

Estimating the Intertemporal Risk–Return Tradeoff Using the Implied Cost of Capital

ĽUBOŠ PÁSTOR, MEENAKSHI SINHA, and BHASKARAN SWAMINATHAN*

ABSTRACT

We argue that the implied cost of capital (ICC), computed using earnings forecasts, is useful in capturing time variation in expected stock returns. First, we show theoretically that ICC is perfectly correlated with the conditional expected stock return under plausible conditions. Second, our simulations show that ICC is helpful in detecting an intertemporal risk–return relation, even when earnings forecasts are poor. Finally, in empirical analysis, we construct the time series of ICC for the G-7 countries. We find a positive relation between the conditional mean and variance of stock returns, at both the country level and the world market level.

THE TRADEOFF BETWEEN RISK and return is a central concept in finance. Finance theory generally predicts a positive risk–return relation, both across assets and over time. For example, the intertemporal capital asset pricing model (ICAPM) of Merton (1973) predicts a positive time-series relation between the conditional mean and variance of market returns. However, the empirical evidence on the sign of the intertemporal risk–return relation is inconclusive.¹

To explain the mixed nature of the evidence, some researchers have shown that the intertemporal mean–variance relation need not be positive theoretically (e.g., Abel (1988), Backus and Gregory (1993), and Whitelaw (2000)). Others have argued that a positive mean–variance relation emerges when the empirical specification includes hedging demands (e.g., Scruggs (1998) and Guo and Whitelaw (2006)). Yet others argue that the relation is highly sensitive to the way conditional variance is measured (e.g., Harvey (2001), Wang (2004), and

*Ľuboš Pástor is at the Graduate School of Business at the University of Chicago, NBER, and CEPR. Meenakshi Sinha is at Cornerstone Research. Bhaskaran Swaminathan is at LSV Asset Management and the Johnson Graduate School of Management at Cornell University. Helpful comments were gratefully received from Ray Ball, George Constantinides, Doug Diamond, Gene Fama, Christian Leuz, Toby Moskowitz, Rob Stambaugh, Pietro Veronesi, Robert Whitelaw, Franco Wong, the anonymous referee, the anonymous associate editor who acted as editor, and the audiences at the University of British Columbia, the University of Chicago accounting workshop, the University of Chicago finance lunch workshop, and the 2005 meetings of the Western Finance Association.

¹ Some studies find a positive relation (e.g., Scruggs (1998), Ghysels, Santa-Clara, and Valkanov (2005), Lundblad (2007), and Ludvigson and Ng (2007)), but others find a negative relation (e.g., Campbell (1987), Turner, Startz, and Nelson (1989), Glosten, Jagannathan, and Runkle (1993), Harvey (2001), Lettau and Ludvigson (2007), and Brandt and Kang (2004)), and yet others find either no significant relation or mixed evidence (e.g., French, Schwert, and Stambaugh (1987), Nelson (1991), Campbell and Hentschel (1992), Chan, Karolyi, and Stulz (1992), and Whitelaw (1994, 2000)).

Ghysels, Santa-Clara, and Valkanov (2005)). For example, Ghysels et al. state that “the main difficulty in testing the ICAPM relation is that the conditional variance of the market is not observable” and that “the conflicting findings of the above studies are mostly due to differences in the approach to modeling the conditional variance” (p. 510).

While estimating the conditional variance of market returns is clearly important, we believe that estimating the conditional mean return is no less important. First moments of returns are generally more difficult to estimate than second moments (Merton (1980)). The conditional mean return is sometimes estimated by projecting future returns onto a set of conditioning variables.² The results produced by this approach tend to be sensitive to the choice of the conditioning variables (Harvey (2001)). Another popular estimate of the conditional mean return in this literature is the realized future return.³ Although realized returns provide unbiased estimates of expected returns, they are notoriously noisy. For example, Elton (1999 p. 1199) argues that “realized returns are a very poor measure of expected returns.” Lundblad (2007) shows that when realized returns proxy for expected returns, a very long sample is needed to detect a positive risk–return relation in simulations.

This paper reexamines the conditional mean–variance relation using a different proxy for the conditional expected return: the implied cost of capital (ICC). The ICC for a given asset is the discount rate (or internal rate of return) that equates the asset’s market value to the present value of all expected future cash flows. The literature on the ICC has developed in part in response to the failure of the standard asset pricing models to provide precise estimates of the firm-level cost of equity capital.⁴ One appealing feature of the ICC as a proxy for expected returns is that it does not rely on noisy realized asset returns.

We focus on the time series of the ICC whereas much of the extant literature analyzes the cross-section. The evidence on the cross-sectional relation between the ICC and risk is mixed. Some studies find a positive relation between the ICC and market beta (e.g., Kaplan and Ruback (1995), Botosan (1997), Gode and Mohanram (2003), and Easton and Monahan (2005)), while others find this relation to be mostly insignificant (e.g., Gebhardt, Lee, and Swaminathan (2001) and Lee, Ng, and Swaminathan (2007)). The ICC seems to be more closely related to stock return volatility than to beta (e.g., Friend, Westerfield, and Granito (1978) and Hail and Leuz (2006)). Botosan and Plumlee (2005) report that some ICC estimates are significantly related to firm risk while others are not. Instead of further analyzing the cross-sectional relation between the firm-level ICC and firm risk, we estimate the time-series relation between the market-level ICC and market risk.

² See, for example, Campbell (1987), Harvey (2001), Whitelaw (1994), and Lettau and Ludvigson (2007).

³ See, for example, French, Schwert, and Stambaugh (1987), Nelson (1991), Chan, Karolyi, and Stulz (1992), Glosten, Jagannathan, and Runkle (1993), Scruggs (1998), Ghysels, Santa-Clara, and Valkanov (2005), and Lundblad (2007).

⁴ See, for example, Fama and French (1997), Elton (1999), and Pástor and Stambaugh (1999).

The accounting literature evaluates the usefulness of the ICC as a proxy for the expected stock return mainly by testing whether the ICC can predict realized returns. The general conclusion is that the ICC is not a very good proxy (e.g., Guay, Kothari, and Shu (2003) and Easton and Monahan (2005)).⁵ However, the literature recognizes that the large amount of noise in realized returns limits the power of the predictability tests. Further complications arise when the expected return varies over time because a high realized return often signals that the expected return is falling rather than that the expected return is high. A different approach is adopted by Botosan and Plumlee (2005), who assess the usefulness of five ICC measures based on their ability to capture the cross-sectional relation between expected returns and risk. In a similar spirit, we judge the ICC based on its ability to detect the time-series relation between expected returns and risk, and find the ICC to be quite useful. We also show analytically that the ICC should be a good proxy for the conditional expected stock return.

In our theoretical analysis, we examine the relation between the ICC and the conditional expected stock return. We show that if dividend growth follows an AR(1) process, the ICC is a linear function of the dividend yield and dividend growth. If, in addition, the conditional expected return also follows an AR(1) process, then the ICC is perfectly correlated with the conditional expected return over time. Therefore, the ICC should be useful in capturing time variation in expected returns.

In our simulation analysis, we analyze the usefulness of the ICC in estimating the intertemporal risk–return tradeoff. We design a simple framework in which the conditional mean and variance of stock returns are positively related. We simulate the time series of the conditional moments and compare the ability of various proxies for the conditional mean to detect the positive mean–variance relation. We find that the relation is much easier to detect using the ICC than using realized returns.

Importantly, the ICC outperforms realized returns even in the absence of information about dividend growth. In that case, the ICC is perfectly correlated with the dividend–price ratio, so its changes are driven mostly by changes in the stock price: Increases in the stock price are accompanied by declines in the ICC, and vice versa. As long as the stock price changes are driven to some extent by changes in expected returns, the ICC should be positively related to the conditional variance. In line with this intuition, we find that the ICC–variance correlation is high, especially when stock returns are driven mostly by changes in expected returns (as opposed to changes in expected cash flows). However, the ICC outperforms realized returns also when only a small fraction

⁵ One exception in this literature is Gode and Mohanram (2003), who find some significant associations between the ICC and future returns at the portfolio level. They find that the ICC estimates based on the Gebhardt, Lee, and Swaminathan (2001) model tend to outperform those based on the Ohlson and Juettner-Nauroth (2005) model in forecasting 1- and 2-year-ahead returns, and that both models perform well in forecasting 3-year-ahead returns. Guay, Kothari, and Shu (2003) also find that the Gebhardt-Lee-Swaminathan ICC estimates have significant predictive power for 2- and 3-year-ahead returns.

of the market return variance is due to time-varying expected returns. In short, the ICC seems well suited for capturing the risk–return tradeoff, even when we have little information about future cash flow.

The accounting literature has developed a variety of approaches to estimating the ICC (e.g., Claus and Thomas (2001), Gebhardt, Lee, and Swaminathan (2001), Easton et al. (2002), Easton (2004), and Ohlson and Juettner-Nauroth (2005)). We do not take a stand on which approach is best; instead, we argue that the whole class of ICC models should be useful in capturing time variation in expected returns. In any ICC model, a large part of the time variation in ICC is due to changes in stock prices. Moreover, there is empirical evidence that changes in stock prices at the market level are driven mostly by changes in expected returns (e.g., Campbell and Ammer (1993)). Therefore, any sensible measure of ICC should capture some of the time variation in expected market return. Our empirical construction of ICC builds on the work of Gebhardt, Lee, and Swaminathan (2001) and Lee, Ng, and Swaminathan (2007), but we also show that using the alternative approach of Easton et al. (2002) leads to the same conclusions.

In our empirical analysis, we estimate the intertemporal relation between the conditional mean and variance of excess market returns in the G-7 countries. We construct monthly estimates of the conditional mean and variance in 1981 to 2002 (for the United States) and 1990 to 2002 (for Canada, France, Germany, Italy, Japan, and United Kingdom). To proxy for the conditional mean, we first compute the ICC for each firm in each month by using analyst forecasts of earnings and historical plowback rates. We then aggregate these cost of capital estimates across firms to compute a market-wide ICC for each country, both equal- and value-weighted. Finally, we subtract the long-term local government bond yield from the ICC to compute the implied risk premium for each country. This implied risk premium is the measure of the conditional mean return that we use in our regression tests.

To estimate the conditional variance of market returns for a given country in a given month, we average squared daily market returns over the previous month. This approach to variance estimation is simpler than some other approaches developed in the literature.⁶ Although we believe that it is important to estimate the conditional variance as precisely as possible, we choose a simple variance estimator to highlight the paper's focus on the conditional expected return.

We find a positive relation between the conditional mean and variance of market returns. Consider the equal-weighted average implied risk premium as a proxy for the expected excess market return at the country level. We find a positive relation between the levels of the implied risk premia and volatility in all G-7 countries, and this relation is statistically significant for five of the seven countries. We also find a positive and statistically significant relation between

⁶ Our estimator corresponds to the simplest variance estimator considered by French, Schwert, and Stambaugh (1987). Andersen et al. (2003) discuss the advantages of realized volatility relative to ARCH, stochastic volatility, and other parametric volatility models.

shocks to the risk premia and shocks to volatility in Canada, France, Germany, United Kingdom, and United States. The evidence based on the value-weighted average implied risk premium is somewhat weaker but still generally supportive of a positive mean–variance relation. We find a positive and significant relation between the levels of the implied risk premia and volatility in four of the seven countries. The relation between the shocks to the premia and to volatility is positive and significant for three countries (France, United Kingdom, and United States). The results are similar whether we use variance or standard deviation to measure volatility.

We also find a positive intertemporal risk–return tradeoff at the global level. There is a positive relation between the world market volatility and the world market implied risk premium, approximated by averaging the implied risk premia across the G-7 countries. There is also a positive relation between several individual country risk premia and the world market volatility. Finally, some country risk premia are positively related to the conditional covariances of the country returns with the world market portfolio. This evidence is consistent with partial international integration of the G-7 financial markets.

It is noteworthy that we find any statistically significant relations at all, given the short length of our samples (22 years for the United States and 13 years for the other six countries) and the fact that we estimate the conditional variance in a simple manner. In contrast, the tests that use realized returns to proxy for expected returns do not find a significant risk–return tradeoff in any of the seven countries. Consistent with our simulation evidence, the ICC seems more powerful than realized returns in capturing time-varying expected returns.

To assess the robustness of our results, we estimate return volatility using the implied volatility from the options market, which is available to us for the U.S. stock market. The results based on implied volatility are even stronger than those based on realized volatility. The mean–variance relation is significantly positive with the *t*-statistics on the order of ten in a 17-year-long sample. Additional tests show that the mean–variance relation remains positive after controlling for hedging demands, and that this relation is not driven by analyst forecast errors. Finally, the mean–variance relation weakens but remains mostly positive when we replace the ICC by the dividend yield, effectively discarding the information about dividend growth contained in analysts' earnings forecasts.

To summarize our contribution, this paper bridges two previously unconnected strands of the literature. The first strand uses the ICC as a proxy for the expected stock return. While most of this literature focuses on the cross-section of the firm-level ICC, we focus on the time series of the market-level ICC. We argue that the ICC is a valuable proxy for the expected stock return despite its (previously documented) failure to reliably predict future stock returns. Our contribution is to show—theoretically, in simulations, as well as empirically—that the ICC is useful in capturing time variation in expected stock market returns. The second strand of the literature focuses on the time-series relation between the conditional mean and volatility of stock market returns. Our

contribution is to estimate this relation by using a novel proxy (ICC) for the conditional expected return. Unlike most of the literature, we find a significantly positive mean–variance relation. We find the positive relation not only in the United States but also in several international markets, as well as at the global market level.

