Volatility Dynamics of Financial Bubbles

Alexander Anemogiannis
California Institute of Technology
1200 E. California Blvd
Pasadena, CA
aanemogi@caltech.edu

Nikolai Kovachki
California Institute of Technology
1200 E. California Blvd
Pasadena, CA
nkovachki@caltech.edu

ABSTRACT

A reasonable definition of bubbles in financial markets is a period of unsustainable, super-exponential growth followed by a sharp crash. They are responsible for some of the largest economic downturns, but their dynamics have yet to be understood. Of the existing research attempts at understanding and predicting the behavior of financial bubbles, the log-periodic power law model is a somewhat popular approach to predicting a bubble's crashing time. It argues that the log of the prices follows a power law drift decorated by log-periodic oscillations that increase in frequency as they approach the crash time. After implementing this model, we believed the oscillations about the power law drift were in fact a stochastic process which shouldn't be captured by a deterministic model. We instead propose that a bubble's growth is better explained by a power law drift in the log of the prices, decorated by a stochastic process with increasing variance.

1. MOTIVATION

Crashes in financial markets can have devastating impacts on the economies of nations and industries. Furthermore, the countless relations among nations arising from globalization can cause these crashes to cascade into global downturns. There are numerous examples of this phenomenon. The 1987 crash known as "Black Monday" affected countries all over the world. It began with a crash in Hong Kong's Hang Seng stock exchange on October 19th and made its way across the world to western Europe and the United States in less than a day. Just to give an idea of the widespread effects of this crash, stock markets in Hong Kong, New Zealand, Australia, the United Kingdom, and the United States had fallen by 45.5%, 60%, 41.8%, 26.5%, and 22.7%, respectively by the end of the month. Similar global effects arose from the United States housing market crash in 2008, which led into the global financial crisis of 2008-2009, having devastating effects that caused national recessions across the world. In the United States alone, the federal government had to spend over $700 billion to bailout failing banks.

The negative impacts resulting from the crashes of financial bubbles can be immense, with devastating impacts on the government and citizens of nations across the world. Despite these disastrous effects, not much work has gone into investigating financial bubbles. This lack of research suggests some sort acceptance of these events, that they're unpredictable and the best we can do is react to them. Several researchers didn't share this sentiment and began to develop a model that would predict a bubble's crash. Before we can reach the point where we can look at a forming bubble and infer the necessary actions needed to mitigate the harmful effects of its crash, we need to first understand how it works and what underlying processes are responsible for its behavior. The goal of our research is to be able to begin characterizing these underlying processes in order to provide this understanding.

2. BACKGROUND & PREVIOUS WORK

Before we can talk about bubble detection, prediction, or characterization, we need to have an understanding of how a bubble actually forms in a microeconomic perspective. In the simplest terms, a bubble is a period in time when the perceived value of some commodity is much higher than its intrinsic value. For a classic example, consider the price of tulips during the early seventeenth century in western Europe. A single tulip was worth an amount equivalent to $10,000 today simply because people believed its value would keep increasing. This sort of belief is generally brought on by self-reinforcement. People believe that the value of some commodity is high or will become high and are then able to convince others of the same thing who then, in turn, convince others and so on the process continues. But all bubbles pop. And what generally makes them pop is the herd mentality which makes them form. Someone gets the idea that the commodity is no longer worth its value, begins offloading, and other quickly follow. We are interested in trying to understand when and why a bubble has reached its critical mass. How large must the herd grow before it can be brought down? Since human behavior is very difficult to model accurately, we study bubbles from the perspective of a model that has already been applied in physical systems.

2.1 Critical Events

Detection and prediction of critical events has generally been studied in the context of natural phenomenons and thus much of the research and proposed models has been done by physicists. But before we discuss how some of these models can be applied to financial markets, we need to make the notion of a critical event a bit more clear. The easiest way to do this is, of course, to consider a physical system. Imagine a balloon being filled with air. If we were to observe the structural integrity of the system in time then everything would seem fine until the very last moment when we have put in just enough air that the balloon bursts. It is this burst that we call a critical event and we are interested in studying its identifiable causes. The amount of air needed to bring about the burst is dependent on numerous factors including the way the balloon was constructed, the materials used, and the speed at which the air flowed in. Simple experiments can be designed to study exactly what precursors lead to the balloon bursting, and we can look for any patterns that may be indicative of the upcoming explosion. It is the detection of these patterns which would hopefully stay invariant under certain changes in conditions that we really care about. There two types of information that we would want such patterns to tell us. The first is weather our system is a state where a critical event is likely to happen. If we were to let the balloon sit or slowly fill it with air, then we shouldn't expect a burst to happen anytime soon. In our balloon example, it is very easy to tell whether we should be expecting a critical event or not, but when we start analyzing financial bubbles then it's no longer obvious what the regime of the market is. The second type is that, given we are in a state where a critical event is likely, we would like an indication of how close we are to the critical point. Our work focuses on examining the model that tries to capture patterns that give predictive information, so we will generally assume that our data is in a regime likely of a critical event.

