Estimating the Expected Marginal Rate of Substitution: A Systematic Exploitation of Idiosyncratic Risk

Robert P. Flood and Andrew K. Rose*

Revised: March 9, 2005

Robert P. Flood
Research Dept, IMF
700 19th St., NW
Washington, DC 20431
Tel: (202) 623-7667
Fax: (202) 589-7667
E-mail: rflood@imf.org

Andrew K. Rose (correspondence)
Haas School of Business
University of California
Berkeley, CA 94720-1900
Tel: (510) 642-6609
Fax: (510) 642-4700
E-mail: arose@haas.berkeley.edu

Abstract

We develop a methodology to estimate the shadow risk free rate or expected intertemporal marginal rate of substitution, "EMRS". Our technique relies upon exploiting idiosyncratic risk, since theory dictates that idiosyncratic shocks earn the EMRS. We apply our methodology to recent monthly and daily data sets for the New York and Toronto Stock Exchanges. We estimate EMRS with precision and considerable time-series volatility, subject to an identification assumption. Both markets seem to be internally integrated; different assets traded on a given market share the same EMRS. We reject integration between the stock markets, and between stock and money markets.

JEL Classification Numbers: G14

Keywords: integration, asset, market, discount, stock.

* Flood is Senior Economist, Research Department, International Monetary Fund. Rose is 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. This is a heavily revised version of a working paper with the same title. For comments, we thank workshop participants at the Carnegie-Rochester conference, Dartmouth, the Federal Reserve Board, Minnesota, Princeton, SMU, and Wisconsin as well as Jon Faust, Marvin Goodfriend, Rich Lyons, Mark Watson, Chris Sims, Ken West, Yangru Wu, and especially David Marshall and an anonymous referee. Rose thanks INSEAD and SMU for hospitality during the course of this research. The data set, sample output, and a current version of this paper are available at http://faculty.haas.berkeley.edu/arose.

1 Introduction

In this paper, we develop and apply a simple methodology to estimate the shadow risk-free rate or expected intertemporal marginal rate of substitution (hereafter “EMRS”). We do this for two reasons. First, it is of intrinsic interest. Second, when different series for the EMRS are estimated for different markets, comparing these estimates provides a natural test for integration between markets. Our method is novel in that it exploits information in asset-idiosyncratic shocks.

While the primary objective of this paper is methodological, we illustrate our technique by applying it to monthly and daily data covering firms from large American and Canadian stock exchanges. Our method delivers EMRS estimates with precision and striking volatility. Estimates from different markets can be distinguished from each other and from the Treasury bill equivalent.

Section 2 motivates our measurement by providing a number of macroeconomic applications. We then present our methodology; implementation details are discussed in the following section. Our empirical results are presented in section 5, while the paper ends with a brief conclusion.

2 Why Should Macroeconomists Care About Asset Market Integration?

We begin with a conventional intertemporal asset pricing condition:

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

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 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).

We adopt the standard definition of asset integration – *two assets are said to be integrated when the systemic and idiosyncratic risks in those assets are priced identically*. 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. Although many moments of m_{t+1} are involved in asset market integration, the object of interest to us in this study is $E_t m_{t+1}$ the time t expectation of the intertemporal marginal rate of substitution (EMRS). We 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 is simple to measure, subject to certain caveats discussed below. Third, cross-market differences in estimated values of $E_t m_{t+1}$ are statistically distinguishable, providing powerful evidence concerning market integration. We are testing only for first-moment equality when many additional moments are used in asset pricing; thus, ours it a test of a necessary condition for integration. If we reject equality of the first moment, we can reject integration, but failing to reject first-moment equality is consistent with (but does not imply) complete integration.

2.1 Motivation

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 there are a number of compelling reasons why most policy-oriented macroeconomists should be interested in asset market integration.

When macroeconomic modeling was based on descriptive structure a generation ago, market integration was not very relevant to macro. Modern macroeconomic models, however, are usually built on the assumption that agents maximize an intertemporal utility function in a stochastic setting (e.g., King and Rebelo, 2000, and Clarida, Gali and Gertler, 1999). In such macro models our equation (1) could be used to characterize bond holdings and might look like:

$$\frac{1}{(1+i_t)} = E_t\left(\frac{\rho u_c(c_{t+1}) * q_t}{u_c(c_t) * q_{t+1}}\right). \tag{2}$$

The corresponding equation for stock holding or to value firm revenues would be:

$$p_t = E_t\left(\frac{\rho u_c(c_{t+1}) * q_t}{u_c(c_t) * q_{t+1}} * x_{t+1}\right). \tag{3}$$

In equations (2) and (3), $0 < \rho < 1$ is a constant; $u(c_t)$ is a concave period flow of utility function with argument period t consumption and $u_c(c_t)$ is its partial derivative with respect to period t consumption. The price level at time t is q_t; the one-period interest rate is i_t; x_{t+1} and p_t are the aggregate counterparts to the symbols defined in equation (1).

Equation (2) is the basis for the modern IS curve; equation (3) is an efficiency condition for investment undertaken in period t. The point is that equations (2) and (3) both use the same stochastic discount rate, $m_{t+1} = \frac{\rho u_c(c_{t+1},..)*q_t}{u_c(c_t,..)*q_{t+1}}$, to evaluate short-term bonds (equation 2) and the payoffs from real productive assets (equation 3).

In our terminology, modern macro models assume markets pricing bonds and real assets to be integrated – both use the same stochastic discount rate. Asset market integration is a transmission channel of monetary policy in these models. In policy models, monetary authorities adjust $(1+i_t)$ in response to the current and expected future state of the economy, see e.g., Clarida, Gali and Gertler (1999). One transmission mechanism has m_{t+1} adjusting in equation (2) with the discount rate transmitting policy to other decisions, as in equation (3). Cross-market and cross-decision equality of intertemporal discount rates is a substantive assumption. Testing for asset market integration between stocks and money markets (which we do below) can be thought of as a specification test for the modern IS curve.

The IS curve is just one manifestation of asset market integration. Agents in many modern macro models use the same intertemporal discount rate to evaluate all intertemporal decisions – consuming vs. saving (the IS curve), enjoying leisure now vs. enjoying it later (the labor supply schedule), investing savings in various assets, and so on. Because these models typically use a single intertemporal discount rate, asset market integration is an essential ingredient.

