Popping the Bitcoin Bubble:
An application of log-periodic power law modeling to digital currency

Alec MacDonell
University of Notre Dame

April 1, 2014

Abstract

The year 2013 witnessed a remarkable increase in public interest and awareness of digital currencies such as Bitcoin. Hailed by some as the currency of the future, these "cryptocurrencies" have gained notoriety for their use in online black markets, but have yet to gain widespread acceptance. Given their novelty and lack of central regulating authorities, digital currencies experience high volatility and uncertainty regarding value. Taking Bitcoin as a representative example, this paper first uses autoregressive moving average (ARMA) functions to explain trading values, then applies log-periodic power law (LPPL) models in an attempt to predict crashes. The results of ARMA modeling show that Bitcoin values react to the CBOE Volatility Index, suggesting that a primary force currently driving Bitcoin values is speculation by investors looking outside traditional markets. In addition, the LPPL models accurately predict ex-ante the crash that occurred in December 2013, making LPPL models a potentially valuable tool for understanding bubble behavior in digital currencies.

*I would like to thank Steven Lugauer for advising me during the research and writing of this paper. His instruction and guidance has been invaluable.

Table of Contents

1 Introduction

Traditional instruction in economics teaches that currency takes one of two forms – commodity money or fiat money. Commodity money, typified by gold, relies on a rare, tangible substance capable of serving as both a store of value and a medium of exchange. Fiat money, (e.g. the U.S. dollar or the Euro) has no intrinsic value, but instead relies on the backing of a central authority that guarantees its worth. The last several years have witnessed the introduction of cryptocurrencies – digital currencies that rely on cryptography to ensure security – which cannot be accurately defined as either commodity or fiat money. They exist online, not in real space, and there is no authority guaranteeing their value. As such, their value is not explicitly tied to one nation's economic success, nor can it be directly influenced by a central authority's fiscal or monetary policy. Intangible, floating, and unregulated, cryptocurrencies seem to occupy a category all their own.

The first successful cryptocurrency was Bitcoin, which began trading in 2009. Since then, dozens of other cryptocurrencies have been introduced, many based on the Bitcoin model. Some are created for a specific purpose – Namecoin, for instance, attempts to create a new domain name system outside of conventional regulation. This would, in theory, make it much harder to control or suppress content on the Internet. Other cryptocurrencies are created as mediums of exchange for specific groups – Dogecoin is a currency used almost exclusively by members of Reddit, a popular social media site. In January 2014, Redditors used Dogecoin to raise $30,000 to help fund the Jamaican bobsled team in their bid for the Sochi Winter Olympics. For the purposes of this investigation, I will focus on Bitcoin as a proxy for digital currencies in general. Not only is Bitcoin the oldest cryptocurrency in

circulation, it is also by far the best known and most valuable. While some digital currencies trade into U.S. dollars at a rate of fractions of a cent per coin, Bitcoin values have risen at times over $1000 per coin, making the movements in the data far more significant.

This paper has two goals: first, to identify the factors that drive the value of Bitcoins, second, to test whether it is possible to successfully anticipate crashes in Bitcoin values. My hypothesis is that the value of the Bitcoin is primarily driven by two variables – the level of public interest over time and the level of investor confidence in traditional markets. The first variable reflects Bitcoin's newness as a currency. The relatively low number of currently available markets limits the potential use of Bitcoins as an actual currency used to buy and sell goods and services. However, as media attention increases, the probability of widespread adoption should increase correspondingly, endowing the Bitcoin with greater potential as a medium of exchange. The second variable looks at Bitcoins as a vehicle for investment. As economic uncertainty increases and traditional investment avenues become less stable, investors may look to shelter their wealth in other places, as seen in the demand for gold during the recent recession. Bitcoins, given their decentralized structure, could feasibly remain insulated from economic downturns and provide a potential store of value.

To test this hypothesis, I used maximum likelihood estimation (MLE) to estimate the parameters of the Bitcoin value model. Variables were added to the model based on their statistical significance and their affect on the Schwarz Bayesian Criterion (SBC). The only exogenous variable that came into the final model was the Chicago Board of Exchange

Volatility Index (VIX). The effect, however, was the opposite of the one stated in my hypothesis – Bitcoin values rise as the VIX falls. The effects of the VIX on Bitcoin values imply that when volatility in traditional markets is low, investors seek out Bitcoin precisely because its high volatility provides opportunities for speculation. No variables representing Bitcoin as a real medium of exchange came into the final model, implying that values are driven almost entirely by investors, not consumers.

Given the extreme volatility of Bitcoin values, a forecast based on my model yields wide confidence intervals. To mitigate this high level of uncertainty, there exists a need to anticipate critical points where values are likely to crash. Several authors (Johansen et al. 2000, Lin et al. 2009, Geraskin & Fantazzini 2010) have had success applying log-periodic power law (LPPL) models to historical financial bubbles to predict the timing of crash events. Based on the idea that crashes are preceded by faster-than-exponential growth, these models filter for several criteria that distinguish this behavior. An application of LPPL models to Bitcoin data over the three last years demonstrated that the models were able to predict the crash that occurred in December 2014. This successful prediction opens up the possibility of using LPPL models to understand future behavior in Bitcoin and other digital currencies.

The rest of the paper will be laid out as follows: Section 2 sets up the historical and theoretical background necessary to understand Bitcoins as a currency. Section 3 presents a summary of the data and the steps taken to produce a stable model. Section 4 discusses the theories on financial bubbles, and Section 5 outlines the specifications of LPPL

modeling. Section 6 explains the application of LPPL models to the December 2013 Bitcoin crash, and Section 7 discusses the results and implications of both the ARMA and LPPL models.

2 Background

Bitcoins were launched in 2009 by an anonymous individual or group under the pseudonym Satoshi Nakamoto. Nakomoto created the original Bitcoin software and published a proof of concept focusing on the system's security. Bitcoin differed from previous digital currencies by eliminating the need for a trusted third party – transactions were verified by the combined power of all participating computers. Bitcoins could not be seized from other users without obtaining their private digital signatures, and a single user could not fraudulently spend Bitcoins without possessing computing power greater than that of all other honest users combined. The project continued as an open-source development, incorporating contributions and ideas from programmers all over the world. Due to the limited number of legal venues for spending Bitcoins, their use generally falls into one of two buckets: first, a medium for anonymously purchasing goods and services (often of questionable legality) over the Internet, and second, a vehicle for investment similar to gold or forex trading.

At the time they began trading on Mt. Gox, the largest online Bitcoin exchange, one Bitcoin was roughly equivalent to a U.S. nickel – 5 cents – in value. Over the next two years, this value rose slowly, breaking $1 in early 2012 and eventually reaching values between $10 and $20. Suddenly, in late March of 2013, Bitcoin trading values began a rapid escalation.

Over the course of one month, Bitcoin values increased sixfold from just under $40 to a high of $230. The spike may have been initiated by an economic crisis in Cyprus, which saw the small nation's credit rating plummet to junk status (Detrixhe 2013). With mounting uncertainty over the future of Cyprus' banks, investors, eager to find other havens for their assets, began to buy into Bitcoins, pushing trading volumes and values to historic highs (Cohan 2013). In the following months, Bitcoin values settled back down to the low $100s, though with greater variance than pre-spike.