We want to perform a similar analysis when considering financial bubbles. Of course this becomes much harder, because we can no longer design experiments to help us collect relevant data nor do we know what data would be most helpful. All we have to work with are index prices, so we can only create models which attempt to capture patterns in this time series data. To get a better idea of what the data looks like and what kind of qualitative observations we can make about it, consider the following plot which shows the logarithm of the price of Hong Kong's Hang Seng index.

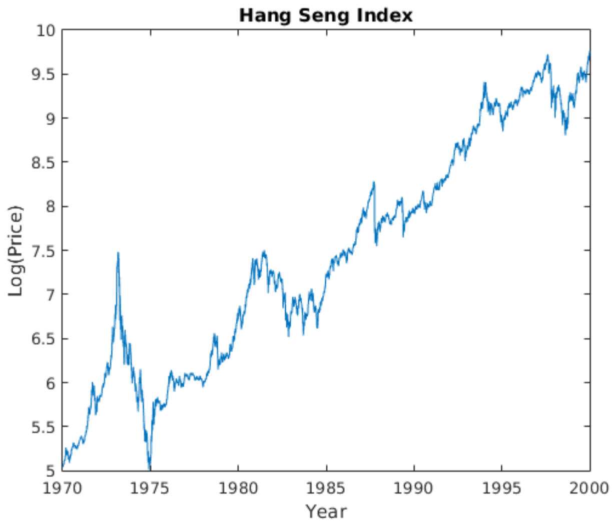

Figure 1: Logarithm of the prices of the Hang Seng Index from 1970 to 2000.

It is fairly obvious the overall growth of the logarithm of the prices is linear which is an observation consistent with the Black-Scholes-Merton model. But what if we considered smaller time scales? There are sections of the plot that exhibit faster than linear growth and are generally followed by a large crash. Such sections of the data are compatible with the idea of self-reinforcement growth which we used to characterize bubbles, and hence we say that a bubble is a period in time when the logarithm of the prices of a certain index grow faster than linearly and are followed by a steep decline which we call the critical or crash time. The most obvious next step is to say that the growth we are observing is instead polynomial. This may seem like a completely unjustified assumption, but let us consider critical systems in nature. In statistical mechanics, we tend to model multi-particle systems that are close to a critical time by a power law (for example the two dimensional Ising model). For further details see [7]. Experimentation has shown that such models are fairly accurate descriptors of real world events, so it is not entirely unreasonable to expect a large multi-agent system, like a financial market, to behave similarly. With this heuristic motivation in mind, we can write down a deterministic model for the prices $p(t)$ of some commodity in a bubble regime, mainly

$$\log(p(t)) = A + B(t_c - t)^\beta \tag{1}$$

where t_c is the time of the crash, A is the logarithm of the price at the crash and B, β are constants such that $B < 0$, $0 < \beta < 1$. If we already knew the critical time t_c then fitting this model to a data set is simply a matter of running a power-regression, but having to also infer t_c makes this a non-convex minimization problem which is not easily solved. We will discuss the problem of estimation further in Section 3.1. For now let us examine how well this model does in predicting t_c. We fit the model to well-known financial bubbles by taking the time series of prices for a certain index starting at the time we believe is the beginning of the bubble (the point where we first observe faster than linear growth in the logarithms) and ending around two weeks before the crash.

Index	Interval	Crash Date	Predicted Crash
S&P 500	1982-1987	10/19/1987	02/04/1988
NASDAQ	1997-2000	03/10/2000	06/17/2000
Hang Seng	1993-1997	06/04/1997	11/02/1997

Figure 2: Critical time predictions for the 1987 Crash, The Dot-Com Bubble, and the 1997 Asian Financial Crisis using the model given by Equation (1).

This model performs quite poorly in its ability to predict a crash, as is clear from the Figure above, but it does well in capturing the mean drift of the market as indicated by strong R^2 values. This result can be interpreted to mean that the crash time is not fully determined by the acceleration of the price. A qualitative observation of a single bubble reveals that there many small crashes within the bubble formation and their frequency increases the closer we get to the actual crash time. The next model we present, proposed by Johansen, Lediot, and Sornette [5], attempts to capture this observation.

2.2 Log-Periodic Power Law Model

The Log Periodic Power Law (LPPL) model has become the de facto model when attempting to predict bubble crash times. It has seen most use at the Financial Crisis Observatory at ETH - Zurich which is a research group attempting to show and exploit inefficiencies in the market during bubble regimes to be used for predictability. The model can be derived in two entirely different ways both of which capture different assumption we can about modern financial markets.

The first method exploits the susceptibility of a fractal structure known as the hierarchical diamond lattice. If we look back to the model given by Equation (1), its striking similarity to the Ising model tells us that, from the perspective of microscopic modeling, all investors are connected in a uniform way. This is due to the way the Ising model treats interactions between particles. But in a real market some investor can be more or less connected to others. Certainly someone working on Wall Street and making hundreds of trades every single day is much more influential to many investors than someone casually investing from home. The diamond hierarchical lattice is a way to capture the connectedness of an investor to other investors in the market. So let us consider an ideal market with no dividends and ignore risk aversion, interest rates, and liquidity constrains. We start with two investors that are linked in this market. When we say linked, practically we mean that the decisions of one of the investors can influence the decisions of the other. Mathematically, we mean an undirected graph with two vertices and an edge between them. Now we substitute this edge with 4 new edges, forming a rhombus and adding two new investors to the graph. We can take each edge and again substitute it for for new edges, adding two additional investors at each substitution and forming a diamond lattice. The following figure (taken from [1]) shows the the first few steps in the procedure of building our investor lattice.

