

UC Santa Barbara

Departmental Working Papers

Title

Obtaining Critical Values for Test of Markov Regime Switching

Permalink

<https://escholarship.org/uc/item/3685g3qr>

Authors

Steigerwald, Douglas G
Bostwick, Valerie K

Publication Date

2012-10-20

Obtaining Critical Values for Test of Markov Regime Switching

Valerie K. Bostwick and Douglas G. Steigerwald

Department of Economics

University of California, Santa Barbara

October 20, 2012

Abstract

For Markov regime-switching models, testing for the possible presence of more than one regime requires the use of a non-standard test statistic. Carter and Steigerwald (forthcoming, *Journal of Econometric Methods*) derive in detail the analytic steps needed to implement the test of Markov regime-switching proposed by Cho and White (2007, *Econometrica*). We summarize the implementation steps and address the computational issues that arise. A new command to compute regime-switching critical values, `rscv`, is introduced and presented in the context of empirical research.

1 Introduction

Markov regime-switching models are frequently used in economic analysis and are prevalent in a variety of fields including finance, industrial organization, and business cycle theory. Unfortunately, conducting proper inference with these models can be exceptionally challenging. In particular, testing for the possible presence of more than one regime requires the use of a non-standard test statistic and critical values that may differ across model specifications.

Cho and White (2007) demonstrate that, due to the unusually complicated nature of the null space, the appropriate measure for a test of more than one regime in the Markov regime-switching framework is a quasi-likelihood ratio (QLR) statistic. They provide an asymptotic null distribution for this test statistic from which critical values should be drawn. Because this distribution is a function of a Gaussian process, the critical values are not easily obtained from a simple closed-form distribution. Moreover, the elements of the Gaussian process underlying the asymptotic null distribution are dependent upon one another. For this reason the critical values depend on the covariance of the Gaussian process and, due to the complex nature of this covariance structure, are best calculated using numerical approximation. In this article we summarize the steps necessary for such an approximation and introduce the new Stata command, `rscv`, which implements the methodology to produce the desired regime-switching critical values for a QLR test of only one regime.

We focus on the case of a simple linear model with Gaussian errors, but the QLR test and the `rscv` command are generalizable to a much broader class of models. This methodology can be applied to models with multiple covariates and non-Gaussian

errors. It is also applicable to regime-switching models where the dependent variable is vector valued, although the difference between distributions must be in only one mean parameter. Although most regime-switching models are thought of in the context of time-series data, we provide an example in Section 5 of how the QLR test can be used in cross-section models. However, there is one notable restriction on the allowable class of regime-switching models. Carter and Steigerwald (2012) establish that the quasi-maximum likelihood estimator created by the use of the quasi-log-likelihood is inconsistent if the covariates include lagged values of the dependent variable. For this reason, the QLR test should be used with extreme caution on autoregressive models.

The article is organized as follows. In Section 2 we describe the unusual null space that corresponds to a test of only one regime versus the alternative of regime-switching. In Section 3 we present the QLR test statistic, as derived by Cho and White, and the corresponding asymptotic null distribution. We also summarize the detailed analysis in Carter and Steigerwald describing the covariance structure of the relevant Gaussian process. In Section 4 we describe the methodology used by the `rscv` command to numerically approximate the relevant critical values. We also present the syntax and options of the `rscv` command and provide sample output. We illustrate use of the `rscv` command with an application from economic literature in Section 5. Finally, we conclude with some remarks on the general applicability of this command and the underlying methods.

2 Null Hypothesis

Specification of a Markov regime-switching model requires a test to confirm the presence of multiple regimes. The first step is to test the null hypothesis of a single regime against the alternative hypothesis of Markov switching between two regimes. If this null hypothesis can be rejected, then the researcher can progress to estimation of Markov regime-switching models with two, or more, regimes. The key to conducting valid inference is then a test of the null hypothesis of a single regime, which yields an asymptotic size equal to or less than the nominal test size.

To understand how to conduct valid inference for the null hypothesis of only a single regime, consider a basic regime-switching model

$$y_t = \theta_0 + \delta s_t + u_t, \tag{1}$$

where $u_t \sim i.i.d.N(0, \sigma^2)$. The unobserved state variable $s_t \in \{0, 1\}$ indicates regimes: in state 0, y_t has mean θ_0 , while in state 1, y_t has mean $\theta_1 = \theta_0 + \delta$. The sequence $\{s_t\}_{t=1}^n$ is generated by a first-order Markov process with $\mathbb{P}(s_t = 1 | s_{t-1} = 0) = p_0$ and $\mathbb{P}(s_t = 0 | s_{t-1} = 1) = p_1$.

