

**A New Approach to Asset Integration:
Methodology and Mystery
Robert P. Flood and Andrew K. Rose***

Draft: February 28, 2004

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

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

Abstract

This paper develops a simple new methodology to test financial market integration. Our technique is tightly based on a general intertemporal asset-pricing model, and relies on estimating and comparing expected discount rates across asset markets. Expected discount rates are allowed to vary freely over time, constrained only by the fact that they are equal across (risk-adjusted) assets. Assets are allowed to have very general risk characteristics, and are constrained only by a linear factor-model of covariances with the discount rate over short time periods. The technique is undemanding in terms of both data and estimation, and includes CAPM, consumption-based models, various ICAPM and other models as special cases. We provide a variety of domestic and international empirical illustrations of our technique, and find surprisingly little evidence of integration. While the S&P 500 market seems typically to be integrated, others are not, including: the NASDAQ, the Toronto Stock Exchange, and three different classes of American bonds. Further, there is little evidence of integration between these apparently deep frictionless financial markets.

JEL Classification Numbers: F32, G15

Keywords: discount, rate, intertemporal, financial, market, expected, price, stock, bond.

* 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. We thank Pedro Rodriguez and Rafael Romeu for assistance with the data and David Bowman, Jon Faust, Cam Harvey, Robert Hodrick, Jonathan Kearns, Rich Lyons, Matt Pritzker, Tony Richards, Mark Rubinstein, Ken Singleton, Antonio Spilimbergo, Richard Stanton, Lars Svensson, Janet Yellen, and seminar participants at the Federal Reserve Board, the International Monetary Fund, Reserve Bank of Australia for comments. This is a shortened version of a paper with the same title that is forthcoming in *Contemporary Issues in International Finance*. Rose thanks INSEAD, the IMF, and the RBA for hospitality during the course of this research. The data set, sample E-Views programs and output, and a current version of this paper are available at <http://faculty.haas.berkeley.edu/arose>.

1: Defining the Problem

What does *securities market integration* mean? We adopt the view that financial markets are integrated when assets are priced by the same stochastic discount rate. More precisely, we define security markets to be integrated if all assets priced on those markets satisfy the pricing condition:

$$p_t^j = E_t(d_{t+1}x_{t+1}^j) \quad (1)$$

where: p_t^j is the price at time t of asset j , $E_t(\cdot)$ is the expectations operator conditional on information available at t , d_{t+1} is the market discount rate for income accruing in period $t+1$ (also widely known as the intertemporal marginal rate of substitution, the growth of marginal utility, the zero-beta return, or a pricing kernel), and x_{t+1}^j is the income received at $t+1$ by owners of asset j at time t (the value of the asset plus any dividends or coupons).¹ The substantive point of our definition is that all assets in a market share the same discount rate. There is no asset-specific discount rate in an integrated market, and no market-specific discount rate in markets that are integrated with each other. We rely only on a completely standard and general intertemporal model of asset valuation.²

Our object of interest in this study is d_{t+1} , the discount rate. More precisely, we are concerned with estimates of the expected market discount rate, $E_t d_{t+1}$, for two reasons. First, learning more about discount rates is of intrinsic interest, and has driven much research (e.g., Hansen and Jagannathan, 1991, who focus on their second moments). The market discount rate is the unobservable *DNA* of intertemporal decisions; characterizing its distribution is a central

task of Economics and Finance. The discount rate ties pricing in a huge variety of asset markets to peoples' saving and investment decisions. While the discount rate itself is unobservable, we can use asset prices and payoffs to characterize aspects of its distribution.

Second, measures of the expected discount rate lead naturally to an intuitive test for integration; in this paper, we propose and implement such a simple test for the equality of $E_t d_{t+1}$ across sets of assets. The logic of our study is as follows: By definition markets are integrated when assets in those markets are priced by the same discount rate. If d_{t+1} is equal across markets then so too must be $E_t d_{t+1}$. We conduct internal and cross-market tests for equality of $E_t d_{t+1}$ estimates inferred from different asset portfolios.

2: Empirical Strategy

We key off the fact that in an integrated market, the discount rate prices all assets held by the marginal asset holder. Indeed what we *mean* by asset market integration is that the same discount rate prices all the assets. In other words, if we could extract d_{t+1} (or its expectation) independently from a number of different asset markets, *they should all be the same if those markets are integrated.*

Consider a generic identity related to (1):

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

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

$$x_{t+1}^j = -[1/E_t(d_{t+1})]COV_t(d_{t+1}, x_{t+1}^j) + [1/E_t(d_{t+1})]p_t^j + \mathbf{e}_{t+1}^j, \quad \text{or}$$

$$x_{t+1}^j = \mathbf{d}_t(p_t^j - COV_t(d_{t+1}, x_{t+1}^j)) + \mathbf{e}_{t+1}^j \quad (3)$$

where $\mathbf{d}_t \equiv 1/E_t(d_{t+1})$ and $\mathbf{e}_{t+1}^j \equiv x_{t+1}^j - E_t(x_{t+1}^j)$, a prediction error.

We then impose two mild restrictions:

- 1) *Rational Expectations*: \mathbf{e}_{t+1}^j is assumed to be white noise, uncorrelated with information available at time t, and
- 2) *Constant Asset-Specific Effects*: $COV_t(d_{t+1}, x_{t+1}^j) = \mathbf{b}_j^0 + \sum^i \mathbf{b}_j^i f_t^i$, for the relevant sample,

where: \mathbf{b}_j^0 is an asset-specific intercept, \mathbf{b}_j^i is a set of I asset-specific factor coefficients and f_t^i a vector of time-varying factors.

With our assumptions, equation (3) becomes a panel estimating equation. We exploit *cross-sectional* variation to estimate $\{E(\mathbf{d}_t)\}$, coefficients that are time varying but common to all assets. These discount rates are the focus of our study. We use *time-series* variation to estimate the asset-specific “fixed effects” and factor loadings $\{\mathbf{b}^0, \mathbf{b}^i\}$, coefficients that are constant across time. Intuitively, these coefficients are used to account for asset-specific systematic risk (the covariances). We treat them as nuisance coefficients, required only to clear the way to produce estimated discount rates.

Estimating (3) for a set of assets $j=1, \dots, J_0$ and then repeating the analysis for the same period of time with a different set of assets $j=1, \dots, J_1$ gives us two sets of estimates of $\{E(d)\}$, a sequence of estimated discount rates. These can be compared directly, using conventional

statistical techniques. In particular, estimated discount rates can be compared either one by one (using t-tests), or jointly (using a likelihood-ratio test). Under the null hypothesis of market integration, the two sets of $E(d)$ coefficients are equal.

Our assumptions are weak. It seems uncontroversial to assume that expectations are rational for financial markets, at least in the sense that pricing errors are not *ex ante* predictable. It also seems reasonable to assume that the firm-specific covariances (of payoffs with the discount rate) are either constant or depend on only a small number of factors; it is certainly standard practice (e.g., Fama and French, 1996). Further, we have to make the latter assumption only for short time periods.

Our methodology has a number of 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 rely heavily on a particular asset-pricing model (e.g., the CAPM used by Bekaert and Harvey, 1995), though standard models are completely consistent with our methodology. Third, we do not need to model the expected discount rate directly. The discount rate need not be determined uniquely, so long as the expectation of the discount rate is unique. Fourth, our strategy requires only two relatively mild assumptions; we need not assume e.g., complete markets or homogeneous investors, or that we can model “mimicking portfolios” well. Fifth, the technique requires only accessible and reliable data on asset prices, returns and time-varying factors (if the latter are employed); no other data is required (e.g., the “world” or “market” portfolio). Sixth, the methodology can be used at very high frequencies and at low frequencies as well (though the latter requires a set of reasonable factors). Seventh, the technique can be used to compare expected discount rates across many different classes of assets including domestic and foreign

stocks, bonds, and commodities. Next, the technique is easy to implement and can be applied with standard econometric packages; no specialized software is required. Finally, the technique is focused on an intrinsically interesting object, the estimated expected discount rate.

3: Relationship to the Literature

We consider the pricing of two assets to be *integrated* when the discount rate (d_{t+1}) used to price next period's payoff to one asset, is the same as the discount rate used to price the same-period payoff to the other asset.. This definition of asset integration holds across asset pairs, asset portfolios, and asset markets. Indeed, it provides our definition of an asset market, which is *a portfolio of assets priced by the same discount rate*. The discount rate accounts fully for aggregate risks, is stochastic across periods, and is fully consistent with all intertemporal models of asset pricing.

An example may help to fix ideas. Consider a representative-agent model of a macro-economy. Suppose that the agent holds a risky asset – say an equity share. The Euler equation characterizing the agent's holding of this asset is:

$$p_t^j = E_t \left(\frac{\mathbf{r}u'(c_{t+1})}{u'(c_t)} x_{t+1}^j \right) \quad (1')$$

where: $\mathbf{r} < 1$ is a constant, $u'(c_t)$ is the marginal utility of consumption at time t and the prices and payoffs are real. In this equation (1'), $d_{t+1} \equiv \frac{\mathbf{r}u'(c_{t+1})}{u'(c_t)}$. We refer to d_{t+1} as the discount rate because it discounts things from time t+1 to today, time t. The discount rate is not necessarily

the constant $r < 1$, but it could be. This discount rate prices risks occurring in x_{t+1}^j that covary with the discount rate.³

Equation (1') illustrates the crucial point that the discount rate has no *special provision* for asset j . In equation (1'), idiosyncratic risks are not priced; only aggregate risks are priced. This is a tautology, but a useful one. The only way a risk connected to asset j will be priced is to the extent that the risk is correlated with aggregate risk. If asset j contains any remaining risk, that risk is idiosyncratic and disappears upon aggregation. When the discount rate is the same across a portfolio of assets, then the asset-specific risks are shared by holders of those assets. When it is different across asset portfolios, then risks connected to those portfolios are not shared across portfolio holders.

Equation (1') is just an example. The stochastic variable we call d_{t+1} arises in any context where people buy something today that is expected to pay off tomorrow. For example, there need be no representative agent, and everyone need not hold the same portfolio. This paper is concerned with characterizing and testing aspects of the distribution of d_{t+1} , an unobservable variable. From observable data on asset prices and payouts, we infer properties of agent's beliefs about the distribution of the discount rate.

Before we recount the most important contributions in the literature relevant to our investigation, we emphasize that we break from the literature at the most fundamental level. The literature is nearly uniform in its concentration on the variance of the discount rate, and its covariance with asset payouts. We develop these moments but only as nuisance coefficients, we need to clear from our path in order to measure the estimated discount rate, $E_t d_{t+1}$. We do not assume d_{t+1} (and therefore $E_t d_{t+1}$) to be equal across all assets or across all portfolios of assets. Instead, we estimate $E_t d_{t+1}$ in panel regression models, and test the proposition of cross-portfolio

equality of $E_t d_{t+1}$. That is, we check for asset integration empirically. Are the $E_t d_{t+1}$ estimates produced from two asset portfolios significantly different from each other? If not, we cannot reject asset integration. If so, we can.

Most of the literature *assumes* $E_t d_{t+1}$ to be equal across assets because it is convenient to do so. Consider the generic asset-pricing equation as applied to a safe government security, sold at a price of \$1 and paying $\$(1+i_t)$ next period. This becomes:

$$1 = E_t(d_{t+1}(1+i_t))$$

Since we have assumed the payment $\$(1+i_t)$ to be risk free, it follows that $1/(1+i) = E_t d_{t+1}$. Of course this is useful for assets other than government-backed ones *only if the discount rate for those non-government assets is identical to the discount rate for government-backed assets*. We test for the equality of expected discount rates across classes of assets, rather than assume that equality.

