A New Test for Asset Bubbles

Adam Check*

June 11, 2014

Abstract

I apply a recently developed Markov Switching Time-Varying Parameter (MS-TVP) model to test for bubbles in asset markets. In particular, I adapt the model put forth in Eo and Kim (2012), which takes advantage of the use of hierarchical priors governing the evolution of time-varying parameters in a Markov switching model, to the Augmented Dickey-Fuller (ADF) test for asset bubbles proposed in Hall et al. (1999). This paper expands on the prior literature in two important directions. First, it introduces Bayesian estimation and inference to Hall et al.'s (1999) ADF bubble test. Next, it allows the parameters in the Hall et al. (1999) model to change upon entering each episode of a high return and slow return regime. I find that for periodically collapsing bubbles generated according to the process introduced in Evans (1991), both the MS-TVP and the Bayesian Hall et al. (1999) tests have similar power to detect bubbles.

*I can be contacted at ajc@uoregon.edu

†I would like to thank participants in the University of Oregon Macro Group, who gave me very helpful comments and suggestions.

1 Introduction

Policymakers at the Federal Reserve believe that it is vital to determine whether a price bubble exists in important asset markets.[1] While the appropriate steps to take after recognizing the existence of the bubble are debatable, in order to take any action it would first be necessary to know that the bubble existed. However, as the transcript of the Federal Open Market Committee meeting from June 29 to June 30, 2005 indicates, even during the peak of the massive U.S. housing bubble, policymakers were in disagreement over whether the housing market was in a bubble. This disagreement highlights an important and surprising deficiency in the asset bubble literature - the lack of existence of a powerful and broadly agreed upon test for asset bubbles.

Due to this deficiency, I propose a new test that generalizes the Markov switching test for explosive roots in Hall et al. (1999). Specifically, I allow the parameters of Hall et al.'s (1999) model to vary over time. In principle, this should allow the detection of multiple bubbles in the same sample, even if the growth rates of the bubbles differ. At the same time, I introduce Bayesian estimation of this model and use Bayesian model comparison to decide between competing models, rather than using the classical confidence interval-based inference used in Hall et al. (1999). This Bayesian perspective allows me to easily test jointly for both switching in the price dynamics and an explosive root, whereas Hall et al. (1999) tested only for the presence of an explosive root.

To investigate the power of my proposed test, I use artificially generated price series that contain periodically collapsing bubbles. I first estimate the constant parameter Hall et al. (1999) model using Bayesian methods and altering Hall et al.'s (1999) testing procedure slightly to jointly test for both an explosive root and Markov switching. Next, I estimate my proposed time-varying parameter generalization. I find that for a test with 5% size, Bayesian estimation and testing of both the constant parameter Hall et al. (1999) model

[1]This is evident from a speech given on February 7, 2013 by Federal Reserve Governor Jeremy Stein entitled "Overheating in Credit Markets: Origins, Measurement, and Policy Responses". A transcript of this speech available at: http://www.federalreserve.gov/newsevents/speech/stein20130207a.pdf

and the more general time-varying parameter model are able to detect these periodically collapsing bubbles nearly 80% of the time.

The idea to econometrically test for bubbles in asset markets has been around for decades, originating shortly after the variance bound tests that were proposed contemporaneously by Shiller (1981) and LeRoy and Porter (1981). These tests attempt to determine whether the observed variance of actual asset prices exceeds the variance bound implied by the frequently used risk-neutral asset pricing equation. Tirole (1985) and Blanchard and Watson (1982) suggest that these variance bounds tests be used to detect bubbles, but Flood et al. (1994) eventually showed that variance bounds tests were actually very poorly suited to test for bubbles, since in the presence of a bubble, the variance would not exceed the bound implied by the test.

Another bubble testing procedure was put forth by West (1987), who suggests the use of a two-step test which estimates the relationship between dividends and stock prices both directly and indirectly. If the estimated relationship differs between the two methods, and the researcher is reasonably certain that they have specified the model correctly, then there exists evidence in favor of a bubble. The major problem with this test is that under different model specifications, researchers have come to different conclusions about the existence of bubbles in U.S. stock prices.

Yet another take on bubble testing was put forward by Diba and Grossman (1988), who use the integration/cointegration properties of dividends and prices to test whether prices take on the integration properties of the dividend process. If they do then there is no bubble, since as suggested by theory, the dynamic properties of the price series depends only on the process followed by the dividends. However, if the price series displays integration patterns that are not shared with the dividend process, then this would suggest that something else is also driving asset prices. If we are sure that dividends are the only relevant fundamental, then we would conclude that there is a bubble in asset prices.

More specifically, Diba and Grossman (1988) use an Augmented Dickey-Fuller (ADF) test in order to test both first differenced prices and first differenced dividends for a unit root. If first differenced prices display a unit root, but first differenced dividends do not, then the price series is consistent with a bubble. Using this test on generated price series data using a very simple bubble process that grows exponentially in every period, they find that their test can detect a bubble 95% of the time.

However, Evans (1991) points out that although the Diba and Grossman (1988) test is appealing and works well on data generated with a relatively simple bubble process, when faced with bubbles that grow at different rates in different time periods it loses almost all of its power, and detects only a handful of bubbles. In order to build on Diba and Grossman's (1988) intuitive and appealing idea for a bubble test, Hall et al. (1999) generalize the Diba and Grossman (1988) test. First, they test for an explosive root in the levels of prices and dividends, rather than testing for a unit root in the first differences. Next, they allow the parameters of this test to alternate between two regimes according to a Markov switching process. This allows the test to capture the fact that the bubble is growing much faster in some periods than in others, and will allow at least a subset of periods to be consistent with a bubble. In practice, Hall et al. (1999) find that their test is capable of detecting the Evans (1991) style bubbles about 60% of the time. While this is great improvement over the Diba and Grossman (1988) test, to the dismay of policymakers who wish to determine whether a particular asset is in a bubble, it will still miss the presence of a bubble nearly 40% of the time.

It is through this lens that I introduce a generalization to Hall et al.'s (1999) test for asset bubbles. First, I bring a Bayesian perspective to the test. Second, I allow for the growth rates of bubbles to change upon each episode of bubble, so that the test does not restrict all bubbles in the sample to have the same growth rate. Third, I test jointly for Markov switching and an explosive root in the price series, while Hall et al. (1999) assume Markov switching under the null, and test only for an explosive root. I find that jointly testing proves important for the detection of bubbles that follow Evans' (1991) bubble generation procedure, as it improves the detection rate to nearly 80%. Allowing for differing growth rates does not meaningfully alter the ability of the test to detect for bubbles in this environment, but this may be simply because in expectation, all bubbles all grow at the same rate under this process.

The rest of the paper is organized as follows. In section two, I present some background on rational bubbles and provide more details for the bubble tests that are most closely related to my new bubble test. In section three, I outline my test, based on Eo and Kim (2012), which allows the parameters of a Markov switching ADF test to vary over time. In section four, I detail estimation procedures for both the Bayesian implementation of Hall et al.'s (1999) model as well as the MS-TVP model. and Bayesian model comparison. In section five I introduce Bayesian model comparison and outline my testing procedure. In section six I present my results. I conclude in section seven.

2 Rational Bubbles, Testing, and Periodically Collapsing Bubbles

2.1 Rational Bubbles

Historically, many bubble tests are designed to detect "rational" bubbles. These include the early variance bounds tests developed by Shiller (1981) and LeRoy and Porter (1981), and implemented by Cochrane (1992); the "two-step" tests developed in West (1987); and the integration/cointegration test developed in Diba and Grossman (1988).[2] However, before we can properly discuss bubble testing, we must specify what we mean by a "rational" bubble.

Rational bubbles are periods during which agents are willing to pay more for an asset than the asset's fundamental value, which is the value implied by the present value of expected future dividend streams. The reason that these agents are willing pay a premium during

[2]See Gürkaynak (2008) for more detail.

a rational bubble is that they anticipate being able to sell the asset for more than the fundamental value at a later date. This bubble is rational in the sense that if everyone shares this belief, then the asset is priced correctly despite the fact that it trades for more than its fundamental value.

We can mathematically formalize the above intuition in a simple asset pricing framework. First, assume that there is an infinitely-lived representative agent, who seeks to maximize expected lifetime utility in an endowment economy:

$$\max_{c_t} E_t \left[\sum_{i=0}^{\infty} \beta^i u(c_{t+i}) \right]$$
$$\text{s.t. } c_{t+i} = y_{t+i} + (p_{t+i} + d_{t+i}) x_{t+i-1} - p_{t+i} x_{t+i}$$

Where y_t is the income of the agent at time t, x_t is the number of units of the asset held by the agent in period t, p_t is the price of the asset in period t, d_t is the dividend paid to those holding the asset at the beginning of period t, and $0 < \beta < 1$ is the discount rate of the representative agent.

