MODEL-FREE DETECTION OF A SPECULATIVE ASSET BUBBLE: EVIDENCE FROM THE WORLD MARKET FOR SUPERSTAR WINES

TOR N. TOLHURST

University of California, Davis

Abstract. Bubbles occur when an asset price deviates from its fundamental value. Economists have proven asset bubbles are consistent with neoclassical theory and can arise in a variety of laboratory settings; yet cogent, direct empirical evidence of a bubble has remained elusive. The challenge for the empiricist is that fundamental values are necessarily unobservable. Rather than estimate fundamental values or impose identical fundamental values on two closely related products, I assume the price difference between two closely related products is probabilistically bounded. By doing so I construct a direct test that is consistent with all structural models, thus avoiding the endemic joint hypothesis testing problem. I find strong evidence of a bubble in the price of the Bordeaux wine Lafite Rothschild, relative to other fine wines exchanged on global secondary markets. Assuming two moments exist so Cheybychev's inequality applies, the magnitude of observed deviations is such that the probability of a random, false detection is less than 0.1 percent.

Keywords. Asset pricing, rational bubbles, probability inequalities, wine.

E-mail address: tntolhurst@ucdavis.edu.
Date: September 28, 2018.
I thank Julian Alston, Jean-Marie Cardebat, Dalia Ghanem, Rachael Goodhue, Rich Sexton, Shu Shen, Aaron Smith, and Jeffrey Williams, as well as seminar and conference participants at the International Network for Economic Research 2017, Australasian Agricultural and Resource Economics Society 2018, and Agricultural and Applied Economics Association 2018 meetings, for their helpful comments, discussion, and suggestions. I thank the London International Vinter's Exchange for providing data.

1. Introduction

Asset bubbles damage and scar the real economy: bubbles distort price signals and thus interfere with the coordination function of prices; bubble fervor promotes recklessness and facilitates fraud; financial crises are often precipitated by bubble bursts; and bubbles otherwise increase uncertainty and erode trust in markets. Economic agents also behave in intriguing ways during bubbles. This intriguing behavior—reflected in the eloquent titles of prominent works, from *Extraordinary Popular Delusions and the Madness of Crowds* (Mackay, 1852) to *Irrational Exuberance* (Shiller, 2015)—has stimulated intellectual work in behavioral economics and the efficient market hypothesis. However, empirical evidence of an asset bubble has been elusive: a recent article by Giglio, Maggiori, and Stroebel (2016) developed a model-free bubble test to avoid the joint hypothesis test problem endemic to the empirical literature, but failed to find evidence of a bubble in two asset markets, the United Kingdom and Singapore housing markets, commonly thought to have experienced bubbles.

In this manuscript I use a test similar to, but more general than, Giglio, Maggiori, and Stroebel's to look for a bubble in a particular asset class, fine wines exchanged on global secondary markets. I find compelling evidence of a bubble. In the wake of the 2008 financial crisis, fine wine received attention as a promising alternative investment in both the popular press (e.g. Lister, 2011; Rabinovitch, 2010) and academia (e.g. Cardebat and Figuet, 2010; Faye, Le Fur, and Prat, 2015; Le Fur, Ameur, and Faye, 2016). Beginning in 2010, some industry insiders began to conjecture there was a bubble in one particular Bordeaux wine, *Chateau Lafite-Rothschild* (e.g. Stimpfig, 2010; Temperton, 2011; Authers, 2012). It is easy to see why in Figure 1, which plots monthly prices of first growth Bordeaux wines and clearly illustrates a boom in Lafite prices. These wines are widely considered to be of equivalent quality. They are the only wines classified as first growth in Bordeaux: Lafite, Margaux, Haut Brion, and Latour since 1855 and Mouton since 1973. Prior to 2006, their prices were roughly equivalent: June 2005, they all cost roughly £75 per bottle with Margaux £79,

