Session 1:
Financial Market Linkages

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

*Andrew K. Rose**

In this paper, I develop and apply a simple methodology to estimate the expected (intertemporal) marginal rate of substitution (EMRS). The EMRS is an economic variable of considerable 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.

My primary objective is expositional, and this paper 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 section 2. The empirical results are presented in section 3, and the paper ends with a brief conclusion.

---

* I thank conference participants at the Bank of Canada and, in particular, Eric Santor, Lawrence Schembri, and Gregor Smith for comments. This study draws heavily on my paper, "Estimating the Expected Marginal Rate of Substitution," co-written with Robert P. Flood.

1 The Methodology

It is easy to motivate this exercise. Asset-market integration is a topic of continuing interest in international finance; see, for example, Adam et al. (2002). It is of special interest in Europe, where continuing monetary and institutional integration have led to lower barriers to asset trade inside the European Union. But economists from other countries should also be interested; after all, the "open" in "small open economy" refers typically to asset-market integration, and is a critical assumption for macroeconomists modelling 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 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 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, $E_t()$ is the expectations operator conditional on information available at t, m_{t+1} is the time-varying intertemporal 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 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 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.[1] 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 $COV_t()$ denotes the conditional covariance operator. It is useful to rewrite this as

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

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 MRS should be equal inside integrated financial markets. My work below is concerned essentially 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 ε_{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, $\delta_t \equiv 1/E_t(m_{t+1})$,

1. 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*.

and asset-specific period risk premiums, the covariance term. Equation (4) is given economic content by adding two assumptions:

(i) *rational expectations*: ε^j_{t+1} is assumed to be uncorrelated with information available at time *t*, and

(ii) *covariance model*: $COV_t(m_{t+1}, x^j_{t+1}/p^j_t) = \beta^j_0 + \sum_i \beta^j_i f_{i,t}$, for the relevant sample,

where β^j_0 is an asset-specific intercept, β^j_i is a set of *I* asset-specific factor coefficients, and $f_{i,t}$ is 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 (4) 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.

In practice, many empirical asset-pricing modellers 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, for example, the treasury bill rate; it is not estimated at all. While this simplifies empirical work considerably, it 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 premiums.

1.2 New stuff

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

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

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 $\tilde{p}_t^j$,

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

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

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

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

The first term inside the brackets, $(p_t^i / \tilde{p}_t^j)$, 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, $p_{t+1}^i / \tilde{p}_t^j$. 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, $(p_t^i / \tilde{p}_t^j)$, is orthogonal to systematic risk, $COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$, equation (7) can be decomposed as

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

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, ε_{t+1}^j, is a forecasting error that 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 idiosyncratic phenomena. Since both terms are orthogonal to the regressor that represents idiosyncratic risk, $(p_t^i / \tilde{p}_t^j)$, the coefficients of interest, $\{\delta\}$, can be consistently estimated via equation (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 EMRS, 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 centrepiece 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 premiums, 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 EMRS—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, I also reject integration.

This normalization has the advantage—in common with the strategy of Flood and 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 $COV_t(m_{t+1}, x^j_{t+1}/\tilde{p}^j_t)$, such as, e.g., the capital asset-pricing model 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 $COV_t(m_{t+1}, x^j_{t+1})$, 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 $COV_t(m_{t+1}, x^j_{t+1})$. 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 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 classes of intertemporal decisions, including savings 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 focuses on estimating an intrinsically interesting object, the (inverse of the) EMRS.

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 $\hat{p}_t^j$. (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):

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

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:

$$\hat{p}_t^j \equiv p_t^j {}^{*}\exp(\hat{a}_j / \hat{b}_j, {}^{*}\ln(\bar{p}_t / \bar{p}_{t-1})). \tag{10}$$

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.[2] 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 estimate $\hat{p}_t^j$, but a cross-sectional approach is also possible. None of these assumptions is critical; they simply seem to work in practice.[3] But while this setup has delivered sensible results, I stress that one only needs *some* model for $\hat{p}_t^j$, not this particular one.

