices (relative to dividends) seem to forecast lower dividends in this sample, so returns must account for more than 100 per cent of the variation in price-dividend ratios.

The evidence just reviewed implies that the observed excess volatility of the present value of future real cash flows is really just a phenomenon of time-varying expected excess stock market returns, as exhibited in the well documented long-horizon forecastability of excess returns using the price-dividend ratio.[2]

This has at least two profound implications for analysis of a strategic rebalancing regime.

First, the "market portfolio" will not have static weight between different asset classes.

Ignoring changes in the relative size of asset classes due to new issuance, since asset classes have different sensitivities to time-varying risk premia, their relative share of the "market portfolio" will change as risk premia change,

Second, observed financial-market excess volatility requires models with some mechanism for generating time-varying risk premia.

Comparing analytical frameworks

The early work on quantitative asset pricing asked whether analytical frameworks or models could account for one or maybe even a few of the asset-pricing facts in isolation. Over the last couple of years, the focus has been on whether the models can simultaneously account for a wide variety of phenomena in financial markets. This recent research has made important progress: researchers now have a much more consistent explanation of the size and time variation of risk premia across different asset classes. By carefully documenting dimensions along which existing models don't perform as well, significant progress is being made in understanding where the theory needs improvement.

In modern asset-pricing theory, a pricing kernel accounts for asset returns. The reverse is also true: asset returns contain information about the pricing kernel that gave rise to them. Mean excess returns on equity, bonds and other assets correspond to properties of the pricing kernel, specifically its dispersion and horizon dependence. These properties of models and their pricing kernels are important to distinguish them and judge their usefulness.

Hansen and Jagannathan (1991) compare the properties of theoretical models to those implied by observed returns. In their case, the properties are the mean and variance of the pricing kernel. The Hansen-Jagannathan bounds state that the ratio of the standard deviation of a pricing kernel to its mean must exceed the Sharpe ratio attained by any portfolio.

Backus et al. (2011) suggest entropy and horizon dependence as measures that capture a wider set of asset-pricing statistics. Entropy, as a measure of dispersion, is a generalization of the variance. The entropy of a pricing kernel is an upper bound on mean excess returns. Horizon dependence has no counterpart in the Hansen-Jagannathan methodology. It is a measure of how entropy varies with the investment horizon. As with entropy, its magnitude can be inferred from asset prices: negative (positive) horizon dependence is associated with an upward-sloping (downward-sloping) mean yield curve and positive (negative) mean yield spreads.

Two modelling frameworks have emerged as standard for accounting for volatility of financial markets and predictability of excess returns and satisfy the above-mentioned bounds: The habit-specification

2 Since Fama and French (1988) documented that the price-dividend ratio predicts realised returns, many papers have addressed this question in various contexts. Subsequent papers can be loosely categorised into two groups. On the one hand, reasearchers have introduced more sophisticated statistical tests and new estimation methods to complement the OLS regression results. Some argue that careful statistical analysis provides little support for the predictability of returns by the price-dividend ratio (Nelson and Kim, 1993; Stambaugh, 1999; Valkanov, 2003; Goyal and Welch, 2008; Ang and Bekaert, 2007). Others have introduced more powerful ways to test for return predictability, for example, by looking jointly at returns and dividends (Cochrane, 2008), by using a latent-variables approach (Pastor and Stambaugh, 2009; Binsbergen and Koijen, 2010; Rytchkov, 2008; Lacerda and Santa-Clara, 2010), by accounting for structural shifts in the mean of the price-dividend ratio (Lettau and Van Nieuwerburgh, 2008), and by imposing weak restrictions on the signs of coefficients and return forecasts (Campbell and Thompson, 2008). These papers suggest that it is imprtant to go beyond simple predictive regressions as they can be misleading.

of Campbell and Cochrane (1999) and the recursive risk-preference specification of Epstein and Zin (1989), as applied by, among others, Bansal and Yaron (2004).[3]

The key assumption in Campbell and Cochrane (1999) is that investors' utility functions depend on the past history of aggregate consumption, so they capture a "catching-up-with-the-Joneses" motive. Investors are more risk-averse in recessions, when their consumption is low relative to past aggregate consumption. They are less risk-averse in booms, when their consumption is high, and so gambling feels less threatening.