Our ideas build especially on Hansen and Jagannathan (1991), Cochrane (2001), and Brandt, Cochrane and Santa-Clara (2002). The application is indirect, but the lineage is clear. All these authors use equations like equation (1') to construct bounds concerning the standard deviations of d and of asset prices. Where they concentrate on second moments we concentrate on first moments. Chabot (2000) independently uses an approach similar to ours to assess stock market integration in the nineteenth century, while Chen and Knez (1995) provide a related application.

We also pull ideas from mainstream Finance's asset-pricing econometrics as summarized by Cochrane (2002). This branch of inquiry was pioneered by Sharpe (1964) and refined

subsequently by many others, including Fama (1970, 1991); Cochrane (2002) provides an excellent survey. Our relation to this empirical literature can be seen from our payoff equation (3), which we repeat as:

$$x_{t+1}^j = \mathbf{d}_t (p_t^j - COV_t(d_{t+1}, x_{t+1}^j)) + \mathbf{e}_{t+1}^j \quad (3)$$

When it is assumed that $\mathbf{d}_t = 1 + i_t$, we obtain

$$x_{t+1}^j - (1 + i_t)p_t^j = -(1 + i_t)COV_t(d_{t+1}, x_{t+1}^j) + \mathbf{e}_{t+1}^j \quad (3')$$

The term $(1 + i_t) \text{cov}(d_{t+1}, x_{t+1}^j)$ is then modeled in this literature as a function (usually linear) of market-wide factors. The factors are market-wide, because it is assumed that idiosyncratic factors are not relevant to $(1 + i_t) \text{cov}(d_{t+1}, x_{t+1}^j)$.⁴

We differ from this Finance standard in three aspects. First, we do not assume $\mathbf{d}_t = 1 + i_t$. Instead, we estimate \mathbf{d}_t each period based on a portfolio of assets we maintain to be integrated; $\hat{\mathbf{d}}_t$ is an estimated or “shadow” risk-free return. We then test sets of $\hat{\mathbf{d}}_t$ from one portfolio against those obtained from another portfolio. If the assets in the two portfolios are priced by the same d , then the $\hat{\mathbf{d}}_t$ will not differ significantly.

Second and less important, we concentrate attention on $\hat{\mathbf{d}}_t$, not on estimates of factor loadings (regression coefficients) estimated in linear models of $(1 + i_t) \text{cov}(d_{t+1}, x_{t+1}^j)$. Because we concentrate on $\hat{\mathbf{d}}_t$, our focus is on the cross-sectional dimension of the panel, e.g., the number

of stocks, rather than the length of the time series. Our time series dimension is short by Finance standards; we use one or two months of daily data. Limiting our time dimension is intended to minimize specification errors resulting from time-variation in factor loadings (and factor), but it also limits also the estimation precision of time-constant parameters.

Third, our estimating equation (3) is non-linear while equation (3'), the Finance standard, is linear when $(1+i_t)\text{cov}(d_{t+1}, x_{t+1}^j)$ is linear. Thus, in our specification the term $d_t \text{cov}(d_{t+1}, x_{t+1}^j)$ – the compound value of bearing asset j 's risk from t to $t+1$ —is time varying because of time variation in the factors explaining $\text{cov}(d_{t+1}, x_{t+1}^j)$ and because of time variation in d_t . When constrained by $d_t=(1+i_t)$, variation in this term is limited because in practice monetary authorities smooth short-term interest rates.

The tradition in International Finance to which we owe our greatest debt is surveyed by Karolyi and Stulz (2002) and known as the world CAPM (WCAPM) literature; see, e.g., Solnik (1974), and the recent contributions by Edison and Warnock (2003) and Goetzmann, Li, and Rouwenhorst (2001). This is very close to our work in spirit, but more specialized in application. Recall that the Capital Asset Pricing Model is a theory of d_{t+1} . In Sharpe's CAPM, $d_{t+1} = a + bR_{t+1}^m$, where a and b are data-determined constants and R_{t+1}^m is the return on the domestic-market portfolio. Turning this model into its international version entails moving from R_{t+1}^m to R_{t+1}^w , where R_{t+1}^w is the return on the world portfolio. The idea is if CAPM, or a variant, is correct and asset markets are integrated internationally, then the process of world asset integration involves d 's moving toward $a + bR_{t+1}^w$ (or the multi-factor equivalent) from some domestic-economy initial position. We encompass single-factor WCAPM and its multi-factor

variants as special cases. Instead of our having to take a strong stand on a , b and the identity of R_{t+1}^m or R_{t+1}^w , we estimate d directly.

Finally, a few words on arbitrage and replication. Replication is a fundamental idea in finance; two identical cash flows should have the same price, thereby precluding arbitrage. To use our notation, two assets with identical x 's should have the same p 's. This of course assumes that the same discount rate is used; if the same discount rate were not applicable to two assets with identical cash flows, they would not have identical prices. Thus equality of discount rates is critical to the replication/no arbitrage pricing technique commonly used in finance.

Our methodology is completely consistent with that of replication; in an arbitrage-free world, if we examined a set of assets with identical x 's, not only would their prices be identical, but our technique would deliver identical expected discount rates. However, though our methodology does not preclude them, it doesn't need or require identical payoffs.

When/if we find different estimated discount rates, it does not necessarily imply a deviation from arbitrage since we do not rely on comparing identical cash flows. Thus, our methodology is really an extension of the concept of market integration, beyond the sphere where one can apply replication arguments. That is, this project is fundamentally about testing asset integration when one cannot readily apply a replication/no-arbitrage argument.

4: Implementation

We begin by estimating a model with firm-specific intercepts and a single time-varying factor. In practice, we divide through by lagged prices (and redefining residuals and coefficients appropriately):

$$x_{t+1}^j / p_{t-1}^j = \mathbf{d}_t((p_t^j / p_{t-1}^j) + \mathbf{b}_j^0 + \mathbf{b}_j^1 f_t) + \mathbf{e}_{t+1}^j \quad (4)$$

for assets $j=1, \dots, J$, periods $t=1, \dots, T$. That is, we allow $\{\mathbf{d}\}$ to vary period by period, while we use a “two-factor” model and let $\{\beta\} = \{\mathbf{b}^0, \mathbf{b}^1\}$ vary asset by asset. We normalize the data by lagged prices since we believe that $COV_t(d_{t+1}, x_{t+1}^j / p_{t-1}^j)$ can be modeled by a simple factor model with time-invariant coefficients more plausibly than $COV_t(d_{t+1}, x_{t+1}^j)$, and to ensure stationarity of all variables.⁵

Equation (4) can be estimated directly with non-linear least squares. The degree of non-linearity is not particularly high; conditional on $\{\mathbf{d}\}$ the problem is linear in $\{\beta\}$ and vice versa.⁶ We also use robust (heteroskedasticity and autocorrelation consistent “Newey West”) covariance estimators.⁷

A few words on our choice of the factor model are in order. Our model is more general than, and subsumes the static Capital Asset Pricing Model (CAPM) for two reasons. First, the CAPM models the discount rate as a linear function of the “market return” and thus delivers an asset-specific correlation that has a time-invariant correlation with “the market.” This implies a constant correlation of the discount rate with the asset, which would be picked up by our asset-specific intercepts $\{\mathbf{b}_j^0\}$. Second, the CAPM (in both single and multi-factor versions) estimates discount-rate covariances conditional on either $E(d) = 1/(1+i)$ or on some other maintained

model for $E(d)$, e.g., $E(d_{t+1}) = E\left(\frac{\mathbf{r}u'(c_{t+1})}{u'(c_t)}\right)$. We need not maintain any particular model of

$E(d)$; it remains a vector of unconstrained coefficients estimated period by period in our methodology.

We choose as our single time-varying factor the square of the market return, that is $[\ln(\text{Index})_t - \ln(\text{Index})_{t-1}]^2$ where Index is e.g., the S&P 500 index when we examine large American stocks (and e.g., the NASDAQ when we examine NASDAQ stocks, etc.).⁸ This seems a natural choice to us; it is a simple function of a relevant aggregate shock that is easily observable. While we think of this as consistent with the spirit of the Intertemporal Capital Asset Pricing Model (ICAPM), it does not seem to be a particularly important issue, and we discuss a few different factor models below.

We start with a moderately high frequency approach. Using daily data allows us to estimate the coefficients of interest $\{d\}$ without assuming that firm-specific coefficients are constant for implausibly long periods of time.⁹

5: A Detailed Illustration: Large American Stocks

Our empirical work begins with an examination of the integration of deep American equity markets. Large American stocks are traded on liquid markets, which we consider *a priori* to be integrated. We begin with daily data over a quiet two-month period, April-May 1999 (about a year before the end of the Clinton bull market).¹⁰ Two months gives us a span of over forty daily observations; this does not appear to stretch the credibility of our assumption of constant asset-specific effects excessively, while still allowing us to test financial market integration for an interesting span of data. We see no reason why higher-frequency data cannot be used.¹¹

Our data set is drawn from the “US Pricing” database provided by Thomson Analytics. We use closing rates for the first (in terms of ticker symbol) one hundred firms from the S&P 500 that did not go ex-dividend during the months in question.¹² The absence of dividend

payments allows us to set $x_{t+1}^j = p_{t+1}^j$; we choose a hundred firms since we split the data set in two to test for integration and think that fifty firm provides a reasonable cross-section.¹³

Our sample period consists of 43 days. Since we lose the first and last observations because of lags (p_{t-1}^j) and leads (p_{t+1}^j), we are left with a total of 4100 observations in our panel data set (100 firms x 41 days). Our data has been checked for transcription errors both visually and with random crosschecking.

We begin by using data from the first 50 firms to estimate discount rates (i.e., estimates of $\mathbf{d}_t \equiv [1/E_t(d_{t+1})]$). We graph our estimated deltas along with a plus/minus two standard error confidence interval in Figure 1.

The expected discount rates seem reasonable. The estimates of delta are close to unity (and are never significantly different at standard confidence levels), with relatively tight confidence intervals.¹⁴ It is interesting to note that they vary considerably over time, consistent with the thrust of Hansen and Jagannathan (1991). The hypothesis of constant delta is rejected at any reasonable significance level.¹⁵

We are less interested in $\{\beta\} = \{\mathbf{b}^0, \mathbf{b}^1\}$, coefficients that are nuisances for our purposes. Still, we note in passing that the $\{\mathbf{b}^0\}$ coefficients are negative while $\{\mathbf{b}^1\}$ estimates are positive; both sets of coefficients are jointly significant but individually insignificant.

What we are really interested in is using our estimates to test for market integration. One easy way to do this is to compare the delta estimates from the 50 firms graphed in figure 1 with those from a different set of S&P firms (but the same time period). Figure 2 portrays the expected discount rates from Figure 1 along with those from another (mutually exclusive) set of 50 S&P firms, again from April-May 1999; we also include the plus/minus two standard error confidence interval (the latter from the second set of firms).