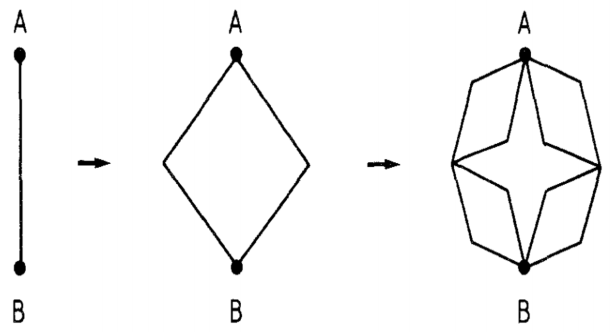

Figure 3: Recursive construction of the hierarchical diamond lattice.

A little bit of combinatorics tells that after n iterations our lattice contains $\frac{2}{3}(2+4^n)$ investors with 4^n links between them. This model was solved by Derrida, De Seze, and Itzykson in 1983 [1]. They arrived at the following solution for the susceptibility of the system

$$\chi = A_0(K_c - K)^{-\gamma} + A_1(K_c - K)^{-\gamma+i\omega} + ...$$

It is quite common in physics to deal with complex numbers, but the same cannot be said for economics, so we will take the real part of the above equation and then its first order approximation to obtain the hazard rate of the market,

$$h(t) \approx A_0(t_c - t)^{-\gamma} + A_1(t_c - t)^{\gamma}\cos(\omega\log(t_c - t) + \phi)$$

Plugging this hazard rate into simple model of rational expectation for a network of investors (for more details see [5]) we obtain

$$\log(p(t)) = A + B(t_c - t)^{\beta}[1 + C\cos(\omega\log(t_c - t) + \phi)] \quad (2)$$

with the constraints: $B < 0$, $0 < \beta < 1$, and $0 \leq \phi \leq 2\pi$. We call the model given by Equation (2), the Log Periodic Power Law. If we think about our observation that oscillations become more and more frequent the closer we get to the crash time, we notice Equation (2) captures exactly this behavior. The cosine term allows for fitting of oscillations and the logarithmic term drives them closer to each other as they approach the bubble's crashing point. We can also arrive at this model from the assumption that a bubble regime has discrete scale invariance because log-periodicity is an observable of systems with discrete scale invariance [3].

Mathematically, scale invariance means that for some observable function $\Phi(x)$ there is a function $\mu : \mathbb{R} \to \mathbb{R}$ such that for arbitrary λ,

$$\Phi(x) = \mu(\lambda)\Phi(\lambda x)$$

Since our invariance is discrete, we say the above equation is valid only for a countable set $\{\lambda_1, \lambda_2, ...\}$ where we can write $\lambda_n = \lambda^n$ with λ being our invariant scale ratio. With this in mind, the most general solution to the above equation is

$$\Phi(x) = x^{\alpha} P\left(\frac{\log(x)}{\log \lambda}\right)$$

where P is an arbitrary function of period 1. We can write down the Fourier expansion of $\Phi(x)$ to obtain

$$\Phi(x) = \sum_{n=-\infty}^{\infty} c_n e^{2\pi n i \frac{\log(x)}{\log(\lambda)}}$$

Hence $\Phi(x)$ is a sum of power laws with an infinitely discrete spectrum of complex exponents, so if we are to take the real part of the first order approximation, we would obtain the right hand side of Equation (2). Now that we know where our model comes from, we discuss methods of estimation and potential short-falls as our log-periodicity may be fitting to observable noise rather than an inherent property of bubbles.

3. MODEL DYNAMICS

3.1 Estimation of the LPPL Model

Estimating the parameters of the LPPL model is not an easy task as its cost function is non-convex and contains many local minima. We use a standard cost function which calculates the average of the sum of squared residuals. Let $\{y_t\}_{t=1}^{N}$ be a time series of the logarithms of the observed prices for some index and set $\theta = (A, B, C, t_c, \beta, \omega, \phi)$. Then

$$f(\theta, t) = A + B(t_c - t)^{\beta} + C(t_c - t)^{\beta}\cos(\omega\log(t_c - t) + \phi)$$
$$C(\theta) = \frac{1}{N}\sum_{t=1}^{N}(y_t - f(\theta, t))^2$$

Hence our problem becomes

$$\min_{\theta} C(\theta)$$

This a minimization problem over a seven dimensional parameter space. However notice that $f(\theta, t)$ is linear in the arguments A, B, C and hence we can reduce the problem to a minimization over a four dimensional parameter space with a standard OLS technique as follows. Define

$$k_t = (t_c - t)^{\beta}$$
$$g_t = (t_c - t)^{\beta}\cos(\omega\log(t_c - t) + \phi)$$