The paper is organized as follows. Section I characterizes the ICC analytically and relates it to the conditional expected return. Section II provides simulation evidence on the usefulness of the ICC in estimating the risk–return tradeoff. Section III describes our data and empirical methodology. Section IV discusses the empirical results. Section V concludes.

I. Implied Cost of Capital (ICC)

The ICC is the discount rate that equates the present value of expected future dividends to the current stock price. When the conditional expected stock return is constant over time, it is equal to the ICC. However, when the expected return varies over time, which is the realistic scenario considered here, the ICC and the expected return are generally different. In this section, we analytically characterize the relation between the two quantities.

One common approach is to define the ICC as the value of r_e that solves

$$P_t = \sum_{k=1}^{\infty} \frac{E_t(D_{t+k})}{(1+r_e)^k}, \quad (1)$$

where P_t is the stock price and D_t are the dividends at time t . For tractability, we propose a slightly different but analogous definition. Campbell and Shiller (1988) develop a useful approximation to the present value formula, which expresses the log price $p_t = \log(P_t)$ as

$$p_t = \frac{k}{1-\rho} + (1-\rho) \sum_{j=0}^{\infty} \rho^j E_t(d_{t+1+j}) - \sum_{j=0}^{\infty} \rho^j E_t(r_{t+1+j}), \quad (2)$$

where r_t is the log stock return, $d_t \equiv \log(D_t)$, $\rho = 1/(1 + \exp(\overline{d} - \overline{p}))$, $k = -\log(\rho) - (1-\rho)\log(1/\rho - 1)$, and $\overline{d} - \overline{p}$ is the average log dividend–price ratio. In this framework, it is natural to define the ICC as the value of $r_{e,t}$ that solves

$$p_t = \frac{k}{1-\rho} + (1-\rho) \sum_{j=0}^{\infty} \rho^j E_t(d_{t+1+j}) - r_{e,t} \sum_{j=0}^{\infty} \rho^j. \quad (3)$$

To provide some insight into the ICC, it is convenient to assume that log dividend growth $g_{t+1} \equiv d_{t+1} - d_t$ follows a stationary AR(1) process:

$$g_{t+1} = \gamma + \delta g_t + v_{t+1}, \quad 0 < \delta < 1, \quad v_{t+1} \sim N(0, \sigma_v^2). \quad (4)$$

Given these dynamics of g_t , the Appendix shows that

$$\sum_{j=0}^{\infty} \rho^j E_t(d_{t+1+j}) = \frac{d_t}{1-\rho} + \frac{\gamma}{(1-\delta)(1-\rho)^2} - \frac{\gamma\delta}{(1-\delta)(1-\rho)(1-\rho\delta)} + \frac{\delta g_t}{(1-\rho)(1-\rho\delta)}. \quad (5)$$

Substituting this equation into equation (3), we obtain

$$p_t = \frac{k}{1-\rho} + d_t + \frac{\gamma}{(1-\delta)(1-\rho)} - \frac{\gamma\delta}{(1-\delta)(1-\rho\delta)} + g_t \frac{\delta}{1-\rho\delta} - \frac{r_{e,t}}{1-\rho}, \quad (6)$$

which can be rearranged into

$$r_{e,t} = k + \frac{\gamma}{1-\delta} + (d_t - p_t)(1-\rho) + \left(g_t - \frac{\gamma}{1-\delta}\right) \frac{\delta(1-\rho)}{1-\rho\delta}. \quad (7)$$

The ICC, $r_{e,t}$, is a simple linear function of the log dividend–price ratio, $d_t - p_t$, and log dividend growth, g_t . (Note some similarity with the well-known constant-parameter Gordon growth model, in which $P = D/(r - g)$, and thus $r = D/P + g$.)

Further insight into the ICC can be obtained by assuming that the conditional expected return, $\mu_t \equiv E_t(r_{t+1})$, also follows a stationary AR(1) process:⁷

$$\mu_{t+1} = \alpha + \beta\mu_t + u_{t+1}, \quad 0 < \beta < 1, \quad u_{t+1} \sim N(0, \sigma_u^2). \quad (8)$$

Under this assumption, the Appendix shows that

$$\sum_{j=0}^{\infty} \rho^j E_t(r_{t+1+j}) = \frac{\alpha}{(1-\beta)(1-\rho)} + \left(\mu_t - \frac{\alpha}{1-\beta}\right) \frac{1}{1-\rho\beta}. \quad (9)$$

Plugging equations (5) and (9) into equation (2), we obtain

$$p_t = \frac{k}{1-\rho} + \frac{\gamma}{(1-\delta)(1-\rho)} - \frac{\alpha}{(1-\beta)(1-\rho)} + d_t + \left(g_t - \frac{\gamma}{1-\delta}\right) \frac{\delta}{(1-\rho\delta)} - \left(\mu_t - \frac{\alpha}{1-\beta}\right) \frac{1}{1-\rho\beta}. \quad (10)$$

The log stock price p_t is a simple function of d_t , g_t , and μ_t . The stock price increases with dividends d_t and dividend growth g_t , and it decreases with expected return μ_t . Note that p_t depends on the deviations of μ_t and g_t from their unconditional means of $\alpha/(1-\beta)$ and $\gamma/(1-\delta)$, respectively. Comparing equations (7) and (10), we have

$$r_{e,t} = \frac{\alpha}{1-\beta} + \left(\mu_t - \frac{\alpha}{1-\beta}\right) \frac{1-\rho}{1-\rho\beta}, \quad (11)$$

⁷ The assumptions that μ_t and g_t follow AR processes seem plausible as both returns and dividend growth are somewhat predictable. These assumptions are commonly made (e.g., Campbell and Shiller, 1988).

which implies that $r_{e,t}$ and μ_t are perfectly correlated. Thus, the ICC is a perfect proxy for the conditional expected return in the time series in an AR(1) framework.

We also consider a modified version of the ICC,

$$r_{e2,t} = k + \frac{\gamma}{1 - \delta} + (d_t - p_t)(1 - \rho). \quad (12)$$

This expression is obtained from equation (7) by setting g_t equal to its unconditional mean of $\gamma/(1 - \delta)$. This definition of $r_{e2,t}$ captures the idea that information about dividend growth is often limited in practice. Note that $r_{e2,t}$ is perfectly correlated with the dividend–price ratio, which is often used to proxy for the expected return. Since dividends tend to vary less than prices, the time variation in $r_{e2,t}$ is driven mostly by the variation in p_t .

II. Simulation

This section builds on the framework developed in Section I. First, we make additional assumptions about the conditional variance of stock returns. We impose a positive relation between the conditional mean and variance, and then we analyze the ability of various proxies for the conditional mean to detect this relation in simulated data. We find that the proxy proposed in this paper, the ICC, is very good at detecting the intertemporal risk–return tradeoff, even in the absence of conditional information about future cash flow.

A. The Variance of Stock Returns

Let $\sigma_t^2 \equiv \text{Var}_t(r_{t+1})$ denote the conditional variance of stock returns. We assume that the conditional variance is related to the conditional mean as follows:

$$\mu_t = a + b\sigma_t^2 + e_t, \quad b > 0, \quad e_t \sim N(0, \sigma_e^2)1_{[e_t \leq \bar{e}_t]}, \quad (13)$$

so that e_t is drawn from a truncated normal distribution with an upper bound of $\bar{e}_t$. The truncation of e_t ensures nonnegativity of the variance draws, as explained below. We assume the risk-free rate of zero, so that μ_t can also be thought of as the expected excess return.

Equation (13) defines the process for σ_t^2 , conditional on μ_t : $\sigma_t^2 = (\mu_t - a - e_t)/b$. In the absence of the truncation of e_t , σ_t^2 would follow an AR(1) process with an autoregressive parameter equal to β . In the presence of the truncation, σ_t^2 follows a process that is approximately autoregressive. Note that the strength of the mean–variance association in equation (13) can be measured as $\sigma_u^2/(\sigma_u^2 + \sigma_e^2)$, which is approximately equal to the fraction of the conditional variance of σ_t^2 that can be explained by the conditional variance of μ_t .

We show in the Appendix that the return variance in Section I can be approximated by

$$\text{Var}_t(r_{t+1}) = \frac{1}{(1 - \rho\delta)^2} \sigma_v^2 + \frac{\rho^2}{(1 - \rho\beta)^2} \sigma_u^2. \quad (14)$$

This expression is detached from the process for σ_t^2 defined in equation (13). To ensure that σ_t^2 can be interpreted as the variance of stock returns, we make σ_v^2 from equation (4) vary over time in a way that equates σ_t^2 from equation (13) to $\text{Var}_t(r_{t+1})$ from equation (14):

$$\sigma_{v,t+1}^2 = (1 - \rho\delta)^2 \left(\sigma_t^2 - \frac{\rho^2}{(1 - \rho\beta)^2} \sigma_u^2 \right). \quad (15)$$

Since $\sigma_{v,t+1}^2$ must be nonnegative, $\sigma_t^2 \geq \bar{\sigma}^2$ must hold in each period, where $\bar{\sigma}^2 = \frac{\rho^2}{(1 - \rho\beta)^2} \sigma_u^2$. To ensure that this inequality holds for each draw of σ_t^2 , we truncate the distribution of e_t in equation (13) at $\bar{e}_t = \mu_t - a - b\bar{\sigma}^2$.

The first term in equation (14) captures the return variance that is due to news about dividend growth. The second term captures the variance due to news about expected future returns. The fraction of the return variance that is explained by the variation in the expected return is therefore given by $\phi_t = \rho^2 \sigma_u^2 / ((1 - \rho\beta)^2 \sigma_t^2)$. Replacing σ_t^2 by its unconditional mean yields an unconditional value of ϕ_t , which we denote by ϕ .

B. The Simulation Procedure

In this subsection, we describe how we simulate the time series of μ_t , σ_t^2 , r_t , $r_{e,t}$, and $r_{e2,t}$, and how we use these time series to analyze the intertemporal risk–return relation.

The parameters are as follows. In equation (8), we choose $\alpha = 0.25\%$ per month and $\beta = 0.8$, which implies the unconditional mean return of 15% per year. In equation (13), we choose $a = 0.5\%$ per month and $b = 2.78$, so the unconditional return variance is $(18\%)^2$ per year. In equation (4), we choose $\gamma = 0.16\%$ per month and $\delta = 0.8$, so the unconditional mean of g_t is 10% per year. We solve for ρ and k numerically and obtain $\rho = 0.9955$ and $k = 0.0291$. In equation (8), σ_u takes five different values (0.34, 0.58, 0.75, 0.89, 1.01)% per month, selected so that the fraction of the return variance that can be explained by the variation in the expected return, earlier denoted by ϕ , takes the values of (0.1, 0.3, 0.5, 0.7, 0.9). For each value of σ_u , the value of σ_e in equation (13) is chosen so that the strength of the mean–variance link (i.e., $\sigma_u^2 / (\sigma_u^2 + \sigma_e^2)$) takes the values of (0.1, 0.3, 0.5, 0.7, 0.9).

The variables g_0 , μ_0 , and σ_0 are initialized at their unconditional values, $d_0 = 0$, and the initial price is computed from equation (10) as $p_0 = f_1(g_0, \mu_0, d_0)$. The following process is repeated in each period t , $t = 1, \dots, T$, conditional on the information up to time $t - 1$:

- (1) Compute $\sigma_{v,t}$ from equation (15).
- (2) Draw g_t from equation (4).
- (3) Construct $d_t = d_{t-1} + g_t$.
- (4) Draw μ_t from equation (8).
- (5) Compute the price, $p_t = f_1(g_t, \mu_t, d_t)$, from equation (10).

- (6) Compute the ICC, $r_{e,t} = f_2(g_t, p_t, d_t)$, from equation (7). Also compute the modified ICC, $r_{e2,t} = f_3(p_t, d_t)$, from equation (12).
- (7) Draw σ_t from equation (13).
- (8) Compute the realized return as $r_t = \log((P_t + D_t)/P_{t-1})$.

This process generates the time series of all variables used in the following subsection.

C. The Simulation Results

In this subsection, we use the time series simulated in Section II.B to estimate the intertemporal relation between the conditional mean and variance of returns. We consider the regression

$$\mu_t = c + d\sigma_t^2 + \epsilon_t, \quad (16)$$

with three proxies for μ_t : $r_{e,t}$, $r_{e2,t}$, and r_{t+1} . The realized return, r_{t+1} , is a common proxy for μ_t in this literature. We examine the performance of $r_{e,t}$ and $r_{e2,t}$ relative to r_{t+1} in detecting $d > 0$, which is imposed in the simulation via $b > 0$ in equation (13).

We consider five sample sizes: $T = 60, 120, 240, 360$, and 600 months. For each sample size, we simulate 5,000 time series of $r_{e,t}$, $r_{e2,t}$, r_{t+1} , μ_t , and σ_t^2 . For each time series, we run the regression (16) and record the estimated slope coefficient $\hat{d}$. We take the average of the 5,000 $\hat{d}$'s to be the true value of d , given the large number of simulations. The “ t -statistic” is computed by dividing the average $\hat{d}$ by the standard deviation of the 5,000 $\hat{d}$'s. In the same manner, we compute the true correlations between σ_t^2 and the three proxies for μ_t .

Table I reports the correlations and their t -statistics.⁸ As the strength of the mean–variance link increases (i.e., as we move from the left to the right in the table), all correlations increase, along with their significance. As T increases (i.e., as we move down the table), the correlations remain about the same, but their significance increases. Neither result is surprising: It is easier to detect a stronger mean–variance link, especially in large samples.