In open economy work, e.g., Obstfeld and Rogoff (2000), capital market integration plays an even more substantive role. It is present in all the ways of closed-economy models, plus it plays a role in cross-currency and other international asset trading possibilities. Integration manifests itself perhaps most clearly in pricing foreign-currency bonds. The foreign- currency bond pricing condition is:

$$1 = E_t(m^*_{t+1} * \frac{s_{t+1}(1+i^*_t)}{s_t}), \tag{4}$$

where m^*_{t+1} is the discount rate used to price foreign currency assets, s_t is the domestic-currency price of foreign exchange, and i^*_t is the foreign-currency interest rate. Another way to write equation (4) is:

$$1 = COV(m^*_{t+1}, \frac{s_{t+1}(1+i^*_t)}{s_t}) + E_t m^*_{t+1} E_t(\frac{s_{t+1}(1+i^*_t)}{s_t}). \tag{5}$$

When pricing of foreign-currency securities is integrated with pricing of domestic-currency securities, $E_t m^*_{t+1} = E_t m_{t+1} = 1/(1+i)$ and equation (5) becomes

$$1 = COV(m_{t+1}, \frac{s_{t+1}(1+i^*_t)}{s_t}) + \frac{1+i^*_t}{1+i_t} E_t(\frac{s_{t+1}}{s_t}), \tag{6}$$

which is familiar from the work of Hodrick (1987). If, however, bond markets are not integrated internationally, equation (5) becomes:

$$1 = COV_t(m^*_{t+1}, \frac{s_{t+1}(1+i^*_t)}{s_t}) + \frac{(1+i^*_t)E_t m^*_{t+1}}{(1+i_t)E_t m_{t+1}} E_t(\frac{s_{t+1}}{s_t}). \tag{7}$$

Without integration, the size of $\theta_t = \frac{E_t m^*_{t+1}}{E_t m_{t+1}}$ is uncertain and varies through time. Tests of uncovered interest parity (UIP) set *both* $COV_t(m^*_{t+1}, \frac{s_{t+1}(1+i^*_t)}{s_t}) = 0$ and $\theta_t = 1$, Hodrick (1987). When such tests reject UIP they are usually interpreted in terms of $COV(m^*_{t+1}, \frac{s_{t+1}(1+i^*_t)}{s_t})$ being non-zero and correlated with interest rate regressors, but this need not be the whole story. For instance, Alvarez, Atkeson, and Kehoe (2002) interpret the UIP puzzle as well as excess volatility results using segmented asset markets, i.e., $\theta_t \neq 1$.

3 Methodology

Consider a standard decomposition of equation (1):

$$p^j_t = E_t(m_{t+1}x^j_{t+1}) = COV_t(m_{t+1}, x^j_{t+1}) + E_t(m_{t+1})E_t(x^j_{t+1}). \tag{8}$$

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

$$x^j_{t+1} = -[1/E_t(m_{t+1})]COV_t(m_{t+1}, x^j_{t+1}) + [1/E_t(m_{t+1})]p^j_t + \varepsilon^j_{t+1}, \text{ or}$$

$$x^j_{t+1} = \delta_t(p^j_t - COV_t(m_{t+1}, x^j_{t+1})) + \varepsilon^j_{t+1} \tag{9}$$

where $\varepsilon^j_{t+1} \equiv x^j_{t+1} - E_t(x^j_{t+1})$, 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 to us. In an integrated market without trading frictions, it is identical for all assets. Our work below is essentially concerned with exploiting and testing this restriction.[1]

3.1 Two Earlier Approaches to Parametric Estimation of Market Integration

It is typical in domestic finance to make equation (9) stationary by dividing the equation by $p_{j,t}$, resulting in:

$$x_{t+1}^{j} / p_{t}^{j} = \delta_t (1 - COV_t(m_{t+1}, x_{t+1}^{j} / p_{t}^{j})) + \varepsilon_{t+1}^{j} \ , \tag{10}$$

where ε_{t+1}^{j} is redefined appropriately. This normalization converts equation (9) 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 (10) is given economic content by adding two assumptions:

1) *Rational Expectations*: ε_{t+1}^{j} is assumed to be uncorrelated with information available at time t, and

2) Covariance Model: $COV_t(m_{t+1}, x_{t+1}^{j} / p_{t}^{j}) \ = \ \beta_0^j + \Sigma_i \beta_i^j f_{i,t}$

where: β_0^j is an asset-specific intercept, β_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 with 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 conditions we wish to *test* rather than assume.

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.

A second approach is provided by Flood and Rose (2003), who follow the spirit of Roll and Ross (1980) in testing for market integration based on cross-market equality of δ_t. Flood and Rose differ slightly from the Finance standard and normalize by p^j_{t-1} instead of p^j_t

$$x^j_{t+1} / p^j_{t-1} = \delta_t((p^j_t / p^j_{t-1}) - COV_t(m_{t+1}, x^j_{t+1} / p^j_{t-1})) + \varepsilon^j_{t+1} \tag{10'}$$

In this equation, the factor loadings $\{\beta\}$ (from the model used to proxy the covariance model) can still be estimated. But in addition, *cross-sectional* variation can be used to *estimate* $\{\delta\}$, the coefficients of interest that represent the EMRS and are time varying but common to all assets. Still, this approach – in common with the traditional approach that relies on (10) – requires correct specification of the covariance term, e.g., in the form of a factor model. If the latter is mis-specified, the $\{\delta\}$ estimates will also be incorrect.

The traditional finance approach allows one to estimate a covariance model and factor loadings (betas) with precision, at the expense of precluding precise estimation of the EMRS (since it is assumed to be the T-bill return or is an imprecisely estimated period constant). The second approach is oriented towards estimating the EMRS, but still requires specification of a covariance model. We now continue further down this road and develop a third approach which is even more geared towards estimating the EMRS; it does not require any explicit covariance model specification.

3.2 A New Strategy

The Finance standard, discussed above, and the approach followed by Flood and Rose (2003) both require that the discount rate pricing assets in equation (1) be identical for assets in an integrated market. This is a strong assumption. There are many discount rates that give identical prices for a given set of payoffs, see e.g. Cochrane (2001) chs. 4-6 or Hansen and Richard (1987). The approach we take presently recognizes that discount rates are not unique. Instead of relying on assumptions that makes them unique (e.g., a representative agent assumption), we impose instead an insurance assumption that gives unique equilibrium values of agents shadow risk free rates, $1/E_t m_{t+1}$.

Let $E_t \overline{m}_{t+1}$ be the expectation of the highest discount rate pricing assets in a portfolio hypothesized to be integrated. Then $1/E_t \overline{m}_{t+1}$ is the lowest shadow risk free rate in the portfolio. Since the shadow risk free rate is the price at which agents willingly bear idiosyncratic risks, the agent whose discount rate expectation is $E_t \overline{m}_{t+1}$ has a comparative advantage in bearing such risks. She/he is a natural "insurance agent" who can increase everyone's expected utility by holding the idiosyncratic risks faced by others in exchange for an insurance premium. Assuming the stock market reaches an equilibrium where such arbitrage is not possible, which we assume, means that either the insurance agent holds all the idiosyncratic risks in a corner solution or that the process of pricing and selling idiosyncratic-risk insurance drives different agents' shadow risk free rates into equality. In either case, the shadow price of holding idiosyncratic risk is unique in an integrated market.