Clearly the two sets of expected discount rates are close when examined day by day; the differences are individually insignificant at conventional levels. It is also simple to test for joint equality of the two sets of deltas. The log-likelihood of our estimate of (4) from the first set of fifty firms is 4192, while that from the second set of fifty firms is 4333. When we pool across all hundred firms and estimate a single set of deltas, the log-likelihood is 8505. Under the null hypothesis of market integration, the deltas should be equal. With normally distributed residuals, twice the difference in the log-likelihoods is distributed as χ^2 under the null with T degrees of freedom. Since $2((4192+4333) - 8505) = 40$ sits virtually at the median of the $\chi^2(41)$ distribution (the p-value is .49), the null hypothesis that the S&P 500 stock market is integrated cannot be rejected during this period of time. All this bolsters our confidence in the methodology. This is especially true since the excess kurtosis commonly observed in daily returns probably makes our critical values (which rely on normality) quite conservative. Further, it is standard in finance to examine portfolios of assets which have considerably less noise than individual assets, making our test even more demanding.

In passing, we note that the point estimates of our expected discount rates do not seem to depend very sensitively on the exact factor model, i.e., the parameterization of $COV_t(d_{t+1}, x_{t+1}^j / p_{t-1}^j)$. We have re-estimated our model without our time-varying factor (i.e., setting $\{\mathbf{b}^1 = 0\}$) and also without our intercepts (i.e., setting $\{\mathbf{b}^0 = 0\}$). Figure 3 portrays the expected discount rates for our default specification and both alternatives, generated from all one hundred S&P firms. While the confidence intervals change across specification, the point estimates do not seem to vary either dramatically or systematically. We stick with our default two-factor model since it is both more general than the alternatives, and delivers the widest

confidence intervals, making it more difficult for us to reject the hypothesis of market integration.

We have redone our analysis for two other two-month periods in 1999: July-August and also October-November. We have also repeated the analysis for the same three two-month periods in 2002; all results are presented in Table 1. Two of the six sample periods seem to present only marginal evidence in favor of the null hypothesis of market integration, while the results for October-November 2002 are inconsistent with the null at all conventional significance levels. While it appears that the hypothesis that the market for S&P 500 stocks is integrated can be rejected for at least one of our six sample periods, we try not to take these results too literally, for two reasons. First, the residuals are unlikely to be normally distributed because of leptokurtosis. Second, we are using individual stocks rather than the portfolios that most finance economists use. Ongoing research indicates that bootstrapped results from portfolios of stocks leads one to conclude that the hypothesis of integration of the S&P cannot be rejected.

Figure 5 portrays expected discount rates for the six different sample periods we examine, all estimated from 100 S&P firms. The time-series volatility of delta is striking and wholly consistent with the spirit of Hansen-Jagannathan (1991).

Using different factor models (that is, different models of for firm covariances $COV_t(d_{t+1}, x_{t+1}^j / p_{t-1}^j)$) does not seem to change our results. Appendix 1 has four sets of analogous results to those in Table 1, derived using four different factor models. Our default model (of Table 1) is a two-factor model with a set of firm-specific intercepts $\{\beta^0\}$ and firm-specific slopes for the square of the market return $\{\beta^1\}$. We have also examined: a) a one factor model with just $\{\beta^0\}$; b) another one factor model with just $\{\beta^1\}$; c) a two-factor model with $\{\beta^0\}$ and firm-specific slopes for the *level* of the market return, and d) a three factor model with

intercepts and firm-specific slopes for both the level and square of the market return. No conclusions of substance are much affected by the precise choice of factor model. We find this robustness reassuring. Still, there is no reason why other factors (e.g., firm size) cannot be used.¹⁶

6: Further Examples

Having illustrated our basic methodology in some depth, we now provide a series of other applications of our technique. These are intended to illustrate the power and breadth of its potential use.

Is the S&P Integrated with the NASDAQ?

Most large American stocks are traded on the floor of the New York Stock Exchange; many smaller stocks are traded electronically on the NASDAQ. It is interesting to compare S&P 500 stocks to the NASDAQ; we test whether S&P 500 equities are integrated with stocks traded on the NASDAQ. To do this, we obtain NASDAQ data that are similar in style to those from the S&P 500 (they are closing rates, also obtained from the US Pricing database of Thomson Analytics). Again, we use the first 100 (in terms of ticker symbol) firms that did not go ex-dividend during the samples we examine.

Are American and Canadian Stock Markets Integrated?

Canada and the United States are similar economies in a number of respects, and there are few barriers to flows of goods, services, or capital between them. We find it interesting to ask if the Canadian stock market is integrated with its American counterpart. To pursue this, we

use closing prices on equities from the Toronto Stock Exchange (TSE) 300, again obtained from Thomson Analytics. We convert these prices in Canadian dollars into American dollars by using closing foreign exchange rates obtained from the Bank of Canada.¹⁷

We have been unable to obtain a full set of 100 Canadian stocks for all 6 periods. In fact we only have data for 63 firms for both 1999m4/5, and 1999m7/8, and 66 firms for 1999m10/11. Since the set of firms rises to 81, 82, and 83 for the three 2002 samples, we choose not to use 1999 data for this example.

Are American Stock and Bond Markets Integrated?

Our methodology can be readily applied to financial markets beyond equity markets, and the most obvious candidates are bond markets. In particular, it is interesting to ask if stock and bond markets are integrated in the sense that the expected discount rates are similar.

We begin by using closing “clean market” prices on AAA class US government, and corporate bonds, taken from DataStream. As always, we choose those that did not go ex-dividend/coupon during the sample. (Since we do not have a complete set of 100 AAA bonds for 1999 – we have only 92/92/53 bonds for the three samples – we again only use 2002 data for the bond examples.) We also use similar data on A+ bonds (from corporations, financial institutions or governments), and BB+/BB bonds (often referred to as “junk bonds”). It should be noted that the quality of the bond price quotes is questionable since some bond markets are illiquid.

Evidence

We present four different sorts of evidence on asset integration. First, we present likelihood-ratio tests of the hypothesis of internal market integration, analogous to those of Table 1. These are presented in Tables 2, 4, 6, 8, and 10 for NASDAQ stocks, TSE stocks, AAA bonds, A+ bonds, and junk bonds respectively. Second, we present likelihood-ratio tests of integration across asset markets. Thus in Table 3 we test the integration of NASDAQ and S&P stocks; tables 5, 7, 9, and 11 are analogues that compare the S&P to the five other assets.

The evidence presented in tables 1 through 11 is purely statistical in the sense that we can reject the null hypothesis of asset integration, or not. But rejection of the hypothesis of market integration may occur for different economic reasons. There is a big difference between two sets of deltas that are similar in magnitude but sufficiently far apart to reject the null of integration under the assumption of normality, and two sets of deltas that are wildly different. Since we are interested in interpreting our results, we present time-series plots of the expected discount rates derived from our six different assets (along with appropriate confidence intervals) in figures 6 through 8. Scatter-plots of expected discount rates are portrayed in figures 9 through 11.

Finally, tables 12 through 17 present quantitative economic measures of the degree of market integration (DMI) using a number of different metrics for the “closeness” of the expected discount rates. There are a number of ways of measuring the closeness of expected discount rates. We have not yet settled on a single summary statistic, but provide a few different measures of the degree of market integration.

The first measure we choose is the mean absolute difference between the expected discount rates. Thus for any two asset classes p and q , we compute $(1 / T) \sum_t | \mathbf{d}_t^p - \mathbf{d}_t^q |$. The second, closely related measure is based on the Grubel-Lloyd measure of intra-industry trade

and is $(1/T)\sum_t 2 | \mathbf{d}_t^p - \mathbf{d}_t^q | / (\mathbf{d}_t^p + \mathbf{d}_t^q)$. For both measures, smaller values indicate closer integration; a value of zero indicates perfect integration. The results for the three different samples in 2002 are available in Tables 12-14. To interpret these numbers, notice that the entry in Table 12 relating S&P to TSE (below the diagonal) is .04. This corresponds to a 4% *daily* interest rate differential, which is large compared to annualized interest rates.

We also use two measures borrowed from Brandt, Cochrane and Santa-Clara (2002), who examine international risk sharing. In particular, we compute

$1 - [\mathbf{s}^2(\ln \mathbf{d}_t^p - \ln \mathbf{d}_t^q) / (\mathbf{s}^2(\ln \mathbf{d}_t^p) + \mathbf{s}^2(\ln \mathbf{d}_t^q))]$ and $1 - [\mathbf{s}^2(\mathbf{d}_t^p - \mathbf{d}_t^q) / (\mathbf{s}^2(\mathbf{d}_t^p) + \mathbf{s}^2(\mathbf{d}_t^q))]$. Results for the 2002 samples are available in Tables 15-17. These measures have properties similar to those of correlations; the measures are unity when the expected discount rates are identical, zero when the expected discount rates are uncorrelated, and equal to minus one if e.g.,

$$\ln \mathbf{d}_t^p = -\ln \mathbf{d}_t^q .$$

Finally we note that all four of our measures ignore sampling imprecision; that is we do not provide confidence intervals for any of the measures.

Results

It is easy to summarize the results that we find beyond the S&P. Without taking critical values too literally (because of leptokurtosis and the fact that we examine assets rather than portfolios), the null hypothesis of integration inside bond markets is rejected. Sometimes the rejections are quite staggering in the sense of likelihood ratio statistics that exceed one thousand (when consistency with the null implies figures below sixty). The evidence for integration across asset classes uniformly rejects the null hypothesis, usually in an overwhelming fashion.¹⁸ A different interpretation is that we have found the cross-sectional analogue to Hansen-

Jagannathan (1991); Hansen and Jagannathan find evidence of time-series dispersion in discount rates whereas we find evidence of cross-sectional dispersion in expected discount rates.

The expected discount rates portrayed in figures 6 through 8 indicate that deltas for different asset classes are usually volatile on a time series basis. But they often differ across asset classes. We can see this clearly by focusing on April-May 2002, a sample of special interest since it is only for this sample that the hypothesis of market integration cannot be rejected for the S&P (though it can for the five other assets). There are a number of days when the expected discount rates of different asset classes are quite different. For instance, day 28 (May 8, 2002) in delta estimates of .78 (se of .06) from the S&P, 1.33 (.26) from the NASDAQ, .99 (.29) from AAA bonds, 1.09 (.24) from A+ bonds, and .55 (.07) from junk bonds. These expected discount rates seem far apart in both statistical and economic senses.

The finding of economically significant difference in deltas is corroborated in Tables 12-17, which tabulate four measures of the degree of market integration (DMI) for the different assets and sample periods. Our measures of DMI vary dramatically from period to period, and there are no obvious groupings of assets that are consistently tightly integrated.

Of course, we have only examined six financial asset classes, and only for six periods of time; our results may not be general. Still, we emphasize that the assets we examine are traded on apparently deep markets with few important frictions. We find the nearly uniform lack of evidence of integration both reassuring (since it implies that our technique is powerful), and puzzling (since we do not understand why these markets are not integrated). Much food for thought!

7: Looking Backward and Looking Forward

In this paper, we hope to have made two contributions. Most importantly, we presented a methodology for testing asset integration. Rather than assume that the “risk-free” rate from short-term government treasury bills is the appropriate discount rate for stock and bond markets, we test this assumption by comparing estimated discount rate derived directly from asset price data. Our technique can be easily implemented using standard data and econometric techniques, while being tightly based upon a standard general theoretical framework. It has demonstrable empirical power to estimate expected discount rates – the estimated inverse of the marginal rate of intertemporal substitution – with precision. In fact, we have been able to reject the hypothesis of equal expected discount rates (and thus market integration) for a number of different financial markets. The assets we consider include the S&P 500, the NASDAQ, the TSE, and a number of American bond markets; none have any obvious substantial trade frictions. We are thus somewhat perplexed by our second contribution, a general lack of evidence supporting asset integration across markets.

