

The Financial Bubble Experiment: Advanced Diagnostics and Forecasts of Bubble Terminations Volume II–Master Document

D. Sornette, R. Woodard, M. Fedorovsky, S. Reimann, H. Woodard, W.-X. Zhou
(The Financial Crisis Observatory)*
*Department of Management, Technology and Economics,
ETH Zurich, Kreuzplatz 5, CH-8032 Zurich, Switzerland*
(Dated: June 1, 2010)

This is the second installment of the Financial Bubble Experiment. Here we provide the digital fingerprint of an electronic document [1] in which we identify 7 bubbles in 7 different global assets; for 4 of these assets, we present windows of dates of the most likely ending time of each bubble. We will provide that document of the original analysis on 1 November 2010.

I. INTRODUCTION

The Financial Bubble Experiment (FBE) aims at testing the following two hypotheses:

- **Hypothesis H1:** Financial (and other) bubbles can be diagnosed in real-time before they end.
- **Hypothesis H2:** The termination of financial (and other) bubbles can be bracketed using probabilistic forecasts, with a reliability better than chance.

In a medical context, H1 corresponds to the diagnostic of cancer and H2 to the forecast of remaining life expectancy.

The motivation of the Financial Bubble Experiment finds its roots in the failure of standard approaches. Indeed, neither the academic nor professional literature provides a clear consensus for an operational definition of financial bubbles or techniques for their diagnosis in real time. Instead, the literature reflects a permeating culture that simply assumes that any forecast of a bubble's demise is inherently impossible.

Because back-testing is subjected to a host of possible biases, we propose the FBE as a real-time advanced forecast methodology that is constructed to be free, as much as possible, of all possible biases plaguing previous tests of bubbles. In particular, active researchers are constantly tweaking their procedures, so that predicted 'events' become moving targets. Only advance forecasts can be free of data-snooping and other statistical biases of ex-post tests. The FBE aims at rigorously testing bubble predictability using methods developed in our group and by other scholars over the last decade. The main concepts and techniques used for the FBE have been documented in numerous papers [2–6] and the book [7].

The FCO research team is currently developing and testing novel estimations methods that will be progressively implemented in future releases.

In the FBE, we propose a new method of delivering our forecasts where the results are revealed only after the predicted event has passed but where the original date when we produced these same results can be publicly, digitally authenticated.

Since our science and techniques involve forecasting, the best test of a forecast is to publicize it and wait to see how accurate it is, whether the wait involves days, weeks or months (we rarely make forecasts for longer time scales). We will do this and at the same time we want to delay the unveiling of our results until after the forecasted event has passed to avoid potential issues of liability, ethics and speculation. Also, we think that a full set of results showing multiple forecasts all at once is more revealing of the quality of our current methods than would be a trickle of one such forecast every month or so. We also want to address the obvious criticism of cherry picking successful forecasts, as explained below. In order to be convincing, our experiment has to report all cases, be they successes or failures.

The digital fingerprint of our first set of bubble forecasts was released on 2 November 2009 (with a hash update on 6 November 2009). We added a new bubble forecast on 23 December 2009. The original forecasts and post-analysis were presented publicly on 3 May 2010 and uploaded to the arXiv server on 14 May 2010. All versions are available at [8].

*Electronic address: dsornette@ethz.ch

This second set of forecasts presents the methodology described in [8] and the digital fingerprint of a single document that identifies and analyses 7 current asset bubbles (H1). For 4 of those 7 bubbles, the document also provides windows of dates of the most likely ending time of each bubble. We will provide that original document of the analysis on 1 November 2010.

II. DESCRIPTION OF THE METHODOLOGY OF THE FINANCIAL BUBBLE EXPERIMENT

Our method for this experiment is the following:

- We choose a series of dates with a fixed periodicity on which we will reveal our forecasts and make these dates public by immediately posting them on our University web site and on the first version of our main publication, which we describe below. Specifically, our first publication of the forecasts was issued on 3 May 2010, with successive deliveries every 6 months. The forecasts of the current document will be presented on 1 November 2010. However, we keep open the option of changing the periodicity of the future deliveries as the experiment unfolds and we learn from it and from feedback of the scientific community.
- We then continue our current research involving analysis of thousands of global financial time series.
- When we have a confident forecast, we summarize it in a simple `.pdf` document.
- We do not make this document public. Instead, we make its digital fingerprint public. We generate two digital fingerprints for each document, with the publicly available 256 and 512 bit versions of the SHA-2 hash algorithm [11] [12]. This creates two strings of letters and numbers that are unique to this file. Any change at all in the contents of this file will result in different SHA-2 signatures.
- We create the first version of our main document (this one), containing a brief description of our theory and methods, the SHA-2 hashes of our forecast and the date (1 November 2010) on which we will make the original `.pdf` document public.
- We upload this main ‘meta’ document to <http://arxiv.org>. This makes public our experiment and the SHA-2 hashes of our forecast. In addition, it generates an independent timestamp documenting the date on which we made (or at least uploaded) our forecast. arxiv.org automatically places the date of when the document was first placed on its server as ‘v1’ (version 1). It is important for the integrity of the experiment that this date is documented by a trusted third party.
- We continue our research until we find our next confident forecast. We again put the forecast results in a `.pdf` document and generate the SHA-2 hashes. We now update our master document with the date and digital fingerprint of this new forecast and upload this latest version of the master document to arxiv.org. The server will call this ‘v2’ (version 2) of the same document while keeping ‘v1’ publicly available as a way to ensure integrity of the experiment (i.e., to ensure that we do not modify the SHA-2 hashes in the original document). Again, ‘v2’ has a timestamp created by arxiv.org.
- Notice that each new version contains the previous SHA-2 signatures, so that in the end there will be a list of dates of publication and associated SHA-2 signatures.
- We continue this protocol until the future date (1 November 2010) at which time we upload our final version of the master document. For this final version, we include the URL of a web site where the `.pdf` documents of all of our past forecasts can be downloaded and independently checked for consistent SHA-2 hashes. For convenience, we will include a summary of all of our forecasts in this final document.

Note that the above method implies two aspects of the same important check to the integrity of our experiment:

1. We will reveal all forecasts, be they successful or not.
2. We will not simply ‘cherry-pick’ the results that we would want the community to see (with a few token, possibly, bad results). We do not have another simultaneous outlet where we are running a similar experiment, since arxiv.org is a very visible international platform.

III. BACKGROUND AND THEORY

Our theories of financial bubbles and crashes have been well-researched and documented over the past 15 years in many papers and books. We refer the reader to the Bibliography. In particular, broad overviews can be found in [2–6]. In short, our theories are based on positive feedback on the growth rate of an asset’s price by price, return and other financial and economic variables, which lead to faster-than-exponential (power law) growth. The positive feedback is partially due to imitation and herding among humans, who are actively trading the asset. This signature is quantitatively identified in a time series by a faster-than-exponential power law component, the existence of increasing low-frequency volatility, these two ingredients occurring either in isolation or simultaneously with varying relative amplitudes. A convenient representation has been found to be the existence of a power law growth decorated by oscillations in the logarithm of time. The simplest mathematical embodiment is obtained as the first order expansion of the log-periodic power law (LPPL) model and is shown in Eq. (1):

$$\ln P = A + B|t - t_c|^\alpha + C|t - t_c|^\alpha \cos[\omega \ln |t - t_c| + \phi] \quad (1)$$

where P is the price of the asset and t is time. There are 7 parameters in this nonlinear equation. Our past work has led to the hypothesis that the LPPL signals can be useful precursors to an ending (change of regime) of the bubble, either in a crash or a less-dramatic leveling off of the growth.

IV. METHODS

A. Bubble identification

As are our theories, our methods are documented elsewhere so we only briefly mention the general technique so that the forecasts that we make public can be better understood. In short, we scan thousands of financial time series each week and identify regions in the series that are well-fit by Eq. (1). We divide each time series into sub-series defined by start and end times, t_1 and t_2 and then fit each sub-series (t_1, t_2). We choose $\max(t_2)$ as the date of the most recent available observation. Many sub-series are created according to the following parameters: $dt_1 = dt_2 = 7$ days, $\min(t_2 - t_1) = 91$ days and $\max(t_2 - t_1) = 1092$ days.

After filtering all fits with an appropriate range of parameters, we select those assets that have the strongest LPPL signatures. To improve statistics, we can calculate the residues between the model and the observations and use the residues to create 10 synthetic datasets (bootstraps) that have similar statistics as the original time series. We fit Eq. (1) to the synthetic data and then extrapolate this entire ensemble of LPPL models to six months beyond our last observation. One of the parameters in the LPPL equation is the “crash” time t_c , which represents the most probable time of the end of the bubble and change of regime. We identify the 20%/80% and 5%/95% quantiles of t_c of the fits of the ensemble consisting of original fits and bootstrap fits. These two sets of quantiles, the date of the last observation and the number of fits in the ensemble are published in our forecasts.