Table I shows a clear ranking among the three proxies for μ_t in terms of their ability to detect the positive risk–return relation. The highest and most significant correlations with σ_t^2 are achieved by $r_{e,t}$ and the lowest by r_{t+1} . This ranking is the same for all five values of T , all five degrees of the mean–variance link, and all five values of ϕ . For example, for $T = 240$ and the 0.5 values of the mean–variance link and ϕ , the correlations achieved by $r_{e,t}$, $r_{e2,t}$, and r_{t+1} are 0.74 ($t = 15.78$), 0.40 ($t = 4.33$), and 0.14 ($t = 1.69$), respectively.

The best performance of $r_{e,t}$ is not surprising, given the perfect correlation between $r_{e,t}$ and μ_t (equation (11)). More interesting is that r_{t+1} is uniformly outperformed not only by $r_{e,t}$ but also by $r_{e2,t}$, the ICC in which g_t is replaced by its unconditional mean. In practice, we (and the equity analysts whose forecasts we use in empirical work) often have little information about future cash flow.

⁸ To fit on one page, the table reports results for only three out of five values for each parameter.

Table I
Simulation Evidence: Correlations between Return Variance
and Proxies for Expected Return

This table reports the time-series correlations between the return variance σ_t^2 and three proxies for expected return μ_t : the implied cost of capital (r_{et}), the implied cost of capital with unknown conditional expected cash flow (r_{e2t}), and realized return r_{t+1} . Each correlation is computed by averaging the estimated correlations across 5,000 simulations. The t -statistics, reported in parentheses, are computed by dividing the corresponding average correlation by the standard deviation of the 5,000 correlations. The degree of the mean–variance link is the fraction of the conditional variance of σ_t^2 that can be explained by the conditional variance of μ_t , or $\sigma_u^2/(\sigma_u^2 + \sigma_e^2)$. The variable ϕ denotes the average fraction of the return variance that can be explained by the variation in expected returns, or the unconditional mean of $\rho^2 \sigma_u^2 / ((1 - \rho\beta)^2 \sigma_t^2)$. The length of the sample period over which the correlations are computed is denoted by T .

ϕ	Degree of Mean–Variance Link								
	0.1			0.5			0.9		
	r_{et}	r_{e2t}	r_{t+1}	r_{et}	r_{e2t}	r_{t+1}	r_{et}	r_{e2t}	r_{t+1}
<i>T</i> = 60 months									
0.1	0.34 (2.70)	0.11 (0.71)	0.03 (0.17)	0.78 (11.25)	0.28 (1.39)	0.08 (0.57)	0.97 (78.49)	0.32 (1.47)	0.11 (0.82)
0.5	0.31 (2.40)	0.15 (1.03)	0.04 (0.26)	0.70 (7.36)	0.39 (2.35)	0.14 (0.87)	0.95 (46.40)	0.47 (2.58)	0.21 (1.42)
0.9	0.29 (2.30)	0.16 (1.15)	0.04 (0.27)	0.68 (6.44)	0.42 (2.73)	0.15 (0.99)	0.94 (39.97)	0.53 (3.18)	0.24 (1.69)
<i>T</i> = 240 months									
0.1	0.37 (5.80)	0.11 (1.48)	0.03 (0.40)	0.81 (24.86)	0.29 (2.61)	0.08 (1.12)	0.97 (175.12)	0.32 (2.69)	0.11 (1.49)
0.5	0.33 (5.08)	0.16 (2.15)	0.04 (0.54)	0.74 (15.78)	0.40 (4.33)	0.14 (1.69)	0.96 (101.50)	0.46 (4.55)	0.20 (2.56)
0.9	0.32 (4.90)	0.17 (2.37)	0.05 (0.59)	0.71 (13.88)	0.43 (4.99)	0.15 (1.92)	0.95 (86.47)	0.51 (5.54)	0.23 (3.05)
<i>T</i> = 600 months									
0.1	0.37 (9.26)	0.11 (2.24)	0.03 (0.63)	0.81 (40.35)	0.28 (3.89)	0.08 (1.78)	0.97 (286.92)	0.32 (4.01)	0.11 (2.37)
0.5	0.33 (8.21)	0.16 (3.32)	0.04 (0.87)	0.74 (25.85)	0.40 (6.55)	0.13 (2.66)	0.96 (169.42)	0.46 (6.78)	0.19 (4.02)
0.9	0.32 (7.81)	0.18 (3.73)	0.05 (0.95)	0.72 (22.38)	0.44 (7.54)	0.15 (3.04)	0.95 (142.43)	0.51 (8.23)	0.23 (4.81)

Our results show that the ICC can help us estimate the intertemporal risk–return relation even in the absence of the cash flow information.

Since the time variation in $r_{e2,t}$ is driven mostly by p_t , the ability of $r_{e2,t}$ to detect the positive risk–return tradeoff stems from the fact that price changes tend to be accompanied by variance changes in the opposite direction. As long as price changes are driven to some extent by changes in μ_t (i.e., $\phi > 0$) and μ_t is positively related to σ_t^2 , the regression of $r_{e2,t}$ on σ_t^2 should detect the positive relation between μ_t and σ_t^2 in a long enough sample. Table I shows that $r_{e2,t}$ works better as ϕ increases, which is not surprising. More important, $r_{e2,t}$ works

well even for relatively low values of ϕ and relatively small sample sizes. For example, for $\phi = 0.3$, $T = 120$, and the mean–variance link of 0.5, the estimated correlation between $r_{e2,t}$ and σ_t^2 is 0.36 ($t = 2.66$). The empirical estimates of ϕ are generally in the neighborhood of 0.7 (e.g., Campbell and Ammer, 1993). For $\phi = 0.7$, $r_{e2,t}$ has a significant correlation with σ_t^2 even for T as low as 60 months and the mean–variance link as low as 0.3. These results show that even $r_{e2,t}$ can be quite useful in estimating the risk–return tradeoff.

In contrast, r_{t+1} performs poorly. Its correlation with σ_t^2 is never significant for $T \leq 60$ months, even when the mean–variance link is 0.9. When the link is 0.5, we need at least a 30-year-long sample to find a significant relation between r_{t+1} and σ_t^2 , and when the link is 0.3, we need a 50-year-long sample. Realized returns seem too noisy to be very useful as proxies for expected returns, consistent with Lundblad (2007). Proxying for expected returns by the ICC allows us to detect a positive risk–return tradeoff in substantially shorter samples than would be required if we used realized returns.⁹ This fact seems useful especially in international markets, in which long return histories are often unavailable.

III. Empirical Methodology

A. The Methodology for Computing Implied Cost of Capital

We compute the implied cost of equity capital for each firm as the internal rate of return that equates the present value of future dividends to the current stock price, following the approach of Gebhardt, Lee, and Swaminathan (2001) and Lee, Ng, and Swaminathan (2007). We use the term “dividends” quite generally to describe the free cash flow to equity (FCFE), which captures the total cash flow available to shareholders, net of any stock repurchases and new equity issues. The stock valuation formula in equation (1) expresses the stock price in terms of an infinite series, but we explicitly forecast FCFE only over a finite horizon and capture the free cash flow beyond the last explicit forecast period in a “terminal value” calculation. In other words, the value of a firm is computed in two parts, as the present value of FCFE up to the terminal period $t + T$ plus the present value of FCFE beyond the terminal period. We compute future FCFE up to year $t + T + 1$ as the product of annual earnings forecasts and one minus the plowback rate:

$$E_t(FCFE_{t+k}) = FE_{t+k} \times (1 - b_{t+k}), \quad (17)$$

where FE_{t+k} and b_{t+k} are, respectively, the forecasts of earnings and the plowback rate for year $t + k$.¹⁰ The plowback rate is the fraction of earnings that is

⁹ We have also estimated the regression (16) in first differences and obtained results that lead to exactly the same conclusions. Exactly the same conclusions are also reached based on the slope coefficients $\hat{d}$ rather than on correlations. The results are robust to reasonably large changes in the parameter specification.

¹⁰ Define net new equity investment as capital expenditures plus change in working capital minus depreciation and amortization minus net new issues of debt. FCFE is net income minus net new equity investment. Defining the plowback rate as net new equity investment over net income, we obtain equation (17).

reinvested in the firm, or one minus the payout ratio. The earnings forecasts for years $t + 1$ through $t + 3$ are based on analyst forecasts, and the forecasts from year $t + 4$ to year $t + T + 1$ are computed by mean-reverting the year $t + 3$ earnings growth rate to its steady-state value by year $t + T + 2$. We assume the steady-state growth rate starting in year $t + T + 2$ is equal to the long-run nominal GDP growth rate, g , computed as the sum of the long-run real GDP growth rate (a rolling average of annual real GDP growth) and the long-run average rate of inflation based on the implicit GDP deflator (more details are provided below).

The assumption that each firm's steady-state growth rate equals the GDP growth rate is imperfect because the shares of firms in the aggregate economy can change over the long run. Alas, it seems difficult to determine ex ante which firms will grow faster or slower than GDP in the long run. However, the steady-state assumption is correct on average, and that is all we need. We work only with market-wide averages of the ICC across firms, so any potential bias in the firm-level ICC should approximately wash out.

We impose an exponential rate of decline to mean-revert the year $t + 3$ growth rate to the steady-state growth rate.¹¹ Specifically, we compute earnings growth rates and earnings forecasts for years $t + 4$ to $t + T + 1$ ($k = 4, \dots, T + 1$) as follows:

$$g_{t+k} = g_{t+k-1} \times \exp[\log(g/g_{t+3})/(T-1)], \quad (18)$$

$$FE_{t+k} = FE_{t+k-1} \times (1 + g_{t+k}). \quad (19)$$

We forecast plowback rates in two stages: (1) we explicitly forecast plowback rates for years $t + 1$ and $t + 2$ (see the next section), and (2) we mean-revert the plowback rates between years $t + 2$ and $t + T + 1$ linearly to a steady-state value computed from the sustainable growth rate formula.¹² This formula assumes that, in the steady-state, the product of the steady-state return on new investments, ROI , and the steady-state plowback rate is equal to the steady-state growth rate in earnings (see Brealey and Myers (2002)); that is, $g = ROI \times b$. We then set $ROI = r_e$ for new investments in the steady state, assuming that competition drives returns on these investments down to the cost of equity. Thus, our main assumptions are that the earnings growth rate reverts to the long-run nominal GDP growth rate, and that the return on new investment, ROI , reverts to the (implied) cost of equity, r_e .

Substituting $ROI = r_e$ in the sustainable growth rate formula and solving for b provides the steady-state value for the plowback rate, $b = g/r_e$. The intermediate plowback rates from $t + 3$ to $t + T$ ($k = 3, \dots, T$) are computed as follows:

¹¹ We choose the exponential rate of decline to be consistent with the empirical evidence that growth rates of earnings mean-revert rapidly (e.g., Chan, Karceski, and Lakonishok (2003)). Given this rapid mean reversion, any potential biases in analysts' short-term earnings forecasts should not have large effects on the long-run growth rates, and therefore also on our estimates of the ICC.

¹² We assume that year $t + k$ plowback affects year $t + k + 1$ earnings growth. We assume a linear decline in the plowback rate because plowback rates appear to mean-revert slower than earnings growth rates.

$$b_{t+k} = b_{t+k-1} - \frac{b_{t+2} - b}{T - 1}. \quad (20)$$

The terminal value at time $t + T$, TV_{t+T} , is computed as the present value of a perpetuity equal to the ratio of the year $t + T + 1$ earnings forecast divided by the cost of equity:

$$TV_{t+T} = \frac{FE_{t+T+1}}{r_e}, \quad (21)$$

where FE_{t+T+1} is the earnings forecast for year $t + T + 1$. Note that the use of the no-growth perpetuity formula does not imply that earnings or cash flows do not grow after period $t + T$. Rather, it simply means that any new investments after year $t + T$ earn zero economic profits. In other words, any growth in earnings or cash flow after year T is value irrelevant.

Substituting equations (17) to (21) into the infinite-horizon, free cash flow valuation model in equation (1) provides the following empirically tractable finite-horizon model:

$$P_t = \sum_{k=1}^T \frac{FE_{t+k}(1 - b_{t+k})}{(1 + r_e)^k} + \frac{FE_{t+T+1}}{r_e(1 + r_e)^T}. \quad (22)$$

We use a 15-year horizon ($T = 15$), following Lee, Ng, and Swaminathan (2007).

A.1 Earnings Forecasts over the First 3 Years

We obtain earnings forecasts for years $t + 1$ and $t + 2$ from the I/B/E/S database. I/B/E/S analysts forecast 1- and 2-year-ahead earnings per share (EPS) for each firm as well as the long-term earnings growth rate (Ltg). We use the consensus (mean) 1- and 2-year-ahead EPS forecasts (FE_{t+1} and FE_{t+2}), and we compute a 3-year-ahead earnings forecast as $FE_{t+3} = FE_{t+2}(1 + Ltg)$.¹³ Firms with growth rates above 100% (below 2%) are assigned values of 100% (2%).