We can derive the Euler equation using a variational argument. Using the consumption good, c_t, as the numeraire, we consider the gains and losses from giving up a unit of consumption in order to buy the asset today, and we define the net one-period rate of return from holding an asset as:

$$r_{t+1} = \frac{p_{t+1} + d_{t+1}}{p_t} - 1$$

Then we can see that:

$$u'(c_t) = E_t \beta (1 + r_{t+1}) u'(c_{t+1})$$

In words, if we give up a unit of consumption today, then we can use the proceeds to buy the asset, but we lose the marginal utility that the unit of consumption would give us today.

However, tomorrow we will be able to consume $(1 + r_{t+1})$ units of the consumption good, since we can consume the one-period return given by the asset. Therefore, tomorrow we will gain the marginal utility of consuming $(1 + r_{t+1})$ units of the consumption good. Taking into account the fact that we will not get this utility until tomorrow, we average across all possible returns by using the expectations operator and discount the future utility using the discount factor β.

Using our definition of the net return, r_{t+1}, we can substitute and solve directly for the price of the asset today:

$$p_t = \beta E_t \left\{ (p_{t+1} + d_{t+1}) \frac{u'(c_{t+1})}{u'(c_t)} \right\}$$

Assuming risk-neutral preferences, we have $u'(c_t) = u'(c_{t+1}) = k \ \forall \ c_t$, where k is a constant. We can then rewrite the above equation as:

$$p_t = \beta E_t(p_{t+1} + d_{t+1}) \tag{1}$$

We can find a particular solution to this first-order difference equation by using the law of iterated expectations and iterating on the above equation:

$$p_t = \sum_{i=1}^{s} \beta^i E_t d_{t+i} + \beta^s E_t p_{t+s}$$

Taking the limit as $s \to \infty$, we can see that as long as $\beta^s E_t p_{t+s} \to 0$, we have the solution:

$$p_t = \sum_{i=1}^{\infty} \beta^i E_t d_{t+i} = F_t$$

This solution is called the *fundamental* solution, since the price today relies only on the expected value of the future dividend stream. This solution is ensured by imposing a transversality condition on the value of the agent's savings, and that is usual practice in asset pricing models with infinitely lived agents.

However, there is another solution for the first order difference equation in equation (1):

$$p_t = F_t + B_t \tag{2}$$

where B_t, the *bubble* component of the solution, is any random variable that satisfies:

$$B_t = \beta E_t B_{t+1} \tag{3}$$

We verify that this is a solution to equation (1) by a direct proof, given below. First, assume that equations (2) and (3) constitute a solution to (1). Then we have:

$$\begin{aligned}
p_t &= \beta E_t(p_{t+1} + d_{t+1}) \\
F_t + B_t &= \beta E_t(F_{t+1} + B_{t+1} + d_{t+1}) \\
\sum_{i=1}^{\infty} \beta^i E_t d_{t+i} + \beta E_t B_{t+1} &= \beta \left(\sum_{i=2}^{\infty} \beta^i E_t d_{t+i} + \beta E_t B_{t+2} + E_t d_{t+1} \right) \\
\sum_{i=1}^{\infty} \beta^i E_t d_{t+i} + \beta E_t B_{t+1} &= \left(\sum_{i=1}^{\infty} \beta^i E_t d_{t+i} + \beta^2 E_t B_{t+2} \right) \\
\beta E_t B_{t+1} &= \beta^2 E_t B_{t+2} \\
E_t B_{t+1} &= \beta E_t B_{t+2} \qquad (4)
\end{aligned}$$

Iterating equation (3) forward one period, and using the law of iterated expectations, we have:

$$\begin{aligned}
B_{t+1} &= \beta E_{t+1} B_{t+2} \\
E_t B_{t+1} &= \beta E_t B_{t+2}
\end{aligned}$$

which shows that the equality in equation (4) always holds, and verifies that equation (2) nests an entire class of solutions, so long as equation (3) also holds. ■

Finally, note that bubbles of this type are usually ruled out by imposing a transversality condition. In fact, Tirole (1982) argues that bubbles can *always* be ruled out in infinitely lived rational expectations models. However, it is common practice in the bubble testing literature to abstract away from this theoretical argument, and work from the assumption that equations (2) and (3) define the asset price.[3] This approach is somewhat justified by the fact that Tirole (1985) shows that rational bubbles can exist in overlapping generations models, and by the fact that Kindleberger (2000) finds numerous examples of asset bubbles throughout modern history. As pointed out by Evans (1991), if these observed bubbles are not "rational", then a desirable feature of a bubble test would be that it has power against many different bubble specifications.

2.2 Imposing Structure on the General Solution

In order to devise a test for rational asset bubbles, it is necessary to put more structure on the problem above. It is particularly helpful to define dynamic equations for both the dividend and the bubble component. For the dividend component, the typical assumption is that d_t is integrated of order one, i.e. it is $I(1)$.[4] Specifically, d_t is usually assumed to follow a random walk with drift:[5]

$$d_t = \mu + d_{t-1} + \varepsilon_t \tag{5}$$

Then there are two possible scenarios: the asset price does not contain a bubble component, B_t, or the asset price does contain a bubble component.

In the absence of a bubble component, B_t, we have $p_t = F_t$, and it can be shown that p_t

[3] See Gürkaynak (2008)

[4] In a model with dynamic growth, the assumption is instead that $\ln(d_t)$ is $I(1)$. However, we will usually be working in the relatively simpler set-up outlined in the text.

[5] We use a random walk for expositional purposes, but the analysis remains the same for any stationary ARMA process for Δd_t. See Evans (1991) for more details.

is also $I(1)$ and that p_t and d_t are cointegrated. Furthermore, we have:

$$E_t F_{t+i} = \frac{\beta}{(1-\beta)}(d_t + \mu i) + \frac{\beta}{1-\beta}\mu$$

which becomes dominated by $\frac{\beta}{(1-\beta)}\mu i$ as i gets large.[6] This implies that the forecast of the fundamental value grows linearly over time, increasing by $\frac{\beta}{(1-\beta)}\mu$ each period, and reflects the unit root in the process for d_t.

In addition, rearranging equation (3) and using the law of iterated expectations, we can show that the time t expectation of the bubble component at time $t+i$ is given by:

$$E_t B_{t+i} = \frac{1}{\beta^i} B_t$$

Therefore, as pointed out in Evans (1991), the conditional expectation of the bubble component grows at rate $\frac{1}{\beta} > 1$. Combining the two components, we can see that as i gets large we have:

$$E_t p_{t+i} \rightarrow \frac{\beta}{(1-\beta)}\mu i + \frac{1}{\beta^i} B_t$$

Provided $B_t > 0$, eventually the exponential growth of the bubble component will overwhelm the linear growth of the fundamental component, and the forecast of the price will explode to infinity at the rate of the growth of the bubble component, $\frac{1}{\beta}$.

With the additional structure we have put on the asset pricing model, we have a testable hypothesis. If the price of the asset grows faster than the underlying fundamentals grow, then there is a bubble in the asset price.[7] It is using this logic that Diba & Grossman designed their test for bubbles in asset markets.

[6] In general, from Beveridge and Nelson (1981), we know that if d_t follows a stationary ARMA process, then as j tends to infinity, $E_t F_{t+j} \rightarrow C_t + jE(\Delta F_t)$ for some C_t. See also Evans (1991).

[7] Of course, this assumes that all fundamentals are observable.

2.3 Diba and Grossman (1988) Test

Diba and Grossman's (1988) idea is to exploit the specification of a rational bubble, noting that if the price of the asset contained a bubble component, the asset price would grow at a rate faster than suggested by the growth rate of the fundamental process. Furthermore, Diba and Grossman (1988) provide the additional insight that if the fundamental process, d_t is $I(1)$, then the fundamental is stationarity in first differences. Therefore, if there is no bubble component, then the first difference of the asset price, Δp_t, would also be stationary.

However, in the presence of an exponentially growing bubble component, differencing prices any finite number of times will not yield a stationary process for Δp_t. In fact, in the presence of a bubble component, the fundamental and the asset price would not be cointegrated. Therefore, Diba and Grossman (1988) propose the following test:[8]

1. Test p_t and d_t for stationarity.
2. If both p_t and d_t are non-stationary test p_t and d_t for cointegration.
3. If they are cointegrated, conclude that there is no bubble.
4. If they are not, then conclude that we cannot rule out the existence of a bubble.

In fact, if we were certain that we had included all relevant measures of fundamentals in d_t, then the lack of stationarity in the price series or the lack of cointegration between the price series and the fundamental series would indicate the presence of a bubble.