The key is to understand the parameter space that corresponds to the null hypothesis. Under the null hypothesis there exists a single regime, with mean θ_* . Hence the null parameter space must capture all the possible regions that correspond to a single regime. The first region corresponds to the assumption that $\theta_0 = \theta_1 = \theta_*$, which carries with it the implicit assumption that each of the two regimes is observed with positive probability: $p_0 > 0$ and $p_1 > 0$. The non-standard feature of the null

space is the inclusion of two additional regions, each of which also correspond to a single regime, with mean θ_* . The second region corresponds to the assumption that only regime 0 occurs with positive probability, $p_0 = 0$, and that $\theta_0 = \theta_*$. Note that in this second region, the mean of regime 1, θ_1 is not identified, so that this region in the null hypothesis does not impose any value on $\theta_1 - \theta_0$. The third region is a mirror image of the second region, where now the assumption is that regime 1 occurs with probability 1: $p_1 = 0$ and $\theta_1 = \theta_*$. The three regions are depicted in Figure 1. The vertical distance measures the value of p_0 and of p_1 and the horizontal distance measures the value of $\theta_1 - \theta_0$. Thus the vertical line at $\theta_1 = \theta_0$ captures the region of the null parameter space that corresponds to the assumption that $\theta_0 = \theta_1 = \theta_*$ together with $p_0, p_1 \in (0, 1)$. The lower horizontal line captures the region of the null parameter space where $p_0 = 0$ and $\theta_1 - \theta_0$ is unrestricted. Similarly, the upper horizontal line captures the region of the null parameter space where $p_1 = 0$ and $\theta_1 - \theta_0$ is unrestricted.




Figure 1: All three regions of the null hypothesis $H_0 : p_0 = 0$ and $\theta_0 = \theta_*$; $p_1 = 0$ and $\theta_1 = \theta_*$; or $\theta_0 = \theta_1 = \theta_*$ together with local neighborhoods of $p_1 = 0$ and $\theta_0 = \theta_1 = \theta_*$.

The additional curves that correspond to the values $p_0 = 0$ and $p_1 = 0$ play a crucial role in guarding against the misclassification of a small group of extremal values as a second regime. In Figure 1 we depict the null space together with local neighborhoods for two points in this space. These two neighborhoods illustrate the different roles of the three curves in the null space. Points in the circular neighborhood of the point on $\theta_1 - \theta_0 = 0$ correspond to processes with two regimes that have only slightly separated means. On the other hand, points in the semicircular neighborhood around the point on $p_1 = 0$ correspond to processes in which there are two regimes with widely separated means, one of which occurs infrequently. As one is often concerned that rejection of the null hypothesis of a single regime is due to a small group of outliers, rather than multiple regimes, including these boundary values reduces precisely this type of false rejection. Consequently, a valid test of the null hypothesis of a single regime must account for the entire null region and include all three curves.

3 Derivation of the Critical Values

3.1 Quasi-Likelihood Ratio Test Statistic

To implement a valid test of the null hypothesis of a single regime, a likelihood ratio statistic is needed. When considering the likelihood ratio statistic for a Markov regime-switching process, Cho and White find that the necessary inclusion of $p_0 = 0$ and $p_1 = 0$ in the parameter space creates significant difficulties in the asymptotic analysis. These difficulties lead them to consider a quasi-likelihood ratio (QLR)

statistic for which the Markov structure of the state variable is ignored and $\{s_t\}$ is instead a sequence of i.i.d. random variables.

This i.i.d. restriction allows Cho and White to consider only the stationary probability, $\mathbb{P}(s_t = 1) = \pi$, where $\pi = p_0/(p_0 + p_1)$. Because $\pi = 1$ if and only if $p_1 = 0$ (and $\pi = 0$ if and only if $p_0 = 0$), the null hypothesis for a test of one regime based on the QLR statistic is expressed with three curves. The null hypothesis is, $H_0 : \theta_0 = \theta_1 = \theta_*$ (curve 1), $\pi = 0$ and $\theta_0 = \theta_*$ (curve 2), $\pi = 1$ and $\theta_1 = \theta_*$ (curve 3). The alternative hypothesis is $H_1 : \pi \in (0, 1)$ and $\theta_0 \neq \theta_1$.