In order to exploit the equality of shadow risk free rates in an integrated portfolio, we need to take a stand on systemic risks vs. idiosyncratic risks. Suppose we observe $\tilde{p}_t^j$, the "systematic price," 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:

$$\ln(p_t^j / p_{t-1}^j) = \alpha_0^j + \sum_{i=1}^{N} \alpha_i^j f_t^i + v_t^j, \qquad (11)$$

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

$$\tilde{p}_t^j = p_{t-1}^j \exp(\alpha_0^j + \sum_{i=1}^{N} \alpha_i^j f_t^i), \tag{12}$$

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

Normalizing by the systematic price delivers:

$$x_{t+1}^j / \tilde{p}_t^j = \delta_t[(p_t^j / \tilde{p}_t^j) - COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)] + \varepsilon_{t+1}^j. \tag{13}$$

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, we 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, $COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$, (13) can be decomposed as

$$x_{t+1}^j / \tilde{p}_t^j = \delta_t(p_t^j / \tilde{p}_t^j) - \delta_t COV_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j) + \varepsilon_{t+1}^j = \delta_t \exp(v_t^j) + u_{t+1}^j \tag{14}$$

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 which is unrelated to all information at time t by rational expectations (though there may be a common cross-sectional element; more on that below). 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 (14). 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.[2]

The basic idea of this study and the essential way it differs from previous work is that we use only asset-idiosyncratic shocks to identify and measure the expected marginal rate of substitution (or rather, its inverse), $\{\delta\}$. This stands typical Finance methodology – the first 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. Our test for asset market integration is simple; we check if the implied prices of carrying idiosyncratic risks – measures of the expected marginal rate of substitution – are equal across assets. If equality of the estimated EMRS cannot be rejected, then our test cannot reject cross-asset integration. If, however, we can reject equality then we also reject integration.

Our 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, we 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 our method and traditional methods is that we 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 our method is that it deals only with observable variables. The stochastic discount rate m_{t+1} is unobservable as are its moments. When we project asset price movements onto a set of aggregate factors, we are 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 our method is that it leaves a highly volatile regressor – idiosyncratic shocks – attached to $\delta_t \equiv 1/E_t(m_{t+1})$.

Our 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, we do not need to model the EMRS directly; we allow it to vary over time in a completely general fashion. Third, the technique requires only accessible and reliable data on asset prices and payoffs. 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. Finally, the technique is focused on estimating an intrinsically interesting object, the (inverse of the) expected marginal rate of substitution.

4 Empirical Implementation

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

$$\ln(p_t^j / p_{t-1}^j) = a_j + b_j * \ln(\overline{p}_t / \overline{p}_{t-1}) + c_j^1 f_t^1 + c_j^2 f_t^2 + c_j^3 f_t^3 + v_t^j \tag{15}$$

where a_j, $\{b_j\}$, and $\{c_j^k\}$ are fixed regression coefficients, $\overline{p}_t$ is the market-wide average price, $\{f_t\}$ are the three Fama-French (1996) factors, and v_t^j is the time-t asset idiosyncratic shock. This equation models the first-difference of the natural logarithm of a particular asset price as a linear function of the price growth of the market, augmented by the familiar Fama-French factors. The three Fama-French factors are: 1) the overall stock market return, less the treasury-bill rate, 2) the performance of small stocks relative to big stocks, and 3) the performance of "value" stocks relative to "growth" stocks. Further details and the data set itself are available at French's website.[3]

Estimates of equation (15) allow us to produce the fitted value of $\hat{p}_t^j$, which we define as:

$$\hat{p}_t^j \equiv p_{t-1}^j * \exp(\hat{a}_j + \hat{b}_j \ln(\overline{p}_t / \overline{p}_{t-1}) + \hat{c}_j^1 f_t^1 + \hat{c}_j^2 f_t^2 + \hat{c}_j^3 f_t^3) \tag{16}$$

We are not particularly attached to this specific model of $\hat{p}_t^j$. For instance, one could employ the Kalman filter to avoid using future data and allow for moving coefficient estimates. The regressors of equation (15) are meant to control for aggregate factors; one can add/change/or subtract some or all of them. We have assumed that the log first-difference of prices is linear in the market and factors; one could change the particular functional form assumption. We have used a time-series approach to estimating $\hat{p}_t^j$ but a cross-sectional approach seems possible. None of these assumptions are critical; they simply seem to work for us in practice.[4] But while this particular setup has delivered sensible results, we stress that one only needs *some* model for $\hat{p}_t^j$, not this precise one.

4.1 Estimation

We are fundamentally interested in estimating $\{\delta\}$ from the following model:

$$x_{t+1}^j / \hat{p}_t^j = \delta_t (p_t^j / \hat{p}_t^j) + u_{t+1}^j \tag{17}$$

for assets j=1,…,J, periods t=1,…,T. We 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. Equation (17) also includes a "generated regressor" 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. Estimation via the Generalized Method of Moments (GMM) allows us to handle both potential econometric issues, while not requiring independent and identically distributed disturbances. Accordingly, we use it below, comparing it with OLS (for simplicity). As instrumental variables for $\{p_t^j / \hat{p}_t^j\}$, we use its lag $\{p_t^j / \hat{p}_{t-1}^j\}$.[5]

4.2 The Data Sets

We 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. We use different frequencies both for intrinsic interest and to check the sensitivity of our techniques. Though these frequencies are standard in Finance, there is nothing special about them, and there is no obvious reason why our methodology could not be used at either higher or lower frequencies. We focus on stock markets, but again see no reason why bond or other markets could not be considered.

Our 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). We have adjusted for stock splits, and checked and corrected the data for errors. We only retain the 389 companies that have data for the full sample span that were traded on the New York Stock Exchange (NYSE).[6] Since we are interested in estimating and comparing implied EMRS across markets, we also include data from another market. We 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 we convert Canadian dollars into American using comparably timed exchange rates.

For the monthly data set, we 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, we have data for 246 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 us to compare our *estimates* of the expected risk-free rate with T-bill returns. We use data on T-bill returns downloaded from the Federal Reserve's website.[7]

5 Results

The focus of this paper is estimating the expected marginal rate of substitution. We begin with an illustration that relies on monthly data from 380 firms in the S&P 500 traded on the NYSE. We have 117 observations between March 1994 and November 2003, since we lose two observations at the beginning and one at the end of the sample, due to leads/lags.

The three graphs on the left of Figure 1 portray estimates of the EMRS from (17), denoted $\{\hat{\delta}_t\}$. These were estimated using two different techniques: GMM and OLS, using only the first 190 firms. The mean of $\{\hat{\delta}_t\}$ is plotted, along with a +/- 2 standard error confidence interval band. The OLS and GMM point estimates are highly correlated. The primary differences between the different estimates lie in the standard errors; both estimators deliver small standard errors, with the GMM standard errors being slightly smaller than those of OLS (but with more period to period volatility).[8] Indeed, we rarely find significant differences between the estimators below, and tend to rely on GMM below.[9]

Even though we estimate the expected MRS from only 190 firms, the results seem 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 117 periods is around 1.011, implying an annual MRS of around 1.14 ($=1.011^{12}$). While somewhat high compared to e.g., Treasury bill returns, this figure is plausible in magnitude. Further, the measures of EMRS are estimated with precision; the confidence intervals are 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 1.11. This considerable volatility in the expected EMRS mirrors our (2003) results as well as the famously high lower-bound of Hansen-Jagannathan (1991).[10]

5.1 Integration within the S&P 500

Do our results depend sensitively on the exact choice of firms chosen? An easy way to check is to estimate $\{\hat{\delta}_t\}$ using data from all 380 firms and look at the differences from the 190-firm 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 380 firms, and $\{\hat{\delta}_t\}$ estimated from only the last 190 firms. The differences are economically small; they average around .004 (for both estimators). They also have (relatively) large standard errors (averaging over .009), so that the individual differences are typically statistically insignificant. 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 NYSE.