Diba and Grossman (1988) use this test on 100 simulated series, each lasting 100 periods and containing an explosive bubble component that evolves according to $B_{t+1} = (1+r)B_t + z_{t+1}$, where $r = 0.05$ and $z_{t+1} \sim$ iid $N(0, \sigma_z)$. They find that their test has high power to detect a bubble, as 95% of their simulated price series are nonstationary in first differences. However, as pointed out in Evans (1991), the process assumed for the evolution of the bubble

[8] They also propose other, similar tests, such as testing the first differences of both series for stationarity. All of their proposed tests have approximately the same power, and were shown in Evans (1991) to have low power against periodically collapsing bubbles.

may be overly simplistic, since it assumes that a bubble will grow at an exponential rate, with a small amount of noise, forever.

Figure 1: Typical Diba and Grossman Bubble

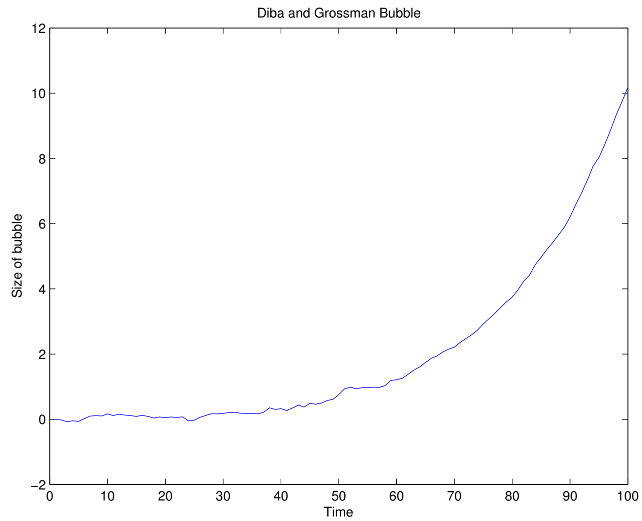

2.4 Evans (1991) Bubble Process

Evans (1991) astutely observes that in the real world, bubbles could not possibly have the form hypothesized by Diba and Grossman (1988). That is, no one thinks that an asset bubble could grow unabated forever. In fact, all of the dozens of examples of historical bubbles cited in Kindleberger (2000) eventually collapsed. Therefore, Evans proposes a process for bubbles that allows them to periodically collapse, and shows that the tests suggested by Diba and Grossman (1988) have very little power to detect this more realistic type of bubble.

Evans (1991) retains the same risk neutral asset market set-up considered in Diba and Grossman (1988), and assumes that dividends evolve according to equation (5). However, he proposes the following formulation for the bubble component:

$$B_{t+1} = \begin{cases} (1+r)B_t u_{t+1} & \text{if } B_t \leq \alpha \\ \left[\delta + \frac{1+r}{\pi}\left(B_t - \frac{\delta}{1+r}\right)\xi_{t+1}\right] u_{t+1} & \text{if } B_t > \alpha \end{cases}$$

Where δ and α are scalars that satisfy $0 < \delta < (1+r)\alpha$, u_t is a sequence of i.i.d. random variables with $E_t u_{t+1} = 1$:

$$u_t = \exp\left(z_t - \frac{\sigma_z^2}{2}\right)$$
$$z_t \sim N(0, \sigma_z^2)$$

and ξ_t is an exogenous i.i.d. Bernoulli process such that:

$$Pr(\xi_t = 0) = 1 - \pi$$
$$Pr(\xi_t = 1) = \pi$$

In words, the bubble process follows a linear switching process. Recall from equation (3) that a rational bubble component must satisfy:

$$B_t = \beta E_t B_{t+1}$$

Therefore, we need to verify that the above process satisfies this requirement. When $B_t \leq \alpha$, we have:

$$E_t B_{t+1} = E_t (1+r) B_t u_{t+1}$$
$$E_t B_{t+1} = (1+r) B_t E_t u_{t+1}$$
$$E_t B_{t+1} = (1+r) B_t$$
$$B_t = \beta E_t B_{t+1}$$

where we have used the fact that in equilibrium, $\frac{1}{\beta} = (1+r)$.

If instead, $B_t > \alpha$, we have:

$$
\begin{aligned}
E_t B_{t+1} &= E_t \left\{ \left[\delta + \frac{1+r}{\pi} \left(B_t - \frac{\delta}{1+r} \right) \xi_{t+1} \right] u_{t+1} \right\} \\
E_t B_{t+1} &= E_t \left\{ \delta u_{t+1} + \frac{1+r}{\pi} \left(B_t - \frac{\delta}{1+r} \right) \xi_{t+1} u_{t+1} \right\} \\
E_t B_{t+1} &= \delta E_t u_{t+1} + \frac{1+r}{\pi} \left(B_t - \frac{\delta}{1+r} \right) E_t \xi_{t+1} u_{t+1} \\
E_t B_{t+1} &= \delta + \frac{1+r}{\pi} \left(B_t - \frac{\delta}{1+r} \right) \pi \\
E_t B_{t+1} &= \delta + (1+r) \left(B_t - \frac{\delta}{1+r} \right) \\
E_t B_{t+1} &= \delta + (1+r) B_t - \delta \\
E_t B_{t+1} &= (1+r) B_t \\
B_t &= \beta E_t B_{t+1}
\end{aligned}
$$

where, in order to go from line 3 to line 4 we have used the fact that ξ_{t+1} and u_{t+1} are each i.i.d. random variables, with $E(\xi_{t+1}) = \pi$ and $E(u_{t+1}) = 1$.

Now that we know that this bubble process conforms to the requirement for a rational bubble, we can analyze some of its properties. If $B_t \leq \alpha$, then the bubble grows at rate $\frac{1}{\beta}$. However, once the size of the bubble exceeds the predetermined level α, then it will grow at the faster rate, $\frac{B_t}{\beta\pi}$ if $\xi_{t+1} = 1$. However, if $\xi_{t+1} = 0$, then the bubble collapses to δu_{t+1}. Therefore, $1 - \pi$ is the probability of the bubble collapsing each period. Once the bubble collapses, it will return to growing at the slower rate, $\frac{1}{\beta}$, until it eventually exceeds the exogenously given scalar α.

After generating 200 price series, using the same fundamentals process as Diba and Grossman (1988), but his newly proposed collapsing bubbles for the bubble component, Evans tests these series for the presence of a bubble by using Diba and Grossman's (1988) proposed unit root and cointegration tests.[9] He finds that these tests perform extremely

[9]See the appendix for full detail and the particular parameter values chosen in Evans (1991).

Figure 2: Typical Evans Bubble

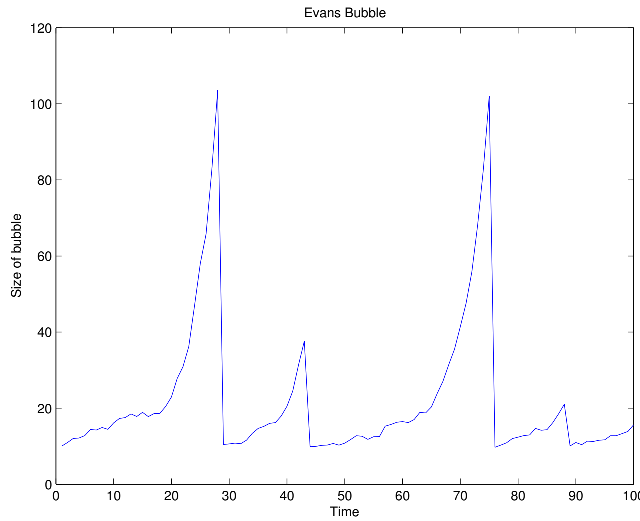

poorly, and detect only a handful[10] of bubbles in these generated price series. Since these collapsing bubbles are a much more plausible bubble generating process than the process set forth in Diba and Grossman (1988), Evans (1991) concludes that the Diba and Grossman (1988) test for asset bubbles is insufficient, and that work should be done to devise a more powerful and flexible bubble test.

2.5 Hall et al.'s (1999) Test for bubbles

Hall et al. (1999) take up the call to action in Evans (1991), and attempt to design a bubble test that is more able to detect the presence of periodically collapsing bubbles than the tests presented in Diba and Grossman (1988). To do so, they return to the Augmented Dickey Fuller (ADF) test that was used in Diba and Grossman (1988), with two main differences. In Diba and Grossman (1988), this ADF test was used to test for the presence of a unit root in p_t or Δp_t. In Hall et al. (1999), the authors instead modify this test to test for an *explosive* root in p_t. Furthermore, Hall et al. (1999) allow the parameters of the test to switch between two regimes: a low return regime and a high return regime.

[10]It is hard to tell from the table presented in Evans (1991). However, it appears that at most three of the 200 bubbles were detected.