2.1 Estimation issues

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

---

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

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

$$x^j_{t+1}/\hat{p}^j_t = \delta_t(p^i_t/\hat{p}^j_t) + u^j_{t+1} \tag{11}$$

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

Using $\hat{p}^j_t$ in place of the unobservable $\tilde{p}^j_t$ might induce important measurement error. Hence, it is natural to consider estimation of equation (11) with instrumental variables (IVs) for consistent estimation of $\{\delta\}$. IV is also known to handle the “generated-regressor” issue, which has long been associated with potentially overstated precision of standard errors; see Shanken (1992) and Cochrane (2001 and website correction[4]). The latter shows that this is typically not important in practice, especially for monthly data. While IV estimation seems natural, estimation via generalized method of moments (GMM) allows me to handle both potential econometric issues, while not requiring independent and identically distributed disturbances. Accordingly, I use and compare two estimation techniques, IV and GMM (I sometimes also use OLS for simplicity). As IVs for $\{p^j_t/\hat{p}^j_t\}$, I use the set of time-varying market-wide average prices, $\{\bar{p}_t\}$.

2.2 The data sets

I employ two 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 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 Collaborative Research Support Program (CRSP) database and consist of month-end prices and returns (including dividends, if any) for all firms in the Standard and Poor (S&P) 500 (as of the end of 2003). I have adjusted for stock splits and have checked and corrected the data. I retain only the 435 companies that have data for the full sample span.[5]

---

4. <http://gsbwww.uchicago.edu/fac/john.cochrane/research/Papers/typos.pdf>.

5. This might lead to selection bias, but of ambiguous sign. Firms that disappear from the sample leave because of either positive idiosyncratic shocks and merger/takeover, or negative ones and bankruptcy/takeover. This issue merits more attention.

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 (TSX). 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 observations on 389 firms from the S&P 500 traded on the NYSE and 152 firms from the TSX. 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 TSX 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 Treasury bill returns. I use data on Treasury bill returns downloaded from the website of the Board of Governors of the Federal Reserve System.

Finally, for purposes of estimation, I place stocks into portfolios, typically in groups of twenty. I do this randomly (alphabetical order of ticker), though I see no reason why one could not group firms on the basis of beta or size, for example. 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 does not hinder my ability to estimate the parameters of interest to us, although 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 EMRS, and I begin with an illustration that relies on monthly data from 400 firms in the S&P 500, grouped into portfolios of twenty firms.[6] I have 118 observations between February 1994 and November 2003, since I lose an observation at either end of the sample, owing to leads/lags.

The three graphs on the left in Figure 1 portray estimates of the EMRS from equation (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 +/–2 standard error

6. These firms are traded mostly on the NYSE, but a few are on the NASDAQ.

Figure 1
Estimated EMRS, portfolios of 20 S&P 500 firms, 1994M2–2003M11: different estimators