For our basic model in (1), the quasi-log-likelihood analyzed by Cho and White is

$$L_n(\pi, \sigma^2, \theta_0, \theta_1) = \frac{1}{n} \sum_{t=1}^n l_t(\pi, \sigma^2, \theta_0, \theta_1), \quad (2)$$

where $l_t(\pi, \sigma^2, \theta_0, \theta_1) := \log((1 - \pi)f(y_t|\sigma^2, \theta_0) + \pi f(y_t|\sigma^2, \theta_1))$ and $f(y_t|\sigma^2, \theta_j)$ is the conditional density with $j = 0, 1$. Define $(\hat{\pi}, \hat{\sigma}^2, \hat{\theta}_0, \hat{\theta}_1)$ to be the parameter values that maximize the quasi-log-likelihood function. Let $(1, \tilde{\sigma}^2, \cdot, \tilde{\theta}_1)$ be the parameter values that maximize L_n under the null hypothesis that $\pi = 1$. The QLR statistic is then

$$QLR_n = 2n \left(L_n(\hat{\pi}, \hat{\sigma}^2, \hat{\theta}_0, \hat{\theta}_1) - L_n(1, \tilde{\sigma}^2, \cdot, \tilde{\theta}_1) \right). \quad (3)$$

3.2 Asymptotic Null Distribution

The asymptotic null distribution of QLR_n will be largely determined by the behavior of the statistic in the area of the null region near to $\pi = 1$. Note that this region of the null space is equivalent to the space near $\pi = 0$, where the identification

problem is reversed (θ_1 is not identified) but the behavior of the asymptotic null distribution is identical. We will also discuss the behavior of the distribution in the area of the null region near to $\theta_0 = \theta_1 = \theta_*$, however that region is less critical, as it only affects the behavior of the asymptotic null distribution when $\pi = 1/2$.

In the null region near to $\pi = 1$ the distribution is complicated by the fact that θ_0 is not identified, so changes in the value of θ_0 will not alter the distribution. In consequence, while we expect $\hat{\theta}_1$ and $\hat{\sigma}^2$ to be close to their population values when $\hat{\pi}$ is close to 1, there is no expectation that $\hat{\theta}_0$ will be close to its true population value.

To derive the asymptotic null distribution for the region corresponding to $\pi = 1$, first consider θ_0 as fixed at some value θ'_0 . The distribution of QLR_n in this region is then simply χ^2 with one degree-of-freedom. Let this distribution, when θ_0 is fixed at θ'_0 , be expressed as $\mathcal{G}(\theta'_0)^2$, where $\mathcal{G}(\theta'_0) \sim N(0, 1)$. As the value θ'_0 is arbitrary, the distribution of QLR_n depends on the stochastic process formed from the sequence of χ^2 random variables, each indexed by a particular value of θ_0 . Moreover, because the estimates of $(\pi, \sigma^2, \theta_1)$ that maximize the likelihood for any two values of $\theta_0 = \{\theta'_0, \theta''_0\}$ are calculated from the same sample, the elements of this sequence must be dependent upon each other. Therefore, the corresponding sequences $\mathcal{G}(\theta'_0)^2$ and $\mathcal{G}(\theta''_0)^2$ are dependent.

Additionally, we must consider the estimation of π . Because maximization of the likelihood can yield a $\hat{\pi} > 1$, the fact that π is a probability implies that $\hat{\pi}$ should be subject to the boundary condition, $\pi \leq 1$. When we impose this boundary condition on $\hat{\pi}$, the asymptotic null distribution of QLR_n is no longer a χ^2 process. To see this,

note first that the boundary condition $\pi \leq 1$ implies that if $\hat{\pi} > 1$, then the estimate of π is truncated back to $\hat{\pi} = 1$ and $QLR_n = 0$. Thus, when the boundary condition is imposed, the asymptotic null distribution of QLR_n has a point mass at 0. To understand the remainder of the null distribution, we have that the estimator $\hat{\pi}$ is asymptotically equal to $1 + c\mathcal{G}(\theta_0)$, where c is a positive constant, which implies that if $\hat{\pi} < 1$ then $\mathcal{G}(\theta_0) < 0$. Hence the remainder of the null distribution is governed by the negative part of the Gaussian process $\mathcal{G}(\theta_0)$.

