Analysis of the Predictive Ability of Information Accumulated over Nights, Weekends, and Holidays

Ilias Tsiakas*
Warwick Business School, UK

Preliminary and Incomplete

July 1, 2004

Abstract

Recent empirical evidence suggests that the weekend and holiday calendar effects are much stronger and statistically significant in volatility as opposed to expected returns. This paper seeks an explanation for this empirical finding by undertaking a comprehensive investigation of the predictive ability of information accumulated over nights, weekends and holidays for a series of global indices. We study this form of seasonal heteroscedasticity by employing a generalized stochastic volatility model, in which the conditional daily volatility is measured in calendar time from open-to-close of the market, and depends on lagged close-to-open returns. We conduct a series of empirical tests and conclude that the information accumulated over weekends and especially holidays is a predictor of subsequent daily volatility. The SV parameters are estimated by implementing a Bayesian MCMC algorithm, which is adjusted for sampling the seasonal volatility level effects. We compute in-sample and out-of-sample density forecasts for assessing the adequacy of the conditional distribution. We also use Bayes factors as a likelihood-based framework for evaluating the SV specifications. Bayes factors account for both estimation and model risk. We conclude by computing volatility forecasts relevant for risk management.

Keywords: Stochastic Volatility, Calendar Effects, Seasonal Heteroscedasticity, Bayesian MCMC Estimation, Bootstrapping, Forecasting.

JEL Classification: C11, C15, C22, C51, C52, C53, G10.

*Accounting and Finance Department, Warwick Business School, University of Warwick, Coventry CV4 7AL, United Kingdom. Tel: +44 (0)2476 524966. Fax: +44 (0)2476 523779. Email: ilias.tsiakas@wbs.ac.uk.

1 Introduction

Recent empirical evidence suggests that the weekend and holiday calendar effects are much stronger and statistically significant in volatility as opposed to expected returns. This paper seeks an explanation for this empirical finding by undertaking a comprehensive investigation of the predictive ability of information accumulated over nights, weekends and holidays for a series of global indices. We study this form of seasonal heteroscedasticity by employing a generalized stochastic volatility model, in which the conditional daily volatility is measured in calendar time from open-to-close of the market, and depends on lagged close-to-open returns. We conduct a series of empirical tests and conclude that the information accumulated over weekends and especially holidays is a predictor of subsequent daily volatility. The SV parameters are estimated by implementing a Bayesian MCMC algorithm, which is adjusted for sampling the seasonal volatility level effects. We compute in-sample and out-of-sample density forecasts for assessing the adequacy of the conditional distribution. We also use Bayes factors as a likelihood-based framework for evaluating the SV specifications. Bayes factors account for both estimation and model risk. We conclude by computing volatility forecasts relevant for risk management.

We compute in-sample and out-of-sample density forecasts for assessing the adequacy of the conditional distribution of the SV specifications. We explicitly test whether modelling the distinct behaviour of different days and months results in better one-step ahead volatility and density forecasts. We also compute Bayes factors, which provide a framework for specification diagnostics and model selection over the set of SV models. Bayes factors account for both estimation risk by integrating out parameter uncertainty and for model risk, which arises from the uncertainty over selecting a model specification. More importantly, the Bayes factor diagnostic measures the statistical cost of dimensionality due to the explicit accounting of all seasonal periodic effects in returns and volatility.

Models of stochastic volatility have been used extensively in theoretical option pricing since the contribution of Hull and White (1987) in generalizing the Black-Scholes option pricing scheme. Like GARCH-type models, they are designed to capture the persistent and hence predictable component of daily volatility (for a comparison of GARCH and SV models see Fleming and Kirby (2003)). However, SV has a fundamental difference with GARCH. The assumption of a stochastic second moment introduces an additional source of risk that cannot be perfectly hedged using $t-1$ information. A GARCH specification describes the conditional distribution of returns as being a function of exclusively past information. In contrast, the SV model specifies the joint conditional distribution of both the return and the volatility process. Intuitively, SV allows for the possibility of random (but rather small) contemporaneous volatility shocks due to news events and policy changes. In other words, there may exist unobserved contemporaneous variables that affect the volatility process, which is not possible in GARCH.

Despite their parsimonious structure, intuitive appeal and popularity in theoretical option pricing, SV models have been much less popular than GARCH in empirical applications. This is exclusively due to the difficulties associated with estimating SV models using conventional Classical econometric methods. Specifically, models of discrete-time stochastic volatility cannot be estimated with likelihood-based methods because the likelihood function is not available analytically.[1] Bayesian es-

[1] In the Classical framework, Sandmann and Koopman (1998) propose a Monte Carlo Likelihood (MCL) method for estimating simple (plain vanilla) SV models. For a Simulated Maximum Likelihood (SML) estimation method of heavy-tailed SV models see Liesenfeld and Jung (2000).

timation offers a substantial computational advantage over any Classical approach because it avoids tackling very difficult, if not intractable, numerical optimization procedures. This has turned the development of fast and efficient Bayesian MCMC algorithms for the estimation of SV models into one of the most promising and challenging tasks of modern time series analysis.[2]

The SV parameters are estimated by implementing the MCMC algorithm of Chib, Nardari and Shephard (2002), which builds on the procedures developed by Kim, Shephard and Chib (1998).[3] The specification and model selection tests are based on the filtering methods of Pitt and Shephard (1999). The marginal likelihood input to the computation of Bayes factors is constructed as in Chib (1995), and Chib and Jeliazkov (2001).

The paper is organized as follows. Section 2 describes the size and statistical significance of the overnight information in six international stock indices. Section 3 discusses the plain vanilla Gaussian SV and the proposed information SV specification (iSV) models, which is designed to assess the predictive ability of information accumulated over nights, weekends and holidays in global financial markets. The SV specifications lead to the testing hypotheses presented at the end of Section 3. A sketch of the MCMC algorithm is offered in Section 4. Section 5 examines the in-sample and out-of-sample conditional dynamics of the SV models and discusses Bayes factors as a diagnostic tool for model selection. Section 6 discusses the results and Section 7 concludes.

2 Index Returns Data

In this paper, we measure the size and the predictive ability of the information accumulated over nights, weekends and holidays, which impacts on the expected returns and volatility of global financial markets. We use daily open and closing price data from 6 international stock indices, which we selected on the basis of two criteria: (i) there are two indices representing each of the three main global trading blocks: USA's Dow Jones Industrial Average (DJIA) and Canada's S&P/TSX Composite from North America; UK's FTSE 100 and Germany's DAX 30 from Europe; and Japan's Nikkei 225 and Hong Kong's Hang Seng from Asia, and (ii) one index is from a major and one from a minor market from each trading block. In other words, we assume that New York is a bigger market than Toronto, London is bigger than Frankfurt, and finally Tokyo is bigger than Hong Kong. This will allow us to test not only if there are differences in the way markets respond to pre-open information by geography but also by relative size.

The start date of the samples and the number of observations for each index are shown in Table 1. The sample size of the available data for both open and closing prices ranges from 20 years in the case of the FTSE 100 (5039 observations) to just 11 years for the Hang Seng (2658 observations). The average sample size is approximately 16 years. The end date for all indices is December 31, 2003, except in the case of Hong Kong for which the available opening price data ends on September 28, 2001. For ease of identification, we will be referring to the indices by their country name. The source of the data is Datastream.

Let P_t^O and P_t^C denote the daily open and closing values of each international price index, respectively. Then, the continuously compounded open-to-close percent *day* returns are constructed

[2] For a general reference on MCMC methods in financial econometrics see Johannes and Polson (2002).

[3] For an alternative reference on Bayesian analysis of SV models see Jacquier, Polson and Rossi (2002).

simply as

$$r_t^D = 100 \log \frac{P_t^C}{P_t^O} \tag{2.1}$$

We isolate the pre-opening information by defining the lagged close-to-open *night* returns as

$$r_{t-1/2}^N = 100 \log \frac{P_t^O}{P_{t-1}^C} \tag{2.2}$$

For convenience we denote days with time integers and nights with half integers. Hence, if Tuesday is day 2 and Wednesday is day 3, then the period from Tuesday close to the Wednesday open is period $2\frac{1}{2}$. Hence the information available to the econometrician at the start of period t is $F_t = \{r_{t-1/2,} r_{t-1}, ..., r_1, r_{1/2}\}$.

We further distinguish four mutually exclusive types of lagged overnight close-to-open information: (i) $r_{t-1/2}^{WN}$ is the close-to-open *weeknight* return for the Monday, Tuesday, Wednesday and Thursday nights, unless there is an intervening holiday; for example, the Monday weeknight $r_{t-1/2}^{WN}$ is the overnight period from the Monday close to the Tuesday open; (ii) $r_{t-1/2}^{WE}$ is the close-to-open *weekend* return measured from the Friday close to the Monday open, unless there is an intervening Friday or Monday holiday; (iii) $r_{t-1/2}^{H}$ is the close-to-open *holiday* return measured from the pre-holiday close to the post-holiday open, if there is no intervening weekend; and finally (iv) $r_{t-1/2}^{LW}$ is the close-to-open long *weekend* return measured from the pre-holiday close to the post-holiday open, if the non-trading holiday is either a Friday or a Monday. All four are defined such that there is no overlap; for example, long weekends count as neither weekends nor holidays. The sum of the returns of all weeknights, weekends, holidays and long weekends gives the night return.

2.1 Descriptive Statistics and Overnight Information

Table 1 presents the descriptive statistics of the day and night returns and volatility. Panel A indicates that all average night and day returns are positive with three notable exceptions. For Hong Kong the average night return is negative. More importantly, for Germany and Japan the mean day return is negative, whereas the mean night return is positive. In addition, for the United States the mean night return is essentially zero. In all cases, the day returns have higher standard deviation than the night returns, but (with the exception of Japan) the night returns have much higher kurtosis. Finally, the correlation of day and lagged night returns ranges from -3% in the UK to 23% in Japan, with five of the six correlations being positive.

Panel B presents the results for absolute returns. Here, we use absolute returns as a simple model-free proxy to daily volatility, which we will formally model in the next section. In all cases, the average absolute day returns are higher than the average absolute night returns. Further, with the exception of Japan, the mean absolute night returns have higher positive skewness and kurtosis than the mean absolute day returns. Finally, the correlation between day and lagged night absolute returns ranges from -5% for Germany to 33% for Canada, with four of the six correlations being positive. Note that the average correlation between day and night across the six countries is 9% for returns and 11% for absolute returns.

Table 2 presents the proportions of volatility due to the information accumulated during the day and during the night. The day and night volatility proportions are defined as

$$vp^D = \frac{1}{T}\sum_{t=1}^{T}\frac{|r_t^D|}{|r_t^D| + |r_{t-1/2}^N|} \tag{2.3}$$

and

$$vp^N = \frac{1}{T}\sum_{t=1}^{T}\frac{|r_{t-1/2}^N|}{|r_t^D| + |r_{t-1/2}^N|} \tag{2.4}$$

For the three minor markets (Canada, Germany and Hong Kong) the volatility proportions are remarkably similar at around 40% for the night proportion and 60% for the day proportion. For the two major markets outside North America (UK and Japan) they become 30%-70%. For the US, the night proportion falls to just 7%.

2.2 Statistical Significance and Bootstrapping

We assess the statistical significance of the day and night effects by performing formal hypothesis tests based on bootstrapping. Our null hypotheses are all one-sided and test whether (for example) the day returns are higher (or in the case of Germany and Japan lower) than the night returns. Each variable is tested against its complement: day vs. night, weekend vs. non-weekend nights, holiday vs. non-holiday nights, and long weekend vs. non-long weekend nights. Each hypothesis is tested at 90%, 95%, and 99% confidence levels. The details on forming one-sided hypothesis tests using bootstrapping are in Appendix A (also see Hansen (2004)). In assessing statistical significance, we have excluded the returns data from the two weeks of October 19-23, 1987 and September 10-14, if applicable. For the US DJIA data we have also excluded September 17, 2001, which was the first trading day in the US after 9/11. This ensures that our results are not driven by very few outliers and guarantees there are no biases depending on (i) whether a data sample starts before or after October 1987, or (ii) how many days a market was closed for 9/11.

The evidence on statistical significance is presented in Table 3. In expected returns, days are statistically different to nights all with at least 95% confidence in four of the six countries, the exceptions being Canada and the UK. The weekend effect is significant in 3 countries, the holiday effect in 2, and the long weekend effect in 1 country index.

The strongest data-based result is displayed in Panel B of Table 3: for all six indices day volatility is statistically different to night volatility with 99% confidence. In addition, for all four countries outside of North America, the weekends are statistically different to their complements. For Canada only long weekends are significant and for the US only holidays. Long weekends are significant in five countries, the only exception being the US. This is not surprising because for most Canadian long weekends, the US market is open and hence Canada responds stongly to the information accumulating in its large neighbour. In contrast, the US appears to be sensitive only to information generated within the US, and therefore does not respond to information accumulated abroad during US weekends and long weekends. In fact, the US holiday, which is the only overnight effect that statistically different to its complement in the US, is smaller that its complement. In other words, the US in general does not respond to overnight information, but during holidays it responds even less! On average, of the three night effects, 1 is statistically significant in expected returns (or 0.67 with at least 95% confidence) versus 1.67 in volatility (or 1.33 with at least 95% confidence).

Figure 1 presents the boostrap distribution of the day and night expected returns for the six indices. Figure 2 does the same for the day and night volatilities and demonstrates that in all countries the day volatility is much higher and very statistically different to the night volatility. Figures 1 and 2 offer a visual inspection of (i) the overlap of the two bootstrap distributions for each day and night mean, and (ii) the position of a sample mean in the bootstrap distribution of its complement. Similarly, figures 3 and 4 plot the boostrap distribution of the strongest night effect against the distribution of its night complement for the expected returns and volatility, respectively. We can see that the strongest mean effects are less strong relative to the strongest volatility effects: there is more overlap in the mean boostrap distributions as opposed to the volatility distributions.

