

*The Expected Marginal Rate of Substitution  
in the United States and Canada*

**Andrew K. Rose\***

Revised: December 1, 2004

Andrew K. Rose  
Haas School of Business  
University of California  
Berkeley, CA 94720-1900  
Tel: (510) 642-6609  
E-mail: arose@haas.berkeley.edu

**Abstract**

This paper applies a new methodology to estimate the expected intertemporal marginal rate of substitution or “EMRS”, using only data on asset prices and returns. A novel feature of the technique is that it relies upon exploiting idiosyncratic risk, since theory dictates that idiosyncratic shocks earn the EMRS. I apply the methodology to two different data sets: monthly data from 1994 through 2003, and daily data for 2003. Both data sets include large stocks traded on the New York and Toronto Stock Exchanges. For both monthly and daily frequencies, I find plausible estimates of EMRS with considerable precision and time-series volatility. I then use these estimates to test for asset integration, both within and between stock markets. I find that both markets seem to be internally integrated in the sense that different assets traded on a given market share the same EMRS. More interestingly, the technique is powerful enough to reject integration between the two stock markets, and between stock and bond markets.

**JEL Classification Numbers:** G14

**Keywords:** stochastic, discount, factor, intertemporal integration, asset, market, price, stock.

\* B.T. Rocca Jr. Professor of International Business, Haas School of Business at the University of California, Berkeley, NBER Research Associate, and CEPR Research Fellow. I thank conference participants at the Bank of Canada, Eric Santor, Larry Schembri, and Greg Smith for comments. This draws heavily on my paper “Estimating the Expected Marginal Rate of Substitution” with Robert P. Flood. The data set, sample output, and a current version of this paper are available at <http://faculty.haas.berkeley.edu/arose>.

In this paper, I develop and apply a simple methodology to estimate the expected intertemporal marginal rate of substitution (hereafter “EMRS”). The EMRS is an economic variable of considerable intrinsic interest. More importantly, when different series for the EMRS are estimated for different markets (such as the New York and Toronto stock exchanges), comparing these estimates provides a natural yet powerful test for integration between markets. The method is novel in that it exploits information in asset-idiosyncratic shocks.

The primary objective of this paper is expositional, and is intended to present a new methodology. Nevertheless, I illustrate the technique by applying it to monthly and daily data covering firms from the largest American and Canadian stock exchanges. It turns out that the method delivers plausible EMRS estimates with considerable precision. Estimates from the Canadian and American markets can be distinguished from each other and from the Treasury bill equivalent.

Section 1 presents the methodology; technical details required to implement the technique are discussed in the following section. The empirical results are presented in section 3, while the paper ends with a brief conclusion.

## **1 The Technique**

It is easy to motivate this exercise. Asset market integration is a topic of continuing interest in international finance, see e.g., Adam et. al. (2002). It is of special interest in Europe where continuing monetary and institutional integration have lead to lower barriers to asset trade inside the EU. But economists from other countries should also be interested; after all, the “open” adjective in “small open economy” typically refers to asset market integration, and is a critical assumption for macroeconomists modeling countries like Canada. Macroeconomists in

developed countries almost always assume that financial markets are integrated both internally, and with real markets (e.g., goods and services markets through consumption and investment decisions). Open economy macroeconomists often assume that financial markets in different countries are integrated; that is, they assume away a financial “border effect.” But the EMRS is also of direct interest in itself. Any forward-looking model that uses intertemporal optimization has agents comparing leisure now vs. leisure later, buying assets now vs. later, and so forth; such comparisons are done through the MRS.

I begin with a completely conventional intertemporal asset pricing condition:

$$p_t^j = E_t(m_{t+1}x_{t+1}^j) \tag{1}$$

where:  $p_t^j$  is the price at time t of asset j,  $E_t(\cdot)$  is the expectations operator conditional on information available at t,  $m_{t+1}$  is the time-varying intertemporal marginal rate of substitution (MRS), used to discount income accruing in period t+1 (also known as the stochastic discount factor, marginal utility growth, or pricing kernel), and  $x_{t+1}^j$  is income received at t+1 by owners of asset j at time t (the future value of the asset plus any dividends or other income).

I adopt a standard definition of asset integration – *two portfolios are said to be integrated when they are priced by the same stochastic discount factor*. Here “priced” means that equation (1) holds for the assets in question. Equation (1) involves the moments of  $m_{t+1}$  and  $x_{t+1}^j$ , not the realized values of those variables. In particular, for integration I do not require realized values of  $m_{t+1}$  to be equated across assets or agents pricing assets.

Although many moments of  $m_{t+1}$  are involved in asset market integration, the object of interest to me in this study is  $E_t m_{t+1}$  the time  $t$  expectation of the intertemporal marginal rate of substitution (EMRS). I concentrate on the first moment for three reasons. First, the expectation of the MRS,  $E_t m_{t+1}$  is intrinsically important; it lies at the heart of much intertemporal macroeconomic and financial economics and is virtually the DNA of modern aggregate economics. Second, it turns out to be simple to measure with high statistical accuracy. Third, cross-market differences in estimated values of  $E_t m_{t+1}$  are statistically distinguishable, providing powerful evidence concerning market integration.<sup>1</sup> I am testing only for first-moment equality when many additional moments are used in asset pricing; thus, this is a test of a necessary condition for integration. If I reject equality of the first moment, I can reject integration, but failing to reject first-moment equality is consistent with (but does not imply) complete integration.

Consider a standard decomposition of equation (1):

$$p_t^j = E_t(m_{t+1} x_{t+1}^j) = COV_t(m_{t+1}, x_{t+1}^j) + E_t(m_{t+1})E_t(x_{t+1}^j). \quad (2)$$

where  $COV_t(\cdot)$  denotes the conditional covariance operator. It is useful to rewrite this as

$$\begin{aligned} x_{t+1}^j &= -[1/E_t(m_{t+1})]COV_t(m_{t+1}, x_{t+1}^j) + [1/E_t(m_{t+1})]p_t^j + \varepsilon_{t+1}^j, \text{ or} \\ x_{t+1}^j &= \delta_t(p_t^j - COV_t(m_{t+1}, x_{t+1}^j)) + \varepsilon_{t+1}^j \end{aligned} \quad (3)$$

where  $\varepsilon_{t+1}^j \equiv x_{t+1}^j - E_t(x_{t+1}^j)$ , a prediction error orthogonal to information at time  $t$ , and  $\delta_t \equiv 1/E_t(m_{t+1})$ . The latter time-series vector is the set of parameters of interest. In an integrated market without trading frictions, it is identical for all assets, since the first moment of the marginal rate of substitution should be equal inside integrated financial markets. My work below is essentially concerned with exploiting and testing this restriction.

### 1.1 Old Stuff

It is typical in domestic finance to make equation (3) stationary by dividing the equation by  $p_{j,t}$ , resulting in:

$$x_{t+1}^j / p_t^j = \delta_t (1 - COV_t(m_{t+1}, x_{t+1}^j / p_t^j)) + \varepsilon_{t+1}^j, \quad (4)$$

where  $\varepsilon_{t+1}^j$  is redefined appropriately. This normalization converts equation (3) into a traditional asset-pricing equation. That is, it breaks one-period asset returns,  $x_{t+1}^j / p_t^j$ , into the risk-free market return,  $\delta_t \equiv 1/E_t(m_{t+1})$ , and asset-specific period risk premia, the covariance term.

