Bootstrap-Based Inference in Models With a Nearly Noninvertible Moving Average Component

Nikolay Gospodinov
Department of Economics, Concordia University, Montreal, Quebec, Canada H3G 1M8
(*gospodin@vax2.concordia.ca*)

This article proposes a bootstrap method for constructing two-sided confidence intervals for the moving average (MA) parameter in nearly noninvertible models. The confidence intervals are obtained by inverting the acceptance region of the likelihood ratio (LR) test reflecting the asymmetry of the likelihood near the noninvertibility boundary. The limiting distribution of the LR statistic is nonpivotal and its quantiles are parameterized as a function of the MA parameter and then approximated by grid bootstrap. The proposed method is used to investigate the parameter instability in inflation and time variability of risk premium in interest rates.

KEY WORDS: Local-to-unity asymptotics; Near unit root process; Resampling method.

1. INTRODUCTION

Many macroeconomic time series can be modeled parsimoniously as autoregressive moving average (ARMA) processes (Schwert 1987). Time aggregation, measurement errors, errors from rational forecasts, outliers, and feedback relationships are among the reasons to justify the presence of a statistically significant moving average (MA) component in the univariate representation of the economic variables. Although these models lead to some computational difficulties, it is possible to conduct valid inference by using the standard statistical theory, provided that the MA parameter is away from the boundary of the noninvertibility region.

Some serious problems arise, however, when the time series contains an MA root near or on the unit circle. This case is of practical interest because an MA coefficient close to unity may indicate overdifferencing of the series under consideration since differencing of a trend stationary process induces noninvertibility. Also, some economic models, such as the consumption model with durable goods (Mankiw 1982) and the Cagan model of hyperinflation with rational expectations (Christiano 1987), can produce a large MA component that requires an appropriate framework for valid statistical inference. Testing for MA unit roots generated a large body of research that has led to the development of powerful tests for detecting noninvertible MA processes (Tanaka 1990; Davis, Chen, and Dunsmuir 1995). These tests, however, are concerned with the very specialized case in which the MA root is exactly equal to 1, and they do not provide information about the sampling uncertainty. Some point and interval estimators, such as median unbiased estimates and two-sided confidence intervals, might prove to be more useful and informative in applied work. Furthermore, a similar statistical apparatus is needed for evaluating the coefficient variability in time-varying parameter models (Stock and Watson 1998), which provide a flexible modeling framework for many economic and financial series.

Construction of interval estimators in models with a large MA component turns out to be a particularly difficult and nonstandard problem. To introduce the main ideas, consider the stochastic process $\{y_t\}_{t=1}^T$ generated from the MA(1) model

$$y_t = e_t - \theta_0 e_{t-1}, \tag{1}$$

where $|\theta_0| \le 1$ and $\{e_t\}_{t=0}^T$ is a martingale difference sequence with respect to the past history of y_t.

Let θ_0 be the parameter of interest and η_0 denote a possibly infinite dimensional nuisance parameter vector that completely characterizes the distribution of e. If $|\theta_0|$ is strictly less than 1, the process y_t is invertible and the maximum likelihood (ML) estimator of θ_0 is asymptotically normally distributed with mean θ_0 and variance $(1-\theta_0^2)/T$. When θ_0 is close to the unit circle, the Gaussian distribution provides a rather inaccurate approximation to the limiting behavior of the ML estimator. Moreover, in finite samples the estimator takes values exactly on the boundary of the invertibility region with positive probability (pile-up effect) when the true MA parameter is in the vicinity of 1. The observed point probability mass at unity results from the symmetry of the likelihood function around 1 and the small sample deficiency to identify all critical points near the unit circle.

The distribution of the ML estimator of θ in the presence of an MA unit root is nonstandard and was derived by Davis and Dunsmuir (1996) (DD). Recasting the MA parameter into a local-to-unity form, $\theta_0 = 1 - c/T$ for a finite constant $c \ge 0$, provides a useful framework for analyzing the limiting behavior of the ML estimator and allows a smooth transition from the normal approximation to the asymptotic distribution in the noninvertible case. The DD asymptotic approximation can be used for hypothesis testing of $\theta = 1$ against $\theta < 1$ as well as construction of one-sided confidence intervals. In some cases, it might be more useful to assess the sampling uncertainty around the point estimate of the MA parameter in both directions. The inversion of the quantiles of the DD limiting distribution of the ML estimator is not appropriate for obtaining

Journal of Business & Economic Statistics
April 2002, Vol. 20, No. 2

two-sided confidence intervals. The same caveat applies to the construction of two-sided confidence intervals by inverting the acceptance region of the likelihood ratio or score tests for an MA unit root. Typically, the inversion of one-sided (upper tail) tests produces one-sided (upper) confidence intervals.

An alternative to the construction of confidence intervals based on asymptotic approximations is the conventional bootstrap method, which is expected to provide a higher-order accuracy than does the normal approximation for MA roots away from the noninvertibility boundary. The validity of the bootstrap approximation requires that the test statistic be asymptotically pivotal. This condition, however, is not satisfied in the MA(1) model with a large MA root. As will be shown later, the DD limiting representation for θ near the unit circle depends on the local-to-unity parameter c, which is not consistently estimable. Therefore, as the MA parameter approaches unity, the conventional bootstrap fails to provide even first-order correct inference.

In this article, a new bootstrap-based procedure is proposed for constructing two-sided confidence intervals in models with a large MA component. The confidence bounds of these intervals are obtained by inverting the acceptance region of a specific test statistic. Because the likelihood function of models with an MA component is highly nonlinear, the likelihood ratio appears to be the most appropriate statistic reflecting the pronounced asymmetry of the likelihood in the neighborhood of 1. The quantiles of the distribution of the test statistic are approximated by the grid bootstrap method proposed independently by Kabaila (1993), Hansen (1999), and Chuang and Lai (2000) and designed to address the inconsistency of the naive bootstrap for highly persistent autoregressive processes. The grid bootstrap explicitly parameterizes the distribution quantiles as a function of c and does not require a consistent estimate of this parameter. Although it is more computationally expensive than the asymptotic and conventional bootstrap methods, the simulations show that the grid bootstrap is characterized by excellent coverage properties and flexibility to handle a large class of nonstandard inference problems. Furthermore, the grid bootstrap method works globally over the whole parameter space whereas the asymptotic approximations are valid only in the vicinity of 1.

The idea to use a family of resampling distributions rather than a single distribution can be traced to Sprott's (1981) discussion of Efron's seminal work on bootstrap (Efron 1981). This family of distributions is indexed by c, and a single member of the family is used for resampling by conditioning on a particular value of c. Chuang and Lai (2000) called this approach a "hybrid resampling method" because it hybridizes exact methods and bootstrap methods for inference. The unit-root bootstrap tests proposed by Basawa, Mallik, McCormick, Reeves, and Taylor (1991) and Ferretti and Romo (1996) that impose the null hypothesis of a unit root in resampling the data can be regarded as a special case of the grid or hybrid bootstrap method. Given the recent advances in the literature on bootstrapping near-integrated autoregressive models and similarity of the inference problems (in particular, the nonpivotalness of the distribution of the largest root), it seems natural to extend the bootstrap-based methods to models with a large MA component. But, as pointed out earlier, the statistical analysis of nearly noninvertible MA models is further complicated by some nonstandard properties of the ML estimator and the construction of confidence intervals with controlled coverage appears to be more delicate.

This article is organized as follows. Section 2 introduces the statistical methodology for construction of confidence intervals by test inversion. Then, we develop asymptotic and bootstrap approximations of the distribution of the likelihood ratio test in MA(1) models that produce confidence intervals with controlled coverage. The consistency and the accuracy of the grid bootstrap method for interval estimation are discussed in Section 3. Section 4 presents the results from a finite-sample simulation study. The empirical part of the article investigates the possible instability in the MA component of U.S. inflation by constructing confidence intervals for the MA parameter in five different monetary policy regimes. In the second application, the risk premium is modeled as an unobserved component that evolves as a random walk over time, and interval estimates of the variability of the risk premium are provided. The series used in both applications exhibit strong conditional heteroscedasticity, which motivates the use of a heteroscedasticity-robust resampling algorithm. Section 6 summarizes the results and outlines some extensions.

2. CONFIDENCE INTERVAL CONSTRUCTION BY TEST INVERSION

In this section, we use the duality between hypothesis testing and interval estimation and consider the construction of confidence intervals by inverting a test statistic R. Let $y_1, y_2, \dots, y_T$ be a random sample with population distribution $F_0(y) = \Pr\{y_t \le y\}$. Suppose that under the null hypothesis, $R_T = R_T(y_1, y_2, \dots, y_T)$ is a continuous random variable with sampling distribution that may depend on an unknown parameter c,

$$G_T(q \mid c_0) = \Pr\{R_T \le q \mid H_0 : c = c_0\}.$$

For a fixed c and confidence level α, there exists a unique $q_\alpha(c)$ such that

$$\Pr\{R_T \le q_\alpha(c) \mid H_0\} \ge \alpha,$$

where $q_\alpha(c)$ is the αth quantile of the distribution of R_T that can be obtained by inverting G_T. More precisely, $q_\alpha(c) = G_T^{-1}(\alpha \mid c) = \inf\{q(c) : \Pr[R_T \le q(c) \mid H_0] \ge \alpha\}$.

Then, the $100\alpha\%$ confidence set for the MA parameter $\theta_0 = 1 - c/T$ is given by the set of values of θ satisfying $R_T \le q_\alpha(c)$, that is,

$$C_\alpha(y) = \{\theta \in \Theta : R_T \le q_\alpha(c)\}$$

or equivalently

$$C_\alpha(y) = \{\theta \in \Theta : y \in A(\theta)\},$$

where Θ is the parameter space and $A(\theta)$ is the acceptance region of the test R_T (Casella and Berger 1990; Lehmann 1986). Although θ is real-valued, we define a confidence set rather than a confidence interval because it is possible that in some cases the inversion of the test statistic leads to disjoint intervals. The endpoints of the confidence set are the infimum and the supremum over $C_\alpha(y)$, respectively. In particular, the two-sided, equal-tailed confidence interval of nominal coverage α is given by $C_\alpha(y) = [\theta_L, \theta_U]$, where the confidence limits are defined to satisfy $\theta_L = \inf\{\theta \in \Theta : \Pr(R_T \le q_\alpha(c) \mid H_0) \ge \alpha\}$ and $\theta_U = \sup\{\theta \in \Theta : \Pr(R_T \le q_\alpha(c) \mid H_0) \ge \alpha\}$.

The construction of confidence intervals by test inversion depends on the formulation of both the null and the alternative hypothesis of R_T. The first class of tests studied in this article are two-sided tests with a null $H_0 : \theta = \theta_i$, evaluated on a grid of points $\theta_i \in \Theta$, against the alternative $H_1 : \theta \ne \theta_i$, where θ is the MA parameter. If R_T is an asymptotically pivotal statistic, then the conventional asymptotic and bootstrap methods, which treat the αth quantile function as constant over all values of θ, are appropriate for constructing valid confidence intervals. Otherwise, one must explicitly parameterize the quantiles of G_T as a function of the MA parameter θ.