The two key assumptions in Bansal and Yaron (2004) are so-called recursive Epstein and Zin (1989) preferences and long-run risks, defined as small but persistent changes in expected consumption growth. They find that, in order to account for asset-pricing fluctuations, investors put relatively more weight on changes in uncertainty associated with long-run growth than on short-run consumption fluctuations.

For questions of dynamic risk sharing and rebalancing, these two frameworks are not equally suitable. Both in models with the classical "power utility" specification and in models with Campbell-Cochrane preferences, the source of fluctuations in investors' price of risk is fluctuations in consumption between two consecutive periods. An implication of this is that with these preference specifications, all dynamic risk-sharing problems – and hence rebalancing problems – become essentially static allocation problems. In models with market clearing, there is no rationale for rebalancing as expected returns change.

With recursive Epstein-Zin preferences, the meaning of risk aversion is more elaborate. Investors care not only about consumption fluctuations, but also about long-term growth and uncertainty associated with this. Hence decision problems become inherently dynamic and investors not only care about distributions over outcomes, but also, among others, when uncertainty is resolved. These preference specifications will potentially give frameworks with which we can analyse dynamic asset-allocation questions and rebalancing.

Preference heterogeneity and rebalancing

Well-documented time-variation in risk premiums have attracted a substantial literature on how investors could exploit the market-timing and intertemporal hedging opportunities implicit in expected return variation unrelated to market betas. Following the approach of Merton (1971), important contributions include Brennan et al. (1997), Campbell and Viceira (1999, 2002). The problem with the approach of these papers is that they don't account for why anybody would take the other side of the rebalancing trades. A useful and durable portfolio theory – in particular if applied by a very large, long-term investor – must be consistent with the fact that all assets must be held be somebody. If one investor sells an asset there must be another investor who buys that asset.

A first approach to the problem of time-variation in risk premiums, strong expected return variation, and rebalancing might be preference heterogeneity. "Preference heterogeneity" might be literally different attitudes towards risk across states and time, but it might also be considered a reduced-form representation of other forms of heterogeneity, like market frictions, differences in regulations, market incompleteness, or non-tradable risks.

There are two complimentary approaches to solving models with preference heterogeneity: The first is to solve the decentralized competitive equilibrium portfolios: Solve each individual investor's constrained optimization problem taking market prices a given; aggregate portfolio demand/supply across all investors; solve for asset prices that clear markets; and calculate individual portfolio choices at equilibrium prices.

3 The Epstein and Zin (1989) preference construct has as its intellectual antecedents the time-consistent recursive preference structures of Kreps and Porteus (1978) and Johnsen and Donaldson (1985), and the Chew (1983, 1989) and Denkel (1986) preference generalisation, which allows the independence assumption of classical expected utility to be relaxed. Backus et al. (2004) provides discussion of this preference class.

The other is to solve for risk sharing problems and consumption allocations. Since portfolios are not an end but a means to deliver consumption at future dates under various contingencies, optimal portfolio strategies can be derived from optimal consumption allocations.

Risk sharing over a random outcome

Consider a static risk-sharing problem with two types of investors and where the heterogeneity among the investors can be summarized by their risk aversion parameter α. Aggregate output is y and the Pareto weights are λ.

Total output $y(s)$ is a random variable with a given outcome for each state of the world s. The probabilities associated with s are $Pr[y = y(s)] = \pi(s)$

The Pareto problem with weight λ on investors of Type 1 and weight (1- λ) on investors of Type 2

$$\max_{\{c^1(s)\},\{c^2(s)\}} \lambda E_t\left[\sum_{\tau=0}^{\infty} \beta_1^{\tau} \frac{(c^1(s))^{\alpha_1}}{\alpha_1} \pi(s)\right] + (1 - \lambda) E_t\left[\sum_{\tau=0}^{\infty} \beta_2^{\tau} \frac{(c^2(s))^{\alpha_2}}{\alpha_2} \pi(s)\right]$$

subject to

$$c^1(s) + c^2(s) = y(s) \qquad s = 1,2, \ldots\ldots S$$

The approximate solutions to the static problem:

$$\log(c^1(s)) \approx \overbrace{\left(\frac{1-\alpha_2}{1-\hat{\alpha}}\right)}^{\text{(constant) share of outstanding equities}} \log(y(s)) + \overbrace{constant}^{\text{(constant) share of outstanding debt}}$$

$$\log(c^2(s)) \approx \underbrace{\left(\frac{1-\alpha_1}{1-\hat{\alpha}}\right)}_{\text{(constant) share of outstanding equities}} \log(y(s)) + \underbrace{constant}_{\text{(constant) share of outstanding debt}}$$

where $\hat{\alpha}$ is a weighted sum of the different investors risk preferences,

$$\hat{\alpha} = \lambda\, \alpha_2 + (1 - \lambda)\alpha_1$$

As we see, all types of investors get a given fraction of realized output. The share does not only depend on the investor's own risk preference, but also on how risk preferences are relative to the other investors' preferences.