The equation that Hall et al. (1999) estimate to conduct their test for an explosive root in the level of the price series, p_t, is given below:

$$\Delta p_t = \mu_0(1 - S_t) + \mu_1 S_t + [\phi_0(1 - S_t) + \phi_1 S_t]p_{t-1} + \sum_{j=1}^{k}[\psi_{0,j}(1 - S_t) + \psi_{1,j}S_t]\Delta p_{t-j} + \sigma_e e_t$$

$$S_t \in \{0, 1\}$$

Here, S_t is an indicator variable, indicating whether we are in a low return regime ($S_t = 1$), or a high return regime (S_t=0). If these regimes were directly observable, then we could estimate the above equation by treating S_t as a dummy variable. It may be helpful to break this equation into its piecewise components before discussing it further:

$$\Delta p_t = \begin{cases} \mu_0 + \phi_0 p_{t-1} + \sum_{j=1}^{k}\psi_{0,j}\Delta p_{t-j} + \sigma_e e_t & \text{if } S_t = 0 \\ \mu_1 + \phi_1 p_{t-1} + \sum_{j=1}^{k}\psi_{1,j}\Delta p_{t-j} + \sigma_e e_t & \text{if } S_t = 1 \end{cases}$$

Here, p_t is the price series of interest,[11] ϕ_i is the AR(1) parameter determining the impact of p_{t-1} on p_t. $\psi_{i,j}$ is the coefficient on the j^{th} lag of the price series, for $j \geq 2$, which determines how p_{t-j} impacts p_t.

Finally, the interpretation of μ_i depends on the estimation of ϕ_i. For $\phi_i < 1$, $\frac{\mu_i}{1-\phi_i}$ is the mean of the price series in regime i. For $\phi_i = 1$, μ_i is the drift (i.e. time trend) of the random walk process for the price series. If $\phi_i > 1$, then μ_i helps determine (along with ϕ_i) the level of the price series that determines whether the explosive price process is exploding to negative or positive infinity.

However, a problem with the procedure outlined above is that the researcher will not know which periods constitute a high return or low return regime, so the regime dummy variable, S_t is unobserved. Therefore, Hall et al. (1999) estimate these regimes using a Markov switching model. In this model, the researcher assumes that the probability of

[11] In applied empirical work (looking at S&P data, for instance), this will actually be the log of the price series. In my simple price generation equations, I can simply use the level of prices.

moving from regime i in period $t-1$ to regime j in period t depends only on what regime the price process was in in period $t-1$. Since there are two regimes in the model, then there are four possible transitions that occur, with probabilities given by p_{ij} for $i, j \in \{0, 1\}$, where p_{ij} is the probability of switching from regime i in period $t-1$ to regime j in period t.

The testing procedure is as follows:

1. Estimate the Markov switching model via Maximum Likelihood estimation.

2. Use bootstrapping to find a one-sided 95% confidence interval against which to test the presence of an explosive root in the high return regime.

3. If this test confirms that the price series, p_t, has an explosive root, but the fundamental series does not, then the price series is consistent with the presence of a bubble.

4. However, if the price series does not have an explosive root, or if both the price series and its corresponding fundamental have an explosive root, then conclude that there is not a bubble.

This generalization of the ADF test goes a long way in allowing Hall et al. (1999) detect the presence of bubbles. Using the equations above, and performing maximum likelihood estimation with bootstrapped errors, Hall et al. (1999) find that the high return regime has an explosive root over 75% of the time. Because the existence of a bubble is only confirmed when the price has an explosive root but dividends do not, Hall et al. (1999) are only able to classify about 60% of their price series as containing a bubble. Compared to the results in Evans (1991), which detected only a handful of bubbles out of 200, this is a great improvement. However, the results may be disappointing to policymakers, as this test still misses about 40% of these seemingly realistic type of bubbles.

Since Hall et al. (1999), there have been more attempts to improve bubble testing. Some recently developed tests, such as Phillips et al. (2011), use recursive tests to test for bubbles and remain agnostic about the structural form of the regime. Using Phillips et al.'s (2011) estimation technique has two desirable features. First, Phillips and Magdalinos (2007) have worked out limit theory for mildly explosive processes, and Phillips et al. (2011) are able to apply that theory to their bubble testing procedure to test directly for explosive roots without the need for bootstrapping, a computationally intensive procedure that Hall et al. (1999) needed to undertake to estimate the distribution of the AR(1) parameter under an explosive processes. Second, it allows the researcher to date-stamp the beginning and ending dates of bubbles, although this can also be done in the Markov Switching framework by using the estimated regime probabilities. However, using the feasible parameters used by Evans (1991) and Hall et al. (1999), Phillips et al.'s (2011) test detects the presence of a bubble only 43.2% of the time.[12]

3 MS-TVP Model

Due to the relatively low power of existing bubble tests, there remains a deficiency in the bubble testing literature. In order to try to increase the power of existing bubble tests, I generalize Hall et al.'s (1999) test by allowing the AR parameters in their MS model to evolve according to a random walk each time the price series enters a high return or low return regime. My test uses Bayesian inference, and therefore admits the use of hierarchical priors, as suggested by Koop and Potter (2007), which makes estimation of this test feasible via Gibbs sampling.

To see where the time-varying nature of my test may be particularly helpful, consider the following example. Suppose that there are three different bubbles in our sample, which is a reasonable number for both Evans' (1991) generated price series and for the entire history of S&P 500 stock data. Assume that two of the three bubbles are very large, with very high growth rates, while the third is relatively small, with a slower growth rate. Then the

[12]This test performs much better for slightly different parameterizations of Evans (1991) bubble generation process that could still be considered realistic.

fixed parameter Hall et al. (1999) test may find the two large bubbles, and estimate a large value for ϕ_0, the AR(1) coefficient in the high growth regime. Since the estimated ϕ_0 is so large, the Hall et al. (1999) test may not detect the third bubble, since it will be more qualitatively similar to the slow growth regime than the bubble regime. However, the new model with time varying parameters should be able to detect this third bubble. Since the AR(1) coefficient in the high growth regime, $\phi_{0,t}$ can change over time, it will be higher during the two large bubbles, and lower, but still explosive, during the third bubble. With this intuition in mind, I will present my generalization of the Hall et al. (1999) test below.

3.1 Eo & Kim's Model Applied to the ADF Test

In Eo and Kim (2012), the authors present a newly developed MS-TVP model in the context of GDP growth and identification of recessions. Eo and Kim (2012) were trying to overcome a similar problem to the example presented above - that Hamilton's (1989) Markov switching model with constant coefficients did a poor job at identifying relatively mild recessions, like the one in 2001. While Eo and Kim (2012) only consider a model that has no autoregressive components, so that it is only a mean switching model, it is easy to generalize their model to one with lags. Then the MS-TVP Augmented Dickey fuller test can be written as:

$$\begin{aligned}\Delta p_t = &\ \mu_{0,\tau}(1 - S_t) + \mu_{1,\tau} S_t + [\phi_{0,\tau}(1 - S_t) + \phi_{1,\tau} S_t] p_{t-1} + \\ &\sum_{j=1}^{k} [\psi_{0,\tau,j}(1 - S_t) + \psi_{1,\tau,j} S_t] \Delta p_{t-j} + \sigma_e e_t\end{aligned}$$

Let $\beta_{0,\tau} = [\mu_{0,\tau}\ \phi_{0,\tau}\ \psi_{0,\tau,1}\ \ldots\ \psi_{0,\tau,k}]'$ and $\beta_{1,\tau} = [\mu_{1,\tau}\ \phi_{1,\tau}\ \psi_{1,\tau,1}\ \ldots\ \psi_{1,\tau,k}]'$ be the vectors of time varying parameters in regime 0 and regime 1, respectively, in a model with $k + 1$

lags. Then the parameters transition according to:

$$\begin{bmatrix} \beta_{0,\tau} \\ \beta_{1,\tau} \end{bmatrix} = \begin{bmatrix} \beta_{0,\tau-1} \\ \beta_{1,\tau-1} \end{bmatrix} + \begin{bmatrix} \omega_{\beta_0,\tau} \\ \omega_{\beta_1,\tau} \end{bmatrix}$$

where $\tau = 1, 2, \cdots, N_0+N_1$, and the $\omega_{x,\tau}$ are white noise shocks particular to each parameter. In words, τ is the number of the particular realization of the regime. For example, if $\tau = 1$ is the first realization of the high return regime, then $\tau = 2$ is the first realization of the low return regime, $\tau = 3$ is the second realization of the high return regime, $\tau = 4$ is the second realization of the low return regime, etc. Therefore N_0 is the number of times the price series has been in the high return regime, N_1 is the number of times the price series has been in the low return regime, and $N_0 + N_1$ is the total number of realizations of all regimes.