The graphs 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, March 1994, is the same when estimated from all 380 firms and only from the last 190. 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).[11] Accordingly, we estimate the distribution for our critical values with a conventional bootstrap. With our bootstrapped results, we find the hypothesis of joint equality

$\{\hat{\delta}_t\}$ for all 117 observations cannot be rejected at any conventional significance level. That is, we cannot reject integration within the S&P 500. While this might only indicate a lack of statistical power in our techniques, we show later on that it is easy to reject equality of $\{\hat{\delta}_t\}$ across substantively different markets.

5.2 Estimates of the Expected Marginal Rate of Substitution and Treasury Bills

The hypothesis of equality of $\{\hat{\delta}_t\}$ cannot be rejected when the 380 firms are split into halves. 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; we 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 4.1% annualized), somewhat lower than but close to the sample average of $\{\hat{\delta}_t\}$.

But while the first moments of our 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 our estimated EMRS. The standard deviation of $\{\vec{\delta}_t\}$ (across time) is .001, which is smaller than that of $\{\hat{\delta}_t\}$ by an order of magnitude! 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.[12]

To summarize, our 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. This variation is grossly at odds with the smooth T-bill return. Unsurprisingly, we can reject equality between our estimates of EMRS and those of the T-bill.

5.3 Common Shocks

A referee has forcefully and correctly pointed out to us that aggregate shocks to the market would show up in our δ_t estimator. The most obvious way to handle this issue is to include a comprehensive set of time-specific intercepts in our estimation of (17). These intercepts represent the period-average of the aggregate shock, along with the period-average risk premium. Omitting intercepts can force these terms into $\{\hat{\delta}_t\}$

When we add a set of intercepts, we find that they are typically jointly insignificant.[13] Further, these estimates turn out in practice to be *negatively* correlated with the market return (though imprecisely estimated). Also, including them reduces the precision of $\{\hat{\delta}_t\}$ considerably. Thus, we exclude the intercepts in our default estimation equation. Still, the volatility of $\{\hat{\delta}_t\}$ and the high correlation between $\{\hat{\delta}_t\}$ and the market return, makes us wary of assuming away aggregate shocks. Hence we check below the robustness of our integration results to inclusion of this set of intercepts, and try to be conservative in interpreting our results.[14]

5.4 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 two different markets, the NYSE and the TSE. We use 380 firms to estimate the EMRS from the NYSE, and 150 firms to estimate the EMRS from the TSE. We estimate $\{\hat{\delta}_t\}$ in the same way as above, using GMM for 117 observations between 1994M3 and 2003M11. To facilitate comparison, we also graph the EMRS implicit in the short Treasury bill return.

For both equity markets, the average value of EMRS seems reasonable, being slightly over unity. These are again estimated with 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 over .04 for both the NYSE and the TSE, in stark contrast to the smooth T-bill return portrayed at the bottom of Figure 2.

Our results for both markets are consistent with our earlier findings. Different estimators (OLS and GMM) deliver economically similar results which are statistically close. There is considerable leptokurtosis. Bootstrapped tests for internal integration indicate that using e.g., half the NYSE firms delivers similar estimates of the EMRS to using all the firms. That is, we find little evidence against internal integration for the stock markets.

We consider these results to be reassuring, given the depth and liquidity of the advanced stock markets we consider. But they might simply indicate a lack of power in our statistical techniques; after all, they are simply not rejecting a necessary (but not sufficient) test for market integration. Accordingly, as a more stringent test, we also now test formally for integration across markets. This is also a subject of considerable intrinsic interest.

We begin comparing the estimated risk-free rate across markets with a series of scatter-plots in Figure 3. The top graph on the left of the Figure compares 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, we 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 .64. However they are not identical; the mean absolute difference between the $\{\hat{\delta}_t\}$ derived from the NYSE and the TSE is .03, and almost 20% are greater than .05.

It is straightforward to formally test the hypothesis that the estimated EMRS are equal across markets. We can use a simple Chow test to test for equality between $\{\hat{\delta}_t\}$ derived from the 380 NYSE firm with those which *also* use the 150 TSE stocks. When we do so, we find strong evidence against integration. The F-test for integration between the NYSE and the TSE is well over 10 when we use OLS and comparably high when we use GMM, strongly rejecting the null hypothesis of integration. This is true whether we include a comprehensive set of time-specific intercepts or not.

Succinctly, while our estimates of the expected MRS are similar for the different markets, they are significantly different in both the economic and statistical senses. That is, we are 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 our methodology is not lacking in statistical power.

5.5 Daily Results

Thus far, we have used a decade of monthly data. We now present results derived from the most recent available year of daily data, 2003. We use closing rates for the 246 days when both markets were open (but again lose 3 observations for lead/lags), converting Canadian dollar quotes from the TSE into American dollars using a comparable exchange rate. We consider the same pair of markets, noting in passing that both the American and Canadian markets close at 4:00pm daily in the same time zone.

Figure 4 is the daily analogue to the monthly estimates displayed in Figure 2. In particular, we plot the mean of the expected MRS for both markets, along with a +/- 2 standard error confidence interval (the T-bill equivalent is also plotted at the bottom of the figure). We use GMM as our estimator, though essentially nothing changes if we use OLS. We present estimate $\{\hat{\delta}_t\}$ using 440 NYSE firms and 220 TSE firms.

As with the monthly data set, the means of the series again seem reasonable; they are just above 1.001 for both markets. 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 both series of the EMRS 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 $\{\vec{\delta}_t\}$.[15]

When we 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, we are unable to reject the hypothesis of equality at any reasonable confidence interval. That is, we (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 3 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 .61 (for either OLS or GMM). The mean absolute difference between the series is around .006, and range to just 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 we test for equality between the estimates of $\{\hat{\delta}_t\}$ portrayed in Figures 3/4, we find the hypothesis rejected for the NYSE against the TSE (the F-test statistic is over 24), and again this conclusion does not change when we include a set of time-specific intercepts.

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

6 Conclusion

In this paper, we have developed a methodology for estimating the expected intertemporal marginal rate of substitution (EMRS). Our 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 us 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.

We apply our methodology to a decade of monthly data and a year of daily data, including data on stocks traded on the New York and Toronto Stock Exchanges. For both data sets, we find estimates of the expected marginal rate of substitution with reasonable means but time-series volatility which is high enough to be puzzling. We cannot generally reject the hypothesis that markets are internally integrated in the sense that different assets traded on a given market seem to have the same expected marginal rate of substitution. However, we find it easy to reject the hypothesis of equal EMRS across markets. This is both of direct interest, and indicates that our technique has statistical power. Still, we do not wish to claim too much for our technique. When we estimate period constants, they are typically jointly insignificant; accordingly, allowing for them does not substantially alter our results. Including them reduces the precision of the estimates of EMRS substantially, but ruling them out *a priori* seems implausible.

There are many possible ways to extend our work. One could add a covariance model to equation (14). A well-specified covariance model should result in more efficient estimates of the EMRS. One could sort stocks into portfolios in some systematic way (e.g., size, industry, or beta), and use portfolios instead of individual equities. Our analysis could be extended to other assets, such as long bonds or commodities. We could imagine more extensive tests for internal and cross-market integration. More or different factors could be added to the first stage regression, equation (15). While our 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 we have been able to reject the hypothesis of integration in the sense of equal expected marginal rates of substitution across markets, we have not explained the reasons for this finding of apparent market segmentation. If our 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.