Equation (4) is given economic content by adding two assumptions:

- 1) *Rational Expectations*:  $\varepsilon_{t+1}^j$  is assumed to be uncorrelated with information available at time  $t$ , and
- 2) *Covariance Model*:  $COV_t(m_{t+1}, x_{t+1}^j / p_t^j) = \beta_0^j + \sum_i \beta_i^j f_{i,t}$ , for the relevant sample,

where:  $\beta_0^j$  is an asset-specific intercept,  $\beta_i^j$  is a set of I asset-specific factor coefficients and  $f_{i,t}$  a vector of time-varying factors. Both assumptions are common in the literature; Campbell, Lo and MacKinlay (1997) and Cochrane (2001) provide excellent discussions. With these two assumptions, equation (10) becomes a panel estimating equation. *Time-series* variation is used to estimate the asset-specific factor loadings  $\{\beta\}$ , coefficients that are constant across time. Estimating these factor loadings is a key objective of this research program.

In practice, many empirical asset pricing modelers set  $\delta(t) = 1 + i(t)$ , where  $i(t)$  is an appropriate short-term riskless interest rate. That is, the EMRS is simply equated to, e.g., the Treasury-bill rate; it is not estimated at all. While this simplifies empirical work considerably, this *assumes* integration between stock and money markets, one of the very assumptions I wish to *test* rather than make.

The first approach to testing asset market integration between a pair of markets makes one of the factors, say the first one, equal to a market identifier. This allows cross-sectional estimation of a market-specific effect each period. For a set of risk factors that are held to price assets in both markets, the market-specific effects should all be zero under the null of integration. Rejecting the joint null hypothesis – but maintaining rational expectations – rejects *either* market integration, *or* the risk pricing model (or both).

Two points are essential to the first approach. First, it is based on the Finance standard where the risk premium is postulated to be a function (usually linear) of a set of aggregate risks. Second, the market integration test is tested as part of a *joint* hypothesis that includes the aggregate risks that model risk premia.

## 1.2 New Stuff

In this paper, I rely on a different normalization. Suppose I observe  $\tilde{p}_t^j$ , which is defined to be the value of  $p_t^j$  conditional on idiosyncratic information (available at time t) being set to zero. Consider the regression:

$$\ln(p_t^j / p_{t-1}^j) = \alpha_0^j + \sum_{i=1}^N \alpha_i^j f_t^i + v_t^j, \quad (5)$$

where the  $f_t^i$  are a set of aggregate factors, e.g., the log of unity plus average price growth, and  $v_t^j$ , the residual, is the idiosyncratic part of asset j price return. From the definition of  $\tilde{p}_t^j$ ,

$$\tilde{p}_t^j = p_{t-1}^j \exp(\alpha_0^j + \sum_{i=1}^N \alpha_i^j f_t^i), \quad (6)$$

which is  $p_t^j$  with its idiosyncratic part set to zero.

Normalizing by  $\tilde{p}_t^j$  delivers:

$$x_{t+1}^j / \tilde{p}_t^j = \delta_t [(p_t^j / \tilde{p}_t^j) - COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)] + \varepsilon_{t+1}^j. \quad (7)$$

The first term inside the brackets,  $(p_t^j / \tilde{p}_t^j)$ , equals  $\exp(v_t^j)$ , which is a function of only idiosyncratic information. The second term,  $COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$  is the covariance of the unknown market discount rate,  $m_{t+1}$ , with the synthetic return,  $x_{t+1}^j / \tilde{p}_t^j$ . Similar to the risk premium assumption in Finance, I assume  $COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$  moves only because of

aggregate phenomena. Since idiosyncratic risk,  $(p_t^j / \tilde{p}_t^j)$  is orthogonal to systematic risk,  $COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$ , (7) can be decomposed as

$$x_{t+1}^j / \tilde{p}_t^j = \delta_t (p_t^j / \tilde{p}_t^j) - \delta_t COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j) + \varepsilon_{t+1}^j = \delta_t \exp(v_t^j) + u_{t+1}^j \quad (8)$$

where  $u_{t+1}^j \equiv \varepsilon_{t+1}^j - \delta_t COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$ . By design, both parts of the composite error term are orthogonal to the only regressor,  $\exp(v_t^j) = p_t^j / \tilde{p}_t^j$ . The first part,  $\varepsilon_{t+1}^j$ , is a forecasting error which is unrelated to all information at time t by rational expectations. The second part,  $COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$  is unaffected by any idiosyncratic phenomena. Since both terms are orthogonal to the regressor that represents idiosyncratic risk,  $(p_t^j / \tilde{p}_t^j)$ , the coefficients of interest,  $\{\delta\}$ , can be consistently estimated via (8). A correct empirical specification of  $COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$  would lead to more *efficient* estimation of  $\{\delta\}$ . However, an empirical specification of  $COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$  is unnecessary for *consistent* estimation.

The basic idea of this study and the essential way it differs from previous work is that I use asset-idiosyncratic shocks to identify and measure the expected marginal rate of substitution (or rather, its inverse),  $\{\delta\}$ . This stands typical Finance methodology – the approach discussed above – on its head. In traditional asset-pricing Finance, idiosyncratic risk is irrelevant and orthogonal to the center-piece measures of aggregate risk. By their nature, idiosyncratic risks are easy to insure against and hence carry no risk premium. While idiosyncratic shocks carry no information about individual asset risk premia, they are loaded with information relevant to market aggregates. The new test for asset market integration is simple; I check if the implied prices of carrying idiosyncratic risks – measures of the expected marginal rate of substitution –

are equal across portfolios. If equality of the estimated EMRS cannot be rejected, then the test cannot reject cross-portfolio integration. If, however, I can reject equality then I also reject integration.

This normalization has the advantage – in common with the strategy of Flood-Rose (2003) – that it allows estimation of  $\{\delta\}$ . However, it does not rely directly on a correctly specified asset pricing model. That is, I do not explicitly rely on a model of  $COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$ , (such as, e.g., the CAPM used by Bekaert and Harvey, 1995).

The essential difference between this method and traditional methods is that I substitute a representation of price movements plus an orthogonality condition, for a model of  $COV_t(m_{t+1}, x_{t+1}^j)$ , which incorporates a similar orthogonality condition. The advantage of the method is that it deals only with observable variables. The stochastic discount rate  $m_{t+1}$  is unobservable as are its moments. When I project asset price movements onto a set of aggregate factors, I am taking the same stand on relevant aggregates that others take when they model  $COV_t(m_{t+1}, x_{t+1}^j)$ . The advantage of this method is that it leaves a highly volatile regressor – idiosyncratic shocks – attached to  $\delta_t = 1/m_{t+1}$ .

The new methodology has a number of other strengths. First, it is based on a general intertemporal theoretical framework, unlike other measures of asset integration such as stock market correlations (see the discussion in e.g., Adam et. al. 2002). Second, I do not need to model the EMRS directly; I allow it to vary over time in a completely general fashion. Third, the technique requires only accessible and reliable data on asset prices and returns. Fourth, the methodology can be used at a full range of frequencies. Fifth, the technique can be used to compare estimates of EMRS across many different classes of intertemporal decisions, including saving decisions that involve domestic and foreign stocks, bonds, and commodities. Sixth, the

technique is easy to implement and can be applied with standard econometric packages; no specialized software is required. Finally, the technique is focused on estimating an intrinsically interesting object, the (inverse of the) expected marginal rate of substitution.