For estimation, it will be helpful to both collect all parameters in a vector and rewrite the system in terms of time t instead of τ. First, let $\beta_t = [\beta'_{0,t}\ \beta'_{1,t}]'$ be the vector of all estimated coefficients at time t. Then the system can be written in state space form. First the Measurement Equation (ME):

$$p_t = H_t\beta_t + e_t$$

$$H_t = [(1 - S_t)p_{t-1}\ S_tp_{t-1}\ (1 - S_t)\Delta p_{t-1} \ldots S_tp_{t-k}]$$

$$e_t \sim NID(0, \sigma_e^2)$$

Likewise, we can write the State Equation (SE):

$$\beta_t = F\beta_{t-1} + \omega_t$$

$$F = I_{\text{length}(\beta)}$$

$$\omega_t \sim MVN\left(\begin{bmatrix} 0 \\ 0 \\ \vdots \\ \vdots \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} d_{10,t}\sigma^2_{\mu_0} & 0 & 0 & 0 & \dots & 0 \\ 0 & d_{01,t}\sigma^2_{\mu_1} & 0 & 0 & \dots & 0 \\ 0 & 0 & d_{10,t}\sigma^2_{\phi_0} & 0 & \dots & 0 \\ 0 & 0 & 0 & \ddots & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & d_{01,t}\sigma^2_{\psi_{1,k}} \end{bmatrix}\right)$$

Here, $d_{ij,t}$ is a dummy variable which equals one when $t-1 = i$ and $t = j$. Therefore, the shocks, ω_t, are heteroscedastic - a shock to the regime i parameters only occurs when the price series enters an episode of regime i, and equals 0 otherwise. However, because the disturbance term in the ME is conditionally Gaussian, we are still able to use the Kalman Filter to estimate the parameters of this state space model.

4 Estimation

I use Markov Chain Monte Carlo (MCMC) Bayesian estimation techniques to estimate both Hall et al.'s (1999) model and the time-varying generalization outlined above. In both estimations, I use a Normal prior on the regression coefficients, a Gamma prior on the inverse of the variance parameter, and a Beta prior on the regime transition probabilities. These priors are conditionally conjugate, so they admit the use of the Gibbs sampler, which is the most efficient form of the more general Metropolis-Hastings MCMC estimation technique.

As stated above, in my estimation I use a Normal prior on all of the regression coefficients, including the AR(1) coefficient, $\phi_{i,t}$, at all points in time. There is a vast literature outlining the sensitivity with respect to the choice of prior in testing the root of an AR process, with primary contributions coming from Sims (1988), Berger and Yang (1994), and Lubrano (1995). Xia and Griffiths (2012) demonstrate that for Bayesian posterior confidence interval based tests, a uniform prior over the AR(1) coefficient rejects the null hypothesis of a unit root too infrequently, but performs much better when using Bayesian model comparison. Of all the priors considered by Xia and Griffiths (2012), this uniform prior is most similar to the Normal prior used in this paper.

4.1 Estimation of Hall et al.'s Model

In order to estimate the model presented in Hall et al. (1999), I use the Gibbs sampler with relatively tight priors on the AR(1) coefficients, but diffuse priors on μ_i, the variance parameter, and the transition probabilities. The Gibbs sampler for this model is standard for a two-state Markov Switching regression with autoregressive parameters. For the sake of brevity, I omit a detailed description of the sampler here, but the interested reader can find a detailed exposition of the Gibbs sampler for this model in Kim and Nelson (1999).

Recall that the Hall et al. (1999) model is given by:

$$\Delta p_t = \begin{cases} \mu_0 + \phi_0 p_{t-1} + \sum_{j=1}^{k} \psi_{0,j} \Delta p_{t-j} + \sigma_e e_t & \text{if } S_t = 0 \\ \mu_1 + \phi_1 p_{t-1} + \sum_{j=1}^{k} \psi_{1,j} \Delta p_{t-j} + \sigma_e e_t & \text{if } S_t = 1 \end{cases}$$

Since I will be testing for an explosive root, the prior on the AR(1) coefficient is of extreme importance. I place a tight prior on the AR(1) coefficient in both regimes. In the high return regime (regime 0), the prior on ϕ_0 is a Normal distribution with mean zero and standard deviation set to 0.05 that is truncated to lie above zero (i.e. the AR(1) parameter in the high return regime is restricted to be consistent with explosive growth). This relatively small standard deviation reflects our prior knowledge that even when a root is only very slightly explosive, the process blows up quickly. Therefore, I believe that even in the presence of a bubble, the explosive root will not be very large. This is reflected by my prior, which holds that I have roughly 95% confidence that the AR(1) parameter is less than 1.1.

In the low return regime (regime 1), the prior on the AR(1) parameter, ϕ_1, is also a Normal distribution with mean zero and standard deviation set to 0.05. However, this distribution is truncated from above at ϕ_0, in order ensure uniqueness of the likelihood function of the MS model. This prior again reflects the fact that I believe that the AR(1) parameter is near 0, and if it is explosive, it is probably only very slightly explosive.

4.2 Estimation of MS-TVP Model

As shown in Eo and Kim (2012), with standard prior distributions on the parameters that are being estimated, the conditional posterior distributions can each be derived analytically, so we can use the Gibbs sampler to estimate this model. Since this model is less well known, I will present an overview of the steps involved in the Gibbs sampler below. However, since my estimation procedure is nearly identical to the procedure discussed at length in Eo and Kim (2012), I will keep this overview relatively brief.

As in Hall et al.'s (1999) model, the priors on the AR(1) coefficients are very important. In order to facilitate estimation of the more flexible switching model, I do not center the prior distributions for the initial conditions for the AR coefficients at zero.[13] Instead, for the AR(1) parameter in the high return regime, I set the prior for the initial condition to 0.05, and in the low return regime, I set the prior for the initial condition at -0.05, each with a small variance.

The steps for the Gibbs sampler are as follows:

Step 0:

Initialize the hyperparameters of the model, $\tilde{\Omega} = [\sigma_e^2 \ \sigma_{\mu_0}^2 \ \sigma_{\mu_1}^2 \ldots \ \sigma_{\psi_{1,k}}^2]'$, the time-varying parameters, $\tilde{\beta}_T = [\beta_1 \ \beta_2 \ \ldots \ \beta_T]'$, and transition matrix $\tilde{P} = \begin{bmatrix} p_{00} & 1-p_{00} \\ 1-p_{11} & p_{11} \end{bmatrix}$

Step 1:

Generate the regime for each time period, $\tilde{S}_T = [S_1, S_2, \ldots, S_T]'$, conditional on $\tilde{\beta}_T$, $\tilde{\Omega}$, $\tilde{P}$,

[13] This issue is discussed in more detail in section 6

and data $\tilde{Y}_T$. This is based on the multi-move sampler developed by Carter and Kohn (1994) and explained in Kim and Nelson (1999).

Step 2:

Based on the state space model, generate the time-varying parameters:

$\tilde{\beta_T} = [\mu_{0,T}\ \mu_{1,T}\ \phi_{0,T}\ \phi_{1,T}\ \psi_{0,1,T}\ \psi_{1,1,T}\ \ldots\ \psi_{1,k,T}]'$

conditional on $\tilde{\Omega}$, $\tilde{S}_T$, $\tilde{P}$, and the data, $\tilde{Y}_t = [p_1\ p_2\ \ldots p_T]'$. This can be done by exploiting the state space form of the model to run a Carter and Kohn (1994) algorithm utilizing the Kalman filter.

Step 3:

Generate the hyperparameters of the model, $\tilde{\Omega}$, conditional on $\tilde{\beta_T}$, $\tilde{P}$, $\tilde{S}_T$ and $\tilde{Y}_T$. This is done by exploiting the fact that conditional on the other parameters of the model, each of the state space equations are line by line OLS, as is the measurement equation.

Step 4:

Generate the matrix of transition probabilities, $\tilde{P}$, conditional on $\tilde{S}_T$.

5 Testing Procedure

The estimation method laid out above gives us parameter estimates the for MS-TVP model. However, even after obtaining these estimates, I still need to assess the performance of each test, both in absolute terms and relative to the performance of the Hall et al. (1999) test. To do so, I use two sets of competing models, and use Bayesian model comparison to determine which of the two models is more likely.

To assess the performance of the MS-TVP test, I generate 201 time series each consisting of 100 periods according to Evans' (1991) dynamic price and bubble equations. Then, to test for the presence of an explosive root, I specify two competing models. The first comparison is between the Hall et al. (1999) model and a null model. The second is between the MS-TVP model and a null model. In order to assess the performance of the MS-TVP model relative to the Hall et al. (1999) model, I compare the performance of each test against their corresponding null models.