Dynamic risk sharing with power utility

Now consider a dynamic risk-sharing problem where investor heterogeneity is summarized by heterogeneity in their risk aversion parameter α. The objective functions of the two types of investors are:

$$W_t = E_t\left[\sum_{\tau=0}^{\infty} \beta_1^{\tau} \frac{(c_{t+\tau}^1)\alpha_1}{\alpha_1}\right] and\ V_t = E_t\left[\sum_{\tau=0}^{\infty} \beta_1^{\tau} \frac{(c_{t+\tau}^2)\alpha_2}{\alpha_2}\right]$$

The Pareto problem with weight on investors of Type 1 and weight (1-) on investors of Type 2

$$\max_{\{c_t^1\},\{c_t^2\}} \lambda E_t\left[\sum_{\tau=0}^{\infty} \beta_1^{\tau} \frac{(c_{t+\tau}^1)^{\alpha_1}}{\alpha_1}\pi(s)\right] + (1-\lambda)E_t\left[\sum_{\tau=0}^{\infty} \beta_2^{\tau} \frac{(c_{t+\tau}^2)^{\alpha_2}}{\alpha_2}\pi(s)\right]$$

subject to

$$c_t^1 + c_t^2 = y_t \qquad \forall t$$

The approximate solutions to the dynamic problem is the same as the solutions to the static problem:

$$\log(c^1(s)) \approx \overbrace{\left(\frac{1-\alpha_2}{1-\hat{\alpha}}\right)}^{\text{(constant) share of outstanding equities}} \log\big(y(s)\big) + \overbrace{constant}^{\text{(constant) share of outstanding debt}}$$

$$\log(c^2(s)) \approx \underbrace{\left(\frac{1-\alpha_1}{1-\hat{\alpha}}\right)}_{\text{(constant) share of outstanding equities}} \log\big(y(s)\big) + \underbrace{constant}_{\text{(constant) share of outstanding debt}}$$

where

$$\hat{\alpha} = \lambda\, \alpha_2 + (1-\lambda)\alpha_1$$

All the dynamics are still exogenous. Irrespectively of heterogeneity, all types of investors hold a constant fraction of the total outstanding equities and bonds. Applied on a rebalancing question, these frameworks would imply no rationale for rebalancing a portfolio.

Dynamic risk sharing with long-run risk

As we have argued, models with standard power utility are not able to account for asset-market dynamics. We will therefore rather consider dynamic risk-sharing problems within preference heterogeneity as an explanation for portfolio heterogeneity and rebalancing within a frictionless model that satisfy the above-mentioned bounds. A generalization of expected utility that is capable of richer dynamics.

$$U_t = \left[(1-\beta)c_t^{\rho} + \beta\mu_t(U_{t+1})^{\rho}\right]^{1/\rho}$$
$$\mu_t(U_{t+1}) = (E_t U_{t+1}^{\alpha})^{1/\alpha}$$

Interpretations of the two curvature parameters, ρ and α are: $1/(1-\rho)$ is the elasticity of substitution across time of $c_{t+1}(s)$ for c_t . $1/(1-\alpha)$ is the elasticity of substitution across states at the same point in time of $c_{t+1}(s')$ for $c_{t+1}(s)$.

In a representative investor economy the market pricing kernel will be equal to the marginal rate of substitution

$$m_{t+1} = \beta \underbrace{\left(\frac{c_{t+1}}{c_t}\right)^{\rho-1}}_{\text{Short-run risk}} \underbrace{\left(\frac{U_{t+1}}{\mu_t(U_{t+1})}\right)^{\alpha-\rho}}_{\text{Long-run risk}}$$

Note that compared with the power utility models reviewed above, the key to account for asset-market facts is the added term: the long-run risk. Note also the role of recursive preferences: if α = ρ, the second term disappears and there’s no role for predictable consumption growth or volatility.