The second class of tests includes one-sided upper-tail tests with a null $H_0 : \theta = 1$ against $H_1 : \theta < 1$, which are concerned with the local behavior of θ in the neighborhood of 1. If the quantiles of the distribution of these test statistics are monotone increasing and continuous in θ, then they can be inverted to produce median unbiased and confidence limit estimates. The median unbiased estimate of θ is

$\theta_{MU} = q_{.5}^{-1}(R_T)$ such that

$$\Pr\{\theta \le q_{.5}^{-1}(R_T)\} = \Pr\{\theta \ge q_{.5}^{-1}(R_T)\} = .5,$$

and the $100\alpha\%$ one-sided confidence interval for θ is given by

$$C_{1,\alpha}(y) = \{\theta \in \Theta : \theta \le q_\alpha^{-1}(R_T)\} = [\underline{\theta}, \theta_U],$$

where $\theta_U = q_\alpha^{-1}(R_T)$ is the upper confidence limit and $\underline{\theta} = -1$ in the MA(1) case. For completeness, the $100\alpha\%$ two-sided confidence interval for θ, obtained by inverting the test statistic of $\theta = 1$, is given by

$$C_{2,\alpha}(y) = \{\theta \in \Theta : q_{(1-\alpha)/2}^{-1}(R_T) \le \theta \le q_{(1+\alpha)/2}^{-1}(R_T)\}.$$

2.1 Asymptotic Confidence Intervals

To derive the limiting theory for constructing confidence intervals for the MA parameter, we impose the following restriction.

Assumption 1. In model (1), $e_t \sim \text{iid}(0, \sigma^2)$ with $E(e_t^4) < \infty$.

The Gaussian quasi log-likelihood of the MA(1) model concentrated with respect to σ^2 (with all constant terms omitted) is given by

$$l_T(\theta) = -\frac{1}{2}\log|\Omega(\theta)| - \frac{T}{2}\log \mathbf{y}'\Omega^{-1}(\theta)\mathbf{y}, \qquad (2)$$

where $\Omega(\theta)$ is the covariance matrix of the data vector $\mathbf{y} = (y_1, y_2, \ldots, y_T)'$. When the MA parameter is in interior of the invertibility region, the ML estimate $\hat\theta$ is asymptotically normally distributed. However, as the MA parameter approaches unity, the exact finite-sample distribution of the ML estimator is more closely approximated by its limiting distribution at 1 rather than the Gaussian distribution. For this reason, we adopt the local-to-unity representation of the MA parameter. Let $\theta_0 = 1 - c/T$ be the true value of $\theta = 1 - \beta/T$, where $c, \beta \ge 0$ are constants. This is the appropriate nesting (normalization) because the ML estimator is T-consistent for the MA parameter in the neighborhood of 1 (Sargan and Bhargava 1983).

This article is primarily concerned with the construction of two-sided, equal-tailed confidence intervals for the MA parameter θ. We argue that intervals with best coverage properties and precision are obtained from inverting the likelihood ratio statistic $\text{LR}(\theta_0) = 2[l_T(\hat\theta) - l_T(\theta_0)]$ that tests the hypothesis $H_0 : \theta = \theta_0$ versus $H_1 : \theta \ne \theta_0$. The LR statistic captures the significant asymmetry of the likelihood in the neighborhood of 1 and is well behaved even when the MA root is on the unit circle. By contrast, the standard Lagrange multiplier and Wald tests are symmetric by construction and encounter some problems when the MA parameter is estimated to lie exactly on the noninvertibility boundary. For small values of θ_0, it is known that the $\text{LR}(\theta_0)$ statistic is approximately χ^2-distributed, whereas for values of θ_0 close to 1, the asymptotic approximation of the distribution of $\text{LR}(\theta_0)$ is given by the following result, which is a restatement of Theorem 2.1 in Davis et al. (1995) and Theorem 2.1 in Davis and Dunsmuir (1996).

Lemma 1. For model (1) with $\theta_0 = 1 - c/T$, $\theta = 1 - \beta/T$ and under Assumption 1,

$$2[l_T(\hat\theta) - l_T(\theta_0)] \xrightarrow{d} Z_c(\tilde\beta_c^{gl}) - Z_c(c), \qquad (3)$$

where

$$Z_c(\beta) = \sum_{k=1}^{\infty} \frac{\beta^2(\pi^2k^2 + c^2)X_k^2}{\pi^2k^2(\pi^2k^2 + \beta^2)} + \sum_{k=1}^{\infty}\log\left(\frac{\pi^2k^2}{\pi^2k^2 + \beta^2}\right), \qquad (4)$$

$$Z_c(c) = \sum_{k=1}^{\infty}\frac{c^2X_k^2}{\pi^2k^2} + \sum_{k=1}^{\infty}\log\left(\frac{\pi^2k^2}{\pi^2k^2 + c^2}\right), \qquad (5)$$

$\tilde\beta_c^{gl}$ is the global maximizer of $Z_c(\beta)$, $X_k \sim \text{nid}(0, 1)$, and $\xrightarrow{d}$ denotes weak convergence on the space of continuous functions on $[0, \infty)$.

Proof. See Appendix A for the proof.

Some nonstandard features of the result in Lemma 1 require further attention. First, the global maximizer $\tilde\beta_c^{gl}$ is a function of the true local-to-unity parameter c and hence the limiting distribution of $\text{LR}(\theta_0)$ is nonpivotal. Moreover, c is not consistently estimable and one cannot substitute a good estimate of c in the limiting representations and obtain the corresponding quantiles. Instead, one can construct a grid of points for the local-to-unity parameter c and evaluate the asymptotic representation at each of these points. Then, the acceptance region of the LR test can be inverted to obtain interval estimators for c and θ.

Second, the asymptotic results in the literature for the nearly noninvertible MA(1) model were derived largely for the local maximizer of l_T, and the limiting behavior of the global maximizer is less known. The local maximizer of Z_c

is given by $\tilde{\beta}_c^{l\max} = \inf\{\beta \geq 0 : \beta H_c(\beta) = 0$ and $\beta H_c'(\beta) + H_c(\beta) < 0\}$, where $H_c(\beta) = \sum_{k=1}^{\infty} 4(\pi^2k^2 + c^2)X_k^2/(\pi^2k^2 + \beta^2)^2 - \sum_{k=1}^{\infty} 4/(\pi^2k^2+\beta^2)$. Also, the computation of the local maximizer is typically much less involved than the search for a global maximum. For these reasons, we work with the local maximizer for constructing asymptotic-based confidence intervals. Davis et al. (1995) found that the limiting distributions of the local and global maximizers of $Z_c(\beta)$ differ, but this difference is not reflected in the asymptotic distribution of the LR statistic.

In the notation of the previous section, the confidence set constructed by inversion of the LR(θ_0) test is $C_\alpha^{\text{LR}}(y) = \{\theta \in \Theta : \text{LR}(\theta_0) \leq q_\alpha^{\text{LR}}(c)\}$. It is worth pointing out that although the LR test statistic seems to be the most appropriate and informative statistic in this context, it is not clear if it still possesses its optimality properties (Ploberger 1999). This is because the second derivative of the likelihood function converges in distribution to a random quantity rather than converging in probability to a constant matrix as in the usual case.

In addition, we consider the t-statistic of $H_0: \theta = \theta_0$ versus $H_1: \theta \neq \theta_0$, which in the MA(1) case is the square root of the Wald statistic. Its corresponding $100\alpha\%$ confidence interval is given by $C_\alpha^t(y) = \{\theta \in \Theta : q_{(1-\alpha)/2}^t(c) \leq t(\theta_0) \leq q_{(1+\alpha)/2}^t(c)\}$, where q^t denotes the corresponding quantile from the distribution of the t-statistic. The problem with the t-statistic is that if the estimated MA root is estimated to lie exactly on the unit circle, which happens quite often when the parameter is close to the noninvertibility boundary because of the pile-up effect, the score of the likelihood function (2), evaluated at 1, is identically 0. Then, the use of the outer product method results in a very large estimate of the sampling variance of the parameter and the t-statistic collapses to 0. Therefore, tests that use an estimate of the variance such as the t-statistic and the Wald statistic are not appropriate for statistical inference in the neighborhood of unity. Moreover, the two-sided Lagrange multiplier test also does not appear to be a good candidate for inversion in this region of the parameter space. The Lagrange multiplier test behaves well for values of the MA root away from unit circle but not for values very close to 1 because at $\theta = 1$ the likelihood function has a zero slope.

2.2 Bootstrap Confidence Intervals

Before the properties of the bootstrap-based confidence intervals can be discussed, some notation is needed. Let $(y_1, y_2, \ldots, y_T)$ again be a random sample with population distribution $F_0(y) = \Pr\{y_t \leq y\}$, which belongs to a family of distribution functions Ξ and is generally unknown. Also, let F_e and $F_{e,T}$ denote the population and empirical distribution function of the unobserved errors e_t in (1). The bootstrap methods operate conditionally on the realized sample by replacing the unknown distribution function with an estimator. An estimator of F_e is given by the empirical distribution function of the ML residuals $\hat{e}_1, \hat{e}_2, \ldots, \hat{e}_T$, $\widehat{F}_{e,T} = \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}\{\hat{e}_t \leq e\}$. The corresponding empirical distribution function of the data is denoted as $\widehat{F}_T$. The bootstrap principle replaces the pair $(F_0, \widehat{F}_T)$ by the pair $(\widehat{F}_T, F_T^*)$, where F_T^* is the distribution function of a (bootstrap) sample $y_1^*, y_2^*, \ldots, y_T^*$ drawn from $\widehat{F}_T$. To see this, we can rewrite the sampling distribution of the test statistic $R_T(\hat{\theta}, \theta_0)$ under the null as $G_T(q \mid F_0) = \Pr\{R_T[\theta(\widehat{F}_T), \theta(F_0)] \leq q \mid F_0\}$. Then, the bootstrap analog of this distribution is $G_T^*(q \mid \widehat{F}_T) = \Pr\{R_T[\theta(F_T^*), \theta(\widehat{F}_T)] \leq q \mid \widehat{F}_T\}$ and the bootstrap confidence set for θ is $C_\alpha^*(y) = \{\theta \in \Theta : R_T \leq q_\alpha^*(\theta)\}$, where $q_\alpha^*(\theta)$ is the αth bootstrap quantile function.

In MA(1) models, the conventional bootstrap method draws a sequence of residuals $\{e_t^*\}_{t=1}^T$ with replacement from $\widehat{F}_{e,T}$ and generates bootstrap data using the ML estimate of θ. For estimation, Bose (1990) showed that the conventional bootstrap approximation of the distribution of the MA parameter estimate increases the accuracy from $O_p(T^{-1/2})$ to $o_p(T^{-1/2})$ almost surely for θ strictly less than 1. The bootstrap-based inference in the MA(1) model, however, has not been fully investigated in the literature. It is known that for θ away from the noninvertibility boundary, the parameter estimate is approximately normally distributed and an appropriate studentization would provide asymptotic pivotalness and render a first-order asymptotic coverage. Furthermore, if the studentized statistic admits an Edgeworth expansion, a higher-order asymptotic coverage can be obtained.

When the MA parameter is close to 1, however, the asymptotic distribution of the ML estimator derived by Davis and Dunsmuir (1996) depends on the local-to-unity parameter, which is not consistently estimable. This implies that the conventional bootstrap is not even first-order correct. For this reason, we use the grid bootstrap method of Hansen (1999) designed originally to deal with the nonpivotalness of the distribution of the AR parameter in near-integrated models. In the present context, this method explicitly allows the quantile functions of the distribution of the test statistic to vary with θ and does not require a consistent estimate of the local-to-unity parameter. The grid bootstrap partitions the relevant part of the parameter space of θ into n equally spaced points $\vartheta = \{\theta_i\}_{i=1}^n$ and generates data under the null $H_0: \theta = \theta_i$ or, equivalently, $H_0: \beta = c_i$.