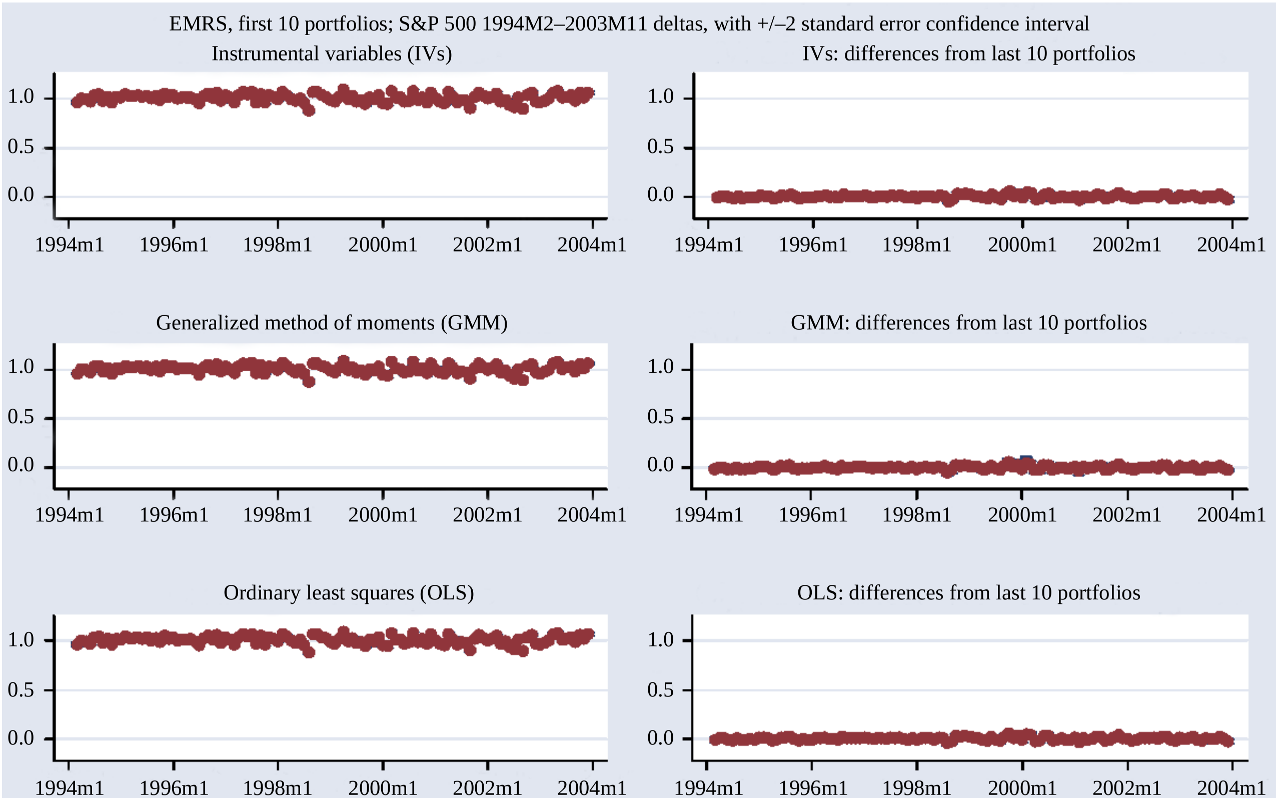

confidence interval band. The OLS and GMM point estimates are identical (by construction) and are highly correlated with the IV estimates. The primary differences between the different estimates lie in the standard errors; all three estimators deliver small standard errors, with the GMM standard errors being slightly smaller than those of either IV or OLS (but with more period-to-period volatility).[7] Indeed, I rarely find significant differences between the three estimators, and thus tend to rely on IV.[8]

Even though I estimate the expected MRS from only ten portfolios, the results appear 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, for example, Treasury bill returns, this figure is certainly plausible in magnitude. Furthermore, 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 EMRS 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 EMRS mirrors my (2003) results with Flood as well as the famously high lower bound of Hansen-Jagannathan (1991).[9]

3.1 Integration within the S&P 500

Do the results depend on the choice of portfolios? An easy way to check is to estimate $\{\hat{\delta}_t\}$, using data from all twenty portfolios and to examine 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 EMRS for the three estimators. In particular, the graphs on the right portray the difference between $\{\hat{\delta}_t\}$ estimated from all twenty portfolios, and $\{\hat{\delta}_t\}$

---

7. 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.

8. The exception is bootstrapped tests for integration, where I tend to use OLS for computational simplicity.

9. 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 not surprising 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 no substantial difference to my results.

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 EMRS. Figure 1 thus delivers little evidence of significant departures from integration inside the S&P 500.

The columns 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 EMRS 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 that 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 will show 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 $\vec{\delta}_t \equiv (1 + i_t)$, where i_t is the monthly return on nominal Treasury bills. The sample average of $\{\vec{\delta}_t\}$ is around 1.003 (around 3.7 per cent 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 Treasury bill equivalent are similar, the second moments are not. The Treasury bill rate has considerably lower time-series volatility than the estimated EMRS. The standard deviation of $\{\vec{\delta}_t\}$ (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 Treasury bill equivalent, it is unsurprising that the hypothesis of equality between the two can formally be rejected at any reasonable level of significance.[10]

To summarize, my estimates of the time-varying expectation of MRS 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 Treasury bill return. Unsurprisingly, I can reject equality between my estimates of EMRS and those of the Treasury bill.

3.3 The TSX

What of different markets? Figure 2 provides estimates of the EMRS (along with a +/–2 standard error confidence interval) derived from the NYSE and the TSX. 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 TSX have seven stocks. I estimate $\{\hat{\delta}_t\}$ in the same way as above, using IVs 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 time-series volatility. The standard deviation (over time) of $\{\hat{\delta}_t\}$ is .04 for both the NYSE and the TSX. While this is consistent with received wisdom in finance (e.g., Hansen and Jagannathan 1991), it contrasts starkly with the smooth Treasury 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 that 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 TSX portfolios delivers significantly different estimates of the expected MRS from using all twenty TSX portfolios. That is, I find no evidence against internal integration for either market.

I consider these results reassuring, given the depth and liquidity of the two stock markets I consider. But the results might simply indicate a lack of

---

10. The F-test statistic for equality between the EMRS and the Treasury 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.