## 2 Taking it to the Data

In practice  $\tilde{p}_t^j$  is an unobservable variable. Thus, I use an observable statistical counterpart derived from an empirical model, denoted  $\hat{p}_t^j$  (I note that this induces measurement error, an issue I handle below). I do this in a straightforward way, using simple time-series regressions that link individual asset-price returns to the average. In particular, I estimate the following J time series regressions via ordinary least squares (OLS):

$$\ln(p_t^j / p_{t-1}^j) = a_j + b_j * \ln(\bar{p}_t / \bar{p}_{t-1}) + v_t^j \quad (9)$$

where  $a_j$  and  $b_j$  are fixed regression coefficients,  $\bar{p}_t$  is the market-wide average price and  $v_t^j$  is the time-t asset idiosyncratic shock. This equation has a natural and intuitive interpretation; it models the first-difference of the natural logarithm of a particular asset price as a linear function of the price growth of the market. Estimates of equation (9) allow me to produce the fitted value of  $\hat{p}_t^j$ , which I define as:

$$\hat{p}_t^j \equiv p_{t-1}^j * \exp(\hat{a}_j + \hat{b}_j \ln(\bar{p}_t / \bar{p}_{t-1})) \quad (10)$$

I am not sentimentally attached to this specific model of  $\hat{p}_t^j$ . I could instead employ the Kalman filter to avoid using future data and allow for moving coefficient estimates.

Alternatively one could add additional regressors to equation (9) to control for more aggregate factors.<sup>2</sup> I have assumed that the log first-difference of prices is linear in the market; one could change the particular functional form assumption. I have used a time-series approach to estimating  $\hat{p}_t^j$  but a cross-sectional approach is also possible. None of these assumptions are critical; they simply seem to work in practice.<sup>3</sup> But while this particular setup has delivered sensible results, I stress that one only needs *some* model for  $\hat{p}_t^j$ , not this precise one.

## 2.1 Estimation Issues

I am fundamentally interested in estimating  $\{\delta\}$  from the following model:

$$x_{t+1}^j / \hat{p}_t^j = \delta_t (p_t^j / \hat{p}_t^j) + u_{t+1}^j \quad (11)$$

for assets  $j=1, \dots, J$ , periods  $t=1, \dots, T$ . I allow  $\{\delta_t\}$  to vary arbitrarily period by period.

Using  $\hat{p}_t^j$  in place of the unobservable  $\tilde{p}_t^j$  might induce important measurement error. Hence it is natural to consider estimation of equation (11) with instrumental variables (IV) for consistent estimation of  $\{\delta_t\}$ . IV is also known to handle the “generated regressor” issue which has long been known to be associated with potentially overstated precision of standard errors; see Shanken (1992) and Cochrane (2001 and website correction). The latter show that this is not typically very important in practice, especially for monthly data. While IV estimation seems natural, estimation via the Generalized Method of Moments (GMM) also allows me to handle both potential econometric issues, while also not requiring independent and identically

distributed disturbances. Accordingly, I use and compare two different estimation techniques: IV, and GMM (I sometimes also use OLS for simplicity). As instrumental variables for  $\{p_t^j / \hat{p}_t^j\}$ , I use the set of time-varying market-wide average prices  $\{\bar{p}_t\}$ .

## 2.2 The Data Sets

I employ two different data sets. The first is a decade of monthly data, spanning 1994M1 through 2003M12, while the second is a year of daily data for 2003. I use different frequencies both for intrinsic interest and to check the sensitivity of my technique. Though these frequencies are standard in Finance, there is nothing special about them, and there is no obvious reason why the methodology could not be used at either higher or lower frequencies. I focus on stock markets, but again see no reason why bond and other markets could not be considered.

My American stock data were extracted from the CRSP data base and consist of month-end prices and returns (including dividends, if any) for all firms in the S&P 500 (as of the end of 2003). I have adjusted for stock splits, and checked and corrected the data for errors. I only retain the 435 companies that have data for the full sample span.<sup>4</sup>

Since I am interested in estimating and comparing implied EMRS across markets, I also include data from the Canadian market. In particular, I add comparable data for the firms in the S&P/TSX Composite Index of the Toronto Stock Exchange (TSE). This data set is extracted from Datastream, and I convert Canadian dollars into American using comparably timed exchange rates. Canadian equity may be more concentrated in firms that specialize in commodities, and may also be more likely to be closely held by families. It is thus interesting to compare the Canadian EMRS with its American counterpart.

For the monthly data set, I have 120 monthly observations on 389 firms from the S&P 500 traded on the NYSE and 152 firms from the TSE. For the daily data set, I have data for 247 business days when both the Canadian and American stock exchanges were open, on 440 NYSE firms and 223 TSE firms.

It has been traditional since at least Fama and MacBeth (1973) to use yields on short-horizon treasury bills to proxy the risk-free rate, and it is natural for me to compare the *estimates* of the expected risk-free rate with T-bill returns. I use data on T-bill returns downloaded from the Federal Reserve's website.

Finally, for purposes of estimation, I group stocks into portfolios, typically twenty. I do this randomly (alphabetical order of ticker), though I see no reason why one could not group firms on the basis of e.g., beta, size, or whatnot. I use portfolios partly to remain within the Finance tradition followed since Fama and MacBeth (1973). But using portfolios also makes my task more difficult, since portfolios have lower idiosyncratic risk and reduce cross-sectional dimensionality. It turns out not to destroy my ability to estimate the parameters of interest to us, though there is no obvious reason why individual securities could not be used in place of portfolios.

### **3 Results**

The focus of this paper is estimating the expected marginal rate of substitution. I begin with an illustration that relies on monthly data from 400 firms in the S&P 500, grouped into portfolios of twenty firms.<sup>5</sup> I have 118 observations between February 1994 and November 2003, since I lose an observation at either end of the sample, due to leads/lags.

The three graphs on the left of Figure 1 portray estimates of the EMRS from (11), denoted  $\{\hat{\delta}_t\}$ . These were estimated using three different techniques: IV, GMM, and OLS. In each case, I use data from only the first ten portfolios. The mean of  $\{\hat{\delta}_t\}$  is plotted, along with a  $\pm 2$  standard error confidence interval band. The OLS and GMM point estimates are identical (by construction), and are extremely highly correlated with the IV estimates. The primary differences between the different estimates lie in the standard errors; all three estimators deliver small standard errors, with the GMM standard errors being slightly smaller than those of either IV or OLS (but with more period to period volatility).<sup>6</sup> Indeed, I almost never find significant differences between the three estimators below, and thus tend to rely on IV below.<sup>7</sup>

Even though I estimate the expected MRS from only ten portfolios, the results seem quite sensible. Most of the estimates of the (inverse of the) expected monthly MRS are just over unity. The sample average of  $\{\hat{\delta}_t\}$  over the 118 periods is around 1.0085, implying an annual MRS of slightly over 1.1 ( $=1.0085^{12}$ ). While somewhat high compared to e.g., Treasury bill returns, this figure is certainly plausible in magnitude. Further, the measures of EMRS are estimated with precision; the confidence interval is barely distinguishable from the means in the plots. Still, the most striking feature of the expected MRS is not its mean, but its volatility over time. The standard deviation of  $\{\hat{\delta}_t\}$  is around .04 for all estimators, and the point estimates vary over the decade between .88 and 1.09. This considerable volatility in the expected EMRS mirrors my (2003) results with Flood as well as the famously high lower bound of Hansen-Jagannathan (1991).<sup>8</sup>

### **3.1 Integration within the S&P 500**

Do the results depend sensitively on the exact choice of portfolios chosen? An easy way to check is to estimate  $\{\hat{\delta}_t\}$  using data from all twenty portfolios and look at the differences from the ten-portfolio estimates. This is done on the right side of Figure 1, which graphs the mean and confidence intervals of the expected MRS for the three estimators. In particular, the graphs on the right portray the difference between  $\{\hat{\delta}_t\}$  estimated from all twenty portfolios, and  $\{\hat{\delta}_t\}$  estimated from only the last ten portfolios. The differences are economically small; they average around .003 (for all three estimators). They also have large standard errors (of around .011), so that the differences do not appear to be statistically significant. In an integrated market, all securities should deliver the same expected marginal rate of substitution. Figure 1 thus delivers little evidence of significant departures from integration inside the S&P 500.