This allows us to write $f(\theta, t)$ as

$$f(\theta, t) = A + Bk_t + Cg_t$$

The above equation gives rise to a matrix formulation, mainly we define,

$$X = \begin{pmatrix} 1 & k_1 & g_1 \\ \vdots & \vdots & \vdots \\ 1 & k_N & g_N \end{pmatrix}$$

and $y = (y_1, ..., y_N)^T$ then

$$\begin{pmatrix} A \\ B \\ C \end{pmatrix} = (X^T X)^{-1} X^T y$$

This leaves us with a non-convex optimization problem over a four dimensional parameter space. There are many non-linear programming methods that can yield solutions. The one most commonly used in the literature is known as Taboo search, but, in practice, any alternative grid search method works just as well. First we strengthen our constraints, so the oscillations are not too slow or too fast as suggested by Lin [6], as follows

$$0.1 \le \beta \le 0.9$$
$$6 \le \omega \le 15$$
$$0 \le \phi \le 2\pi$$
$$N \le t_c < \infty$$

These give us a relatively small grid for the solution space of (β, ω). So let $S_1, ..., S_n \subset [0.1, 0.9] \times [6, 15]$ be disjoint squares of equal side length such that $\cup_{i=1}^{n} S_i = [0.1, 0.9] \times [6, 15]$. Then pick candidate solutions $x_1 \in S_1, ..., x_n \in S_n$. Now we use each $x_1, ..., x_n$ as a starting value for the Levenberg-Marquardt nonlinear least squares algorithm to obtain solutions $(\phi, t_c)_1, ..., (\phi, t_c)_n$. We again use these values as starting points for Levenberg-Marquardt to re-calculate $(\beta, \omega)_1, ..., (\beta, \omega)_n$. We can continue to do this until we are confident that each solution converges to a minima and then we take the solution which gives the smallest sum of squared residual as our final solution. We found no significant improvement when picking $n > 25$, and the algorithm generally converges for around 2 to 3 epochs.

Jacobsson [4] has suggested an alternative method of estimation via a genetic algorithm, which starts with a population of random solutions in the four dimensional space and breeds and mutates new populations until a minimum is found. More recently, Fantazzini [2] has shown that it is easier to fit the LPPL model to "anti-bubbles" (where we simply reverse our time series to be $\{y_t\}_{t=N}^{1}$) and proceed to use a 3-step MLE approach via the Broyden-Fletcher-Goldfarb-Shanno algorithm with a quadratic step length method. In practice, we found that all three algorithm perform about as well, but the grid search takes less time to converge as compared to the genetic and MLE approaches. The following figure shows the LPPL model fit to the S&P 500 index before the crash of October 1987.

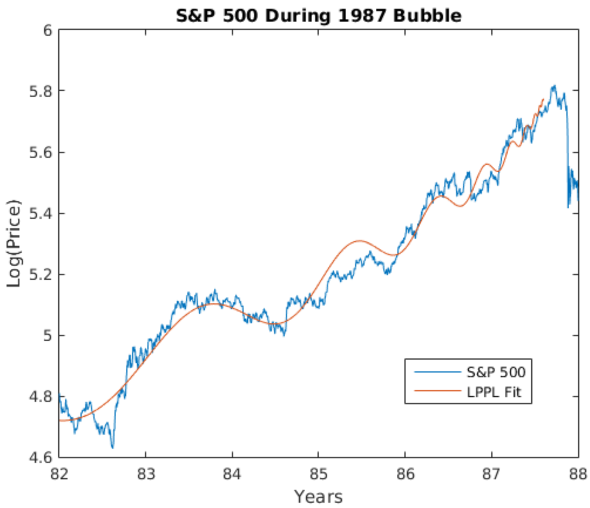

Figure 4: The LPPL model fitted to the S&P500 index from 1982 to 1987. The model parameters are $t_c = 87.8046$, $\omega = 10.8986$, $\beta = 0.5886$, and $\phi = 3.7631$.

The above fit predicts that the crash would happen on October 21th, 1987 which is only two days away from the actual crash which occurred on October 19th, 1987. We have tested this model across numerous other bubbles in various global financial markets, and it tends to perform fairly well, being, on average, around twenty days off of the actual crash.

3.2 Residual Analysis

After we implemented the LPPL model, we began to consider what ways we could improve it. The idea we ended up pursuing was that the oscillations about the power law drift belong to a stochastic process. This completely shifted our implementation of LPPL since a stochastic process shouldn't be captured by a deterministic model. Our claim is that when looking at the logarithm of a stock's prices, a bubble's behavior can be explained by a deterministic power law drift component decorated by a stochastic process which has increasing volatility as it approaches the time of the bubble's crash.

3.2.1 Sample Variance across Bubbles

Our first approach was as follows. We found 39 bubbles in the composite indices of stock exchanges across the world. When isolating bubbles in the time series we examined, we chose intervals which best matched our definition for bubbles: they exhibited super-exponential growth by having convex growth in the log of their prices, and a sharp decline followed this growth. We then normalized the prices for each bubble to be between 0 and 1 (by dividing all the data points by the maximum value in the set), so we can group data points across different bubbles. We then fitted each bubble to a power law and saved the residuals from this fit to give us the stochastic process about the power law drift. The time series of residuals were then aligned at 0 (as though we're aligning the bubbles at their origin) and we made sample populations consisting of points defined by their distance from the origin (see Figure 5). For a given distance, we calculated the sample variance for a population consisting of points (that were the given distance away from their origin) across all sets of residuals. This approach supported our claim. Figure 5 indicates an increase in the variance of the stochastic process underlying the residuals as the distance from the bubble's origin increases.