Figure 2
Estimated EMRS, 1994M2–2003M11: different markets

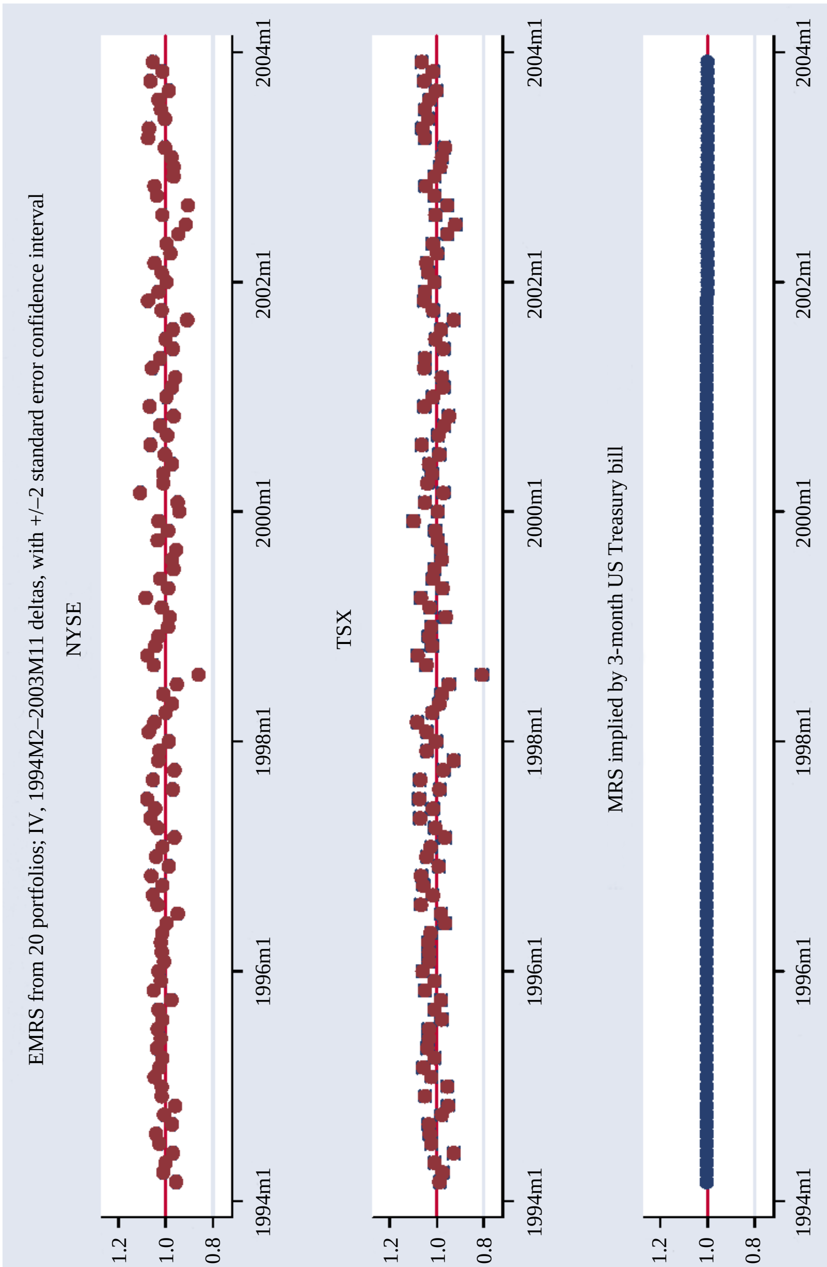

power in my statistical techniques; after all, they are simply not rejecting a necessary (but not sufficient) test for market integration. Accordingly, as a more stringent measure, I also test formally for integration across markets.

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 TSX (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 Treasury bill rate on the ordinate. Clearly, the estimates of the EMRS from the TSX 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 TSX is .02, and almost 3 per cent 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 is 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 NYSE portfolios 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 that also use the eight TSX 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 TSX 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.

While my estimates of the EMRS are similar for the two 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 interesting, since there are few obvious reasons for this segmentation. Moreover, the estimates 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 TSX into American dollars, using a comparable exchange rate. I consider the same two markets, noting that both the American and Canadian markets close at 4:00 p.m. 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 Figure 5). In particular, I plot the mean of the EMRS for both markets, along with a +/–2 standard error confidence interval (the Treasury bill equivalent is also plotted at the bottom of the figure). I use IV as my estimator, although essentially nothing changes if I use OLS or GMM. In each case, I estimate $\{\hat{\delta}_t\}$ using twenty portfolios; each NYSE portfolio has 22 stocks, and each TSX portfolio 11.