A.2 Plowback Rates

For each U.S. firm, we compute the plowback rate (b_t) for the first 3 years as one minus the firm's most recent net payout ratio (p_t). To compute p_t , we first compute net payout (NP_t) as gross payout (i.e., dividends plus share repurchases) minus any issuance of new stock: $NP_t = D_t + REP_t - NE_t$, where D_t is

¹³ If only a subset of these forecasts are available, we try to infer the missing forecasts from the available ones. For example, if Ltg is not available, we estimate it from the growth rate implied by the consensus forecasts in years 1 and 2, $g_{t+3} = FE_{t+2}/FE_{t+1} - 1$. If FE_{t+2} is not available but FE_{t+1} and Ltg are, we compute $FE_{t+2} = FE_{t+1}(1 + Ltg)$. If neither Ltg nor FE_{t+2} are available, we compute Ltg from the ratio of FE_{t+1} to the most recent realized earnings and then compute $FE_{t+2} = FE_{t+1}(1 + Ltg)$ again.

the amount of common dividends paid by the firm in year t (COMPUSTAT item D21), REP_t is the amount of common and preferred stock purchased by the firm in year t (item D115), and NE_t is the amount of common and preferred stock sold by the firm in year t (item D108). We then compute the net payout ratio, p_t , as NP_t/NI_t , where NI_t is the firm's net income in year t (item D18).¹⁴ To ensure that our computations are based on publicly available information, we require the fiscal year-end to be at least 3 months prior to the date of computation of the cost of equity.

For the other G-7 countries, we use a simpler approach to estimate the payout ratio due to data limitations. If dividends and positive earnings are available for the prior fiscal year, we use the dividend payout ratio. For firms with negative earnings, we divide dividends by typical long-run earnings, estimated to be 6% of total assets. The long-run return on assets in the United States is 6%. See also Gebhardt, Lee, and Swaminathan (2001).

Given our forecasts of earnings and plowback rates, we compute the ICC as r_e from equation (22) for each firm at each month-end. To trim the outliers, we delete the top 0.5% and the bottom 0.5% of the ICC values in each month. We then compute the country-level ICC as an equal-weighted or value-weighted average of the individual firms' ICCs. The value weights are based on market values at the most recent year-end. Finally, we compute the implied risk premium for each G-7 country as the ICC minus the local risk-free rate.

B. Data

We obtain return data from CRSP (for U.S. firms) and Datastream (for non-U.S. firms), accounting data from COMPUSTAT (U.S.) and Worldscope (non-U.S.), and analyst forecasts from I/B/E/S (for both U.S. and non-U.S. firms). To ensure a reasonable number of firms in each country, we limit our analysis to the period of January 1981 to December 2002 for the United States, and January 1990 to December 2002 for the other six countries.

We require non-U.S. firms to have monthly price and share outstanding numbers available in I/B/E/S. For U.S. firms, monthly data on market capitalization are obtained from CRSP. We require the availability of the following data items: common dividend, net income, book value of common equity, fiscal year-end date, and currency denomination. These items come from the most recent fiscal year ending at least 6 months (3 months in the case of the United States) prior to the month in which the cost of capital is computed. As discussed above,

¹⁴ The NP_t/NI_t ratio is available for 65% of firms. If this ratio is not available, we compute p_t as NP_t/FE_t , where FE_t is the earnings forecast from I/B/E/S as of December of the previous year for the fiscal year ending in year t . The NP_t/FE_t ratio is used for 6% of firms. For the remaining 29% of firms for which even that ratio is unavailable, p_t is computed as the median NP_t/NI_t across all firms in the corresponding industry-size portfolio. The industry-size portfolios are formed each year by first sorting firms into 48 industries based on the Fama–French classification and then forming three equal-number-of-firms portfolios based on market cap within each industry. If p_t is above one or below -0.5 , we set it equal to the median NP_t/NI_t of the industry-size portfolio. Industry-size portfolios with a median NP_t/NI_t below -0.5 are given a value of -0.5 .

for U.S. firms, we also require data on share repurchases and new stock issuance to compute the net payout ratio. We exclude ADRs, closed-end funds, REITs, and firms with negative common equity. We use I/B/E/S to obtain monthly data on 1-year and 2-year consensus EPS forecasts and estimates of the long-term growth rate, all in local currency.

To measure market returns, we use monthly returns on the CRSP value-weighted index for the United States and monthly local-currency returns on the MSCI index for the other countries. Data on nominal GDP growth rates are obtained from the Bureau of Economic Analysis and the World Bank. Each year, we compute the “steady-state” GDP growth rate as the historical average of the GDP growth rates using annual data up to that year. For the United States, our GDP data begin in 1930. For France, Italy, Japan, and the United Kingdom, GDP growth rates begin in 1961. For Canada and Germany, these data begin in 1966 and 1972, respectively.

For non-U.S. firms, I/B/E/S reports analyst forecasts, price, and shares outstanding within a few days after the 15th of each month. Therefore, we compute the ICC for non-U.S. firms as of mid-month. For consistency, we compute monthly returns from the first trading day after the 15th of the previous month to the first trading day after the 15th of the current month. Each month, we also estimate the conditional variance and standard deviation of market returns using mid-month to mid-month daily returns. For U.S. firms, we obtain month-end price data from CRSP, and compute monthly returns and volatilities from the beginning to the end of the month.

For each country, we compute the implied risk premium as the ICC minus the yield to maturity on the local 10-year government bond (obtained from Datastream).¹⁵ The only exception is Italy, for which we use the 7-year bond because the 10-year bond data begin later. For the United States, we use month-end bond yields since we compute the month-end ICC. For the other countries, we use mid-month yields to match the timing of the ICC estimates.

To compute realized excess returns, we subtract the local 1-month risk-free rate from realized returns. For the risk-free rate, we use monthly returns on a 1-month Treasury bill for the U.S., Canada, and the United Kingdom. Data on U.S. T-bill rates are obtained from Kenneth French’s website and the T-bill rates for Canada and the United Kingdom are obtained from Datastream. For the other four countries, the T-bill rates are not available for the full sample period, so we use the interbank 1-month offer rates provided by the British Bankers Association (BBA), also obtained from Datastream.¹⁶ Datastream provides two

¹⁵ Claus and Thomas (2001) also use the 10-year risk-free rate to construct the ICC-based risk premium.

¹⁶ For Italy, both short-term and long-term risk-free rates are available starting from July 1990. For this reason, all regressions for Italy start in July 1990. We also compare the interbank rates with the 1-month T-bill rates in countries where both rates are available. In Canada, the average spread of the interbank rate over the T-bill rate is 0.23% annualized over the period 1990 to 2003. In the United Kingdom, the spread is about 0.24%. We also compare the interbank rates to the euro-currency rates and find only marginal differences. For example, in the United Kingdom, the interbank rates are about 0.01% higher than the euro-pound rates.

Table II
Summary Statistics

The table contains the summary statistics for the monthly time series of the implied risk premia and the return volatility for each of the G-7 countries. The statistics are computed over the period 1981 to 2002 for the United States and over 1990 to 2002 for the other countries (except for Italy, where the sample begins in July 1990). The implied risk premium is computed as the implied cost of equity capital minus the yield on the country's 10-year government bond (except for Italy, where the 7-year bond is used). The firm-level premia are equal-weighted (Panel A) or value-weighted (Panel B) across all firms in the given country. Standard deviation of returns is the annualized standard deviation of the daily market returns computed over the previous month, where the "market" is the CRSP value-weighted index for the United States and the MSCI index for the other countries. All quantities are reported in percent per year. At the bottom, we report the average across months of the number of firms for which the implied cost of capital is available.

	CAN	FRA	GER	ITA	JAP	UK	USA
Panel A: Equal-Weighted Implied Risk Premia							
Mean	8.15	5.69	4.79	4.21	5.41	6.99	4.57
Std deviation	2.49	2.28	4.63	4.25	5.57	2.43	1.38
Minimum	1.93	0.48	−4.68	−5.62	−3.91	2.00	0.92
Maximum	12.78	10.64	16.39	12.57	15.18	12.67	7.76
Panel B: Value-Weighted Implied Risk Premia							
Mean	4.73	2.86	1.56	0.61	3.12	2.89	2.65
Std deviation	1.58	0.99	1.90	3.56	3.29	1.13	0.97
Minimum	0.25	0.24	−4.68	−7.56	−3.36	−0.67	−0.18
Maximum	9.21	5.46	7.25	6.70	9.09	6.08	6.17
Panel C: Standard Deviation of Returns							
Mean	13.69	18.76	19.41	20.84	19.44	14.80	13.72
Std deviation	7.65	8.27	10.47	7.78	8.18	6.86	7.60
Minimum	5.67	6.36	5.72	6.94	7.88	6.40	4.49
Maximum	42.05	54.39	62.31	49.95	51.86	47.59	78.33
Number of firms	275	308	279	115	960	787	1,795

series on interbank offer rates, one provided by BBA and another that originates within the country. We use the former since there is a longer time series of data available for the BBA series in most countries. The rates on the two series are very similar for most countries except for Japan, where the rate provided by BBA is 0.03% below the local interbank rate. We use the BBA series for Japan since the data go back to 1989, whereas the data on local rates start only in 1995.

Table II provides summary statistics on the implied risk premia and return volatilities (annualized monthly standard deviations computed from daily returns) for the G-7 countries. The average equal-weighted risk premium varies from 4.2% in Italy to 8.2% in Canada. The value-weighted averages are smaller, ranging from 0.6% in Italy to 4.7% in Canada. These estimates are similar to those found in Lee, Ng, and Swaminathan (2007). The average standard




Figure 1. USA: Implied market risk premium and market return volatility. Panel A plots the monthly time series of the implied market risk premium for the United States, computed as the difference between the implied cost of capital and the yield to maturity on the 10-year Treasury bond. The implied cost of capital is computed as an equal-weighted (dashed line) or value-weighted (solid line) average of the implied costs of capital across all U.S. firms. Panel B plots the monthly time series of the realized market return volatility, computed from daily data within the month.

deviation of returns varies from 13.7% in the United States and Canada to 20.8% in Italy. The table also provides the average number of firms per month in each country. The United States has the highest average number of firms (1,795) and Italy has the lowest (115).

Figures 1 and 2 plot the monthly time series of the implied risk premiums in all seven countries. Figure 3 reports the country return volatilities. The equal-weighted U.S. risk premium in Figure 1 fluctuates between 1% and 8%




Figure 2. Implied market risk premium. Each panel plots the monthly time series of the implied market risk premium, computed as the difference between the implied cost of capital and the yield to maturity on the 10-year local government bond. The implied cost of capital is computed as an equal-weighted (dashed line) or value-weighted (solid line) average of the implied costs of capital across all firms in the given country.

from 1981 to 2002, with most of the values falling in the 4% to 6% range. The value-weighted U.S. premium fluctuates between zero and 6%, but mostly between 2% and 4%, consistent with Claus and Thomas (2001). The largest changes in the premium tend to occur in months with large absolute stock returns. In most countries, the implied risk premium rises in the 1990s. This rise is due in part to the increasing cash flow expectations in the 1990s and in part to the declining risk-free rates. When the risk-free rates are added back to plot the ICC, the upward trend remains apparent only for Germany and Japan, and the ICC in the United States exhibits a clear decline.




Figure 3. Market return volatility. Each panel plots the monthly time series of the realized market return volatility in the given country, computed as the annualized standard deviation of daily returns within the month.

Several studies find that analyst forecasts tend to be systematically biased upward. Given this bias, the true risk premia may well be lower than those reported in Figures 1 and 2. However, since we are interested in the time variation in the risk premia, the bias has no effect on our results if it is constant over time. Even if the bias varies over time, it has no effect on our results as long as its time variation is uncorrelated with market return volatility. In order to artificially create our results, the bias would have to be significantly positively correlated with market volatility. We do not find any correlation between analyst forecast errors and market volatility in our subsequent analysis in Section IV.E.

IV. Empirical Results

This section presents our empirical findings. For each G-7 country, we regress the conditional mean return on the conditional market volatility in various specifications. Since Merton’s ICAPM postulates a positive relation between the conditional mean and variance of market returns, variance seems to be more relevant than standard deviation as a measure of market volatility. Nonetheless, we consider not only variance (σ_t^2) but also standard deviation (σ_t), to assess the sensitivity of our results.

In most of our analysis, we ignore any potential hedging demands (Merton (1973)), as do Merton (1980) and others. However, in Section IV.D, we show that including popular proxies for hedging demands has little effect on our results for the U.S. market.

A. Volatility and Realized Returns

We begin by using the realized excess market return at time $t + 1$, r_{t+1} , as a proxy for the expected excess market return at time t . We regress this proxy on market volatility Vol_t (σ_t^2 or σ_t):

$$r_{t+1} = a + b Vol_t + e_{t+1}. \quad (23)$$

This regression is an empirical analogue of the simulated regression in equation (16).

The results are disappointing. There is no evidence of any relation between volatility and next period’s realized return. In monthly data covering the same time period as the rest of our analysis (1981 to 2002 for the United States and 1990 to 2002 for the other countries), the estimates of b are not significantly different from zero in any of the seven countries. In fact, in three countries, the estimates of b are negative. Across the seven countries and two Vol_t choices, the highest t -statistic is 1.17 and the highest adjusted- R^2 is 0.15%, confirming that this month’s volatility has very little predictive power for next month’s return.¹⁷ When we extend the data back to July 1926 for the United States, we find a positive but still insignificant estimate of b , with a t -statistic well below one. Using quarterly data leads to the same conclusions. We construct compounded quarterly excess returns by subtracting compounded quarterly T-bill returns from compounded quarterly market returns in all G-7 countries. We regress this quarter’s excess return on last quarter’s realized volatility, estimated using daily return data within the quarter. We find that most of the country-level slope coefficients are negative rather than positive, and no coefficient is significantly positive. All of these results are consistent with our simulation evidence that it is difficult to detect a mean–variance relation in tests that use realized returns to proxy for expected returns.

¹⁷ The full table of results is available in the NBER working paper version of this article. Given the insignificant nature of the results in this section, we do not tabulate them here, to save space.