Recursive preferences have also important implications for preferences over the dynamics of risk. With $\alpha < \rho$ investors prefer early resolution of uncertainty. Equivalent to preferring early resolution of uncertainty is to prefer risk concentrated in the variance of the conditional mean – and equivalent to preferring late resolution of uncertainty is to prefer risk concentrated in the conditional variance.

Now consider a dynamic risk-sharing problem where investor heterogeneity is summarized by heterogeneity in their risk aversion parameters ρ and α. The objective functions of the two types of investors are:

$$W_t = \left[(1-\beta_1)c_{1,t}^{\rho_1} + \beta_1 \mu_{1,t}(W_{t+1})^{\rho_1}\right]^{\frac{1}{\rho_1}}$$

$$\mu_{1,t}(W_{t+1}) = \left[E_t W_{t+1}^{\alpha_1}\right]^{\frac{1}{\alpha_1}}$$

$$V_t = \left[(1-\beta_2)c_{2,t}^{\rho_2} + \beta_2 \mu_{2,t}(V_{t+1})^{\rho_2}\right]^{\frac{1}{\rho_2}}$$

$$\mu_{2,t}(V_{t+1}) = \left[E_t V_{t+1}^{\alpha_2}\right]^{\frac{1}{\alpha_2}}$$

The investors marginal rates of substitution are

$$m_{1,t+1} = \beta_1 \left(\frac{c_{1,t+1}}{c_{1,t}}\right)^{\rho_1-1} \left(\frac{W_{t+1}}{\mu_{1,t}(W_{t+1})}\right)^{\alpha_1-\rho_1}$$

$$m_{2,t+1} = \beta_2 \left(\frac{c_{2,t+1}}{c_{2,t}}\right)^{\rho_2-1} \left(\frac{W_{t+1}}{\mu_{2,t}(W_{t+1})}\right)^{\alpha_2-\rho_2}$$

In competitive markets these are equalized

$$m_{1,t+1} = m_{2,t+1}$$

i.e.

$$\beta_1 \left(\frac{c_{1,t+1}}{c_{1,t}}\right)^{\rho_1-1} \left(\frac{W_{t+1}}{\mu_{1,t}(W_{t+1})}\right)^{\alpha_1-\rho_1} = \beta_2 \left(\frac{c_{2,t+1}}{c_{2,t}}\right)^{\rho_2-1} \left(\frac{W_{t+1}}{\mu_{2,t}(W_{t+1})}\right)^{\alpha_2-\rho}$$

The portfolio dynamics – the dynamic portfolio strategies of the two investor types – that support this condition might, however, be very different.

Now there might be important insights into the rationales for rebalancing. If asset prices move mainly due to changes in the conditional variance – i.e. about uncertainty about future – and as they do in most structural studies of asset markets – then the investors who are less sensitive to changes in the conditional variance would find it optimal to overweight equities relative to the market capitalization weights when conditional variance and expected returns are high and to underweight equities relative to the market capitalization weights when conditional variance and expected returns are low.

"Preference heterogeneity"

Markets are obviously not complete and frictionless. Preference-parameter heterogeneity might therefore not be differences in "preferences" in a narrow definition of the word, but rather represents differences in how types of investors are affected by market frictions and sources of incompleteness.

Cochrane (2010) distinguishes between three categories of frictions: segmented markets, institutional finance or intermediated markets, and liquidity premiums.

In segmented markets, risks are shared only between investors in a specific group, not across groups. This feature limits risk-bearing activity and therefore leads to the emergence of premia that are not related to aggregate risks. Thus one basic consequence of segmented markets is that risks are not shared across investors as they are in the standard model. Differences in risk-bearing activities might also be due to differences in non-tradeable assets.

Intermediated markets – or institutional finance – apply to a different, vertical, separation of investor from payoff. Investors use delegated managers to handle their assets. Then, frictions or principal/agent problems in the delegated management relationship spill over into market prices for the assets.

A long tradition in asset pricing recognizes that some assets have higher or lower discount rates in compensation for greater or lesser liquidity. Defining liquidity – there may be several different kinds – modelling it and understanding it deeply are still open questions.