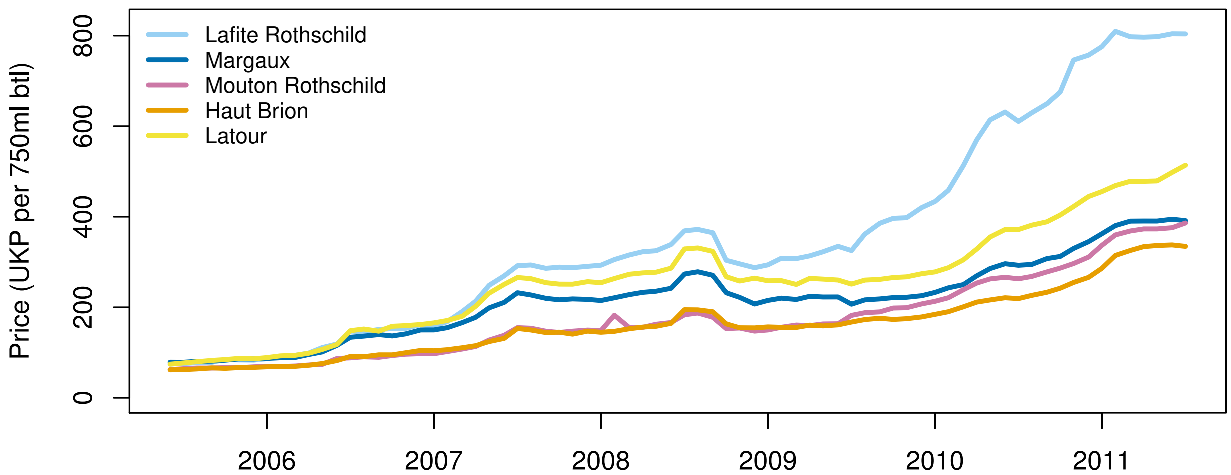

FIGURE 1. Monthly average price of a ten-vintage portfolio of the five *premier cru* Bordeaux wines, June 2005 – July 2011.

Lafite £75, Latour £74, Mouton £62, and Haut Brion £62.[1] Prices moved mostly together until June 2009. Then Lafite boomed: at the peak in February 2011, Lafite cost £810 per bottle, compared to £469 for Latour, £380 for Margaux, £360 for Mouton, and £314 for Haut Brion. Returns to Lafite over the nearly 70-month boom were astounding: 51.3 percent annualized for 5.75 years. As would be expected in a bubble, the relative price of Lafite then dropped precipitously. By January 2017, Lafite was down to £490 and again aligned with other first growth prices: Latour £495, Margaux £356, Mouton £350, and Haut Brion £337.[2] The ultimate purpose of this manuscript is to to determine whether or not the price changes in Lafite provide direct evidence of a bubble. I will argue they do.

The possibility of bubbles has been well-established by economists.[3] Rational bubble theory shows explosive solutions for the market-clearing price are compatible with a neo-classical setup (Samuelson, 1958; Diamond, 1965; Blanchard, 1979; Blanchard and Watson,

[1]All prices are real in 2015 terms, deflated by the UK CPI reported by the Great Britain Office for National Statistics. Available at: `https://www.ons.gov.uk/economy/inflationandpriceindices/timeseries/d7bt/mm23`.

[2]The annualized rate of return from June 2005 to Jan. 2017 was a healthy but more reasonable 17.5 percent.

[3]The literature on bubbles is of course vast and rich. Note that while the literature demarcates a number of "bubble-like" behaviors—such as sunspots (Le Roy and Porter, 1981), fads (Camerer, 1989), information bubbles (Bikhchandani, Hirshleifer, and Welch, 1992), and irrational bubbles (Le Roy, 2004)—I limit the scope of the paper to focus on the leading conception, the rational bubble. For more complete reviews of the literature emphasizing theory see Camerer (1989) and Le Roy (2004) or for empirics West (1988), Flood and Hodrick (1990) and Gürkaynak (2008). For review-type discussions of bubbles in the context of the 2007–2008 financial crisis, see especially Bhattacharya and Yu (2008), O'Hara (2008), and Carvalho, Martin, and Ventura (2012).

1982; Tirole, 1982, 1985; Froot and Obstfeld, 1991). The mechanism behind rational bubbles is a form of self-fulfilling prophecy: agents expect a bubble return, which leads to a realized bubble return, which leads to further expected bubble returns, continued participation in the bubble trade, and so on. Experimental evidence, reviewed by Palan (2013), has largely supported the theoretical existence of rational bubbles (Smith, Suchanek, and Williams, 1988; Lei, Noussair, and Plott, 2001; Dufwenberg, Lindqvist, and Moore, 2005; Kirchler, Huber, and Stöckl, 2012; Palfrey and Wang, 2012; Moinas and Pouget, 2013; Eckel and Füllbrunn, 2015; Andrade, Odean, and Lin, 2015). Bubbles arise in laboratory settings even with knowledgeable traders who experienced bubbles in the past (Hussam, Porter, and Smith, 2008). Despite the common association of bubbles with economically irrational behavior, economists have shown bubbles are possible with rational, full-informed agents. Rational bubbles are the most important conceptualization of bubbles and the theoretical cornerstone of the empirical bubble literature.