B. Volatility and the Implied Risk Premia

Next, we consider three regression specifications with the implied risk premium μ_t :¹⁸

$$\mu_t = a + b \text{Vol}_t + e_t, \quad (24)$$

$$\Delta\mu_t = a + b\Delta\text{Vol}_t + e_t, \quad (25)$$

$$\epsilon_{\mu,t} = a + b\epsilon_{V,t} + e_t, \quad (26)$$

where μ_t is the implied risk premium at the end of month t , $\Delta\mu_t = \mu_t - \mu_{t-1}$, $\epsilon_{\mu,t}$ is the residual from an AR(1) model estimated for μ_t in the full sample, and $\epsilon_{V,t}$ is the analogous residual from an independent AR(1) process for volatility (σ_t^2 or σ_t). Regressions (25) and (26) examine the relations between proxies for shocks to volatility and shocks to expected returns. Tests involving shocks may be more powerful than tests involving levels because any persistent biases in the estimates of the conditional mean and volatility should not influence the monthly shocks. To correct the standard errors for potential autocorrelation, we use 12 Newey–West lags in regression (24) and 1 lag for the other specifications. We use more lags for regression (24) because μ_t is highly persistent.

Panel A of Table III presents the results in the case in which the country-level implied premium is an equal-weighted average of the individual firm premia. First, consider regression (24). Using σ_t to measure volatility, the risk–return relation is positive ($b > 0$) for all G-7 countries, and the relation is statistically significant in all countries but Italy and Japan. Using σ_t^2 for volatility, the risk–return relation is again significantly positive for five of the seven countries. In regressions (25) and (26), we find a statistically significant positive relation between shocks to the risk premia and shocks to volatility in Canada, France, Germany, the United Kingdom, and the United States. Only in Italy, the country with the lowest number of firms, does the slope coefficient have the wrong sign (statistically insignificant). We find it striking that our results are statistically significant in so many cases, despite the relatively short samples used in the estimation (22 years for the United States and 13 years for the other six countries).

Panel B of Table III is an equivalent of Panel A, with the equal-weighted country risk premia replaced by the value-weighted ones. As in Panel A, regression (24) uncovers a positive mean–variance relation. This relation is statistically significant in France, Germany, Italy, and the United Kingdom, and it is insignificantly positive in Canada and the United States. The regressions based on shocks find a significantly positive relation in France, the United Kingdom, and the United States. The value-weighted evidence is somewhat weaker than the equal-weighted evidence. Across the seven countries, the correlations

¹⁸ Two additional specifications are considered in the NBER working paper version. The results from the two specifications are very similar to those presented here.

Table III
Country-Level Implied Premium Regressed on Country Market Volatility

This table reports the slope coefficients from the regressions in equations (24) through (26) for each of the G-7 markets. The sample periods begin in January 1981 for the United States, in July 1990 for Italy, and in January 1990 for the remaining five countries. Each country's sample period ends in December 2002. In Panel A, the country premia μ_t are obtained as equal-weighted averages of the premia across all firms in that country. In Panel B, the premia are value-weighted. Return volatility σ_t is the realized market volatility from month $t - 1$, estimated from daily data on the given country's MSCI index (except for the United States where we use the CRSP value-weighted index) and annualized. The column labels σ_t^2 and σ_t denote the regressor. R^2 denotes adjusted- R^2 . The t -statistics are adjusted for residual autocorrelation and heteroskedasticity using the Newey–West correction.

	σ_t^2							σ_t						
	CAN	FRA	GER	ITA	JAP	UK	USA	CAN	FRA	GER	ITA	JAP	UK	USA
Panel A: Equal-Weighted Implied Risk Premia														
	Levels: $\mu_t = a + b \sigma_t^{(2)} + e_t$													
$\hat{b}$	0.36	0.24	0.38	0.23	−0.02	0.47	0.09	0.16	0.04	0.24	0.11	0.05	0.23	0.07
t	3.83	6.31	6.66	2.20	−0.13	7.23	1.78	3.98	8.39	5.04	1.74	0.63	14.06	3.18
R^2	0.19	0.22	0.24	0.04	−0.01	0.34	0.09	0.25	0.26	0.28	0.03	−0.00	0.40	0.16
	Differences: $\Delta\mu_t = a + b \Delta\sigma_t^{(2)} + e_t$													
$\hat{b}$	0.05	0.06	0.03	0.00	0.02	0.11	0.03	0.02	0.03	0.02	0.00	0.01	0.05	0.02
t	2.02	3.47	3.12	−0.17	1.88	5.11	2.38	1.96	2.88	2.41	−0.13	1.78	5.43	2.83
R^2	0.03	0.08	0.03	−0.01	0.01	0.18	0.08	0.02	0.06	0.02	−0.01	0.01	0.19	0.08
	AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{v,t} + e_t$													
$\hat{b}$	0.06	0.07	0.03	−0.01	0.01	0.12	0.05	0.03	0.03	0.02	−0.01	0.01	0.06	0.03
t	2.45	4.15	2.36	−0.37	0.97	4.52	10.53	2.31	3.10	2.01	−0.37	1.16	5.12	4.01
R^2	0.03	0.10	0.02	−0.01	−0.00	0.18	0.15	0.03	0.07	0.01	−0.01	−0.00	0.19	0.12
Panel B: Value-Weighted Implied Risk Premia														
	Levels: $\mu_t = a + b \sigma_t^{(2)} + e_t$													
$\hat{b}$	0.08	0.12	0.16	0.20	−0.04	0.20	0.04	0.04	0.07	0.09	0.09	0.01	0.09	0.03
t	1.01	9.17	5.26	2.40	−0.47	8.21	1.17	1.06	5.33	3.59	1.75	0.17	5.97	1.42
R^2	0.02	0.30	0.26	0.04	−0.00	0.27	0.02	0.03	0.30	0.26	0.03	−0.01	0.28	0.04
	Differences: $\Delta\mu_t = a + b \Delta\sigma_t^{(2)} + e_t$													
$\hat{b}$	0.03	0.03	0.01	0.00	0.02	0.07	0.02	0.01	0.01	0.01	0.00	0.01	0.04	0.01
t	1.09	2.90	0.66	0.17	1.81	2.13	1.83	1.15	2.23	0.65	0.20	1.62	2.71	1.77
R^2	0.00	0.04	−0.00	−0.01	0.01	0.05	0.03	0.00	0.03	−0.00	−0.01	0.01	0.08	0.03
	AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{v,t} + e_t$													
$\hat{b}$	0.02	0.04	0.02	−0.01	0.01	0.09	0.03	0.01	0.02	0.01	0.00	0.00	0.05	0.02
t	1.03	3.34	1.08	−0.21	0.64	2.95	6.99	1.14	2.51	1.06	−0.25	0.73	3.47	2.71
R^2	−0.00	0.06	0.01	−0.01	−0.00	0.08	0.07	0.00	0.05	0.01	−0.01	−0.00	0.10	0.05

between the premium and volatility range from 13% to 60% for the equal-weighted premium and from 10% to 45% for the value-weighted premium.¹⁹

Should we pay more attention to the results in Panel A, where the ICC is equal-weighted across firms in computing the country-level ICC, or to the results in Panel B, where the ICC is value-weighted? Equal-weighting typically pays disproportionate attention to small firms, but it would be misleading to argue that the results in Panel A are driven by small firms. The firms in our sample are a subset of firms in any given country and this is not a random

¹⁹ The full correlation table is available in the NBER working paper version of this article.

subset because firms that satisfy our data requirements (which include analyst forecasts) tend to be among the largest firms in their countries. As a result, value-weighting focuses on the largest among these already large firms, which overweights the largest firms relative to the country's market portfolio. Equal-weighting pays more attention to smaller firms in our large-firm subsets, which partly compensates for the absence of truly small firms in our sample. It is not clear whether value-weighting or equal-weighting produces an aggregate expected return that is closer to the expected return on the country's true market portfolio, so we consider both the panels of Table III informative.

The regressions in Table III are estimated separately for each individual country. To test if the estimated positive mean–variance relation is jointly significant across the G-7 countries, we estimate a multivariate seemingly unrelated regression (SUR) model involving all seven countries for each of the three regression specifications. A joint F -test of the hypothesis that all seven slope coefficients are equal to zero rejects the null for each specification.

Overall, the results in Table III show a positive relation between the conditional mean and volatility of the country-level market returns. These results confirm our simulation findings that a positive intertemporal mean–variance relation, if present, is easier to detect by using the ICC than by using the future realized return as a proxy for the expected return.

C. Robustness: Implied Volatility

So far, we have estimated the conditional return volatility by the volatility realized over the previous month. This approach involves nontrivial estimation error, which biases our results against finding a mean–variance relation. In this subsection, we consider an alternative volatility estimator: the implied volatility from the options market. Implied volatility data are available to us for the U.S. stock market over the period January 1986 through December 2002. We use the month-end series of the VXO index, which is based on the S&P 100 options. The data are obtained from the CBOE.²⁰

Panel A of Table IV contains the results from regressions (24) through (26). The estimated risk–return relation is clearly positive. For example, consider regression (25), in which first differences in the implied premium are regressed on first differences in implied volatility. Across the four specifications (σ_t^2 and σ_t , equal-weighted and value-weighted implied premium), the t -statistics for the slope coefficient range from 9.77 to 10.47. Based on the residuals in μ_t and $\sigma_t^{(2)}$ (regression (26)), the t -statistics range from 9.24 to 11.24. This level of statistical significance is striking, given the relatively short sample period. It appears that implied volatility contains less estimation error than realized volatility.

²⁰ The VXO index used to be known as the VIX index until CBOE modified the VIX methodology in September 2003 (it switched to the S&P 500 index options and changed the index formula).

Table IV
U.S. Implied Premium Regressed on U.S. Implied Volatility, with and without Hedging Demands

This table reports the slope coefficients from the regressions in equations (24) through (26) for the U.S. stock market between January 1986 through December 2002. Return volatility σ_t is the implied volatility from month t , measured as the VXO index. Panel A includes no proxies for hedging demands (H_t) in the estimation, whereas Panel B includes five such proxies: the excess return on the 30-year Treasury bond, the default spread, the term spread, the detrended risk-free rate, and the dividend–price ratio. The column label “EW” (VW) denotes an equal-weighted (value-weighted) implied risk premium, and σ_t^2 or σ_t denote the regressor. The t -statistics are adjusted for residual autocorrelation using the Newey–West correction.

	σ_t^2		σ_t	
	EW	VW	EW	VW
Panel A: No Hedging Demands Included				
Levels (eq. (24)): $\mu_t = a + b\sigma_t^{(2)} + e_t$				
$\hat{b}$	0.11	0.04	0.06	0.02
t	4.37	1.89	3.95	1.21
R^2	0.25	0.06	0.27	0.03
Differences (eq. (25)): $\Delta \mu_t = a + b\Delta \sigma_t^{(2)} + e_t$				
$\hat{b}$	0.09	0.07	0.06	0.05
t	10.47	9.77	10.17	9.84
R^2	0.35	0.32	0.34	0.32
AR(1) Residuals (eq. (26)): $\epsilon_{\mu,t} = a + b\epsilon_{V,t} + e_t$				
$\hat{b}$	0.09	0.08	0.06	0.05
t	11.24	9.64	10.37	9.24
R^2	0.38	0.31	0.35	0.29
Panel B: Hedging Demands Included				
Eq. (24) with Hedging Demands: $\mu_t = a + b\sigma_t^{(2)} + hH_t + e_t$				
$\hat{b}$	0.09	0.04	0.05	0.01
t	4.58	1.77	3.89	1.14
R^2	0.46	0.17	0.44	0.16
Eq. (25) with Hedging Demands: $\Delta \mu_t = a + b\Delta \sigma_t^{(2)} + h\Delta H_t + e_t$				
$\hat{b}$	0.03	0.02	0.02	0.02
t	4.21	3.42	4.22	3.80
R^2	0.68	0.65	0.69	0.65
Eq. (26) with Hedging Demands: $\epsilon_{\mu,t} = a + b\epsilon_{V,t} + h\epsilon_{H,t} + e_t$				
$\hat{b}$	0.08	0.06	0.05	0.04
t	7.47	5.97	6.56	5.53
R^2	0.41	0.35	0.38	0.34

D. Robustness: Hedging Demands

According to Merton (1973), the conditional expected excess market return depends not only on the conditional variance of market returns but also on hedging demands, i.e., on the market’s covariance with the state variables that

capture investment opportunities. Scruggs (1998) and Guo and Whitelaw (2006) argue that hedging demands are important in uncovering a positive mean–variance relation. Although we find this relation even without including hedging demands in our estimation, it seems useful to test whether the relation survives the inclusion of commonly used proxies for hedging demands.

We model hedging demands as a linear combination of five macroeconomic variables that have been used in prior studies. The first variable is the excess return on the 30-year U.S. Treasury bond, obtained from CRSP. This variable is motivated by Scruggs (1998), who uses long-term government bond excess returns as a catch-all proxy for hedging demands. The other four variables follow Guo and Whitelaw (2006): the default spread (Baa–Aaa yield spread, obtained from the St. Louis Fed), the term spread (30-year minus 1-month Treasury yield spread, obtained from CRSP), the detrended risk-free rate (the 1-month T-bill rate in excess of its 12-month moving average), and the dividend–price ratio (extracted from the value-weighted CRSP market return series with and without dividends).

We add all five variables to the right-hand side of each regression from Panel A of Table IV for the U.S. market, and report the results in Panel B. The inclusion of the hedging demand proxies has a relatively small effect on the estimated coefficients. In both panels of Table IV, the mean–variance relation is highly statistically significant in the same set of 10 out of 12 specifications. We conclude that the positive risk–return tradeoff in the United States is robust to controlling for popular proxies for hedging demands.