Within all these areas, there is a large and rapidly growing academic literature that can be used in order to identify and critically discuss the unique characteristics of a hypothetical large investor within a government organization. Some references:

Non-tradeable assets: Danthine et al. (1992), Fama and French (1996), Constantinides and Duffie (1996), Vissing-Jørgensen (2002)

Segmented markets: Mitchell et al. (2007), Froot and O'Connell (2008)

"Institutional finance", "intermediated markets": Gârleanu and Pedersen (2011),

Brunnermeier (2009), Brunnermeier and Pedersen (2009), He and Krishnamurthy (2008)

Liquidity premium, trading volume: Krishnamurthy (2002), Boudoukh and Whitelaw (1991), Longstaff (2004), Acharya and Pedersen (2005)

Market segmentation and intermediated markets are important to account for cross-sectional risk factors and premiums. In order to account for and identify comparative advantages and disadvantages with respect to time-varying expected returns, institutional features and contributions that shed light on time-varying liquidity might be a better guide.

According to Cochrane (2010, p. 60), liquidity, trading and market microstructure have often been ignored in mainstream finance. Prices are set and perceived as if trading were unimportant. In this view, trading and microstructure, though interesting, don't add that much to the big picture. Cochrane continues to state that variation in discount rates due to "liquidity", whether individual or systemic, will be seen to be much more important than we now think.

Institutional features might be such as rules and regulations governing insurance companies and other financial institutions, which forces these institutions to reduce risk exposure when risk (and expected returns) are high, further contributing to falling prices and higher expected returns. Not being governed by such rules and regulations might, within the analytical frameworks discussed in the previous section, be interpreted as having a smaller distance between ρ and α, in other words have weaker preference for resolution of uncertainty, and a comparative advantage in taking counter-cyclical long-term risk.

Not being subject to margin requirements, or other issues that have been proven important to understand liquidity and market microstructure, might also be translated into having a smaller distance between r and a, in other words a weaker preference for resolution of uncertainty. Large, long-term investors, who that are not governed by such rules, regulations and margin requirements, might find it profitable to take the other side of these trades.

However, several hedge funds and active managers, including university endowments, describe part of their strategy as providing liquidity, presumably making just the trade described above, in expectation of higher returns. It is therefore not obvious that the magnitudes are such that it is possible to base a strategy on the rules and regulations of insurance and similar financial firms. Also, many hedge funds, university endowments, and active managers sold in panic along with everyone else in the Fall of 2008. The need for careful scrutiny of internal incentive and decision systems for a counter-cyclical rebalancing rule should not be understated.

Future research

What can account for the behavior of risk premia in stock and bond markets, both over time and cross-sectionally across classes of assets? For investment professionals, finding practical answers to these questions is the fundamental purpose of financial economics, as well as its principal reward. For academic researchers, the progression of empirical evidence on these questions has presented a continuing challenge to asset pricing theory. Identifying the important sources of risks that individuals and firms face and how these in turn affect allocations and prices constutes a road map for future inquiry. These risks might include, for example, aggregate productivity shocks, uninsurable, non-tradable risk, and/or financial frictions faced by both individuals and firms.

As argued in this note, generalized recursive preferences (e.g. Kreps and Porteus, 1978; Epstein and Zin, 1989; Weil, 1989), provide a framework to not only identify the market price of short-term fluctuations, but also deviations from normality and time dependence. More than one identified risk factor allows not only studying preference heterogeneity and a rationale for rebalancing, but also what risks a rebalancing rule exposes the asset owner for. These questions are yet to be satisfactory answered.

Variation in discount rates might due to systemic liquidity, trading and market microstructure. Whether systematic variations in individual-investor liquidity might a source of "preference heterogeneity" is still an open question.

These and other analytical framework might not only be used to address the issues of rebalancing. Several researchers and papers use features of this framework or extension of it to address various issues and markets. These include research on global asset allocation, foreign exchange markets, the term structure of interest rates, credit spreads, derivative markets, co-integration and portfolio choice, the value premiums, robust control and learning, and inflation risk premia.

Concluding remarks

Long-horizon forecastability of excess returns is a well-documented phenomenon. Almost certainly, it has the same sources as observed excess volatility of the market value of real cash flows. It is, however, not immediately apparent whether a large investor with a long time horizon have a comparative advantage in exploiting the market-timing and intertemporal-hedging opportunities implicit in time-varying expected returns. In particular, the questions are what are the risks associated with rebalancing a portfolio, and why might other investors find it optimal to take the other side of a rebalancing trade?