On October 2, 2013, the U.S. government shut down the Silk Road, an online black market known as the "Amazon.com of illegal drugs" (NPR 2011). The FBI arrested the site's founder and seized over 26,000 Bitcoins, worth $3.6 million, generating widespread speculation on the future of the currency and the government's stance on its legality. Finally, in mid-November, the U.S. Department of Justice acknowledged Bitcoins as a "legal means of exchange" and expressed a generally positive opinion of the currency. Trading values climbed immediately, soon breaking $1000.

Bitcoin values hit an all-time high on December 4, 2013, closing at $1238. Over the next three days, values collapsed, dropping back down below $700. There was a brief recovery period, followed by further drops until December 18, when values bottomed out at $541. Since that time, values have fluctuated between $600 and $1000, and the amount of data since the crash is still too small to assume any trend.

Though gaining acceptance, Bitcoin is a nascent currency and its fate is undetermined. The underlying system is cryptographically secure, but there remain weak points that a malicious attacker might seek to exploit. If a hacker were to compromise the security of an individual's computer or one of the exchanges on which Bitcoins are traded, he or she might be able to steal Bitcoins from other users. As responses to this threat become more sophisticated, Bitcoin values will likely experience some stabilization.

3 Data and ARMA Model Construction

3.1 Data Summary and Stationarity

In order to maintain consistent timing, I chose to work with historical daily closing prices from Mt. Gox. As the largest online market for Bitcoin trading, Mt. Gox should best represent Bitcoin trading as a whole. The initial data set has 163 observations with a weekly frequency, beginning July 18, 2010 and ending August 25, 2013. Values range from

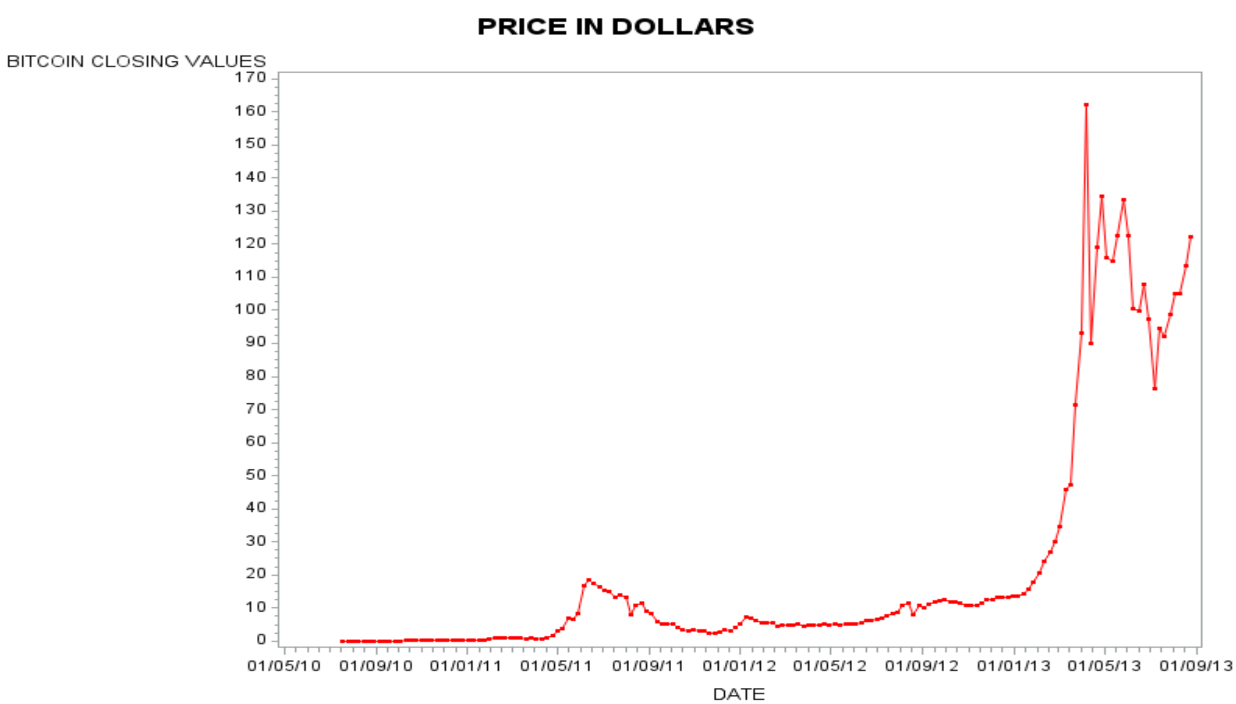

a minimum of \$.05 to a maximum of \$162.30, with a mean of \$21.53 and a standard deviation of \$36.83. From March 2013 onwards, there is a great deal of volatility in the data, as mentioned above, this is largely driven by increased speculation on the part of investors.

The Wold decomposition theorem states that any stationary discrete random process can be decomposed into two uncorrelated processes, one autoregressive (AR) and the other a moving average (MA). In order to achieve stationarity, the data must have homoscedastic errors and a constant unconditional mean.

Initially, the Augmented Dickey-Fuller (ADF) Test failed to reject the null of a unit root, unsurprising given the visible upward trend in the data. After first-differencing the data set, the ADF strongly rejected the null, indicating that the unit root had been successfully differenced out. I then performed a log-likelihood test, which returned a lower SBC for logged variables regardless of first-differencing. After a log transformation, the series was determined to be stationary and ready for ARMA modeling.

3.2 Univariate ARIMA Model and Forecast

The final univariate ARMA model is as follows:

Let $w_t = \Delta \ln(y_t)$.

$$w_t = 0.04386 + 0.33776\, w_{t-3} + 0.15052\, w_{t-30} + .23361\, \varepsilon_{t-40} + \varepsilon_t$$

In order to obtain this model, I first ran an ARIMA procedure in SAS and obtained an initial SBC of -59.9447. A repetition of the ARIMA procedure without an intercept resulted in a

higher SBC, so the intercept was retained in the model. Observations of the trend and correlation analysis showed likely correlation at a lag of 3 weeks, 10 weeks, 20 weeks, 30 weeks, and 40 weeks. Testing combinations of these lags as both AR and MA terms showed that a model with AR terms of lag 3 and 30 and an MA term of lag 40 provided the lowest SBC at -66.2772.

For the next twelve periods (y_{T+1} through y_{T+12}), the forecast shows a steady rise from \$128.54 at y_{T+1} to \$190.91 at y_{T+6} and \$275.23 at y_{T+12}. However, given the volatility of the data, the 90% confidence intervals are wide, making overall confidence in the forecast low. In y_{T+1} the 90% confidence interval has a low of \$95.01 and a high of \$173.89, a range of \$78.88. By y_{T+12}, the low is \$84.55 and the high is \$895.93, a range of \$811.38. While confidence in the point forecast is low, these intervals do suggest an upward trend going forward from y_T.

