Math Finan Econ
DOI 10.1007/s11579-016-0170-z

The lifetime of a financial bubble

Yoshiki Obayashi[1] · Philip Protter[2] · Shihao Yang[3]

Received: 15 June 2015 / Accepted: 21 March 2016

Abstract We combine both a mathematical analysis of financial bubbles and a statistical procedure for determining when a given stock is in a bubble, with an analysis of a large data set, in order to compute the empirical distribution of the lifetime of financial bubbles. We find that it follows a generalized gamma distribution, and we provide estimates for its parameters. We also perform goodness of fit tests, and we provide a derivation, within the context of bubbles, that explains why the generalized gamma distribution might be the natural one to expect for the lifetimes of financial bubbles.

Keywords Financial bubbles · Bubble lifetimes · Strict local martingales · Generalized gamma distributions

JEL Classification G100 · G120 · C810

1 Introduction

We combine a mathematical analysis of financial bubbles and a statistical procedure for determining when a given stock is in a bubble, with an analysis of a large data set, in order to compute the empirical distribution of the lifetime of financial bubbles. We find that it follows a generalized gamma distribution, and we provide estimates for its parameters. We also perform goodness of fit tests.

Financial bubbles are often described as when the price of a risky asset (in this paper, we study stocks) has a market price that exceeds the price a rational person would pay for the stock, known as the *fundamental price*. The price a rational person would pay is typically

✉ Philip Protter
pep2117@columbia.edu

1 Applied Academics LLC, New York, NY, USA

2 Statistics Department, Columbia University, New York, NY 10027, USA

3 Statistics Department, Harvard University, Cambridge, MA 02138, USA

Published online: 30 March 2016

considered to be the conditional expectation of the future cash flows of the stock, considered under a risk neutral measure. This definition is problematic, for while one can observe market prices, it is not really possible to calculate future cash flows.

This is where mathematics comes to the rescue. In a series of research papers (selected) [4–6,11,18,19,24] it has become clear (and was proved in [18,19]) that on a finite time horizon the market price exceeds the fundamental price if and only if the market price is a strict local martingale under a selected risk neutral measure (this arbitrariness of a choice of risk neutral measure has been removed to some extent recently; see [11]). Therefore to determine whether or not a given stock has bubble pricing, we need only to check whether or not the market price process, under a risk neutral measure, is a strict local martingale, or alternatively a true martingale.

Let us describe this in a bit more detail. We begin with a complete probability space $(\Omega, \mathcal{F}, P)$ and a filtration $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$ satisfying the "usual hypotheses." We let $r = (r_t)_{t\geq 0}$ be at least progressively measurable, and it denotes the instantaneous default-free spot interest rate, and

$$B_t = \exp\left(\int_0^t r_u du\right) \tag{1}$$

is then the time t value of a money market account. We work on a time interval $[0, T]$ where T is a finite fixed time. Working with one chosen and fixed stock, we let $D = (D_t)_{0\leq t<\tau} \geq 0$ be the dividend process, and we assume it is a semimartingale. We let $S = (S_t)_{0\leq t\leq T}$ be nonnegative and denote the semimartingale price process of the stock. Since S has càdlàg paths, it represents the price process *ex cash flow*. By ex cash flow we mean that the price at time t is after all dividends have been paid, including the time t dividend. We also assume that S is such that there is an absence of arbitrage opportunities, in the sense of Delbaen and Schachermayer [8]. Let $\Delta \in \mathcal{F}_\tau$ be the time τ terminal payoff or liquidation value of the asset. We assume that $\Delta \geq 0$.

Next let W be the *wealth process* associated with the market price of the risky asset plus accumulated cash flows:

$$W_t = \mathbf{1}_{\{t<\tau\}} S_t + B_t \int_0^{t\wedge\tau} \frac{1}{B_u} dD_u + \frac{B_t}{B_\tau}\Delta\mathbf{1}_{\{\tau\leq t\}}. \tag{2}$$

Note that all cash flows are invested in the money market account. We assume the market is incomplete, so that there are an infinite number of risk neutral measures. For a given risk neutral measure Q we take conditional expectations in (2) and rearrange the terms, this translates into:

$$S_t^\star = E_Q\left(\int_t^{\tau\wedge T^\star} \frac{1}{B_u} dD_u + \frac{\Delta}{B_\tau}\mathbf{1}_{\{\tau\leq T^\star\}}\middle| \mathcal{F}_t\right) B_t. \tag{3}$$

The superscript $\star$ will be used systematically to denote fundamental values.

Definition 1 We define β_t by

$$\beta_t = S_t - S_t^\star,$$

the difference between the market price and the fundamental price. (In a well functioning market, this difference is 0.) The process β is called a *bubble*.

Note that the absence of arbitrage keeps β nonnegative.

The next theorem (see [19] or [27] for proofs and more details) allows us not to consider the *fundamental price* of our stock given in (3), which is vague and ultimately unknowable, with a precise mathematical criterion provided by Theorem 1 below:

Theorem 1 *A risky asset price process S is undergoing bubble pricing on the compact time interval $[0, T]$ if and only if under the chosen risk neutral measure the bubble process β is not a martingale but is a strict local martingale. This is equivalent to the price process S being a strict local martingale (since $S^\star$ is always a martingale) and not a martingale.*

With Definition 1, which involves $(S_t^\star)_{0\le t\le T}$ and hence implicitly the risk neutral measure, and despite the recent work [11] where as many objects as possible are defined without reference to a risk neutral measure, there is nevertheless the issue of which risk neutral measure should one use. Fortunately, in simple cases that are nevertheless sophisticated enough to be useful, we can finesse this issue. For example, if we model the dynamic evolution of the stock price as a solution of a stochastic differential equation of the form

$$dX_t = \sigma(X_t)dB_t + b(X_t, Y_t)dt; \quad X_0 = 1 \tag{4}$$

where b is controlled as to not be too big, and Y represents a stochastic process reflecting relevant market forces, then under any one of the infinite choice of risk neutral measures Q we have that the drift disappears via a Girsanov type transformation, to get

$$dX_t = \sigma(X_t)dB_t; \quad X_0 = 1 \tag{5}$$

and therefore it does not matter which risk neutral measure we use! (See for example [17,27] for details of this procedure.) Moreover in current work, J. Jacod and P. Protter [16] have been able to extend this idea to equations with stochastic volatility, although there are many caveats to this procedure:

$$dX_t = \sigma(X_t, \nu_t)dB_t + b(X_t, Y_t)dt; \quad X_0 = 1 \tag{6}$$

The key result we use in this paper is taken from [17], although here we refine it to a large extent. The idea is the following: We use the work of [10,13,14] as well as our own work, to estimate from data the coefficient $\sigma(x)$ in equation (4), getting $\hat{\sigma}$, which is equivalent to estimating it for (5). We then use a Reproducing Kernel Hilbert Space (RKHS) method of interpolation to smooth the estimate to form a graph of $\hat{\sigma}$. Next we use the theory developed in [9,21,25] that says if $\sigma(x)$ satisfies the two conditions

$$\begin{aligned} \int_0^\varepsilon \frac{x}{\sigma(x)^2}dx &= \infty \\ \int_\varepsilon^\infty \frac{x}{\sigma(x)^2}dx &< \infty \end{aligned} \tag{7}$$

then $X > 0$ and also X is a strict local martingale. If on the other hand the integral in the inequality (7) equals ∞, then X is a martingale. Therefore it is the asymptotic behavior of σ as $x \to \infty$ that is the key in deciding whether or not X is a strict local martingale. To determine this, we extend $\hat{\sigma}$ smoothly to the entire half line $[0, \infty)$, using an RKHS based extrapolation procedure, coupled with an optimization criterion, and then we use this to determine whether or not we are experiencing bubble pricing. This is described in detail in [17], so we do not reproduce it here.