Alvarez, Fernando, Andrew Atkeson, and Patrick Kehoe, 2002, "Money, Interest Rates, and Exchange Rates in Endogenously Segmented Markets" *Journal of Political Economy* 110-1, 73-112.

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.

Clarida, Richard, Jordi Gali, and Mark Gertler, 1999, "The Science of Monetary Policy: A New Keynesian Perspective" *Journal of Economic Literature* 37-2, 1661-1707.

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.

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

Hodrick, Robert , 1987, *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*, Harwood, New York.

King, Robert and Sergio Rebelo , 2000, “Resuscitating Real Business Cycles” in *Handbook of Macroeconomics* , John Taylor and Michael Woodford eds., North-Holland, 927-1007.

Obstfeld, Maurice and Kenneth Rogoff, 2000, “New Directions for Stochastic Open Economy Models” *Journal of International Economics* 50, 117-153.

Roll, Richard, and Stephen A. Ross, 1980, “An Empirical Investigation of the Arbitrage Pricing Theory” *Journal of Finance* 35-5, 1073-1103.

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

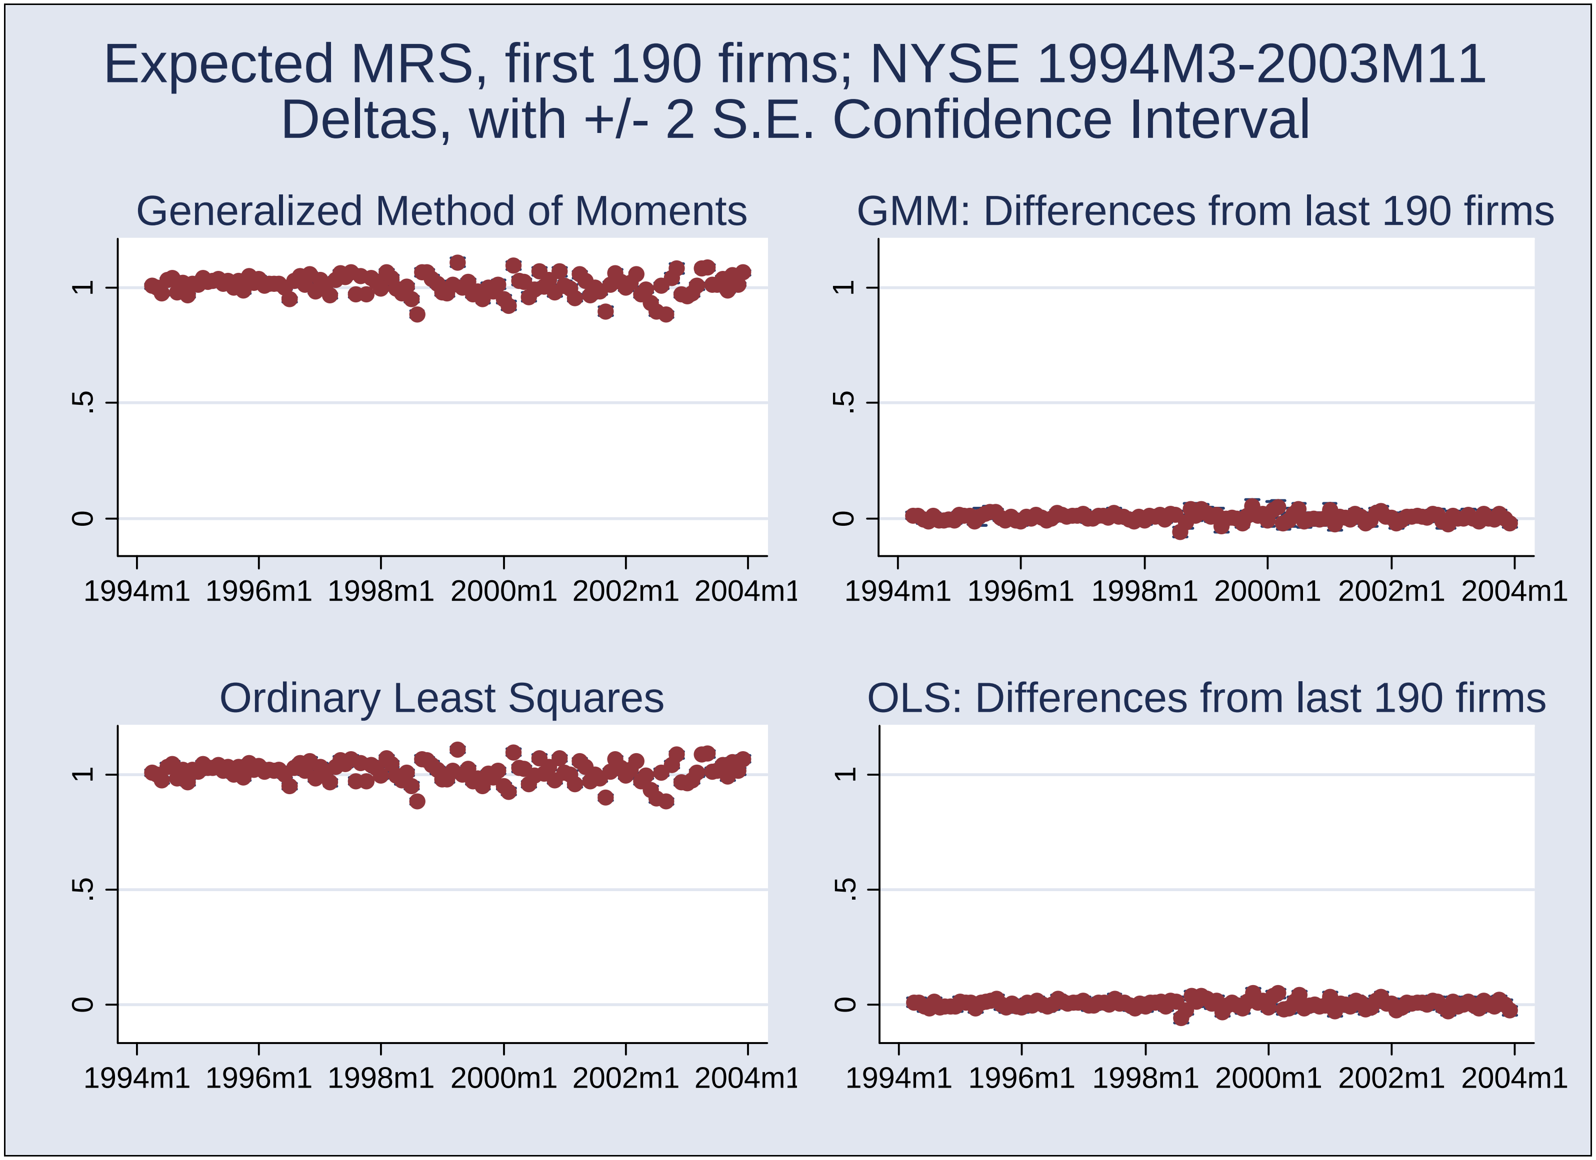

Figure 1: Estimated Expected MRS, 380 NYSE firms, 1994M3-2003M11: Different Estimators