E. Robustness: Analyst Forecast Errors

The ICC is measured with error, in part because analyst forecasts may not perfectly capture the market's cash flow expectations. For example, if analysts are too optimistic, the ICC is too high.²¹ If the analysts' forecast errors are somehow positively related to market volatility, they could create a false appearance of a positive mean–variance relation. In this section, we conduct two tests of this conjecture.

In the first test, we replace the analysts' ex ante forecasts of earnings by the ex post realized earnings. This approach avoids any potential analyst bias but it introduces complications of its own because ex post earnings are not known to investors in advance. Since realized earnings equal expected earnings plus error, realized earnings are more volatile than expected earnings. As a result, using ex post earnings injects noise in the estimation, which is likely to reduce the statistical significance of our results.

We introduce the ex post realized earnings in the estimation in a way that parallels our baseline methodology. We use realized earnings in years $t + 1$ and $t + 2$ in place of the consensus analyst forecasts for those years. We then use the growth in realized earnings between years $t + 1$ and $t + 2$ as an estimate of the

²¹ Easton and Sommers (2006) quantify the bias in the ICC that results from biased analyst forecasts.

growth rate in year $t + 3$ to forecast earnings in year $t + 3$. Finally, we exponentially mean-revert the year $t + 3$ earnings growth rate to the long-run nominal GDP growth rate. This procedure, which uses realized earnings for years $t + 1$ and $t + 2$ but not for years $t + 3$ and higher, eliminates any potential analyst bias while minimizing the noise due to the use of ex post information. Plowback rates are estimated in the same way as before. We estimate the ICC in each month from 1981 to 2002 for the same sample of U.S. firms as in our baseline methodology. As before, we compute both the equal- and value-weighted average U.S. market risk premia.

The equal-weighted premium has a time-series mean of 4.5% per year, which is slightly lower than the mean of the premium based on the ex ante forecasts (4.6%; see Table II). However, the premium's standard deviation is 2.4%, which is almost twice as large as its ex ante counterpart (1.4%). For the value-weighted premium, the mean is also close to the ex ante mean but the standard deviation is almost three times larger (2.7% vs. 1% in Table II). This is not surprising since ex post earnings are much more volatile than ex ante earnings forecasts. The correlation between the equal-weighted (value-weighted) risk premia based on the ex ante versus ex post earnings is 51% (42%).

Table V presents the regression results for the ex post premium estimates in 24 different specifications (three regressions, equal- vs. value-weighted premium, σ_t^2 vs. σ_t , and realized vs. implied volatility). In every specification, the sign of the estimated risk–return relation is positive. The relation is statistically significant only in the residual regressions (26), namely, in all four equal-weighted specifications and in three value-weighted specifications. This reduction in significance is not surprising given the additional noise involved in using ex post earnings, as discussed earlier. However, since the coefficient estimates are uniformly positive and similar in magnitude to those in Table III, it seems unlikely that the estimated mean–variance relation is due to biases in analyst forecasts.

In the second test, we investigate whether analyst forecast errors are related to market volatility. We compute the forecast errors for each firm and month as the ratio of the difference between the consensus 1-year-ahead analyst forecast of EPS and the corresponding actual EPS to the 1-year-ahead forecast. We average the forecast errors across firms in each month to compute a market-wide forecast error for each of the G-7 countries. To eliminate outliers, we delete the top 0.5% and the bottom 0.5% of forecast errors in each country and month. Finally, we run country-level regressions of the forecast errors on the levels and changes in market variance, using all available data (1981 to 2002 for the United States and 1990 to 2002 for the other countries). We do not find a significant relation between analyst forecast errors and market volatility in any of the G-7 countries, regardless of whether the forecast errors are equal-weighted or value-weighted across firms.²² Thus, our finding of a positive mean–variance relation does not appear to be driven by analyst forecast errors.

²² To save space, we do not tabulate the results. In additional tests, we find no significant relation between absolute values of analyst forecast errors and market volatility in any of the G-7 countries.

Table V
U.S. Implied Premium Regressed on U.S. Market Volatility, Implied Premium Based on Ex Post Not Ex Ante Earnings

This table reports the slope coefficients from the time-series regressions in equations (24) through (26) for the U.S. stock market. The implied risk premium is computed by using realized ex post earnings instead of analyst forecasts of earnings as in our baseline methodology. In Panel A, return volatility σ_t is the realized market volatility from month $t - 1$, estimated from daily data and annualized, and the sample period is January 1981 to December 2002. In Panel B, σ_t is the implied volatility from month t , measured as the VXO index, and the sample period is January 1986 (when VXO data begin) to December 2002. The column label “EW” (VW) denotes an equal-weighted (value-weighted) implied risk premium, and σ_t^2 or σ_t denote the regressor. The t -statistics are adjusted for residual autocorrelation using the Newey–West correction.

	σ_t^2		σ_t	
	EW	VW	EW	VW
Panel A: Realized Volatility				
Levels: $\mu_t = a + b \sigma_t^{(2)} + e_t$				
$\hat{b}$	0.09	0.07	0.07	0.03
t	1.77	1.32	1.67	0.72
R^2	0.03	0.01	0.04	0.01
Differences: $\Delta \mu_t = a + b \Delta \sigma_t^{(2)} + e_t$				
$\hat{b}$	0.00	0.00	0.01	0.01
t	0.23	0.45	1.20	1.29
R^2	0.00	0.00	0.01	0.01
AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{V,t} + e_t$				
$\hat{b}$	0.02	0.02	0.02	0.01
t	2.61	2.01	2.69	2.03
R^2	0.03	0.02	0.03	0.02
Panel B: Implied Volatility				
Levels: $\mu_t = a + b \sigma_t^{(2)} + e_t$				
$\hat{b}$	0.07	0.07	0.03	0.01
t	1.40	0.94	0.96	0.31
R^2	0.04	0.02	0.02	0.00
Differences: $\Delta \mu_t = a + b \Delta \sigma_t^{(2)} + e_t$				
$\hat{b}$	0.01	0.02	0.01	0.02
t	1.44	1.55	1.67	1.64
R^2	0.01	0.01	0.01	0.01
AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{V,t} + e_t$				
$\hat{b}$	0.03	0.03	0.02	0.02
t	2.88	2.23	2.43	1.91
R^2	0.04	0.02	0.03	0.02

F. Robustness: Model for Computing the ICC

Our methodology, described in Section III.A, is not the only plausible approach for computing the ICC. In this section, we investigate the sensitivity of our results for the U.S. market to two reasonable departures from our baseline methodology.

First, we change the horizon in equation (22) from $T = 15$ to $T = 10$ and $T = 20$ years. The resulting equal-weighted implied premia have time series means of 3.72% per year for $T = 10$ and 5.28% for $T = 20$, and their standard deviations are 1.36% per year for $T = 10$ and 1.41% for $T = 20$. For the value-weighted premia, the means are 1.94% for $T = 10$ and 3.25% for $T = 20$, and the standard deviations are 0.90% for $T = 10$ and 1.05% for $T = 20$. Although the choice of T clearly affects the level of the implied premium, it barely affects the time variation in the premium (which is the focus of this paper) because the implied premia for $T = 10$ and $T = 20$ are both over 99% correlated with their $T = 15$ counterparts. It is therefore not surprising that the regression results for $T = 10$ and $T = 20$ are so similar to those for $T = 15$ in Table III that they are not worth reporting.

In the second and more fundamental departure, we replace our methodology by the approach of Easton et al. (2002), which simultaneously estimates the ICC and the growth rate of residual earnings beyond the forecast horizon. Instead of making assumptions about the long-term growth rate, this approach estimates the growth rate from cross-sectional regressions. The approach requires estimation at the portfolio level, which suits us because we are interested in the ICC for the U.S. market as a whole. We implement this approach by estimating monthly iterated cross-sectional regressions of 4-year cum-dividend earnings on the price-to-book ratio from 1981 to 2002 on the same sample of U.S. firms as in Table III.²³ There is no equal- or value-weighting; the approach produces only one ICC estimate at the market level. The estimates of the long-run growth rate have a time-series mean of 8.4%. The average implied risk premium has a mean of 4.2% and standard deviation of 2%, both of which are comparable to their equal-weighted counterparts in Table II. In addition, the implied U.S. market premium is highly (88%) correlated with the equal-weighted implied premium obtained by our baseline methodology.

Table VI reports the results from regressions (24) through (26). As before, we find a significant positive relation between the implied risk premium and market volatility, regardless of whether we measure volatility as σ_t or σ_t^2 and whether we use realized or implied volatility. For example, based on the residual specification (26), the t -statistics range from 2.53 to 5.17 across the four versions of volatility. It is comforting that using another sensible approach to computing the ICC leads to the same conclusions. As explained in the introduction, we believe that the whole class of ICC models, not just our baseline model, should be useful in capturing time variation in expected returns.

G. Integration of International Financial Markets

So far, we have tested the conditional mean–variance relation separately for each country. This section analyzes the risk–return tradeoff from the global perspective.

²³ As in Easton et al. (2002), we use an initial ICC estimate of 12% to compute the ratio of aggregate 4-year cum-dividend earnings to book value of equity. See Section 2.1 of that paper for further details.

Table VI
U.S. Implied Premium Regressed on U.S. Market Volatility, Implied Premium Based on the Easton et al. (2002) Approach

This table reports the slope coefficients from the time-series regressions in equations (24) through (26) for the U.S. stock market. The implied risk premium of the U.S. stock market is computed by using the approach of Easton et al. (2002) instead of our baseline methodology. In Panel A, return volatility σ_t is the realized market volatility from month $t - 1$, estimated from daily data and annualized, and the sample period is January 1981 to December 2002. In Panel B, σ_t is the implied volatility from month t , measured as the VXO index, and the sample period is January 1986 (when VXO data begin) to December 2002. The column labels σ_t^2 and σ_t denote the regressor. The t -statistics are adjusted for residual autocorrelation using the Newey–West correction.

	σ_t^2	σ_t
Panel A: Realized Volatility		
Levels: $\mu_t = a + b \sigma_t^{(2)} + e_t$		
$\hat{b}$	0.11	0.09
t	2.50	2.57
R^2	0.05	0.10
Differences: $\Delta \mu_t = a + b \Delta \sigma_t^{(2)} + e_t$		
$\hat{b}$	0.02	0.01
t	2.87	2.02
R^2	0.03	0.02
AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{V,t} + e_t$		
$\hat{b}$	0.03	0.01
t	3.80	2.53
R^2	0.05	0.02
Panel B: Implied Volatility		
Levels: $\mu_t = a + b \sigma_t^{(2)} + e_t$		
$\hat{b}$	0.10	0.06
t	1.92	1.81
R^2	0.07	0.09
Differences: $\Delta \mu_t = a + b \Delta \sigma_t^{(2)} + e_t$		
$\hat{b}$	0.05	0.03
t	4.89	4.31
R^2	0.11	0.08
AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{V,t} + e_t$		
$\hat{b}$	0.06	0.03
t	5.17	4.41
R^2	0.12	0.09

We compute the aggregate risk premium by averaging the equal-weighted or value-weighted country risk premia across the G-7 countries. We refer to this premium as the world market risk premium even though the G-7 markets account for only about 70% of the world market capitalization as of 2002. We compute the world market volatility from the daily returns on the MSCI value-weighted world market index in the previous month. Because of the reporting differences on I/B/E/S across the G-7 countries (see Section III.B), we compute

monthly volatility in two ways: from the beginning of the month to the month-end, as well as mid-month to mid-month. We conduct three tests.

In our first test, we assess the strength of the risk-return relation at the world market level by regressing the world market premium on the world market volatility. We run the same regressions ((24) to (26)) as we did at the country level. We do this for two definitions of the risk premium (equal-weighted and value-weighted) and two definitions of market volatility (σ_t^2 and σ_t), which gives us four variations of each of the three regressions.

Table VII shows a strong positive relation between the world market levels of the premium and volatility. This relation is statistically significant in all specifications. The relation between the shocks to the premium and volatility is also positive in each specification, but it is significant only for the equal-weighted risk premia. On balance, this evidence supports a positive intertemporal risk-return tradeoff at the world market level.

Under the assumptions of Merton (1973), the coefficient of proportionality between the conditional mean and variance can be interpreted as relative risk aversion. Our estimates of the slope coefficient in the level regression range from 0.28 to 0.67 across the four basic specifications. These coefficients may understate the true average level of risk aversion in the economy, for at least two reasons. First, return volatility is measured with error, and the resulting attenuation bias makes the estimated slope coefficient smaller than the true coefficient. Second, under assumptions more general than Merton's, these slope coefficients need not represent risk aversion (e.g., Backus and Gregory (1993), Campbell (1993), and Veronesi (2000)).

In our second test, we examine the cross-market risk-return relation by regressing the seven country-level risk premia on the world market variance. Panel A of Table VIII shows that the level of the equal-weighted premium is positively and significantly related to the world market volatility in six of the seven countries. The table also shows that shocks to the equal-weighted premia are positively related to the volatility shocks, but this relation is significant only in France, the United Kingdom, and the United States. The results based on the value-weighted risk premia are only slightly weaker. The level relation is positive and significant in four countries, and the relation (26) is significant in three countries (France, United Kingdom, United States).