It is widely documented that many economic time series exhibit conditional heteroscedasticity. The standard bootstrap, which resamples with replacement from $\widehat{F}_{e,T}$, does not take the heteroscedasticity into account and destroys potentially important information in the structure of the residuals. For this reason, we combine the grid bootstrap with a weighted resampling scheme without replacement proposed by Wu (1986) and known also as the wild bootstrap (Härdle and Mammen 1993). This method constructs bootstrap data by setting $e_t^* = \nu_t \hat{e}_t$, where ν_t is drawn from a distribution F_ν that satisfy $E(\nu_t) = 0$, $E(\nu_t^2) = E(\nu_t^3) = 1$. In particular, we set $\nu_t = (\delta_1 + \varphi_{t,1}/\sqrt{2})(\delta_2 + \varphi_{t,2}/\sqrt{2}) - \delta_1\delta_2$, where $\varphi_{t,i} \sim N(0,1)$, $\delta_1 = (3/4 + \sqrt{17}/12)^{1/2}$ and $\delta_2 = (3/4 - \sqrt{17}/12)^{1/2}$. Although the weighted resampling without replacement is robust to heteroscedasticity of general form, it has been of limited use in time series models because it holds all regressors fixed across replications; this is very restrictive in the presence of lagged dependent variables. Fortunately, the structure of the MA(1) model allows one to employ the weighted resampling scheme in mimicking the dynamics of the series. As an alternative, we can explicitly parameterize the form of the conditional heteroscedasticity, for example, as a generalized autoregressive conditionally heteroscedastic (GARCH) process, and then recursively generate bootstrap series from the conditional mean and the conditional variance functions. The validity of this approach, however, may depend on the correct specification of the scedastic function.

The grid bootstrap procedure is valid under the condition of functional independence between the parameter of interest θ_0 and the nuisance parameter vector η_0, which is satisfied in the MA(1) model (Kabaila 1993). Then, by assuming that $\eta_0 = \hat{\eta}$, we draw a sequence of residuals $\{e_t^*\}_{t=1}^T$ as described in the previous paragraph, generate bootstrap samples $y_t^* = e_t^* - \theta_i e_{t-1}^*$ at the grid point θ_i, and calculate the statistic R_T^*. These steps are repeated B times, and the 100αth element of the ordered statistic R_T^* gives the critical value of the test at θ_i. Applying this algorithm at each grid point and connecting the respective critical values produces an approximation of the αth quantile function of the finite-sample distribution of R_T. Finally, we can use the test-inversion method for constructing interval estimators of the form $C_\alpha^*(y) = \{\theta \in \Theta : R_T \leq q_\alpha^*(c)\}$, which will be at least first-order asymptotically correct.

The construction of confidence intervals for the MA parameter by inverting the LR test for the sequence of hypotheses $\theta = \theta_i, i = 1, 2, \ldots, 41$, is illustrated in Figure 1. The data are generated from a Gaussian MA(1) model with a true parameter value $\theta = .95$ and sample size of 100. The ML estimate of θ for this particular realization is .945. Along with the LR statistic, Figure 1 plots the .9th quantile functions from the grid and conventional bootstrap as well as the DD and χ^2 asymptotics. The dependence of the grid bootstrap and DD quantiles on the values of θ and the asymmetry of the LR statistic are apparent from the figure. Therefore, the lack of pivotalness in the distribution of the LR test makes the assumption of constant quantiles used in the conventional bootstrap and χ^2 approximations inappropriate when θ is near the noninvertibility boundary. Projecting the intersections of the LR statistic with the corresponding quantile onto the horizontal axis produces the lower and the upper endpoints of the confidence interval for θ. Another conceptually similar approach involves construction and inversion of the grid bootstrap p-value function $p^*(\theta_0) = \Pr\{R_T^* \geq R_T \mid \theta = \theta_0\}$ (Davidson and MacKinnon 1999).

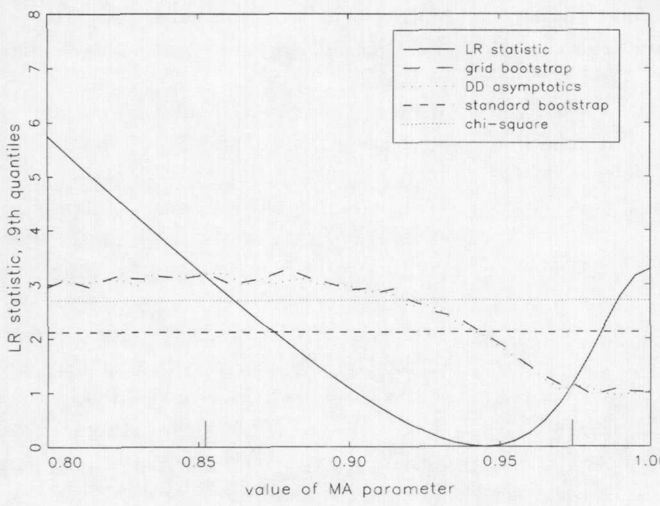

Figure 1. Confidence Intervals by LR Test Inversion. The data used in the figure were generated from a standard normal MA(1) model with $\theta_0 = .95$ and sample size of 100. The ML estimate is $\hat{\theta} = .945$. The solid line is the LR test statistic $2[l_T(\hat{\theta}) - l_T(\theta_i)]$, evaluated on a grid of 40 equally spaced values of θ_i from .80 to 1. The figure also plots the χ^2 asymptotic, Davis-Dunsmuir asymptotic, conventional bootstrap, and grid bootstrap .9th quantile functions from the respective distributions. The grid bootstrap 90% two-sided confidence interval is obtained by projecting the intersections of the LR statistic with the grid bootstrap critical value function onto the horizontal axis, indicated by the two spikes at .86 and .97.

3. CONSISTENCY AND ACCURACY OF THE GRID BOOTSTRAP

The conditions for the consistency of the bootstrap distribution $G_T^*(q \mid \widehat{F}_T)$ as an estimator of the true distribution of the statistic R_T were given by Bickel and Freedman (1981), Beran and Ducharme (1991), and Horowitz (in press). Let d_Ξ be a metric on the space Ξ of distribution functions and $F_0 \in \Xi$.

The first condition states that for any $\delta > 0$, $\lim_{T\to\infty} \Pr\{d_\Xi(\widehat{F}_T, F_0) > \delta\} = 0$; that is, the sampling distribution of the data $\widehat{F}_T$ is a consistent estimator of the true distribution F_0. Second, the asymptotic distribution of the statistic R_T, $G_\infty(q \mid F)$, is continuous in q for any $F \in \Xi$. This condition requires a continuity of the mapping from the sampling distribution of the data to the limiting distribution of the test statistic. Finally, for any sequence $\{H_T\} \in \Xi$ belonging to a neighborhood of the true distribution F_0 such that $\lim_{T\to\infty} d_\Xi(H_T, F_0) = 0$, the distribution $G_T^*(q \mid H_T)$ converges weakly to its limiting distribution $G_\infty(q \mid F_0)$.

Some of these conditions for consistency of the bootstrap estimator are clearly violated in models with a large MA component. For instance, when the MA parameter is exactly on the unit circle, the asymptotic representation of the ML estimator changes discontinuously from the Gaussian to the DD distribution. Moreover, some inference problems arise from imposing the invertibility restriction, in which case the MA parameter lies exactly on the boundary of the parameter space (Andrews 2000). The discontinuity problem can be circumvented by adopting the local-to-unity framework that allows a smooth transition from the normal distribution to the limiting representation at unity. Unfortunately, this introduces another difficulty arising from the presence of a nuisance parameter that is not consistently estimable. As a result, the conventional bootstrap would not provide a consistent estimate of the true distribution of the test statistic R_T.

The grid bootstrap approximates the distribution $G_T(q \mid F_0)$ for a sequence of null hypotheses $H_0 : \theta = \theta_i$ (or, in terms of the local-to-unity parameters, $\beta = c_i$) for $\theta_i \in \vartheta$ by generating bootstrap data imposing the null. This method does not require a consistent estimate of the parameter c and is asymptotically valid. Furthermore, it appears that imposing the null when the parameter is exactly on the boundary of the parameter space alleviates the problem pointed out by Andrews (2000). Also, it is worth noting that when the MA parameter is exactly on the invertibility boundary, the ML estimate is median unbiased because the pile-up probability in this case is .6575 (Sargan and Bhargava 1983).

As argued here, we prefer to work with the LR test of $\theta = \theta_0$ and invert its bootstrap acceptance region to obtain confidence intervals. The main result that establishes the consistency of the grid bootstrap approximation when the limiting distribution of the LR statistic depends on an unknown parameter is given in the following theorem.

Theorem 1. For any fixed $c, \delta > 0$, and $F_0 \in \Xi$, the grid bootstrap provides a consistent estimator of the limiting distribution of the $LR(\theta_0)$ statistic:

$$\lim_{T\to\infty} \Pr\left\{\sup_q \left|G_T^*(q \mid \widehat{F}_T, c) - G_\infty(q \mid F_0, c)\right| > \delta\right\} = 0.$$

Proof. see Appendix A for the proof.

An important feature of the result in Theorem 1 is that the grid bootstrap is a consistent estimator of the distribution of $\mathrm{LR}(\theta_0)$ over the whole parameter space Θ, not only in the vicinity of 1. The result is valid for both restricted and estimated residuals as long as the latter are consistently estimated. Also, it immediately follows from Theorem 1 that the grid bootstrap approximation is first-order accurate, whereas the conventional bootstrap fails to achieve that.

Chandra and Ghosh (1979) showed that under normality of the ML estimator and some smoothness and moment conditions, the χ^2 distribution approximates the finite sample distribution of the LR statistic up to order $O(T^{-1})$,

$$\Pr(\mathrm{LR} \le x) = G\chi^2(x) + O(T^{-1}),$$

where $G\chi^2(x)$ is the χ^2 distribution function. Although Chandra and Ghosh (1979) established this result for iid data, they argued that the iid structure can be relaxed and a similar result can be derived for dependent data. Then, using arguments similar to that of Hall (1992, chap. 3), it can be shown that the bootstrap approximation provides an improvement over the χ^2 asymptotics,

$$\Pr(\mathrm{LR} \le x) - \Pr(\mathrm{LR}^* \le x) = o_p(T^{-1}),$$

and the coverage error of the $100\alpha\%$ two-sided intervals, obtained by inverting the bootstrap $\mathrm{LR}(\theta_0)$ statistic, is of order $o_p(T^{-1})$ (Carpenter 1998).

When $\theta_0 = 1 - c/T$, the asymptotic normality of the ML estimator of θ is lost. To establish the possible higher accuracy of the bootstrap confidence intervals, we need to show first that the distribution function of the LR statistic, $G_T(q \mid c, \eta_0)$, admits an Edgeworth expansion. The technical conditions for developing an Edgeworth expansion of the distribution of $\mathrm{LR}(\theta_0)$ when θ is near or on the unit circle have not been derived yet. For this reason, it can be claimed only that in this region of the parameter space the bootstrap approximation achieves the first-order asymptotic accuracy.

4. MONTE CARLO STUDY

4.1 Experimental Design

To evaluate the finite sample performance of the discussed methods for constructing confidence intervals of the MA parameter, we conducted a Monte Carlo simulation study. The design of the experiment is as follows. The data are generated from an MA(1) model,

$$y_t = e_t - \theta_0 e_{t-1} \text{ with } T = 100$$

and

$$\theta_0 = .6, .8, .9, .95, .99, 1,$$

where $e_t = \sigma_t \xi_t$, $\xi_t \sim \mathrm{nid}(0, 1)$, $\sigma_t^2 = (1 - \gamma_1 - \gamma_2) + (\gamma_1 \xi_{t-1}^2 + \gamma_2)\sigma_{t-1}^2$ with two specification for the scedastic function: DGP1, $\gamma_1 = \gamma_2 = 0$ (iid errors), and DGP2, $\gamma_1 = .3$, $\gamma_2 = .6$ (GARCH(1, 1) errors). The simulation experiment examines the coverage probabilities and the length of the confidence intervals at nominal level $\alpha = .9$. The number of Monte Carlo replications is 2,000. Note that another set of simulations confirmed that the results reported in the following are robust to different specifications of the error distribution, such as t distribution with 5 degrees of freedom, $\xi_t \sim t_5(0, 1)$, and a mixture of normals, $\xi_t \sim [.9N(-1, 1) + .1N(9, 1)]$.