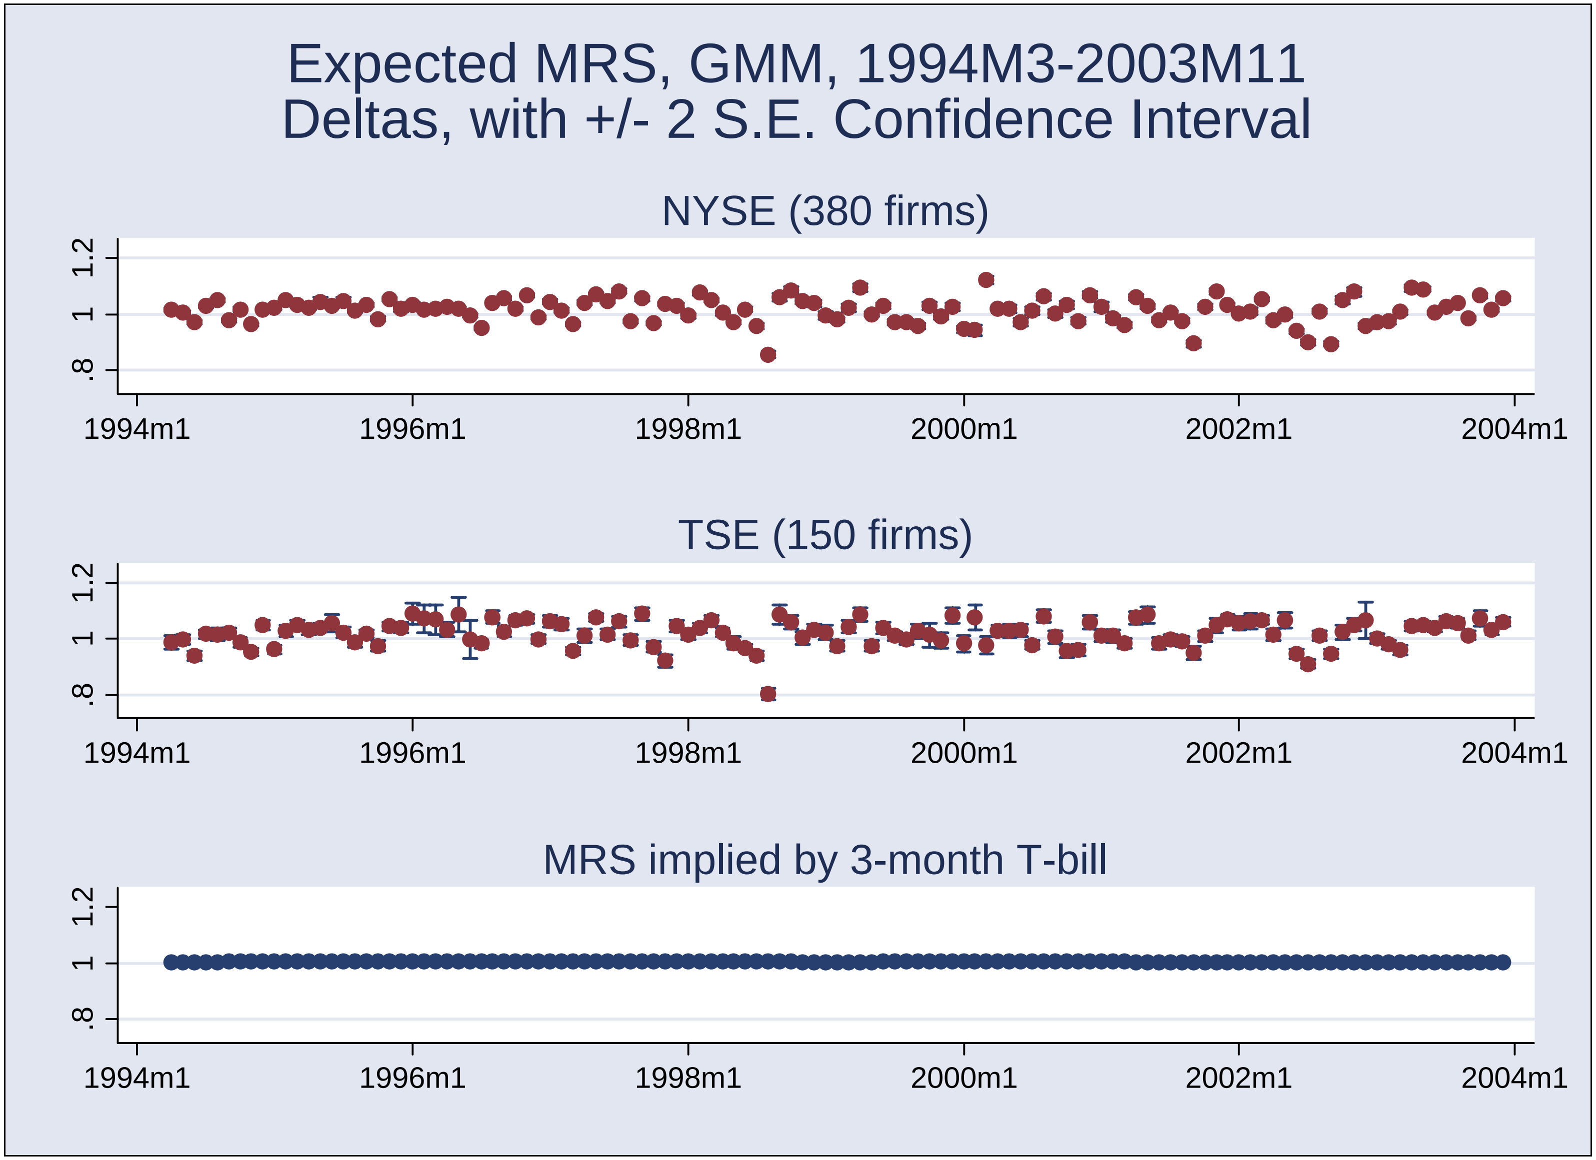

Figure 2: Estimates of Expected Marginal Rate of Substitution, 1994M3-2003M11: Different Markets