In our third test, we analyze the relation between the implied country risk premia and the conditional covariances with the world market portfolio. This analysis is motivated by Chan, Karolyi, and Stulz (1992), who find that the U.S. risk premium is positively related to the conditional covariance of U.S. stocks with a foreign index but unrelated to its own conditional variance. For each G-7 country, we regress the risk premium on the conditional covariance between the country returns and the world market returns as well as on the world market variance. The conditional covariances with the world market portfolio are estimated from daily returns in the previous month. Panel B of Table VIII reports the results. In the levels specification, we find a significantly positive relation between the risk premia and the conditional covariances in five of the seven countries, consistent with the international CAPM. In the other two

Table VII
World-Level Implied Premium Regressed on World Market Volatility

This table reports the slope coefficients from the time-series regressions in equations (24) through (26) for the aggregate world stock market in July 1990 to December 2002. The world implied risk premium μ_t is approximated by averaging the annualized implied risk premia across the G-7 countries. The world market volatility σ_t is computed from the daily returns on the MSCI World index in the previous month and it is also annualized. In Panel A, volatility is computed from the beginning of the month to the month-end. In Panel B, it is computed mid-month to mid-month. The column label “EW” (VW) denotes equal-weighted (value-weighted) country-level implied premia, and σ_t^2 or σ_t denote the regressor. R^2 denotes adjusted- R^2 . The t -statistics are adjusted for residual autocorrelation and heteroskedasticity using the Newey–West correction.

	σ_t^2		σ_t	
	EW	VW	EW	VW
Panel A: Month-End Volatility				
Levels: $\mu_t = a + b \sigma_t^{(2)} + e_t$				
$\hat{b}$	0.67	0.30	0.23	0.09
t	3.67	2.39	3.27	2.01
R^2	0.17	0.12	0.19	0.12
Differences: $\Delta\mu_t = a + b \Delta\sigma_t^{(2)} + e_t$				
$\hat{b}$	0.09	0.03	0.03	0.01
t	2.72	1.24	2.45	0.98
R^2	0.11	0.02	0.08	0.01
AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{V,t} + e_t$				
$\hat{b}$	0.12	0.05	0.03	0.02
t	3.49	1.92	3.04	1.49
R^2	0.15	0.05	0.12	0.03
Panel B: Mid-Month Volatility				
Levels: $\mu_t = a + b \sigma_t^{(2)} + e_t$				
$\hat{b}$	0.60	0.28	0.22	0.09
t	3.98	2.80	3.26	2.09
R^2	0.15	0.12	0.18	0.11
Differences: $\Delta\mu_t = a + b \Delta\sigma_t^{(2)} + e_t$				
$\hat{b}$	0.05	0.02	0.02	0.01
t	2.06	0.76	2.02	0.95
R^2	0.03	0.00	0.03	0.00
AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{V,t} + e_t$				
$\hat{b}$	0.06	0.02	0.02	0.01
t	2.03	0.96	2.00	1.11
R^2	0.04	0.01	0.03	0.01

specifications, the relation is estimated to be positive for all seven countries but it is almost never statistically significant. Compared to Panel A, including the conditional covariance in the regression causes the market variance to lose its significantly positive coefficient for all seven countries, similar to the result that Chan, Karolyi, and Stulz (1992) find for the United States.

Table VIII
Country-Level Implied Premium Regressed on World Market Second Moments

This table reports the slope coefficients from the regressions in the panel headings for each of the G-7 markets. The sample period is July 1990 to December 2002. In Panel A, the country premia μ_t are regressed on the world market variance σ_t^2 , estimated from daily data in month $t - 1$ on the MSCI World index and annualized. In Panel B, each country's premium μ_t is regressed on $\sigma_{CW,t}$, the conditional covariance between the country market return and the world market return, which is estimated from daily data in month $t - 1$. All covariances are scaled up by a factor of 252. The column label "EW" (VW) denotes equal-weighted (value-weighted) country-level implied premia. R^2 denotes adjusted- R^2 . The t -statistics are adjusted for residual autocorrelation and heteroskedasticity using the Newey-West correction.

	EW							VW						
	CAN	FRA	GER	ITA	JAP	UK	USA	CAN	FRA	GER	ITA	JAP	UK	USA
Panel A: Regressor: World Market Variance														
	Levels: $\mu_t = a + b \sigma_t^2 + e_t$													
$\hat{b}$	0.29	0.41	1.04	0.72	0.88	0.62	0.15	0.01	0.21	0.41	0.55	0.37	0.26	0.07
t	2.09	3.62	3.84	2.64	3.32	4.77	1.83	0.12	3.36	2.78	2.73	1.87	5.33	1.20
R^2	0.04	0.10	0.16	0.09	0.08	0.22	0.07	-0.01	0.14	0.15	0.08	0.04	0.17	0.03
	Differences: $\Delta \mu_t = a + b \Delta \sigma_t^2 + e_t$													
$\hat{b}$	0.02	0.09	0.05	-0.01	0.02	0.11	0.06	0.00	0.03	0.00	-0.01	0.01	0.07	0.04
t	0.63	2.84	1.37	-0.32	0.54	4.25	2.61	0.06	1.57	-0.00	-0.31	0.51	1.86	1.96
R^2	-0.00	0.05	0.01	-0.01	-0.01	0.09	0.09	-0.01	0.01	-0.01	-0.01	-0.01	0.02	0.04
	AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{V,t} + e_t$													
$\hat{b}$	0.05	0.11	0.05	0.01	0.03	0.12	0.09	0.00	0.05	0.03	0.01	0.02	0.10	0.06
t	1.47	3.13	1.01	0.10	0.81	4.00	11.76	0.09	2.20	0.49	0.12	0.88	2.86	7.70
R^2	0.01	0.06	0.01	-0.01	-0.00	0.09	0.16	-0.01	0.02	-0.00	-0.01	-0.00	0.05	0.08
Panel B: Regressor: Covariance between Country and World, Controlling for World Market Variance														
	Levels: $\mu_t = a + b \sigma_{CW,t} + c \sigma_t^2 + e_t$													
$\hat{b}$	1.20	1.16	1.98	1.42	-1.83	1.58	0.31	0.54	0.53	1.04	1.27	-1.05	0.68	0.08
t	2.96	3.39	5.64	2.59	-5.07	2.66	1.17	1.50	4.14	5.57	2.78	-5.28	2.66	0.61
R^2	0.20	0.23	0.30	0.17	0.28	0.22	0.09	0.07	0.29	0.38	0.16	0.23	0.28	0.03
	Differences: $\Delta \mu_t = a + b \Delta \sigma_{CW,t} + c \Delta \sigma_t^2 + e_t$													
$\hat{b}$	0.14	0.10	0.06	0.00	0.04	0.13	0.04	0.12	0.07	0.14	0.04	0.07	0.09	0.03
t	1.67	1.69	1.15	0.01	1.54	1.27	0.82	1.95	1.45	1.58	0.47	2.48	0.92	0.60
R^2	0.02	0.06	0.00	-0.01	-0.00	0.10	0.09	0.00	0.01	0.01	-0.01	0.01	0.02	0.04
	AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{C,t} + c \epsilon_{V,t} + e_t$													
$\hat{b}$	0.12	0.13	0.08	0.01	0.02	0.11	0.05	0.12	0.10	0.17	0.05	0.02	0.11	0.03
t	1.58	1.98	1.25	0.05	0.47	0.98	0.78	1.95	1.61	1.74	0.41	0.69	0.98	0.63
R^2	0.02	0.08	0.00	-0.01	-0.01	0.09	0.15	0.00	0.03	0.02	-0.01	-0.01	0.05	0.07

Overall, our evidence shows that the country risk premia are affected by foreign asset returns, which suggests that financial markets are at least partially integrated.

H. Dividend Yield versus Implied Cost of Capital

In Section I, we show that when dividend growth follows an AR(1) process, the ICC is a simple function of the dividend yield (D/P) and dividend growth (see equation (7)). The first component, D/P, is commonly used to capture time variation in the expected return. The second component, dividend growth, reflects

Table IX
Country-Level Dividend Yield Premium Regressed on Country
Market Volatility

This table reports the slope coefficients from the regressions (24) through (26) for each of the G-7 markets. The sample periods begin in January 1981 for the United States, in July 1990 for Italy, and in January 1990 for the remaining five countries. Each country's sample period ends in December 2002. In Panel A, the country dividend yield premia (D/P minus the T-bill rate) μ_t are obtained as equal-weighted averages of the premia across all firms in that country. In Panel B, the premia are value-weighted. Return volatility σ_t is the realized market volatility from month $t - 1$, estimated from daily data on the given country's MSCI index (except for the United States where we use the CRSP value-weighted index) and annualized. The column labels σ_t^2 and σ_t denote the regressor. R^2 denotes adjusted- R^2 . The t -statistics are adjusted for residual autocorrelation and heteroskedasticity using the Newey–West correction.

	σ_t^2							σ_t						
	CAN	FRA	GER	ITA	JAP	UK	USA	CAN	FRA	GER	ITA	JAP	UK	USA
Panel A: Equal-Weighted Dividend Yield														
Levels: $\mu_t = a + b \sigma_t^{(2)} + e_t$														
$\hat{b}$	0.20	0.15	0.12	0.77	−0.00	0.27	0.08	0.09	0.07	0.42	0.02	0.13	0.07	
t	4.28	6.62	4.90	1.32	−0.04	5.03	1.09	4.94	5.83	3.88	1.34	0.58	6.74	1.85
R^2	0.23	0.15	0.13	0.07	−0.01	0.14	0.02	0.31	0.17	0.15	0.08	0.00	0.17	0.06
Differences: $\Delta \mu_t = a + b \Delta \sigma_t^{(2)} + e_t$														
$\hat{b}$	0.01	0.03	0.02	0.00	0.00	0.01	0.01	0.01	0.02	0.01	0.01	0.00	0.00	0.01
t	0.98	2.35	1.46	0.05	0.72	0.92	2.40	1.09	1.82	1.19	0.23	0.60	0.67	1.62
R^2	−0.00	0.01	0.02	−0.01	−0.00	−0.00	0.02	−0.00	0.02	0.01	−0.01	−0.00	−0.00	0.01
AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{v,t} + e_t$														
$\hat{b}$	0.02	0.04	0.02	0.03	0.01	0.01	0.02	0.01	0.03	0.01	0.03	0.01	0.01	0.01
t	1.60	3.54	1.96	0.51	1.47	1.72	7.04	1.65	2.58	1.53	0.66	1.56	1.26	2.36
R^2	0.00	0.02	0.04	−0.01	0.01	0.00	0.04	0.00	0.03	0.02	−0.01	0.01	0.00	0.02
Panel B: Value-Weighted Dividend Yield														
Levels: $\mu_t = a + b \sigma_t^{(2)} + e_t$														
$\hat{b}$	0.16	0.16	0.12	0.21	−0.02	0.27	0.05	0.08	0.09	0.07	0.10	0.01	0.13	0.04
t	2.77	8.00	6.59	2.41	−0.31	5.05	1.11	3.14	5.50	4.33	1.93	0.28	6.74	1.46
R^2	0.14	0.20	0.22	0.06	−0.00	0.14	0.01	0.21	0.22	0.24	0.05	−0.01	0.17	0.02
Differences: $\Delta \mu_t = a + b \Delta \sigma_t^{(2)} + e_t$														
$\hat{b}$	0.01	0.01	0.00	−0.00	0.00	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.01
t	0.91	1.40	0.73	−0.04	0.52	0.87	1.77	1.07	1.07	0.62	0.17	0.41	0.64	1.11
R^2	−0.00	0.01	−0.00	−0.01	−0.01	−0.00	0.02	−0.00	0.00	−0.00	−0.01	−0.01	−0.00	0.01
AR(1) Residuals: $\epsilon_{\mu,t} = a + b \epsilon_{v,t} + e_t$														
$\hat{b}$	0.01	0.02	0.01	0.00	0.00	0.01	0.02	0.01	0.01	0.00	0.00	0.00	0.01	0.01
t	1.07	2.03	1.41	−0.15	0.84	1.68	5.92	1.19	1.55	1.08	−0.06	0.80	1.23	1.77
R^2	−0.00	0.03	0.01	−0.01	−0.00	0.00	0.05	0.00	0.02	0.01	−0.01	−0.00	0.00	0.02

the market's expectation of future cash flow. If analyst forecasts are useful in estimating future cash flow, the ICC should better capture time variation in the expected return than D/P does. In this section, we reestimate the risk–return relation by using D/P instead of the ICC as a proxy for the expected return.

Table IX is the counterpart of Table III, with the ICC replaced by D/P. For each country, we construct monthly D/P by equal-weighting (Panel A) or value-weighting (Panel B) the dividend yields of all firms in that country. Firm-level dividend yield is computed as the ratio of all dividends paid in the most recent fiscal year to the market capitalization at the end of the current month. For the premium, we use the difference between the country-wide D/P and the local risk-free rate. The results show that D/P is positively related to volatility in several countries, as predicted by our simulation, but the relation is not as

strong as that observed when using the ICC. In Panel A of Table IX, the relation between D/P and market variance is significantly positive for four countries based on levels and for two or three countries based on shocks. In contrast, in Panel A of Table III, the relation is significantly positive for five countries in all three regression specifications. Similarly, the value-weighted results are slightly stronger in Table III than in Table IX.

This evidence leads to two conclusions. First, since the results based on D/P are weaker than those based on the ICC, analysts' earnings forecasts seem to contain useful information about expected returns. Second, since even the results based on D/P are significant in many specifications, the intertemporal risk–return relation seems reliably positive.

We also repeat the analysis in Table IX replacing D/P with the forecasted earnings-to-price ratio (E/P). Forecasted earnings are the analysts' consensus forecast of earnings 1 year ahead. In untabulated results, we find that E/P is generally even more positively related to market volatility than is D/P. For most countries, including the United States and the United Kingdom, the ICC is more closely related to volatility than is E/P, supporting our earlier conclusions.