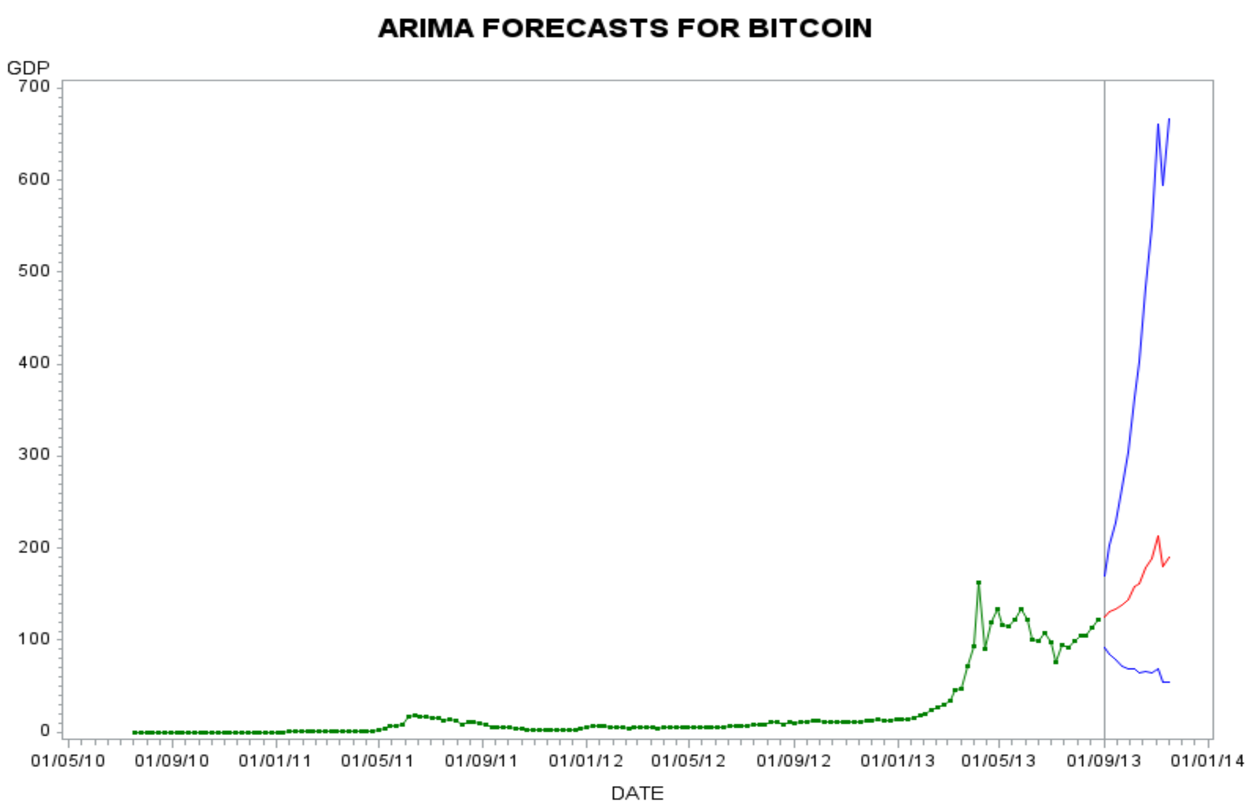

3.3 Multivariate Data

In order to generate tighter confidence intervals and explain more of the variance in Bitcoin values, I attempted to bring several exogenous variables into the model.

First, as a proxy for the level of public attention to Bitcoins, I took the number of news articles mentioning Bitcoins between July 16, 2010 and August 23, 2013 as recorded by Bloomberg Professional Service's news tracker. I chose to record weekly values on Friday for two reasons. First, weekend news mentions were dramatically fewer. Second, I did not want the data to reflect short term trading decisions based on news reports released that day, but rather to reflect trend shifts resulting from a general level of media attention at the time.

To measure investor confidence in the overall economy, I used the Chicago Board of Exchange Market Volatility Index (VIX), often referred to as the "fear index." I also brought in the value of gold bullion in U.S. dollars, hypothesizing that gold would have similar appeal to Bitcoins in the case of investors seeking to shield wealth from economic uncertainty. Both sets of data were taken from the Federal Reserve Economic Data (FRED) collected by the St. Louis Fed.

I included the value of the Euro relative to the U.S. dollar to account for the possibility that there might be a North American bias among Bitcoin investors. If, for instance, the majority of Bitcoin speculators lived in the United States, a decrease in the value of the dollar

relative to other currencies might provide an incentive to shelter wealth in other places, thereby driving value in Bitcoins.

Finally, I created a dummy variable to account for the "gold rush" behavior of the Bitcoin in March and April 2013. This variable takes a value of 0.0 for the entirety of the data set except for the weeks from January 20, 2013 (when the growth in values began to accelerate) to April 7, 2013 (when the values hit their local maximum).