Almost all strategic investment decisions are, either explicitly or implicitly, based on analytical frameworks or models. A challenge is to be aware of the modelling frameworks and know their strengths and weaknesses. Theoretical foundations for how to optimally time the market and adjust to time-varying expected returns, ought to be based on models that can account for time-varying expected excess stock market returns. For a large investor, investment strategies should also be consistent with the facts that other investors are doing as best they can and that all assets must be owned by somebody.

Models that can identify prices associated not only with short-run fluctuations, but also long-run growth and time-varying uncertainty associated with long-run growth are not only able to generate time-varying expected returns as observed in the data, but also to give a rationale for dynamic portfolio rebalancing. Models with so-called recursive preferences show that time-varying risk premia, and hence time-varying expected returns, are driven primarily by uncertainty associated with long-run asset returns.

Rebalancing to exploit time-varying risk premia is far from an arbitrage opportunity. In fact, within the frameworks that have been sketched in this note, it implies increasing exposure to long-run risk in periods when long-run risk is high, and reducing exposure to long-run risk in periods when long-run risk is low. Taking the other side of this trade would mean trying to reduce the fluctuations in a portfolio's exposure to uncertainty associated with long-run economic trend growth.

Heterogeneity in preferences over short-run risk, long-run risk, and uncertainty associated with long-run risk is closely associated with heterogeneity in preferences over timing of resolution of uncertainty. The concept of resolution of uncertainty is related to, but should not be confused with length of the time horizon. More precisely, it is willingness and ability to hold on to assets in periods when the uncertainty about the long-run growth rate of future payoffs is high. Hence, heterogeneity in this form of risk-bearing capacity might account for why it is optimal for some investors to buy and other investors to sell in periods when uncertainty about the future growth rate, and hence also risk premia, are high.

Investors who have the least preference for early resolution of uncertainty might find it optimal with an investment strategy where their risk exposure is counter-cyclical. When long-run risk increases, the uncertainty associated with each equity claim increases and equity prices decrease relative to bond prices. In such a situation when risk associated with the un-rebalanced portfolio has increased, long-run investors might find it optimal to rebalance towards equities and further increase the risk exposure and expected returns of the portfolio.

There might be a combination of model parameters that would imply constant portfolio weights in bonds and equities, but that would generally not be the case. Constant portfolio weights might be considered an approximation of an optimal portfolio rule for an investor who is willing to take larger risks and accept a wider range of outcomes than most other investors.

Preference heterogeneity in this analytical framework is not necessarily only "preference" heterogeneity in the strict definition of the word. A rationale for such dynamic risk-taking behaviour might be found in the academic literature that studies market frictions, non-tradable risks, segmented markets and institutional finance.

In particular, not being subject to the same risk-regulations and margin requirements governing other financial institutions might be isomorphic to having a weaker preference for early resolution of uncertainty than other investors, hence a comparative advantage for taking counter-cyclical long-run risk and a rationale for a rebalancing rule.

References

Acharya, V. V. and Pedersen, L. H. (2005). Asset pricing with liquidity risk. Journal of Financial Economics, 77(2):375–410.

Ang, A. and Bekaert, G. (2007). Stock return predictability: Is it there? Review of Financial Studies, 20(3):651–707.

Backus, D., Chernov, M., and Zin, S. E. (2011). Sources of entropy in representative agent models. NBER Working Papers 17219, National Bureau of Economic Research.

Backus, D., Routledge, B. R., and Zin, S. E. (2004). Exotic preferences for macroeconomists. In Gertler, M. and Rogoff, K., editors, NBER Macroeconomics Annual 2004. MIT Press, Cambridge, MA.

Bansal, R. and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. Journal of Finance, 59(4):1481–1509.

Bollerslev, T., Engle, R. F., and Wooldridge, J. M. (1988). A capital asset pricing model with time-varying covariances. Journal of Political Economy, 96(1):116–31.

Brennan, M. J., Schwartz, E. S., and Lagnado, R. (1997). Strategic asset allocation. Journal of Economic Dynamics and Control, 21(8-9):1377–1403.

Brunnermeier, M. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. Review of Financial Studies, 22:2201–2238.

Brunnermeier, M. K. (2009). Deciphering the liquidity and credit crunch 2007-2008. Journal of Economic Perspectives, 23(1):77–100.

Campbell, J., Lo, A., and MacKinlay, A. (1997). The econometrics of financial markets. Princeton Univ. Press, Princeton, NJ.