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 TSX, respectively. These magnitudes seem intuitively reasonable, if somewhat high; they are roughly comparable in order of magnitude to the Treasury bill interest rate, which averaged just over 1 per cent in 2003. The series of EMRS are also estimated with considerable precision 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 Treasury 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 Treasury bill equivalent $\{\hat{\delta}_t\}$.[11]

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 TSX 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 approximately .006, and ranges 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 and 4, I find the hypothesis rejected for the NYSE against the TSX (the *F*-test statistic is over 17). Bootstrapping the critical values does not reverse these conclusions.

11. The *F*-test statistic for equality between the EMRS and the Treasury bill return is over 150.

Figure 3
IV daily estimates of EMRS, 2003

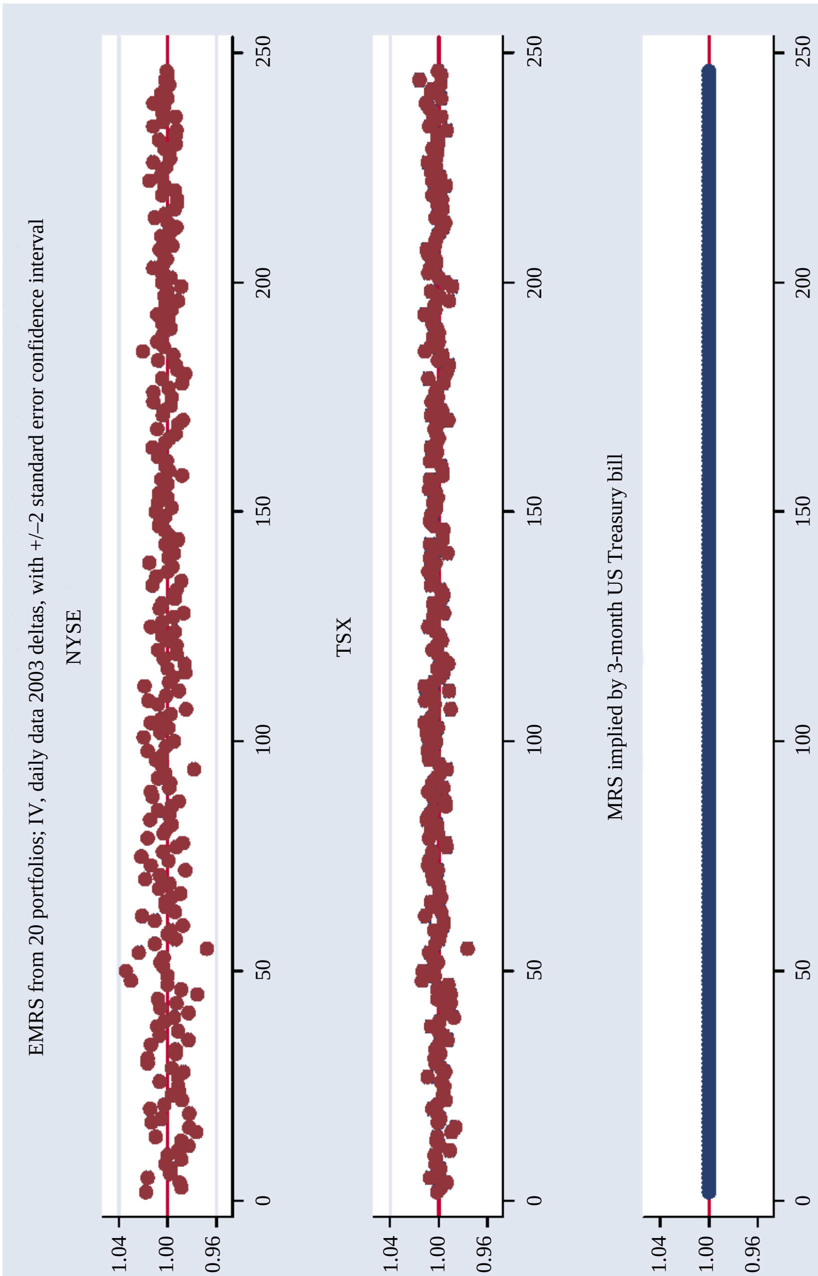

Figure 4
Scatter plots of estimated EMRS across markets

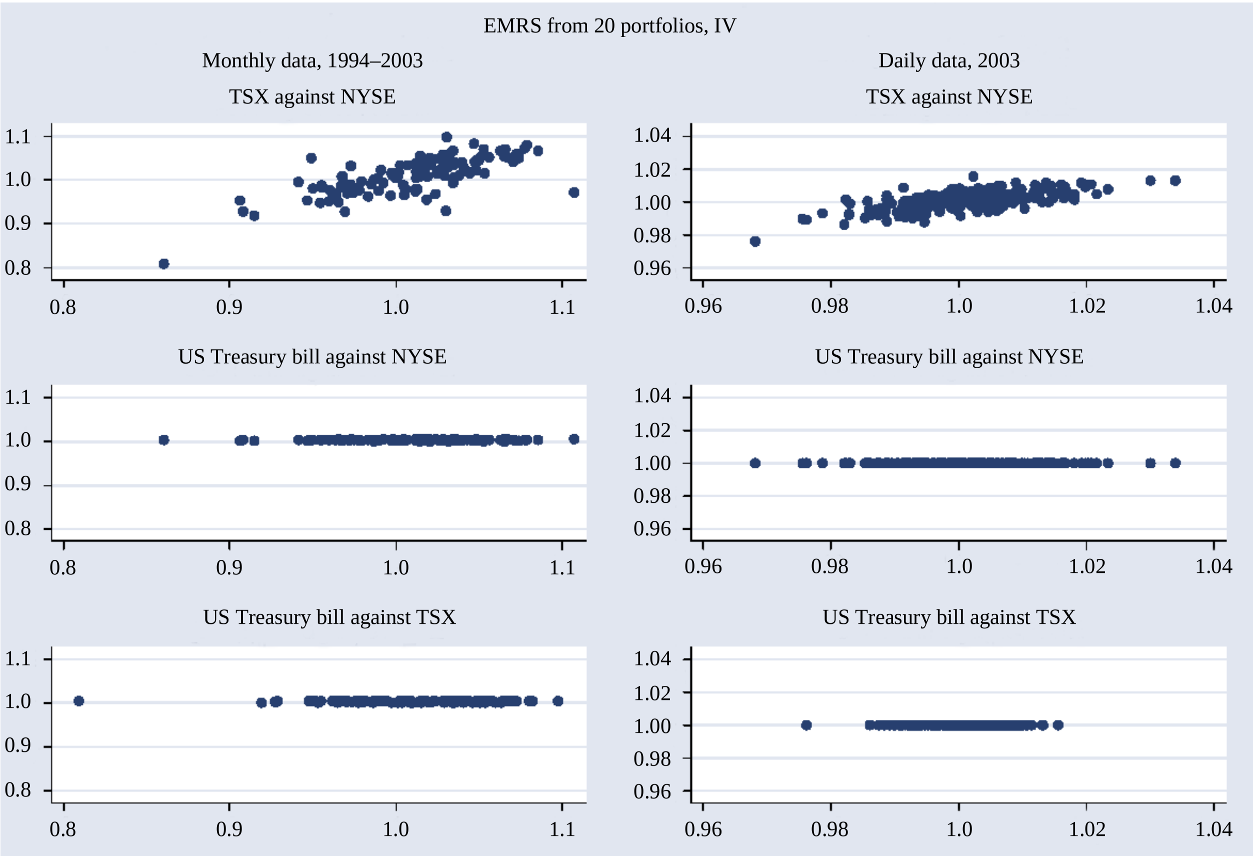

Figure 5
GMM daily estimates of EMRS, 2003

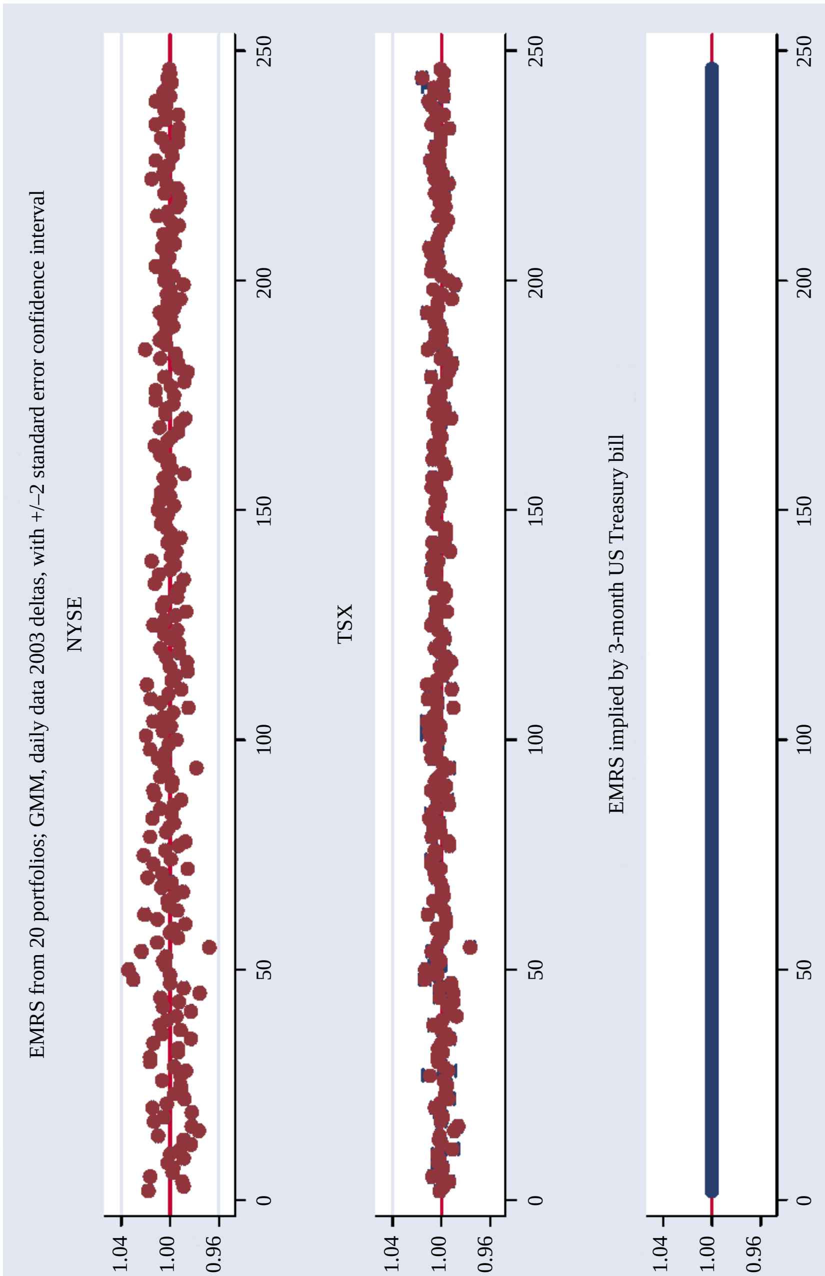