Let Θ define the set of all possible values of θ_0 . The procedure of first maximizing L_n for a fixed value of θ_0 and then obtaining the supremum over Θ , yields the asymptotic null distribution (Cho and White, 2007, Theorem 6(a), p. 1692)

$$QLR_n \Rightarrow \sup_{\Theta} (\min [0, \mathcal{G}(\theta_0)])^2. \quad (4)$$

We have thus far only considered the behavior of QLR_n in a neighborhood of the null region corresponding to $\pi = 1$. When the error term is normally distributed, as in our basic model (1) with $u_t \sim i.i.d.N(0, \sigma^2)$, then the asymptotic null distribution of QLR_n is not determined solely by this null region. If θ_0 is sufficiently close to θ_1 and $\pi = \frac{1}{2}$, then the asymptotic null distribution has an additional term (Cho and White, 2007, Theorem 6(b), p. 1692),

$$QLR_n \Rightarrow \max \left[[\max(0, G)]^2, \sup_{\Theta} [\mathcal{G}(\theta_0)_-]^2 \right], \quad (5)$$

where G is a standard Gaussian random variable that is correlated with $\mathcal{G}(\theta_0)$ and $\mathcal{G}(\theta_0)_- := \min [0, \mathcal{G}(\theta_0)]$.

The critical value for a test based on the statistic QLR_n thus corresponds to a quantile for the largest value over $\max(0, G)^2$ and $\sup_{\Theta} [\mathcal{G}(\theta_0)_-]^2$. In order to determine this quantity one must account for the covariance among the elements of $\mathcal{G}(\theta_0)$ as well as their covariance with G . The structure of this covariance is derived in detail by Carter and Steigerwald. We next present a summary of these derivations for the single equation linear model given in (1).

3.3 Gaussian Process Covariance

Since the Gaussian process $\mathcal{G}(\theta_0)$ arises from the behavior of QLR_n in a neighborhood of the null region $\pi = 1$, Carter and Steigerwald first consider the score of π evaluated at $(1, \sigma^2, \theta_0, \theta_*)$ (the population values under the null hypothesis that $\pi = 1$),

$$\mathcal{S}(\theta_0) = \frac{\partial}{\partial \pi} l_t \Big|_{(1, \sigma^2, \theta_0, \theta_*)}.$$

Note that because $\mathcal{S}(\theta_0) \sim N(0, \mathcal{V}(\theta_0))$, the standardized process $\mathcal{G}(\theta_0)$ is given by the scaled score function

$$\mathcal{G}(\theta_0) = \mathcal{V}(\theta_0)^{-\frac{1}{2}} \mathcal{S}(\theta_0).$$

Because the process $\mathcal{G}(\cdot)$ is a Gaussian process, the dependence among the elements of $\mathcal{G}(\cdot)$ is captured by the covariance among the elements of $\mathcal{G}(\cdot)$. If θ_0 and θ'_0 denote two distinct elements of the process $\mathcal{G}(\cdot)$, then the covariance $\mathbb{E}[\mathcal{G}(\theta_0) \mathcal{G}(\theta'_0)]$ is derived from the covariance $\mathbb{E}[\mathcal{S}(\theta_0) \mathcal{S}(\theta'_0)]$ as

$$\mathbb{E}[\mathcal{G}(\theta_0) \mathcal{G}(\theta'_0)] = \mathcal{V}(\theta_0)^{-\frac{1}{2}} \mathcal{V}(\theta'_0)^{-\frac{1}{2}} \mathbb{E}[\mathcal{S}(\theta_0) \mathcal{S}(\theta'_0)]. \quad (6)$$

Carter and Steigerwald use this fact to derive the covariance for the single equation linear model in (1) where the score is

$$\mathcal{S}(\theta_0) = 1 - \exp \left[\frac{(\theta_0 - \theta_*)}{\sigma^2} \left(y_t - \frac{\theta_0 + \theta_*}{2} \right) \right].$$

They derive the asymptotic variance of $\mathcal{S}(\theta_0)$,

$$\mathcal{V}(\theta_0) = \left(e^{\frac{1}{\sigma^2}(\theta_0 - \theta_*)^2} - 1 - \frac{(\theta_0 - \theta_*)^2}{\sigma^2} - \frac{(\theta_0 - \theta_*)^4}{2\sigma^4} \right)^{-1},$$

and show that the covariance of the score, $\mathbb{E}[\mathcal{S}(\theta_0) \mathcal{S}(\theta'_0)]$, equals $\mathcal{V}(\theta_0) \mathcal{V}(\theta'_0)$ times the following term

$$e^{\frac{1}{\sigma^2}(\theta_0 - \theta_*)(\theta'_0 - \theta_*)} - 1 - \frac{(\theta_0 - \theta_*)(\theta'_0 - \theta_*)}{\sigma^2} - \frac{(\theta_0 - \theta_*)^2(\theta'_0 - \theta_*)^2}{2\sigma^4}.$$