Campbell, J. Y. (1991). A variance decomposition for stock returns. Economic Journal, 101:157–179.

Campbell, J. Y. and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. Journal of Political Economy, 107:205–251.

Campbell, J. Y. and Shiller, R. J. (1989). The dividend-price ratio and expectations of future dividends and discount factors. Review of Financial Studies, 1(3):195–228.

Campbell, J. Y. and Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? Review of Financial Studies, 21(4):1509–1531.

Campbell, J. Y. and Viceira, L. (1999). Consumption and portfolio decisions when expected returns are time varying. Quarterly Journal of Economics, pages 433–495.

Campbell, J. Y. and Viceira, L. (2002). Strategic Asset Allocation: Portfolio Choice for Long- Term Investors. Clarendon Lectures in Economics. Oxford University Press.

Cochrane, J. H. (1992). Explaining the variance of price-dividend ratios. Review of Financial Studies, 5(2):243–80.

Cochrane, J. H. (2005). Asset Pricing, Revised Edition. Princeton University Press, Princeton, NJ.

Cochrane, J. H. (2008). The dog that did not bark: A defense of return predictability. Review of Financial Studies, 21(4):1533–1575.

Cochrane, J. H. (2010). Discount rates. manuscript, University of Chicago Booth School of Business.

Constantinides, G. M. and Duffie, D. (1996). Asset pricing with heterogeneous consumers. Journal of Political Economy, 104:219–40.

Danthine, J.-P., Donaldson, J. B., and Mehra, R. (1992). The equity premium and the allocation of income risk. Journal of Economic Dynamics and Control, 16:509–532.

Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. Econometrica, 57(4):937–69.

Fama, E. F. and French, K. R. (1988). Dividend yields and expected stock returns. Journal of Financial Economics, 22:3–26.

Fama, E. F. and French, K. R. (1989). Business conditions and expected returns on stocks and bonds. Journal of Financial Economics, 25:23–49.

Fama, E. F. and French, K. R. (1996). Multifactor explanations of asset pricing anomalies. Journal of Finance, 51(1):55–84.

Gârleanu, N. and Pedersen, L. H. (2011). Margin-based asset pricing and deviations from the law of one price. Review of Financial Studies, 24(6):1980–2022.

Goyal, A. and Welch, I. (2008). A comprehensive look at the empirical performance of equity premium prediction. Review of Financial Studies, 21(4):1455–1508.

Hansen, L. P. and Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. Journal of Political Economy, 99(2):225–62.

He, Z. and Krishnamurthy, A. (2008). Intermediary asset pricing. NBER Working Papers 14517, National Bureau of Economic Research.

Ilmanen, A. (2011). Expected Returns: An Investor's Guide to Harvesting Market Rewards. John Wiley & Sons.

Johnsen, T. and Donaldson, J. B. (1985). The structure of intertemporal preferences under uncertainty and time consistent plans. Econometrica, 53:1451–1458.

Kreps, D. M. and Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. Econometrica, 46:185200.

Krishnamurthy, A. (2002). The bond/old-bond spread. Journal of Financial Economics, 66(2-3):463–506.

LeRoy, S. F. and Porter, R. D. (1981). The present-value relation: Tests based on implied variance bounds. Econometrica, 49(3):555–74.

Lettau, M. and Nieuwerburgh, S. V. (2008). Reconciling the return predictability evidence. Review of Financial Studies, 21(4):1607–1652.

Mehra, R. and Prescott, E. C. (1985). The equity premium: A puzzle. Journal of Monetary Economics, 15(2):145–161.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous time case. Review of Economics and Statistics, 51:247–257.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous time model. Journal of Economic Theory, 3:373–413.

Mitchell, M., Pedersen, L. H., and Pulvino, T. (2007). Slow moving capital. American Economic Review, 97(2):215–220.

Mossin, J. (1968). Optimal multiperiod portfolio policies. Journal of Business, 41:215–229.

Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. Review of Economics and Statistics, 51:239–246.

Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? American Economic Review, 71(3):421–36.

Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. Journal of Political Economy, 110(4):825–853.

Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. Journal of Monetary Economics, 24(3):401–421.

Norges Bank Investment Management (NBIM)
Bankplassen 2
Postboks 1179 Sentrum
N-0107 Oslo

Tel.: +47 24 07 30 00
www.nbim.no