In short, there is strong quantitative and visual evidence that in the six selected global indices there is clear misspecification if we do not explicitly condition on the overnight information and distinguish the weeknight effect from the weekend, holiday and long weekend overnight effects. In the next section, we will formally model the effect of overnight information on conditional returns and volatility and we will investigate its forecasting.

3 Stochastic Volatility

3.1 The Plain Vanilla SV model

In the stochastic volatility (SV) framework, the plain vanilla SV model typically presents the benchmark against which model comparisons are conducted. According to this plain vanilla SV benchmark, the daily index returns are assumed to follow a univariate discrete-time AR(1) process and are driven by Gaussian innovations:

$$r_t^D = \alpha + \beta_D r_{t-1}^D + \varepsilon_t v_t, \quad \varepsilon_t \sim NID\left(0, 1\right) \tag{3.1}$$

The persistence of the stochastic conditional volatility v_t is captured by the latent log-variance process h_t^D, which is modelled as a dynamic Gaussian variable:

$$v_t = \exp\left(h_t^D/2\right) \tag{3.2}$$

$$h_t^D = \mu + \Gamma' X_t + \phi\left(h_{t-1}^D - \mu\right) + \sigma\eta_t, \quad \eta_t \sim NID\left(0, 1\right) \tag{3.3}$$

In the plain vanilla model, return and volatility innovations are independent: $\{\varepsilon_t\} \perp \{\eta_t\}$. Further, the model assumes (and the estimation algorithm imposes) $|\rho|, |\phi| < 1$ so that both returns and their volatility are stationary processes.[4] Finally, the simple SV specification reduces $\Gamma' X_t = \gamma_D r_{t-1}^D$. In words, if $\gamma_D < 0$ the lower the return shock, the higher the conditional variance in the next few periods. This simple specification allows for a level component measured by $\gamma_D r_{t-1}^D$, as well as a dynamic component measured by $\phi^j \gamma_D r_{t-j}^D$.

[4] In practice, the stationarity restrictions are never violated for daily returns data.

3.2 Information Stochastic Volatility (iSV)

The information SV (iSV) specification conditions both the return and the volatility equations on the information contained in the returns of (i) the lagged open-to-close period (preceding day), and (ii) the lagged close-to-open period (preceding weeknight, weekend, holiday, or long weekend). Specifically, we model the returns as follows:

$$r_t^D = \alpha + \beta_D r_{t-1}^D + \beta_N r_{t-1/2}^N + \varepsilon_t v_t, \quad \varepsilon_t \sim NID\,(0,1) \tag{3.1}$$

$$\beta_N r_{t-1/2}^N = \beta_{WN} r_{t-1/2}^{WN} D_{t-1/2}^{WN} + \beta_{WE} r_{t-1/2}^{WE} D_{t-1/2}^{WE} + \beta_H r_{t-1/2}^{H} D_{t-1/2}^{H} + \beta_{LW} r_{t-1/2}^{LW} D_{t-1/2}^{LW} \tag{3.2}$$

where $\left\{D_{t-1/2}^{WN}, D_{t-1/2}^{WE}, D_{t-1/2}^{H}, D_{t-1/2}^{LW}\right\}$ is the set of dummy variables in $\{0,1\}$, taking the unit value if the previous close-to-open period is a weeknight, weekend, holiday or long weekend, respectively.

Similarly, the persistent and stochastic conditional volatility v_t conditions on lagged day, weeknight, weekend, holiday, and long weekend returns:

$$v_t = \exp\left(h_t^D/2\right) \tag{3.3}$$

$$h_t^D = \mu + \gamma_D r_{t-1}^D + \gamma_N r_{t-1/2}^N + \phi\left(h_{t-1}^D - \mu\right) + \sigma\eta_t, \quad \eta_t \sim NID\,(0,1) \tag{3.4}$$

$$\gamma_N r_{t-1/2}^N = \gamma_{WN} r_{t-1/2}^{WN} D_{t-1/2}^{WN} + \gamma_{WE} r_{t-1/2}^{WE} D_{t-1/2}^{WE} + \gamma_H r_{t-1/2}^{H} D_{t-1/2}^{H} + \gamma_{LW} r_{t-1/2}^{LW} D_{t-1/2}^{LW} \tag{3.5}$$

The main assumptions of the plain vanilla SV model are also valid for the iSV specification. Here, the MCMC algorithm must provide estimates of the three sets of parameters $\theta = \{\theta_1, \theta_2, \theta_3\}$, where $\theta_1 = \{\alpha, \beta_D, \beta_{WN}, \beta_{WE}, \beta_H, \beta_{LW}\}$ is the set of parameters of the conditional mean, $\theta_2 = \left\{\mu, \phi, \sigma^2\right\}$ is the set of parameters of the plain vanilla Gaussian log-variance process, and $\theta_3 = \{\gamma_D, \gamma_{WN}, \gamma_{WE}, \gamma_H, \gamma_{LW}\}$ is the set of volatility parameters which measure the relative effect of overnight information. All θ parameters are time invariant.

For comparative purposes we aslo estimate a simpler iSV specification without separate night effects. Specifically, for the iSV_n specification, the return equation conditions on a single overnight effect $\beta_N r_{t-1/2}^N$. Similarly for the volatility equation which assumes as single source of vernight information $\gamma_N r_{t-1/2}^N$.

3.3 Testing Hypothesis

It is important to explore two types of tests: (i) hypothesis tests regarding the statistical significance of the parameter estimates which capture the predictive ability of the information accumulated in the close-to-open periods, and (ii) assessing the overall performance of the plain vanilla SV versus the two informational SV specifications (iSV and iSV_n) by ranking them using the toolkit discussed in the next two sections.

The information SV framework enables the formal testing of two hypotheses on parameter significance. The first hypothesis examines whether the overnight close-to-open information has statistically significant predictive power on both returns and volatility. The second one tests for the need to differentiate between the weeknight, weekend, holiday and long weekend effects. These are summarized by two separate null hypotheses:

Hypothesis 1: Overnight information has no predictive abillity in returns and volatility:

$$
\begin{aligned}
H_0^A \quad &: \quad \beta_N r_{t-1/2}^N = \gamma_N r_{t-1/2}^N = 0 && (3.1)\\
\text{or} \qquad \beta_{WN} &= \beta_{WE} = \beta_H = \beta_{LW} = 0 && (3.2)\\
\text{and} \qquad \gamma_{WN} &= \gamma_{WE} = \gamma_H = \gamma_{LW} = 0 && (3.3)
\end{aligned}
$$

$$H_1 : \quad otherwise$$

Hypothesis 2: There is no need to differentiate between WN, WE, H, LW

$$
\begin{aligned}
H_0^B \quad &: \quad \beta_{WN} = \beta_{WE} = \beta_H = \beta_{LW} = \beta_N && (3.4)\\
\text{and} \qquad \gamma_{WN} &= \gamma_{WE} = \gamma_H = \gamma_{LW} = \gamma_N && (3.5)
\end{aligned}
$$

$$H_1 : \quad otherwise$$

In addition to tests on the parameter estimates, we also assess the conditional dynamics of all SV specifications by computing one-step ahead density forecasts. We evaluate the relative performance of each specification by ranking all SV models according to the likelihood-based Bayes factor criterion. Note that all SV specifications are run on a set of adjusted samples, which have excluded the daily returns data for the two full weeks encompassing the October 1987 Crash and the 9/11 attacks for the reasons explained in the previous section. The next two sections discuss the set of tools we use for estimation, assessing the SV conditional dynamics and model evaluation.

4 Bayesian MCMC Estimation

We perform Bayesian MCMC estimation by constructing a Markov chain whose limiting distribution is the target posterior density of interest. This Markov chain is a Gibbs sampler in which all parameters are drawn sequentially from their full conditional posterior distribution. The chain is then iterated a large number of times and the sampled draws, beyond a burn-in period, are treated as variates from the target distribution. For the *iSV* model, the Bayesian MCMC algorithm produces estimates of the posterior means of (i) the set parameters of the return equation $\theta_1 = \{\alpha, \beta_D, \beta_{WN}, \beta_W, \beta_H, \beta_{LW}\}$, (ii) the set of parameters of the plain vanilla Gaussian log-variance process $\theta_2 = \{\mu, \phi, \sigma^2\}$, and (iii) the set of volatility parameters which measure the relative effect of overnight information $\theta_3 = \{\gamma_D, \gamma_{WN}, \gamma_W, \gamma_H, \gamma_{LW}\}$.

The key to estimating the *iSV* models is the efficient sampling of the overnight level effects θ_3 in the conditional variance. The paper uses the simple Gibbs step addition of Tsiakas (2004) to the Chib, Nardari and Shephard (2002) algorithm, in which the γ_j vector is drawn conditional on the log-volatilities $\{h_t\}$ using a precision-weighted average of prior information and the conditional likelihood. The details on sampling γ_j are summarized in Appendix B.

4.1 MCMC Diagnostics

The mean of the MCMC draws is an asymptotically efficient estimator of the posterior mean of θ (see Geweke (1989)). The Numerical Standard Error (NSE) is the square root of the asymptotic variance of the estimator:

$$NSE = \sqrt{\frac{\widehat{S}}{M}} \tag{4.1}$$

$$\widehat{S} = \widehat{\gamma}_0 + 2\sum_{j=1}^{B_M} K(z)\,\widehat{\xi}_j \tag{4.2}$$

Here, $M = 5,000$ is the number of iterations (beyond the initial burn-in of 1,000 iterations), $j = 1, ..., B_M = 500$ lags is the set bandwidth, $z = \frac{j}{B_M}$, and $\widehat{\xi}_j$ is the sample autocovariance of the MCMC draws of each estimated parameter cut according to the Parzen kernel $K(z)$.

The NSE diagnostic is distinct from the MCMC standard deviation reported in the legend of Tables 4 and 5. The latter is simply a measure of the variation in the MCMC parameter draws. In contrast, NSE is a measure of variation of the posterior mean estimate across many MCMC chains that can be potentially run. In other words, NSE measures how much difference one should expect in the estimate of the posterior mean if the experiment were to be repeated, and hence provides a measure of convergence.

The Relative Numerical Inefficiency (RNI) is given by

$$RNI = 1 + 2\sum_{j=1}^{B_M} K(z)\,\rho(j) \tag{4.3}$$

where $\rho(j)$ is the autocorrelation in the MCMC draws at lag j for the parameter of interest. RNI accounts solely for the variance inflation (inefficiency) due to the serial correlation of the MCMC parameter draws (see Geweke (1992) for the details). In general, the lower the serial correlation, the lower the number of iterations needed to attain a given level of numerical accuracy. For example, if RNI were to be halved, one would need half the number of iterations to attain the same level of numerical accuracy. The relatively low RNI values reported in legend of Tables 4 and 5 reflect the efficiency of the Metropolis-Hastings algorithm used to sample θ_2.

4.2 Volatility Estimates

The conditional dynamics of the SV model are essentially driven by the persistent, latent and Gaussian log-volatility process $\{h_t\}$. The tools of Chib, Nardari and Shephard (2002) allow the simulation of three distinct estimates of the $\{h_t\}$ vector. First is the smoothed volatility. The MCMC chain samples from the density $h^{(i)} \mid F_T, \theta^{(i)}$. In words, it samples the $\left\{h_t^{(i)}\right\}$ vector at a given iteration $i = 1, ..., M$ conditional on the information F_T from the full dataset (hence smoothed) and the parameter vector draw $\theta^{(i)}$.

Second is the filtered volatility. The Auxiliary Particle Filter of Pitt and Shephard (1999) samples from the density $h_t^j \mid F_t, \theta$. In words, it generates a $j = 1, ..., M = 2,000$ vector of log-volatilities (the "particles") at each t, given the information set F_t and the true values of θ proxied by the MCMC posterior mean estimates. This is a non-trivial task performed by an Auxiliary Sampling-Importance Resampling algorithm. The SV application of the algorithm is also detailed in Chib, Nardari and Shephard (2002).

Third is the one-step ahead predictive volatility. This samples from $h_{t+1}^j \mid F_t, \theta$. Given a vector of $j = 1, ..., M = 2,000$ particles from the filtered density $h_t^j \mid F_t, \theta$ it is straightforward to compute the one-step ahead vector of particles from the predictive density using the Gaussian evolution equation:

$$h_{t+1}^j = \mu + \Gamma' r_{t-1/2}^{CO} + \phi \left(h_t^j - \mu \right) + \sigma \eta_{t+1}^j, \qquad \eta_{t+1}^j \sim NID(0,1) \tag{4.4}$$

4.3 Log-Likelihood

The likelihood function of SV models is not available analytically and hence must be simulated.

$$L(\theta; r) = f(r_1, ..., r_T \mid F_0, \theta) = \prod_{t=1}^{T} f(r_t \mid F_{t-1}, \theta) \tag{4.5}$$

Specifically, the log-likelihood function is evaluated under the predictive density as

$$\log \widehat{L} = \sum_{t=1}^{T} \log \widehat{f}(r_t \mid F_{t-1}, \theta) = \sum_{t=1}^{T} \log \widehat{f}_t(r_t \mid h_t, \theta) \tag{4.6}$$

where h_t is the one-step ahead predictive volatility $h_t \mid F_{t-1}, \theta$, and θ is taken as the posterior mean estimate from the MCMC simulations.

5 Performance Evaluation

A statistical model will not be empirically successful unless it is well specified. For example, a risk manager may be interested in the average probability with which an event arises. This is determined by a well-specified unconditional distribution. More importantly, however, managing day-to-day risk involves making decisions conditional on all available information at time t. This requires a well-specified conditional distribution. This section formally tests whether the iSV model (i) performs well in capturing the in-sample and out-of-sample conditional dynamics of the daily index returns and volatility period-by-period, and (ii) is better specified than the SV benchmark using Bayes factors as the criterion for model selection.