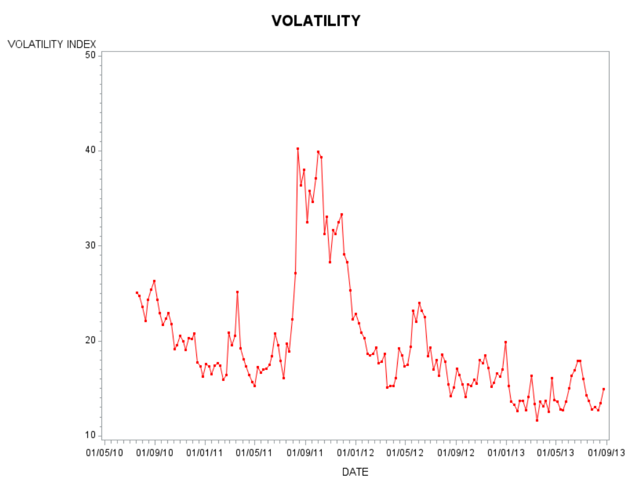

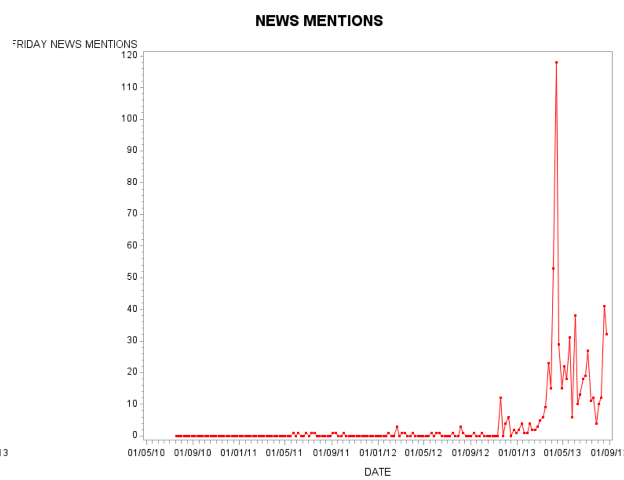

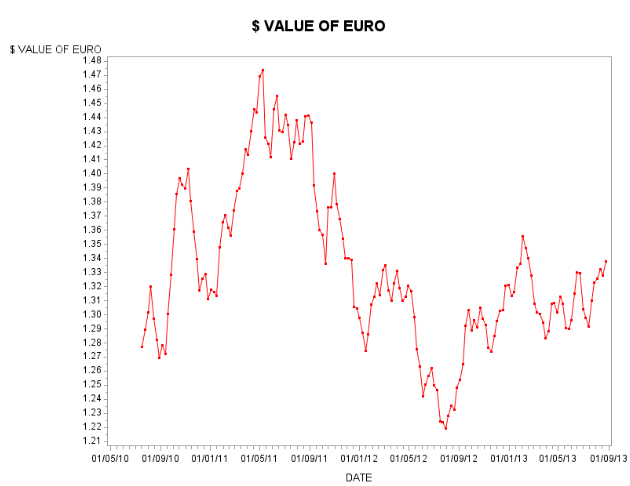

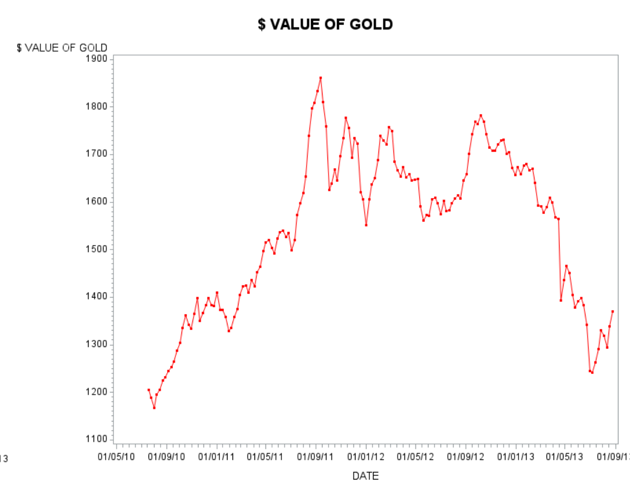

3.4 Multivariate ARIMA Model and Forecast

In order to generate a multivariate model for Bitcoin values, all exogenous variables were pre-whitened, meaning they were made stationary and given their own univariate ARMA model. For all variables except the dummy, both log transformations and first-differencing were necessary. Once stationarity was established, these variables were run against Bitcoin to determine cross-correlations.

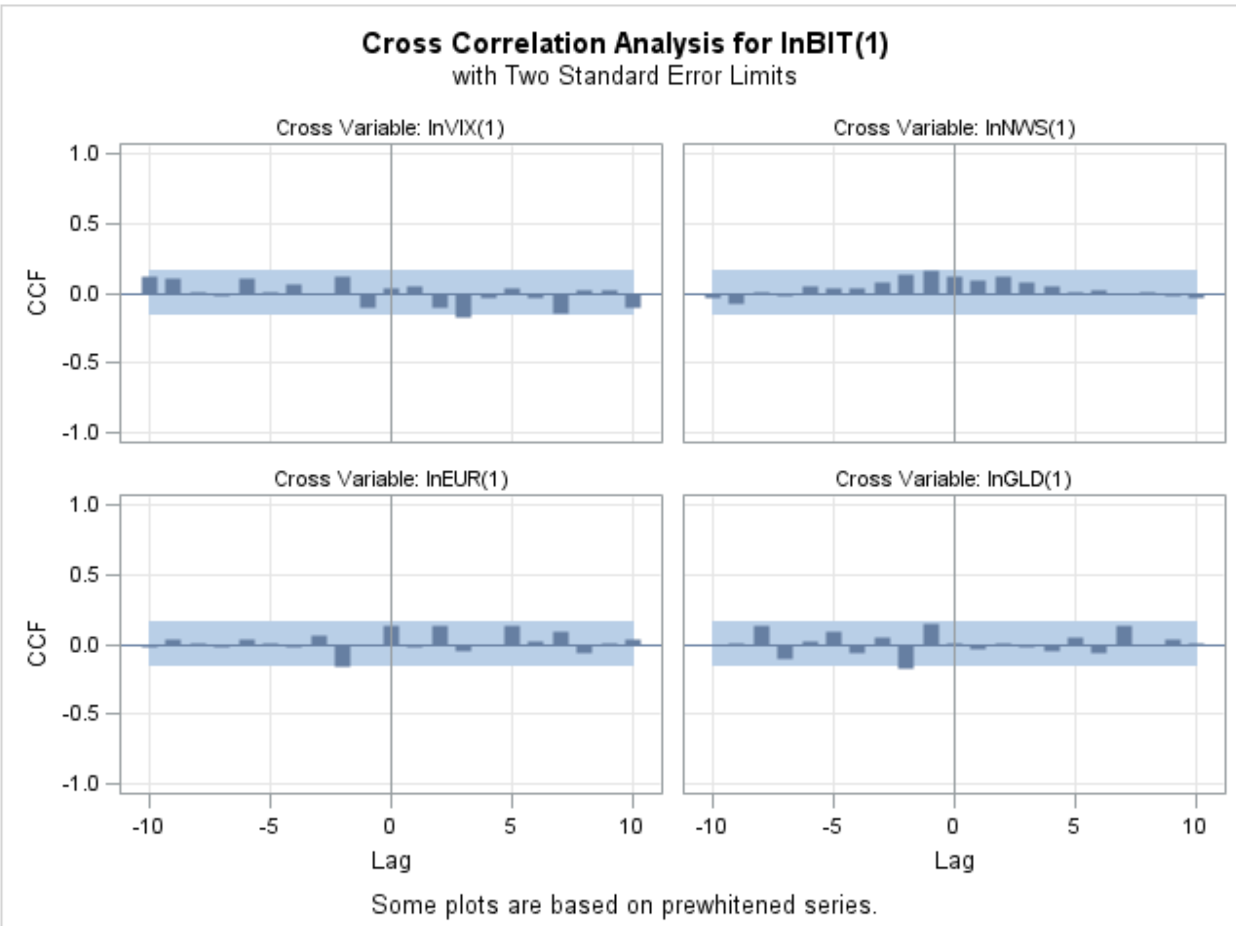

The plots shown above are a representation of the cross-correlation between Bitcoin values and the exogenous variables taken at different lags. The right side of each chart, labeled with positive numbers, refers to the lags of the exogenous variables, while the left

side marks negative lags – in other words, future values of the exogenous variables and their cross-correlation on the present value of Bitcoin. For the purposes of forecasting Bitcoin values, only the right side matters, though it is interesting to note that both gold and Euro values show correlation at a lag of -2, suggesting that Bitcoin values might be useful if one were seeking forecasts for those variables.

Out of these four variables, only the VIX shows any significant correlation, with a noticeable spike at a lag of 3 weeks.

After testing any likely lags of these variables in the ARIMA procedure, the VIX lagged 3 periods was, in fact, the only variable that lowered the SBC. I therefore kept the VIX in the model and rejected the other exogenous variables.

I next added the dummy variable to the model, which further lowered the SBC. However, the addition of this intervention reduced the statistical significance of the first autoregressive term from the univariate model, so it was removed from the model. The SBC reached -79.509, a clear improvement over the univariate value of -66.2772. The final model is as follows:

Let $w_t = \Delta\ln(y_t)$, $v_t = \Delta\ln(VIX_t)$, and d_t=intervention dummy

$$w_t = 0.0315 + 0.2443\, w_{t-30} - 0.4567\, v_{t-3} + 0.1561\, d_t + 0.3195\, \varepsilon_{t-40} + \varepsilon_t \qquad (1)$$

Parameter	Estimate	Standard Error	t Value
Mean	0.0315	0.0135	2.33
MA(t-40)	0.31945	0.09525	3.35
AR(t-30)	0.24427	0.08586	2.85
VIX(t-3)	-0.45673	0.12725	-3.59
Dummy	0.15611	0.051	3.06

The forecast again shows a steady rise, though this time the increases are smaller than in the univariate model– \$124.32 at y_{T+1} to \$133.28 at y_{T+6} and \$186.25 at y_{T+12}. The 90% confidence intervals are tighter than the univariate, but still very wide. At y_{T+1}, the low is \$93.48 and the high is \$165.36. By y_{T+12}, the low is \$68.03 and the high is \$509.95, a spread of \$441.92. This range is nearly 50% smaller than the forecast generated in the univariate model, but overall confidence in the point forecast is still low. Assuming an even probability distribution across the range of values within the confidence limits, a directional estimate would definitely suggest a rise over the twelve periods following y_T.