To illustrate the relationship between rational bubbles and the empirical bubble literature, let us consider a simple version following Flood and Hodrick (1990). Note rational bubble theory does not rely on the assumptions that follow; for example, agents can recognize the bubble will eventually burst (Camerer, 1989), face multiple sources of uncertainty (Doblas-Madrid, 2012), or endogeneous credit constraints (Miao and Wang, 2018). Let P_t be the market-clearing asset price, and D_t its dividend, in period t. Assume homogeneous risk-neutral agents share a time-varying information set, $\mathcal{I}_t$, and a time-invariant, positively-valued discount rate r. The bubble component of the asset price is denoted B_t and nonnegative given free-disposal. There is an asset bubble when $B_t > 0$. The first-order equilibrium condition of the utility-maximization problem is:

$$P_t = \sum_{i=1}^{\infty} \mathbb{E}[D_{t+i}|\mathcal{I}_t](1+r)^{-i} + B_t, \tag{1}$$

which can be satisfied by a positive bubble component value, $B_t > 0$, provided:

$$B_t = \mathbb{E}[B_{t+1}|\mathcal{I}_t](1+r)^{-i}. \tag{2}$$

That is, an asset bubble can satisfy the equilibrium condition (1) provided the bubble component is expected to grow proportional to the discount rate. This is a "rational bubble." The term

$$F_t := \sum_{i=1}^{\infty} \mathbb{E}[D_{t+i}|\mathcal{I}_t](1+r)^{-i}$$

is referred to as the asset's fundamental value.

The fundamental value distinguishes the two classes of empirical bubble tests: direct and indirect. Direct tests incorporate B_t into the hypothesis test using, for example, coefficient restrictions (Flood and Garber, 1980) or two-step, Hausman-style tests (West, 1987). While the ideal direct test would provide the most credible evidence for a bubble, direct tests are plagued by the fact that F_t is unobservable—F_t must be modeled and apparent evidence of a bubble could be reinterpreted as unobserved changes in fundamentals. Indirect tests circumvent the challenge of unobservable F_t by examining observed prices P_t for statistically anomalous behavior (e.g. explosive unit roots). Given the difficulties with unobservable fundamentals and advancements in statistical methods, indirect tests have received the most attention over the past 30-odd years.[4] Prominent examples of indirect tests include variance-bounds (Shiller, 1981; Le Roy and Porter, 1981), stationarity (Hamilton and Whiteman, 1985; Kleidon, 1986), cointegration (Diba and Grossman, 1987, 1988a,b), vector autoregression (Campbell and Shiller, 1987, 1988a,b), Markov switching (Hall, Psaradakis, and Sola, 1999; Al-Anaswah and Wilfling, 2011), and most recently, right-sided unit roots (Phillips and Yu, 2011; Phillips, Wu, and Yu, 2011; Phillips, Shi, and Yu, 2014, 2015a,b) based on the asymptotic theory of Phillips and Magdalinos (2007). A recent application of the Phillips *et al.* approach is to the cryptocurrency Bitcoin (Cheung, Roca, and Su, 2015). However, rejection of the null hypothesis in an indirect test is not equivalent to detection of a rational bubble: Giglio, Maggiori, and Stroebel (2016) find statistical evidence of a bubble using the indirect Phillips *et al.* test, but find a precisely estimated no-bubble null result on the same data with their direct test (which overcomes the joint hypothesis testing problem).

[4]Empirical applications of heterogeneous agent models (Harrison and Kreps, 1978), reviewed in Xiong (2013), are increasingly prominent in the literature but subject to the same joint hypothesis testing pitfalls.

I take a similar approach to Giglio, Maggiori, and Stroebel by using closely related assets to conduct a model-free, direct bubble test. I compare Lafite to a set (and various subsets) of closely related fine Bordeaux wines. One benchmark specification is a composite wine made up of the four other Bordeaux first growth wines, which are widely agreed to be of similar style and quality to Lafite. A second benchmark specification uses data-driven synthetic control methods to find weights on the 46 other fine Bordeaux wines which minimize the mean squared prediction error in a pre-bubble period. While by any measure the composite and synthetic wine are closely related to Lafite, I assume their fundamental values are bounded rather than identical, a weaker assumption than Giglio, Maggiori, and Stroebel's, which makes it more difficult to detect a bubble. To demonstrate the strength of my finding, I take a number of additional steps which also make it more difficult to find a statistically significant asset bubble. Whereas Giglio, Maggiori, and Stroebel fail to reject $B = 0$ using a t-test, I use Cheybychev's inequality to probabilistically bound the fundamental values: a five and ten percent statistical threshold under Cheybychev's would be equivalent to t-test p-values of 9.5×10^{-5} and 0.0023, respectively. Assuming only that the data-generating process has at least two finite moments so that Cheybychev's inequality applies, I find multiple observations that are statistically anomalous. Even with a recursively-updating test which allows for the fundamental values and their relationship to be time-varying, I detect three consecutive bubble periods in late-2009 where the magnitude of observed deviations is more than three standard deviations away from the mean—implying a false detection probability of less than 0.1 percent. Note that if I used less conservative threshold, such as the mean difference in fundamental values (Baur and Glover, 2014), the bubble would be detected with time-varying fundamentals for more than 24 consecutive periods. Thus, I argue my findings provide the first compelling, direct, model-free empirical evidence of an asset bubble.