5.1 Density Forecasts and Conditional SV Dynamics

Kim, Shephard and Chib (1998) form a set of diagnostic tests for assessing the adequacy of the conditional distribution of SV models using the simple Rosenblatt (1952) transformation. This defines SV residuals as

$$u_{t+1} = \Pr(R_{t+1} \leq r_{t+1} \mid h_{t+1}, \theta) \sim UID[0,1] \tag{5.1}$$

where r_{t+1} is the ex-post realized return and θ is the posterior mean estimate. The probability is evaluated using the ex-ante forecasted cumulative distribution function, where $\{h_{t+1}\}$ is the one-step ahead predictive log-volatility $h_{t+1} \mid F_t, \theta$. The uniform residuals are then mapped to a normal distribution simply because there is a larger battery of specification tests available for a normal random variable. Then, under the null that the model is correctly specified, $n_{t+1} = f^{-1}(u_{t+1})$ should be Gaussian white noise. Note that the normalized residuals $\{n_{t+1}\}$ contain the same information as the uniform $\{u_{t+1}\}$.

Using these normalized residuals as a basis for diagnostic testing is not restricted to SV models. Berkowitz (2001) suggests the use of n_{t+1} (also known as density forecasts) for evaluating the performance of generic risk models. An important advantage of density forecasts is that there are as many of them as data observations. In contrast, for example, the popular value-at-risk (VaR) calculation measures the frequency of tail events and hence produces too few tail observations for reasonable sample sizes. Further, note that density forecasts also account for the size of observations, not just their frequency. In short, therefore, density forecasts offer statistical power that is missing in VaR calculations, while using information from the entire conditional distribution, not just a single quantile.

Table 6 reports the specification tests for the in-sample and out-of-sample $\{n_{t+1}\}$ diagnostics. We follow the notation of Bowman and Shenton (1975) and define $\sqrt{b_1} = m_3/m_2^{3/2}$ and $b_2 = m_4/m_2^2$, where m_j is the jth centralized sample residual moment. Then, we define $SKEW$ and $KURT$ to be the asymptotic standard normal test statistics of $\sqrt{b_1}$ and b_2 respectively:

$$SKEW = \sqrt{\frac{T}{6}}\sqrt{b_1} \sim N(0,1) \tag{5.2}$$

$$KURT = \sqrt{\frac{T}{24}}(b_2 - 3) \sim N(0,1) \tag{5.3}$$

The identification of skewness and excess kurtosis in the density forecasts is very important. As it becomes increasingly difficult to capture the properties of higher order moments, misspecification occurs more often at the tails of the predictive density. It is straightforward to prove that if the observed data r_t is fat-tailed relative to the SV model, then the density forecasts will be fat-tailed relative to the standard normal.[5]

5.2 Model Risk and Bayes Factors

Model risk arises from the uncertainty over selecting a model specification. Bayes factors can account for model risk by providing a framework for specification diagnostics over a set of given models. Specifically, a Bayes factor offers a summary of the evidence provided by the data in favour of a scientific theory represented by a statistical model.[6] Consider the two competing hypotheses (models) M_1 and M_2. Using Bayes theorem, it is straightforward to show that the Bayes factor B_{21} (in favour of model M_2) is the ratio of posterior to prior odds

$$B_{21} = \frac{p(M_2 \mid r)}{p(M_1 \mid r)}\frac{\pi(M_1)}{\pi(M_2)} \tag{5.4}$$

[5] For the proof see Berkowitz (2001).

[6] See Kass and Raftery (1995) for a review of Bayes factors.

and is computed as the ratio of the marginal likelihoods

$$B_{21} = \frac{p(r \mid M_2)}{p(r \mid M_1)} \tag{5.5}$$

The marginal likelihood of a model is defined as

$$p(r \mid M_j) = \int_\theta p(r, \theta \mid M_j)\, d\theta = \int_\theta p(r \mid \theta, M_j)\, \pi(\theta \mid M_j)\, d\theta \tag{5.6}$$

Note that the marginal likelihood is an averaged and not a maximized likelihood. This implies that the Bayes factor is an automatic "Occam's Razor" in that it integrates out parameter uncertainty.[7] Further, the marginal likelihood is simply the normalizing constant of the posterior density. Suppressing the model index M_j for simplicity, the marginal likelihood can be written as

$$p(r) = \frac{f(r \mid \theta)\, \pi(\theta)}{\pi(\theta \mid r)} \tag{5.7}$$

where $f(r \mid \theta)$ is the likelihood, $\pi(\theta)$ the prior density, $\pi(\theta \mid r)$ the posterior density, and θ is evaluated at the posterior mean estimate θ^*. Since θ is drawn in the context of Gibbs MCMC sampling, the posterior density $\pi(\theta \mid r)$ is computed using the technique of reduced conditional MCMC runs of Chib (1995). For the parameter blocks of θ (the log-variance parameters and the degrees of freedom) which are sampled in the MCMC chain by implementing a Metropolis-Hastings algorithm, the posterior density is computed as in Chib and Jeliazkov (2001).

To assess the information provided by a Bayes factor, it is useful to consider twice its natural logarithm so as to be on the same scale as the likelihood ratio statistics. To make the interpretation more familiar, Table 7a presents the range of the values of $2\ln(B_{21})$ that constitute evidence against the null hypothesis M_1. Finally, note that model comparisons based on Bayes factors are asymptotically equivalent to evaluations based on the Schwartz (or equivalently the BIC) criterion.[8]

6 Results and Discussion

6.1 Parameter Estimates

See Tables 4-5.

6.2 Bayes Factor Diagnostics and Density Forecasts

See Tables 6-7.

[7] Occam's razor is just the principle of parsimony. For an econometrician, the most useful statement of the principle is "among two competing theories, which make exactly the same prediction, the simplest one is best".

[8] The Schwartz criterion is defined as $S = \log p\left(r \mid \widehat{\theta}_2, M_2\right) - \log p\left(r \mid \widehat{\theta}_1, M_1\right) - \frac{1}{2}(d_2 - d_1)\log(T)$, where d_j is the dimension of θ_j. As $T \to \infty$ the Schwartz criterion satisfies $\frac{S - \log B_{21}}{\log B_{21}} \to 0$ and thus may be viewed as a rough approximation to the log of the Bayes factor. Note that $BIC = -2S$. Again, see Kass and Raftery (1995) for the details.

7 Concluding Remarks

To be completed

References

Bauwens, L., W.B. Omrane, and P. Giot, (2003). “News Announcements, Market Activity and Volatility in the Euro/Dollar Foreign Exchange Market,” *Journal of International Money and Finance* (forthcoming)

Berkowitz, J., (2001). “Testing Density Forecasts with Applications to Risk Management,” *Journal of Business and Economic Statistics*, **19**, 465-74.

Bollerslev, T., and E. Ghysels, (1996). “Periodic Autoregressive Conditional Heteroscedasticity,” *Journal of Business and Economic Statistics*, **14**, 139-151.

Bowman, K.O., and L.R. Shenton, (1975). “Omnibus Test Contours for Departures from Normality Based on $\sqrt{b_1}$ and b_2,” *Biometrika*, **62**, 243-250.

Chib, S. (1995). “Marginal Likelihood from the Gibbs Output,” *Journal of the American Statistical Association*, **90**, 1313-1321

Chib, S., and I. Jeliazkov, (2001). “Marginal Likelihood from the Metropolis-Hastings Output,” *Journal of the American Statistical Association*, **96**, 270-281.

Chib, S., F. Nardari, and N. Shephard, (2002). “Markov Chain Monte Carlo Methods for Stochastic Volatility Models,” *Journal of Econometrics*, **108**, 281-316.

Fleming, J., and C. Kirby, (2003). “A Closer Look at the Relation between GARCH and Stochastic Autoregressive Volatility,” *Journal of Financial Econometrics*, 1, 365-419.

Geweke, J., (1989). “Bayesian Inference in Econometric Models Using Monte Carlo Integration,” *Econometrica*, **57**, 1317-1339.

Geweke, J., (1992). “Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments.” in *Bayesian Statistics 4*, eds. J.M. Bernardo, J. Berger, A.P. Dawid, and A.F.M. Smith, Oxford, U.K.: Oxford UP, 169-193.

Hansen, B., (2004). “The Bootstrap” in *Econometrics,* Graduate lecture notes, Economics Department, University of Wisconsin, 70-87.

Jacquier, E., N. Polson, and P. Rossi, (2002). “Bayesian Analysis of Stochastic Volatility Models,” *Journal of Business and Economic Statistics*, **20**, 69-87.

Kass, R.E., and A.E. Raftery, (1995). “Bayes Factors,” *Journal of the American Statistical Association*, **90**, 773-795.

Kim, S., N. Shephard, and S. Chib, (1998). “Stochastic Volatility: Likelihood Inference and Comparison with ARCH models,” *Review of Economic Studies*, **65**, 361-393.

Lakonishok, J., and S. Smidt, (1988). "Are Seasonal Anomalies Real? A Ninety-Year Perspective," *Review of Financial Studies*, **1**, 403-425.

Pitt, M.K., and N. Shephard, (1999). "Filtering via Simulation Based on Auxiliary Particle Filters," *Journal of the American Statistical Association*, **94**, 590-599.

Rosenblatt, M., (1952). "Remarks on a Multivariate Transformation," *Annals of Mathematical Statistics*, **23**, 470-2.

Sandmann, G., and S.J. Koopman, (1998). "Estimation of Stochastic Volatility Models via Monte Carlo Maximum Likelihood," *Journal of Econometrics*, **87**, 271-301.

Tsiakas, I., (2004a). "Is Seasonal Heteroscedasticity Real? An International Perspective," manuscript, Warwick Business School, University of Warwick, UK.

Tsiakas, I., (2004b). "Periodic Stochastic Volatility and Fat Tails," manuscript, Warwick Business School, University of Warwick, UK.

A APPENDIX A

A.1 One-Sided Hypothesis Testing using Bootstrapping[9]

Our methodology is based on bootstrapping and the tools described below can be applied to the investigation of seasonality in both expected returns (by focusing on average daily returns) and volatility (by analyzing average daily absolute returns). Our formal hypotheses are all one-sided. Since we know whether each calendar effect is higher or lower than its complement, we wish to test the statistical significance of that difference only in the direction it is actually observed.

We test the one-sided hypothesis

$$\begin{aligned} H_0 &: \theta_1 = \theta_2 \quad or \quad \theta_1 - \theta_2 = 0 \\ H_1 &: \theta_1 < \theta_2 \quad or \quad \theta_1 - \theta_2 < 0 \end{aligned} \tag{A.1}$$

at size α. We construct the test statistic

$$T = \frac{\widehat{\theta}_1 - \widehat{\theta}_2}{s\left(\widehat{\theta}_1 - \widehat{\theta}_2\right)} \tag{A.2}$$

and reject in favour of H_1 if $T < c$. The standard error in testing for the difference between two sample means for unequal sample sizes, different population variances, and independent groups is computed as

$$s\left(\widehat{\theta}_1 - \widehat{\theta}_2\right) = \sqrt{\frac{\widehat{\sigma}_1^2}{T_1 - 1} + \frac{\widehat{\sigma}_2^2}{T_2 - 1}} \tag{A.3}$$

The critical value c is selected so that $\Pr(T < c) = a$ or $c = q_\alpha$, where q_α is the quantile of the empirical distribution of test statistic T at the significance level α. Since q_α is unknown, a bootstrap

[9] For the details on hypothesis testing using bootstrapping see Hansen (2003).

test replaces it with the bootstrap estimate q_{α}^{B} and the test rejects if $T < q_{\alpha}^{B}$. Similarly, if the alternative is $H_1 : \theta_1 > \theta_2$ or $\theta_1 - \theta_2 > 0$, the bootstrap test rejects if $T > q_{1-\alpha}^{B}$.

Computationally, the critical value can be estimated from a bootstrap simulation by sorting the bootstrap t-statistics

$$T = \frac{\left(\widehat{\theta}_{1,b}^{*} - \widehat{\theta}_{2,b}^{*}\right) - \left(\widehat{\theta}_1 - \widehat{\theta}_2\right)}{s\left(\widehat{\theta}_1^{*} - \widehat{\theta}_2^{*}\right)} \tag{A.4}$$

where $\widehat{\theta}_{1,b}^{*}$ is the sample mean of θ_1 in the b'th of a total of B bootstrap samples. It is important to note that the bootstrap test statistic is centered at the estimate $\widehat{\theta}_1 - \widehat{\theta}_2$, and the standard error $s\left(\widehat{\theta}_1^{*} - \widehat{\theta}_2^{*}\right)$ is calculated on the bootstrap samples as

$$s\left(\widehat{\theta}_1^{*} - \widehat{\theta}_2^{*}\right) = \sqrt{Var\left(\widehat{\theta}_1^{*} - \widehat{\theta}_2^{*}\right)}, \qquad Var\left(\widehat{\theta}_1^{*} - \widehat{\theta}_2^{*}\right) = \frac{1}{B}\sum_{b=1}^{B}\left\{\left(\widehat{\theta}_{1,b}^{*} - \widehat{\theta}_{2,b}^{*}\right) - \left(\overline{\widehat{\theta}}_1^{*} - \overline{\widehat{\theta}}_2^{*}\right)\right\}^2 \tag{A.5}$$

where $\overline{\widehat{\theta}}_1^{*}$ is the average of the bootstrap means across all the B bootstrap samples. Note that even though we generate the same number of bootstrap samples B for both variables, $\widehat{\theta}_1$ and $\widehat{\theta}_2$ (and hence $\widehat{\theta}_{1,b}^{*} - \widehat{\theta}_{2,b}^{*}$) are constructed using different original sample sizes T. We set $B = 10,000$ bootstrap samples. These t-statistics are then sorted to find the estimated quantiles q_{α}^{B} or $q_{1-\alpha}^{B}$.[10]

B APPENDIX B: Sampling the informational Level Volatility Effects

We perform Bayesian MCMC estimation by constructing a Markov chain whose limiting distribution is the target posterior density of interest. This Markov chain is a Gibbs sampler in which all parameters are drawn sequentially from their full conditional posterior distribution. The chain is then iterated a large number of times and the sampled draws, beyond a burn-in period, are treated as variates from the target distribution. For the *iSV* model, the Bayesian MCMC algorithm produces estimates of the posterior means of (i) the set parameters of the return equation $\theta_1 = \{\alpha, \beta_D, \beta_N, \beta_W, \beta_H, \beta_{LW}\}$, (ii) the set of parameters of the plain vanilla Gaussian log-variance process $\theta_2 = \left\{\mu, \phi, \sigma^2\right\}$, and (iii) the set of volatility parameters which measure the relative effect of overnight information $\theta_3 = \{\gamma_D, \gamma_N, \gamma_W, \gamma_H, \gamma_{LW}\}$.