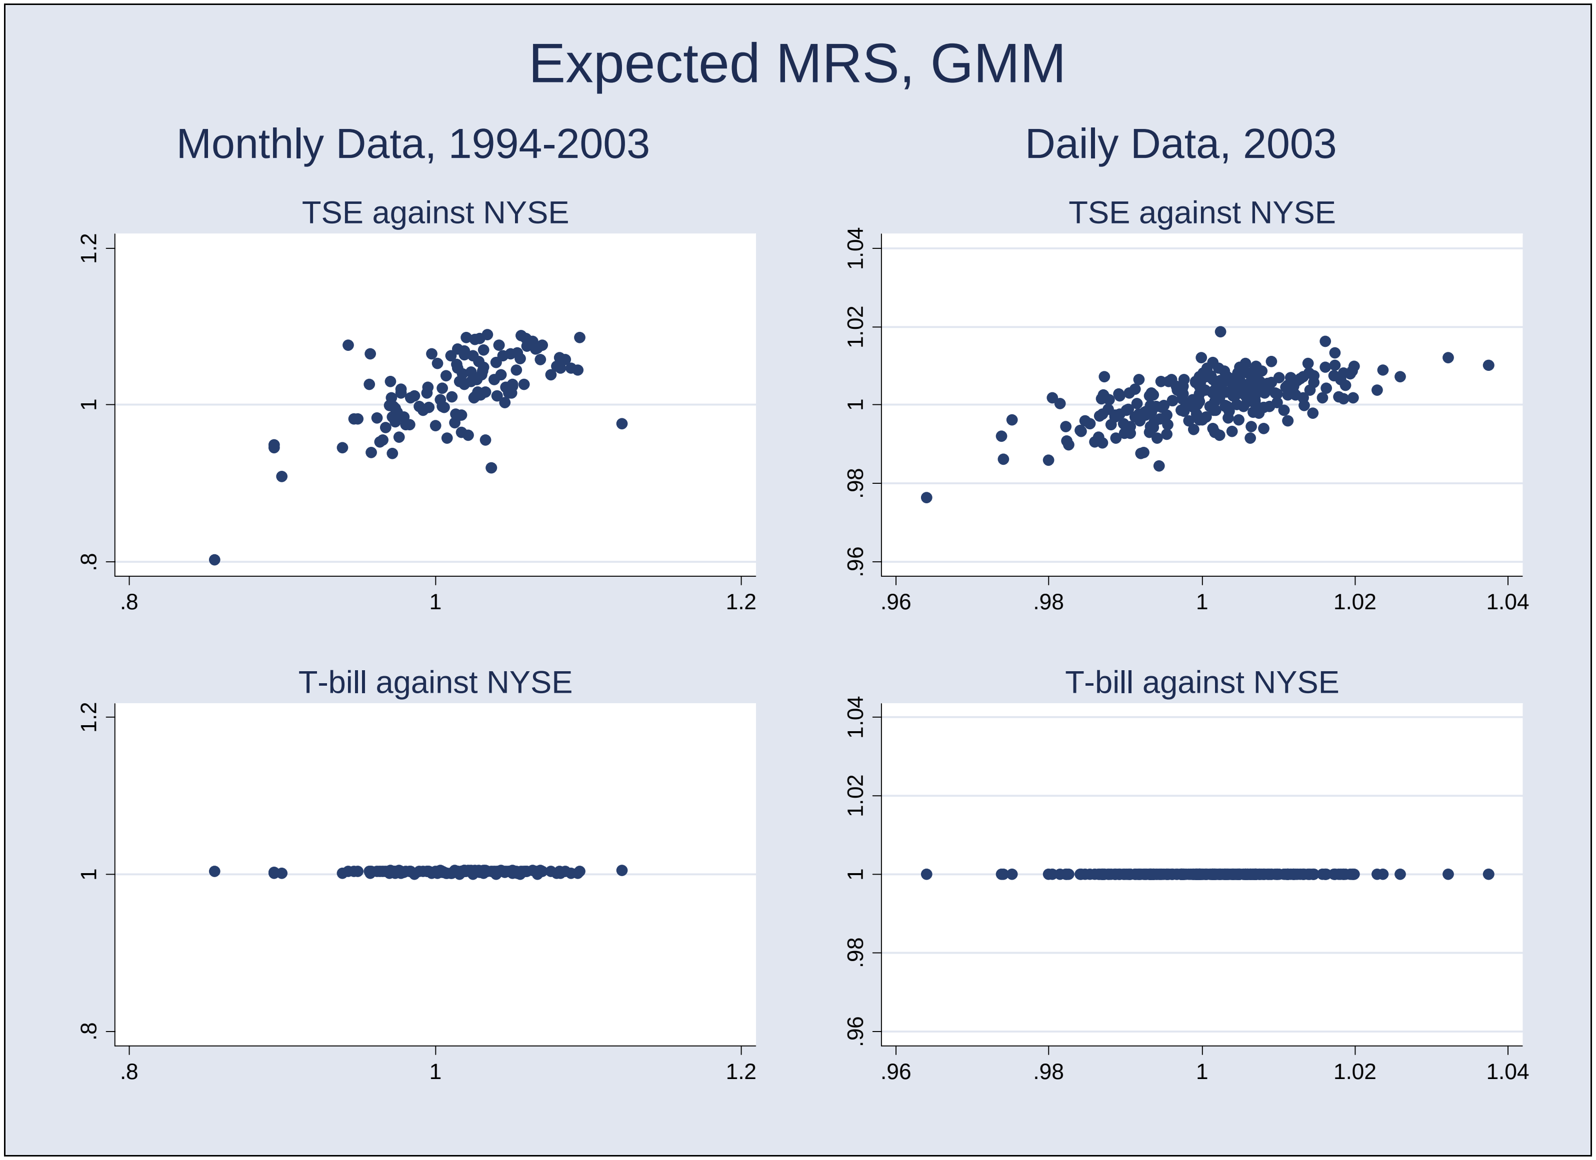

Figure 3: Scatter-plots of Estimated Expected MRS across Markets

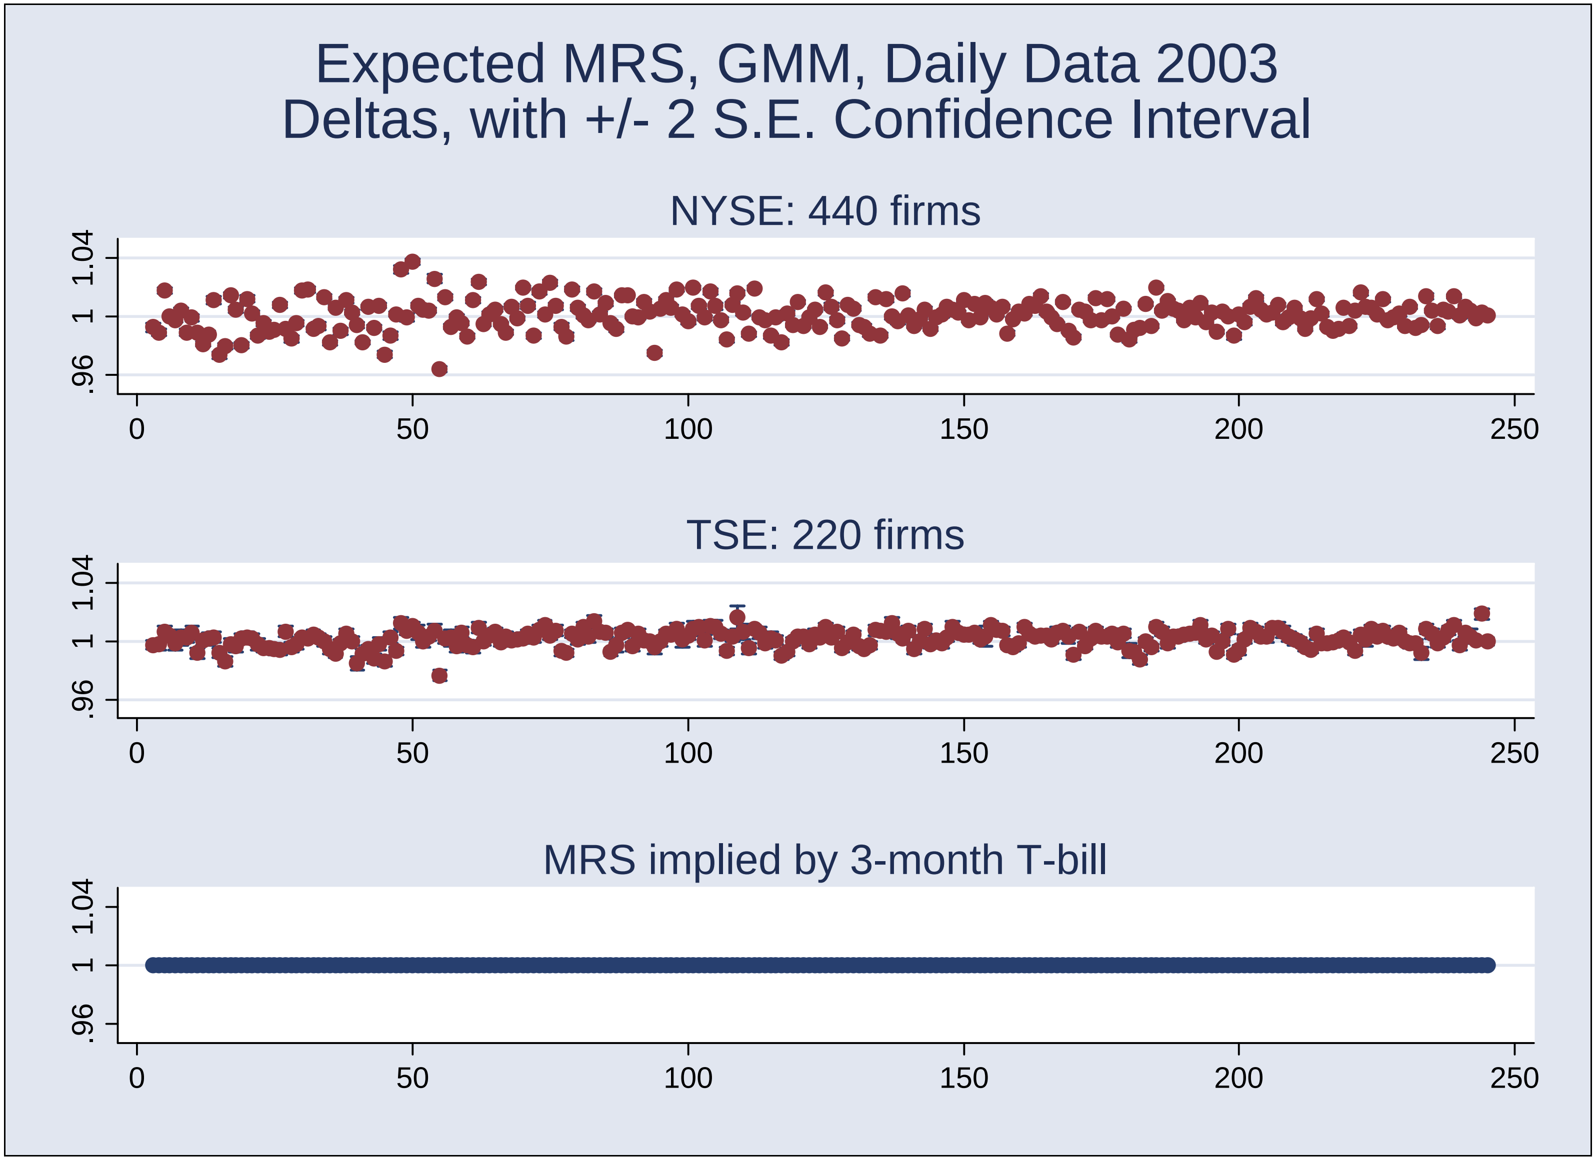

Figure 4: Daily Estimates of Expected Marginal Rate of Substitution, 2003

Figure Legends

1. Figure 1: Estimated Expected MRS, 380 NYSE firms, 1994M3-2003M11: Different Estimators
2. Figure 2: Estimates of Expected Marginal Rate of Substitution, 1994M3-2003M11: Different Markets
3. Figure 3: Scatter-plots of Estimated Expected MRS across Markets
4. Figure 4: Daily Estimates of Expected Marginal Rate of Substitution, 2003

Endnotes

[1] The assumption of representative agents is sufficient to deliver uniqueness of the EMRS and is common in important areas of macroeconomics, finance and international finance. However, it is not necessary; uniqueness obtains with other assumptions, such as a complete set of Arrow-Debreu markets or a set of free-entry insurance markets. With heterogeneity, the estimates we produce below can be unique, but need not equal the EMRSs of all agents. Accordingly, when we refer to "the" EMRS, we are implicitly assuming uniqueness. We discuss this below. Hansen and Richard (1987) provide a related discussion.

[2] Our estimates are consistent if we assume joint log-normality of exp(v) and $\mathrm{cov}_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$. This condition is not *necessary* for consistency of our estimator, but it is *sufficient* and understood easily. Under this assumption, the log idiosyncratic shocks (v) are assumed normal and uncorrelated with the log aggregate shocks generating $\mathrm{cov}_t(m_{t+1}, x_{t+1}^j / \tilde{p}_t^j)$. In future work, we hope to weaken this condition in various ways, for instance by using only temporary idiosyncratic shocks. An alternative set of sufficient conditions for consistency is laid out in an unpublished appendix available on the internet.

[3] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