The remainder of this manuscript proceeds as follows. Section 2 details the empirical test, section 3 describes the data, and section 4 provides the empirical evidence for a bubble in Lafite. The penultimate section discusses alternative explanations for the observed price anomolies, all of which are found to be unsatisfactory. For example, while these wines are anything but ordinary, they are widely available in top-tier wine retailers and on global auction markets. The final section concludes.

2. Empirical Test

I divide this section into two subsections. The first explains the empirical test in light of rational bubble theory and the model-free class of tests advanced by Giglio, Maggiori, and Stroebel (2016). The second describes empirical considerations for the test, which uses parity bounds (Baulch, 1997) within a recursive rolling-window (Phillips, Wu, and Yu, 2011).

2.1. **Theoretical Foundation.** Consider a more general version of the rational bubble theory outlined above. Allow for a time-varying stochastic discount rate, r_t, which can be concisely expressed as the discount factor $\delta_{t,t+i} := \prod_{j=0}^{i-1}(1 + r_j)^{-1}$. The stochastic nature of the discount rate moves the term $\delta_{t,t+i}$ inside the expectation, so that (1) becomes:

$$P_t = \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} D_{t+i} | \mathcal{I}_t] + B_t \tag{1'}$$

and the fundamental component becomes $F_t := \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} D_{t+i} | \mathcal{I}_t]$.

Le Roy (2004) highlights three implications of the rational bubble model: (i) the value of a firm which does not, and is never expected to, pay dividends is entirely a bubble; (ii) on freely-disposable assets, the bubble component must be nonnegative, $B_t \geq 0$; and (iii) rational bubbles of this sort require a model with infinite periods—otherwise backward induction, combined with the observation that $B_T = 0$ at the finite terminus period T, suggests $B_t = 0$ for all t. Implication (iii), first set out as the transversality condition in the dynamic rational expectations equilibrium of Tirole (1982), forms the basis for the Giglio, Maggiori, and Stroebel (2016) model-free test. The practical challenge of this approach is measuring the value of financial claims very far in the future, which Giglio, Maggiori, and Stroebel circumvent by selecting similar assets governed by two different contract terms. Specifically, Giglio, Maggiori, and Stroebel attribute the difference in prices for houses under extremely-long maturity (more than 700 years) leasehold and freehold contracts to B_t; the key difference between the contracts being that leaseholds are finite-maturity and freeholds infinite-maturity.[5]

2.1.1. *Model-Free Test: Equal Fundamental Values and Different Time-Horizons.* Consider two otherwise similar assets. One of the assets is governed under a finite-term lease, the other is governed by an infinite lease; denote the price of the infinite-term asset P^{∞} and the finite-term asset P^{T}. Applying (1′) to the infinite- and finite-term assets, we have:

$$(1'') \qquad P_t^{\infty} = \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} D_{t+i}^{\infty} | \mathcal{I}_t] + \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} B_{t+i} | \mathcal{I}_t],$$

$$\text{and} \qquad P_t^{T} = \sum_{i=1}^{T} \mathbb{E}[\delta_{t,t+i} D_{t+i}^{T} | \mathcal{I}_t] + \sum_{i=1}^{T} \mathbb{E}[\delta_{t,t+i} b_{t+i} | \mathcal{I}_t].$$

Under the transversality condition, it must be that $\sum_{i=1}^{T} \mathbb{E}[\delta_{t,t+i} b_{t+i} | \mathcal{I}_t] = 0$. Let us further assume T is large (they assume $T \geq 700$ in their empirical application) and, other than their time-horizon, the products are otherwise equivalent. Specifically, the assets are sufficiently similar to share the same relevant information sets and discount factors, and the present value of the stream of future dividends are equivalent (given the present value of payments infinitely far in the future is zero). In other words, the two assets have *equivalent fundamental values*:

$$(3) \qquad \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} D_{t+i}^{\infty} | \mathcal{I}_t] = \sum_{i=1}^{T} \mathbb{E}[\delta_{t,t+i} D_{t+i}^{T} | \mathcal{I}_t]$$

Finally, define $B_t := \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} B_{t+i} | \mathcal{I}_t]$ as the bubble component value of the asset. When the assets have equivalent fundamental values, the price difference between the two assets directly gives B_t in present-value terms: $P_t^{\infty} - P_t^{T} = B_t$. A no-bubble null hypothesis can then be tested directly:

$$H_0 : \; P_t^{\infty} - P_t^{T} = B_t = 0.$$

Further, this direct test is model-free: all structural valuation models agree $B_t = 0$ under the null and no model is imposed to estimate the unobservable fundamental values. Instead,

[5]Rational bubbles can arise in finite-period games—even when all agents know the asset will eventually be worthless—under asymmetric information (Conlon, 2004).