The key to estimating the *iSV* models is the efficient sampling of the overnight level effects θ_3 in the conditional variance. The paper uses the simple Gibbs step addition of Tsiakas (2004) to the Chib, Nardari and Shephard (2002) algorithm, in which the γ_j vector is drawn conditional on the log-volatilities $\{h_t\}$ using a precision-weighted average of prior information and the conditional likelihood. The details on sampling γ_j are summarized in Appendix B.

The parameters of all SV models examined in this paper are estimated using the Bayesian MCMC tools of Chib, Nardari and Shephard (2002), which build on the procedures developed by Kim, Shephard and Chib (1998). The algorithm constructs a Markov chain whose limiting distribution is the target posterior density of interest. Here, the Markov chain is a Gibbs sampler where all

[10] For more details on hypothesis testing using bootstrapping see Hansen (2004).

parameters are drawn sequentially from their full conditional posterior distribution. The Gibbs sampler is iterated a large number of times and the sampled draws, beyond a burn-in period, are treated as variates from the target distribution.

The information SV specification require estimation of a relatively high-dimensional parameter vector. In particular, for the iSV model, the Bayesian MCMC algorithm produces estimates of the posterior means of (i) the set parameters of the return equation $\theta_1 = \{\alpha, \beta_D, \beta_N, \beta_W, \beta_H, \beta_{LW}\}$, (ii) the set of parameters of the plain vanilla Gaussian log-variance process $\theta_2 = \left\{\mu, \phi, \sigma^2\right\}$, and (iii) the set of volatility parameters which measure the relative effect of overnight information $\theta_3 = \{\gamma_D, \gamma_N, \gamma_W, \gamma_H, \gamma_{LW}\}$. The key to estimating the high-dimensional iSV model is the efficient sampling of the the overnight volatility effects. This is done using a simple Gibbs step of Tsiakas (2004) where the θ_3 vector is drawn conditional on the log-variance vector $\left\{h_t^D\right\}$ using a precision-weighted average of prior information and the conditional likelihood.

B.1 A brief sketch of the MCMC algorithm

1. Initialize θ, s, λ, ν and transform the data into $r_t^* = \log\left(\frac{\omega_t}{\lambda_t^{-1}} + c\right)$, $c = 0.001$ to put the model in state-space form. The "offset" constant c eliminates the inlier problem.

2. Sample all the log-volatility parameters from their full conditional posterior density: $\theta_2 \mid r^*, s, \gamma$. This posterior is not available analytically. We use the Kalman filter to compute the log-likelihood of transformed data r_t^* as a function of θ (conditional on s_t) and then optimize this conditional log-posterior. We generate a proposal from a t-distribution $t\left(m, V, \xi\right)$ where m is the mode, V is the inverse of the negative Hessian and ξ a tuning parameter. The proposal is then accepted according to the Metropolis-Hastings algorithm. The optimization step makes this an independence chain M–H algorithm and goes a long way in reducing the autocorrelation in the draws of the MCMC chain. For more details on the M-H algorithm see Chib and Greenberg (1995).

3. Sample the seasonal coefficients in the log-variance equation from their full conditional posterior $\gamma \mid r^*, D, h, s$ using the Gibbs step detailed below separately.

4. Sample the log-volatility vector $\{h_t\}$ in one block from the full conditional posterior distribution: $h \mid r^*, s, \theta$. This step uses the de Jong and Shephard (1995) simulation smoother which is an algorithm designed for efficient sampling of the state vector in a state-space model.

5. Sample the degrees of freedom parameter of the conditional distribution from the full conditional density: $\nu \mid r, h, \theta$. Again, we optimize the conditional log-posterior with respect to ν and then use the mode and a scaled inverse of the negative Hessian to generate a proposal that is accepted according to the Metropolis-Hastings algorithm. This independence chain M-H algorithm is also crucial in contributing to low Relative Numerical Inefficiencies for the parameters of interest.

6. Sample $\lambda \mid r, h, \theta$ directly from its posterior:

$$\lambda_t \mid r_t, h_t, \theta \sim Gamma\left(\frac{\nu+1}{2}, \frac{2}{\nu + \omega_t^2/v_t^2}\right) \tag{B.6}$$

7. Sample all the conditional mean coefficients (including the seasonal coefficients in the mean) $\theta_1, \delta \mid r, D, h, \lambda$ simply using a precision-weighted average of a set of normal priors and the normal conditional likelihood. Then update the transformed data $r_t^* = \log\left(\frac{\omega_t^2}{\lambda_t^{-1}} + c\right)$, $c = 0.001$.

8. Finally, sample the mixture indicator variable $s \mid r^*, h, \theta$ directly from its posterior:

$$\Pr(s_t \mid r_t^*, h_t) \propto \Pr(s_t)\, \phi\left(r_t^* \mid h_t + \gamma' D_t + m_{st}, v_{st}^2\right), \qquad t \leq T \tag{B.7}$$

9. Go to step 2 and iterate.

Table 1
Descriptive Statistics of Day and Night Returns and Volatility

Panel A: Daily Returns (%)

	Canada (*TSX Comp*)		*Germany* (*DAX 30*)		*HongKong* (*Hang Seng*)		*Japan* (*Nikkei 225*)		*UK* (*FTSE 100*)		*USA* (*DJIA*)	
Start	1/1/1986		1/1/1989		1/1/1991		1/1/1986		1/1/1984		1/1/1992	
Obs (*T*)	4528		3752		2658		4434		5039		3018	
	DAY	*NIGHT*	*DAY*	*NIGHT*	*DAY*	*NIGHT*	*DAY*	*NIGHT*	*DAY*	*NIGHT*	*DAY*	*NIGHT*
Mean	0.014	0.009	−0.025	0.055	0.052	−0.007	−0.041	0.037	0.017	0.013	0.039	$8e-5$
Std Dev	0.708	0.515	1.23	0.715	1.21	1.20	1.33	0.362	0.963	0.500	1.03	0.155
Min	−9.94	−8.90	−9.10	−9.64	−8.57	−12.0	−16.1	−2.75	−5.89	−9.50	−7.44	−2.06
Max	5.82	5.42	7.40	4.94	7.00	12.8	12.4	2.42	5.90	6.09	6.21	4.37
Skewness	−0.810	−2.15	−0.194	−0.965	−0.347	−0.150	−0.090	0.223	−0.196	−1.80	−0.275	7.14
Kurtosis	17.9	42.09	9.07	19.6	7.28	15.9	12.4	7.56	6.84	66.6	7.67	244.0
$Corr\,(r_t, r_{t-1})$	0.072	−0.028	−0.029	−0.051	−0.056	−0.002	−0.026	0.041	−0.015	−0.062	−0.007	−0.030
$Corr\,(r_t, r_{t-2})$	−0.002	−0.048	−0.017	−0.027	−0.021	−0.061	−0.073	−0.039	−0.006	−0.039	−0.016	−0.027
$Corr\,(r_t, r_{t-3})$	0.009	−0.003	0.002	−0.015	0.051	0.057	0.021	−0.021	−0.055	−0.001	−0.011	0.043
$Corr\,(r_t, r_{t-10})$	0.015	0.019	−0.028	0.011	0.023	−0.003	0.023	0.019	−0.009	0.023	0.027	0.063
$Corr\,(r_t, r_{t-25})$	0.010	0.029	0.001	−0.009	−0.003	0.004	0.026	0.018	0.002	0.014	−0.049	−0.043
$Corr\left(r_t^D, r_{t-1/2}^N\right)$	0.04		0.13		0.12		0.23		−0.03		0.05	

Panel B: Daily Absolute Returns (%)

	DAY	*NIGHT*	*DAY*	*NIGHT*	*DAY*	*NIGHT*	*DAY*	*NIGHT*	*DAY*	*NIGHT*	*DAY*	*NIGHT*
Mean	0.479	0.304	0.797	0.470	0.875	0.755	0.933	0.240	0.693	0.254	0.740	0.045
Std Dev	0.521	0.416	0.937	0.541	0.835	0.939	0.948	0.273	0.669	0.431	0.721	0.148
Min	0.00	0.00	0.00	0.00	0.00	0.00	0.001	0.00	0.00	0.00	0.00	0.00
Max	9.94	8.90	9.10	9.64	9.22	12.8	16.1	2.75	5.90	9.50	7.44	4.37
Skewness	3.99	5.99	2.67	3.92	2.40	3.51	3.21	2.31	2.38	7.25	2.50	12.5
Kurtosis	40.2	77.7	13.1	40.3	13.9	27.6	28.51	10.4	11.9	101.0	14.4	276.0
$Corr\,(r_t, r_{t-1})$	0.293	0.326	0.367	0.323	0.131	0.374	0.213	0.343	0.215	0.351	0.165	0.271
$Corr\,(r_t, r_{t-2})$	0.249	0.310	0.418	0.305	0.171	0.321	0.211	0.337	0.266	0.259	0.197	0.321
$Corr\,(r_t, r_{t-3})$	0.229	0.296	0.413	0.314	0.156	0.312	0.212	0.373	0.258	0.251	0.199	0.279
$Corr\,(r_t, r_{t-10})$	0.224	0.221	0.380	0.257	0.098	0.211	0.146	0.343	0.217	0.232	0.189	0.313
$Corr\,(r_t, r_{t-25})$	0.156	0.166	0.357	0.256	0.067	0.183	0.101	0.348	0.177	0.160	0.142	0.256
$Corr\left(r_t^D, r_{t-1/2}^N\right)$	0.33		−0.05		0.21		0.10		−0.008		0.06	

The Day variable denotes all the open to close periods and the Night variable contains all the lagged close to open periods. In computing these descriptive statistics, we use the entire sample of available data. We include all zero night returns for which the day return is non-zero and vice versa. The end date for all indices is December 31, 2003, except in the case of Hong Kong for which the available opening price data ends at 28/9/2001.

Table 2
Volatility Proportions

	NIGHT-VP	*DAY-VP*
Canada	40%	60%
Germany	43%	57%
Hong Kong	43%	57%
Japan	26%	74%
UK	29%	71%
USA	7%	93%

We define $DAY\text{-}VP$ as the average volatility proportion over days: $vp^D = \frac{1}{T}\sum_{t=1}^{T}\frac{|r_t^D|}{|r_t^D|+|r_{t-1/2}^N|}$. Similarly the $NIGHT\text{-}VP$ is defined as $vp^N = \frac{1}{T}\sum_{t=1}^{T}\frac{|r_{t-1/2}^N|}{|r_t^D|+|r_{t-1/2}^N|}$.

Table 3
Statistical Significance of Days, Nights, Weekends, Holidays and Long Weekends

Panel A: Expected Returns

	Canada	*Germany*	*Hong Kong*	*Japan*	*UK*	*USA*
Day	0.017 (.008, .034)	-0.022^{***} (−.054, .010)	0.053^{**} (0.014, 0.091)	0.038^{***} (−.070, −.004)	0.021 (−.007, .043)	0.042^{**} (.011, .072)
Night	0.011 (−.002, .023)	0.055^{***} (.035, .073)	-0.005^{**} (−.044, .033)	0.037^{***} (.028, .046)	0.015 (.004, .026)	0.002^{**} (−.004, .005)
Weekend	-0.016^{**} (−.044, .011)	0.027 (−.028, .080)	0.007 (−.096, .107)	0.046 (.020, .071)	-0.055^{***} (−.088, −.022)	0.011^{*} (−.001, .027)
Holiday	0.101 (−.035, .234)	0.153 (−.012, .321)	0.307^{*} (−.054, .694)	0.126^{***} (.069, .186)	0.032 (−.147, 0.205)	−0.004 (−.018, 0.008)
Long Weekend	0.029 (−.069, .132)	−0.032 (−.232, .165)	−0.094 (−.412, .228)	0.151^{***} (.074, .231)	−0.006 (−.124, .097)	0.003 (−.019, .027)

Panel B: Volatility

	Canada	*Germany*	*Hong Kong*	*Japan*	*UK*	*USA*
Day	0.475^{***} (.463, .487)	0.793^{***} (.769, .818)	0.875^{***} (.849, .902)	0.927^{***} (.902, .946)	0.688^{***} (.673, .703)	0.739^{***} (.717, .760)
Night	0.300^{***} (.290, .310)	0.470^{***} (.456, .485)	0.752^{***} (.723, .782)	0.240^{***} (.233, .247)	0.249^{***} (.241, .259)	0.045^{***} (.041, .050)
Weekend	0.299 (.277, .321)	0.544^{***} (.503, .588)	0.861^{***} (.785, .941)	0.283^{***} (.264, .303)	0.278^{**} (.250, .309)	0.044 (.032, .060)
Holiday	0.309 (.217, .412)	0.546 (.430, .671)	0.959 (.688, 1.26)	0.245 (.201, .294)	0.335 (.200, .492)	0.020^{***} (.011, .034)
Long Weekend	0.450^{***} (.379, .527)	0.583^{*} (.448, .730)	1.04^{**} (.815, 1.29)	0.350^{***} (.291, .411)	0.332^{*} (.244, 432)	0.053 (.033, .075)

The expected returns are proxied by the average daily returns and volatility is proxied by the average absolute daily returns. The Day variable denotes all the open to close periods and the Night variable contains all the lagged close to open periods. Weekend is the Friday close to a Monday open, Holiday is the pre-holiday close to the post-holiday open, and Long Weekend is close to open period which contains both a weekend and a (Friday or Monday) holiday. Statistical significance is assessed by forming one-sided bootstrap-based hypothesis tests, which compare each variable to its complement: day vs. night, weekend vs. non-weekend nights, holiday vs. non-holiday nights, and long weekend vs. non-long weekend nights. The numbers in parenthesis are the 5% and 95% quantiles for the means generated by 10,000 bootstrap samples. The superscripts *, **, and *** indicate that the relevant one-sided null hypothesis is rejected at significance level $\alpha = 10\%$, $\alpha = 5\%$, and $\alpha = 1\%$, respectively. In assessing statistical significance, we have excluded the returns data from the two weeks of October 19-23, 1987 and September 10-14, if applicable. For the US DJIA data we have also excluded September 17, 2001, which was the first trading day in the US after 9/11.