Confidence intervals are constructed by using six asymptotic approximations: three Gaussian approximations, the limiting distribution of $\mathrm{LR}(\theta_0)$ and $\mathrm{LR}(1)$ for θ on or near the unit circle, and the asymptotic representation of the score-type test of $\theta = 1$ derived by Tanaka (1990).

As pointed out in the introduction, when the MA parameter is inside the invertibility region, the ML estimator of θ_0 is asymptotically normally distributed with mean θ_0 and variance $(1 - \theta_0^2)/T$. The first two Gaussian approximations are based on the $N(0, 1)$ quantities $Z_1 = (\hat\theta - \theta_0)/\sqrt{(1 - \theta_0^2)/T}$ and $Z_2 = (\hat\theta - \theta_0)/\mathrm{s.e.}(\hat\theta)$, where $\mathrm{s.e.}(\hat\theta)$ denotes the sampling standard error of $\hat\theta$. The lower and upper confidence limits, θ_L and θ_U, for Z_1 and Z_2 are obtained analytically by solving equation $\Pr\{Z_i \le z_{(1+\alpha)/2}\}$, $i = 1, 2$, for θ_0, where $z_{(1+\alpha)/2}$ is the $(1 + \alpha)/2$th quantile of the standard normal distribution. The third approximation for constructing asymptotic confidence intervals inverts the χ^2 acceptance region of the LR statistic, $\mathrm{LR}(\theta_0) = 2[l_T(\hat\theta) - l_T(\theta_0)]$. The endpoints of the confidence intervals are calculated numerically and are given by $\theta_L = \max[-1, \inf\{\theta : \mathrm{LR}(\theta_0) \le \chi_\alpha^2\}]$ and $\theta_U = \min[\sup\{\theta : \mathrm{LR}(\theta_0) \le \chi_\alpha^2\}, 1]$.

The fourth approximation is based on the limiting distribution of $\mathrm{LR}(\theta_0)$ derived in Lemma 1. Because this asymptotic representation cannot be computed exactly, the corresponding quantiles are obtained by simulation. First, solve for the local maximizer $\tilde\beta_c^{l\max}$ from the expression for $H_c(\beta)$, following the procedure proposed by Davis and Dunsmuir (1996). The infinite series $H_c(\beta)$ is truncated at $k = 1,000$ and $H_c(0)$ is computed. If $H_c(0) \le 0$, then $\tilde\beta_c^{l\max}$ is set to 0; if $H_c(0) > 0$, then $\tilde\beta_c^{l\max}$ is the smallest nonnegative root of $H_c(\beta)$ found by grid search and linear interpolation between the grid points. The number of replications is 10,000. The true local-to-unity parameter c takes on 51 integer values from 0 to 50, which for sample size $T = 100$ corresponds to $\theta = 1, 0.99, 0.98, \ldots, 0.5$. Then the local maximizer $\tilde\beta_c^{l\max}$ is used to evaluate the asymptotic representations of the two-sided $\mathrm{LR}(\theta_0)$ test of $\theta = \theta_0$. Finally, we invert the αth quantile of the limiting distributions of $\mathrm{LR}(\theta_0)$ to determine the upper and the lower interval endpoints. If there is no corresponding value in the table for the lower quantile, then the lower limit is set to the leftmost point on the grid for θ.

In the Monte Carlo comparison, we also investigate the properties of two popular one-sided tests for an MA unit root with $H_0: \theta = 1$ and $H_1: \theta < 1$: the likelihood ratio test LR(1) of Davis et al. (1995) and the score-type test developed by Tanaka (1990). It immediately follows from Lemma 1 that $\mathrm{LR}(1) = 2[l_T(\theta) - l_T(1)]$ converges in distribution to the global maximizer of the expression given in (4), which is again a function of c. Analogously to the $\mathrm{LR}(\theta_0)$ test, we evaluate the limiting distribution using the local maximizer $\tilde{\beta}_c^{l\max}$, and the $100\alpha\%$ confidence interval contains all the values of the parameter that satisfy the condition $q^{-1}_{(1-\alpha)/2}[\mathrm{LR}(1)] \leq \theta \leq q^{-1}_{(1+\alpha)/2}[\mathrm{LR}(1)]$.

The locally best invariant unbiased score-type test for $\theta = 1$ against $\theta < 1$ (Tanaka 1990) is based on the second derivative of the log-likelihood evaluated at $\theta = 1$ and is given by $S_T = [\mathbf{y}'\Omega^{-2}(1)\mathbf{y}]/[T\mathbf{y}'\Omega^{-1}(1)\mathbf{y}]$. As $T \to \infty$ and $\theta = 1 - c/T$, S_T converges in distribution to $\sum_{k=1}^{\infty}(1/\pi^2k^2 + c^2/\pi^4k^4)X_k^2$, where $X_k \sim \mathrm{nid}(0, 1)$. This asymptotic representation of the S_T test is simulated (10,000 Monte Carlo replications) on a grid of integer values for the local-to-unity parameter c from 0 to 50, and the corresponding quantiles are inverted to locate the lower and the upper confidence limits.

In the conventional bootstrap, we use the other (or Efron's) percentile, Hall's percentile, the percentile-t and the percentile-LR methods for constructing confidence intervals. The two-sided equal-tailed confidence intervals are $[\theta^*_{(1-\alpha)/2}, \theta^*_{(1+\alpha)/2}]$, $[\hat{\theta} - (\theta^*_{(1+\alpha)/2} - \hat{\theta}), \hat{\theta} - (\theta^*_{(1-\alpha)/2} - \hat{\theta})]$, $[\hat{\theta} + t^*_{(1-\alpha)/2}\mathrm{s.e.}(\hat{\theta}), \hat{\theta} + t^*_{(1+\alpha)/2}\mathrm{s.e.}(\hat{\theta})]$, and $[\inf\{\theta : \mathrm{LR} \leq q^*_\alpha\}, \sup\{\theta : \mathrm{LR} \leq q^*_\alpha\}]$, respectively, where * indicates the bootstrap analogs of the corresponding estimators and statistics. The number of bootstrap repetitions is 499.

Finally, for the grid bootstrap we construct a grid of 20 evenly spaced points on $[\hat{\theta} - 5\ \mathrm{s.e.}(\hat{\theta}), \min\{\hat{\theta} - 5\ \mathrm{s.e.}(\hat{\theta}), 1\}]$. The number of bootstrap replications at each grid point is 99. Because this number is relatively small, we smooth the bootstrap quantile functions nonparametrically by kernel regression. The kernel regression uses the Nadaraya-Watson estimator with Epanechnikov kernel and bandwidth of $0.4\sigma_\vartheta$, where $\vartheta = \{\theta_1, \ldots, \theta_{20}\}$ is the constructed sequence of grid points for the MA parameter. The grid bootstrap confidence intervals are obtained by five methods: Hall's percentile, percentile-t, inverting the S_T test of $\theta = 1$, LR test of $\theta = 1$, and LR test of $\theta = \theta_0$.

4.2 Simulation Results and Discussion

The results from the Monte Carlo simulation with iid errors are presented in Tables 1, 2, and 3. Table 1 reports the coverage rate and the precision, measured by the median length, of the two-sided, equal-tailed confidence intervals constructed using the different methods. For small values of the MA parameter the normal distribution is expected to provide a relatively accurate approximation of the distribution of the ML estimator. The Gaussian approximations 1 and 2 have coverage rates close to the nominal level but still undercover by 3% to 4% mainly as a result of underestimating the variance of the ML estimator (up to 40% at $\theta_0 = .8$). The conventional bootstrap methods seem to provide a correction for the underestimation of the parameter variance in the asymptotic Gaussian methods and produce wider intervals with better coverage. The inversion of the $\chi^2(1)$ quantile of the LR statistic works better than the other Gaussian methods, and bootstrapping the test statistic (bootstrap percentile-LR) ensures a coverage rate that is closer to the nominal level. The grid bootstrap methods have similar coverage rates as the conventional bootstrap because in this region of the parameter space the quantiles of the distribution of the different statistics are very close to constant functions of θ.

In the intermediate case $\theta_0 = .8$, the coverage properties of the normal approximations, except the one based on the inversion of the LR statistic, start to deteriorate. The coverage rate of the conventional bootstrap methods also decreases. At this value of the MA parameter the advantages of the grid bootstrap methods become more evident. As θ_0 approaches 1, the Gaussian approximation 1 performs poorly and provides rather misleading information. The second Gaussian approximation, which uses sampling information, has better coverage but the length of the interval becomes prohibitively large for $\theta_0 = .99$ and $\theta_0 = 1$ and the confidence intervals are completely uninformative. Therefore, none of these normal approximations provides a useful guide for obtaining reliable interval estimators for the MA parameter.

The asymptotic approximation based on inversion of the LR statistic using the $\chi^2(1)$ critical value has good coverage and precision although it overcovers in the neighborhood of 1. As predicted by the theoretical discussion in the previous section, the conventional bootstrap fails to provide even first-order accuracy and in some cases is characterized with extremely bad coverage properties. Because the S_T and LR tests of the hypothesis $\theta = 1$ against $\theta < 1$ are one-sided, upper-tail tests, they are not suitable for constructing two-sided confidence intervals, and their coverage probabilities tend to be closer to .95 rather than the nominal level of .9. In fact, the lower point of the interval obtained from these methods is almost always to the left of the true value and causes the observed overcoverage. The grid percentile and grid percentile-t methods are performing well but not as well as the grid bootstrap for the LR test of $\theta = \theta_0$, and they also overcover by several percentage points. It is worth pointing out that the median length from all methods, except those using an estimate of the sampling variance, is decreasing as θ_0 approaches the unit circle because the MA parameter is more precisely estimated (T rather than $\sqrt{T}$ consistent) when it is close to the noninvertibility boundary.

The only two methods that perform consistently well for all values of θ_0 for both coverage and precision are the DD asymptotic and the grid bootstrap approximations of the distribution of the LR test of $\theta = \theta_0$. In all cases, the coverage rate is very close to the nominal level and the lengths of the confidence intervals are smaller than the lengths from the confidence intervals with compatible coverage.