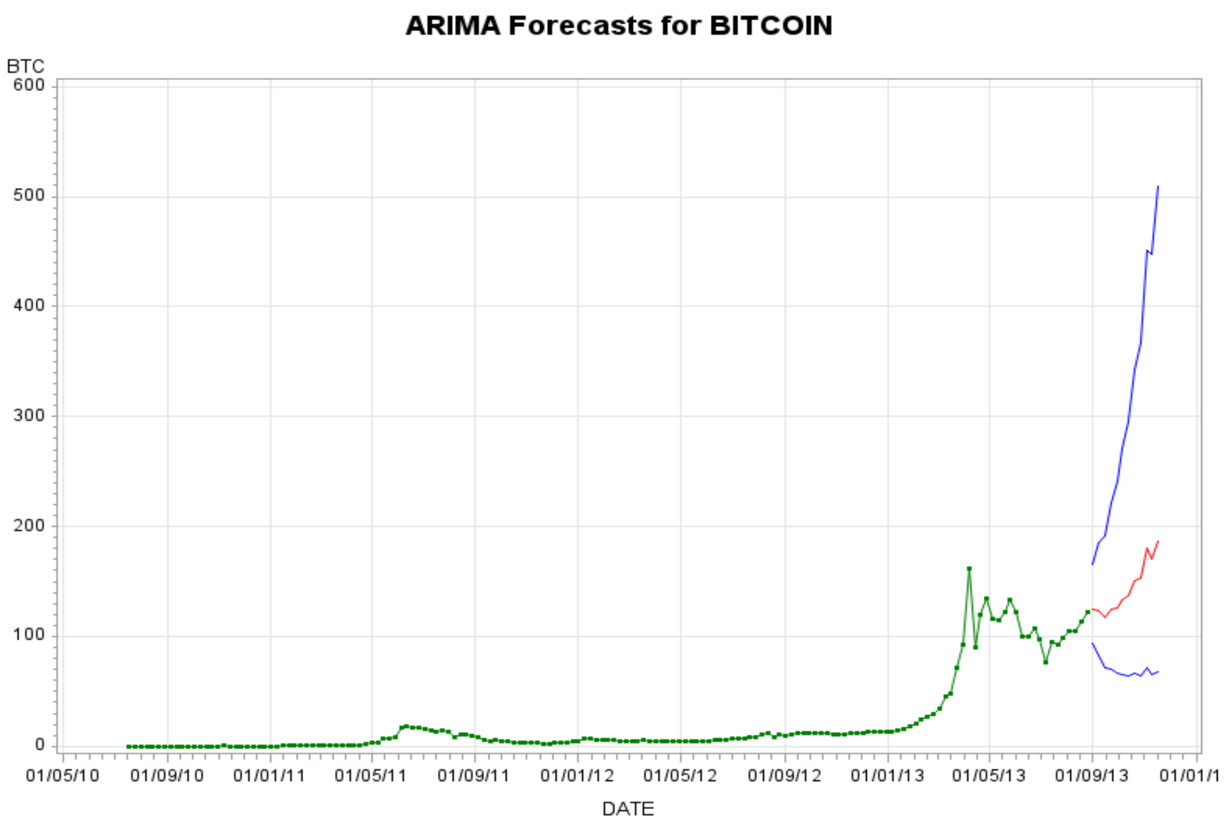

3.4 Holdout Sample

Because the original data sample ended on August 25, several months of new data have become available. This holdout sample will provide a good test for the accuracy of the forecast created above.

Parameter	Estimate	Standard Error	t Value
Mean	0.03934	0.01343	2.93
MA(t-40)	0.20864	0.08815	2.37
AR(t-30)	0.16109	0.08226	1.96
VIX(t-3)	-0.45059	0.13171	-3.42
Dummy	0.16692	0.05240	3.19

After updating the data to include 22 weeks of new observations, the ARIMA procedure produced the results shown above. The signs on all terms remained consistent. The mean, VIX, and intervention dummy terms experienced little change, while the AR and MA terms decreased slightly. The change in the ARMA terms is due to the dramatic growth of prices from October to December 2013, and would likely experience some correction if another intervention dummy were built in to account for the shift. Overall, the model is fairly stable and its terms robust to change.

The forecast from the updated model now has confidence intervals so extreme as to be nearly meaningless beyond a directional estimate that values will rise. If, then, one cannot predict with great certainty the exact value of Bitcoin at any given time, the best one can hope to do is to avoid crashes. To this end, I move now into a discussion of financial bubbles and apply a predictive model to the Bitcoin data set.

4 Financial Bubbles

Frequently referenced in media coverage of economic crises, financial bubbles are a defining characteristic of boom and bust cycles in markets. As generally understood, bubbles are periods of rapid expansion and subsequent collapse, usually accompanied by high trading volumes. Any tradable asset can experience bubble behavior – stocks, currencies, houses, even tulip bulbs. Theories on the root causes of bubbles are varied, with some economists denying that bubbles occur at all (Gürkaynak 2008). To better understand the concept, it will be helpful to briefly review these different views before specifying a model.

One popular theory suggests that excess monetary liquidity can lead to bubble formation (Meier 2009). When investors have a great deal of cash to invest in a limited pool of assets, they have greater incentive to bid the price of the underlying asset above its fundamental value. Low interest rates compound the issue by allowing investors to leverage their capital; this leverage both increases the appreciation of asset prices and raises the stakes in the event of collapse, as banks would also be negatively impacted by a downturn. However, even if the effect of liquidity is accepted, there still remains the question of why investors would behave in a manner that leads to catastrophic collapse, especially since many authors argue that investor behavior during bubble periods can still be considered rational (Garber 1990).

A. *Greater Fool Theory*

One posited explanation of the social phenomena behind bubbles is the so-called Greater Fool Theory, where investors knowingly purchase over-valued assets and rely on the probability that someone else will come along and pay an even higher price (Dreman 1993). As the story goes, this behavior continues until the asset is sold to a greatest fool who can find no new buyers, at which point the value plummets. Support for this theory might be found in the idea that individuals generally consider their own abilities to be greater than those of the population average, leading to the belief that there will always be another investor with less financial knowledge to serve as the greater fool.

B. *Bounded Rationality*

Others have suggested that bubbles are the result of behavior under conditions of bounded rationality. In this conception, individuals have limited information and time, leading them to make decisions that are satisfactory but not necessarily optimal. For instance, investors might understand the pricing model of an asset, but have trouble applying it in practice.

C. *Moral Hazard*

The moral hazard theory suggests that investor behavior may become irresponsible when insulated from risk. If, for instance, an investor knows that in the case of a crisis he will receive a bailout, he has perverse incentives to pursue high-risk investments while ignoring their risk-weighted value. This hypothetical investor's decisions are perfectly rational from an individualistic perspective. It could be argued that the recent U.S. housing bubble fell into the category of moral hazard bubble formation – knowing that safety nets

were in place, investment bankers bought up large amounts of risky debt and greatly increased their use of leverage.

D. Herding

Also referred to as institutionalization, the herding theory of bubble formation suggests that investors behave in imitation of the group trend, much as judges in the Keynesian Beauty Contest choose based on their expectations of others' behavior rather than their own opinions. This theory explains both the sustained growth of a bubble over a period of time and the dramatic collapse as the herd of investors begins to sell rather than buy.