Table 4

Posterior Means of all *iSV* parameters

Panel A: Conditional Mean Parameters

	Canada	*Germany*	*Hong Kong*	*Japan*	*UK*	*USA*
α	0.025*** (.013, .037)	−0.007 (−.022, .009)	0.057*** (.025, .090)	−0.053*** (−.075, −.030)	0.032*** (.015, .051)	0.047*** (.024, .070)
β_D	0.130*** (.104, .156)	−0.064*** (−.091, −.036)	−0.056*** (−.089, −.023)	−0.026** (−.051, −.001)	−0.031** (−.057, −.006)	0.007 (−.024, .039)
β_{WN}	0.039* (−.006, .084)	0.180*** (.150, .210)	0.045** (.003, .086)	0.714*** (.615, .810)	−0.058** (−.106, −.012)	0.251** (.013, .493)
β_{WE}	0.162*** (.068, .256)	0.262*** (.214, .312)	0.114*** (.047, .179)	0.395*** (.266, .527)	−0.142*** (−.227, −.059)	−0.003 (−.060, .063)
β_H	0.361* (−.078, .789)	0.142* (−.028, .317)	0.018 (−.203, .235)	1.98*** (1.37, 2.61)	0.503** (.064, .777)	0.479 (−1.08, 2.04)
β_{LW}	0.151** (.001, .304)	0.375*** (.221, .496)	0.043 (−.084, .181)	1.24*** (.926, 1.55)	−0.255** (−.477, −.031)	0.998* (−.102, 2.09)

Panel B: Log-Variance Parameters

	Canada	*Germany*	*Hong Kong*	*Japan*	*UK*	*USA*
μ	−1.18***. (−1.37, −.990)	−0.907*** (−1.42, −.438)	0.100 (−.101, .313)	−0.094 (−.275, .084)	−0.399*** (−.562, −.239)	−0.173* (−.381, .042)
ϕ	0.972*** (.963, .980)	0.992*** (.988, .996)	0.968*** (.954, .980)	0.968*** (.960, .976)	0.981*** (.976, .986)	0.978*** (.970, .985)
σ^2	0.041*** (.031, .052)	0.025*** (.019, .031)	0.036*** (.025, .050)	0.042*** (.033, .053)	0.019*** (.015, .023)	0.023*** (.017, .030)
γ_D	−0.076*** (−.100, −.051)	−0.043*** (−.058, −.027)	−0.017* (−.038, .003)	−0.087*** (−.103, −.072)	−0.057*** (−.073, −.041)	−0.098*** (−.119, −.078)
γ_{WN}	−0.164*** (−.243, −.085)	−0.096** (−.166, −.024)	0.006 (−.040, .050)	−0.068 (−.198, .057)	−0.141** (−.240, −.042)	0.123 (−.192, .430)
γ_{WE}	−0.491*** (−.643, -.342)	−0.303*** (−.385, −.222)	−0.059 (−.149, .031)	−0.208** (−.381, −.030)	−0.122** (−.216, −.030)	−1.46*** (−2.08, −.683)
γ_H	0.132 (−.708, .988)	0.221 (−.199, .656)	0.096 (−.141, .341)	0.762** (.108, 1.43)	0.858* (−.017, 1.77)	0.046 (−1.57, 1.69)
γ_{LW}	−0.412** (−.710, −.104)	−0.637** (−1.19, −.059)	0.195** (.036, .367)	−0.202 (−.669, .266)	−0.229 (−.745, .265)	0.250 (−.908, 1.43)

The posterior means are the Bayesian MCMC estimates. The MCMC chain run for 5,000 iterations after an initial burn-in of 1,000 iterations. The numbers in parenthesis indicate the 5% and 95% percentiles of the MCMC draws. The MCMC results are based on a set of adjusted samples, which have excluded the returns data from the two weeks of October 19-23, 1987 and September 10-14, if applicable. For the US DJIA data we have also excluded September 17, 2001, which was the first trading day in the US after 9/11.

Table 5
Night vs. Weeknight Effects

	Conditional Mean		**Log-Variance**	
	iSV_n	iSV	iSV_n	iSV
	β_N	β_{WN}	γ_N	γ_{WN}
Canada	0.073^{***}	0.039^{*}	-0.249^{***}	-0.164^{***}
	(.033, .112)	(−.006 .084)	(−.319, −.179)	(−.243, −.085)
Germany	0.201^{***}	0.180^{***}	-0.177^{**}	-0.096^{**}
	(.174, .226)	(.150, .210)	(−.227, −.126)	(−.166, −.024)
Hong Kong	0.065^{***}	0.045^{**}	0.006	0.006
	(.031, .098)	(.003, .086)	(−.032, .043)	(−.040, .050)
Japan	0.676^{***}	0.714^{***}	-0.057	-0.068
	(.601, .753)	(.615, .810)	(−.157, .046)	(−.198, .057)
UK	-0.084^{***}	-0.058^{**}	-0.128^{***}	-0.141^{**}
	(−.125, −.043)	(−.106, −.012)	(−.192, −.065)	(−.240, −.042)
USA	0.232^{**}	0.251^{**}	-0.089	0.123
	(.018, .449)	(.013, .493)	(−.374, .200)	(−.192, .430)

Table 6
Conditional SV Dynamics

		Log-Likelihood	**Density Forecasts** $\{n_{t+1}\}$			
		$\log -L$	VAR	$SKEW$	$KURT$	$BL(30)$
Canada	SV	−5704.2	1.52	5.86	82.9	79.9
	iSV_n	−5670.2	1.53	2.54	83.4	80.9
	iSV	−5663.8	1.49	−1.44	73.4	81.3
Germany	SV	−9538.5	2.93	13.9	92.0	210.6
	iSV_n	−10030.3	3.00	13.1	89.1	201.1
	iSV	−10358.4	3.15	14.2	87.7	212.7
Hong Kong	SV	−4686.3	1.27	−2.72	44.8	53.7
	iSV_n	−4681.3	1.25	−5.29	36.7	48.7
	iSV	−4648.4	1.24	−5.46	36.1	50.1
Japan	SV	−8774.4	1.50	6.99	87.8	65.4
	iSV_n	−8520.8	1.49	7.41	89.2	52.4
	iSV	−8508.1	1.47	7.70	90.3	50.1
UK	SV	−7506.0	1.40	−0.03	66.0	114.9
	iSV_n	−7474.5	1.38	−2.32	62.0	113.7
	iSV	−7470.8	1.38	−2.39	62.0	114.6
USA	SV	−4451.1	1.21	−3.03	29.2	62.7
	iSV_n	−4467.1	1.22	−3.39	30.9	62.9
	iSV	−4441.0	1.20	−2.98	30.6	63.7

BL(30) is the Box-Ljung statistic at 30 lags. Note that $\chi^2(30;.90) = 40.3$ and $\chi^2(30;.95) = 43.8$.

Table 7a
Interpreting Bayes Factors

$2\ln(B_{21})$	B_{21}	Evidence against Model M_1
0 to 2	1 to 3	Not worth more than a bare mention
2 to 6	3 to 20	Positive
6 to 10	20 to 150	Strong
> 10	> 150	Very strong

Table 7b
Bayes Factors $(2\ln B_{ij})$

	iSV_n vs. SV	iSV vs. SV	iSV vs. iSV_n
Canada	52.9	41.9	−11.0
Germany	−1000	−1683	−683
Hong Kong	−5.19	31.1	36.3
Japan	496	496	−0.60
UK	47.1	30.3	−16.7
USA	−39.3	−0.20	−39.1