V. Conclusion

This paper estimates the intertemporal risk–return tradeoff using the ICC as a proxy for the expected market return. We show in simulations as well as empirically that the ICC outperforms the realized return in detecting a risk–return tradeoff. Using the ICC, we find evidence of a positive relation between the conditional mean and variance of market returns in the G-7 countries, at both the country level and the world market level.

Most studies on this subject find either no relation or a negative relation between the conditional mean and variance of the U.S. market returns. The few recent studies that report a positive relation attribute their findings to a superior estimator of the conditional variance (Ghysels, Santa-Clara, and Valkanov (2005)), to the inclusion of hedging demands (Scruggs (1998), and Guo and Whitelaw (2006)), or to a longer sample (Lundblad (2007)). In contrast, our study finds a positive mean–variance relation in an international framework without a long sample, without proxies for hedging demands, and without a sophisticated variance estimator. We attribute our results solely to our expected return proxy, the ICC.

The ICC is negatively related to market prices, by construction. Thus, the fact that it reveals a positive mean–variance relation is related to the empirical fact that changes in market prices are negatively correlated with changes in market volatility (e.g., Black (1976)). We believe that the negative price–volatility relation is due at least in part to time-varying expected returns: Increases in volatility lift expected returns, driving prices down. However, we show that volatility is more closely related to the ICC than to the dividend yield, which implies that there is more to the ICC than just its negative correlation with stock prices.

The evidence of a positive intertemporal relation between the mean and variance of market returns supports the basic prediction of several asset pricing

models (e.g., Merton (1973) and Campbell (1993)). In addition, this relation has important practical implications for financial decision makers. For example, the joint dynamics of the conditional mean and variance matter for portfolio selection. Also, since second moments of returns are generally easier to measure than first moments, imposing a positive mean–variance relation a priori may improve the first moment estimates by incorporating the sample information about the second moments (e.g., Pástor and Stambaugh (2001)). Further implications of the mean–variance relation for inference and decision making can be examined in future work.

Future work can also explore other applications of the ICC. For example, while we examine the market-level ICC, it would also be interesting to analyze the time variation in the firm-level ICC. However, the market-level ICC is likely to be more precise than the firm-level ICC for two reasons. First, some potential biases in the firm-level ICC (such as steady-state growth rates that differ from the economy-wide growth rate) should approximately wash out in the market-wide average. Second, at the firm level, unlike at the market level, cash flow news dominates discount rate news in explaining the variance of stock returns (Vuolteenaho (2002)). Therefore, at the firm level, it becomes more important for analysts' earnings forecasts to accurately capture the market's expectations.

Given the emerging consensus in the finance profession that expected stock returns vary over time, it is important to find reliable ways of measuring this time variation. This paper shows that the ICC is useful in capturing time-varying expected returns. Therefore, we believe that the ICC will find new applications in the near future.

Appendix

Proofs of equations (9) and (5): By iterating equation (8), we obtain $E_t(r_{t+1+j}) = \alpha \frac{1-\beta^j}{1-\beta} + \beta^j \mu_t$. Equation (9) follows:

$$\begin{aligned} \sum_{j=0}^{\infty} \rho^j E_t(r_{t+1+j}) &= \sum_{j=0}^{\infty} \rho^j \left(\alpha \frac{1-\beta^j}{1-\beta} + \beta^j \mu_t \right) \\ &= \frac{\alpha}{1-\beta} \sum_{j=0}^{\infty} \rho^j (1-\beta^j) + \mu_t \sum_{j=0}^{\infty} \rho^j \beta^j \\ &= \frac{\alpha}{1-\beta} \sum_{j=0}^{\infty} \rho^j - \frac{\alpha}{1-\beta} \sum_{j=0}^{\infty} \rho^j \beta^j + \mu_t \sum_{j=0}^{\infty} \rho^j \beta^j \\ &= \frac{\alpha}{(1-\beta)(1-\rho)} + \left(\mu_t - \frac{\alpha}{1-\beta} \right) \frac{1}{1-\rho\beta}. \end{aligned}$$

To prove equation (5), note that $E_t(g_{t+k}) = (1-\delta^k) \frac{\gamma}{1-\delta} + \delta^k g_t$, $E_t(d_{t+k}) = d_t + \sum_{i=1}^k E_t(g_{t+i})$, and $\sum_{j=0}^{\infty} j \rho^j = \frac{\rho}{(1-\rho)^2}$. Equation (5) follows easily in a manner similar to equation (9). ■

Proof of equation (14): Using a first-order Taylor approximation to $r_{t+1} = \log(P_{t+1} + D_{t+1}) - \log(P_t)$, Campbell and Shiller (1988) show that

$$r_{t+1} \approx k + \rho p_{t+1} + (1 - \rho)d_{t+1} - p_t. \quad (\text{A1})$$

Substitute for p_{t+1} from equation (10) and pool together all terms known at time t :

$$r_{t+1} \approx k'_t + d_{t+1} + g_{t+1} \frac{\delta \rho}{(1 - \rho \delta)} - \mu_{t+1} \frac{\rho}{1 - \rho \beta}. \quad (\text{A2})$$

The return variance can be approximated by taking the variance of the right-hand side:

$$\begin{aligned} \text{Var}_t(r_{t+1}) &= \sigma_{v,t+1}^2 + \frac{\delta^2 \rho^2}{(1 - \rho \delta)^2} \sigma_{v,t+1}^2 + \frac{\rho^2}{(1 - \rho \beta)^2} \sigma_u^2 \\ &\quad + 2 \frac{\rho \delta}{(1 - \rho \delta)} \sigma_{v,t+1}^2 - 2 \frac{\rho}{1 - \rho \beta} \sigma_{uv} - 2 \frac{\rho^2 \delta}{(1 - \rho \delta)(1 - \rho \beta)} \sigma_{uv} \\ &= \frac{1}{(1 - \rho \delta)^2} \sigma_{v,t+1}^2 + \frac{\rho^2}{(1 - \rho \beta)^2} \sigma_u^2 - \frac{2\rho}{(1 - \rho \delta)(1 - \rho \beta)} \sigma_{uv}, \end{aligned} \quad (\text{A3})$$

where σ_{uv} is the covariance between u_t and v_t . For simplicity, we assume that $\sigma_{uv} = 0$, so

$$\text{Var}_t(r_{t+1}) = \frac{1}{(1 - \rho \delta)^2} \sigma_{v,t+1}^2 + \frac{\rho^2}{(1 - \rho \beta)^2} \sigma_u^2, \quad (\text{A4})$$

which proves equation (14). ■

REFERENCES

- Abel, Andrew B., 1988, Stock prices under time-varying dividend risk: An exact solution in an infinite-horizon general equilibrium model, *Journal of Monetary Economics* 22, 375–393.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, 2003, Modeling and forecasting realized volatility, *Econometrica* 71, 529–626.
- Backus, David, and Allan W. Gregory, 1993, Theoretical relations between risk premiums and conditional variances, *Journal of Business and Economic Statistics* 11, 177–185.
- Black, Fischer, 1976, Studies of stock market volatility changes, in *1976 Proceedings of the American Statistical Association, Business and Economic Statistics Section*, 177–181.
- Botosan, Christine A., 1997, Disclosure level and the cost of equity capital, *The Accounting Review* 72, 323–349.
- Botosan, Christine A., and Marlene A. Plumlee, 2005, Assessing alternative proxies for the expected risk premium, *The Accounting Review* 80, 21–53.
- Brandt, Michael W., and Qiang Kang, 2004, On the relationship between the conditional mean and volatility of stock returns: A latent VAR approach, *Journal of Financial Economics* 72, 217–257.
- Brealey, Richard A., and Stewart Myers, 2002, *Principles of Corporate Finance*, Seventh Edition (McGraw-Hill, Irwin).
- Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373–399.

- Campbell, John Y., 1993, Intertemporal asset pricing without consumption data, *American Economic Review* 83, 487–512.
- Campbell, John Y., and John Ammer, 1993, What moves the stock and bond markets? A variance decomposition for long-term asset returns, *Journal of Finance* 48, 3–37.
- Campbell, John Y., and Ludger Hentschel, 1992, No news is good news: An asymmetric model of changing volatility in stock returns, *Journal of Financial Economics* 31, 281–331.
- Campbell, John Y., and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.
- Chan, Louis K. C., Jason Karceski, and Josef Lakonishok, 2003, The level and persistence of growth rates, *Journal of Finance* 58, 643–684.
- Chan, Louis K. C., Andrew Karolyi, and René M. Stulz, 1992, Global financial markets and the risk premium, *Journal of Financial Economics* 32, 137–167.
- Claus, James, and Jacob Thomas, 2001, Equity premia as low as three percent? Evidence from analysts' earnings forecasts for domestic and international stock markets, *Journal of Finance* 56, 1629–1666.
- Easton, Peter D., 2004, PE ratios, PEG ratios, and estimating the implied expected rate of return on equity capital, *The Accounting Review* 79, 73–96.
- Easton, Peter D., and Steven J. Monahan, 2005, An evaluation of accounting-based measures of expected returns, *The Accounting Review* 80, 501–538.
- Easton, Peter, and Gregory Sommers, 2006, Bias in expected rates of return implied by analysts' earnings forecasts, Working paper, Notre Dame.
- Easton, Peter, Gary Taylor, Pervin Shroff, and Theodore Sougiannis, 2002, Using forecasts of earnings to simultaneously estimate growth and the rate of return on equity investment, *Journal of Accounting Research* 40, 657–676.
- Elton, Edwin J., 1999, Expected return, realized return, and asset pricing tests, *Journal of Finance* 54, 1199–1220.
- Fama, Eugene F., and Kenneth R. French, 1997, Industry costs of equity, *Journal of Financial Economics* 43, 153–193.
- French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and variance, *Journal of Financial Economics* 19, 3–29.
- Friend, Irwin, Randolph Westerfield, and Michael Granito, 1978, New evidence on the capital asset pricing model, *Journal of Finance* 33, 903–917.
- Gebhardt, William R., Charles M. C. Lee, and Bhaskaran Swaminathan, 2001, Toward an implied cost of capital, *Journal of Accounting Research* 39, 135–176.
- Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2005, There is a risk-return tradeoff after all, *Journal of Financial Economics* 76, 509–548.
- Glosten, Lawrence R., Ravi Jagannathan, and David E. Runkle, 1993, On the relation between the expected value and the variance of the nominal excess return on stocks, *Journal of Finance* 48, 1779–1801.
- Gode, Dan, and Partha Mohanram, 2003, Inferring the cost of capital using the Ohlson-Juettner model, *Review of Accounting Studies* 8, 399–431.
- Guay, Wayne R., S. P. Kothari, and Susan Shu, 2003, Properties of implied cost of capital using analysts' forecasts, Working paper, MIT.
- Guo, Hui, and Robert F. Whitelaw, 2006, Uncovering the risk-return relation in the stock market, *Journal of Finance* 61, 1433–1463.
- Hail, Luzi, and Christian Leuz, 2006, International differences in the cost of equity capital: Do legal institutions and securities regulation matter? *Journal of Accounting Research* 44, 1–47.
- Harvey, Campbell R., 2001, The specification of conditional expectations, *Journal of Empirical Finance* 8, 573–637.
- Kaplan, Steven N., and Richard S. Ruback, 1995, The valuation of cash flow forecasts: An empirical analysis, *Journal of Finance* 50, 1059–1093.
- Lee, Charles M. C., David Ng, and Bhaskaran Swaminathan, 2007, Testing international asset pricing models using implied costs of capital, *Journal of Financial and Quantitative Analysis*, forthcoming.

- Lettau, Martin, and Sydney C. Ludvigson, 2007, Measuring and modeling variation in the risk-return tradeoff, in Yacine Ait-Sahalia and Lars P. Hansen, eds. *Handbook of Financial Econometrics*, forthcoming.
- Ludvigson, Sydney C., and Serena Ng, 2007, The empirical risk-return relation: A factor analysis approach, *Journal of Financial Economics* 83, 171–222.
- Lundblad, Christian, 2007, The risk return tradeoff in the long-run: 1836–2003, *Journal of Financial Economics* 85, 123–150.
- Merton, Robert C., 1973, An intertemporal asset pricing model, *Econometrica* 41, 867–888.
- Merton, Robert C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.
- Nelson, Daniel B., 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* 59, 347–370.
- Ohlson, James A., and Beate E. Juettner-Nauroth, 2005, Expected EPS and EPS growth as determinants of value, *Review of Accounting Studies* 10, 349–365.
- Pástor, Ľuboš, and Robert F. Stambaugh, 1999, Costs of equity capital and model mispricing, *Journal of Finance* 54, 67–121.
- Pástor, Ľuboš, and Robert F. Stambaugh, 2001, The equity premium and structural breaks, *Journal of Finance* 56, 1207–1239.
- Scruggs, John T., 1998, Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: A two-factor approach, *Journal of Finance* 53, 575–603.
- Turner, Christopher M., Richard Startz, and Charles R. Nelson, 1989, A Markov model of heteroskedasticity, risk, and learning in the stock market, *Journal of Financial Economics* 25, 3–22.
- Veronesi, Pietro, 2000, How does information quality affect stock returns? *Journal of Finance* 55, 807–837.
- Vuolteenaho, Tuomo, 2002, What drives firm-level stock returns? *Journal of Finance* 57, 233–264.
- Wang, Leping, 2004, On the intertemporal risk-return relation: A Bayesian model comparison perspective, Working paper, Wharton.
- Whitelaw, Robert F., 1994, Time variations and covariations in the expectation and volatility of stock market returns, *Journal of Finance* 49, 515–541.
- Whitelaw, Robert F., 2000, Stock market risk and return: An equilibrium approach, *Review of Financial Studies* 13, 521–547.