The case when the true MA parameter is exactly equal to 1 deserves more attention. It was shown in Section 3 that the grid bootstrap eliminates the discontinuity of the mapping from the sampling distribution of the data to the asymptotic distribution of the test statistic. However, the construction of two-sided confidence intervals when the parameter is exactly on the upper boundary of the parameter space suffers from the generic problem that the confidence interval never misses

Table 1. 90% Two-Sided, Equal-Tailed Confidence Intervals (iid errors)

Methods	Coverage rate	Median length	Coverage rate	Median length	Coverage rate	Median length
	$\theta_0 = .60$		$\theta_0 = .80$		$\theta_0 = .90$	
ASYMPTOTIC						
Gaussian 1	.8665	.2603	.8385	.1956	.7660	.1421
Gaussian 2	.8645	.2672	.8585	.2042	.9255	.1611
χ^2–*LR* test of $\theta = \theta_0$	.8870	.2635	.8795	.2046	.8740	.1584
S_T test of $\theta = 1$	.9225	.3993	.9530	.4024	.9530	.3574
DD–*LR* test of $\theta = 1$	.9485	.3472	.9770	.3382	.9690	.2554
DD–*LR* test of $\theta = \theta_0$	.9000	.2187	.9090	.2118	.8880	.1588
BOOTSTRAP						
Other percentile	.8570	.2730	.8335	.2260	.7845	.1612
Percentile	.8765	.2730	.8710	.2260	.7320	.1611
Percentile-*t*	.9070	.2884	.8305	.2343	.6350	.1236
Percentile-*LR*	.8830	.2615	.8735	.2085	.7700	.1563
GRID BOOTSTRAP						
Percentile	.9045	.2845	.9075	.2315	.8665	.2034
Percentile-*t*	.9050	.2858	.9050	.2320	.8950	.1966
S_T test of $\theta = 1$	.9060	.6403	.9090	.3874	.9150	.2652
LR test of $\theta = 1$	.9025	.5122	.9050	.3310	.9020	.2357
LR test of $\theta = \theta_0$	.8900	.2676	.8930	.2081	.8850	.1566
	$\theta_0 = .95$		$\theta_0 = .99$		$\theta_0 = 1.00$	
ASYMPTOTIC						
Gaussian 1	.9265	.0988	.7970	.0527	.6220	.0527
Gaussian 2	.9820	.1481	.9670	2.0000	.9560	2.0000
χ^2–*LR* test of $\theta = \theta_0$	.9675	.1214	.9750	.0930	.9745	.0933
S_T test of $\theta = 1$	.9400	.2181	.9455	.1425	.9490	.1425
DD–*LR* test of $\theta = 1$	.9595	.2279	.9460	.5100	.9590	.5100
DD–*LR* test of $\theta = \theta_0$	.9025	.1241	.8920	.0970	.8930	.0966
BOOTSTRAP						
Other percentile	.9945	.0943	.9920	.0718	.9905	.0716
Percentile	.5500	.0694	.3295	.0000	.5585	.0000
Percentile-*t*	.4455	.0464	.2790	.0000	.4495	.0000
Percentile-*LR*	.7190	.0886	.9605	.0494	.9565	.0493
GRID BOOTSTRAP						
Percentile	.9405	.1718	.9585	.1105	.9655	.1105
Percentile-*t*	.9190	.1870	.9545	.3000	.9440	.3000
S_T test of $\theta = 1$	.9290	.2145	.9440	.1504	.9450	.1500
LR test of $\theta = 1$	.9375	.1843	.9480	.1263	.9430	.1263
LR test of $\theta = \theta_0$	.9010	.1228	.9050	.0957	.9005	.0953

NOTE: Gaussian 1 and 2 = Gaussian approximations based on $(\hat{\theta} - \theta_0)/\sqrt{(1 - \theta_0^2)/T}$ and $(\hat{\theta} - \theta_0)/\text{s.e.}(\hat{\theta})$.

Table 2. 90% One-Sided, Upper-Tailed Confidence Intervals (iid errors)

	Coverage rate					
Methods	$\theta_0 = .60$	$\theta_0 = .80$	$\theta_0 = .90$	$\theta_0 = .95$	$\theta_0 = .99$	$\theta_0 = 1.00$
ASYMPTOTIC						
Gaussian 1	.9080	.9035	.8910	.8840	.7575	.6220
Gaussian 2	.9270	.9350	.9415	.9550	.9300	.9120
χ^2–*LR* test of $\theta = \theta_0$	.9145	.9035	.9195	.9345	.9430	.9455
S_T test of $\theta = 1$	1.0000	1.0000	.9520	.9140	.8950	.9005
DD–*LR* test of $\theta = 1$	1.0000	.9775	.9215	.9080	.8950	.9160
DD–*LR* test of $\theta = \theta_0$	.9285	.9185	.9355	.9275	.7920	.7910
BOOTSTRAP						
Other percentile	.9305	.9510	.9680	.9830	.9800	.9800
Percentile	.9150	.9185	.9240	.9425	.9060	.8720
Percentile-*t*	.9450	.9720	.9835	.9810	.6455	.5065
Percentile-*LR*	.9040	.9040	.9270	.9435	.9405	.9290
GRID BOOTSTRAP						
Percentile	.9090	.9040	.8975	.9155	.9245	.9220
Percentile-*t*	.9090	.9060	.9015	.9150	.9050	.8865
S_T test of $\theta = 1$	.9235	.8990	.9050	.9015	.9035	.8935
LR test of $\theta = 1$	.8980	.8970	.8995	.9065	.8990	.8885
LR test of $\theta = \theta_0$	.9095	.9095	.9260	.9185	.8350	.8140

NOTE: Gaussian 1 and 2 = Gaussian approximations based on $(\hat{\theta} - \theta_0)/\sqrt{(1 - \theta_0^2)/T}$ and $(\hat{\theta} - \theta_0)/\text{s.e.}(\hat{\theta})$.

Table 3. *Median Unbiased Estimates* (iid *Errors*)

Methods	*Median bias*	*RMSE*	*MAE*	*Median bias*	*RMSE*	*MAE*
		$\theta_0 = .60$			$\theta_0 = .80$	
ML estimate	.0055	.0881	.0690	.0088	.0720	.0553
ASYMPTOTIC						
S_T test of $\theta = 1$	.1505	.1781	.1519	.0228	.0937	.0802
DD–LR test of $\theta = 1$	.0757	.1158	.0933	.0045	.0810	.0656
GRID BOOTSTRAP						
Percentile	−.0012	.0870	.0682	.0000	.0708	.0543
Percentile-*t*	−.0010	.0873	.0684	.0023	.0726	.0562
S_T test of $\theta = 1$	.2642	.3046	.2776	.0756	.1543	.1291
LR test of $\theta = 1$	.0149	.1793	.1396	.0046	.1072	.0838
		$\theta_0 = .90$			$\theta_0 = .95$	
ML estimate	.0112	.0573	.0455	.0153	.0414	.0345
ASYMPTOTIC						
S_T test of $\theta = 1$	.0026	.0711	.0579	−.0004	.0485	.0390
DD–LR test of $\theta = 1$	−.0013	.0621	.0496	−.0008	.0439	.0369
GRID BOOTSTRAP						
Percentile	.0009	.0575	.0448	.0017	.0439	.0364
Percentile-*t*	.0075	.0628	.0485	.0072	.0657	.0443
S_T test of $\theta = 1$	.0067	.0956	.0721	.0016	.0537	.0417
LR test of $\theta = 1$	.0014	.0653	.0514	.0031	.0429	.0362
		$\theta_0 = .99$			$\theta_0 = 1.00$	
ML estimate	.0100	.0302	.0191	.0000	.0332	.0154
ASYMPTOTIC						
S_T test of $\theta = 1$	.0011	.0323	.0234	−.0111	.0357	.0225
DD–LR test of $\theta = 1$	.0100	.0303	.0227	.0000	.0334	.0196
GRID BOOTSTRAP						
Percentile	.0100	.0356	.0193	.0000	.0389	.0197
Percentile-*t*	.0100	.0882	.0446	.0000	.0936	.0439
S_T test of $\theta = 1$	.0078	.0325	.0144	−.0054	.0357	.0217
LR test of $\theta = 1$	.0040	.0288	.0209	−.0014	.0323	.0197

NOTE: RMSE = root mean square error; MAE = mean absolute error.

to the left. The simulation results in Table 1 indicate that the coverage rate of the confidence intervals obtained from the grid bootstrap LR method is not significantly different from the nominal level, and hence we somewhat circumvent this problem in the MA(1) model. This is because, as mentioned earlier, although the lower confidence limit is always to the left of the true value, the LR statistic reflects the asymmetry of the likelihood function in this region and allows the upper endpoint to undercover.

Tables 2 and 3 summarize some additional information about the properties of the different interval estimators concerning the coverage probabilities of the 90% one-sided confidence intervals, the bias and variability of the median unbiased estimates. As expected, the inversion of the DD–LR test of $\theta = 1$ and the S_T test of $\theta = 1$ gives one-sided, upper-tail confidence intervals with coverage rates very close to the nominal level of .9 when θ_0 is near the unit circle. We also experimented with inverting the DD limiting distribution of the ML estimator of θ, which is given in parts (i) and (ii) of the preliminary lemma in Appendix A. It is interesting to note that the one-sided intervals obtained from this method have a very good coverage even for values of the MA parameter as low as .6 and .8. Table 2 shows that this property is not shared by the DD–LR test of $\theta = 1$.

The confidence intervals based on inversion of the LR test of $\theta = \theta_0$, $C_{1,\alpha}(y) = \{\theta \in \Theta : \text{LR}(c) \leq q_{2\alpha-1}(\theta)\}$, have good coverage but undercover for θ_0 close to one. This would imply that the two-sided LR intervals, obtained as an intersection of two one-sided intervals, undercover the upper confidence limit and overcover the lower confidence limit. Again, this is because the LR test adjusts the right confidence endpoint to ensure better coverage for the two-sided interval.

Finally, the best median unbiased estimates of θ presented in Table 3 are obtained by inverting the grid bootstrap LR test of $\theta = 1$. The DD limiting distribution of the LR(1) test produces small median bias for $\theta_0 = .9, .95,$ and 1, but it does not work well for $\theta_0 = .6$ and .99. Similarly, the inversion of the S_T test gives reasonable median unbiased estimates for $\theta_0 \geq .95$ but its performance deteriorates significantly for $\theta_0 < .9$.

In summary, based on the results from the finite-sample experiment with iid errors, we recommend the grid bootstrap LR test of $\theta = \theta_0$ for constructing two-sided intervals and the grid bootstrap LR test of $\theta = 1$ for obtaining one-sided intervals and median unbiased estimates.

Table 4 examines the robustness of the bootstrap algorithm to presence of conditional heteroscedasticity. Table 4 reports the coverage rate and median length of two-sided confidence intervals for the MA(1) model with $\theta_0 = .95, .99, 1.00,$ and GARCH(1,1) errors. Because the limiting distribution of the LR test is derived under the assumption of iid errors, it is not surprising to see that the coverage of the DD asymptotic approximation deteriorates compared to the homoscedastic case. The invalidity of the χ^2 asymptotic theory and the naive bootstrap when the MA parameter is near or on the unit circle is confirmed by their poor coverage rates. By contrast,

Table 4. 90% Two-Sided, Equal-Tailed Confidence Intervals (GARCH(1,1) Errors)

Methods	*Coverage rate*	*Median length*	*Coverage rate*	*Median length*	*Coverage rate*	*Median length*
	$\theta_0 = .90$		$\theta_0 = .95$		$\theta_0 = 1.00$	
ASYMPTOTIC						
χ^2–*LR* test of $\theta = \theta_0$	.9415	.1198	.9520	.0918	.9530	.0899
DD–*LR* test of $\theta = \theta_0$	.8475	.1226	.8785	.0948	.8810	.0927
BOOTSTRAP						
Percentile-*LR*	.7310	.0921	.9535	.0535	.9525	.0529
GRID BOOTSTRAP						
LR test of $\theta = \theta_0$	.9000	.1388	.9030	.1094	.9055	.1080

the grid bootstrap that employs a weighted resampling scheme exhibits excellent coverage and precision properties.

5. APPLICATIONS

5.1 Parameter Instability in U.S. Inflation

This application investigates the effects of the different monetary policy operating procedures of the Federal Reserve on monthly U.S. inflation for the period January 1959 to December 1998. It is widely documented that the U.S. inflation rate contains a large MA component that might lead to significant size distortions of the standard unit root tests and overrejections of the null hypothesis for presence of a unit root in the level of inflation. The results from some studies (for example, Ng and Perron 1997) suggest that the inflation rate can be better characterized as a nonstationary process but with a strong tendency for mean reversion. The observed nonstationary behavior of inflation may well result from parameter shifts over time rather than a presence of unit root. To investigate if the instability of inflation can be caused by shifts in the MA parameter, we work with the first differences of inflation. By adopting this approach, we practically wash out the potential mean shifts in inflation and keep the autoregressive parameter fixed at 1 by imposing a unit root in the autoregressive polynomial. All the remaining instability should result from instabilities in the MA part.