For the standard Hall et al. (1999) model, I first estimate the model outlined in section 4.1, restricting the AR(1) process to be explosive in at least one of the two regimes. Next, I estimate a model that restricts the price series to behave as it would in the absence of a bubble. In the absence of a bubble, the price series would be driven only by the underlying fundamental, so two features would change. First, and most obviously, the explosive root in the price series would be replaced by a unit root. However, it is also the case that there would be no regime switching, since in the Evans's (1991) bubble generation procedure, regime switching is only a feature of the bubble component. Therefore, the second model simplifies to Bayesian linear regression, with the series restricted to be nonexplosive, i.e. I restrict $\phi \leq 0$. Table 1 summarizes these two competing models.

Table 1: Assessing Hall et al.'s (1999) Test

	Model Equations	Model Restrictions
Model 1	$\Delta p_t = \mu_0 + \phi_0 p_{t-1} + \sum_{j=1}^{k} \psi_{0,j} \Delta p_{t-j} + \sigma_e e_t$ if $S_t = 0$	$\phi_0 > 0$ (Explosive)
	$\Delta p_t = \mu_1 + \phi_1 p_{t-1} + \sum_{j=1}^{k} \psi_{1,j} \Delta p_{t-j} + \sigma_e e_t$ if $S_t = 1$	$\phi_1 \leq \phi_0$ (Identifying Restriction)
Model 2	$\Delta p_t = \mu + \phi p_{t-1} + + \sum_{j=1}^{k} \psi_j \Delta p_{t-j} + \sigma_e e_t$	$\phi \leq 0$ (Non-explosive)

Next, I do the same for the MS-TVP model. The only difference is that for this model, I assume that prices follow an AR(1) process in order to reduce the risk of overfitting that is an ubiquitous concern when using time-varying parameter models. Therefore, when assessing the performance of the MS-TVP model, I use the following two models:

Table 2: Assessing the MS-TVP Test

	Model Equations	Model Restrictions
Model 1	$\Delta p_t = \mu_{0,t} + \phi_{0,t} p_{t-1} + \sigma_e e_t$ if $S_t = 0$	$\phi_{0,t} > 0$ (Explosive)
	$\Delta p_t = \mu_{1,t} + \phi_{1,t} p_{t-1} + \sigma_e e_t$ if $S_t = 1$	$\phi_{1,t} \leq \phi_{0,t}$ (Identifying Restriction)
Model 2	$\Delta p_t = \mu + \phi p_{t-1} + \sigma_e e_t$	$\phi \leq 0$ (Non-explosive)

For both assessments, I use Bayesian model comparison to compare Model 1 to Model 2. In order to do so, I first estimate the marginal density of each model using the decomposition given in Chib (1995):

$$m(Y_T) = \frac{f(Y_T|\tilde{\theta})\pi(\tilde{\theta})}{\pi(\tilde{\theta}|Y_T)} \tag{6}$$

where $\theta = [\Omega, P]$, i.e. it is the collection of all of the hyperparameters that are being estimated. In equation (1), $f(Y_T|\tilde{\theta})$ is the sampling density, $\pi(\tilde{\theta})$ is the prior density of θ, and $\pi(\tilde{\theta}|Y_T)$ is the posterior density of θ, each evaluated at $\theta = \tilde{\theta}$ where $\tilde{\theta}$ is any fixed value of θ in the posterior distribution. Chib (1995) recommends setting $\tilde{\theta}$ equal to a value that occurs with great frequency, such the posterior mean or median, in order to achieve the most accurate approximation. In my application, I use the posterior mean.

Recall that in the MS-TVP model, the shocks in the transition equation are heteroscedastic, since they only occur when there is a shift in regimes. Therefore, there is not an easily computable analytical expression for the sampling density, $f(Y_T|\tilde{\theta})$, in the MS-TVP model. I estimate it using a simple particle filter, based on Fernández-Villaverde and Rubio-Ramírez (2004), which can be used to approximate the marginal density of any parameterized non-linear state space model. In order to decrease the impact of numerical estimation error, I set the number of particles high enough so that to the first decimal place, the filter produces identical estimates.[14]

To compute an estimate of $\pi(\tilde{\theta}|Y_T)$, I use the method suggested in Chib (1995), which consists of running the Gibbs sampler successively, each time holding an additional element of θ, the collection of all estimated hyperparameters, at its posterior mean, $\tilde{\theta}$. Once I have the numerically estimated values for both the sampling density and the posterior density, I can compute the marginal density using the formula in equation (1).

After computing an estimate of the marginal density for each model, I compare the two models by computing the posterior odds ratio. The posterior odds ratio simply describes

[14] In practice, I used 20,000 particles.

how likely one model is relative to another. For example, if the posterior odds ratio for model one compared to model two is 3.0 (so that the odds are 3:1), then I would say, given the data, model one is $\frac{3}{3+1} - \frac{1}{3+1} = 50\%$ more likely to have generated it than model two.

In general, the posterior odds ratio of model one compared to model two can be given as:

$$\frac{pr(m_1|Y_T)}{pr(m_2|Y_T)} = \frac{m_1(Y_T)}{m_2(Y_T)} \frac{pr(m_1)}{pr(m_2)}$$

where m_i denotes model i, and $\frac{pr(m_1)}{pr(m_2)}$ is the prior odds ratio. In my case I assume equal prior probability between the two models, so the latter term drops out and equation (1) becomes:

$$\frac{pr(m_1|Y_T)}{pr(m_2|Y_T)} = \frac{m_1(Y_T)}{m_2(Y_T)} = B_{12}$$

This makes clear that once I have numerically computed my estimates of the marginal densities, I have all I need to compare the two models.

Finally, I need to set criteria that determines when I will prefer one model to another. To do this, I calibrate the testing procedure, by first running competing models on all generated series of fundamentals. After running these competing models on data that I know has been generated according to the second model, I compute the odds ratios, and find the odds ratio for which I would incorrectly prefer model one five percent of the time. In other words, I calibrate the test such the size of the test is approximately five percent.

For example, when comparing the performance of the MS-TVP model with the non-explosive linear regression model, this occurs at an odds ratio equal to 2.16. Therefore, if a given price series has an odds ratio greater than 2.16, and its corresponding fundamental series has an odds ratio less than 2.16, this suggests the presence of a bubble in the price series. In other words, since the price series is explosive but the fundamental series is nonexplosive, this suggests that there is another component aside from the fundamental that is driving the asset price.

However, if both the price series and the fundamental series had an odds ratio greater than 2.16, this would not suggest the existence of a bubble. Since both series are determined to be explosive, it does not suggest that the something other than the fundamental is driving the movements in the price of the asset.

6 Results

6.1 Estimation of Generated Data

Although this procedure is fairly straightforward, I did experience practical difficulties during estimation of the MS-TVP model. When estimating a Markov switching model, we need to enforce some type of inequality restriction on a subset of the parameters to ensure uniqueness of the likelihood. In other words, the likelihood function is symmetric - it makes no difference to the likelihood function whether we label the high return regime as "regime zero" or "regime one".

As my procedure requires a forward run of the Kalman filter,[15] if I have an "unlucky" draw of regimes and hyperparameters, the time-varying parameters may wander outside their restricted region. When I attempt to enforce this restriction on the backward draw via rejection sampling, my sampler may have to sample billions (or more) of times in order to find parameters that fit the restriction. Note that Koop and Potter (2011) point out a similar problem in a time-varying parameter vector autoregression. In their estimation, they compare a multi-move algorithm, which is similar to the Carter-Kohn algorithm I use in estimation of my model, and a single-move algorithm. They find that when they use the multi-move algorithm, their rejection rates are as high as 99.97%. Although in theory the single-move algorithm mixes at a slower rate, the rejection rate is substantially lower than the multi-move algorithm, so in practice Koop and Potter (2011) suggest using a single-move algorithm.

[15]See Step 2 in the previous section.

However, Koop and Potter (2011) perform their analysis within a Metropolis-Hastings setting on a slightly different estimation procedure, so it is not straightforward to generalize their results to my estimation technique. Therefore, to attempt to solve this problem in my model, I set the variances time-varying parameters equal to a small constant instead of estimating them.[16]

Because estimation along with approximation of the marginal density is relatively computationally intensive, I try to strike a balance between accuracy and timeliness. Autocorrelation functions and running mean plots on randomly selected time series simulations suggested the use of at least 5,000 burn-in draws. To be conservative, I chose to use 10,000 burn-in draws. However, in order to also ensure a feasible speed of estimation,[17] I chose a relatively modest 20,000 post burn-in draws.[18]

6.2 Power of MS and MS-TVP Tests

In my estimation, I seek to investigate both the absolute and relative power of the Hall et al. (1999) test and the MS-TVP test. The exact priors used can be found in the appendix. Below, I present the results from the model assessment, first comparing the Hall et al. (1999) model restricted to have an explosive root with a nonexplosive model with no switching, and then comparing the MS-TVP test with an explosive root to a nonexplosive model with no switching. As described in the previous section, I calibrate these tests to have size of .05, i.e. 5% of the time it will incorrectly prefer the model with switching when the true model is the stationary linear process. My results are presented in table 3.