Using these ideas we can in effect determine when a stock enters into a bubble, and when it exits from a bubble, and therefore determine its lifetime. This paper reports on the use of a procedure such as this, and from the procedure and the analysis of tick data for thousands of stocks over a 14-year period, we are able to determine the empirical distribution of the lifetime of financial bubbles. The bubble lifetimes follow a generalized gamma distribution, a standard distribution for lifetime models, and one that arises naturally in survival analysis.

An outline of this paper is as follows. In Sect. 2 we describe our methodology in detail. In Sect. 3 we present the results, and in Sect. 4 we give some caveats concerning our procedure. In Sect. 5 we give a theoretical explanation as to why we should not be very surprised to find that the generalized gamma is the empirical distribution of the bubble lifetimes, and in Sect. 6 we conclude.

2 Methodology

2.1 The actual model

In the introduction, we make the assumption that the stock price evolves according to a stochastic differential equation of the form

$$dX_t = \sigma(X_t)dB_t + b(X_t, Y_t)dt; \quad X_0 = 1 \tag{8}$$

This is a fairly good model for short periods of time, but it is not realistic for a long period of time, and should be replaced (at least) by adding some time dependence, to obtain a model of the form

$$dX_t = \sigma(t, X_t)dB_t + b(t, X_t, Y_t)dt; \quad X_0 = 1 \tag{9}$$

Such a model, however, is a bit too general to lend itself to a good analysis. We make a compromise by assuming the stock prices locally evolve according to (8) and globally evolve according to (9). More precisely we make the assumption that for a compact time interval $[0, T]$ we have a finite partition $0 = t_1, \ldots, t_n = T$, and the volatility coefficient σ has the form:

$$\sigma(t, x) = \sigma_1(x)1_{[t_1,t_2]}(t) + \sum_{i=2}^{n} \sigma_i(x)1_{(t_i,t_{i+1}]}(t) \tag{10}$$

This is in effect a type of regime change model. Note that when we are in a time interval such as the partition interval $(t_i, t_{i+1}]$, we are free to use the criterion (7), under the risk neutral measure so that the drift is irrelevant, to determine if a bubble is detected, or not.

Note that were we to use a model such as (8) as our global model, then once we detected that X, the solution of

$$dX_t = \sigma(X_t)dB_t \text{ under the risk neutral measure } Q, \tag{11}$$

was a strict local martingale, it would have to remain one for all time. Indeed, one can show, using the positivity of X and the strong Markov nature of (11), that the expectation of X under a risk neutral measure Q, is monotonically and strictly decreasing for all time. (See the working paper of Jacod and Protter [16].) From the standpoint of financial bubbles, such a claim would be that once a stock enters a bubble it is in one for all time, and of course this is silly.

Therefore we replace the assumption (8) with its more general form (9), but with the caveat that $\sigma(t, x)$ is of the special form (10). This will allow us to detect, from data, when the stock price enters a bubble, and also when it exits the bubble. It will not, however, allow us to predict the end time of a bubble. Nevertheless, once we know the (empirical) distribution of the lifetime of a bubble, we will be able to make statements of the nature that, at t units of time after the birth of a bubble,

$$P(\text{remaining life of the bubble } > \alpha | \text{bubble is alive now}) = P(L > \alpha + t | L > t)$$
$$= \frac{P(L > \alpha + t)}{P(L > t)},$$

where L is the random variable giving the lifetime of a bubble. In Sects. 2.3 and 2.4.5 we break down our procedure into steps.

2.2 The Data

First we describe the data set we used. We use TAQ (Trade and Quote) data obtained from Wharton Data Research Services (WRDS), for the period from January 1, 2006 to July 31, 2013, a period of little over 7.5 years, on SP500 component stocks. We later expanded this to include the years from 2000 to 2013 on the top 3000 stocks with highest market capital, which consists of in total approximately 3500 different stocks. So the data set is quite extensive. For the processing of the data we first used the Amazon Cloud Drive, and later we switched to the high speed research computing facilities of Harvard University. We used a subsampling procedure with the (noisy) data, modeled along the lines proposed by Zhang et al [31].

2.3 The Raw Mathematical Treatment

Our actual global model is (9), but at a given point in time we assume that (8) holds locally. This can be achieved, for example, by assuming the volatility function is slowly evolving in time or is piece-wise constant (10). We will focus on the local model (8) in this subsection, and on the global model (9) in the next.

As given in Definition 1, a bubble exists at a time t if the process $\beta_t = S_t - S_t^\star \neq 0$. This is equivalent, due to the assumption of an absence of arbitrage opportunities, to $\beta_t = S_t - S_t^\star > 0$. The bubble process β is the difference between the market price and the fundamental price, and if it is not zero, then the theory recalled in Theorem 1 tells us that β is a strict local martingale under a risk neutral measure, and since $S^\star$ is a martingale, we infer that there is a bubble if and only if the market price S is a strict local martingale under a risk neutral measure. Note that because our model follows an SDE of the form (4), we have that irrespective of the risk neutral measure Q actually chosen, the equation under Q is of the form

$$dX_t = \sigma(X_t)dB_t; \quad X_0 = 1 \tag{12}$$

and with minimal assumptions on the coefficient g we have the weak uniqueness of of the solution of (12), so under all risk neutral measures the distribution of X is the same. Thus to determine whether or not β is a strict local martingale, we need to determine if S is a strict local martingale, which in turn means we need to check to see whether or not the condition (7) holds.

This seems simple, but it is really not. First, we need to determine the volatility coefficient $x \mapsto \sigma(x)$, via an estimation from the data. Actually the function should be $(t, x) \mapsto \sigma(t, x)$, but we make the assumption that $\sigma(t, x)$ is locally constant in the time variable, so in any sufficiently short given time interval we are actually approximating a function of the type $t \mapsto \sigma(x)$ as given in (12). We break the procedure into steps:

The description of our technique given below in these four steps is given under the assumption we are in a regime where the coefficient is of the form $\sigma(x)$, and not $\sigma(t, x)$; using the notation of (10), this means we are implicitly assuming we are working on one of the partition time intervals $(t_i, t_{i+1}]$ given in the description (10).