B. Post-analysis

Once the .pdf documents with the full description of the forecasts are made public, the question arises as how to evaluate the quality of the diagnostics and how these results help falsify the two hypotheses? In a nutshell, the problem boils down to qualifying (and quantifying) what is meant by (i) a successful diagnostic of the existence of a bubble and (ii) a successful forecast of the termination of the bubble. In the end, one would like to develop statistical tests to falsify the two hypotheses stated above, using the track record that the present financial bubble experiment has the aim to construct. For instance, Chapter 9 of (Sornette, 2003) suggests a number of options, including the “statistical roulette”, Bayesian inference and error diagrams. Our main goal with this FBE is to timestamp our forecasts as we simultaneously continue our search for adequate measures to qualify the quality of our forecasts.

This quantification is an active, ongoing subset of our research. For our first set of forecasts, we quantified the quality of the forecasts with three measures that were presented in the final analysis:

- **Drawdown analysis:** Drawdown analysis simply identifies the largest drawdown observed between t_2 (date of forecast) and the date of the public ‘unveiling’ of the original forecasts. That is, we identify the largest drawdown in all available data after t_2 . A drawdown is simply defined as the largest peak-to-trough drop in price in a given region.

- **Fraction of up days in a running window:** We calculate one day close-to-close returns for each asset and mark them as positive (up) or non-positive (zero or down). The ratio of up days relative to the sum of up and down days in a running window of 30, 60 or 90 days is plotted on top of the price observations.
- **Derivative of observations:** Another measure of the change of regime is provided by an estimation of the local growth rate. We use the Savitzky-Golay smoothing algorithm to calculate the first derivative of the observations, using a third order polynomial fit centered within windows of 120 and 180 days.

We are developing other measures that will be used in the analysis to be presented on 1 November 2010.

V. BUBBLE FORECASTS

The checksums of the analysis document [1] that contains the names of the 7 assets are shown in Table I. This document showing all 7 assets and 4 forecasts, as well as analysis of each identified bubble, will be uploaded to <http://www.er.ethz.ch/fco/> on 1 November 2010.

Document name	
SHA256SUM	d8b1345dca3a1ff3952d5f8f74595b83accb7b8bcefd163a7552512b5b4cda8e
SHA512SUM	3f529ca27ea8f06934b3ecb01f07b08d648f3d98dbc1253ebb70e8c52a368a9d441f641af0c621f208b509a102caf75337ce321e732d9e8c6cd584434f50880

TABLE I: Checksums of Financial Bubble Experiment forecast document.

-
- [1] D. Sornette, R. Woodard, M. Fedorovsky, S. Reimann, H. Woodard, and W.-X. Zhou (The Financial Crisis Observatory), “The Financial Bubble Experiment: advanced diagnostics and forecasts of bubble terminations, Vol. II–assets document,” (2010).
- [2] Z.-Q. Jiang, W.-X. Zhou, D. Sornette, R. Woodard, K. Bastiaensen, and P. Cauwels, *Journal of Economic Behavior & Organization* **74**, 149 (2010), ISSN 0167-2681, <http://www.sciencedirect.com/science/article/B6>
- [3] A. Johansen, D. Sornette, and O. Ledoit, *J. Risk* **1**, 5 (1999).
- [4] A. Johansen and D. Sornette, *Brussels Economic Review* **49** (2006), (<http://arXiv.org/abs/cond-mat/0210509>).
- [5] D. Sornette and A. Johansen, *Quant. Financ.* **1**, 452 (2001).
- [6] D. Sornette and W.-X. Zhou, *Int. J. Forecast.* **22**, 153 (2006).
- [7] D. Sornette, *Why Stock Markets Crash: Critical Events in Complex Financial Systems* (Princeton University Press, Princeton, 2003).
- [8] D. Sornette, R. Woodard, M. Fedorovsky, S. Reimann, H. Woodard, and W.-X. Zhou (The Financial Crisis Observatory), “The Financial Bubble Experiment: advanced diagnostics and forecasts of bubble terminations,” (2009), <http://arxiv.org/abs/0911.0454>.
- [9] http://en.wikipedia.org/wiki/SHA_hash_functions.
- [10] <http://eprint.iacr.org/2008/270.pdf>.
- [11] http://en.wikipedia.org/wiki/SHA_hash_functions
- [12] <http://eprint.iacr.org/2008/270.pdf>