The rows on the right of Figure 1 compare  $\{\hat{\delta}_t\}$  on a period by period basis for a given estimator. That is, the figures implicitly ask whether the expected MRS for, say, February 1994, is the same when estimated from all twenty portfolios and only from the last ten. This is interesting because equality of  $\{\hat{\delta}_t\}$  derived from different assets is a necessary (but not sufficient) condition for market integration. But it is also interesting to compare the entire set of estimated EMRS simultaneously; that is, to test formally for joint equality. If the disturbances –  $\{\hat{u}_t^j\}$  – were normally distributed, this test would be easy to compute via a standard F-test. However and unsurprisingly, there is massive evidence of non-normality in the form of fat tails (leptokurtosis). Accordingly, I estimate the distribution for my critical values with a conventional bootstrap. With the bootstrapped results, I find the hypothesis of joint equality  $\{\hat{\delta}_t\}$  for all 118 observations cannot be rejected at any conventional significance level (for any estimator). That is, I cannot reject integration within the S&P 500. While this might only

indicate a lack of statistical power in the technique, I show later on that it is easy to reject equality of  $\{\hat{\delta}_t\}$  across substantively different markets.

### 3.2 Stock and Bond Markets

The hypothesis of equality of  $\{\hat{\delta}_t\}$  cannot be rejected when the twenty stock portfolios are split up. But are the estimated EMRS similar to treasury-bill returns? No. It is easy to generate the risk-free rate using an actual interest rate; I simply create  $\bar{\delta}_t \equiv (1 + i_t)$  where  $i_t$  is the monthly return on nominal treasury bills. The sample average of  $\{\bar{\delta}_t\}$  is around 1.003 (around 3.7% annualized), somewhat lower than but close to the sample average of  $\{\hat{\delta}_t\}$ .

But while the first moments of the estimated risk-free rate and the T-bill equivalent are similar, the second moments are not. The T-bill rate has considerably lower time-series volatility than the estimated EMRS. The standard deviation of  $\{\bar{\delta}_t\}$  (across time) is .001, which is smaller than that of  $\{\hat{\delta}_t\}$  by a factor of over thirty! Since the estimated risk-free rate is so much more volatile than the T-bill equivalent, it is unsurprising that the hypothesis of equality between the two can formally be rejected at any reasonable level of significance.<sup>9</sup>

To summarize, my estimates of the time-varying expectation of marginal rate of substitution are intuitively plausible in magnitude, and precisely estimated. They also display considerable volatility over time. While this variation is consistent with the literature, it is grossly at odds with the smooth T-bill return. Unsurprisingly, I can reject equality between my estimates of EMRS and those of the T-bill.

### 3.3 The TSE

What of different markets? Figure 2 provides estimates of the expected marginal rate of substitution (along with a +/- 2 standard error confidence interval) derived from the NYSE and the TSE. In both cases, I use twenty portfolios to estimate the expected risk-free rate. The number of available stocks differs by exchange; I use portfolios of nineteen stocks each from the NYSE, but those from the TSE have seven stocks. I estimate  $\{\hat{\delta}_t\}$  in the same way as above, using instrumental variables for 118 observations between 1994M2 and 2003M11. To facilitate comparison, I also graph the MRS implicit in the short Treasury bill return.

For both markets, the average value of EMRS seems reasonable, being slightly over unity. These are again estimated with considerable precision; the confidence interval can hardly be distinguished from the mean. But again, the single most striking feature of the estimates is their considerable time-series volatility. The standard deviation (over time) of  $\{\hat{\delta}_t\}$  is .04 for both the NYSE and the TSE. While this is consistent with received wisdom in Finance (e.g., Hansen and Jagannathan, 1991), it contrasts starkly with the smooth T-bill return portrayed at the bottom of Figure 2.

My results for both markets are tolerable and consistent with my earlier findings. Different estimators (OLS, GMM, and IV) deliver economically similar results which are statistically close (Hausman tests sometimes reject equality and sometimes do not). There is considerable leptokurtosis. And bootstrapped tests for internal integration indicate no evidence that using e.g., ten TSE portfolios delivers significantly different estimates of the expected MRS from using all twenty TSE portfolios. That is, I find no evidence against internal integration for either market.

I consider these results to be reassuring, given the depth and liquidity of the two stock markets I consider. But they might simply indicate a lack of power in my statistical techniques;

after all, they are simply not rejecting a necessary (but not sufficient) test for market integration. Accordingly, as a more stringent test, I also now test formally for integration across markets. This is also a subject of considerable intrinsic interest.

I begin comparing the estimated risk-free rate across markets with a series of scatter-plots in Figure 4. The top two graphs on the left of the figure compare monthly estimates of  $\{\hat{\delta}_t\}$  from the TSE (on the y-axis) against those derived from the NYSE (on the x-axis). At the bottom-left, I also provide a comparable graph using the T-bill rate on the ordinate. Clearly, the estimates of the expected MRS from the TSE are correlated with that from the NYSE; the correlation coefficient is .73. However they are not identical; the mean absolute difference between the  $\{\hat{\delta}_t\}$  derived from the NYSE and the TSE is .02, and almost 3% are greater than .1.

It is straightforward to formally test the hypothesis that the estimated EMRS are equal across markets. One way to do this is to test for equality between the estimates graphed in Figure 4. While this is perfectly acceptable (and the method I use for the daily results below), I note that the portfolios used to estimate  $\{\hat{\delta}_t\}$  graphed in Figure 4 have different numbers of stocks. Thus they have different degrees of estimation precision. To “balance the playing field” I construct twenty portfolios NYSE with nineteen stocks each. I can then use simple Chow tests to test for equality between  $\{\hat{\delta}_t\}$  derived from the twenty NYSE portfolios with those which *also* use the eight TSE portfolios (again, of nineteen stocks). When I do so, I find strong evidence against integration. The F-test for integration between the NYSE and the TSE is over 8, strongly rejecting the null hypothesis of integration (even allowing for non-normally distributed disturbances that I explore through the bootstrap). Evidently there is at least one source of divergence between the national stock markets of the United States and Canada.

Succinctly, while my estimates of the expected MRS are similar for the three different markets, they are significantly different in both the economic and statistical senses. That is, I am able to reject the hypothesis of equal EMRS across markets, and thus market integration. This result is intrinsically interesting, since there are few obvious reasons for this market segmentation. Moreover, they indicate that my methodology is not lacking in statistical power.

### 3.4 Daily Results

Thus far, I have used a decade of monthly data. I now present results derived from the most recent available year of daily data, 2003. I use closing rates for the 245 days when both markets were open, converting Canadian dollar quotes from the TSE into American dollars using a comparable exchange rate. I consider the same three markets, noting in passing that both the American and Canadian markets close at 4:00pm daily in the same time zone.

Figure 3 is the daily analogue to the monthly estimates displayed in Figure 2 (the GMM analogue is in Appendix Figure A1). In particular, I plot the mean of the expected MRS for all three markets, along with a  $\pm 2$  standard error confidence interval (the T-bill equivalent is also plotted at the bottom of the figure). I use IV as my estimator, though essentially nothing changes if I use OLS or GMM. In each case, I present estimate  $\{\hat{\delta}_t\}$  using twenty portfolios; each NYSE portfolio has 22 stocks, and each TSE portfolio eleven.