the assumptions of rational bubble theory were used to find two assets which have equivalent fundamental values. As there is no specified structural model, the joint hypothesis problem is circumvented (c.f. Giglio, Maggiori, and Stroebel, 2016, pg. 1048).[6]

While the test is direct and circumvents the problem of joint hypothesis testing, it has limitations. First, as with many indirect and direct bubble tests, the model-free approach does not necessarily identify different types of bubbles (e.g. rational versus irrational bubble). It is a direct test which is consistent with all structural valuation models of the fundamental value—an improvement over prior direct tests—and consistent with rational bubble theory, but it may also be consistent with other bubble theories. For example, violations of the transversality condition used in Giglio, Maggiori, and Stroebel are also consistent with bubbles in a myopic-rational-expectations (Tirole, 1982) or heterogeneous belief (Harrison and Kreps, 1978) model. If distinguishing between the different ways bubbles can arise is important, say if a policymaker is trying to deflate or control a possible bubble, the model-free approach is not appropriate.[7] Second, the approach cannot detect a bubble in both assets. Third, few assets (if any) are governed by similarly delineated contract terms as housing is and it is not readily obvious how to apply the test to other assets. While I propose one way to apply the approach to other assets below, this limitation stands in stark contrast to indirect tests, which require only an asset of interest which has a sufficiently long price series to be tested. Fourth, and perhaps more importantly, identification of fundamental values depends on assumptions that cannot be tested directly. These four limitations also apply to my model-free test.

[6]An important practical consideration relevant to Giglio, Maggiori, and Stroebel (2016), but not for my empirical application, is how to ensure price differences are calculated using only comparable assets—for example, one bedroom freehold units to one bedroom leasehold units—over multiple attributes. To do so, they condition housing prices on attributes using hedonic matching methods. Obviously estimating (3) conditional on observable attributes implies an empirical model is imposed, so the use of the label "model-free" can be confusing. They demarcate "model" as referring to a structural valuation model for the fundamental values, while an "empirical specification" is used to condition housing prices on attributes (c.f. Giglio, Maggiori, and Stroebel, 2016, pg. 1064).

[7]Delineating the underlying nature of bubbles is not a widely pursued question in the literature and outside the scope of this essay.

2.1.2. *Model-Free Test: Generalized Setup.* The model-free test could be applied to two similar products which differ in ways other than their time-horizons. Consider two products with related fundamental values and infinite time-horizons. The price of one of the products, call it the superstar product, is denoted by P^*. The price of the second product is denoted by P°. Assume both products could be subject to a common bubble b_t and the superstar product could have an additional bubble component term B_t. This gives a setup very similar to (1″):

$$(1''') \qquad P_t^* = \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} D_{t+i}^* | \mathcal{I}_t] + \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} b_{t+i} | \mathcal{I}_t] + \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} B_{t+i} | \mathcal{I}_t],$$

$$\text{and} \qquad P_t^\circ = \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} D_{t+i}^\circ | \mathcal{I}_t] + \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} b_{t+i} | \mathcal{I}_t],$$

the main differences being the upper bound on the summation terms (now all infinite) and the inclusion of the common bubble terms. For convenience, I introduce the notation:

$$\Delta P_t := P_t^* - P_t^\circ,$$

$$\Delta F_t := \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} D_{t+i}^* | \mathcal{I}_t] - \sum_{i=1}^{\infty} \mathbb{E}[\delta_{t,t+i} D_{t+i}^\circ | \mathcal{I}_t],$$

so that $\Delta P_t = \Delta F_t + B_t$. Note that in ΔP_t the common bubble terms cancel out, making explicit the limitation that the approach cannot detect a bubble common to both assets. Also note that assuming $\Delta F_t = 0$ and testing if B_t is statistically different from zero corresponds to Giglio, Maggiori, and Stroebel (2016). When the fundamental value is subject to parity bounds, as I introduce next, $\Delta F_t = 0$ is a special case.

2.1.3. *Model-Free Test: Parity-Bounded Fundamental Values.* Assume ΔF_t is a random variable with an unknown distribution with at least two finite moments concentrated between a lower and upper bound, $\underline{\Delta F_t}$ and $\overline{\Delta F_t}$, respectively. Fundamental values may vary, but deviations from $\mathbb{E}[\Delta F_t]$ outside of these bounds occur with probability no greater than α (the probability of false detection, analogous to size in a statistical test). Under a no bubble null, $\Delta P_t = \Delta F_t$, so:

$$\text{if} \begin{cases} \Delta P_t \in [\underline{\Delta F_t}, \overline{\Delta F_t}] & \text{there is no bubble,} \\ \Delta P_t \notin [\underline{\Delta F_t}, \overline{\Delta F_t}] & \text{there is a bubble.} \end{cases}$$

That is, a bubble occurs if ΔP_t is observed outside of the parity bounds.