We have chosen to interpret our finding as indicating a lack of integration; but our tests are conditional upon a model of asset covariances. While we find the strength and robustness of our findings with respect to the exact covariance model reassuring, our implicit model makes us cautious in our conclusions. While our technique has a number of strengths, it clearly does not resolve the issue of asset integration. While we can reject the hypothesis of asset integration for certain interesting samples, our technique does not shed light on the economic cause(s) for these rejections. Do asset markets seem to be segmented because of artificial barriers (e.g., capital controls or taxes), asymmetric information, or some other phenomena? This remains an interesting topic for future work.

A number of easy extensions occur to us immediately. It would be interesting to pursue both higher- and lower-frequency approaches, to see if there is more evidence of market integration within say individual days, or across decades, and whether market integration is growing. It would also be interesting to examine asset integration before and after periods of extreme financial turbulence. We would also like to group assets into portfolios (as is the norm in the finance profession), and check for kurtosis explicitly. We plan to pursue these topics in future work.

References

- Adam, Klaus, Tullio Jappelli, Annamaria Menichini, Mario Padula, and Marco Pagano (2002) “Analyse, Compare, and Apply Alternative Indicators and Monitoring Methodologies to Measure the Evolution of Capital Market Integration in the European Union” University of Salerno manuscript.
- Bekaert, Geert and Campbell R. Harvey (1995) “Time-Varying World Market Integration” *Journal of Finance* 50-2, 403-444.
- Brandt, Michael W., John H. Cochrane and Pedro Santa-Clara (2001) “International Risk-Sharing is Better than you Think” NBER Working Paper #8404.
- Chabot, Benjamin (200) “A Single Market? The Stock Exchanges of the United States and London: 1866-1885” University of Michigan working paper.
- Chen, Zhiwu and Peter J. Knez (1995) “Measurement of Market Integration and Arbitrage” *Review of Financial Studies* 8-2, 287-325.
- Cochrane, John H. (2001) *Asset Pricing*, Princeton University Press.
- Edison, Hali and Frank Warnock (2003) “U.S. Investors’ Emerging Market Equity Portfolios: A Security-Level Analysis” IMF manuscript
- Fama, Eugene (1970) “Efficient Capital Markets: A Review of Theory and Empirical Work” *Journal of Finance* 25, 383-417.
- Fama, Eugene (1991) “Efficient Markets: II” *Journal of Finance* 46, 1585-1618.
- Fama, Eugene and Kenneth R. French (1996) “Multifactor Explanations of Asset Pricing Anomalies” *Journal of Finance* 51-1, 55-84.
- Goetzmann, William N., Lingfeng Li, and K. Geert Rouwenhorst (2001) “Long-Term Global Market Correlations” NBER WP #8612.
- 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.
- Karolyi, G.A., and R. Stulz (2002), “Are Financial Assets Priced Globally or Locally?” forthcoming *Handbook of the Economics of Finance* (George Constantinides, Milton Harris and Rene M. Stulz, eds,) North Holland.
- Solnik, Bruno (1974) “An Equilibrium Model of the International Capital Market” *Journal of Economic Theory* 8, 500-24.

Log Likelihoods	April-May 1999	July-Aug. 1999	Oct.-Nov. 1999
First 50 Firms	4192	4819	4191
Second 50 Firms	4333	4899	4358
All 100 Firms	8505	9687	8526
Test Statistic (df) P-value	40 (41) .49	62 (42) .98	46 (41) .73
	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
First 50 Firms	5091	4108	3794
Second 50 Firms	5130	4326	4072
All 100 Firms	10197	8403	7825
Test Statistic (df) P-value	48 (43) .72	62 (43) .97	82 (42) 1.00

Table 1: Tests of Market Integration inside the S&P 500, Two-Factor Model

Log Likelihoods	April-May 1999	July-Aug. 1999	Oct.-Nov. 1999
First 50 Firms	3343	3646	2048
Second 50 Firms	3354	3808	3415
All 100 Firms	6676	7424	4999
Test Statistic (df) P-value	42 (41) .57	60 (42) .96	928 (41) 1.00
	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
First 50 Firms	3747	3427	3023
Second 50 Firms	4169	3085	3045
All 100 Firms	7848	6457	6032
Test Statistic (df) P-value	136 (43) 1.00	110 (43) 1.00	72 (42) .997

Table 2: Tests of Market Integration inside the NASDAQ, Two-Factor Model

Log Likelihoods	April-May 1999	July-Aug. 1999	Oct.-Nov. 1999
100 S&P Firms	8505	9687	8526
100 NASDAQ Firms	6676	7424	4999
Combined	14,715	16,483	12,084
Test Statistic (df) P-value	932 (41) 1.00	1256 (42) 1.00	2882 (41) 1.00
	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
100 S&P Firms	10197	8403	7825
100 NASDAQ Firms	7848	6457	6032
Combined	17,387	14,323	13,368
Test Statistic (df) P-value	1316 (43) 1.00	1074 (43) 1.00	978 (42) 1.00

Table 3: Tests for Market Integration between S&P 500 and NASDAQ, Two-Factor Model

Log Likelihoods	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
First 50 Firms	4588	4224	4012
Second 50 Firms	3253	2991	2994
All 100 Firms	7740	7156	6919
Test Statistic (df) P-value	202 (43) 1.00	118 (43) 1.00	174 (42) 1.00

Table 4: Tests of Market Integration inside the TSE, Two-Factor Model

Log Likelihoods	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
100 S&P Firms	10197	8403	7825
100 TSE Firms	7740	7156	6919
Combined	17,661	15,294	14,573
Test Statistic (df) P-value	552 (43) 1.00	530 (43) 1.00	342 (42) 1.00

Table 5: Tests for Market Integration between S&P 500 and TSE, Two-Factor Model

Log Likelihoods	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
First 50 AAA Bonds	9220	7245	7893
Second 50 AAA Bonds	10,468	8908	8626
All 100 AAA Bonds	19,113	15,518	16,294
Test Statistic (df) P-value	1150 (43) 1.00	1270 (43) 1.00	450 (42) 1.00

Table 6: Tests of Market Integration inside AAA Bonds, Two-Factor Model

Log Likelihoods	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
100 S&P Firms	10197	8403	7825
100 AAA Bonds	19,113	15,518	16,294
Combined	22,484	18,557	17,556
Test Statistic (df) P-value	14,000 (43) 1.00	5000 (43) 1.00	13,000 (42) 1.00

Table 7: Tests for Market Integration between S&P and AAA Bonds, Two-Factor Model

Log Likelihoods	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
First 50 A+ Bonds	8771	7341	6770
Second 50 A+ Bonds	9563	8249	8571
All 100 A+ Bonds	18,158	15,269	14,612
Test Statistic (df) P-value	352 (43) 1.00	642 (43) 1.00	1458 (42) 1.00

Table 8: Tests of Market Integration inside A+ Bonds, Two-Factor Model

Log Likelihoods	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
100 S&P Firms	10197	8403	7825
100 A+ Bonds	18,158	15,269	14,612
Combined	22,497	18,525	17,520
Test Statistic (df) P-value	12,000 (43) 1.00	10,000 (43) 1.00	10,000 (42) 1.00

Table 9: Tests for Market Integration between S&P and A+ Bonds, Two-Factor Model

Log Likelihoods	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
First 50 Junk Bonds	6731	5934	6481
Second 50 Junk Bonds	7613	6098	6522
All 100 Junk Bonds	14,091	11,981	12,975
Test Statistic (df) P-value	506 (43) 1.00	102 (43) 1.00	56 (42) .93

Table 10: Tests of Market Integration inside Junk Bonds, Two-Factor Model

Log Likelihoods	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
100 S&P Firms	10197	8403	7825
100 Junk Bonds	14,091	11,981	12,975
Combined	20,467	18,181	17,485
Test Statistic (df) P-value	7,600 (43) 1.00	4400 (43) 1.00	6,600 (42) 1.00

Table 11: Tests for Market Integration between S&P and Junk Bonds, Two-Factor Model

	S&P 500	NASDAQ	TSE	AAA Bonds	A+ Bonds	Junk Bonds
S&P 500	-	.07	.04	.19	.06	.17
NASDAQ	.06	-	.09	.15	.10	.23
TSE	.04	.08	-	.23	.03	.15
AAA Bonds	.16	.13	.19	-	.24	.35
A+ Bonds	.06	.09	.03	.21	-	.15
Junk Bonds	.17	.23	.15	.33	.15	-

Table 12: Degree of Market Integration, April-May 2002

Mean Absolute Difference of Deltas below diagonal; Grubel-Lloyd Measure above diagonal

	S&P 500	NASDAQ	TSE	AAA Bonds	A+ Bonds	Junk Bonds
S&P 500	-	.13	.05	.05	.05	.07
NASDAQ	.12	-	.12	.11	.15	.17
TSE	.05	.11	-	.04	.04	.05
AAA Bonds	.05	.10	.03	-	.03	.05
A+ Bonds	.05	.13	.04	.03	-	.02
Junk Bonds	.07	.16	.05	.05	.02	-

Table 13: Degree of Market Integration, July-Aug. 2002

Mean Absolute Difference of Deltas below diagonal; Grubel-Lloyd Measure above diagonal

	S&P 500	NASDAQ	TSE	AAA Bonds	A+ Bonds	Junk Bonds
S&P 500	-	.09	.07	.10	.17	.04
NASDAQ	.08	-	.14	.11	.18	.10
TSE	.07	.13	-	.09	.13	.06
AAA Bonds	.09	.10	.08	-	.11	.07
A+ Bonds	.15	.17	.12	.10	-	.15
Junk Bonds	.04	.09	.06	.06	.13	-

Table 14: Degree of Market Integration, Oct.-Nov. 2002

Mean Absolute Difference of Deltas below diagonal; Grubel-Lloyd Measure above diagonal

	S&P 500	NASDAQ	TSE	AAA Bonds	A+ Bonds	Junk Bonds
S&P 500	-	-.58	.57	-.65	-.59	.23
NASDAQ	-.67	-	-.22	.74	.45	-.59
TSE	.55	-.24	-	-.26	-.29	.04
AAA Bonds	-.64	.80	-.23	-	.81	-.52
A+ Bonds	-.56	.46	-.29	.72	-	-.29
Junk Bonds	.27	-.59	.06	-.58	-.27	-

Table 15: Degree of Market Integration, April-May 2002

Brandt et al measure in logs below diagonal; in levels above diagonal

	S&P 500	NASDAQ	TSE	AAA Bonds	A+ Bonds	Junk Bonds
S&P 500	-	-.09	-.20	.10	.07	.08
NASDAQ	-.10	-	.52	-.14	-.17	-.11
TSE	-.21	.54	-	-.07	-.04	-.01
AAA Bonds	.10	-.15	-.06	-	.96	.96
A+ Bonds	.07	-.18	-.04	.95	-	.99
Junk Bonds	.08	-.12	.00	.95	.99	-