First, we use the procedure proposed by Bai and Perron (1998) to determine endogenously the break dates for the mean of inflation. To do this, we leave all the autocorrelation in the residuals and estimate the dates of structural shifts in the mean. This procedure allows us to obtain consistent estimates of the break dates although it is not valid for testing the number of breaks. Setting the number of break dates to four, the dates of mean shifts were estimated as 1967:05, 1973:01, 1978:12, and 1981:09. The identified break dates split the sample into five subperiods roughly corresponding to five monetary regimes in which the Federal Reserve followed different operating procedures (Strongin 1995), namely,

1959–1966: free reserves targeting before the modern Federal funds market

1967–1972: free reserves and bank credit growth targeting

1973–1978: money growth and Federal funds targeting

1979–1981: nonborrowed reserves targeting

1982–1998: borrowed reserves and Federal funds targeting.

The dynamics of the monthly inflation rate and the sample mean for each of the five subperiods are plotted in Figure 2. After identifying the break dates, we take the first differences of the inflation series. The Bayesian information criterion selects the MA(1) model as the optimal parsimonious representation of changes in inflation. The estimated MA parameter for the whole sample is .755 with a standard error of .030. The 90% confidence interval obtained from the grid bootstrap with LR test inversion is [.674, .812]. Note that the other grid bootstrap methods as well as the conventional bootstrap produce very similar values of the lower and upper confidence limits. Because the U.S. inflation data admit a number of outliers, it is interesting to check the robustness of the estimation results with respect to these outliers. The biggest outliers in the change of inflation series are in August and September 1973, when the changes in monthly inflation are 1.79% and −1.35%, which appear to be rather large compared to the standard deviation of .25% for the series. By removing these two observations and reestimating the model, we obtain an estimate for the MA parameter of .743 (standard error .031), which is an indication that the results from the fitted model are not sensitive to the presence of these outliers.

As a next step, we estimate an MA(1) model for each subsample and construct confidence intervals for the MA parameter by bootstrapping the LR test of $\theta = \theta_0$ on a grid of 20 points. If the confidence intervals across different subsample

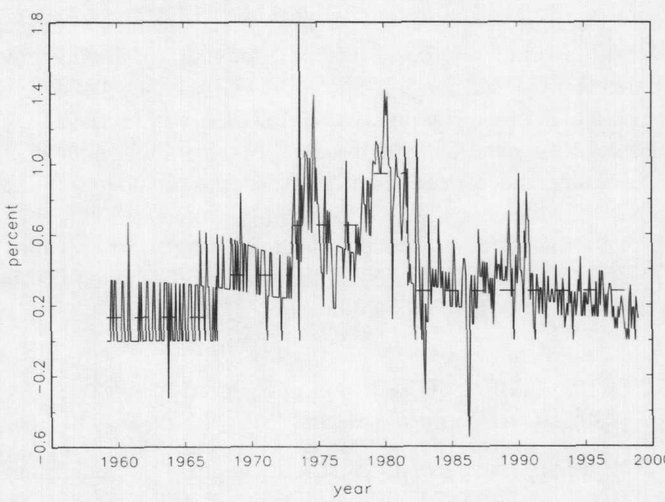

Figure 2. Monthly U.S. Inflation and Subsample Means. The four break dates are determined using the Bai-Perron (1998) procedure.

Table 5. ML Estimates and Grid Bootstrap Confidence Intervals for MA Parameter of Inflation Changes

Sample period	Coefficient	Std. error	90% CI	80% CI
1959 : 03 – 1998 : 12	.755	.030	[.674, .812]	[.697, .801]
1959 : 03 – 1967 : 04	.958	.037	[.880, .991]	[.901, .981]
1967 : 05 – 1972 : 12	.851	.068	[.718, .923]	[.756, .915]
1973 : 01 – 1978 : 11	.797	.074	[.591, .900]	[.648, .889]
1959 : 03 – 1978 : 11	.797	.040	[.715, .852]	[.738, .842]
1978 : 12 – 1981 : 08	.330	.169	[–.002, .737]	[.073, .634]
1981 : 09 – 1998 : 12	.857	.038	[.631, .944]	[.642, .936]

NOTE: CI = confidence interval.

pairs do not overlap, this is evidence that the parameter instability in inflation is caused by changes in the MA component. The ML estimates, their standard errors, and the grid bootstrap confidence intervals based on inversion of the LR statistic for the MA parameter of inflation in each period are presented in Table 5.

The results from Table 5 show that in four of the five subsamples, except for the 1978:12–81:08 period, the estimates of the MA parameter are close to each other and range between .797 and .958. The constructed confidence intervals clearly suggest that there is no evidence of structural instability in the MA parameter in these regimes. This finding would imply that the identified breaks in the inflation dynamics must result either from mean shifts due to some large disturbances such as the oil shock or from changes in the persistence of inflation in its autoregressive part. Because there is no significant change in the MA estimates in the first three subsamples, the results for the combined period 1959:03–78:11 are reported.

The results for the 1979–1981 (Volcker) regime support the hypothesis of a change in the monetary policy operating procedures. In this period, the MA parameter is closer to 0 and the inflation series can be better modeled as a random walk process. The 80% confidence intervals of θ in this subsample do not overlap with the confidence intervals for the 1959–1978 period. This provides some evidence that the differences in the MA parameter and hence the effects of the monetary policy rules on the price dynamics are statistically significant at 20% significance level.

The dynamic behavior of the inflation changes in the period after 1981 also appears to be statistically different from the new operating procedures regime at 20% significance level. By contrast, the 90% confidence intervals do not provide any statistical evidence for a change in the MA parameter before and after the Volker period. Note, however, that the inversion of the LR statistic tends to produce considerably wider confidence intervals for the last two subsamples compared to the grid percentile and the grid percentile-t methods. This may reflect some additional uncertainty (for instance, model misspecification) and asymmetry in the likelihood function that is not captured by the nonlikelihood statistics.

5.2 Time-Varying Risk Premium in Term Structure of Interest Rates

The term structure of interest rates relates the equilibrium yields on bonds of different maturities (for an introduction to term structure models, see Campbell, Lo, and MacKinlay 1997). The relationship between the yields to maturity on a two-period discount bond, $r_{2,t}$, and a one-period discount bond, $r_{1,t}$, is given by

$$r_{1,t+1} = f_{1,t} - \gamma_t + (r_{1,t+1} - E_t r_{1,t+1}), \tag{6}$$

where $f_{1,t} = 2r_{2,t} - r_{1,t}$ is the one-period forward rate and γ_t denotes a possibly time-varying risk premium.

If we assume that the risk premium is constant over time and the expectation errors are nonsystematic, we can test if the forward rate is an unbiased predictor of the future spot rate. In a regression framework, this is equivalent to a test of the restriction $\alpha_1 = 1$ in the levels regression $r_{1,t+1} = \alpha_0 + \alpha_1 f_{1,t} + \epsilon_{t+1}$ or $\beta_1 = 1$ in the differences equation $\Delta r_{1,t+1} = \beta_0 + \beta_1(f_{1,t} - r_{1,t}) + \varepsilon_{t+1}$ (Zivot 2000). In both cases, there is ample empirical evidence for rejection of the forward unbiasedness hypothesis. Two possible explanations, suggested in the literature, are a time-varying term premium and a peso effect. The peso problem may arise, for example, from the anticipation of investors of a change in the monetary regime that does not occur. As a result of this, it may appear that the agents behave irrationally by making systematic mistakes. This possibility is not explored further here. Instead, we focus on the properties and the dynamic behavior of the term premium.

A time-varying risk premium leads to misspecification and omitted variable bias in the levels and differences equations. Typically, the time-varying risk premium is assumed to be stationary. A stationary risk premium, however, cannot explain the rejection of $H_0 : \alpha_1 = 1$ in the levels equation because the omitted variable bias has only a second-order effect in a model with (near-) integrated regressors. For this reason, we adopt an unobserved component framework in which the latent risk premium follows a random walk. Evans and Lewis (1994) suggested that if the term premium contains a small random walk component, one may be able to reconcile the empirical rejections of the unbiasedness hypothesis with the economic theory.

The analysis in this section builds on the work by Stock and Watson (1998), who considered point and interval estimation of the coefficient variance in a time-varying parameter model and used it to investigate a possible decline of the U.S. gross domestic product growth rate due to productivity slowdown. The present approach differs from the one proposed by Stock and Watson (1998) in two respects. First, we use a bootstrap rather than an asymptotic approximation to obtain interval estimators. The interval estimators obtained by the grid bootstrap are expected to possess better properties in small samples than the asymptotic approximation. The applicability of the grid bootstrap procedure described in Section 2.2

is validated by the equivalent representation of the local level model as a restricted MA(1) model. Second, because we are also interested in the upper confidence limit of the coefficient variance, the method of inverting a two-sided LR test would provide shorter intervals (that is, less trend variability) and coverage closer to the nominal level. This claim is supported by the simulation results presented in Table 1.

By rearranging equation (6), we decompose the observed forward errors $f_{1,t} - r_{1,t+1} \equiv x_t$ into a risk premium $f_{1,t} - E_t r_{1,t+1} \equiv \gamma_t$ and expectation errors $E_t r_{1,t+1} - r_{1,t+1} \equiv u_t$. Then, we can use the local level model, which is a very flexible reparameterization of the MA(1) model of the form

$$x_t = \gamma_t + u_t,$$
$$\gamma_t = \gamma_{t-1} + \tau \xi_t,$$

where γ_t is an unobserved, time-varying parameter, u_t and ξ_t are mutually uncorrelated white noise disturbances, and τ is the signal-to-noise ratio.

Taking differences and defining $\Delta x_t = y_t$, we obtain $y_t = \tau\xi_t + \Delta u_t$. It is straightforward to show that this model possesses the same autocorrelation structure as the restricted MA(1) model $\Delta x_t = e_t - \theta e_{t-1}$ with the constraint that $0 \leq \theta \leq 1$. In fact, there exists a one-to-one mapping between the parameters of the two representations τ and θ, namely, $\tau = \sqrt{(1-\theta)^2/\theta}$ and $\theta = (\tau^2 + 2 - \sqrt{\tau^4 + 4\tau^2})/2$, which are monotonic in θ and τ, respectively. This implies that testing for an MA unit root $\theta = 1$ is equivalent to testing for constancy of γ, that is, $H_0 : \tau = 0$ against $H_1 : \tau > 0$. The local-to-unity parameterization of the MA parameter $\theta_0 = 1 - c/T$ in the restricted MA(1) model corresponds to a local-to-zero parameterization of the signal-to-noise ratio $\tau = \lambda/T$ in the local level model. Stock and Watson (1998) used this nesting and several tests for parameter constancy to obtain median unbiased estimates of λ and τ. In this application, we use the Nyblom test for parameter constancy (Nyblom 1989) with $H_0 : \lambda = 0$. The Nyblom statistic has the form $L_T = \sum_{t=1}^{T} \zeta_t^2/\hat{\sigma}^2$, where $\tilde{x}_t = x_t - \bar{x}$ is the demeaned process, $\zeta_t = \sum_{i=1}^{t} \tilde{x}_i$ is a partial sum, and $\hat{\sigma}^2 = T^{-1}\sum_{t=1}^{T} \tilde{x}_t^2$. This test is asymptotically equivalent to Tanaka's S_T test for an MA unit root discussed in Section 4.1.

For the likelihood-based interval estimates, we work with the MA(1) model and then obtain the corresponding values for τ using the monotonic relationship $\tau = \sqrt{(1-\theta)^2/\theta}$. This transformation procedure is valid only if the method for interval estimation of θ is transformation respecting. (For a definition of a transformation respecting method, see Hall 1992, p. 128.) In the MA(1) model with $\theta \leq 1$, the proposed test-inversion methods are transformation respecting because the likelihood-based test statistics are independent of the nuisance parameter vector η_0.