Which theory provides the best explanation? Though opinion is divided on the subject, a strong case can be made for herding. Levine & Zajac (2007) found no experimental evidence to support the Greater Fool or bounded rationality theories. Bubbles still formed in experimental markets even when investors were able to accurately price the asset and no investor considered himself superior to his peers. The moral hazard theory, while compelling as a narrative, fails to explain bubbles that occur in the absence of mitigated risk, and therefore provides an incomplete conception of bubble behavior. The herding theory provides the best model for analyzing Bitcoins for at least two reasons. First, experimental markets have confirmed that investors exhibit herding behavior (Levine & Zajac 2007). Second, the herding theory can be modeled mathematically in a manner that provides a close fit to historical bubbles (Geraskin & Fantazzini 2010; Lin et al. 2009). These models will form the basis for the remainder of this paper.

5 The LPPL Model

5.1 Price Dynamics

Log-periodic power law (LPPL) models have their origins in statistical physics and are used to determine when a system is reaching a critical point. While the original applications were in areas such as magnetism and seismology, the first use in economics was by Didier Sornette in 1996, applying the search for critical points to the October 1987 market crash. This line of research developed into the Johansen-Ledoit-Sornette (JLS) model, which is the model I have chosen to work with.

For simplicity's sake, the model assumes an asset with no dividends and ignores interest rates, risk aversion, information asymmetry, and liquidity constraints. The asset follows a martingale process:

$$E_t[p(t')] = p(t) \tag{2}$$

Where $p(t)$ stands for the price of the asset at time t and $t' > t$. Since the asset pays no dividends, its fundamental value is $p(t) = 0$, therefore any positive values of $p(t)$ represent an appreciation over intrinsic value, i.e. a bubble. This equality also establishes that there can be no arbitrage within this system.

The next step is to characterize the agents participating in the system. Johansen et al. (2000) propose that all traders are connected via a global network where each agent is influenced locally by his or her nearest neighbors. These agents can choose either to buy or to sell. The agent's decisions are the product of two influences – his own individual

instincts and the opinions of his neighbors. When individual instinct dominates, the overall system will be characterized by randomness and disorder, with approximately balanced numbers of agents selling and buying. Once imitation begins to dominate, a crash will result as soon as the general opinion shifts toward sell. Incidentally, this conception provides a reasonable explanation for why crashes can occur in the absence of an easily identifiable external shock to the system.

The ultimate goal is to explain the macro-level probability of a crash using micro behavior, which will be done in the following steps. The crucial variable is the hazard rate $h(t)$, which represents the probability at any time t that a crash will occur given that it has not already.

5.2 Macroscopic Modeling

Using mean field theory (Goldenfield 1992), the imitative process among traders allows the hazard rate to be modeled by the following equation:

$$\frac{dh}{dt} = Ch^{\delta} \tag{3}$$

where $C > 0$ is a constant, and $\delta > 1$ represents the average number of interactions between traders minus one. This equation models the effective behavior of all traders by representing the average interaction. It can easily be seen that as the number of interactions increases, the hazard rate must also increase. Integrating (3) yields:

$$h(t) = \left(\frac{h_0}{t_c - t}\right)^{\alpha}, \qquad \alpha = \frac{1}{\delta - 1} \tag{4}$$

where t_c represents the critical time of the system crash.

Since the model assumes a not-zero probability of a crash occurring, a jump process j is defined, equal to zero before the crash and one afterwards. Since the critical time t_c is unknown, it is defined in probabilistic terms: The cumulative distribution function is $Q(t)$, and probability density function is $q(t) = dQ/dt$. The hazard rate was earlier defined as the probability a crash would occur in the next period given it has not already, which can be expressed mathematically as $h(t) = q(t)/[1\text{-}Q(t)]$. Making the simplifying assumption that the price falls by a fixed percentage $\kappa \in (0,1)$ in the case of a crash, the price is given by:

$$dp = \mu(t)p(t)dt - \kappa p(t)dj \qquad (5)$$

In order to satisfy the martingale condition in equation (2), $E_t[dp] = \mu(t)p(t)dt - \kappa p(t)h(t)dt = 0$, which can be simplified to $\mu(t) = \kappa h(t)$. Plugging this back into (5), the behavior of the price before the crash (when j=0 still obtains) can be modeled by the following differential equation:

$$\log\left[\frac{p(t)}{p(t_0)}\right] = \kappa \int_{t_0}^{t} h(t')dt' \qquad (6)$$

According to (6), the higher the probability of a crash, the faster the price must increase to maintain the martingale condition. This makes perfect sense, as investors generally demand a higher rate of return on risky investments.

5.3 Microscopic Modeling

The JLS model assumes a group of traders indexed by integers i=1,…,I where $N(i)$ other agents are directly connected to each agent i. Each agent takes one of two states: "buy,"

denoted s_i = +1, or "sell," s_i = -1. Johansen et al. (2000) model the state of each trader by the following equation:

$$s_i = \operatorname{sign}\left(K \sum_{j \in N(i)} s_j + \sigma \varepsilon_i\right) \tag{7}$$

where the function $sign(x)$ is equal to +1 for $x > 0$ and -1 for $x < 0$, K is a positive constant, and ε_i is an i.i.d. variable with a standard normal distribution. In this model, K represents the tendency towards imitation, and σ represents the tendency towards idiosyncratic behavior. The value of K relative to σ determines the balance between order and disorder; as noted earlier, when the effect of order dominates a crash will occur. Since equation (7) only models the state of one agent at one point in time, new influences can enter the system and change the state of an agent and his neighbors. For this model, there exists a critical point K_c that determines the state of the system. When K is less than K_c, the system is characterized by disorder – sensitivity to influence is low, and imitation only occurs among small groups. As K approaches K_c, order appears, imitation carries over long distances, and the sensitivity of the system goes to infinity. To model the system, imagine a two-dimensional grid where each trader has four neighbors at equal distances. The susceptibility χ of the system, that is, the chance that a large group of agents suddenly takes the same state, can be modeled by the following power law:

$$\chi \approx A(\, K_c - K)^{-\gamma} \tag{8}$$

where A is a positive constant and γ is the critical exponent of susceptibility. This model is drawn from the two-dimensional Ising model, used in statistical physics to describe interactions that determine molecular spin. K cannot be modeled directly, since the dynamics driving it are unknown. Instead, as long as we assume it evolves smoothly, K can be tied to the time t. If the critical time t_c is defined as the time when $K(t_c)$ reaches K_c, the behavior of K before it reaches K_c can be approximated by $K_c - K(t) \approx constant \times (t_c - t)$. Johansen et al. (2000) state that the hazard rate around the critical time behaves in the same way as the susceptibility, yielding the following equation:

$$h(t) \approx B\times(t_c - t)^{-\alpha} \tag{9}$$

where B is a positive constant. However, since it relies on a relatively simple model of interaction, this model is not yet sufficient to explain global trading behaviors. Johansen et al. (2000) propose instead a hierarchical diamond lattice, which takes into account that trading occurs on multiple levels (e.g. institutional, individual) and that interaction occurs across all levels. To conceive of such a structure, imagine two traders linked to one another by a straight line. Then replace the line with a diamond, adding two new vertices between the original traders, which represent two new traders. The system now has four traders and four links. Next repeat the process, turning each direct link into a diamond in the same manner. A visualization of such a process is represented below.