Table 16: Degree of Market Integration, July-Aug. 2002

Brandt et al measure in logs below diagonal; in levels above diagonal

	S&P 500	NASDAQ	TSE	AAA Bonds	A+ Bonds	Junk Bonds
S&P 500	-	.53	-.67	.58	-.59	.59
NASDAQ	.51	-	-.52	.38	-.47	.33
TSE	-.68	-.49	-	-.57	.57	-.53
AAA Bonds	.60	.40	-.59	-	-.91	.98
A+ Bonds	-.60	-.47	.55	-.95	-	-.85
Junk Bonds	.61	.33	-.55	.97	-.87	-

Table 17: Degree of Market Integration, Oct.-Nov. 2002

Brandt et al measure in logs below diagonal; in levels above diagonal

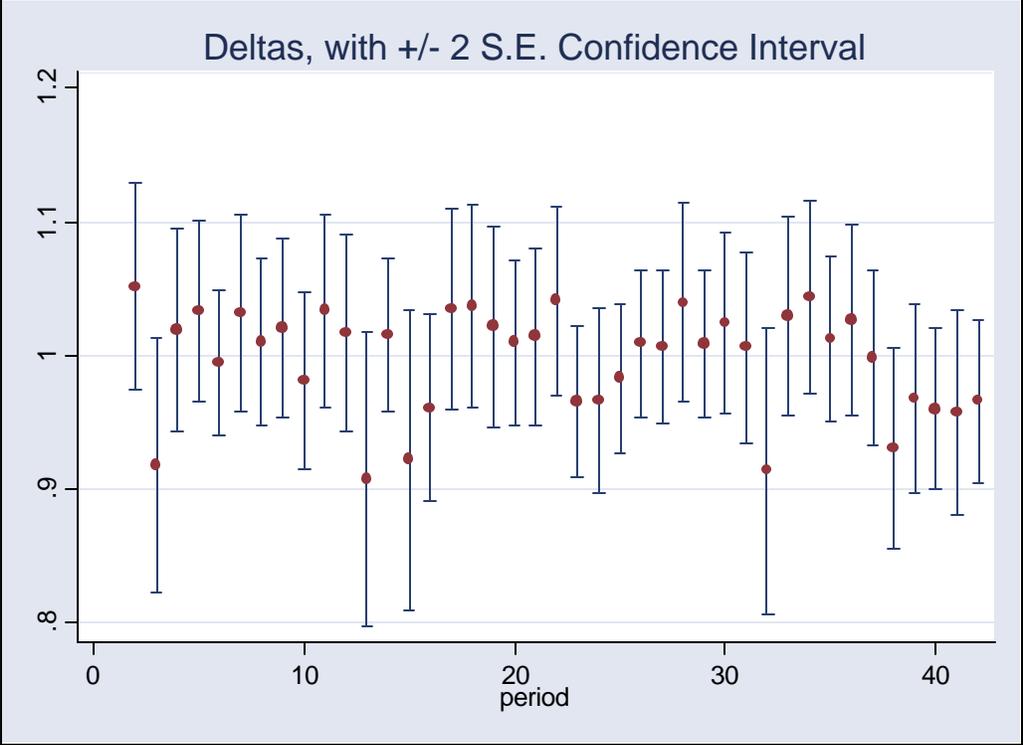


Figure 1: Expected Discount Rates from Fifty S&P 500 Firms, April-May 1999

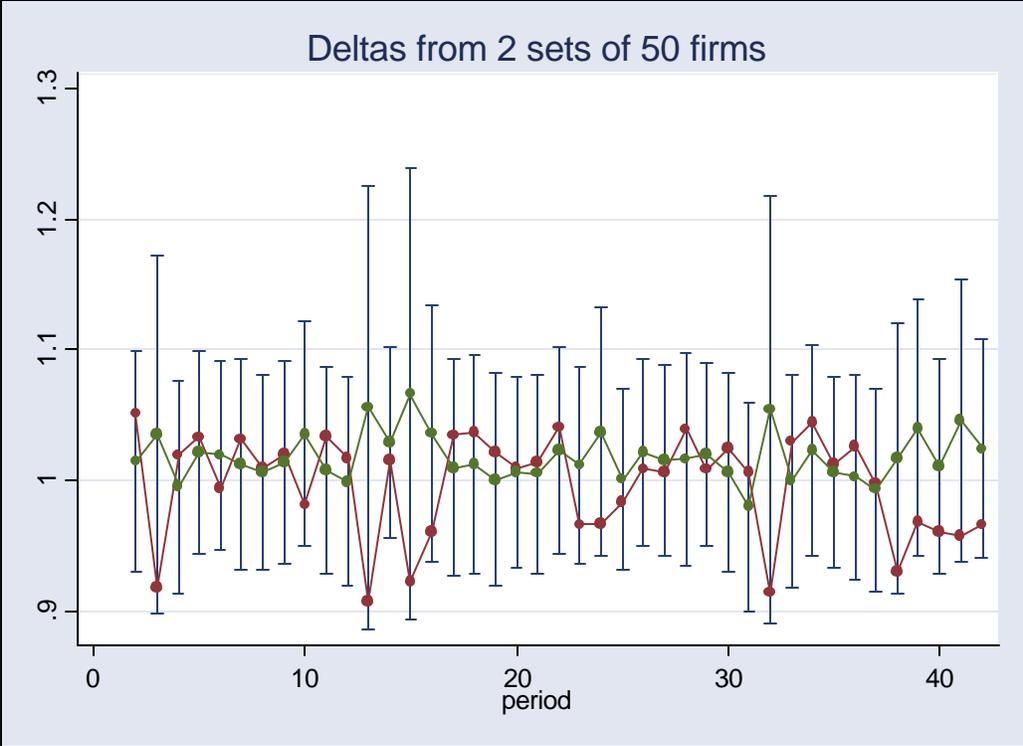


Figure 2: Expected Discount Rates from Two Sets of S&P 500 Firms, April-May 1999

Deltas from 100 S&P firms, 1999 April-May

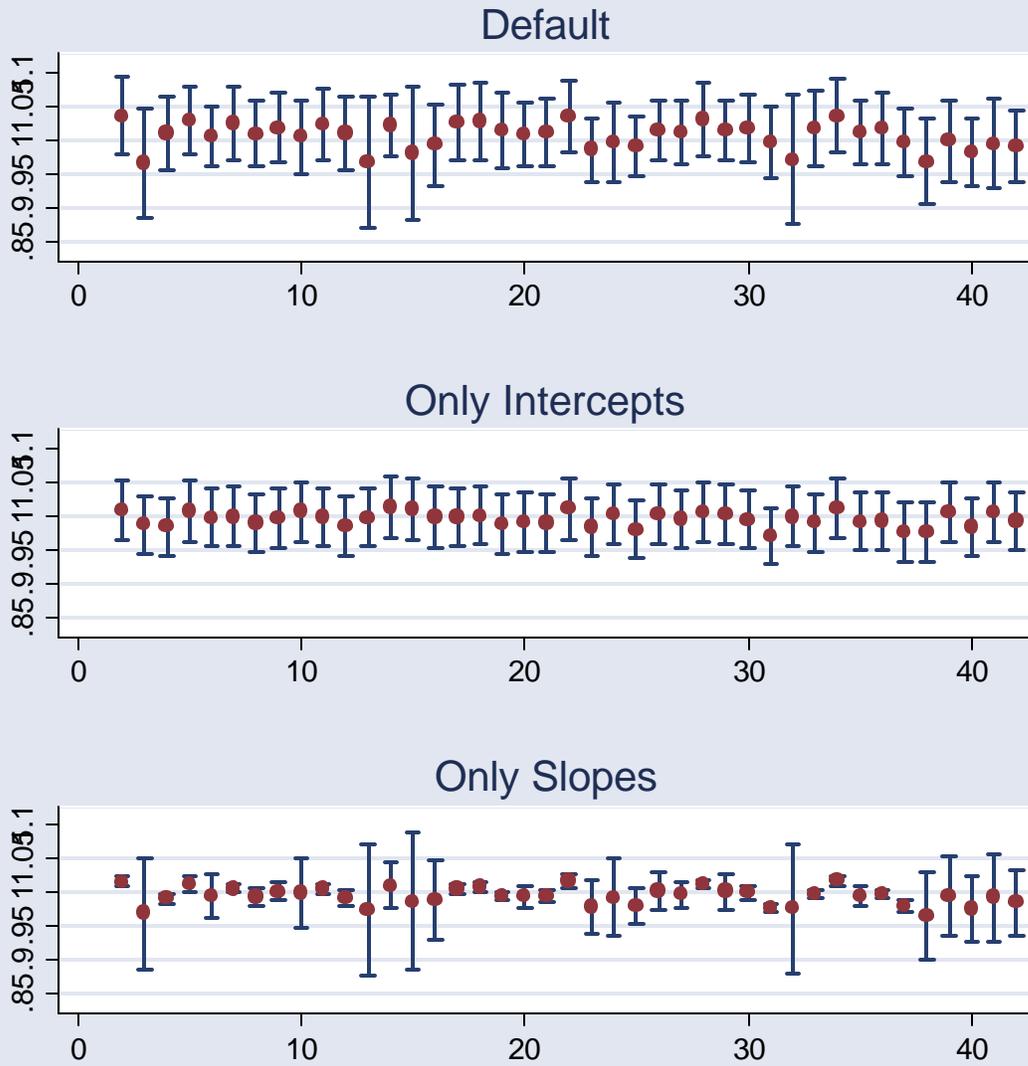
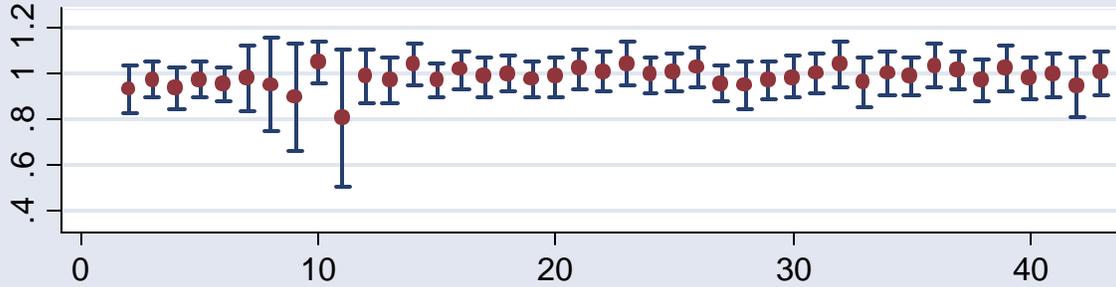


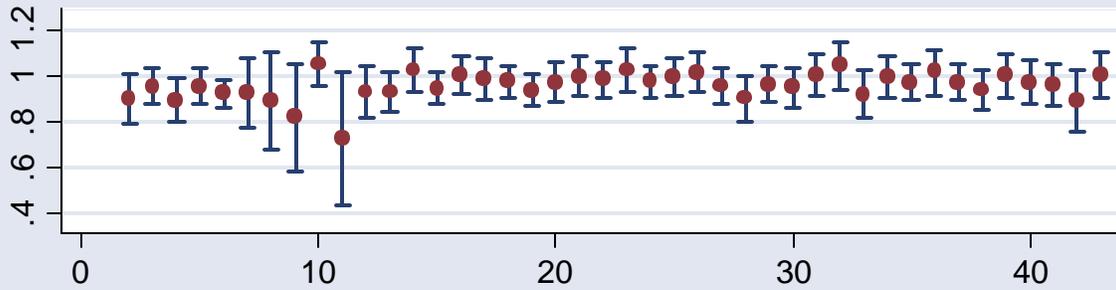
Figure 3: Deltas from Different Factor Models, 100 S&P Firms, April-May 1999