Monthly data for one-month spot and forward rates on U.S. government securities for the period January 1965 to February 1991 were taken from McCulloch and Kwon (1993). Figure 3 plots the dynamics of the forward errors $f_{1,t} - r_{1,t+1}$. The series exhibits very little serial correlation that can be ignored in the subsequent analysis but strong conditional heteroscedasticity. A local level model with a time-varying conditional variance still has white noise disturbances, and its reduced form can be approximated reasonably well by an MA(1) model with conditionally heteroscedastic errors (Harvey, Ruiz, and Sentana 1992). The possible conditional heteroscedasticity, however, needs to be taken into account in mimicking the dynamics of the model for the bootstrap approximation.

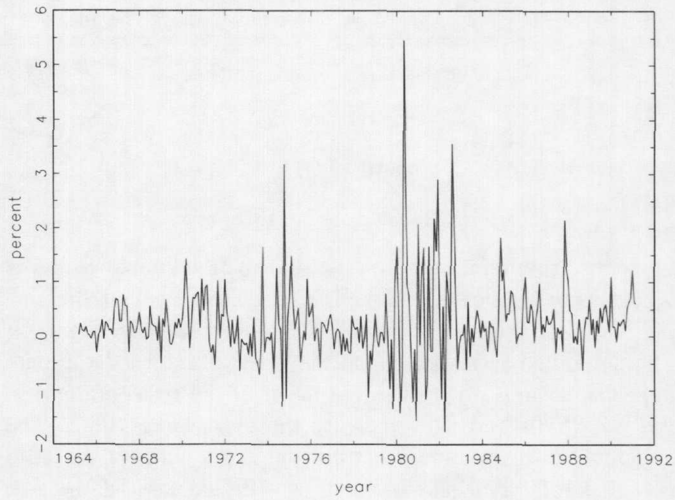

Figure 3. Forward Errors in One-Month T-Bill Rates. The forward errors are computed as a difference between the one-month-ahead forward rate in period t and the one-month spot rate at time t + 1.

The results from constructing median unbiased estimates and 90% two-sided, equal-tailed confidence intervals for the local-to-zero parameter $\lambda = \tau T$ and the signal-to-noise ratio are summarized in Table 6. The ML estimate of λ is 4.231 with sample size $T = 312$. The grid bootstrap method uses 20 grid points with 999 bootstrap replications at each point.

The median unbiased estimates of λ produced by the inversion of the asymptotic and bootstrap one-sided L_T tests are 6.91 and 5.60, respectively. These values imply a small but nonzero variance estimate for the latent risk premium. It is evident also that the median unbiased estimates take into account the pile-up effect and shift the initial ML estimate of λ further away from 0. The two-sided confidence intervals, except the one based on the bootstrap L_T test, include 0 and we cannot reject the null that the risk premium does not contain a random walk component or, equivalently, that the forward errors are a stationary process ($\theta = 1$). This follows immediately from the result that if a stationary series is modeled as $x_t = \gamma + u_t$, where u_t is a stationary invertible ARMA(p, q) process, then $\Delta x_t = u_t - u_{t-1}$ possesses a unit root in the MA polynomial.

As expected from the theoretical discussion and the simulation results, the inversion of the LR test of $\lambda = \lambda_0$ produces shorter intervals than those based on inversion of one-sided tests of $\lambda = 0$, which implies a smaller upper bound of the variability of the risk premium. The confidence intervals for the signal-to-noise ratio from all estimators include .05. Based on the simulation results of Evans and Lewis (1994), a time-varying risk premium that evolves as a random walk over time with small variability (5% of the total variance) can explain the puzzling behavior of the excess returns. Finally, the median unbiased estimate from the bootstrap L_T test of $\lambda = 0$ and the upper confidence limit of λ from the LR test of $\lambda = \lambda_0$ were used to derive a smoothed estimate (using the Kalman filter)

Table 6. Median Unbiased Estimates and 90% Two-Sided Confidence Intervals for λ

Methods	MUE of λ	Implied STNR	CI of λ	CI of STNR
Asymptotic L_T test of $\lambda = 0$	6.911	.0222	[0, 27.732]	[0,.0889]
Grid bootstrap L_T test of $\lambda = 0$	5.594	.0179	[.362,28.190]	[0,.0904]
Grid bootstrap LR test of $\lambda = \lambda_0$			[0, 25.391]	[0,.0814]

NOTE: MUE = median unbiased estimate; CI = confidence interval; STNR = signal-to-noise ratio.

of the risk premium based on the information contained in the whole sample, that is, $\gamma_{t|T}$. The extracted level of the risk premium and the two smoothed series are plotted in Figure 4. The smoothed series based on the median unbiased estimate indicates an upward shift in the mean of the risk premium in the late 1970s and slow increase throughout the 1980s. The smoothed estimate obtained from the upper limit of the 90% confidence interval shows much steeper increase of the risk premium in the period 1979–1982 with a peak in October 1981. Figure 4 suggests that the nonstationary behavior of the risk premium may have been driven by a one-time mean shift rather than a slowly evolving random walk component.

6. CONCLUDING REMARKS

In this article, a method was proposed for obtaining interval estimators for the MA component in near-noninvertible models. It was shown by simulations that inverting the acceptance region of the likelihood ratio statistic is the most appropriate Gaussian approximation for interval construction when the MA parameter is well into the interior of the invertibility region. As the MA parameter approaches unity, the lack of asymptotic pivotalness of the standard test statistics renders conventional bootstrap methods inappropriate and first-order incorrect. For this reason, a modified bootstrap-based method was suggested for constructing confidence intervals when the MA root is near or on the unit circle. An advantage of this method over the asymptotic approximations is that it works globally over the entire parameter space. The obtained two-sided confidence sets for the MA parameter in (nearly) noninvertible models are characterized by excellent coverage and precision. The construction and the properties of one-sided (upper-tail) confidence intervals and median unbiased estimates also were investigated. This statistical procedure was applied to study the possible instability in the MA component of U.S. inflation that may reflect some changes in conducting the monetary policy over the 1959–1998 period. Also, interval estimators were constructed for the variability of the risk premium in interest rates that may help explain the empirical rejection of the unbiasedness of the forward rate as a predictor of the future spot rate.

An interesting finding of this study is that the Davis-Dunsmuir asymptotics provide a very reliable guide to conduct inference in MA(1) models with a large MA root and iid errors. Unfortunately, the Davis-Dunsmuir asymptotic representation cannot be readily extended to more general models and the appropriate distribution theory has not been developed. This puts significant limitations to the use of this asymptotic approximation for econometric analysis because most of the economic time series possess richer dynamic behavior than the one imposed by the MA(1) model. Therefore, it will prove useful if we could generalize the proposed grid bootstrap method to ARMA(p, 1) models or stochastic component models with autocorrelated errors. In principle, we can do this by constructing a grid only for the parameter of interest and replacing the nuisance parameters by their ML estimates. This flexibility of the grid bootstrap is one of its major advantages, and its extension to a broader class of economic processes appears to be a promising direction for future research.

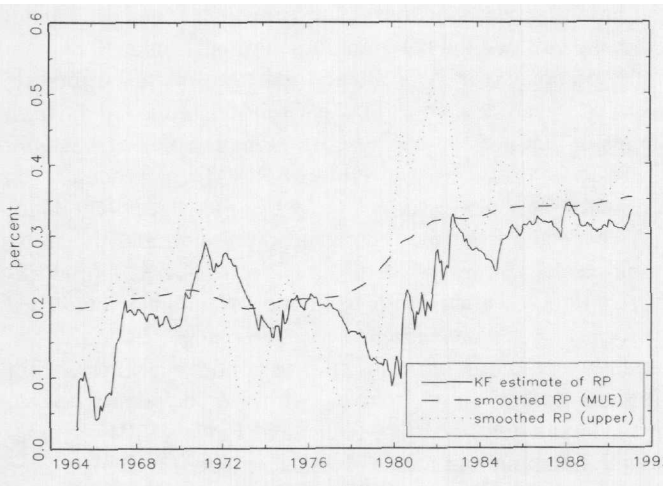

Figure 4. Trend Estimates of the Risk Premium (RP) in One-Month T-Bill Rates. The figure plots the Kalman filter estimate of the risk premium and two smoothed estimates based on the median unbiased estimate (MUE) and the upper 90% confidence limit of the signal-to-noise ratio.

ACKNOWLEDGMENTS

This research was based on Chapter 1 of my Ph.D. dissertation in the Department of Economics at Boston College. I am indebted to Bruce Hansen for his invaluable advice, suggestions, and encouragement and Serena Ng for insightful comments and discussions. I thank the editor, an associate editor, an anonymous referee, Jushan Bai, Arthur Lewbel, and the participants at the 1999 North American Summer Meeting and 1999 European Meeting of the Econometric Society for helpful comments and suggestions.

APPENDIX A: MATHEMATICAL PROOFS

Preliminary Lemma

For model (1) with $\theta_0 = 1 - c/T$ and $\theta = 1 - \beta/T$, and under Assumption 1, we have that

(i) $$[l'_T(\beta), l''_T(\beta)] \xrightarrow{d} \left[\frac{\beta}{4} H_c(\beta), \frac{\beta}{4} H'_c(\beta) + \frac{1}{4} H_c(\beta)\right], \qquad \text{(A.1)}$$

where $H_c(\beta)$ is defined in the text,

(ii) $$\hat{\beta}^{l\max} \xrightarrow{d} \tilde{\beta}_c^{l\max}, \tag{A.2}$$

where $\hat{\beta}^{l\max} = \inf\{\beta \geq 0 : l'_T(\beta) = 0 \text{ and } l''_T(\beta) < 0\}$ is the local maximum of $l_T(.)$ closest to 0 and $\tilde{\beta}_c^{l\max}$ is the local maximizer of $Z_c(\beta) = \int_0^{\beta} (\delta/4) H_c(\delta)\, d\delta$,

(iii) $$T(\hat{\theta} - 1) \xrightarrow{d} -\tilde{\beta}_c^{gl}, \tag{A.3}$$

where $\tilde{\beta}_c^{gl}$ is the global maximizer of $Z_c(\beta)$,

Proof. See the proof of Theorem 2.1 by Davis and Dunsmuir (1996) for (i) and (ii) and the proof of Theorem 2.1 by Davis et al. (1995) for (iii).

Proof of Lemma 1

From (A.1), (A.3), and the continuous mapping theorem, it follows that

$$l_T(\hat{\theta}) \xrightarrow{d} \int_0^{\tilde{\beta}_c^{gl}} \frac{r}{4} H_c(r)\, dr.$$

Also, for a fixed c,

$$l_T(\theta_0) \xrightarrow{d} \int_0^{c} \frac{s}{4} H_c(s)\, ds.$$

Substituting for $H_c(.)$ and evaluating the integral expressions deliver the desired result.

Proof of Theorem 1

To prove the consistency of the bootstrap estimator of the distribution of $\mathrm{LR}(\theta_0)$, we need to verify each of the three conditions stated in Section 3.