Step 1 In this step we estimate the volatility coefficient function $x \mapsto \sigma(x)$ from Eq. (4). To do this we use a slight modification of techniques developed by Florens-Zmirou [10] and Jacod [13,14]. This already is described in detail in the paper [17], and we do not repeat the description here. We do mention that this is a non parametric estimation of a function, so it is not simple. The method of Florens-Zmirou is a local time based estimator, but the work of Jacod improves it quite a bit. We note that we this technique can estimate $\sigma(x)$ only for those x that are in the range of the market price process S. That is, we can only estimate $x \mapsto \sigma(x)$ for those values of x that are assumed at some point by the market price process S. Since the world is finite, the range of S is *a fortiori* bounded. This is a problem because to use the condition (7) we need to know the asymptotic behavior of $\sigma(x)$ as x increases to ∞. We deal with this problem in Step 3 below.

Step 2 Typically the estimate of the function σ, denoted $\hat{\sigma}$, is noisy, so we use a smoothing procedure, via Reproducing Kernel Hilbert Space (RKHS) smoothing. We smooth by interpolation our estimate of σ within the bounded interval where we have observations, and in this way we lose the irregularities of non parametric estimators. This is a standard procedure in Statistics, analogous to linear regression; however the RKHS technique is better adapted to fitting graphs of function which need not be linear. Again, we refer the reader to [17] for details.

Step 3 From Step 2, we now have $\hat{\sigma}$, smoothed, on a compact space interval where price observations are available. But as explained above in Step 1, we need to know the asymptotic behavior of σ (and hence of $\hat{\sigma}$) in order to determine whether or not (7) holds. So we need to extend $\hat{\sigma}$ into the rest of the positive real line $[0, \infty)$.
We do this via the choice of a certain extrapolating RKHS, which—once chosen—determines the tail behavior of our volatility $\hat{\sigma}$. If we let $(H_m)_{m\in\mathbb{N}}$ denote our family of RKHS procedures, then any given choice of m, call it m_0, allows us to interpolate *perfectly* the original estimated points, and thus provides a candidate *RKHS* H_m with which we extrapolate $\hat{\sigma}$. But this represents a choice of m_0 and not an estimation. If we were to stop at this point the method would be too arbitrary. So we do not arbitrarily choose m_0. Instead we choose the index *m using the data available*. In this sense we are using the data twice, reminiscent of a bootstrap technique. To do this we evaluate different RKHS's in order to find the most appropriate one *given the arrangement of the finite number of grid points* from our observations. The *RKHS* method has been used in similar ways in other scientific disciplines. See for example [12], where the authors use it for the reconstruction of functions from scattered data in certain linear functional spaces.

Step 4 We have now estimated the coefficient $\sigma(x)$, cleaned it of the noise, and extended it to the entire half line $x \in [0, \infty)$. To see if the price process S is a strict local martingale or just a martingale on a compact time interval $[a, b]$, we next need to see if

$$\int_{\varepsilon}^{\infty} \frac{x}{\sigma(x)^2} dx < \infty \tag{13}$$

in which case S is a strict local martingale, or alternatively the integral term of (13) equals ∞, in which case S is a true martingale and we do not have a bubble during $[a, b]$. To do this, we check the asymptotic behavior $\lim_{x\to\infty} \frac{x}{\sigma(x)^2}$ to see if $\sigma(x)^2$ tends to infinity sufficiently fast as $x \to \infty$ so that (13) holds, or does not go to infinity sufficiently fast for (13) to hold.

2.4 A refined mathematical treatment

While the methodology works in theory, and as shown in [17] works well for the extreme bubbles of the dot com period, for the type of massive data study we have undertaken, several problems arise. The methodology is designed to estimate the birth of bubbles, since during the birth the stock price process changes from a martingale under the risk neutral measure to a strict local martingale, and this is what we detect. This is fine for detecting bubbles, but does not work nearly as well for determining lifetimes. A similar phenomenon occurs at the death of the bubble, but during the lifetime of the bubble, the volatility can be more quiescent. We deal with this via two procedures, one is via a hidden Markov model procedure, and the second is the imposition of a filter. We now describe these.

2.4.1 The basic idea

Section 2.3 treats the local model (8). We next combine and smooth the raw signal obtained from Sect. 2.3 in the context of the global model (9).

First we smooth the raw daily signal using a statistical method based on a Hidden Markov Model. The Hidden Markov Model will assign a whole period to be positive if the raw data is mostly positive for this period with only a few fleeting negative signals, and conversely it will assign a whole period to be negative if the raw data is mostly negative for this period with only a few fleeting positive signals. Such a Hidden Markov Model captures the heuristic of the slow evolution in terms of the bubble nature of X, which in turn depends on the evolution of $\sigma(t, X_t)$. In effect the Hidden Markov Model corrects the false positive and false negative signals, leading to a more accurate signal. Details about the Hidden Markov Model we use are presented in Sect. 2.4.2.

But there is a second problem. In analyzing the data we find occasional small signals indicating bubbles that are isolated, and without subsequent signals. And on the other hand, we find a few bubbles that are stable at the peak so no signal could be detected between two apparent signals for bubble birth and bubble death. Therefore, in a second, more brutal smoothing effort, we ignore small isolated signals and the extend bubble signals when the stock price is at its peak between bubble birth and bubble death. In order to do this we impose a 5 % filter, such that we do not count a bubble birth or death unless the price changes (respectively upward or downward) by at least 5 %. Thus we need the conflation of a bubble signal, plus a price move of 5 % or more to indicate a change in bubble status. We explain this procedure in Sect. 2.4.3. Note that the imposition of the filter biases our results concerning the empirical distribution of the lifetimes: probably we are not including some short-lived bubbles.

2.4.2 The hidden Markov model procedure

The first problem is that for general stocks, the procedure is quite noisy, in the sense that it detects briefly the existence of bubbles with great frequency. These are essentially false positives, caused by fleeting bouts of high volatility. To cope with this kind of noise, we first used a hidden Markov model (HMM) indicator, which in effect smoothed out the bubble detection procedure, and resolved the issue. The HMM is modified to accommodate the nature of financial bubbles, where the parameter is tuned in order to make it less likely to have consecutive short-lived bubbles.

In contrast to the bubble detection procedure outlined in the Introduction section, our HMM procedure has not been published elsewhere, so we provide a description of it here.

Let $\mathtt{B}$ indicate that one is in a bubble, and $\mathtt{N}$ indicate the opposite, that one is not in a bubble. Also let $h_{1:T}$ represent a T-dimensional vector: $h_{1:T} = (h_1, h_2, \ldots, h_T)^{\intercal}$, where $h_i \in \{\mathtt{B}, \mathtt{N}\}$ which are the true, yet latent, bubble states on T different days. $\nu_{1:T} = (\nu_1, \nu_2, \ldots, \nu_T)^{\intercal}$, where $\nu_i \in \{+, -\}$ is the T-dimensional vector of the observed signal on T different days, where $+$ indicates a positive signal, and $-$ indicates the absence of a positive signal.