Deltas from S&P firms, 2002 Oct-Nov

First 50 Firms



Second 50 Firms



All 100 Firms

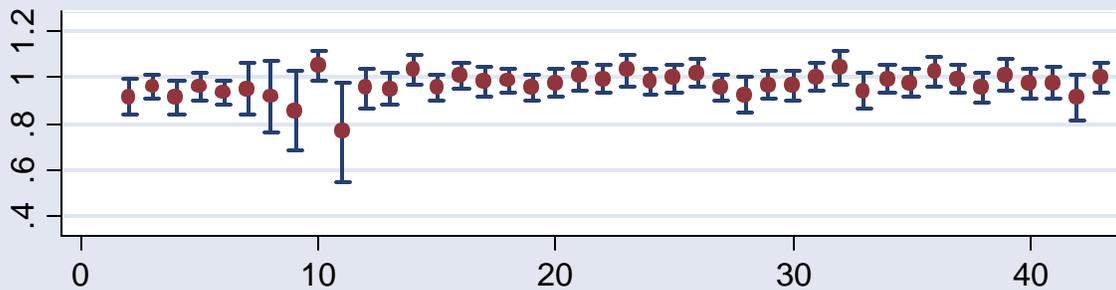


Figure 4: Expected Discount Rates from Different Sets of S&P 500 Firms, Oct-Nov 2002

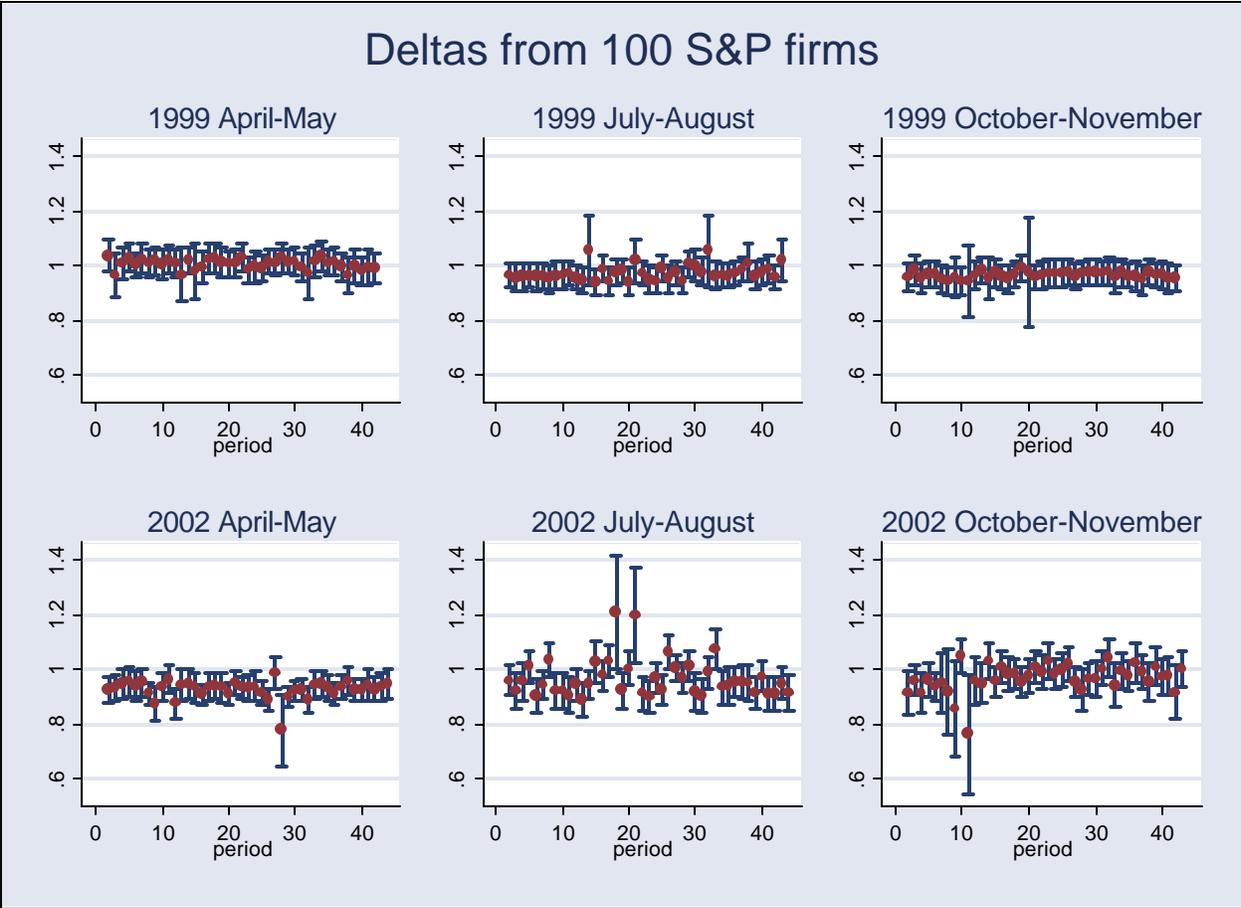


Figure 5: Expected Discount Rates from Sets of 100 S&P 500 Firms,

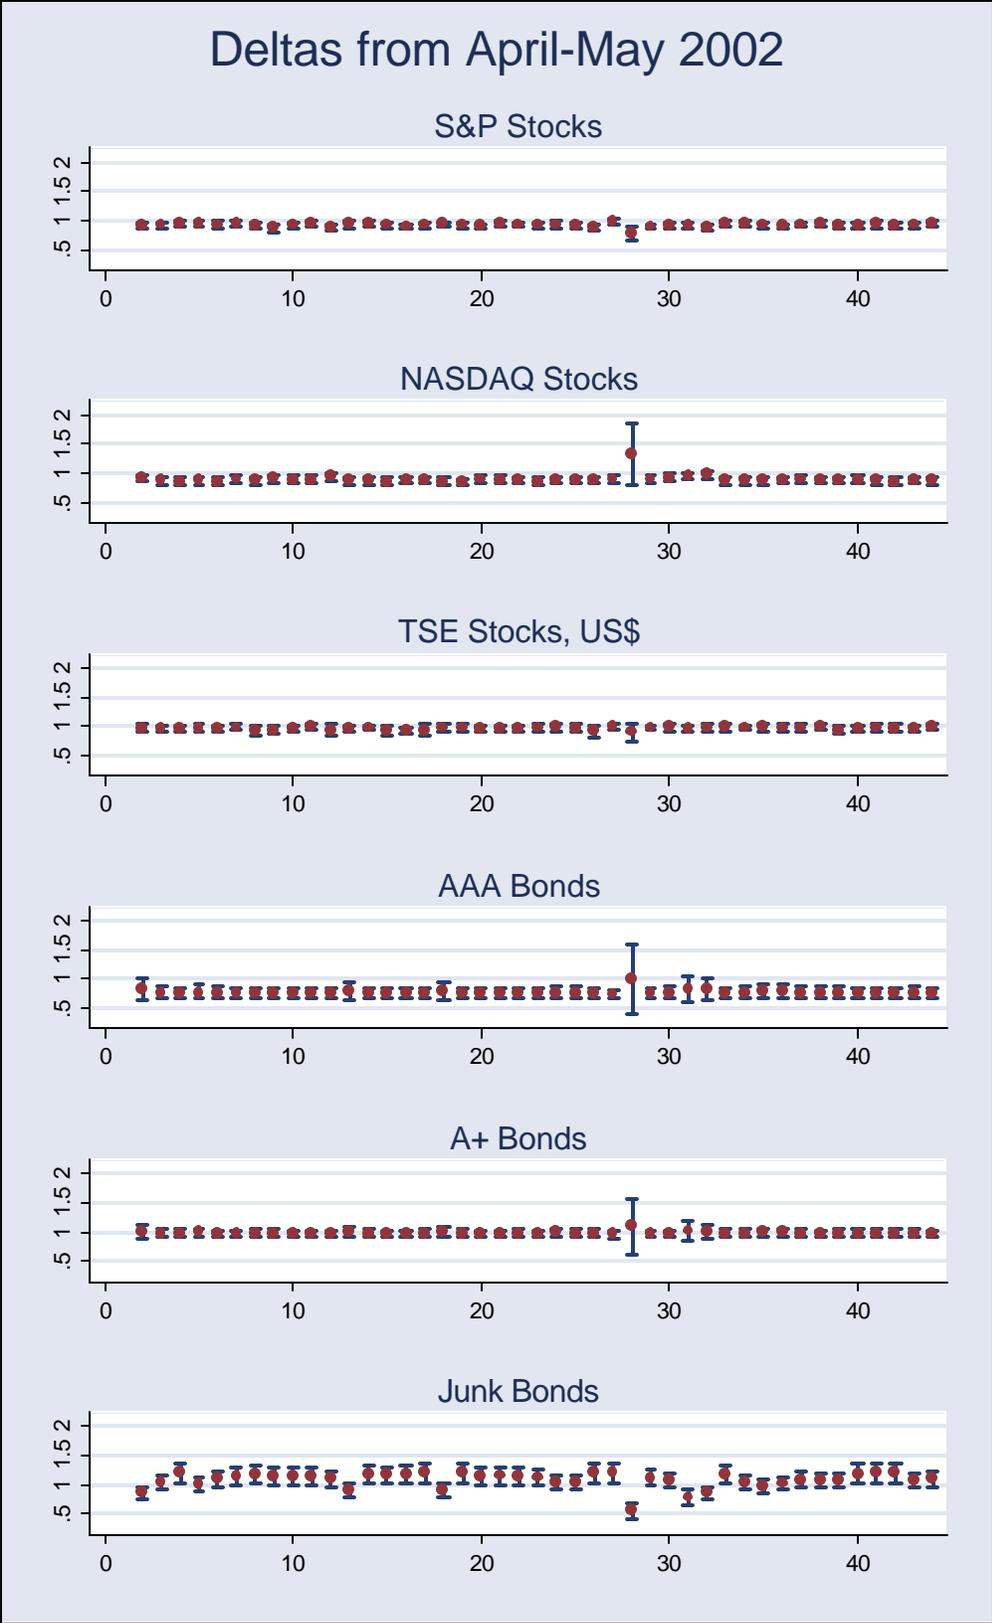


Figure 6: Expected Discount Rates from Different Assets, 2 Factor Model, April-May 2002