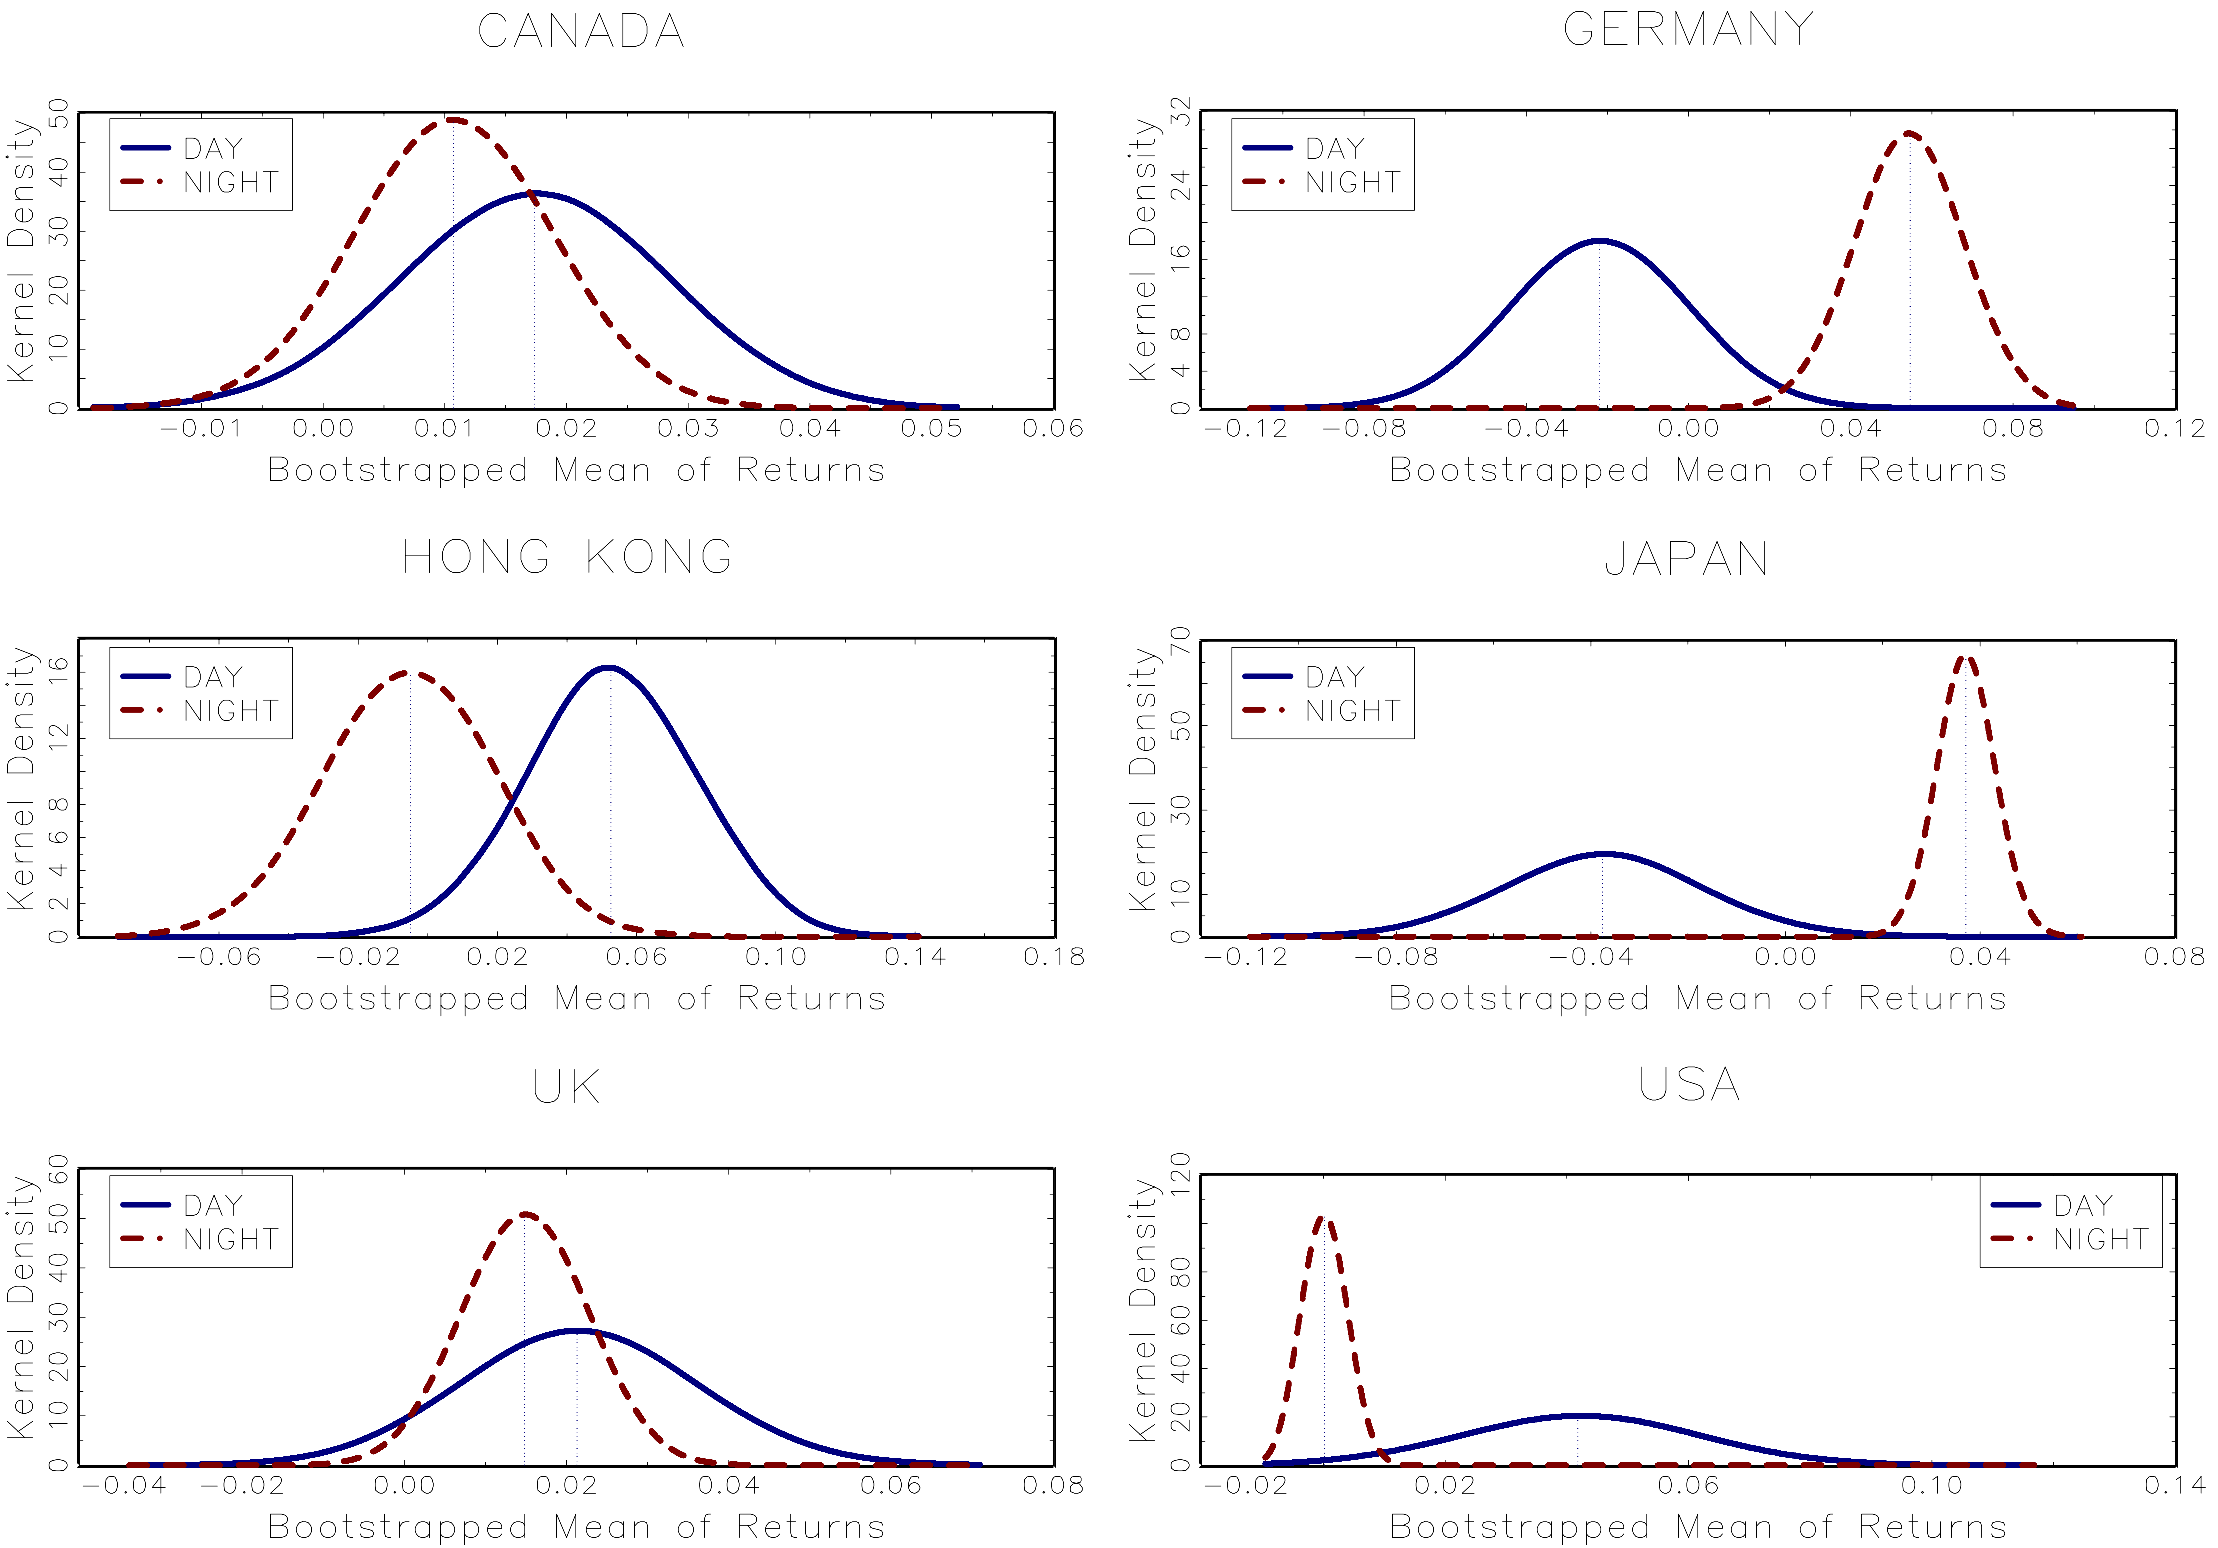
CANADA
GERMANY
HONG KONG
JAPAN
UK
USA
DAY
NIGHT
Kernel Density
Bootstrapped Mean of Returns

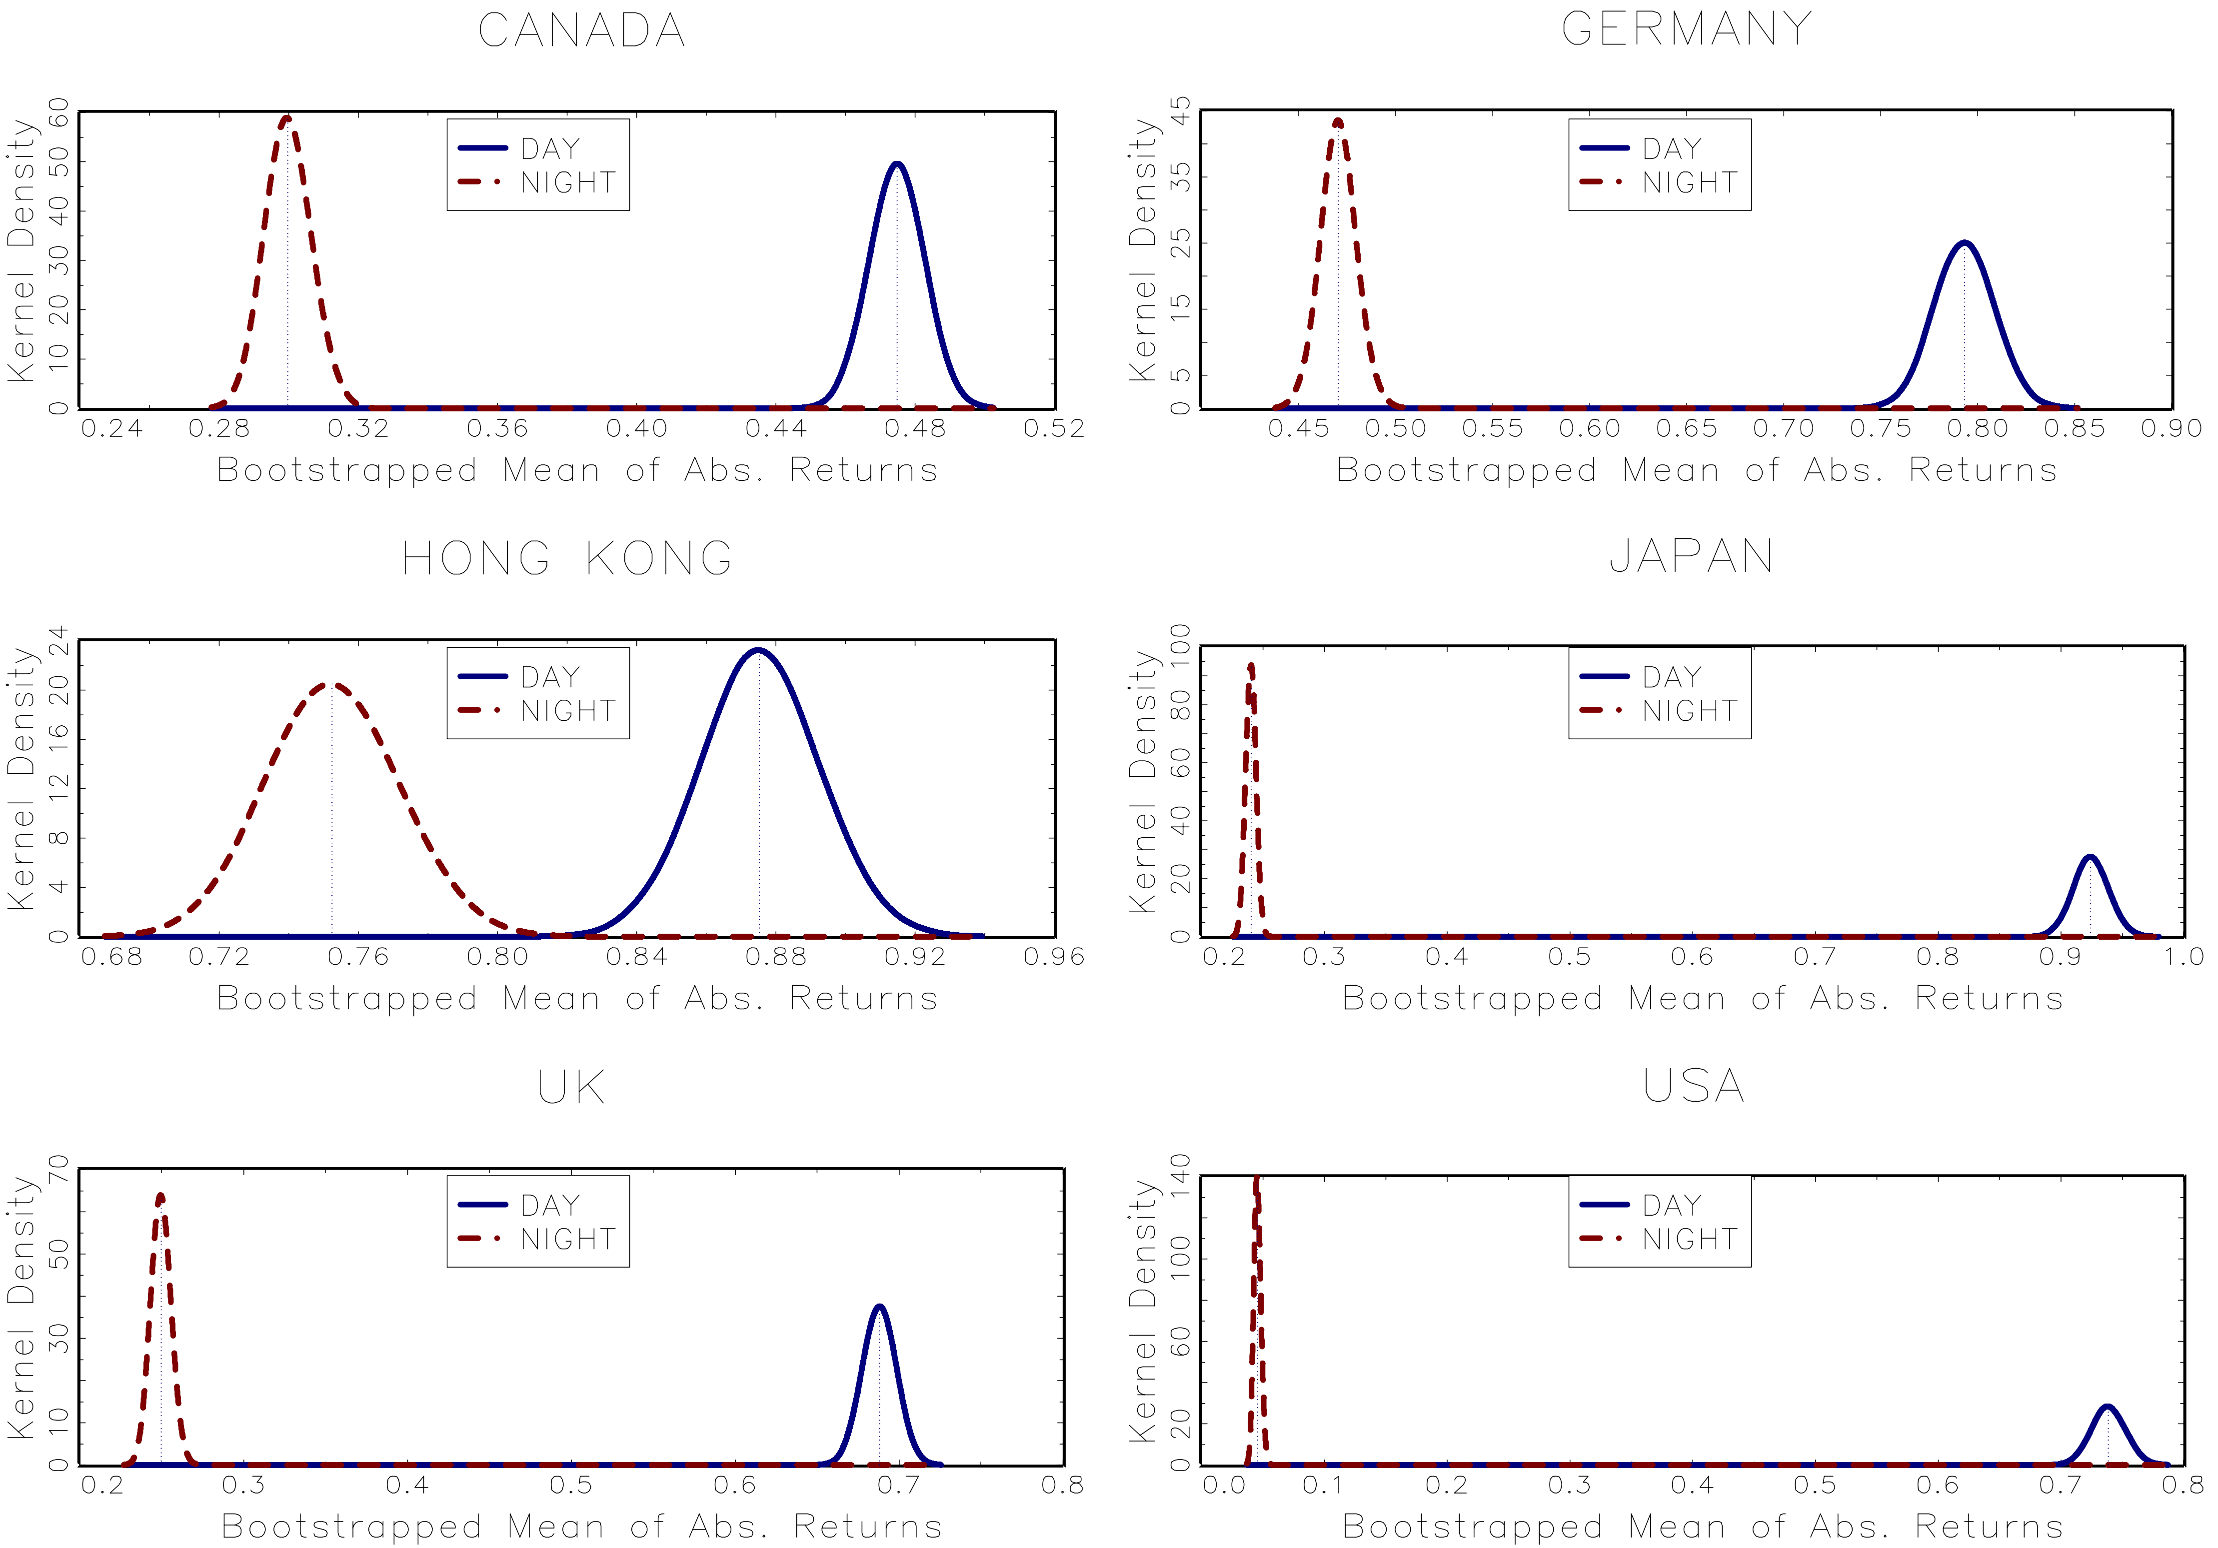
CANADA
GERMANY
HONG KONG
JAPAN
UK
USA
DAY
NIGHT
Kernel Density
Bootstrapped Mean of Abs. Returns
0.24
0.28
0.32
0.36
0.40
0.44
0.48
0.52
0.45
0.50
0.55
0.60
0.65
0.70
0.75
0.80
0.85
0.90
0.68
0.72
0.76
0.80
0.84
0.88
0.92
0.96
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1.0
0.0
0.1