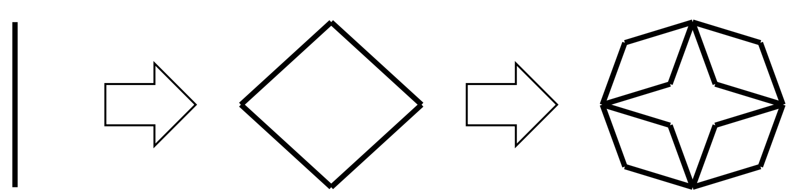

After n iterations of this process, a total of $(2/3) \times (2 + 4^n)$ traders and 4^n links is reached. The least-connected traders have only two neighbors, while the original traders have 2^n neighbors. The rest fall somewhere between these two extremes. A version of this model was solved by Derrida et al. (1983). As above, there exists a critical point K_c. While $K < K_c$, the susceptibility remains finite, approaching infinity as K towards K_c. The most important difference between this and the simpler model is that the critical exponent of susceptibility γ can be a complex number. The general solution is as follows:

$$\begin{aligned}\chi &\approx \mathrm{Re}\left[A_0(K_c - K)^{-\gamma} + A_1(K_c - K)^{-\gamma + i\omega} + \cdots\right] \\ &\approx A_0'(K_c - K)^{-\gamma} + A_1'(K_c - K)^{-\gamma}\cos[\omega \log(K_c - K) + \psi] + \cdots\end{aligned} \tag{10}$$

where A_0, A_1, ω, and ψ are real numbers, and $\mathrm{Re}(x)$ represents the real part of a complex number. The power law in (8) is now characterized by oscillations whose frequency accelerates as the model approaches the critical point. The oscillations are referred to as "log-periodic" since they are periodic in the logarithm of $(K_c - K)$, with frequency $\omega/2\pi$. Applying the same logic used to obtain (9), this model can be rewritten in terms of the hazard rate $h(t)$ and the time t:

$$h(t) \approx B_0(t_c - t)^{-\alpha} + B_1(t_c - t)^{-\alpha}\cos[\omega \log(t_c - t) + \psi'] \tag{11}$$

As above, the hazard rate increases dramatically as t approaches the critical time, but this increase is now decorated by a sequence of accelerating oscillations. Applying equation (11) to (6), the asset price before the crash can be modeled as follows:

$$\ln[p(t)] \approx \ln[p(c)] - \tfrac{\kappa}{\beta}\{B_0(t_c - t)^{\beta} + B_1(t_c - t)^{\beta}\cos[\omega \ln(t_c - t) + \phi]\} \tag{12}$$

Equation (12) can be rewritten as follows:

$$\ln[p(t)] \approx A + B_0(t_c - t)^{\beta}\{1 + C\cos[\omega\ln(t_c - t) + \phi]\} \tag{13}$$

where A is equal to the log of the price at the critical time $\ln[p(c)]$, B represents the increase in $\ln[p(t)]$ prior to the crash, C represents the magnitude of the oscillations, β quantifies the power law, ω represents the frequency of the oscillations, and ϕ is a phase parameter. Equation (13) is called the Log Periodic Power Law (LPPL) and describes the growth of asset prices leading up to a crash.

6 Application to December 2013 Bitcoin Crash

6.1 Fitting LPPL Models

With six parameters ($A, B, C, \beta, \omega,$ *and* ϕ) to fit, estimation of LPPL models can be difficult. Johansen et al. (2000) propose the following method. First, approximate values are chosen for $\beta, \omega,$ and ϕ – the nonlinear parameters. Equation (13) can now by rewritten as $y_i = A + Bf_i + Cg_i$, and $A, B,$ and C can be obtained using ordinary least squares. Plugging these estimates back into the model, a search algorithm is then applied to find the values for $\beta, \omega,$ and ϕ that minimize the sum of squared residuals. Geraskin & Fantazzini (2010) suggest that a better fit can be obtained by first reversing the original time series so that the most recent data point becomes t_1; the LPPL model is then applied to this "anti-bubble." Because the variance of the time series increases as it approaches the critical point, reversing the time series minimizes the effect of the non-stationary component.

In order to identify whether or not the results of the modeling indicate the formation of a bubble, Lin et al. (2009) impose the following conditions:

$$B < 0$$

$$0 < \beta < 1$$

$$6 < \omega < 13$$

$$|C| < 1$$

Residuals are stationary

The first two conditions indicate "faster-than-exponential" acceleration of the log-price, the third condition expresses that the oscillations are neither too slow (such that they would simply become part of the trend) or too fast (such that they would fit the random component of the data). The fourth condition, proposed by Bothmer and Meister (2003) ensures that the hazard rate always remains positive. The final condition checks whether there is still movement in the data not adequately explained by the model. In a test on the S&P 500 Index from 1950 until 2008, Lin et al. (2009) found that the conditions listed above were met for 3% of the 563 windows tested. By contrast, when tested against a set of 2000 synthetically generated GARCH time series, less than 0.2% were found to obey the LPPL bubble conditions, demonstrating a very low rate of false positives.

6.2 The December 2013 Bitcoin Bubble

On December 4, 2013, Bitcoin prices peaked, closing at $1,237.96. The next day, prices fell by 10%, two days later they had fallen by 43%. Though the prices briefly recovered, by December 18 the closing value was down to $541.00, a 56% drop. A quick look at the data set shows that acceleration began in mid-October.

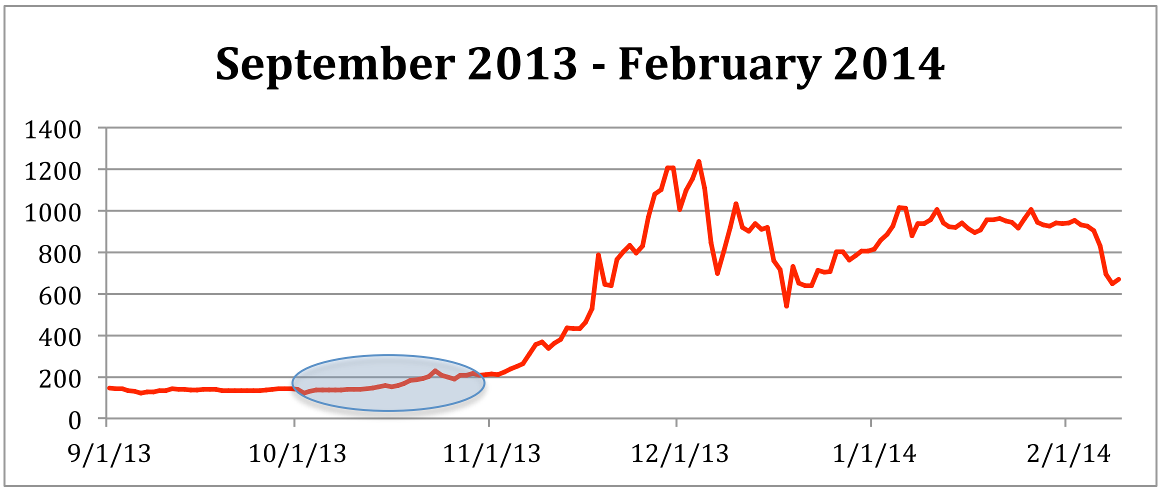

Paradoxically, the most readily identifiable stimulus would have been the seizure of Silk Road and its cache of Bitcoins by the U.S. government in early October. If this were the case, the rise would likely be the result of a growing consensus that the government would ultimately take a positive stance on Bitcoin, as it eventually did in November.