In brief, the daily data set produces similar results to those of my monthly data set. The estimates of the EMRS seem intuitively reasonable, and display volatility consistent with that in the literature, but far in excess of the Treasury bill. While I can never reject the hypothesis of internal integration,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

Conclusion

I have developed a methodology for estimating the EMRS that relies on exploiting the fact that idiosyncratic risk, which does not alter any risk premiums, should deliver a return equal to the market's expectation of the MRS. This enables me to estimate the expected risk-free rate from equity-price data. Comparing the rates estimated from different markets also provides a natural test for market integration, since integrated markets should share a common EMRS.

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 EMRS 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 EMRS. However, I find it easy to reject the hypothesis of equal EMRS across markets. This is of interest and would indicate that the technique has considerable statistical power.

There are many ways to extend the work. One could add a covariance model—e.g., the well-known factor model 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). 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 Treasury bill rate. Most importantly, while I have been able to reject the hypothesis of integration in the sense of equal EMRS across markets, I have not explained the reasons for this apparent market segmentation. If my result stands up to scrutiny, this important task remains.

References

Adam, K., T. Jappelli, A. Menichini, M. Padula, and M. 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. Photocopy.

Bekaert, G. and C.R. Harvey. 1995. "Time-Varying World Market Integration." *Journal of Finance* 50 (2): 403–44.

Campbell, J.Y., A.W. Lo, and A.C. MacKinlay. 1997. *The Econometrics of Financial Markets*. Princeton: Princeton University Press.

Cochrane, J.H. 2001. *Asset Pricing*. Princeton: Princeton University Press.

Fama, E.F. and K.R. French. 1996. "Multifactor Explanations of Asset Pricing Anomalies." *Journal of Finance* 51 (1): 55–84.

Fama, E.F. and J.D. MacBeth. 1973. "Risk, Return, and Equilibrium: Empirical Tests." *Journal of Political Economy* 81 (3): 607–36.

Flood, R.P. and A.K. Rose. 2003. "Financial Integration: A New Methodology and an Illustration." National Bureau of Economic Research Working Paper No. 9880.