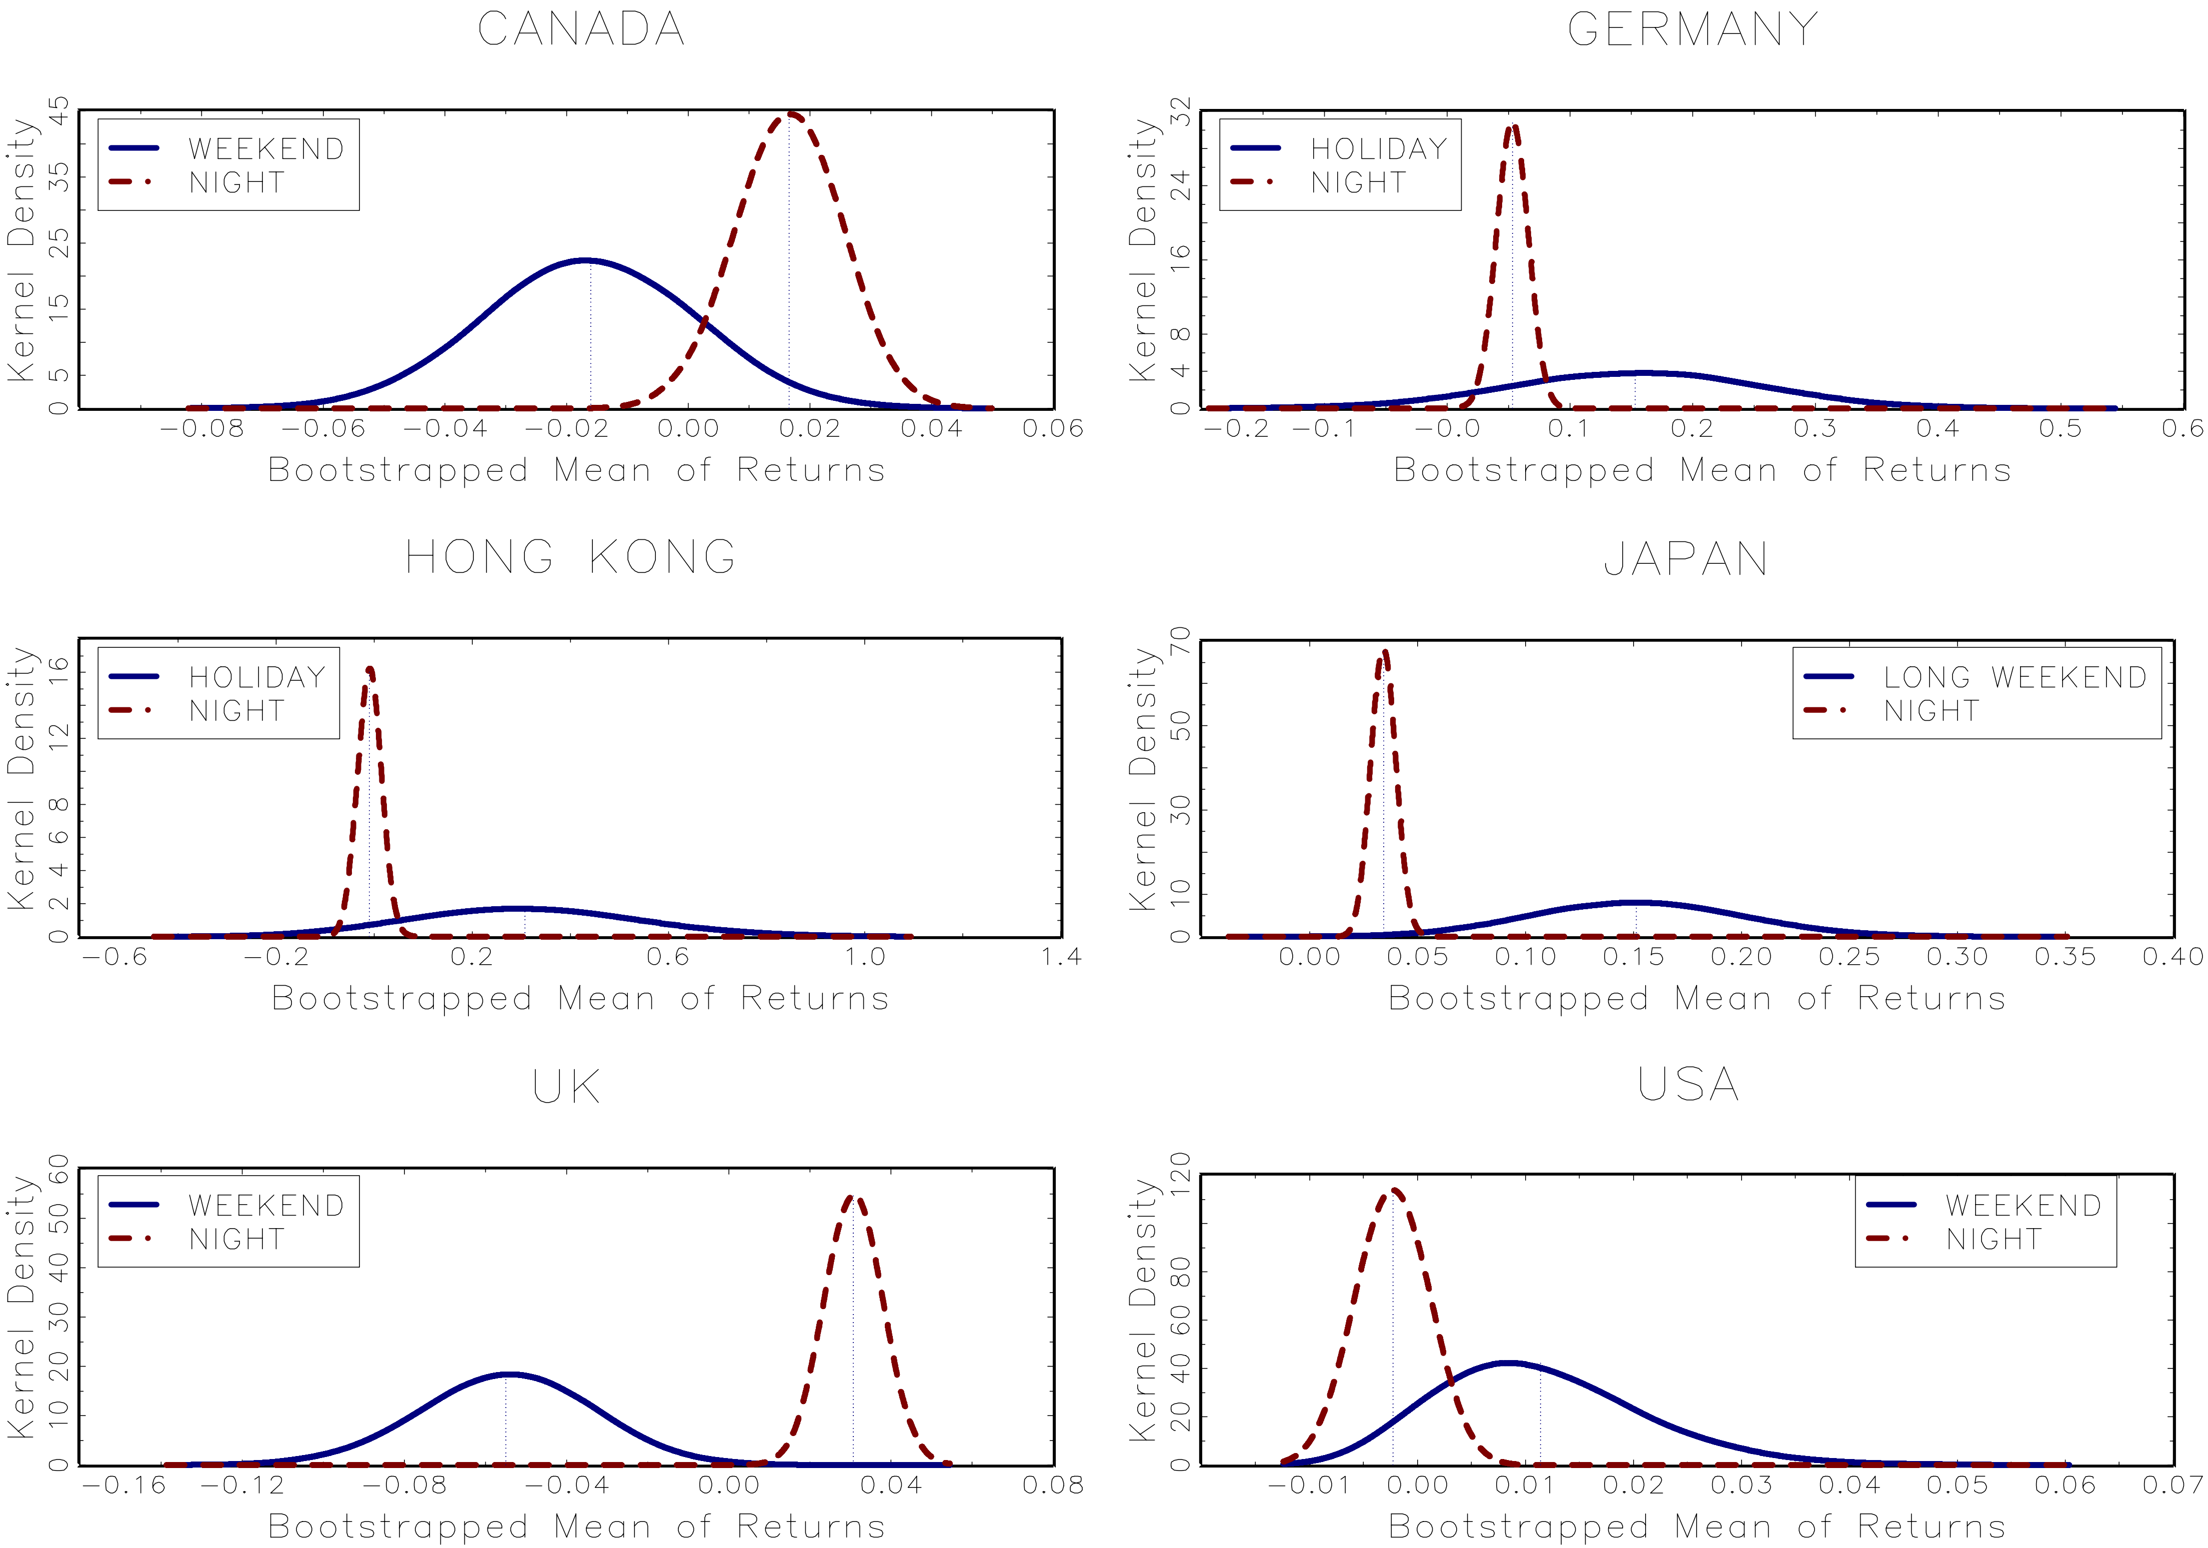
CANADA
WEEKEND
NIGHT
Kernel Density
Bootstrapped Mean of Returns
GERMANY
HOLIDAY
NIGHT
Kernel Density
Bootstrapped Mean of Returns
HONG KONG
HOLIDAY
NIGHT
Kernel Density
Bootstrapped Mean of Returns
JAPAN
LONG WEEKEND
NIGHT
Kernel Density
Bootstrapped Mean of Returns
UK
WEEKEND
NIGHT
Kernel Density
Bootstrapped Mean of Returns
USA
WEEKEND
NIGHT
Kernel Density
Bootstrapped Mean of Returns