Because the regime-specific parameters θ_0 and θ_* enter $\mathbb{E}[\mathcal{S}(\theta_0) \mathcal{S}(\theta'_0)]$ and $\mathcal{V}(\cdot)$ only through $\eta = \frac{\theta_0 - \theta_*}{\sigma}$, the covariance of the Gaussian process can therefore be written as

$$\mathbb{E}[\mathcal{G}(\theta_0) \mathcal{G}(\theta'_0)] = \frac{e^{\eta\eta'} - 1 - \eta\eta' - \frac{(\eta\eta')^2}{2}}{\left(e^{\eta^2} - 1 - \eta^2 - \frac{\eta^4}{2} \right)^{\frac{1}{2}} \left(e^{(\eta')^2} - 1 - (\eta')^2 - \frac{(\eta')^4}{2} \right)^{\frac{1}{2}}}, \quad (7)$$

where $\eta' = \frac{\theta'_0 - \theta_*}{\sigma}$. The quantity $\sup_{\Theta} [\mathcal{G}(\theta_0)_-]^2$ that appears in the asymptotic null distribution is determined by this covariance. Since the regime-specific parameters enter (7) only through η , a researcher does not need to specify the parameter space Θ to calculate $\sup_{\Theta} [\mathcal{G}(\theta_0)_-]^2$. All that is required is the set H that contains the

number of standard deviations that separate the regime means.

Finally, in order to fully capture the behavior of the asymptotic null distribution of QLR_n , we must also account for the covariance between G and $\mathcal{G}(\theta_0)$. Cho and White show that $\text{Cov}(G, \mathcal{G}(\theta_0)) = \left(e^{\eta^2} - 1 - \eta^2 - \frac{\eta^4}{2}\right)^{-\frac{1}{2}} \eta^4$.

4 The `rscv` Command

The syntax of the `rscv` command is

```
rscv [ , ll(value) ul(value) r(value) q(value) ]
```

where `ll` and `ul` are the lower and upper limits of the interval H , `r` is the number of simulation replications, and `q` is the desired quantile. If no options are specified, `rscv` returns the critical value for a size 5 percent QLR test corresponding to a regime separation of ± 1 standard deviation calculated over 100,000 replications. In this section, we describe the simulation process used to obtain these critical values and how each of the `rscv` command options affect those simulations.

Due to the complexity of the covariance structure underlying the asymptotic null distribution, the quantiles needed to obtain critical values for the QLR test are best calculated using numerical approximation. For a QLR test with size 5 percent, the critical value corresponds to the .95 quantile of the limit distribution given on the right side of either (4) or (5). Because the dependence in the process $\mathcal{G}(\theta_0)$ renders numeric integration infeasible, we construct the quantile by simulating independent replications of the process. As the covariance of $\mathcal{G}(\theta_0)$ depends only on an index η , we do not need to simulate $\mathcal{G}(\theta_0)$ directly. Instead we simulate $\mathcal{G}^A(\eta)$, which we will

construct to have the same covariance structure as $\mathcal{G}(\theta_0)$. The process $\mathcal{G}^A(\eta)$ will therefore provide us with the correct quantile, while relying solely on the index, η .

To construct $\mathcal{G}^A(\eta)$ for the covariance structure in (7) recall that, by a Taylor-series expansion, $e^\eta = 1 + \eta + \frac{\eta^2}{2!} + \dots$. Hence, for $\{\epsilon_j\}_{j=0}^\infty \sim i.i.d.N(0, 1)$:

$$\sum_{j=3}^{\infty} \frac{\eta^j}{\sqrt{j!}} \epsilon_j \sim N\left(0, e^{\eta^2} - 1 - \eta^2 - \frac{\eta^4}{2}\right).$$

Using this fact, our simulated process is constructed as

$$\mathcal{G}^A(\eta) = \left(e^{\eta^2} - 1 - \eta^2 - \frac{\eta^4}{2}\right)^{-\frac{1}{2}} \sum_{j=3}^{J-1} \frac{\eta^j}{\sqrt{j!}} \epsilon_j,$$

where J determines the accuracy of the Taylor-series approximation. Note that the covariance of this simulated process, $\mathbb{E}[\mathcal{G}^A(\eta)\mathcal{G}^A(\eta')]$, is identical to the covariance structure of $\mathcal{G}(\theta_0)$ in (7).