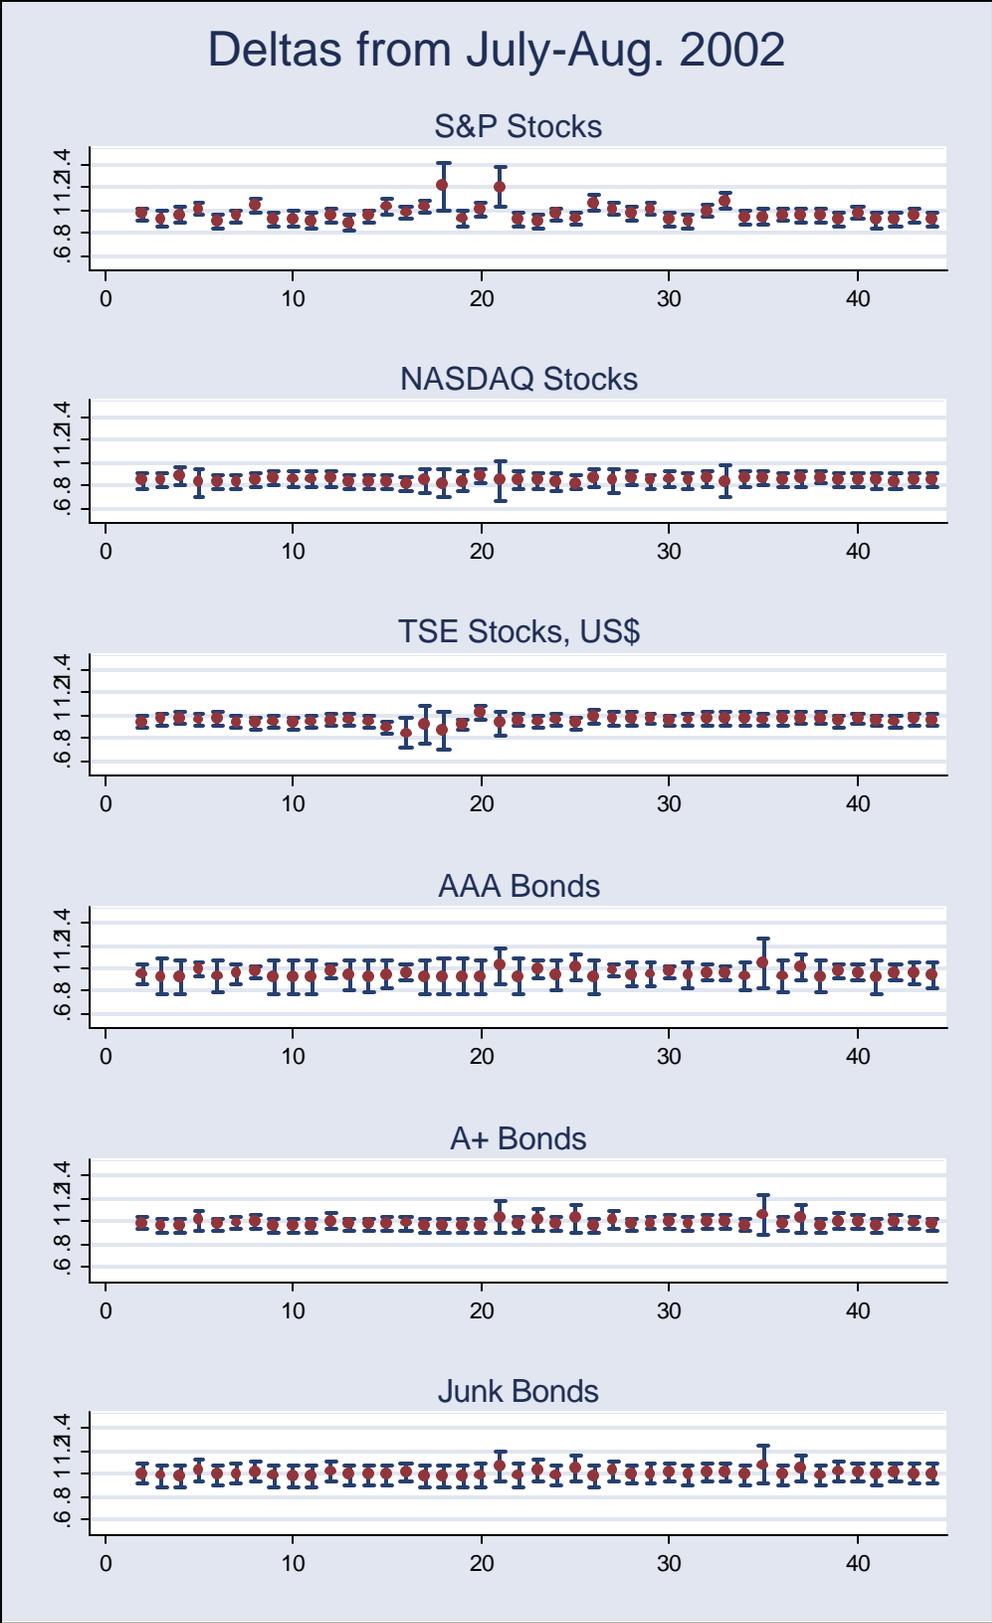


Figure 7: Expected Discount Rates from Different Assets, 2 Factor Model, July-Aug. 2002

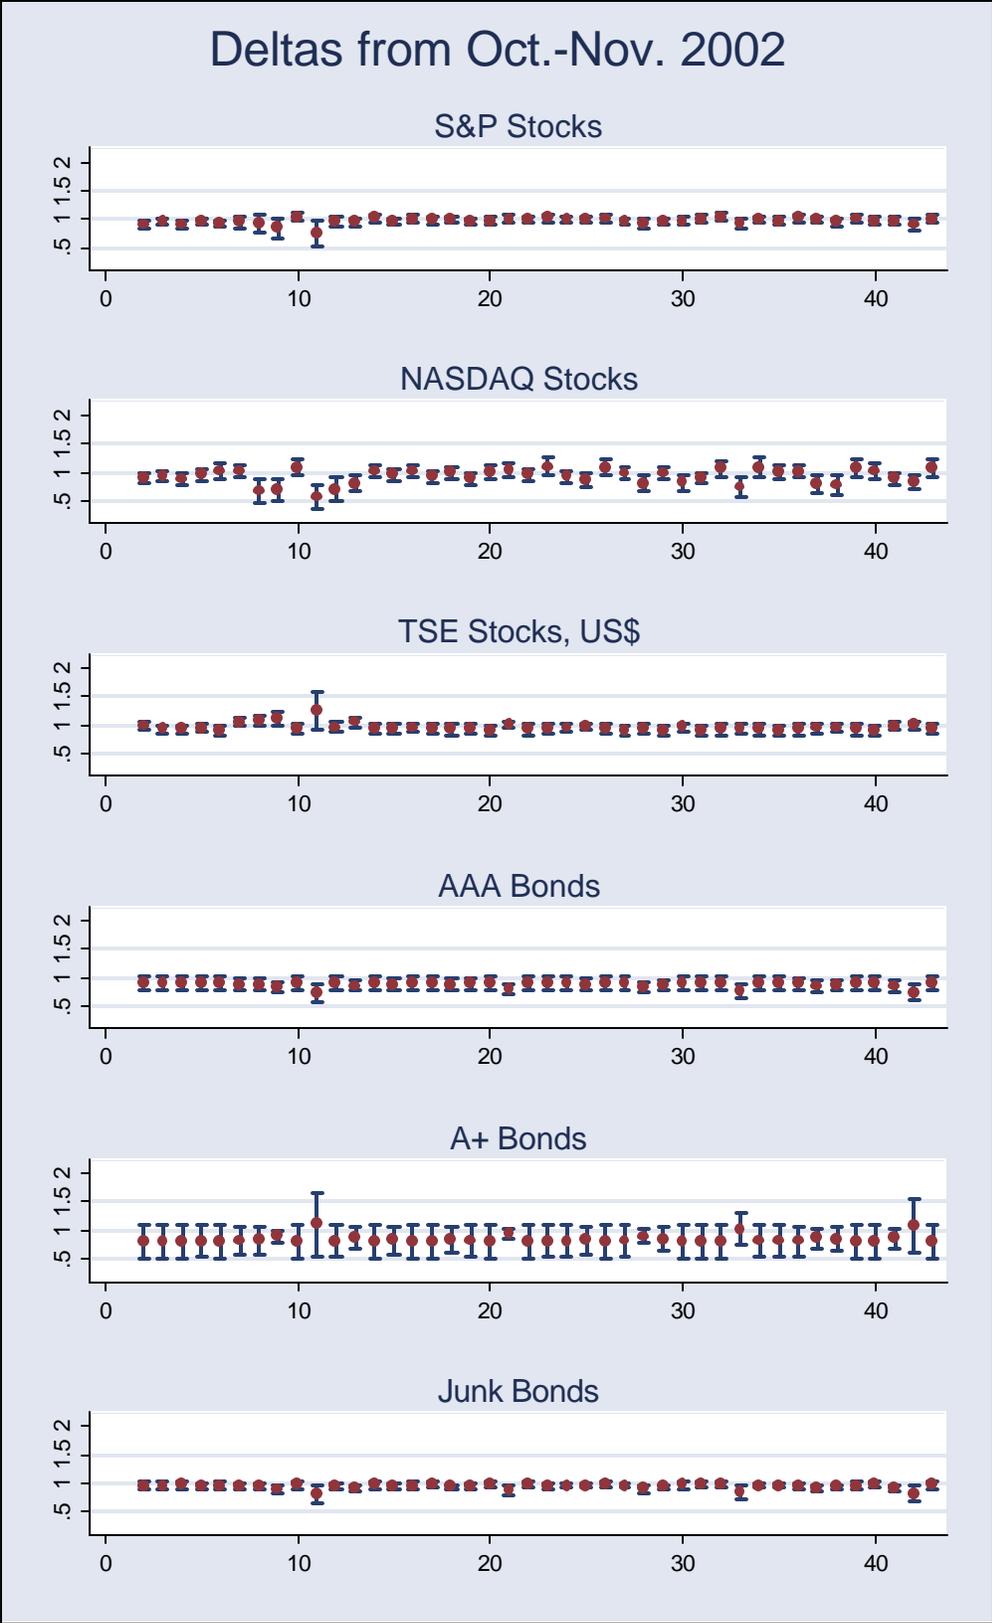


Figure 8: Expected Discount Rates from Different Assets, 2 Factor Model, Oct.-Nov. 2002

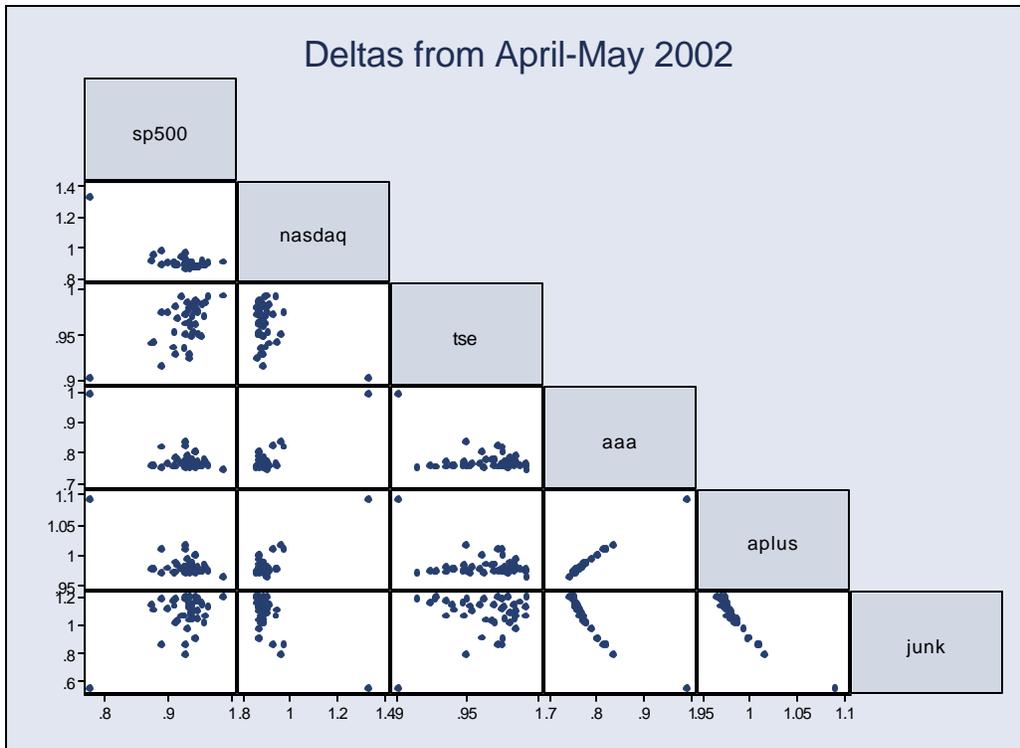


Figure 9: Expected Discount Rates from Different Assets, 2 Factor Model, April-May 2002