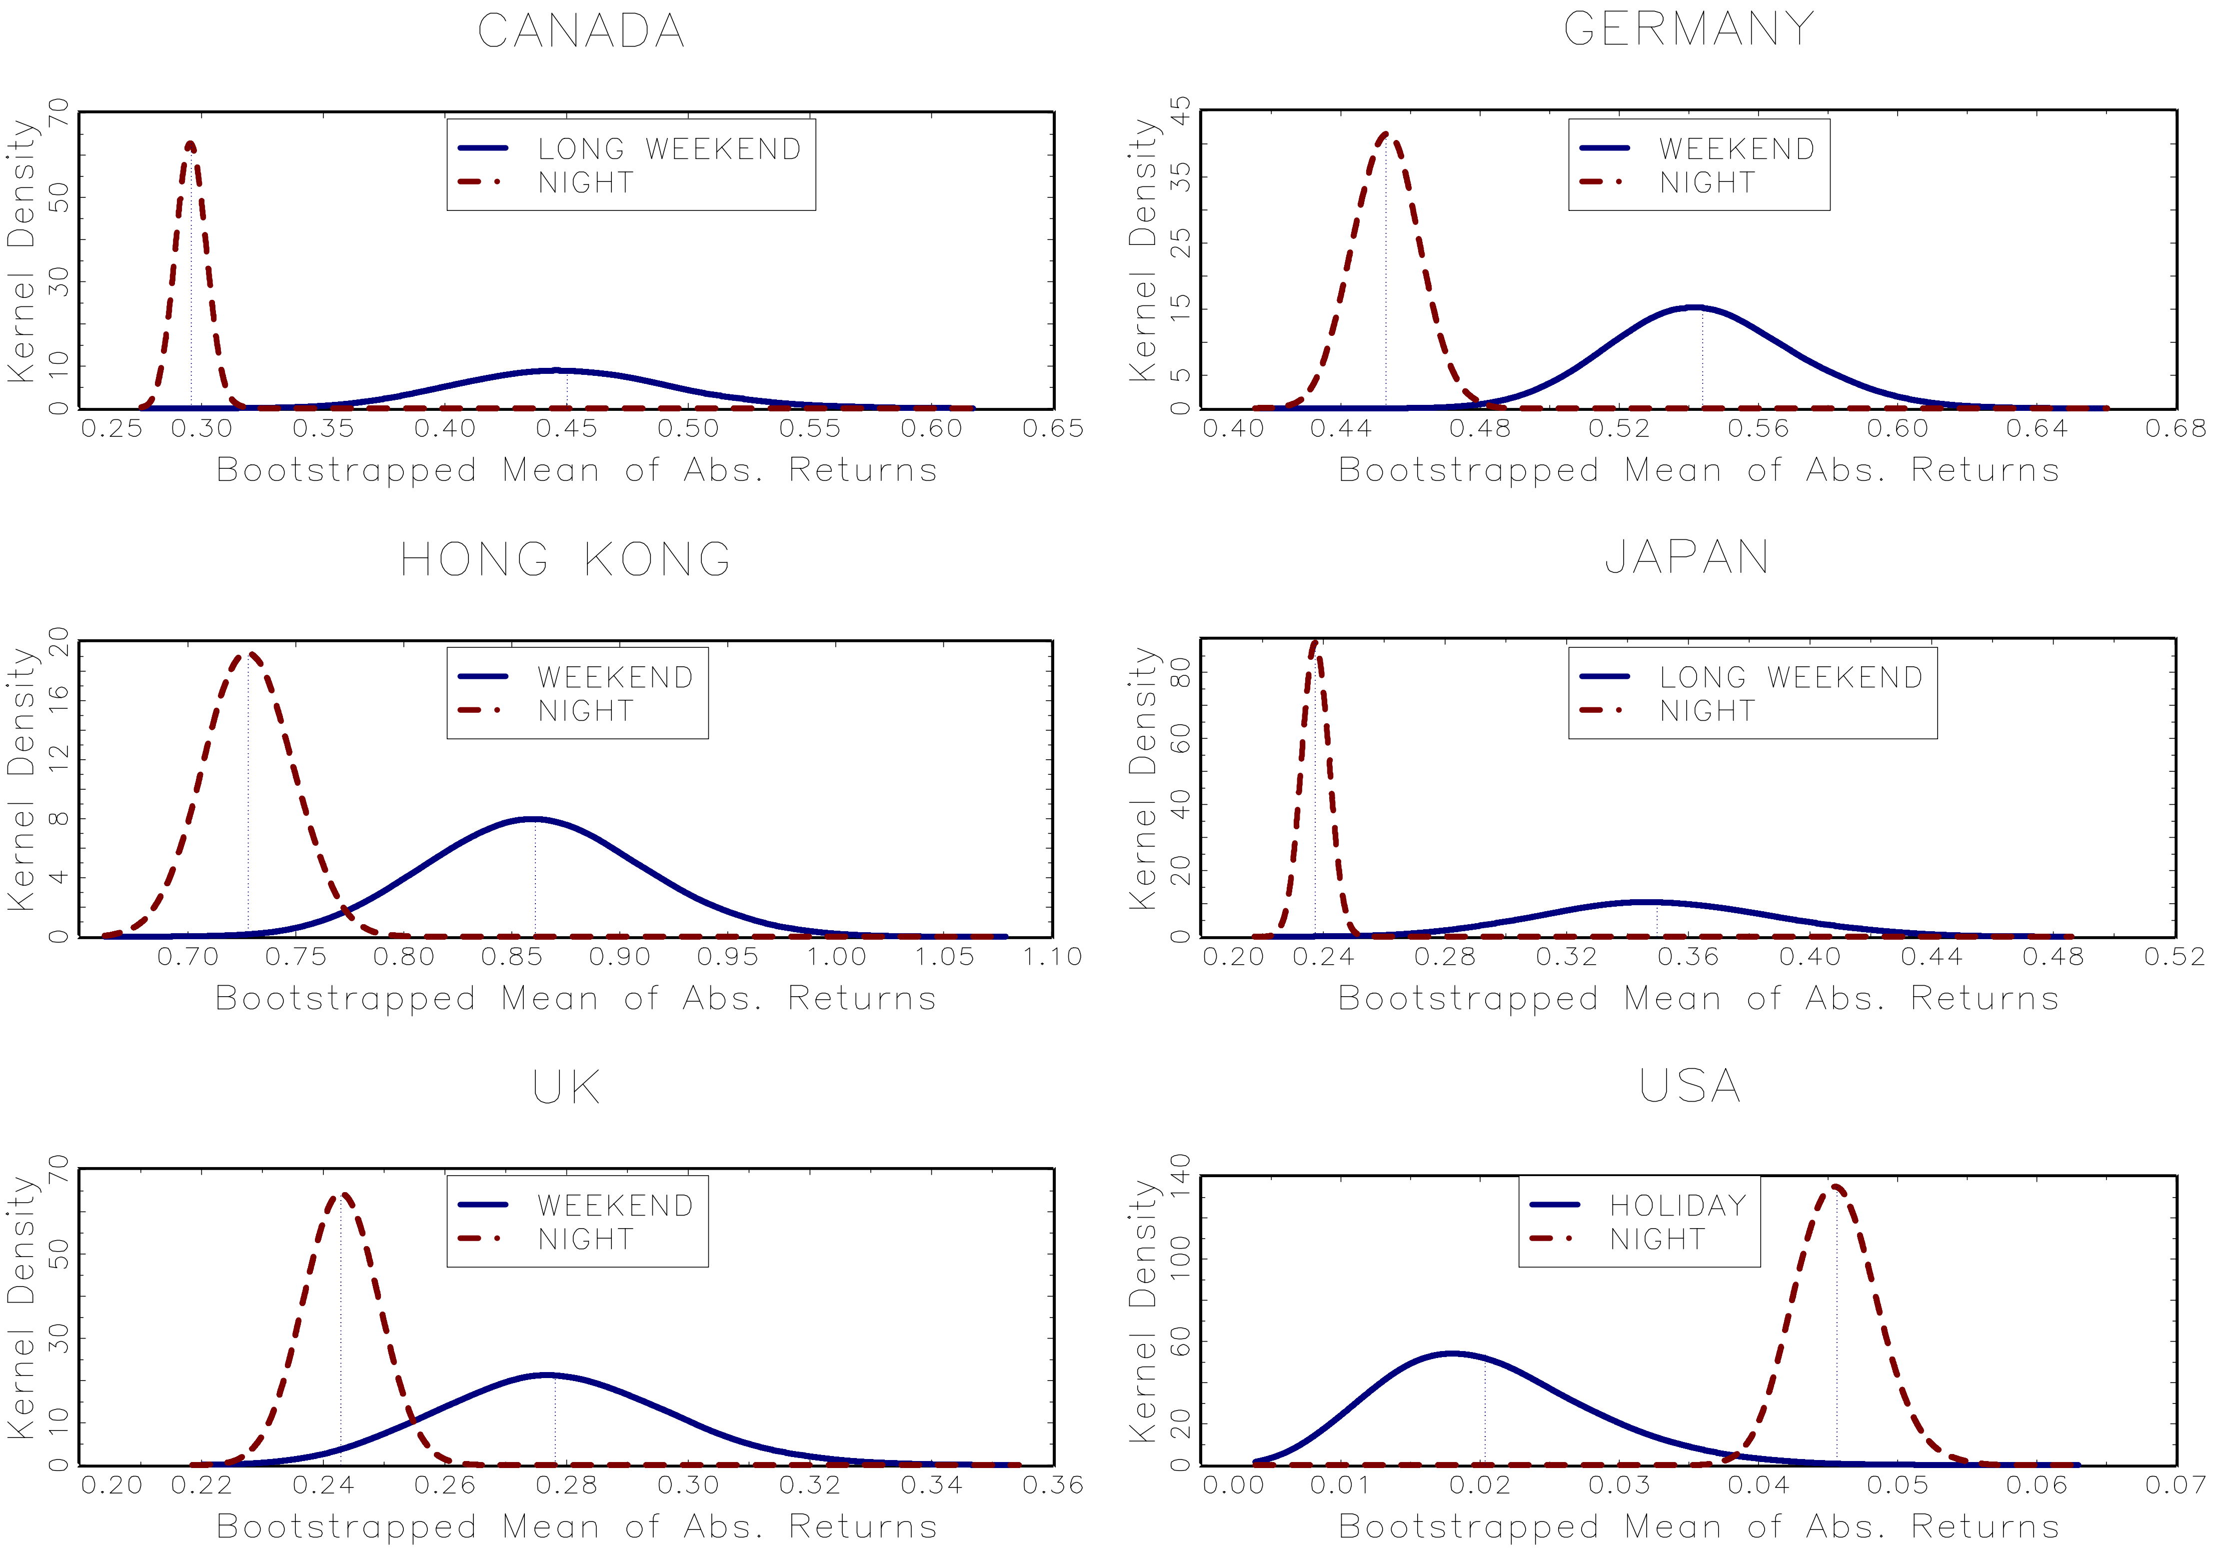
CANADA
LONG WEEKEND
NIGHT
Kernel Density
0 10 30 50 70
0.25 0.30 0.35 0.40 0.45 0.50 0.55 0.60 0.65
Bootstrapped Mean of Abs. Returns
GERMANY
WEEKEND
NIGHT
Kernel Density
0 5 15 25 35 45
0.40 0.44 0.48 0.52 0.56 0.60 0.64 0.68
Bootstrapped Mean of Abs. Returns
HONG KONG
WEEKEND
NIGHT
Kernel Density
0 4 8 12 16 20
0.70 0.75 0.80 0.85 0.90 0.95 1.00 1.05 1.10
Bootstrapped Mean of Abs. Returns
JAPAN
LONG WEEKEND
NIGHT
Kernel Density
0 20 40 60 80
0.20 0.24 0.28 0.32 0.36 0.40 0.44 0.48 0.52
Bootstrapped Mean of Abs. Returns
UK
WEEKEND
NIGHT
Kernel Density
0 10 30 50 70
0.20 0.22 0.24 0.26 0.28 0.30 0.32 0.34 0.36
Bootstrapped Mean of Abs. Returns
USA
HOLIDAY
NIGHT
Kernel Density
0 20 60 100 140
0.00 0.01 0.02 0.03 0.04 0.05 0.06 0.07
Bootstrapped Mean of Abs. Returns