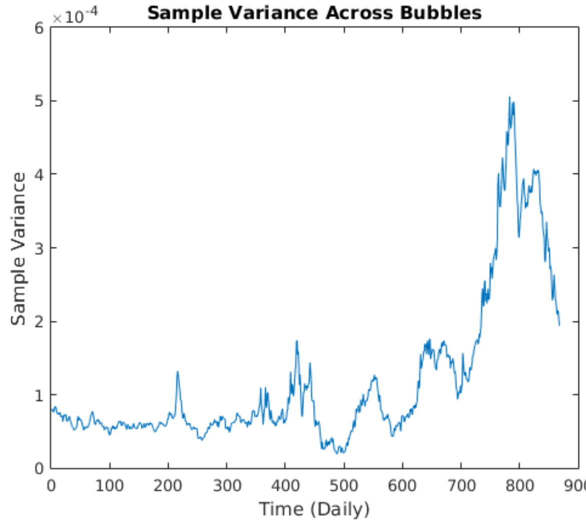

Figure 5: Sample variances across the residuals of 39 bubbles.

These results, however, should be taken with a grain of salt due to complications with identifying the origin of a bubble. When isolating bubbles in stock indices, we could easily identify their apparent crashes, but their origins were much more difficult to select. The goal is to capture the first point at which faster than linear growth occurs. But daily dynamics of stock prices are very volatile so it makes picking an exact origin almost impossible. One could ask why we did not align the bubbles at their crash given that's a much easier point to identify. Such an approach would, defeat the purpose because we would not expect small bubbles to be as volatile at their crash as large bubbles.

3.2.2 Dynamic Linear Models

In order to avoid complications from the noteworthy error in choosing the origin of a bubble, we decided to take another approach using Dynamic Linear Models (DLMs).

DLM Overview.

DLMs are a Bayesian approach to modeling processes given by a time series. They create a linear environment with Gaussian noise which makes it simple to estimate variables using Bayesian conditioning. A DLM consists of an underlying process that evolves linearly, with each evolution (or step through discrete time) introducing Gaussian noise into the process. We can make observations of the underlying at any time, but each observation will introduce another source of Gaussian noise. The general structure of a DLM is given below:

$$\begin{aligned} y_t &= \mu_t + v_t, && v_t \sim N(0, V_t) \\ \mu_t &= \phi\mu_{t-1} + w_t, && w_t \sim N(0, W_t) \end{aligned}$$

where x_t are the hidden states, y_t are the observations, w_t and v_t are Gaussian noise in state evolution and observation, respectively [8].

DLMs are normally used to predict the values of the hidden states over the observed time series. After first selecting the structure (observation and hidden state equations) and prior distributions for the DLM, posterior distributions for the DLM variables can be estimated using Kalman filtering. We took the following steps for the Kalman filtering part of our DLM implementation.

1. Carefully choose the structure (observation and state space equations) to provide an appropriate setting for modeling the desired process. Note that any part of the process not captured in μ_t will appear in v_t.

2. Assume a reasonable prior for the variables in the DLM structure: y_{t-1}, μ_{t-1}

3. Update the observation and hidden state variables using their past values and the update equations from the DLM structure, giving $\mu_t|D_{t-1}, y_t|D_{t-1}$

4. After observing y_t, replace $y_t|D_{t-1}$ with the observed y_t and now use the updated information to make a posterior estimate for the hidden state, $\mu_t|D_t$.

This process gives the following iterative updating equations that allow us to Kalman filter forward from $t = 0$ to $t = T$, giving us the hidden state values at each step. The predefined linear relationships are used to forecast future values, observations are made and compared with their predicted value, and then the error in this prediction weights the update of the hidden process to give a more accurate prediction. This is done in the update equations given below:

$$\begin{aligned} &\mu_t|D_{t-1} \sim N(a_t, R_t) && a_t = \phi m_{t-1}, && R_t = \phi^2 C_t + W \\ &y_t|D_{t-1} \sim N(f_t, Q_t) && f_t = a_t, && A_t = R_t/Q_t \\ &\mu_t|D_t \sim N(m_t, C_t) && Q_t = R_t + V, && C_t = R_t - A_t^2 Q_t \\ &m_t = (1 - A_t)a_{t-1} + A_t e_t, && e_t = y_t - f_t \end{aligned}$$

After filtering forward, the DLM variables will have a posterior distribution at each time t conditional on the DLM values at $t-1$ and the observations up to time t (D_t). Even better estimates can be found using retrospective smoothing to update the DLM posterior distributions from forward filtering. Retrospective smoothing conditions the posterior distribution at time t on additional information: all future values of the hidden states $\{\mu_t\}_{t+1}^T$ and all information revealed through observations D_T. We took the following additional steps to get better estimates of the posteriors for the variables in our DLM structure:

1. Use Kalman filtering to find $p(\mu_T|D_T) \sim N(m_T, C_T)$.

2. Characterize posterior distribution for prior state variables using the equations from the DLM structure:
$p(\mu_t|\mu_{t+1}, ..., \mu_T, D_T) = p(\mu_t|\mu_{t+1}, D_t) \sim N(m_t^*, C_t^*)$
$m_t^* = m_t + B_t(\mu_{t+1} - a_{t+1}), \quad B_t = \phi C_t/R_{t+1}$
$C_t^* = C_t(1 - \phi B_t)$

3. Beginning with $\mu_{T-1}|\mu_T, D_T \sim N(m_T^*, C_T^*)$, iteratively update μ_t for $t = T-1, ..., 1$, and then sample from the updated posterior distribution to get the estimates for the hidden state process and variances v_t, w_t.