The first conclusion that I draw from these results is that compared to the results in Hall et al. (1999), the Bayesian implementation of the test which jointly tests for both switching

[16]While this fixes the issue in most cases, the issue remains in samples that have a very high variance (about 5% of all price series). When the issue remains, I count that particular series as "no bubble" series.

[17]For each time series, estimation entails the estimation of both competing models, as well as approximating the likelihoods of each model. This takes about 30 minutes in Matlab on a 2013 Macbook Pro with a 2.7 Ghz. Intel i7 processor.

[18]Experimenting on a few randomly selected time series, my results were not very sensitive the increasing the number of post burn-in draws.

Table 3: Power of Bayesian MS Test with 5% Size

	Calibrated Odds Ratio	Percent Containing Bubble
Bayesian Hall et al. (1999)	45.40	78.61%
MS-TVP	2.16	79.60%

and an explosive root displays higher power to detect a the presence of a bubble. Hall et al. (1999) are only able to detect a bubble in approximately 60% of the price series. The second conclusion is that for this particular specification, the more general MS-TVP model does not add much power to the detect a bubble.

The first result, that the Bayesian implementation has more power than the classical estimation in Hall et al. (1999), may be partially driven by the fact that I am jointly testing for both switching and an explosive root, while Hall et al. (1999) test only for an explosive root. A model with an explosive root and switching may represent the price process better than a linear model, even if a switching model with a nonexplosive root might provide a fit of the price series data that is superior to both of these. This intuition is supported by my initial attempt at implementing Bayesian testing in the Hall et al. (1999) model, which found that a version of the test that tested only for an explosive root only detected a bubble about 40% of the time.[19] However, since the underlying dividend series does not have switching, I believe that testing only for an explosive root would be incorrect, and that the test presented in this paper fully exploits the specification of the periodically collapsing bubble process found in Evans (1991).

The second result, that the MS-TVP model does not seem to provide superior bubble detection in this model, is actually quite intuitive. Since the growth rates of all of the bubbles in this model are identical conditional on the realized values of the shocks, the standard Hall et al. (1999) test should be expected to perform relatively well compared to the more general model. Under an alternate time-varying specification of the bubble growth rate, or even in

[19]However, these estimations were conducted with slightly different priors. To the extent that the results are sensitive to the priors, this would fail to be an apples-to-apples comparison.

an application to real world price series that may contain several bubbles, my priors are that the MS-TVP model would perform better relative to the Hall et al. (1999) test.

Finally, for purposes of intuition, I believe it is helpful to see what one of my generated and estimated price series actually looks like. Following Evans (1991), I chose the price series with median variance. The first image in Figure 3 plots the actual asset return and the estimated asset return for all periods. The second image plots the probability of being in the low return regime for all periods. The fact that this probability is virtually zero except for four brief periods suggests that this price series is in a bubble in almost all time periods. For this price series, model comparison suggests that this price series contains a bubble.

In Figure 4, I present the same graphs, estimated instead in the MS-TVP model. In addition, I present the time path of the explosive root, $\phi_{0,t}$, which shows that following the first large bubble collapse, the next bubble contains a more explosive root. Following the collapse of the second bubble, the explosive root is smaller.

Figure 3: Median Bubble (Standard MS Model)

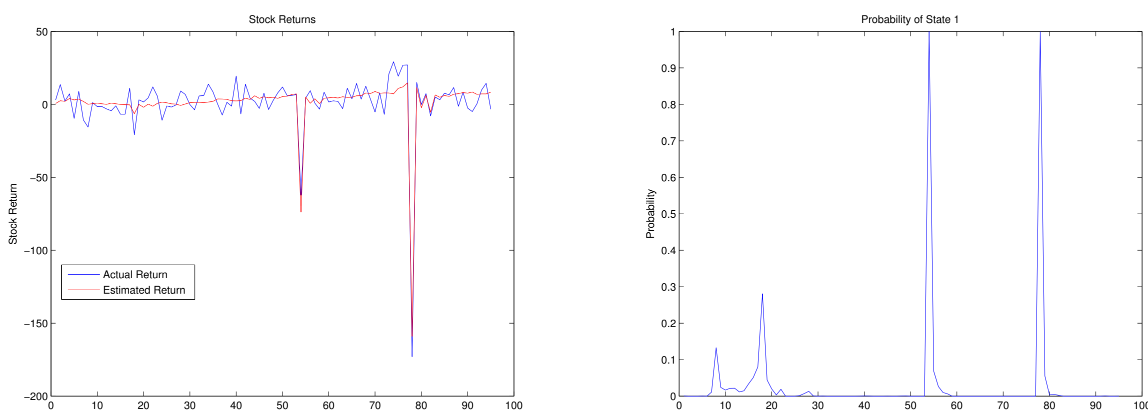

Figure 4: Median Bubble (MS-TVP Model)

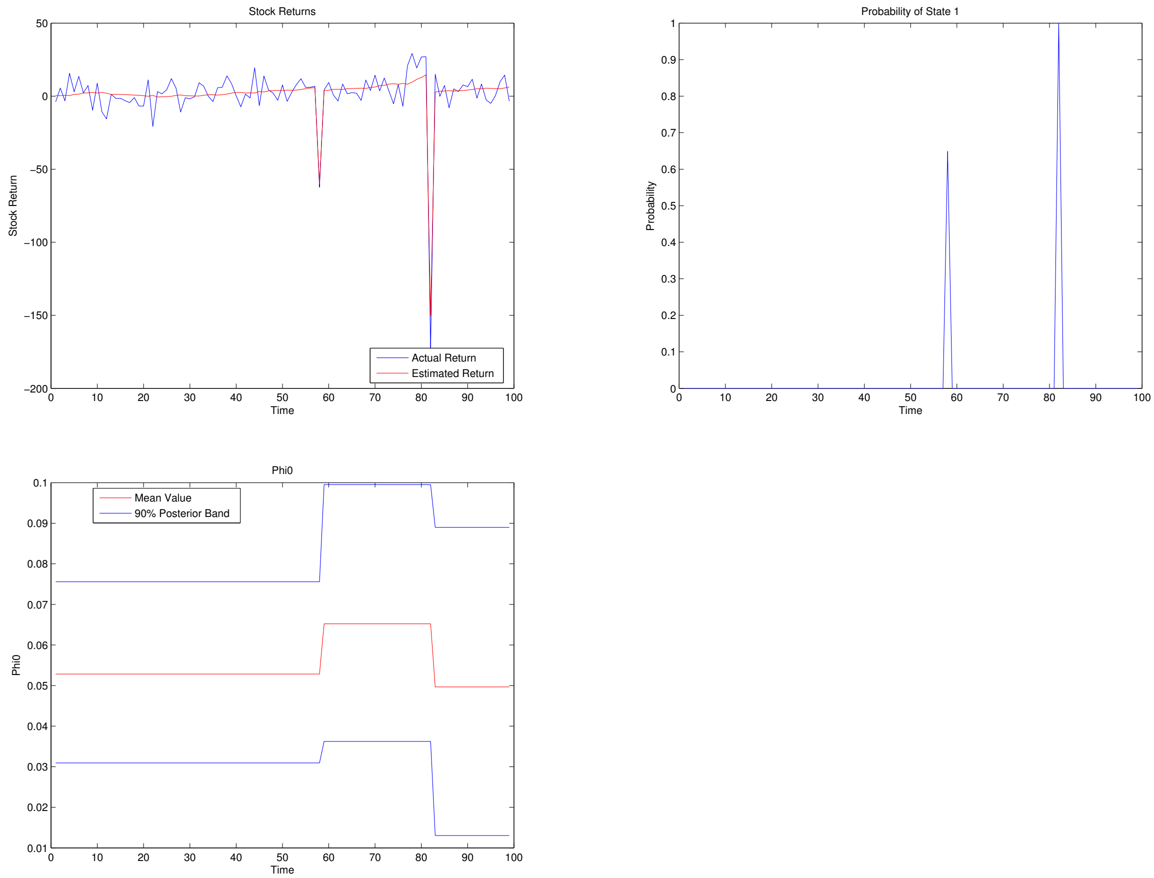

7 Conclusion

I have introduced Bayesian estimation of Hall et al.'s (1999) Markov switching test for asset bubbles, and introduced a new Markov switching time-varying parameter (MS-TVP) model for testing for asset bubbles. This model combines features of Hall et al.'s (1999) test, and the MS-TVP model developed in Eo and Kim (2012). This test, which generalizes the test found in Hall et al. (1999), provides roughly the same power to detect bubbles of the form introduced in Evans (1991). However, the test may perform better in other environments, specifically in a model where the growth rates of bubbles structurally vary over time, as opposed to differing only based on a particular realization of exogenous shocks. This alternate specification may also better approximate reality, as there is no ex ante reason to believe that all bubbles should grow at nearly identical rates.