The parity bounds idea, and more informative bounds, come from the Baulch (1997) spatial market integration test. Consider a product-space analog of Baulch (1997) with our two related products, specifically related in the sense that one product can be transformed into the other at cost π_t^*. Product transformation need not be physical: perhaps the most compelling explanation in the context of products is that π_t^* is the premium consumers are willing to pay for superstar status (or compensation they are willing to accept for consuming the non-superstar wine in place of the superstar wine). By an arbitrage argument, Baulch (1997) would assert the market for the superstar and non-superstar good are integrated provided: (i) when $P_t^\circ + \pi_t^* \geq P_t^*$, no transformation occurs, but (ii) when $P_t^\circ + \pi_t^* < P_t^*$ the non-superstar product is transformed into a superstar product until $P_t^\circ + \pi_t^* = P_t^*$. The price of the two related products can vary independently within the bounds $0 \leq \Delta P_t \leq \pi_t^*$ in the absence of a market anomaly *whilst remaining at parity*, hence parity bound. Therefore I assume $\underline{\Delta F_t} = 0$ and $\overline{\Delta F_t} = \pi_t^*$, where π_t^* to be defined later will come from a concentration inequality and the zero lower bound follows from free-disposal implying $B_t \geq 0$, assuming consumers never pay less for superstar status.[8]

2.1.4. *Applicability to Wine Markets.* Wines are not traditional financial assets, however rational bubble theory can be applied. According to Tirole (1982, 1985), rational bubble

[8]Baulch (1997) proposed the parity bounds model in the context of spatial arbitrage—specifically, the movement of a homogeneous good across two geographically distant markets—and the challenges associated with tests for market integration. Given the unrelated context, many details of Baulch (1997) are not pertinent. Perhaps the most important difference between Baulch (1997) and the proposed approach is that the former uses exogenous information to estimate π_t^*. Specifically, Baulch uses data to estimate transport costs in one period, then extrapolates single-period transport costs into a time-series using a consumer price index, and finds probabilistic bounds using parametric estimates. Given difficulties with specifying the information required to estimate a threshold, I will use previous data and non-restrictive, nonparametric assumptions on the data-generating process to define the upper-bound.

theory applies to assets that have resale value and are inelastically supplied in the short-run—all true for fine Bordeaux wines. All the data I use comes from secondary markets, so clearly the assets have resale value. Wines are inelastically supplied in the short-run as output is constrained to yields on a particular estate in a particular vintage. Expansions to production are restricted by law. Bordeaux estates can expand their production by purchasing neighboring vineyards, though inventory from previous vintages acquired in the acquisition cannot be sold under the acquiring estate's label. Alternatively, a Bordeaux estate could plant new vineyards on a contiguous property, but properties must be within the fixed area defined by appellation laws and even then they must wait multiple (usually seven plus) years for the vineyard to achieve quality production.

In light of the Le Roy (2004) criteria, wine does not pay pecuniary dividends. Dividends and discount rates for holding wine may be nonpecuniary: at least a portion of the divided value might reflect the increased utility from consuming the wine at a later period or the appreciation in quality of the wine due to aging, while at least a portion of the discount rate could represent forebearance cost of deferring consumption from one period to the next. Even if the fundamental value of wine was entirely nonpecuniary, nonpecuniary value is not explicitly ruled out in rational bubble theory and regardless, nonpecuniary value can be readily exchanged for money on markets. Wine is technically not infinitely-lived, but practically it may be considered so. Ultra-premium wines such as first growth Bordeaux retain their value for decades. Using an aggregation database, a cursory search retrieves 17 vintages of Lafite prior to 1900 going back to 1812 and vintages going back to 1892 for Haut Brion, 1884 for Margaux, 1865 for Latour, and 1853 for Mouton. Examples beyond the first growths include vintages going back to 1825 for Yquem and 1849 for Ausone.[9] All of these very old wines are valued at thousands of dollars.[10] Thus, I feel confident assuming wines of this quality-level have practically infinite lives. Besides, the wineries under study (especially

[9] `www.wine-searcher.com` on Aug. 16, 2018.

[10] Two other notable examples of very old (perhaps dating back to the 18th century) and valuable wines are the controversial Jefferson bottles (see: `https://www.newyorker.com/magazine/2007/09/03/the-jefferson-bottles`), and champagne found in a Baltic Sea shipwreck (see: `https://www.theguardian.com/world/2010/jul/18/champagne-found-sea-oldest-vintage`).

the five premier cru) are centuries old and, as I use a rolling ten-vintage portfolio, I would argue I testing for a bubble in the "brands" (e.g. Lafite, Latour, etc.) rather than a specific wine-vintage pair. In the case of Lafite and Latour, for example, their names have been prominent since the mid-17th and early-18th century, respectfully. There is no reason for consumers to expect these names to go away within the foreseeable future.