As with the monthly data set, the means of the series again seem reasonable; they are 1.001 and 1.001 for the NYSE and TSE respectively. These magnitudes seem intuitively reasonable, if somewhat high; they are roughly comparable in order of magnitude to the T-bill interest rate, which averaged just over 1% in 2003. The series of EMRS are also estimated with considerable precision manifest in tight confidence intervals. There is again evidence of

leptokurtosis. Still, the most striking feature of all three series of the estimated MRS is their volatility over time. This is especially true when one compares them with the virtually flat T-bill return. It is little surprise then that the hypothesis that the daily estimates of  $\{\hat{\delta}_t\}$  derived from S&P 500 stock prices are statistically far from the T-bill equivalent  $\{\bar{\delta}_t\}$ .<sup>10</sup>

When I check for internal integration within a market (such as S&P 500 stocks traded on the NYSE) by comparing estimates of  $\{\hat{\delta}_t\}$  derived from different sets of portfolios, I am unable to reject the hypothesis of equality at any reasonable confidence interval. That is, I (unsurprisingly) find no evidence against integration *within* markets.

However, as with the monthly data, integration *across* markets is another story. The scatter-plots of the estimated daily EMRS at the right side of Figure 4 are analogous to those with monthly data immediately to the left. The TSE delivers  $\{\hat{\delta}_t\}$  that are positively correlated with those from the NYSE; the correlation coefficient is .69. The mean absolute difference between the series is around .006, and range to over .02. While these may seem small, they are economically large since they are at a daily frequency. In any case, the series are statistically distinguishable. When I test for equality between the estimates of  $\{\hat{\delta}_t\}$  portrayed in Figures 3/4, I find the hypothesis rejected for the NYSE against the TSE (the F-test statistic is over 17). Bootstrapping the critical values does not reverse these conclusions.

In brief, the daily data set produces similar results to those of my monthly data set. The estimates of the expected MRS seem intuitively reasonable, and display volatility consistent with that in the literature, but far in excess of the T-bill. While I can never reject the hypothesis of internal integration, I always reject the hypothesis of integration across markets in the sense of equal EMRS.

## 4 Conclusion

In this paper, I have developed a methodology for estimating the expected intertemporal marginal rate of substitution. The technique relies on exploiting the general fact that idiosyncratic risk, which does not alter any risk premia, should deliver a return equal to the market's expectation of the marginal rate of substitution. This enables me to estimate the expected risk-free rate from equity price data, an object that is intrinsically interesting. Comparing the rates estimated from different markets also provides a natural test for market integration, since integrated markets should share a common expected MRS.

I apply the methodology to a decade of monthly data and a year of daily data, including data on stocks traded on the New York Stock Exchange and the Toronto Stock Exchange. For both data sets, I find intuitive estimates of the expected marginal rate of substitution with reasonable means and considerable volatility over time. I cannot reject the hypothesis that markets are internally integrated in the sense that different portfolios traded on a given market seem to have the same expected marginal rate of substitution. However, I find it easy to reject the hypothesis of equal EMRS across markets. This is both of direct interest, and indicates that the technique has considerable statistical power.

There are many possible ways to extend the work. One could add a covariance model – e.g., the well-known factor model first developed by Fama and French (1996) – to equation (8). A well-specified covariance model should result in more efficient estimates of the EMRS. Alternatively, one could sort stocks into portfolios in some systematic way (e.g., size, industry, or beta). More factors could be added to the first stage regression, equation (9). More generally, while the use of the  $\tilde{p}_t^j$  normalization has advantages, others might be used instead. One could

test for excess returns that should be possible if EMRS diverges across markets, and if the EMRS is not equal to the t-bill rate. Most importantly, while I have been able to reject the hypothesis of integration in the sense of equal expected marginal rates of substitution across markets, I have not explained the reasons for this finding of apparent market segmentation. If my result stands up to scrutiny, this important task remains.

## References

Adam, Klaus, Tullio Jappelli, Annamaria Menichini, Mario Padula, and Marco Pagano (2002) “Analyse, Compare, and Apply Alternative Indicators and Monitoring Methodologies to Measure the Evolution of Capital Market Integration in the European Union” University of Salerno manuscript.

Bekaert, Geert and Campbell R. Harvey (1995) “Time-Varying World Market Integration” *The Journal of Finance* 50-1, 403-444.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay (1997) *The Econometrics of Financial Markets* (Princeton: University Press).

Cochrane, John H. (2001) *Asset Pricing*, Princeton University Press.

Fama, Eugene F. and Kenneth R. French (1996) “Multifactor Explanations of Asset Pricing Anomalies” *Journal of Finance* 51-1, 55-84.

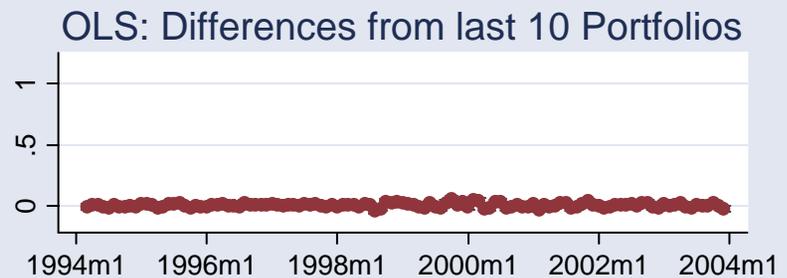
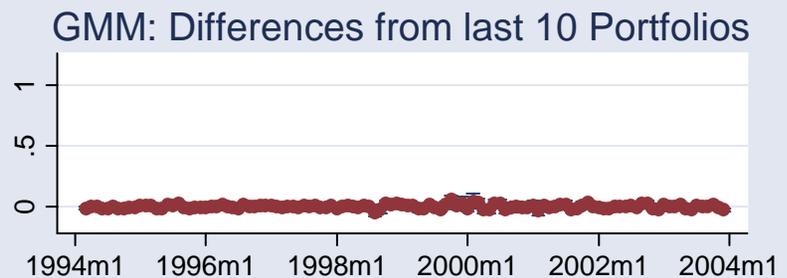
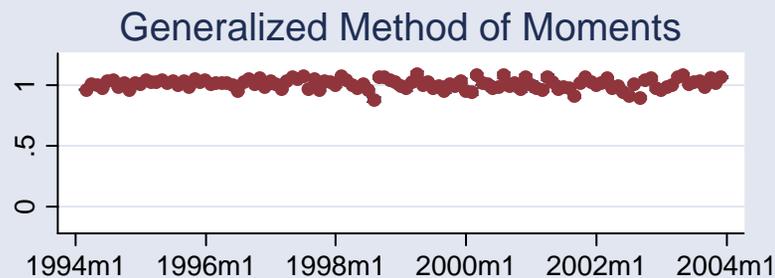
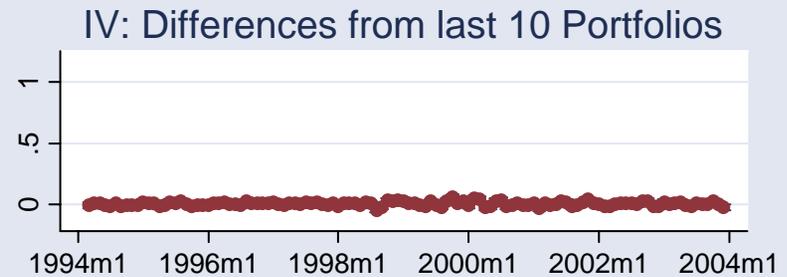
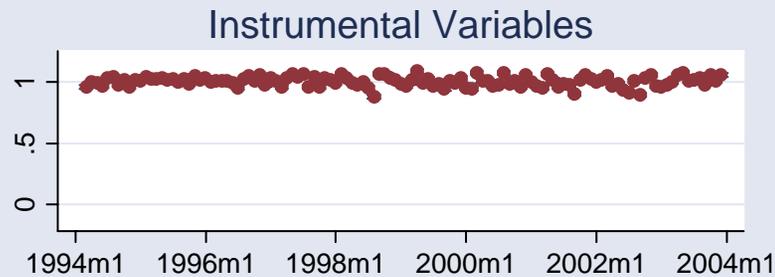
Fama, Eugene F. and James D. MacBeth (1973) “Risk, Return and Equilibrium: Empirical Tests” *Journal of Political Economy* 81-3, 607-636.