In future research, I first like to extend my model to a real world application, such as testing the entire past history of the S&P 500 for periods of bubbles. I would also like to compare the same models above on a different measure - nowcasting performance. Nowcasting can be thought of as forecasting the present state of the world for some object of interest that may only be observable with a lag. In this case, forecasting whether a particular asset is in a bubble today. Although the tests above have power to detect a bubble given a full sample of 100 periods of data, it may be the case that these bubble tests can only detect a bubble after it has collapsed. However, nowcasting of bubbles is of particular interest to policymakers, as the debate in 2005 surrounding the existence of the housing bubble in the Federal Reserve Board of Governors (2005) transcripts indicate.

8 Appendix

8.1 Bayesian Estimation of Hall et al.'s (1999) MS Model

Table 4: Priors for Bayesian MS Model

Parameter	Distribution	Mean	SD
μ_0	Normal	1.5	100
μ_1	Normal	1.5	100
ϕ_0	Normal	0.0	0.05
ϕ_1	Normal	0.0	0.05
ψ_j	Normal	0.0	0.50
h	Gamma	2.0	$\sqrt{8}$
p_{00}	Beta	0.5	0.289
p_{11}	Beta	0.5	0.289

8.2 Price Series Generation

I use the exact same procedure as Evans (1991) and Hall et al. (1999).

Equations:

$$P_t = (1+r)^{-1} E_t(P_{t+1} + d_{t+1})$$

$$F_t = \sum_{j=1}^{\infty} (1+r)^{-j} E_t d_{t+j}$$

$$B_t = (1+r)^{(}-1) E_t B_{t+1}$$

$$P_t = F_t + B_t$$

Where P_t is the stock price at time t, d_t is the value of the dividend at time t, F_t is the fundamental value at time t, B_t is the value of the bubble at time t, and r is the net interest rate.

The dividend follows a random walk with drift:

$$d_t = \mu + d_{t-1} + \varepsilon_t$$
$$\varepsilon_t \sim N(0, \sigma^2)$$

The bubble component, B_t follows the following path:

$$B_{t+1} = \begin{cases} (1+r)B_t u_{t+1} & \text{if } B_t \leq \alpha \\ \left[\delta + \frac{1+r}{\pi}\left(B_t - \frac{\delta}{1+r}\right)\xi_{t+1}\right] u_{t+1} & B_t > \alpha \end{cases}$$

Where δ and α are scalars that satisfy $0 < \delta < (1+r)\alpha$.

$$u_t = exp(z_t - \frac{\sigma_z^2}{2})$$
$$z_t \sim N(0, \sigma_z^2)$$

u_t is a sequence of i.i.d. random variables with $E_t u_{t+1} = 1$, and ξ_t is an exogenous i.i.d. Bernoulli process such that:

$$Pr(\xi_t = 0) = 1 - \pi$$
$$Pr(\xi_t = 1) = \pi$$

i.e. the bubble process follows a linear switching process. Note that if $\xi = 0$, then the bubble collapses to δu_{t+1}. Therefore, $1 - \pi$ is the probability of the bubble collapsing each period. The bubble will continue to grow at the slower rate until it exceeds the exogenously given scalar α.

I use the following parameters:

Table 5: Parameters

Parameter	Value
α	1
δ	0.5
σ_z^2	0.0025
σ_ε^2	0.1574
μ	0.0373
r	0.05
d_0	1.3
B_0	0.5
π	0.85
n	100

Finally, in order to make the variance of the bubble contribute 75% of the variance of the price series, I scale the size of the bubble by κB_t each period, where $\kappa = 20$.

References

Berger, J. O. and Yang, R. Y. (1994). Non informative priors and Bayesian testing for the AR(1) model. *Econometric Theory*, 10:461–482.

Beveridge, S. and Nelson, C. R. (1981). A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the Business Cycle. *Journal of Monetary Economics*, 7:151–174.

Blanchard, O. and Watson, M. (1982). Bubbles, rational expectations, and financial markets. In Wachter, P., editor, *Crises in the Economic and Financial Structure*, pages 295–315. Lexington Books, Lexington, MA.

Carter, C. K. and Kohn, R. (1994). On gibbs sampling for state space models. *Biometrika*, 81(3):541–553.

Chib, S. (1995). Marginal Likelihood from the Gibbs Output. *Journal of the American Statistical Association*, 90(432):1313–1321.

Cochrane, J. H. (1992). Explaining the Variance of Price-Dividend Ratios. *Review of Financial Studies*, 5(2):243–80.

Diba, B. T. and Grossman, H. I. (1988). Explosive Rational Bubbles in Stock Prices? *American Economic Review*, 78(3):520–30.

Eo, Y. and Kim, C.-J. (2012). Markov-Switching Models with Evolving Regime-Specific Parameters: Are Post-War Booms or Recessions All Alike? Working Papers 2012-04, University of Sydney, School of Economics.

Evans, G. W. (1991). Pitfalls in Testing for Explosive Bubbles in Asset Prices. *American Economic Review*, 81(4):922–30.

Federal Reserve Board of Governors (2005). *Meeting of the Federal Open Market Committee on June 29-30, 2005.*

Fernández-Villaverde, J. and Rubio-Ramírez, J. F. (2004). Sequential Monte Carlo Filtering: An Example. Note, University Pennsylvania.

Flood, R., Hodrick, R., and Kaplan, P. (1994). An evaluation of recent evidence on stock price bubbles. In Flood, R. and Garber, P., editors, *Speculative Bubbles, Speculative Attacks, and Policy Switching*, pages 105–133. MIT Press, Cambridge, MA.

Gürkaynak, R. S. (2008). Econometric Tests Of Asset Price Bubbles: Taking Stock . *Journal of Economic Surveys*, 22(1):166–186.

Hall, S. G., Psaradakis, Z., and Sola, M. (1999). Detecting Periodically Collapsing Bubbles: A Markov-Switching Unit Root Test. *Journal of Applied Econometrics*, 14(2):143–54.

Hamilton, J. D. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57(2):357–84.

Kass, R. E. and Raftery, A. E. (1995). Bayes Factors. *Journal of the American Statistical Association*, 90(430):773–795.

Kim, C.-J. and Nelson, C. R. (1999). *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, volume 1 of *MIT Press Books*. The MIT Press.

Kindleberger, C. P. (2000). *Manias, Panics, and Crashes*. John Wiley & Sons, 4 edition.

Koop, G. and Potter, S. M. (2007). Estimation and Forecasting in Models with Multiple Breaks. *Review of Economic Studies*, 74(3):763–789.

Koop, G. and Potter, S. M. (2011). Time varying VARs with inequality restrictions. *Journal of Economic Dynamics and Control*, 35(7):1126–1138.

LeRoy, S. F. and Porter, R. D. (1981). The Present-Value Relation: Tests Based on Implied Variance Bounds. *Econometrica*, 49(3):555–74.

Lubrano, M. (1995). Testing for unit roots in a Bayesian framework. *Journal of Econometrics*, 69(1):81–109.

Pesaran, M. H., Pettenuzzo, D., and Timmermann, A. (2006). Forecasting Time Series Subject to Multiple Structural Breaks. *Review of Economic Studies*, 73(4):1057–1084.

Phillips, P. C. and Magdalinos, T. (2007). Limit theory for moderate deviations from a unit root. *Journal of Econometrics*, 136(1):115–130.

Phillips, P. C. B., Wu, Y., and Yu, J. (2011). Explosive Behavior in the 1990s NASDAQ: When Did Exuberance Escalate Asset Values? *International Economic Review*, 52(1):201–226.

Shiller, R. J. (1981). Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? *American Economic Review*, 71(3):421–36.

Sims, C. A. (1988). Bayesian skepticism on unit root econometrics. *Journal of Economic Dynamics and Control*, 12(2-3):463–474.

Sims, C. A. and Uhlig, H. (1991). Understanding Unit Rooters: A Helicopter Tour. *Econometrica*, 59(6):1591–99.

Tirole, J. (1982). On the Possibility of Speculation under Rational Expectations. *Econometrica*, 50(5):1163–81.

Tirole, J. (1985). Asset Bubbles and Overlapping Generations. *Econometrica*, 53(6):1499–1528.

West, K. D. (1987). A Specification Test for Speculative Bubbles. *The Quarterly Journal of Economics*, 102(3):553–80.

Xia, C. and Griffiths, W. (2012). Bayesian unit root testing: The effect of choice of prior on test outcomes. In Terrell, D. and Millimet, D., editors, *Advances in Econometrics: 30th Anniversary Edition*, volume 30, pages 27–57. Emerald Group Publishing Limited, Bingley, UK.