We assume that $h_{1:T}$ is a time-homogeneous first order Markov chain, and that ν_i depends only on h_i through a fixed emission distribution. The conditional distribution for the joint latent state is

$$P(h_{1:T}|\nu_{1:T}) = P(h_{1:T}, \nu_{1:T})/P(\nu_{1:T}), \quad P(h_{1:T}, \nu_{1:T}) = \prod_{t=1}^{T} P(\nu_t|h_t)P(h_t|h_{t-1})$$

with the last equality coming from the Markov property. So the "most likely hidden path" can be found by maximizing *a posteriori* as follows:

$$\begin{aligned} \arg\max_{h_{1:T}} P(h_{1:T}|\nu_{1:T}) &= \arg\max_{h_{1:T}} P(h_{1:T}, \nu_{1:T}) \\ &= \arg\max_{h_{1:T}} \prod_{t=1}^{T} P(\nu_t|h_t)P(h_t|h_{t-1}). \end{aligned}$$

Assume $P(h_t|h_{t-1})$ to be time-homogeneous and to depend on the matrix A given below, the Markov chain transition matrix:

$$\begin{aligned} P(h_t = \mathtt{B}|h_{t-1} = \mathtt{B}) = a_{11}, \\ P(h_t = \mathtt{N}|h_{t-1} = \mathtt{B}) = a_{12}, \\ P(h_t = \mathtt{B}|h_{t-1} = \mathtt{N}) = a_{21}, \\ P(h_t = \mathtt{N}|h_{t-1} = \mathtt{N}) = a_{12}. \end{aligned}$$

$P(\nu_t|h_t)$ is also assumed to be time-homogeneous and depend only on matrix B, the emission probability matrix:

$$\begin{aligned} P(\nu_i = +|h_i = \mathtt{B}) = b_{11}, \\ P(\nu_i = -|h_i = \mathtt{B}) = b_{12}, \\ P(\nu_i = +|h_i = \mathtt{N}) = b_{21}, \\ P(\nu_i = -|h_i = \mathtt{N}) = b_{22}. \end{aligned}$$

Since $(\nu_t)_{t=1}^{T}$ is observed, if we know matrix A and B, for each series of $(h_t)_{t=1}^{T}$, we can evaluate the above joint probability equation so we can find the series that maximize it. Computationally we can use Viterbi algorithm to simplify calculations (see, for example, [3]).

The above procedure assumes knowledge about the matrices A and B, but in fact they have to be estimated from data. We can write out the likelihood function as

$$L(A, B) = P(\nu_{1:T}) = \sum_{h_T} P(h_T, \nu_{1:T}) = \sum_{h_T} \alpha(h_T), \quad \text{where } \alpha(h_T) := P(h_T, \nu_{1:T})$$

and $\alpha(h_t)$ can be computed recursively as

$$\begin{aligned}
\alpha(h_t) &= P(h_t, \nu_{1:t}) \\
&= \sum_{h_{t-1}} P(h_t, h_{t-1}, \nu_{1:t}) \\
&= \sum_{h_{t-1}} P(\nu_t | h_t, h_{t-1}, \nu_{1:(t-1)}) P(h_t | h_{t-1}, \nu_{1:(t-1)}) P(h_{t-1}, \nu_{1:(t-1)}) \\
&= \sum_{h_{t-1}} P(\nu_t | h_t) P(h_t | h_{t-1}) \alpha(h_{t-1})
\end{aligned}$$

Since we can evaluate the likelihood for any matrix A and B, we can obtain the profile likelihood estimator by fixing the matrix A and then finding the maximizer of B on a grid. Using this HMM we obtain what we call the *smoothed bubble signal*.

The next problem is to calibrate the parameters of the matrices A and B. We make the simplifying assumption that A and B are homogeneous for all ticker symbols (ie, for all stocks). This could be improved perhaps using a hierarchical Bayesian model, but we do not do so here. We then use a profile MLE technique to estimate the parameters, which is an *ad hoc* compromise between the MLE estimate and the economic meaning of the transition probability of true bubble status. That is, we use the MLE for signal error probability, but with a pre-specified bubble transition probability:

The transition rate 0.005 is illuminated by the predominant belief (prior) that bubbles probably last a bit less than one year on average. Conditional on this pre-specified A, we then use a randomly selected 500-stock 3-year subset of data as a training set to estimate the emission probability matrix by a standard MLE approach.

Note that no estimations report standard error because we have so much data that the model-based standard error for estimation is essentially zero.

2.4.3 The 5 % Filter

In addition to the false positive noise question resolved by the imposition of HMM, another problem arose. Sometimes we would have a burst of activity when a bubble was born and again when it died, but in between in some cases the bubble would rest quiescent with respect to our test described in Sect. 2.3, even though the bubble was clearly continuing. To correct this problem we added a 10 % filter: We only recorded bubbles birth when there was a price rise of at least 10 %; analogously the death had to occur with a drop of at least 10 %. This worked well, but the high filter omitted too many bubbles, or at least could have been doing so. So instead we dropped the filter to 5 % for both birth and death of bubbles, and this too worked well, and included bubbles omitted by the 10 % filter. We then stitch together as one bubble a period where in the middle the test might indicate no bubble (See Sect. 2.4.5 for detailed "stitching").

The 5 % filter actually serves a dual purposes: (a) it filters out some bubble signal caused by noise, (b) it classifies each positive signal to be either bubble birth or bubble death.

(a) The first purpose is easy to understand, since inevitably we will have occasional false positive signals, even after HMM (Hidden Markov Model) smoothing. The 5 % filter ensures that beside the super-linear growth of the volatility curve $\sigma(x)$, the price process itself also changed rapidly in absolute value from the onset of the positive signal. As such, we remove the false positive signals caused by small-scale fluctuation of the price process that have no particular economic meaning but could be captured in the shape of volatility curve $\sigma(x)$, since our estimate of the growth rate of $\sigma(x)$ is scale-free.

(b) The second purpose has deeper implications and adds an important piece of information to the bubble detection methodology. Up to now, the bubble detection method is based on the volatility curve, which gives positive bubble signals when the volatility as a function of price level has a sufficiently strong super-linear growth. Such a phenomenon usually happens in the beginning phase of a bubble, when the price rises rapidly, or in the ending phase of a bubble, when the price is likely to plunge. However, we have found that in most of the cycles of financial bubbles, the price level can stay high for a period of time, either reflecting the over-confidence of the market, or the delay in the flow of information.

When a stock is at the peak of a bubble, the price stably stays high and the volatility-based detection method tends to give a negative signal. Therefore, when a bubble is in the birth phase, our detection method can give a positive signal alert, and when a bubble is at a high plateau in price, our detection method tends to give a negative signal, and when a bubble is in the death phase, our detection method will give a positive signal alert again. However, in practice when we are inferring the bubble status, we can only observe the detection signal and we don't know the actual bubble status, thus when we observe a positive signal, we know that it is very likely that either a bubble is forming or a bubble is dying, but based on the detection signal alone we cannot tell which case is occurring. For example, when the detection method gives a positive signal on a particular price process $\{S(t) : 0 \leq t \leq T\}$, if we reverse the price process and define $\tilde{S}(t) = S(T-t)$, the detection method will also give a positive signal on $\{\tilde{S}(t) : 0 \leq t \leq T\}$. However, the underlying bubble status of the price process $S(t)$ and $\tilde{S}(t)$ can be very different: if $S(t)$ resembles a price process during a bubble birth, $\tilde{S}(t)$ must resemble a price process during bubble death. Since the detection method alone cannot differentiate bubble birth or bubble death, we need an extra dimension of information, which is why we need to look at the price change and impose the 5 % filter. We do this by observing the price S_{t_0} at the earliest time we have detected bubble birth, and then declare a bubble once the price rises at least 5 % after this initial signal.