Fixed y_t, μ_t.
The first DLM structure we tried allowed v_t to vary between a low and high volatility state throughout the time series.

$$\begin{aligned} y_t &= \mu_t + v_t, && v_t \sim (1-\pi)N(0, V) + \pi N(0, \kappa V) \\ \mu_t &= \phi\mu_{t-1} + w_t, && w_t \sim N(0, W) \end{aligned} \tag{3}$$

where y_t is fixed to the price observations, μ_t is fixed to the power law drift, and v_t, w_t are independent, zero-mean Gaussian random variables representing the observation and state evolution errors, respectively. pi is the probability of the process being in a high-volatility state where the observational variance is scaled by $\kappa > 1$.

We began with the prior assumptions that y_t is the log of the stock prices, μ_t is the fitted power law to y_t, w_t is drawn from the sample variance of μ_t, and v_t is the residuals of $y_t - \mu_t$. We then did a 2000 iteration "burn-in" of the DLM to get rid of any bias from our prior assumptions and allow the DLM to reach a stationary state. This was done by updating the model variables on each iteration using forward filtering and backwards sampling (FFBS). We used FFBS to update the DLM variables for another 10,000 iterations and averaged over these values to generate the plot in Figure 6.

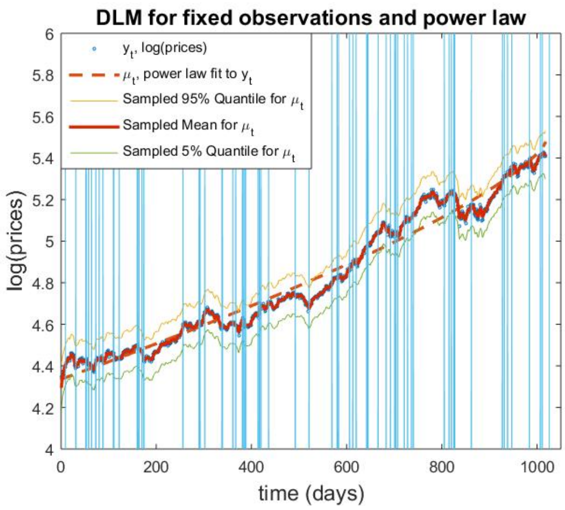

Figure 6: Sampled hidden states for DLM (3)

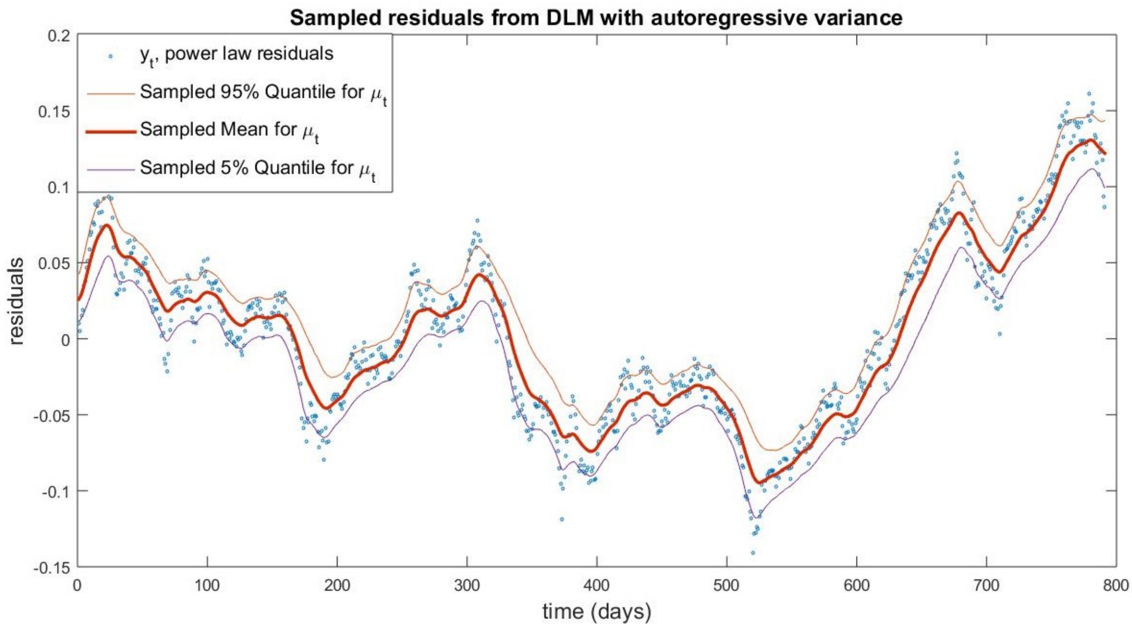

Figure 7: Sampled hidden states for DLM (4)

The vertical blue lines at time t indicate v_t was in a high volatility state. We hoped to be able to observe some sort of indication that the variance of the error in the observation process v_t is more likely to be in a high volatility state as you move away from the bubble's origin, but we failed to make any such connection and restructured the DLM to attempt to provide a better setting to capture increasing volatility in v_t.