Given this timeline of growth, I chose to examine 60-day windows with end dates from October 4 through December 3 and a step length of five days between windows. For each of the twelve windows, the critical time t_c was tested for all values between one and seven days after the last observation to see if a crash was predicted for the coming week. Though many of the samples met several of the LPPL conditions, only two matched all of the

criteria. These particular windows ended on December 3 and predicted that the crash would occur either three or six days later (December 6 or December 9). As noted above, the actual crash took began on December 4 and reached a local minimum on December 6, making the prediction acceptably accurate.

7 Conclusions

I originally hypothesized that Bitcoin values would be driven both by their worth as a medium of exchange and by their use as a vehicle for investment. However, after testing several variables in the model, the only exogenous variable shown to have any statistical significance was the CBOE Volatility Index. This suggests that, at present, investor behavior matters, but consumer behavior does not.

There are several possible explanations. It could simply be the case that investors make up the largest percentage of Bitcoin users. Since there are relatively few vendors that accept Bitcoins as legal tender, use of Bitcoins to buy goods and services may still be too minimal to have an effect on Bitcoin values. On the other hand, it could be the case that there are quite a few consumers who use Bitcoins as a currency, but their spending behavior and demographics cannot be tracked. We know that online black markets exist and that they use Bitcoin as a medium of exchange; nevertheless, there is no good way to accurately track spending in black markets. Finally, it could be the case that news mentions are not a good enough proxy for consumer participation, and therefore do not really parallel the effects of consumer behavior on Bitcoin values.

Regarding the effects of investor behavior, the model suggests that as the VIX rises, Bitcoin values fall. While this contradicts the original hypothesis that fear and uncertainty would drive investors to shelter their money in Bitcoin, it can still be explained in a manner consistent with investor behavior. The inverse relationship between Bitcoin and VIX values suggests that when the VIX is low, that is, when there is relatively little volatility in the S&P 500, investors push money towards Bitcoins. As noted earlier, Bitcoins have a great deal of volatility, which potentially allows for greater return on investment. Investors might therefore choose to speculate on Bitcoins not because they seek to shelter money, but rather because they wish to gamble in the hopes that Bitcoins will offer a higher return than the market.

Ultimately, the most notable aspect of this forecast is the uncertainty. Confidence intervals are very wide, so the overall confidence in the point forecast is low. With such high volatility, the best an investor or user of Bitcoins could hope for is to have advance warning of dramatic crashes. Fortunately, log-periodic power law models appear to be capable of distinguishing critical periods in Bitcoin trading. Using only data from before the event, LPPL models successfully predicted that a crash was impending. To be fair, the models overshot the date of the crash by two days, but, as Bitcoin trading continues, larger data sets should allow for more accurate predictions.

The successful application of LPPL modeling to Bitcoin also holds promise for use on other digital currencies, most of which experience similar volatility. There is abundant room for further research in this arena – whether digital currencies will remain simply a vehicle for

investment or take on roles as true currencies, what implications the introduction of online currencies has for traditional currencies, whether multiple online currencies can coexist or one will come to dominate, and which, if any, of these cryptocurrencies will ultimately survive.

Appendix

Time window from October 4 to December 3 (testing for crash December 6) that meets the LPPL specifications:

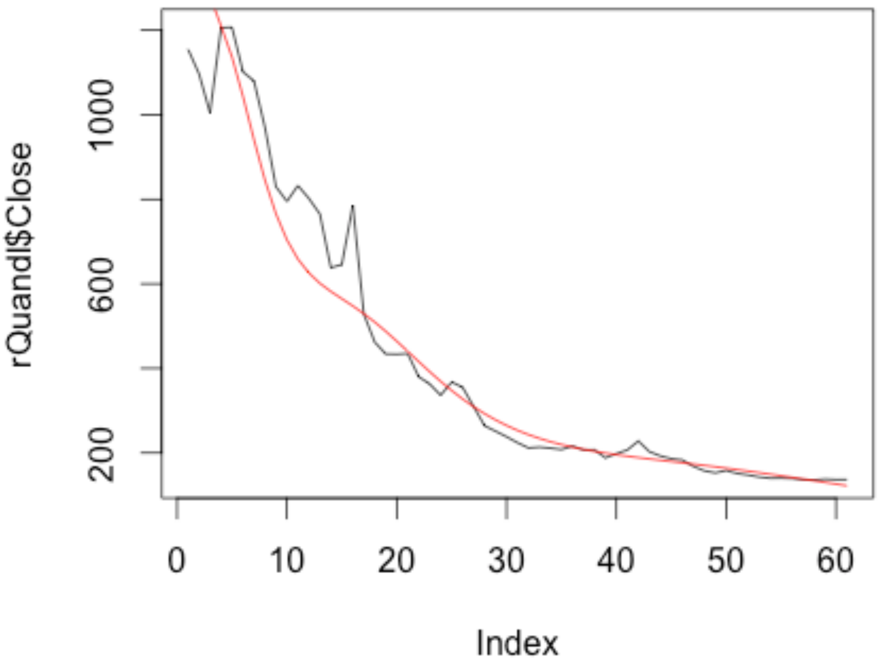

LPPL criteria	
B	-0.76
β	0.4
ω	6.5
C	-0.02

References

Cohan, Peter. "Are Bitcoins Safer Than Cyprus?" *Forbes*. Forbes Magazine, 02 Apr. 2013. Web.

Derrida, B., De Seze, L. and Itzykson, C. 1983. Fractal structure of zeros in hierarchical models. *Journal of Statistical Physics* 33, 559.

Detrixhe, John. "Cyprus Rating Cut Two Steps to B by Fitch on Bank-Bailout Costs." *Bloomberg.com*. Bloomberg, 25 Jan. 2013. Web.

Dreman, David. 1993. "One More for the Road?," *Forbes*. New York, 363.

Garber, Peter. 2001. *Famous First Bubbles: The Fundamentals of Early Manias*. Cambridge, MA: MIT Press.

Geraskin, P., Fantazzini, D. 2010. Everything you always wanted to know about log-periodic power laws for bubble modeling but were afraid to ask, *European Journal of Finance*. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1752115

Goldenfeld, Nigel. 1992. *Lectures on Phase Transitions and the Renormalization Group*. Reading, MA: Addison-Wesley Publishing Company.

Gürkaynak, R. 2008. Econometric tests of asset price bubbles: taking stock. *Journal of Economic Surveys.* Available at http://www.federalreserve.gov/pubs/feds/2005/200504/200504pap.pdf

Johansen, A., Ledoit, O. and Sornette, D. 2000. Crashes as critical points, *International Journal of Theoretical and Applied Finance*, 3, no. 2: 219-255. Available at http://arxiv.org/abs/cond-mat/9810071

Levine, S.S. and Zajac, E.J. 2007. The institutional nature of price bubbles. Available at: http://ssrn.com/abstract=960178.

Lin, L., Ren., R.E. and Sornette, D. 2009. A Consistent Model of Explosive Financial Bubbles With Mean-Reversing Residuals, *Quantitative Finance*, in press. Available at http://arxiv.org/abs/0905.0128

Meier, Simone, and Christian Vits. "Weber Says ECB Has Used Room to Cut Interest Rates (Update3)." *Bloomberg.com*. Bloomberg, 23 June 2009. Web.

NPR. "Silk Road: Not Your Father's Amazon.com." *NPR*. NPR, 12 June 2011. Web.

Sornette, D., Johansen, A. and Bouchaud, J.-P. 1996. Stock market crashes, precursors, and replicas. *Journal of Physics I France* 6, 167-175.