2.4.4 Data related quality of our estimates of the volatility coefficient

It is natural that there is more data when $\sigma(x)$ is small than when it is large. Let us describe our procedure. The time unit is one business day. To get the detection signal on each day t, we need to estimate the growth rate of volatility curve $\sigma_t(x)$. We use a sliding 21-day moving window that immediately precedes day t to estimate $\sigma_t(x)$. Within each day during business hours 9:30 AM to 4:00 PM, we use the tick data sub-sampled at 5 minutes interval as suggested by [31], so in each day we have 79 observations of the price process. Therefore in total we have $21 \times 79 = 1659$ observations to estimate the curve $\sigma_t(x)$ on each day t. We also exclude the overnight change in price when estimating $\sigma_t(x)$ in order to combat unobserved information flow outside business hours. Empirically we find that such an estimating schedule produces a fairly stable estimate with a relatively small standard error. As time continues, we update the estimate $\sigma_t(x)$ on a daily basis, incorporating newly available price observations and discarding obsolete observations more than 21 days old, so as to capture the most recent dynamics in $\sigma_t(x)$.

On each day we estimate $\sigma(x)$ at price level x using the price observations in the neighborhood of x ([19]). In the implementation of bubble detection procedure, we use only estimates of $\sigma(x)$ at level x whose neighborhood has at least 100 price observations, and discard the esti-

Table 1 pre-specified A, bubble state transition probability

	B	N
B	0.995	0.005
N	0.005	0.995

Table 2 MLE of B, signal error probability, conditional on the pre-specified A in Table 1

	+	−
B	0.90	0.10
N	0.01	0.99

mate based on fewer observations. As such, we ensure that we only use accurate estimates of $\sigma(x)$ to determine whether or not the function is increasing sufficiently super-linearly, which is a critical intermediate step of our bubble detection procedure.

2.4.5 Step-by-step procedure

We continue to describe our procedure, now taking into account the time dependence of the coefficient $\sigma(t, x)$. This continues the description of Sect. 2.3.

Step 5 Repeat Step 1–Step 4 in Sect. 2.3 for each day, with the data schedule outlined in 2.4.4. Therefore, for each day in history, we have a raw bubble detection signal indicating whether that particular day is in a bubble or not. Denote the signal on day t as $\nu_t \in \{+, -\}$, where + indicates a positive signal, and − indicates the absence of a positive signal.

Step 6 Use the hidden Markov model (Sect. 2.4.2) with parameter matrix Tables 1 and 2 to smooth the noisy raw signal $(\nu_t)_{t=1,2,\ldots,T}$, and obtain the smoothed signal $(h_t)_{t=1,2,\ldots,T}$, where $h_t \in \{\mathtt{B}, \mathtt{N}\}$ indicates the presence of bubble signal (B) or not (N). The smoothed signal $(h_t)_{t=1,2,\ldots,T}$ has a strong autocorrelation and will switch between B and N only occasionally. We thus define a smoothed signal period to be $[t_1, t_2]$ such that $h_t = \mathtt{B}, \forall t \in [t_1, t_2]$ and $h_{t_1-1} = h_{t_2+1} = \mathtt{N}$.

Step 7 Apply the 5 % filter on all the smooth signal periods. Define a bubble birth period to be a smoothed signal period $[t_1, t_2]$ such that $X_{t_2}/X_{t_1} - 1 \geq 0.05$. Define a bubble death period to be a smoothed signal period $[t_1, t_2]$ such that $X_{t_2}/X_{t_1} - 1 \leq -0.05$.

Step 8 Define the entire bubble period to be the beginning of a bubble birth period to the end of the subsequent bubble death period. Therefore, the *bubble lifetime* for a given bubble in a given stock is the length of the time interval between the bubble's birth and the bubble's death.

3 The results

Using the procedure described in Sect. 2 we examined over 3500 stocks from 2000 to 2013, using tick data with sub-sampling inspired by the paper of Zhang et al [31]. We found bubbles and determined the empirical distribution of their lifetimes. The empirical distribution turned out to be the generalized gamma distribution. The generalized gamma distribution was first proposed by Stacy [29] and has the density

$$f_G(t) = \frac{\lambda p(\lambda t)^{p\kappa-1} e^{-(\lambda t)^p}}{\Gamma(\kappa)}$$

and p and κ are shape parameters. If $\kappa = 1$ then the Generalized Gamma reduces to the Weibull. It also extends the log normal, the exponential, and of course the gamma distribution. Prentice [26] has further re-parameterized the generalized gamma distribution so that the new form includes inverse gamma and inverse Weibull distribution as well. This form of generalized gamma is also stable for maximized likelihood estimates and has the density

$$f_G(t; \alpha, \sigma, q) = \frac{|q|}{\sigma t \Gamma\left(q^{-2}\right)} \left[q^{-2}\left(e^{-\alpha} t\right)^{q/\sigma}\right]^{q^{-2}} \exp\left[-q^{-2}\left(e^{-\alpha} t\right)^{q/\sigma}\right] \tag{14}$$

where α is a location parameter, σ is a scale parameter, and q is a shape parameter that can take negative values. If $q = -1$ then the Generalized Gamma reduces to the inverse Weibull. For more on the distribution, including a detailed discussion of its hazard rates, see [1].

We performed an MLE estimate for the parameters, given our extensive data set. We obtained

- The MLE estimate for $\alpha = 5.2478868$
- The MLE estimate for $\sigma = 0.9016828$
- The MLE estimate for $q = -0.2339065$

We include a graph (cf Fig. 1 below) of the histogram of the empirical distribution, with the generalized gamma density in red superimposed (Fig. 2).

We performed a likelihood ratio test for the sufficiency of the generalized gamma versus an expanded distribution, the generalized F distribution, and we found that the p-value for the extra parameter in the generalized F distribution is 99.93 %. This indicates that the generalized gamma distribution fits the data very well, and that further expending the