We must also account for the covariance between G and $\mathcal{G}(\theta_0)$. Cho and White establish that this covariance corresponds to the term in the Taylor-series expansion for $j = 4$. For this reason we set $G = \epsilon_4$ so that $\text{Cov}(G, \mathcal{G}(\theta_0)) = \text{Cov}(G, \mathcal{G}^A(\eta))$. The critical value that corresponds to (5) for a test size of 5 percent is therefore the .95 quantile of the simulated value

$$\max\left\{[\max(0, \epsilon_4)]^2, \max_{\eta \in H} [\min(0, \mathcal{G}^A(\eta))]^2\right\}. \quad (8)$$

The `rscv` command executes the numerical simulation of (8) by first generating the series $\{\epsilon_j\}_{j=0}^J \sim i.i.d.N(0, 1)$. For each value in a discrete set of $\eta \in H$, it then

constructs $\mathcal{G}^A(\eta) = \left(e^{\eta^2} - 1 - \eta^2 - \frac{\eta^4}{2}\right)^{-\frac{1}{2}} \sum_{j=3}^{J-1} \frac{\eta^j}{\sqrt{j!}} \epsilon_j$. The command then obtains the value $m_r = \max \left\{ [\max(0, \epsilon_4)]^2, \max_{\eta} [\min(0, \mathcal{G}^A(\eta))]^2 \right\}$ corresponding to the right side of (5) for each replication (indexed by r). Let $\{m_{[r]}\}_{r=1}^R$ be the vector of ordered values of m_r calculated in each replication. The command `rscv` returns the critical value for a test with size q from $m_{[(1-q)R]}$.

For each replication, `rscv` calculates $\mathcal{G}^A(\eta)$ at a fine grid of values over the interval H . To do so, we must specify three quantities: the interval H (which must encompass the true value of η), the grid of values over H (given by the grid mesh), and the number of desired terms in the Taylor-series approximation, J . To understand the interplay in specifying these three quantities, suppose that θ_0 is thought to lie within 3 standard deviations of θ_1 . The interval is $H = [-3.0, 3.0]$ and with a grid mesh of 0.01, the process is calculated at the points $(-3.00, -2.99, \dots, 3.00)$. Because the process is calculated at only a finite number of values the accuracy of the calculated maximum increases as the grid mesh shrinks. For this reason the command `rscv` implements a grid mesh of 0.01, as recommended in Cho and White (2007, p. 1693).

Given the grid mesh of 0.01 and the user-specified interval H , we must determine the appropriate value of J . To do so, consider the approximation error, $\xi_{J,\eta} = \left(e^{\eta^2} - 1 - \eta^2 - \frac{\eta^4}{2}\right)^{-\frac{1}{2}} \sum_{j=J}^{\infty} \frac{\eta^j}{\sqrt{j!}} \epsilon_j$. We want to ensure that, as J increases, the variance of $\xi_{J,\eta}$ is decreasing towards zero. Carter and Steigerwald show that, for large J , $\text{Var}(\xi_{J,\eta}) \leq e^{2J \log \eta - J \log J}$. They therefore recommend that $(\max_H |\eta|)^2 / J \leq 1/2$. Cho and White select $J = 150$ and consider a maximal value of $\eta = 5$, so $\eta^2 / J \leq 1/6$. The command `rscv` implements a value of J such that, for the user-specified interval H , $(\max_H |\eta|)^2 / J \leq 1/2$.

The `rscv` command also allows the user to specify the number of simulation replications and the desired quantile. Note that for large values of H and the default number of replications ($r = 100,000$), the `rscv` command may require more memory than a 32-bit operating system can provide. In this case, the user may need to specify a smaller number of replications in order to calculate the critical values for the desired interval, H . Critical values derived using fewer simulation replications may be stable only to a single significant digit. Table 1 depicts the results of `rscv` for a size 5 percent test over varying values of `ll`, `ul`, and `r`.