Flood, Robert P. and Andrew K. Rose (2003) “Financial Integration: A New Methodology and an Illustration” NBER Working Paper #9880.

Hansen, Lars Peter and Ravi Jagannathan (1991) “Implications of Security Market Data for Models of Dynamic Economies” *Journal of Political Economy* 99-2, 225-262.

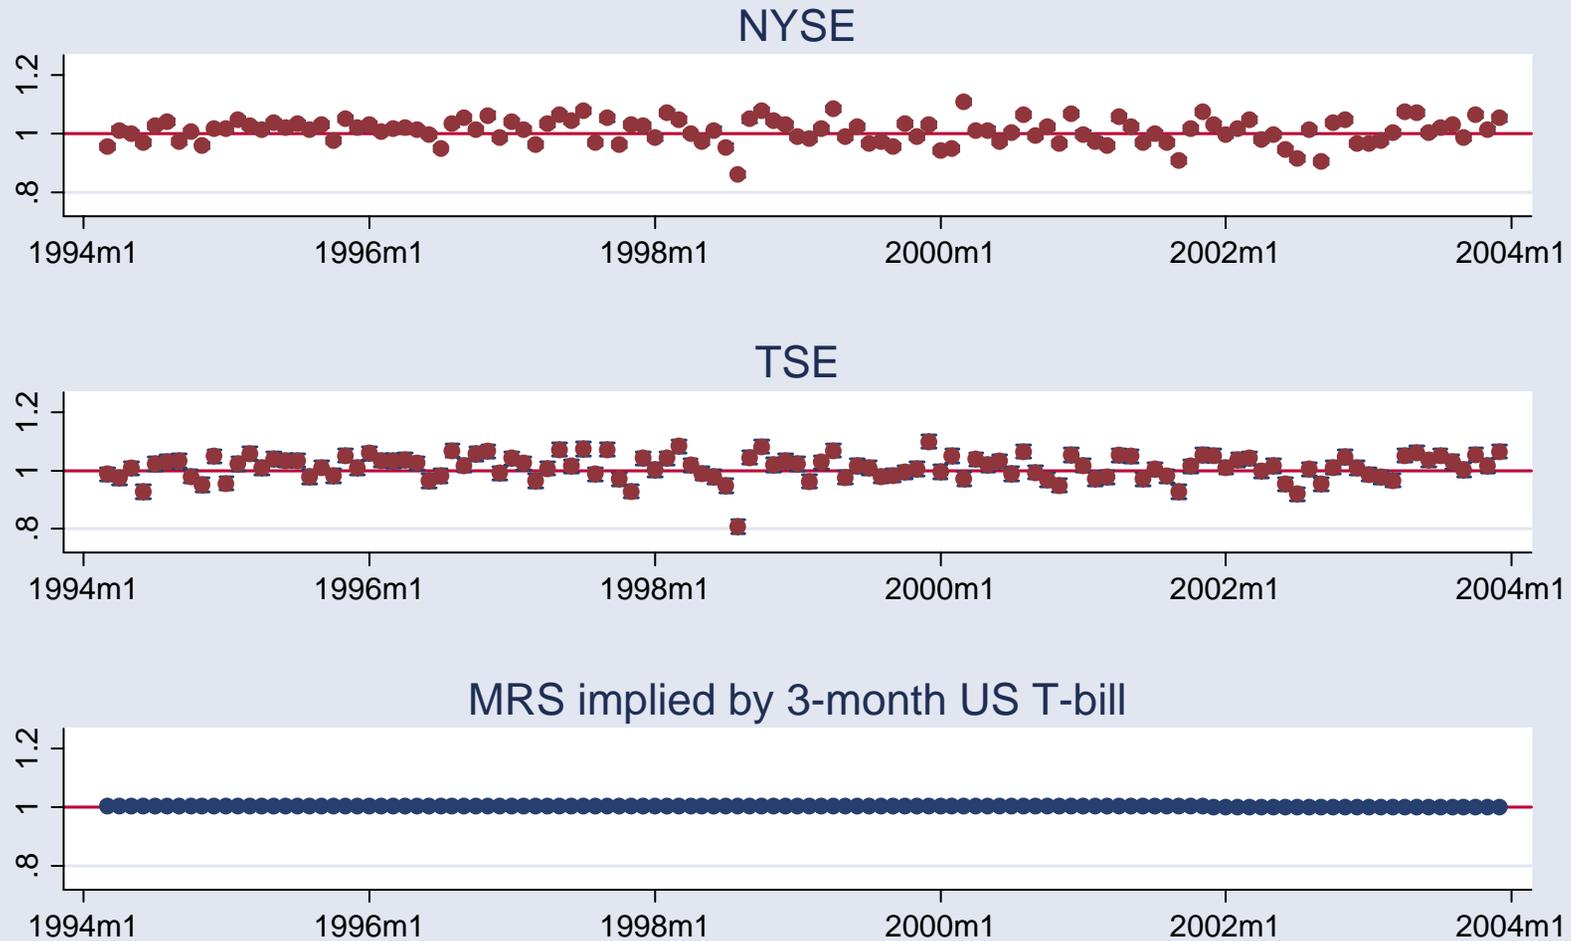
Shanken, Jay (1992) “On the Estimation of Beta-Pricing Models” *Review of Financial Studies* 5-1, 1-33.