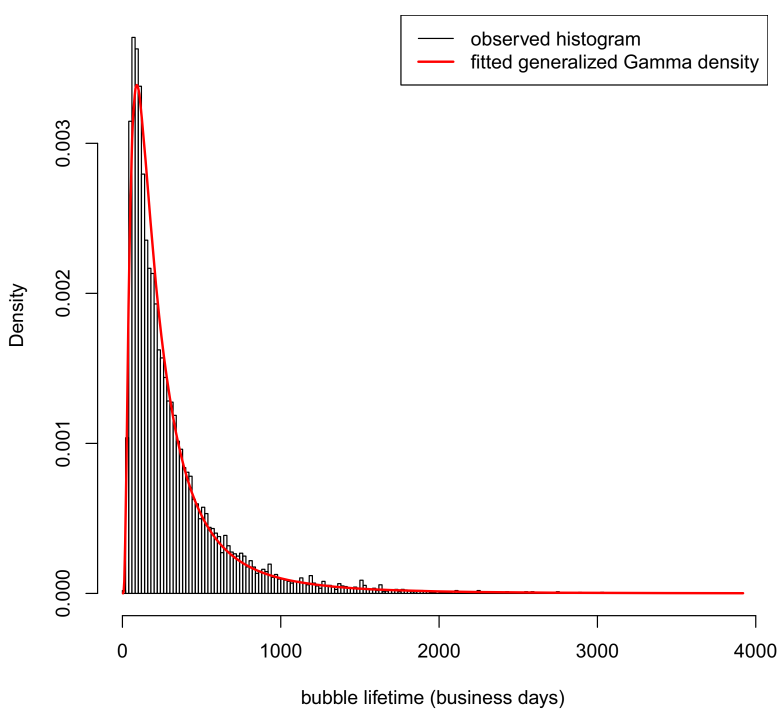

Fig. 1 Histogram of bubble lifetimes

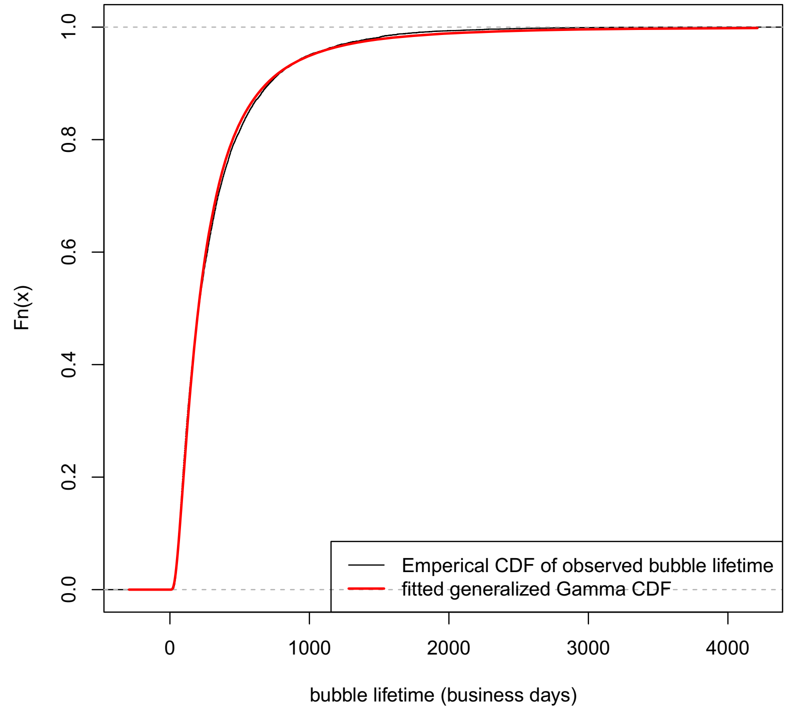

Fig. 2 Goodness of fit graphs

parametric distribution is not necessary. We also performed several likelihood ratio tests for possible redundancy in the generalized gamma distribution. Specifically, we compared the generalized gamma distribution with the Gamma, Weibull, Ammag (see, e.g. [7]), log normal, Inverse Gamma, Inverse Weibull and the Inverse Ammag distributions, and found that the extra parameter in the generalized gamma is significant in all comparisons, even with a Bonferroni correction for multiple testing. We thus conclude that the generalized gamma distribution is sufficient to provide an adequate fit to the bubble lifetime data.

We also remark that a routine application of Glivenko-Cantelli theorem (see e.g. [15]) shows that in this case the empirical distribution converges (uniformly, a.s.) to the true distribution as the number of stocks n tends to ∞, meaning that for large n (such as we have here), the empirical distribution is quite close to the true distribution of bubble lifetimes.

3.1 Frequency of bubbles occurring

We note that during a 13 year period we detected 13,060 bubbles in all of the stocks in our study. This comes to roughly four bubbles per stock during this 13 year period. This seems to us as a surprising result: Bubbles seem to be much more frequent than we thought at the outset of this project.

3.2 Results for specific stocks

We give here some examples of long-lived and short-lived bubbles. We do this via graphs. The green areas are the bubble birth periods, and the purple areas are the bubble deaths. Yellow means that the bubble has begun, and continues via the hidden Markov procedure with the 5 % filter until its death, which is purple. Gray means there is indeed a bubble signal, but it is suppressed by the 5 % change requirement. We begin with two examples of short-lived

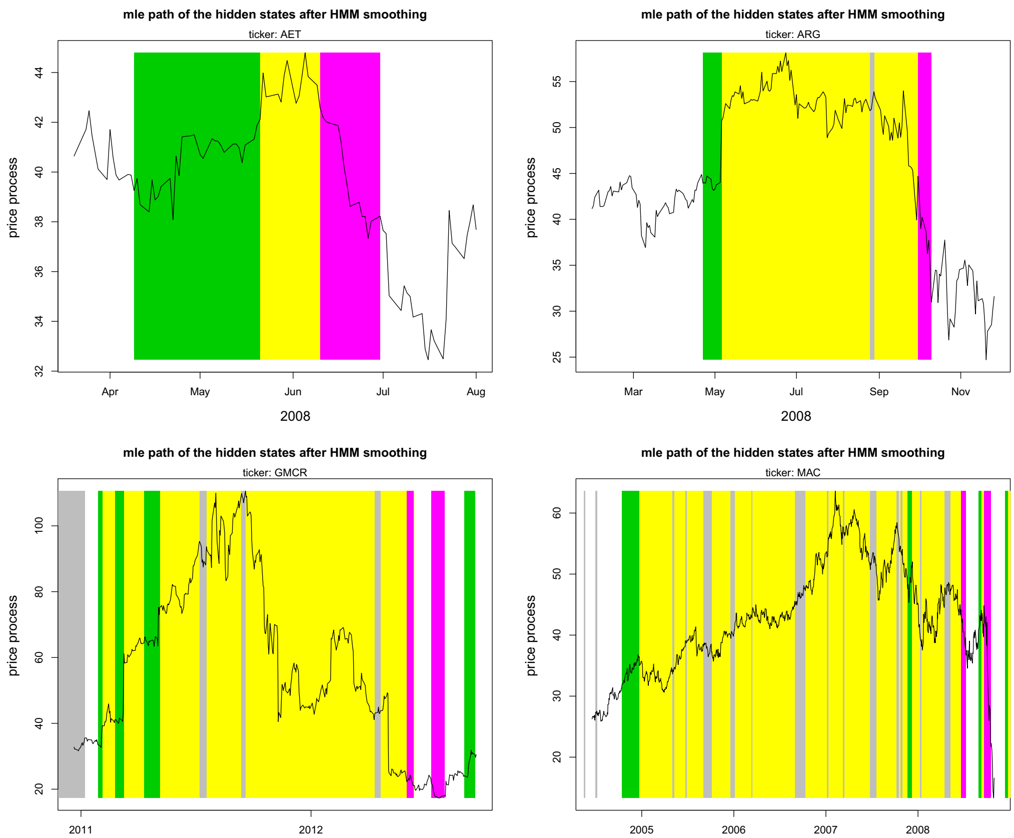

Fig. 3 Bubble lifetime examples for specific stocks. *Top panels* are short-lived bubbles, and *bottom panels* are long-lived bubbles. The *green* areas are the bubble birth periods, and the *purple* areas are the bubble deaths. *Yellow* means that the bubble has begun, but continues via the hidden Markov procedure with 5 % filter until its death, which is *purple*. *Gray* means there is indeed a smoothed bubble signal, but it is suppressed by the 5 % change requirement

bubbles and then give two examples of long-lived bubbles. Graphs are in Fig. 3. The graphs show the stock prices and corresponding bubble states.