Autoregressive v_t.
The second DLM structure we attempted was

$$\begin{aligned} y_t &= \mu_t + v_t, & v_t &\sim N(0V_t) \\ & & V_t &= \alpha_0 + \alpha_1 V_{t-1} + \eta_t \\ & & \eta &\sim \text{iid}N(0, \sigma_n) \\ \mu_t &= \phi\mu_{t-1} + w_t, & w_t &\sim N(0, W) \end{aligned} \tag{4}$$

where the observations y_t are now the residuals of log(prices) and the power law fit to the log(prices). Our hope was to smooth out μ_t by capturing the volatility of the residuals in v_t, which has an autoregressive variance.

We used the same update equations as listed in the overview, only we added another step to update the autoregressive variance V_t after sampling$\{\mu_t\}_1^T$. The process to sample v_t from an autoregressive V_t is shown below:

1. $\hat{v}_t = y_t - \mu_t$
2. $\hat{\alpha} = \mathbf{YX'}(\mathbf{XX'})^{-1}$
3. $\alpha = \hat{\alpha} + \sigma$
4. $V_t = \alpha_1 V_{t-1} + \alpha_0$
5. $v_t \sim N(0, V_t)$

where $Y = \{\hat{v}_t\}_1^{T-1}$, $X = \{\hat{v}_t\}_2^T$, $\sigma = \Gamma^{-1}\left(\frac{T-1}{2}, \frac{ee'}{\frac{T-1}{2}}\right)$, and

$e = \mathbf{Y} - \hat{\alpha}\mathbf{X}$

The results of this DLM are shown in Figures (7) and (8). Figure (7) shows the smoothed hidden state process for the residuals, and Figure(8) shows that the variance in the residual process increases with time from the bubble origin. These results support our claim that the stochastic process about the power law drift has increasing variance as the distance from the bubble origin increases.

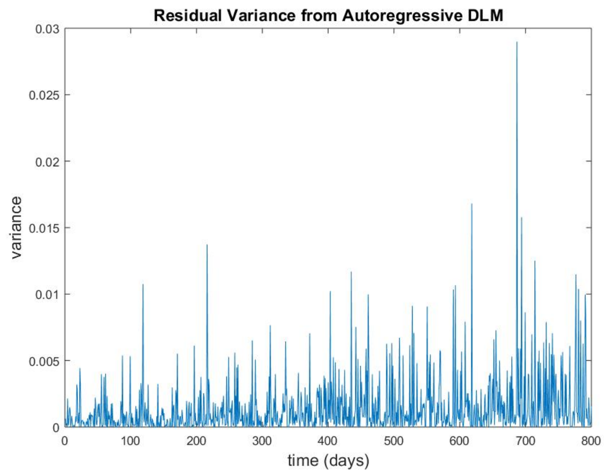

Figure 8: Sampled hidden states for DLM (4)

3.3 Stochastic Models

We can propose several models for the stochastic process governing a bubble regime. It is widely accepted the returns of a commodity, outside of a bubble, follow a Geometric

Brownian Motion. More precisely, if we let S_t be the price of some stock at time t then S_t evolves in time according to the stochastic differential equation,

$$dS_t = \mu S_t dt + \sigma S_t dW_t \tag{5}$$

where μ and σ are constants representing the mean drift and mean volatility respectively, while W_t is the standard Wiener Process. The above equation implies that the logarithmic returns of a stock are linear in μ and quadratic in σ, mainly for any $t > 0$, we have

$$\log(S_t) - \log(S_0) = \left(\mu - \frac{\sigma^2}{2}\right) t + \sigma W_t$$

assuming $W_0 = 0$. This implies that we if we fix some starting price, we should observe linear growth in the logarithm of the prices following. If we think back to the Hang Seng composite index shown in Figure (1) then that is exactly what we see in the long run. However during each bubble regime, we observe faster than linear growth. To account for this, as well as the increasing frequency of oscillations prior to the crash time, we propose the following two stochastic models for S_t while the market is in a bubble regime,

$$dS_t = \beta B(t_c - t)^{\beta - 1} S_t dt + \sigma S_t W_t - \delta S_t N_t \tag{6}$$

$$dS_t = \frac{\partial f(\theta, t)}{\partial t} S_t dt + \sigma S_t W_t \tag{7}$$

where $f(\theta, t)$ was defined at the beginning of Section 3.1, δ is a constant, and N_t is a jump process. The model in Equation (6) says the mean drift in the market during a bubble is characterized by a power law, and we account for the increasing number of oscillations by including a jump process. The model in Equation (7) claims that drift of the market is a log periodic power law. If we subscribe to the idea that the LPPL model captures something inherent to the bubble regime, then Equation (7) gives us a full description of the underlying stochastic process, while if we believe that LPPL simply captures increasing noise then Equation (6) is a much more likely candidate for the underlying stochastic process. Based on our result about increasing variance in the residuals, we believe that Equation (6) accurately captures the stochastic process of bubble regimes. We are, however, yet to find a statistical method for comparing the two models apart from qualitative observation which is not only non-rigorous but also incredibly difficult as stochastic models are very sensitive to parameter changes. The following figure shows a few simulations of Equation (6) with the parameters tuned to a bubble in the SWE market.