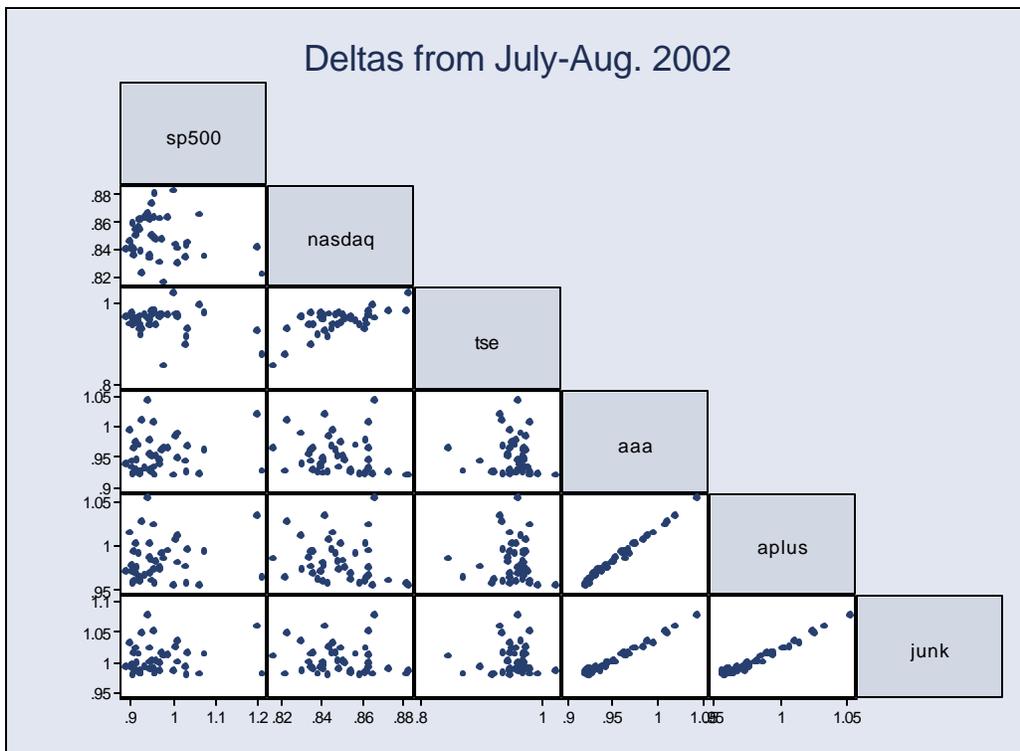


Figure 10: Expected Discount Rates from Different Assets, 2 Factor Model, July-Aug. 2002

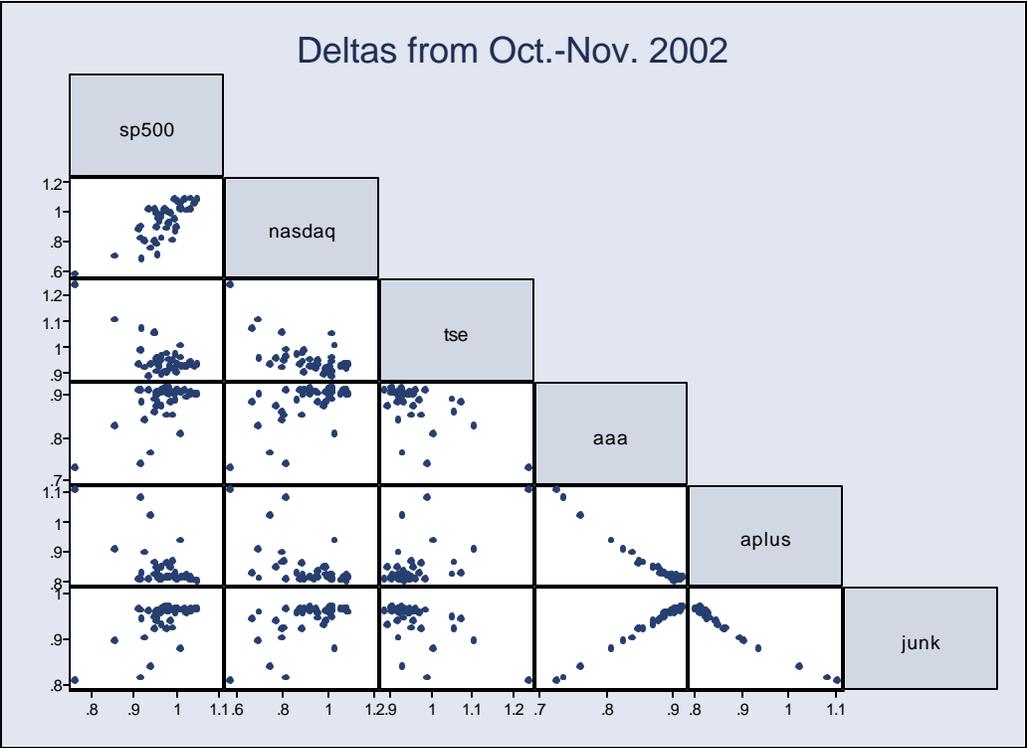


Figure 11: Expected Discount Rates from Different Assets, 2 Factor Model, Oct.-Nov. 2002

Endnotes

¹ Any liquidity or other services accruing to the asset holders are also included in x , though in this paper we ignore such phenomena.

² To our knowledge this Euler equation is present in all existing equilibrium asset pricing models.

³ Note that d_{t+1} is a real discount rate – one that discounts the real payoff, x_{t+1}^j , and turns it into a real price p_t^j . The units of the discount rate depend on the units of the asset payoff, and the units used for the asset price. If next period's payoff had been in American dollars, then we would have used an American dollar discount rate. Had the payoff been in Canadian dollars, we would have used a Canadian dollar discount rate. Changes in discount rate units to say American dollars are made by inflating the payoff and current period asset price by the appropriate American price levels for periods $t+1$ and t respectively, and then undoing that operation in the discount rate – as below:

$$Q_t p_t^j = E_t \frac{r u'(c_{t+1}) Q_t}{u'(c_t) Q_{t+1}} Q_{t+1} x_{t+1}^j$$

Here $d_{t+1} = \frac{r u'(c_{t+1}) Q_t}{u'(c_t) Q_{t+1}}$, an American dollar discount rate. In practice, at high frequencies we assume $Q_t = Q_{t+1}$.

We maintain the d_{t+1} notation for all discounting. In this paper, all prices and payoffs will be quoted in or converted to U.S. -dollar units.

⁴ Such equations are usually deflated by p_t^j in practice.

⁵ Without any normalization, the covariances are proportional to prices, so that with constant covariances, we end up with a term proportional to the price in the residual. That is, dividing by the lagged price makes the residual better behaved. We have experimented with other normalizations, such as the average level of prices at time t \bar{p}_t , and found similar results.

⁶ We prefer not to use a concentrated maximum likelihood estimator, since we would rather not take a stand on the distribution of ε . In passing, we have experimented some with starting values, and never found local maxima to be a problem in practice.

⁷ We impose no constraints on δ so that it need not be e.g., greater than unity; we see no reason why constraints could not be added in future work.

⁸ We use the Toronto Stock Exchange Index for Canadian Stocks and the Lehman Brothers bond index for all bonds.

⁹ This is especially true when we do not include many time-varying factors.

¹⁰ We choose these months to avoid January (and its effect), February (a short month), and March (a quarter-ending month), but test for sample sensitivity extensively below.

¹¹ For instance, we could use data at five-minute intervals for a day, making our assumption of constant asset-specific effects even more plausible; but the question of whether financial markets are integrated over hours (not weeks) is less interesting to us.

¹² The New York Stock Exchange closes at 4:00pm daily, as does the Toronto stock exchange in the same time zone, a fact we use later on. See *The Compact Handbook of World Stock, Derivative and Commodity Exchanges, Year 2001*. Note that Thomson provides price quotes for holidays that we use, in part to ensure consistent samples across markets.

¹³ We include data from firms like: Ace Ltd., Transocean Inc., ADC Telecommunications, AES Corp., AMR Corp., AT&T Corp., Adobe Sys Inc., AMD Inc., Air Prods. & Chems. Inc., and Allegheny Energy Inc.

¹⁴ In practice, there is little cross-sectional dependence left in our residuals (the time dummies seem to pick it all up); a regression of the residuals on a comprehensive set of time dummies yields an R^2 of essentially zero. Still, one could always use GMM in the event of encountering such problems. We have experimented with GMM, and it seems typically to deliver results almost identical to those of least squares in practice. We have also experimented

with different standard errors. Conventional asymptotic standard errors tend to be a little smaller than the Newey-West ones, which are similar to bootstrapped ones. For instance, consider the one-factor model estimated for the 100 S&P firms in April-May 1999. Conventional standard errors average .017 (with little variation), Newey-West standard errors average .022, and bootstrapped standard errors average .023. Also, the bootstrap shows only very trivial bias in our delta estimates.

¹⁵ As shown in Table 1, the log likelihood of our equation estimated on 50 S&P firms is 4192. In April-May 1999, the US 3-month treasury bill rate averaged 4.4%, while the US 30-year treasury bond averaged 5.93%; these are daily rates of 1.00017 and 1.00023 respectively. If delta had been constant at the average riskless interest rate, it might have been expected to average between these values. Yet the log likelihood for the default equation estimated with a constant of this magnitude in place of the deltas is only 4061. Under the null hypothesis of deltas that are constant and equal to the inverse of the riskless interest rate, $2*(4192-4061)$ is distributed as a chi-square with 41 degrees of freedom, grossly inconsistent with the null. When we use all 100 firms, the analogue is $2*(8505-8281)$, again grossly inconsistent with the null.

¹⁶ One could also split the assets by e.g., size, beta, or something else while testing for integration within a market.

¹⁷ They are available from the “Rates and Statistics” section of www.bankofcanada.ca. The close rate is updated at about 4:30pm, some 30 minutes later than both the TSE and NYSE close. This adds some measurement error which is probably small, since a) the C\$/ rate is stable during this period; and b) the C\$/ market tends to be inactive from 4:00pm to 4:30pm EST/EDST. There are a few days when the C\$/ exchange market is closed in Canada; in this case we substitute closing rates from the *Financial Times*.

¹⁸ Our rejections of integration do not seem to stem from assuming that our asset-specific effects are constant for a two-month sample. In appendix 2, we present the analogues to Tables 1 through 3, but computed only with the first-half of the (two-month) sample. The results from the one-month sample are quite similar to those from the two-month sample.