- *Short-lived bubble* Our first example is for the company Aetna, Inc (ticker AET), a large American insurance company. The bubble is from April to July, 2011.
- *Short-lived bubble* Our second example is for Airgas, Inc.(ticker ARG) from May to October, 2008.
- *Long-lived bubble* Our third example is for Keurig Green Mountain (ticker GMCR) from early 2011 to mid 2012.
- *Long-lived bubble* Our fourth example is for The Macerich Company (ticker MAC) from late 2004 to mid 2008. Note that this bubble has two deaths, a first tentative one, and then a dramatic one.

4 Caveats concerning the results

It is of course obvious that our results are only as good as our model is accurate. The strict local martingale approach seems to be accepted as a model of financial bubbles, as it is used by quite a few leaders in the field, as is evidenced by the recent papers [4–6, 11, 19, 20, 24, 27].

The method of detecting bubbles, as published in [17] and [27], in contrast, relies on an extrapolation procedure that from a statistics standpoint, is not a consistent estimator. Nevertheless the method seems to work.

Another possible caveat is the 5 % filter technique, which clearly biases the results by excluding some small and short lived bubbles. As such, the empirical distribution probably has a larger component of short lived bubbles than we see in Fig. 1, and therefore a Weibull distribution might ultimately be a better fit than the generalized gamma, if one were to be able to solve this issue in the future. Nevertheless if one is aware of this filter and interprets the results in the light of its use, we do not see this as a true problem.

5 Why the generalized Γ might be a natural choice

The generalized gamma distribution occurs with some frequency in survival analysis, but it is less well known than other distributions it extends, such as the exponential, the Weibull, the gamma, and the log normal distributions. It is a three parameter distribution and was introduced in 1962 by Stacy [29]. In this paragraph we will present a derivation that shows why it is natural for the distribution of the lifetime of bubbles to follow a generalized gamma distribution.

We take the convention that all bubbles begin at time $t = 0$. We can achieve this by simply translating the bubble birth time to $t = 0$. We uniformly partition $\mathbb{R}_+$ into intervals $[t_{i-1}, t_i)$ of length Δt. Next we let N_i denote the number of bubbles still alive in $[t_{i-1}, t_i)$. Let N be the total number of bubbles in our universe. Then

$$\frac{N_i}{N} \text{ is the proportion of bubbles still alive at time } t_{i-1} \tag{15}$$

We assume that the proportion of bubbles alive decreases geometrically with time, and we express this as

$$\sum_{i=1}^{\infty} \left(\frac{N_i}{N}\right) t_i^{\beta} = K \quad \text{for constants } \beta > 0,\, K > 0 \tag{16}$$

We also assume the death rate of bubbles alive at time t_{i-1} is proportional to a power of t_i. This gives that the likelihood of bubble death increases geometrically with age. Thus if we let g_i denote the number of bubble deaths in $[t_{i-1}, t_i)$, we assume

$$g_i = A t_i^{\alpha-1} \tag{17}$$

so that g_i is proportional to a power of t, and the proportionality constant is A.

We next look for the most probable distribution satisfying (15),(16) and (17). This is similar to derivations of the velocity distribution of a degenerate gas, in statistical mechanics (see, for example, [30, pp. 213–220]), and was explained in detail in an old paper of Lienhard and Meyer [22]. Therefore we do not repeat the details here, but merely present a sketch of the ideas.

Let W be the number of ways bubbles can die in $[t_{i-1}, t_i)$ given that they are alive at time t_{i-1}, for all intervals $[t_{i-1}, t_i)$ over $[0, \infty)$. For example, bubbles can have a dramatic death, or they can die slowly, with a whimper, and one can give descriptions in between. An alternative way of viewing this is that there can also be varying economic explanations for why bubbles die, such as disagreements among different agents as to the state of current conditions; see for example [28]. Then as derived in Sommerfeld [30],

$$W = N! \prod_{i=1}^{\infty} \frac{g_i^{N_i}}{N_i!} \tag{18}$$

We let $\tilde{N}_i$ denote the values of N_i that maximize W. One can then show

$$\frac{\tilde{N}_i}{N} = \frac{\Delta t \left[\beta \left(\frac{\beta K}{\alpha} \right)^{-\alpha/\beta} \right]}{\Gamma\left(\frac{\alpha}{\beta}\right)} t_i^{\alpha-1} \exp\left(-\frac{\alpha}{\beta} \frac{t_i^\beta}{K} \right) \tag{19}$$

The idea for showing (19) is to maximize $\log(W)$ and to use that the maximum occurs when

$$d \log(W) = \sum_{i=1}^{\infty} [\log(A t_i^{\alpha-1} - \log(N_i)] dN_i = 0$$

and then to use Stirling's approximation for factorial. See [22] for details.

Our last step is to use this discrete distribution which we have found to approximate the continuous distribution. Let τ be a stopping time. The probability that a given bubble is still alive in the interval $[t_{i-1}, t_i)$ is given by $P(t_{i-1} \leq \tau < t_i) = \tilde{N}_i/N$. Let the sought density f satisfy

$$\frac{\tilde{N}_i}{N} = \int_{t_{i-1}}^{t_i} f(s) ds = \Delta t f(\xi)$$

by the mean value theorem, for some ξ such that $t_{i-1} \leq \xi \leq t_i$. Next let $\Delta t \to 0$ and use (19) to get

$$f(t) = \left[\frac{\beta}{\Gamma(\alpha/\beta)} \left(\frac{\alpha}{\beta K} \right)^{\alpha/\beta} \right] t^{\alpha-1} \exp\left(-\frac{\alpha}{\beta} \frac{t^\beta}{K} \right), \quad \text{for } t \geq 0 \tag{20}$$

where of course α, β and k are all positive (so that $f \geq 0$). Finally, if we make the change of variable $a = (\beta K/\alpha)^{1/\beta}$ we obtain

$$f(t) = \left(\frac{\beta}{a^\alpha \Gamma(\alpha/\beta)} \right) t^{a-1} \exp(-(t/a)^\beta) \tag{21}$$

which is a more customary expression for the density of the generalized gamma density, and the one originally proposed by Stacy [29].

Remark 2 In the derivation above, we made two major assumptions: (1) That the proportion of bubbles alive decreases theoretically with time (shifting all bubbles to start at time $t = 0$), and (2) the death rate of bubbles alive at time t_{i-1} is proportional to a power of t_i. These are reflected in Eqs. (16) and (17). These are fairly standard assumptions in the studies of lifetime distributions in survival analysis, but do they apply to bubbles? We did not analyze the data intensively in this regard, but the observation that the empirical distribution fits a generalized gamma distribution tends to lend credence to these assumptions. This may seem circular, since we are deriving the distribution as a generalized gamma, but letting the data speak means the two assumptions are largely true. This is a bit unsatisfying since we are not giving economic reasons, rather than simply empirical reasons, but it is still satisfying in a different way, because these are two minimal properties observed from the data, and yet they still imply that we have a generalized gamma, which is what we found in fact to be the case.