2.2. **Empirical Considerations.** We are given $(\Delta P_1, \Delta P_2, \ldots, \Delta P_T)$, a sequence of price differences, and want to test for a bubble in the observation ΔP_{T_1} where $T_1 \leq T$. Assume a subset of the observations $(\Delta P_{t_0}, \ldots, \Delta P_{T_0})$ with $1 \leq t_0 < T_1$ and $T_0 < T_1$ is known to be bubble-free. A bubble is detected when ΔP_{T_1} is above the parity bound π_t^*. The test is conducted by using a concentration inequality to get π_t^* based on parameter estimates from $(\Delta P_{t_0}, \ldots, \Delta P_{T_0})$ and then comparing ΔP_{T_1} to π_t^*. Thus, three items require explanation: (i) benchmark wines used to calculate price differences ΔP_t, (ii) which data form the bubble-free subset $(\Delta P_{t_0}, \ldots, \Delta P_{T_0})$, and (iii) the concentration inequality used to determine the parity bound π_t^*. In all calculations that follow, all prices are expressed in index terms with a value of 100 in July 2005.

2.2.1. *Price Difference Calculations.* Recall $\Delta P_t := P_t^* - P_t^\circ$. Given anecdotal evidence (e.g. Stimpfig, 2010; Temperton, 2011; Authers, 2012), Lafite Rothschild is the superstar wine in $(1''')$, with price P_t^*. I use two approaches for defining the benchmark wine: the four other first growth Bordeaux wines and a synthetic control benchmark. The first growth benchmark is a simple average of Margaux, Mouton, Haut Brion, and Latour prices. These four wines had long been considered to be of equivalent quality to Lafite and with a long history of roughly equivalent prices.

The second approach uses the data-driven synthetic control methods (Abadie and Gardeazabal, 2003; Abadie, Diamond, and Hainmueller, 2010, 2015) to construct a composite "synthetic Lafite." Consider the bubble-free period as the analog to a preintervention period. Weights are assigned to $J = 46$ other Bordeaux wines (which includes the first growths but excludes Lafite and three wines which do not report a complete price history) to minimize the mean squared prediction error in the bubble-free period. That is, select weights $w_1, \ldots, w_J$

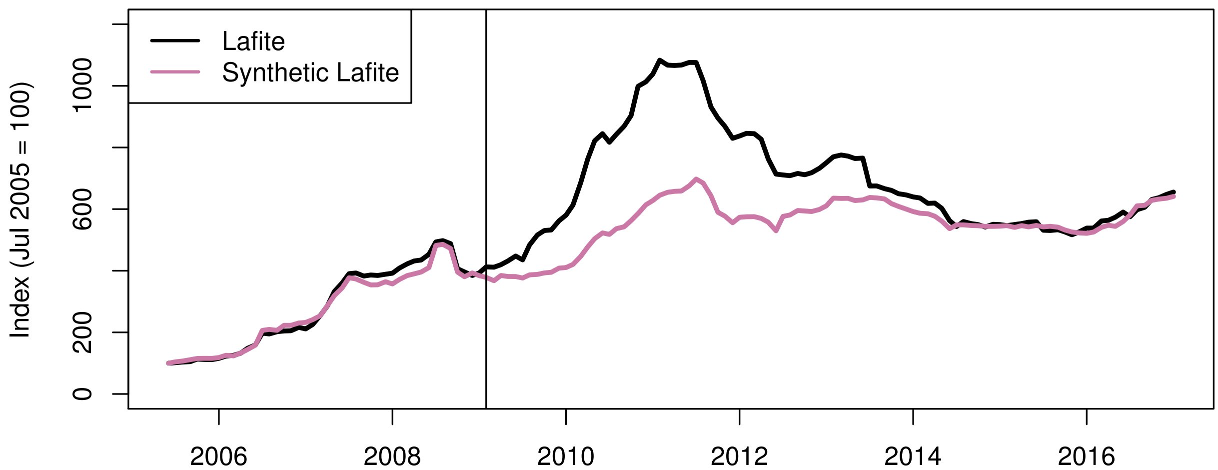

FIGURE 2. Synthetic Lafite constructed with synthetic control methods based on July 2005–Jan. 2009 pre-period.

to minimize $T_0^{-1} \sum_{t=1}^{T_0} \left(P_t^* - \sum_{j=1}^{J} w_j P_{jt}^{\circ} \right)^2$ subject to $0 \leq w_j \leq 1$ for $j = 1, \ldots, J$ and $w_1 + \cdots + w_J = 1$. Intuitively, the method uses data-driven weights to create a series that mimics the behavior of Lafite prices during the bubble-free period. More details on the other Bordeaux wines are provided in Section 3, while further details on synthetic control can be found in Abadie, Diamond, and Hainmueller (2010) and Abadie, Diamond, and Hainmueller (2015). The synthetic control is estimated with prices in index terms to account for different levels in prices. Figure 2 presents the synthetic control estimate of Lafite based on July 2005–Jan. 2009 pre-period as explained next.