# Expected MRS, first 10 portfolios; S&P 500 1994M2-2003M11 Deltas, with +/- 2 S.E. Confidence Interval



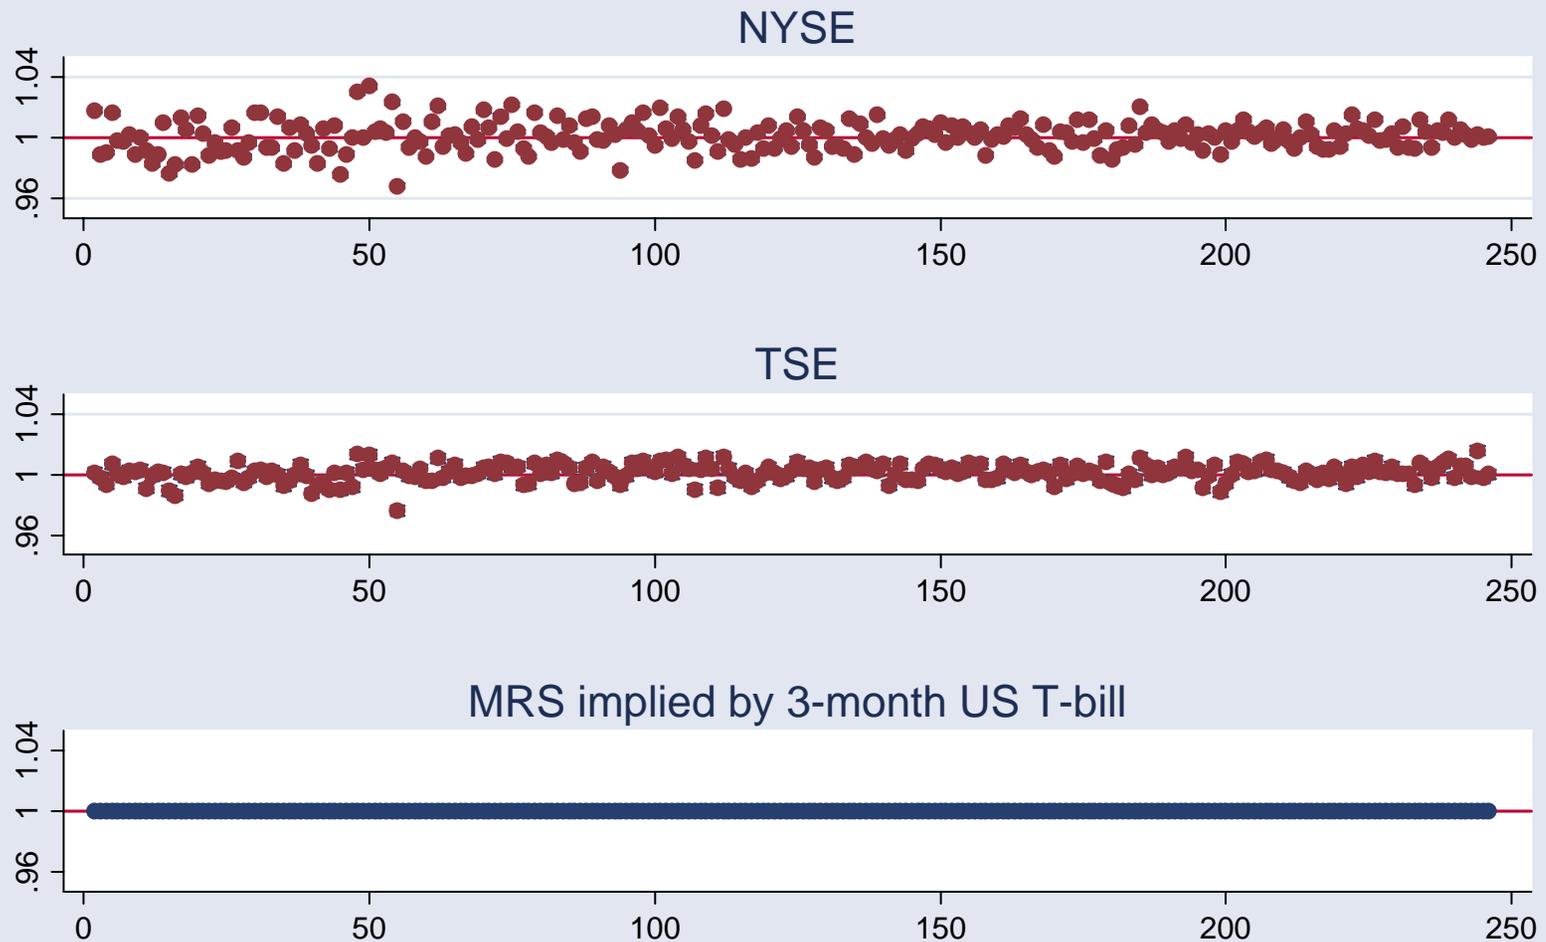
**Figure 1: Estimated Expected MRS, Portfolios of 20 S&P500 firms, 1994M2-2003M11: Different Estimators**

# Expected MRS from 20 portfolios, IV, 1994M2-2003M11 Deltas, with +/- 2 S.E. Confidence Interval



**Figure 2: Estimates of Expected Marginal Rate of Substitution, 1994M2-2003M11: Different Markets**

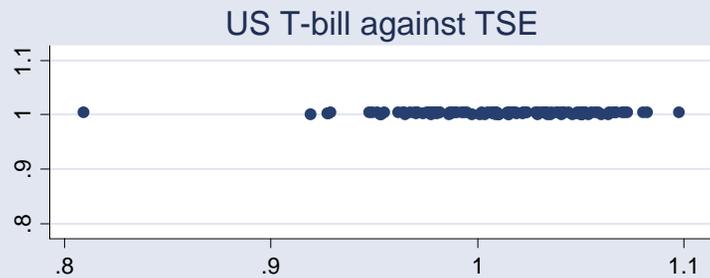
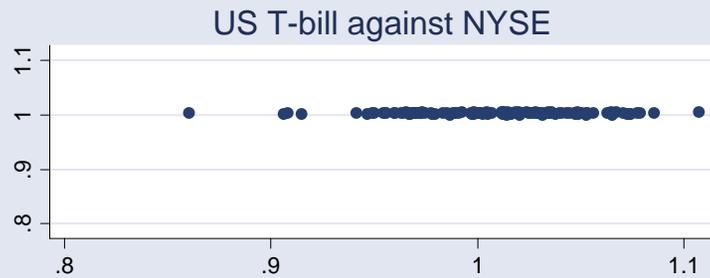
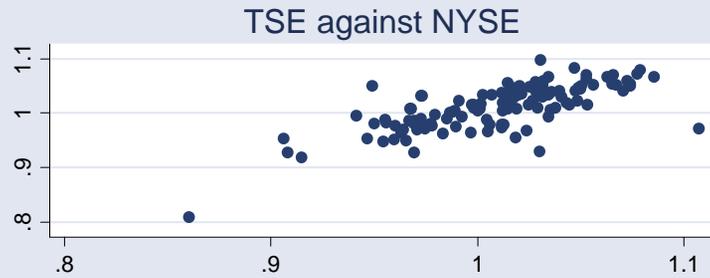
# Expected MRS from 20 portfolios, IV, Daily Data 2003 Deltas, with +/- 2 S.E. Confidence Interval