6 Concluding remarks

In this paper we have used the interpretation of a bubble in a stock as existing when the price process is a strict local martingale under a chosen risk neutral measure, in an incomplete market setting. By using relatively simple models of the stock price process $X = (X_t)_{0\leq t\leq T}$ of the form

$$X_t = 1 + \int_0^t \sigma(X_s)dB_s \tag{22}$$

for σ satisfying the Delbaen-Shirakawa conditions [9] given in (7), we note that X will be a strict local martingale either for all of the risk neutral measures, or none of them, so it does not matter which one is chosen to use. We also note that due to recent work of Jacod and Protter [16] this can be extended (with some important caveats) to equations of the form

$$X_t = 1 + \int_0^t \sigma(X_s, \nu_s)dB_s \tag{23}$$

where ν is a stochastic volatility process. Alternatively, see [2] or [23]. We then use the technique developed in [17], modified in this paper as explained in Sect. 2, to determine when, and for how long, a given stock price process is undergoing bubble pricing. Using a massive data set of over 3000 stocks, for a period exceeding 10 years, we then compute the empirical distribution of the lifetimes of bubbles. We find that it follows a generalized gamma distribution and we give maximum likelihood estimates for the distribution, and perform goodness of fit tests, which show that the fit is good, and the distribution we have found seems to be correct, given the caveats detailed in Sect. 4. Following the ideas of Lienhard and Meyer [22] in their 1967 paper concerning the generalized gamma, we present an explanation (via a derivation argument) as to why it is not surprising that the distribution of bubble lifetimes follows the generalized gamma distribution.

Acknowledgments We wish to thank an anonymous referee and the editor Frank Riedel for helpful comments and observations which have improved the paper. A question from Bob Jarrow led to our reporting the actual number of bubbles we observed during the 13 year period 2000 to 2014 (see Sect. 3). Supported in part by NSF Grant DMS-1308483.

References

1. Agarwal, S., Al-Saleh, J.: Generalized gamma type distribution and its hazard rate function. Commun. Stat. Theory Methods **30**(2), 309–318 (2001). doi:10.1081/STA-100002033
2. Andersen, L.B.G., Piterbarg, V.: Moment explosions in stochastic volatility models. Financ. Stoch. **11**, 29–50 (2007)
3. Barber, D.: Bayesian Reasoning and Machine Learning. Cambridge University Press, Cambridge (2012)
4. Bayraktar, E., Kardaras, C., Xing, H.: Strict local martingale deflators and valuing American call-type options. Financ. Stoch. **16**, 275–291 (2011)
5. Biagini, F., Föllmer, H., Nedelcu, S.: Shifting martingale measures and the slow birth of a bubble. Financ. Stoch. **18**, 297–326 (2014)
6. Cox, A., Hobson, D.: Local martingales, bubbles and option prices. Financ. Stoch. **9**, 477–492 (2005)
7. Cox, C.: The generalized F distribution: an umbrella for parametric survival analysis. Stat. Med. **27**, 4301–4312 (2008)
8. Delbaen, F., Schachermayer, W.: The fundamental theorem of asset pricing for unbounded stochastic processes. Math. Ann. **312**(2), 215–250 (1998)
9. Delbaen, F., Shirakawa, H.: No arbitrage condition for positive diffusion price processes. Asia-Pacific Financ. Mark. **9**, 159–168 (2002)

10. Florens-Zmirou, D.: On estimating the diffusion coefficient from discrete observations. J. Appl. Probab. **30**, 790–804 (1993)
11. Herdegen, M., Schweizer, M.: Economics-Based Financial Bubbles (and Why They Imply Strict Local Martingales), Swiss Finance Institute Research Paper No. 15-05. doi:10.2139/ssrn.2566815, SSRN: http://ssrn.com/abstract=2566815 (2015)
12. Hollebeek, T., Ho, T.S., Rabitz, H.: Constructing multidimensional molecular potential energy surfaces from AB initio data. Annu. Rev. Phys. Chem. **50**, 537–570 (1999)
13. Jacod, J.: Rates of convergence to the local time of a diffusion. Ann. l'Inst. Henri Poincaré, Sect. B **34**, 505–544 (1998)
14. Jacod, J.: Non-parametric Kernel estimation of the coefficient of a diffusion. Scand. J. Stat. **27**, 83–96 (2000)
15. Jacod, J., Protter, P.: Probability Essentials, 2nd edn. Springer, Heidelberg (2004)
16. Jacod, J., Protter, P.: Strict Local Martingale Solutions of Stochastic Differential Equations. Working Paper (2015)
17. Jarrow, R., Kchia, Y., Protter, P.: How to detect an asset bubble. SIAM J. Financ. Math. **2**, 839–865 (2011)
18. Jarrow, R., Protter, P., Shimbo, K.: Asset price bubbles in a complete market. In: Madan, D.B. (ed.) Advances in Mathematical Finance, pp. 105–130. Birkhauser, Boston (2006)
19. Jarrow, R., Protter, P., Shimbo, K.: Asset price bubbles in incomplete markets. Math. Financ. **20**, 145–185 (2010)
20. Kardaras, C., Kreher, D., Nikeghbali, A.: Strict local martingales, bubbles. Ann. Appl. Probab. **25**, 1827–1867 (2015)
21. Kotani, S.: On a condition that one dimensional diffusion processes are martingales. Memoriam Paul-André Meyer. Springer, Berlin (2006)
22. Lienhard, J.H., Meyer, P.L.: A physical basis for the generalized gamma distribution. Q. Appl. Math. **25**(3), 330–334 (1967)
23. Lions, P.L., Musiela, M.: Correlations and bounds for stochastic volatility models. Ann. Inst. Henri Poincaré, (C) Nonlinear Anal. **24**(1), 1–16 (2007)
24. Loewenstein, M., Willard, G.A.: Rational equilibrium asset-pricing bubbles in continuous trading models. J. Econ. Theory **91**, 17–58 (2000)
25. Mijatović, A., Urusov, M.: On the martingale property of certain local martingales. Probab. Theory Relat. Fields **152**, 1–30 (2012)
26. Prentice, R.L.: A log gamma model and its maximum likelihood estimation. Biometrika **61**(3), 539–544 (1974)
27. Protter, P.: A mathematical theory of financial bubbles. In: Benth, F.E., et al. (eds.) Paris-Princeton Lectures on Mathematical Finance 2013, Lecture Notes in Mathematics (2081), pp. 1–108. Springer, Cham (2013)
28. Schneikman, J., Xiong, W.: Overconfidence and speculation bubbles. J. Polit. Econ. **111**(6), 1183–1220 (2003)
29. Stacy, E.W.: A generalization of the gamma distribution. Ann. Math. Stat. **33**, 1187–1192 (1962)
30. Sommerfeld, A.: Lectures on Theoretical Physics: Thermodynamics and Statistical Mechanics, vol. 5. Acadmemic Press, New York (1964)
31. Zhang, L., Mykland, P., Aït-Sahalia, Y.: A tale of two time scales: Determining integrated volatility with noisy high-frequency data. J. Am. Stat. Assoc. **100**(472), 1394–1411 (2005). doi:10.1198/016214505000000169