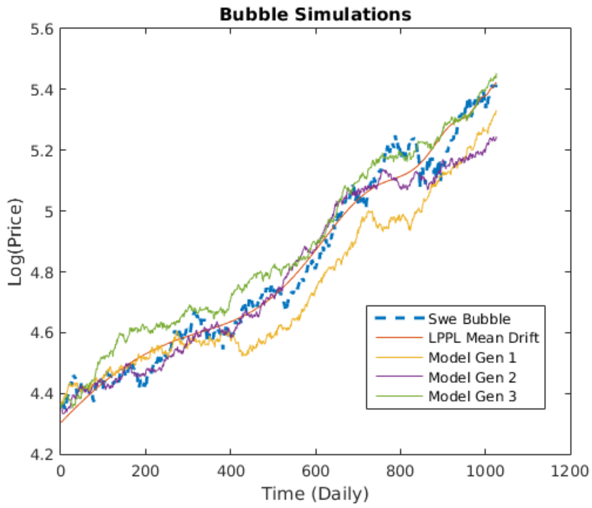

Figure 9: Simulations of Equation (6).

4. CONCLUSIONS

We began our examination of bubble dynamics using the log-periodic power law model; however, after we implemented it, we observed that the log-periodic oscillations it captures are most likely a stochastic process. Specifically, the log-periodic oscillations are a stochastic process with increasing variance which shouldn't be modeled with a deterministic fit.

We first investigated this claim by creating populations consisting of residuals (from the a power law fit) across 39 bubbles that are defined by their distance from the bubble's origin. We then looked at the sample variance of these populations against the distance each population was from the origin and observed that the variance increases with the distance from the bubbles' origins. There was a noticeable amount of error in selecting origins for the bubbles, so we decided to use dynamic linear models (DLMs) as a more reliable source of evidence.

When using a DLM with the structure specified in 4, we were also able to observe that the residual process's variance increases with the distance away from a bubble's origin. This approach allowed us to look at bubbles individually, so we no longer had to worry about error resulting from an inaccurate specification of a bubble's origin.

This evidence leads us to believe that a bubble is best modeled by a process similar to the one given by Equation (6). A way to experimentally verify such a claim is however unknown to us, so we have began looking at the microeconomic modeling theory for an explanation as to why random jumps occur.

5. FUTURE WORK

Improve Characterization of Underlying Process.
We were only able to show that the stochastic process deco-

rating the power law drift has an increasing variance. There are other characteristics about the process (eg. like the frequency of its oscillations about the mean drift) which we may be able to characterize. This would provide a better understanding of the process and improve its potential application.

Look for Predictive Correlations.
Once we understand the time evolution of a bubble's variance, we can try to look for certain indicators that would tell us how close we are to the crash. We would hope that all bubbles evolve in reasonably similar ways according to their volatility. This could lead us to some modifications to the LPPL that would improve its predictive strength.

Adapt Analysis to Use Incomplete Information.
So far all of our analysis is conditional on all of the data in a bubble, from its origin up until just before its crash. This is useful from a more academic perspective simply hoping to understand the underlying processes in a bubble. However, this contingency on complete information prevents a real-time application of our work since only incomplete information is available in practice.

6. ACKNOWLEDGMENTS

The authors would like to thank Professor Adam Wierman and Professor Benjamin Gillen for their guidance and assistance throughout this project.

7. REFERENCES

[1] B. Derrida, L. D. Seze, , and C. Itzykson. Fractal structure of zeros in hierarchical models. *Journal of Statistical Physics*, 33(3):559–569, 1983.

[2] D. Fantazzini. Modelling bubbles and anti-bubbles in bear markets: A medium-term trading analysis. *THE HANDBOOK OF TRADING, McGraw-Hill Finance and Investing*, pages 365–388, 2010.

[3] P. Geraskin and D. Fantazzini. Everything you always wanted to know about log periodic power laws for bubble modelling but were afraid to ask. *European Journal of Finance*, January 2011.

[4] E. Jacobsson. How to predict crashes in financial markets with the log-periodic power law. *ACM Trans. Program. Lang. Syst.*, 15(5):795–825, November 1993.

[5] A. Johansen, O. Ledoit, and D. Sornette. Crashes as critical points. *International Journal of Theoretical and Applied Finance*, 3(1):219–255, January 2000.

[6] L. Lin, R. R.E, and D. Sornette. A consistent model of 'explosive' financial bubbles with mean-reversing residuals. *Department of Mathematical Statistics, Stockholm University*, 2009.

[7] J. Sethna. *Statistical Mechanics: Entropy, Order Parameters and Complexity.* OUP Oxford, Oxford, United Kingdom, 2006.

[8] M. West and J. Harrison. *Bayesian Forecasting and Dynamic Models.* Springer-Verlag, New York, New York, 1997.