**Figure 3: IV Daily Estimates of Expected Marginal Rate of Substitution, 2003**

# Expected MRS from 20 portfolios, IV

Monthly Data, 2004-2003



Daily Data, 2003

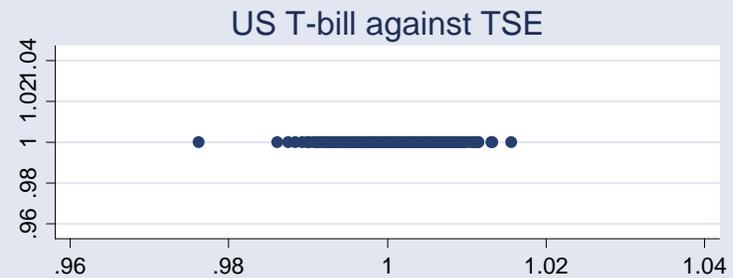
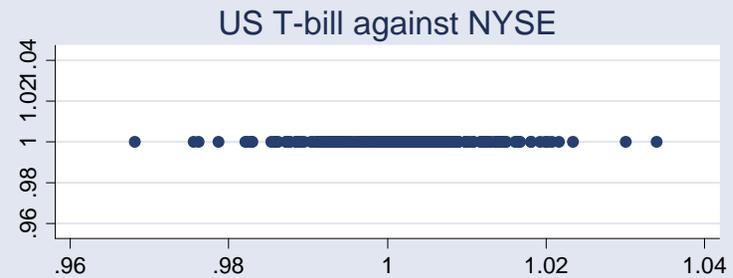
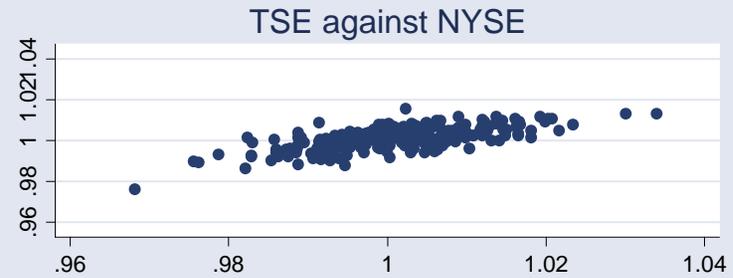
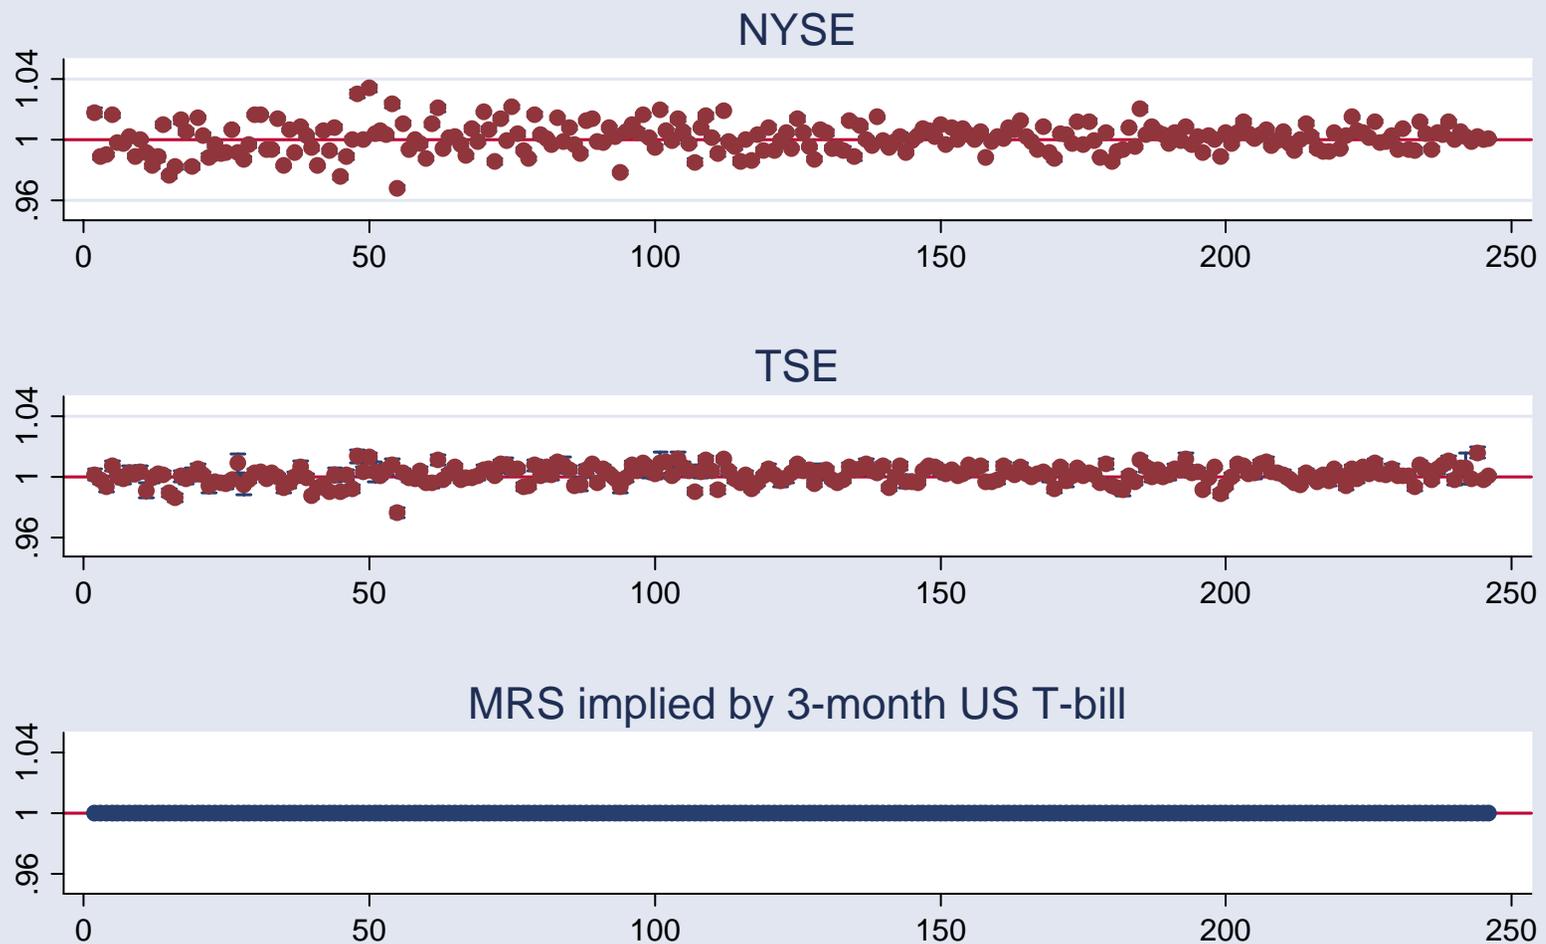


Figure 4: Scatter-plots of Estimated Expected MRS across Markets

# Expected MRS from 20 portfolios, GMM, Daily Data 2003 Deltas, with +/- 2 S.E. Confidence Interval



**Figure A1: GMM Daily Estimates of Expected Marginal Rate of Substitution, 2003**

## Endnotes

<sup>1</sup> In general, there is no guarantee that  $m_{t+1}$  is a unique variable; agents behaving according to equation (1) use the entire perceived distribution of  $m_{t+1}$  to price assets at  $t$ .

<sup>2</sup> I have experimented with additional regressors suggested by Fama and French (1996), and they seem to make no difference to the results in practice.

<sup>3</sup> For instance, the median  $R^2$  from the twenty estimates of equation (9) is a respectable .77, and the lowest of the twenty  $R^2$ s is still .59.

<sup>4</sup> This might lead to selection bias, but of ambiguous sign. Firms that disappear from the sample leave because of either positive idiosyncratic shocks and mergers/takeovers, or negative ones and bankruptcy/takeover. This issue is worthy of more work.

<sup>5</sup> These firms are traded mostly on the NYSE, but a few are on the NASDAQ.

<sup>6</sup> One can compare the differences between the estimators with a Hausman test. In this case, the difference between the OLS and IV estimators turns out to be economically small but marginally statistically significant; the hypothesis of equality is rejected at the .006 confidence interval. Also, robust standard errors (either clustered by portfolio or not) are typically even smaller.

<sup>7</sup> The exception is bootstrapped tests for integration, where I tend to use OLS for computational simplicity.

<sup>8</sup> The estimates of EMRS have essentially no persistence and are also uncorrelated with traditional finance factors, such as the three used by Fama and French (1996). Adding either an intercept or portfolio-specific intercepts to the estimating equation (11) changes results little, which is unsurprising since the former are small and of marginal significance; the same is basically true of time-specific intercepts. Finally, I have added the three time-varying Fama-French factors to the first stage equation (9); this makes any substantial difference to my results.

<sup>9</sup> The F-test statistic for equality between the expected MRS and the T-bill return is over 50; under the null hypothesis of market integration, it has degrees of freedom (118, 2360). Bootstrapping the critical values has no substantive effect on conclusions.

<sup>10</sup> The F-test statistic for equality between the expected MRS and the t-bill return is over 150.