Table 1: CRITICAL VALUES FOR LINEAR MODELS WITH GAUSSIAN ERRORS

	H	$[-1, 1]$	$[-2, 2]$	$[-3, 3]$	$[-4, 4]$	$[-5, 5]$
Replications	100,000	4.9	5.6	6.2	6.7	7.0
	10,000	4.9	5.6	6.2	6.6	7.1

Nominal level 5 percent; grid mesh of .01.

5 Example

We present an example, derived from the economics literature, that delineates the key features of the QLR test and the associated `rscv` command. Most importantly, we made clear the restrictions that must be imposed on the model to form the QLR test statistic. Our example is derived from Bloom, Canning, and Sevilla (2003), who test whether the large differences in income levels across countries are better explained by differences in intrinsic geography or by a regime-switching model of multiple equilibria with poverty traps. To this end, the authors use cross-sectional data to analyze the distribution of per capita income levels for countries with similar

exogenous characteristics and test for the presence of multiple regimes.

Unlike the simple model, (1), that we have considered up until now, Bloom et al. present a model that includes several added complexities that are commonly used in regime-switching applications. These additions include covariates with coefficients that vary across regimes, as well as error variances that are regime-specific. The authors also allow the regime probabilities to depend on the included covariates. We will show how to properly construct the QLR test statistic for this type of complex model and how to use the `rscv` command to derive the corresponding critical values.

Bloom et al. propose a model of regime-switching between two equilibria. Regime 1 occurs with probability $p(x)$ and corresponds to countries that are in a poverty trap equilibrium:

$$y = \mu_1 + \delta_1 x + \epsilon_1, \text{Var}(\epsilon_1) = \sigma_1^2. \quad (9)$$

Regime 2 occurs with probability $1 - p(x)$ and corresponds to countries in a wealthy equilibrium:

$$y = \mu_2 + \delta_2 x + \epsilon_2, \text{Var}(\epsilon_2) = \sigma_2^2. \quad (10)$$

In both regimes, y is log Gross Domestic Product (GDP) per capita and x is absolute latitude, which functions as a catchall for a variety of exogenous geographic characteristics.

This model is slightly different from a Markov regime-switching model in that the authors are looking at different regimes in a cross-section, rather than over time. For this reason, the probability of being in either regime is stationary and the unobserved regime indicator is an i.i.d. random variable. These modifications correspond exactly

to those made by Cho and White to create the quasi-log-likelihood, so that in this model the log-likelihood ratio and the QLR are one and the same.

Given this model, we can write the quasi-log-likelihood as

$$L_n(p, \sigma^2, \delta, \mu_1, \mu_2) = \frac{1}{n} \sum_{i=1}^n l_i(p, \sigma^2, \delta, \mu_1, \mu_2),$$

where $l_i(p, \sigma^2, \delta, \mu_1, \mu_2) := \log((1-p)f(y_i|\sigma^2, \delta, \mu_1) + pf(y_i|\sigma^2, \delta, \mu_2))$ and $f(y_i|\sigma^2, \delta, \mu_j)$ is the conditional density for $j = 1, 2$. The QLR statistic is then

$$QLR_n = 2n \left(L_n(\hat{p}, \hat{\sigma}^2, \hat{\delta}, \hat{\mu}_1, \hat{\mu}_2) - L_n(1, \tilde{\sigma}^2, \tilde{\delta}, \cdot, \tilde{\mu}_2) \right). \quad (11)$$

Note that we have simplified the model in (9) and (10) in three ways. First, we do not allow for the regime probability, p , to depend on the exogenous characteristics, x . Second, we restrict the variance of the error terms to be the same across regimes such that $\epsilon_j \sim N(0, \sigma^2) \forall j$.¹ Finally, note that the model in (9) and (10) includes a covariate, latitude, in addition to the regime-specific intercept. Carter and Steigerwald show how the asymptotic null distribution of the QLR test does not depend on any included covariates, however, it does require that the difference between regime distributions be in only the intercept, μ_j . Therefore, while Bloom et al. allow the coefficient on latitude to vary across regimes, we require a constant coefficient, $\delta = \delta_1 = \delta_2$. These restrictions are all necessary for the asymptotic null distribution in (4) and (5) to hold. Without these restrictions, the asymptotic null

¹Bloom et al. assume normally distributed errors but the QLR test also allows for any error distribution within the exponential family.

distribution of the QLR test statistic is unknown and the critical values derived by the `rscv` command are not applicable.