Hansen, L.P. and R. Jagannathan. 1991. "Implications of Security Market Data for Models of Dynamic Economies." *Journal of Political Economy* 99 (2): 225–62.

Shanken, J. 1992. "On the Estimation of Beta-Pricing Models." *Review of Financial Studies* 5 (1): 1–33.

Discussion

Eric Santor

The motivation for this paper is simple: to estimate the expected (intertemporal) marginal rate of substitution (EMRS). Deviating from most other papers in the literature that use some variation of Fama-French factors to explain asset returns, Rose develops a new, simple methodology. He then applies this methodology to test for market integration between US and Canadian equity markets. As a result, he is able to reject the hypothesis of market integration for the United States and Canada.

His methodology is straightforward. First, he estimates the returns to asset price j as a function of the average market-price growth:

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

where p_t^j is the price of asset j at time t, $\bar{p}_t$ is the average market price, a_j and b_j are fixed-regression coefficients, and v_t^j is an idiosyncratic shock. From these estimates, he constructs a "synthetic" price for each asset j.

$$\hat{p}_t^j \equiv p_{t-1}^j \times \exp\left(\hat{a}_j + \hat{b}_j \ln(\bar{p}_t / \bar{p}_{t-1})\right) \tag{2}$$

Using this synthetic price, the EMRS is estimated from the following equation:

$$x_{t+1}^j / \hat{p}_t^j = \delta\left(p_t^j / \hat{p}_t^j\right) + u_{t+1}^j, \tag{3}$$

where x_{t+1}^j is the income received by the owners of asset j and u_{t+1}^j is the error term. The estimate of δ is the EMRS. Not surprisingly, the estimate of δ depends heavily on the estimated synthetic price.

The model is taken to the data using daily returns from 2003 and monthly returns from 1994–2003 for the Standard and Poor 500 (S&P 500), NASDAQ, and the Toronto Stock Exchange (TSX). The monthly data cover 389 US firms and 152 Canadian firms, and the daily data cover 440 US firms and 223 Canadian firms. The data for Treasury bills are taken from the Federal Reserves' website. The estimation of equation (3) using instrumental variables produced reasonable measures of EMRS. Using these estimates, Rose then tests for market integration. If the measures of EMRS differ across markets, then market integration can be rejected. He finds evidence of integration between the S&P 500 and NASDAQ. However, he rejects integration between the S&P 500 and Treasury bills. Most importantly, he rejects the hypothesis of market integration between the S&P 500 and the TSX. Rose's results are compelling evidence that full integration between the US and Canadian equity markets has yet to occur. However, I have a number of concerns regarding the methodology.

The first is that to compare returns, the Canadian returns are converted to US-dollar returns using the contemporaneous Can$/US$ exchange rate. But could exchange rate effects drive a wedge between the respective EMRSs? For instance, if a common factor affects asset prices in the US and Canadian markets, would this factor also affect the exchange rate in a non-random manner? If movements in the exchange rate are correlated to asset-price movements, this could drive a wedge between the respective measures of EMRS.

My second concern is the assumption of the US investor. It is possible that US investors have different preferences from other investors. More generally, should we expect that the EMRS be equivalent across markets? Most likely, there is considerable investor heterogeneity, which implies that there will be differences in preferences over risk. It is not clear, therefore, that this simple test of market integration is appropriate.

My third concern relates to the methodology of constructing the synthetic price. First, does the model matter? Would the inclusion of Fama-French factors improve the model's efficiency? And should one include exchange rate risk? Second, the results depend on the use of instrumental variables. Does the market-wide average price satisfy the conditions of a valid instrument? Showing first-stage results would be helpful in this regard.

The test of market integration relies on the notion that the portfolios used to estimate the EMRS for each market are comparable. But could heterogeneity in the sample portfolios bias the results? The problem of heterogeneity may be particularly acute in the case of Canada, since there are a large number of family-owned firms listed on the TSX (typified by dual-class share structures) relative to the New York Stock Exchange

(NYSE). A possible solution to this problem of sample portfolio heterogeneity is to consider only Canadian firms that are cross-listed on the NYSE. In a similar vein, estimation of the EMRS may be biased owing to heterogeneous treatment effects. That is, common factors may affect different types of firms differentially. To account for heterogeneity in the sample portfolios, the author could consider a matching-methods approach. This would help to control for sample bias, which may in turn affect the estimation of the EMRS across the two samples.

In conclusion, Rose provides a simple and novel means for estimating EMRS, which is important, given its central role in many macroeconomic and finance models.