For part (i), we need to establish that $\lim_{T\to\infty} \Pr\{d_{\Xi}(\widehat{F}_T, F_0) > \delta\} = 0$ for any $\delta > 0$ and $F_0 \in \Xi$. Define $d_2(F, H) = \inf\left(E|X - Y|^2\right)^{1/2}$ over all joint distributions for the random variables X and Y whose marginal distributions are F and H, respectively. This is the Mallows metric of degree 2 (see Bickel and Friedman 1981, sec. 8). Let $\{e_t\}_{t=1}^T$ and $\{\hat{e}_t\}_{t=1}^T$ denote the original (unobserved) errors and the estimated residuals that are iid sequences with population distribution F_e and sampling distributions $F_{e,T}$ and $\widehat{F}_{e,T}$, respectively. The sequence of residuals $\{\hat{e}_t\}_{t=1}^T$ is calculated as the difference between y_t and its optimal linear projection $E(y_t|y_{t-1}, \ldots, y_1)$ conditional on the past values. It can be shown that this projection error converges in mean square to the true innovations even for $\theta = 1$ (Hamilton 1994, p. 97). Then, Kreiss and Franke (1992) showed that for ARMA models

$$d_2(F_{e,T}, F_e) \to 0 \text{ almost surely as } T \to \infty,$$

$$d_2(\widehat{F}_{e,T}, F_e) \to 0 \text{ in probability as } T \to \infty. \tag{A.4}$$

Next, let $y_t = e_t - (1 - \frac{c}{T})e_{t-1}$ and $y_t^* = e_t^* - (1 - \frac{c}{T})e_{t-1}^*$ be the data-generating and the bootstrap processes where $e_t^* = \nu_t \hat{e}_t$ and $\{\nu_t\}$ is an independent sequence from a common distribution F_ν that satisfies $E(\nu_t) = 0$, $E(\nu_t^2) = E(\nu_t^3) = 1$. By applying Lemmas 8.6 and 8.3 from Bickel and Friedman (1981) and using (A.4), it follows that for a fixed c

$$\begin{aligned} d_2(\widehat{F}_T, F_0)^2 &= \inf\left(E|y^* - y|^2\right) \leq E|y_t^* - y_t|^2 \\ &\leq E(e_t^* - e_t)^2 - \left|1 - \frac{c}{T}\right| E(e_{t-1}^* - e_{t-1})^2 \\ &\leq E(\hat{e}_t - e_t)^2 - E(\hat{e}_{t-1} - e_{t-1})^2 + o_p(1) \\ &\to 0 \text{ in probability as } T \to \infty. \end{aligned} \tag{A.5}$$

This shows that the first condition for consistency of the bootstrap is satisfied. Note that the main result is driven by the fact that the bootstrap data are generated under the null hypothesis $H_0 : \beta = c$. The conventional bootstrap fails to provide a consistent estimator of F_0 because $(\hat{\beta} - c) = O_p(1)$ from (A.3).

For part (ii), the continuity of the asymptotic distribution of the LR statistic is ensured by the local-to-unity framework.

Finally, for part (iii), to show the weak convergence of $G_T^*(q \mid \widehat{F}_T)$ to its limiting distribution $G_\infty(q \mid F_0)$, we rewrite the concentrated log-likelihood in the following canonical form (Anderson and Takemura 1986):

$$M(\rho(\theta)) = -\sum_{t=1}^{T} \log(1 + 2\rho d_t) - T \log\left[\sum_{t=1}^{T} \frac{z_{t,T}^2}{1 + 2\rho d_t}\right], \tag{A.6}$$

where $\rho(\theta) = \frac{-\theta}{1+\theta^2}$ is the correlation coefficient, $d_t = \cos(\frac{\pi t}{T+1})$, $z_{t,T} = \sqrt{\frac{2}{T+1}} \sum_{s=1}^{T} y_s \sin(\frac{\pi st}{T+1})$, and $z_{t,1}, z_{t,2}, \ldots, z_{t,T} \sim \mathrm{nid}(0, \sigma_y^2(1 + 2\rho d_t))$.

The likelihood function depends on the data only through the term z_t in (A.6). Then the bootstrap counterpart $z_{t,T}^*$, generated under the null hypothesis $H_0 : \beta = c$, has the form

$$z_{t,T}^* = \sqrt{\frac{2}{T+1}} \sum_{s=1}^{T} \left(e_s^* - \left(1 - \frac{c}{T}\right) e_{s-1}^*\right) \sin\left(\frac{\pi st}{T+1}\right).$$

Using the same arguments as those in the proof of Proposition A.2 by Davis and Dunsmuir (1996), we can show that for any positive integer m

$$\frac{1}{\sigma_{y^*}\sqrt{1 + 2\rho d_t}} (z_{1,T}^*, z_{2,T}^*, \ldots, z_{m,T}^*) \xrightarrow{d} (x_1, x_2, \ldots, x_m),$$

where $x_1, x_2, \ldots, x_m$ are nid(0,1) variables.

The convergence result given in (iii) follows from (A.5), the continuous mapping theorem and the proofs of Theorem 2.1 by Davis and Dunsmuir (1996) and Theorem 2.1 by Davis et al. (1995).

APPENDIX B: DATA DESCRIPTION

The inflation rate series in Section 5.1 is constructed from the U.S. consumer price index for all urban consumers: all items, 1982–1984 = 100, seasonally adjusted. The data source is the U.S. Department of Labor, Bureau of Labor Statistics. The data were downloaded from the Federal Reserve Bank of St. Louis database at http://www.stls.frb.org/fred/data/cpi/cpiaucsl.

The data series for one-month spot and forward rates on U.S. government securities in Section 5.2 were taken from McCulloch and Kwon (1993). They consist of annualized, continuously compounded, tax-adjusted zero-coupon yield rates derived using a cubic spline discount function.

[Received December 1999. Revised April 2001.]

REFERENCES

Anderson, T. W., and Takemura, A. (1986), "Why Do Noninvertible Estimated Moving Averages Occur?," *Journal of Time Series Analysis*, 7, 235–254.

Andrews, D. W. K. (2000), "Inconsistency of the Bootstrap When a Parameter is on the Boundary of the Parameter Space," *Econometrica*, 68, 399–405.

Bai, J., and Perron, P. (1998), "Estimating and Testing Linear Models With Multiple Structural Changes," *Econometrica*, 66, 47–78.

Basawa, I. V., Mallik, A. K., McCormick, W. P., Reeves, J. H., and Taylor, R. L. (1991), "Bootstrap Test of Significance and Sequential Bootstrap Estimation for Unstable First-Order Autoregressive Processes," *Communications in Statistics-Theory and Methods*, 20, 1015–1026.

Beran, R., and Ducharme, G. R. (1991), *Asymptotic Theory for Bootstrap Methods in Statistics*, Montreal, Canada: Les Publications CRM.

Bickel, P. J., and Freedman, D. (1981), "Some Asymptotic Theory for the Bootstrap," *Annals of Statistics*, 9, 1196–1217.

Bose, A. (1990), "Bootstrap in Moving Average Models," *Annals of the Institute of Statistical Mathematics*, 42, 753–768.

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

Carpenter, J. (1998), "Assessing Parameter Uncertainty via Bootstrap Likelihood Ratio Confidence Regions," *Journal of Applied Statistics*, 25, 639–650.

Casella, G., and Berger, R. L. (1990), *Statistical Inference*, Pacific Grove, CA: Wadsworth and Brooks/Cole.

Chandra, T. K., and Ghosh, J. K. (1979), "Valid Asymptotic Expansions for the Likelihood Ratio Statistic and Other Perturbed Chi-Square Variables," *Sankhya A*, 41, 22–47.

Christiano, L. J. (1987), "Cagan's Model of Hyperinflation Under Rational Expectations," *International Economic Review*, 28, 33–49.

Chuang, C.-S., and Lai, T. L. (2000), "Hybrid Resampling Methods for Confidence Intervals" (with comments), *Statistica Sinica*, 10, 1–50.

Davidson, R., and MacKinnon, J. G. (1999), "Size Distortion of Bootstrap Tests," *Econometric Theory*, 15, 361–376.

Davis, R. A., Chen, M., and Dunsmuir, W. T. M. (1995), "Inference for MA(1) Processes With a Root on or Near the Unit Circle," *Probability and Mathematical Statistics*, 15, 227–242.

Davis, R. A., and Dunsmuir, W. T. M. (1996), "Maximum Likelihood Estimation for MA(1) Processes With a Root on or Near the Unit Circle," *Econometric Theory*, 12, 1–29.

Efron, B. (1981), "Nonparametric Standard Errors and Confidence Intervals" (with discussion), *Canadian Journal of Statistics*, 9, 139–172.

Evans, M. D. D., and Lewis, K. K. (1994), "Do Risk Premia Explain it All? Evidence From the Term Structure," *Journal of Monetary Economics*, 33, 285–318.

Ferretti, N., and Romo, J. (1996), "Unit Root Bootstrap Tests for AR(1) Models," *Biometrika*, 83, 849–860.

Hall, P. (1992), *The Bootstrap and Edgeworth Expansion*, New York: Springer-Verlag.

Hamilton, J. D. (1994), *Time Series Analysis*, Princeton, NJ: Princeton University Press.

Hansen, B. E. (1999), "The Grid Bootstrap and the Autoregressive Model," *Review of Economics and Statistics*, 81, 594–607.

Härdle, W., and Mammen, E. (1993), "Comparing Nonparametric Versus Parametric Regression Fits," *Annals of Statistics*, 21, 1926–1947.

Harvey, A., Ruiz, E., and Sentana, E. (1992), "Unobserved Component Time Series Models With ARCH Disturbances," *Journal of Econometrics*, 52, 129–157.

Horowitz, J. L. (in press), "The Bootstrap," in *Handbook of Econometrics* (Vol. 5), J. J. Heckman and E. E. Leamer (eds.), New York: North–Holland.

Kabaila, P. (1993), "Some Properties of Profile Bootstrap Confidence Intervals," *Australian Journal of Statistics*, 35, 205–214.

Kreiss, J.-P., and Franke, J. (1992), "Bootstrapping Stationary ARMA Models," *Journal of Time Series Analysis*, 13, 297–317.

Lehmann, E. L. (1986), *Testing Statistical Hypotheses*, New York: John Wiley and Sons.

Mankiw, N. G. (1982), "Hall's Consumption Hypothesis and Durable Goods," *Journal of Monetary Economics*, 10, 417–425.

McCulloch, J. H., and Kwon, H.-C. (1993), "U.S. Term Structure Data, 1947–1991," Working Paper 93-6, Ohio State University, Dept. of Economics.

Ng, S., and Perron, P. (1997), "Estimation and Inference in Nearly Unbalanced, Nearly Cointegrated Systems," *Journal of Econometrics*, 79, 54–81.

Nyblom, J. (1989), "Testing for the Constancy of Parameters Over Time," *Journal of the American Statistical Association*, 84, 223–230.

Ploberger, W. (1999), "A Complete Class of Tests When the Likelihood is Locally Asymptotically Quadratic," unpublished manuscript, University of Rochester, Dept. of Economics.

Sargan, J. D., and Bhargava, A. (1983), "Maximum Likelihood Estimation of Regression Models With First-Order Moving Average Errors When the Root Lies on the Unit Circle," *Econometrica*, 51, 799–820.

Schwert, G. W. (1987), "Effects of Model Specification on Tests for Unit Roots in Macroeconomic Data," *Journal of Monetary Economics*, 20, 73–105.

Sprott, D. A. (1981), "Comment on Nonparametric Standard Errors and Confidence Intervals," by B. Efron, *Canadian Journal of Statistics*, 9, 167–168.

Stock, J. H., and Watson, M. W. (1998), "Median Unbiased Estimation of Coefficient Variance in a Time-Varying Parameter Model," *Journal of the American Statistical Association*, 93, 349–358.

Strongin, S. (1995), "The Identification of Monetary Policy Disturbances: Explaining the Liquidity Puzzle," *Journal of Monetary Economics*, 35, 463–497.

Tanaka, K. (1990), "Testing for a Moving Average Unit Root," *Econometric Theory*, 6, 433–4.

Wu, C. F. J. (1986), "Jackknife, Bootstrap and Other Resampling Methods in Regression Analysis" (with discussion), *Annals of Statistics*, 14, 1261–1350.

Zivot, E. (2000), "Cointegration and Forward and Spot Exchange Rate Regressions," *Journal of International Money and Finance*, 19, 785–812.