[4] The average R^2 from our 380 monthly NYSE estimates of equation (15) is .31, with a range of (.03, .61).

[5] We are not wedded to this particular set of instruments, but our (brief) experiments with others have lead to similar results.

[6] 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 may be worthy of more work.

[7] In particular, we use closing secondary bid prices on three-month treasury bills for both daily data sets.

[8] Hausman tests rarely provide indication of any significant difference between the two estimators.

[9] The exception is our bootstrapped tests for integration, where we tend to use OLS for computational simplicity.

[10] We have no implicit model of preferences or technology that could deliver such a volatile EMRS; our task in this paper is simply to attempt to measure it.

[11] This is a well-known phenomena; see, e.g., Campbell, Lo, and MacKinlay (1997).

[12] The F-test statistic for equality between the expected MRS and the T-bill return is over 80; under the null hypothesis of market integration, it has degrees of freedom (117, 40,000).

[13] For instance, the chi-squared test for setting all the intercepts to zero is insignificant at the .5 confidence level, when we use GMM and 380 NYSE stocks, assuming normality. Because of non-normality, we have also verified our key results with the bootstrap.

[14] We note in passing that there is an observational equivalence between common shocks and unobservable changes in EMRS. This is an identification problem similar to the difficulty of differentiating bubbles from expected unobserved regime switches, long noted by Flood and Garber, Hamilton and Whiteman, and others. Implicitly, we have assumed away the existence of aggregate shocks to achieve identification. This assumption makes us uncomfortable, despite our intercept tests.

[15] The F-test statistic for equality between the expected MRS and the t-bill return is over 150.