Modifying the quasi-log-likelihood in this way does not diminish the validity of the QLR as a test of a single regime for the model in (9) and (10). Note that under the null hypothesis of one regime there is necessarily only one error variance, only one coefficient for each covariate, and a regime probability equal to 1. Thus, under the null hypothesis, the QLR test will necessarily have the correct size even if the data is accurately modeled by a more complex system. Following a rejection of the null hypothesis using this restricted quasi-log-likelihood, the researcher can then confidently proceed to estimate a model with regime-specific variances and coefficients, if desired.²

Using the same Penn World Table³ and CIA World Factbook⁴ data as in Bloom et al., we fit a model under the null hypothesis of a single regime and under the alternative of two regimes and estimate the corresponding log-likelihood values. These estimates and the resulting QLR test statistic are shown in Table 2.

Finally, we use the command `rscv` to calculate the critical value for the QLR test of size 5 percent. We allow for the possibility that the two regimes are widely separated and set $H = [-5, 5]$. The command and output are shown below.

²With a more complex data generating process these restrictions may however lead to an increase in the probability of failing to reject a false null hypothesis and, hence, a decrease in the power of the QLR test.

³Summers, R., and A. Heston. (1991). "The Penn World Table (Mark 5): An Expanded Set of International Comparisons, 1950-1988," *Quarterly Journal of Economics* 106, 327-368.

⁴The World Factbook 2009. Washington, DC: Central Intelligence Agency, 2009. <https://www.cia.gov/library/publications/the-world-factbook/appendix/appendix-f.html>. Latitude data for countries appearing in the 1985 Penn World Tables and missing from the CIA World Factbook comes from <https://www.google.com/>.

Table 2: QLR TEST OF TWO REGIMES VS. ONE REGIME

	Single Regime	Two Regimes	
		Regime I	Regime II
Constant (μ_1, μ_2)	6.928	6.533	7.815
Latitude (δ)	0.041		0.045
SD of error (σ)	0.802		0.599
Probability of Regime I (p)			0.228
Log likelihood (L_n)	-182.1		-180.0
QLR_n			4.3

```
. rscv ,ll(-5) ul(5) r(100000) q(0.95)
7.051934397
```

Given that this critical value of 7.05 exceeds the QLR statistic of 4.3, we cannot reject the null hypothesis of a single regime.

This result is consistent with the findings of Bloom et al., although they use a different method to obtain the necessary critical values. They report a likelihood ratio and the corresponding critical values for a restricted version of their model where the regime probabilities are fixed (p does not depend on x). Using this restricted model, the authors do not reject the null hypothesis of a single regime. At the time that Bloom et al. were published, researchers had yet to successfully derive the asymptotic null distribution for a likelihood ratio test of regime-switching. For this reason, the authors employ Monte Carlo methods to generate their critical values using random data generated from the estimated relationship given by the model in (9) and (10). The primary disadvantage of this approach is that the derived critical values are then dependent upon the authors' assumptions concerning the underlying

data generating process.

Bloom et al. go on to report a likelihood ratio test of a single regime model against the unrestricted model with latitude-dependent regime probabilities. Using the unrestricted model, the likelihood ratio and simulated critical values allow the authors to reject the null hypothesis in favor of the alternative of two regimes. Because the null distribution derived by Cho and White applies only to the restricted QLR presented in (11), we are unable to use the QLR test, and hence the `rscv` command, to obtain the critical values necessary to evaluate this unrestricted test statistic.

6 Discussion

For the case of a simple linear model with Gaussian errors, we provide a methodology and a new command, `rscv`, to construct critical values for a test of regime-switching. Despite the complexity of the underlying methodology, the execution of `rscv` is relatively simple and merely requires the researcher to provide a range for the standardized distance between regime means. We demonstrate in Section 5 both how these methods can be generalized to a very broad class of models and the restrictions necessary to properly estimate the QLR statistic and utilize the `rscv` critical values.

References

Bloom, D., D. Canning, and J. Sevilla. 2003. Geography and Poverty Traps. *Journal of Economic Growth* 8: 355–378.

Carter, A., and D. Steigerwald. 2012. Testing for Regime Switching: A Comment. *Econometrica* 80(4): 1809–1812.

———. forthcoming. Markov Regime-Switching Tests: Asymptotic Critical Values. *Journal of Econometric Methods* .

Cho, J., and H. White. 2007. Testing for Regime Switching. *Econometrica* 75: 1671–1720.