2.2.2. *Bubble-Free Data Subset.* I again use two approaches to select the bubble-free data. In the first, I simply assume all observations prior to 2009 form a bubble-free pre-period; that is, ΔF_t is assumed to be fully described by pre-2009 data and otherwise time-invariant. In the second, I use a rolling window (forward recursive) approach commonly used to test for bubbles that simulates how one would test for a bubble concurrent with new data appearing from one period to the next (e.g. Phillips and Yu, 2011; Baur and Glover, 2014). The rolling window approach is advantageous in that it allows the relationship between unobservable fundamentals, ΔF_t, to be time-varying. However, this comes at the cost of an important disadvantage: by updating from one iteration to the next, when the bubble value is nonzero but not sufficiently large to exceed the π_t^* threshold, it becomes more difficult to detect a bubble in the subsequent period. Thus, the rolling window may fail to detect a bubble even when a bubble is present.[11] Regardless, I detect a bubble under both approaches. I use pre-2009 observations as a bubble-free period for the first approach and a 24-month rolling window for the second. Robustness checks demonstrate results are not sensitive to either of these choices.

2.2.3. *Concentration Inequality.* In order to operationalize this approach, I use Cheybychev's inequality to define an upper bound for the parity bounds. While other inequalities could be used, I use Cheybychev's because it is applicable to any distribution with finite first two moments and thus imposes very little structure on the underlying ΔF_t data-generating process.[12] Formally, consider testing for a bubble in the random variable ΔF_t, which has two (possibly time-dependent) finite moments denoted by μ_t and σ_t, respectively. Cheybychev's inequality is:

$$p(|\Delta F_t - \mu_t| \geq \theta\sigma_t) \leq \frac{1}{\theta^2}$$

for any real number $\theta > 0$. Cheybychev's inequality implies the probability of an observation θ standard deviations from the mean is θ^{-2}, so if $\theta = \sqrt{20}$, there is no more than a five percent probability of an observation $\sqrt{20} \approx 4.47$ standard deviations away from the mean. I use thresholds of $\theta = \sqrt{20}$ and $\theta = \sqrt{10}$ corresponding to false detection rates of five and ten percent, respectively. This is a very conservative bound—with a standard normal distribution the probability of a realization $\sqrt{20}$ ($\sqrt{10}$) standard deviations above the mean is only 3.2×10^{-6} (0.00078)—which lends the useful property of a low false detection rate.

The parity bound from the concentration inequality is then defined as $\pi_t^* := \mu_t + \theta\sigma_t$. Implementation of the test will depend on the method for selecting bubble-free data. Say we want to test for a bubble at period T_1. Replace the moments with their sample analogues. With ΔF_t assumed time-invariant, $\pi_{T_1}^*$ is calculated based on observations $(\Delta P_1, \ldots, \Delta P_{T_0})$,

[11] Borrowing the languages of a statistical test, the rolling window increases the probability of type II error, implying the basic approach has higher power.

[12] Assuming the existence of at least two finite moments does technically imply a distributional restriction on the prices (and hence their components) because it rules out some fat-tailed distributions such as the Cauchy.

where $T_0 < T_1$ is demarcates the pre-determined bubble-free period. With ΔF_t assumed time-varying, $\pi^*_{T_1}$ is calculated based on a rolling window of observations $(\Delta P_{T_1-H-1}, \ldots, \Delta P_{T_1-1})$ where H is the length of the rolling window. Thus, I detect a bubble when the current price premium exceeds 4.472 standard deviations of prices in the previous H quarters; prices must escalate rapidly from one period to the next and must exceed a high threshold to detect a bubble, given the implied probabilities of the multiplicative factor. In either case, a bubble occurs in period T_1 when:

$$\text{if} \begin{cases} \Delta P_{T_1} \leq \hat{\pi}^*_{T_1} & \text{there is no bubble} \\ \Delta P_{T_1} > \hat{\pi}^*_{T_1} & \text{there is a bubble} \end{cases}$$

2.2.4. *Advantages and Limitations.* The method I use has several advantages, including: (i) the method is model-free in the sense of Giglio, Maggiori, and Stroebel (2016) as discussed below; (ii) the upper bound is nonparametric and conservative; (iii) the method is quick-detection in that only one possibly bubbly observation is required; (iv) the forward recursive basis on which the method is calculated allows for a gradually time-varying relationship between the two products; (v) the method is computationally simple in that it requires only the calculation of a sample mean and standard deviation; (vi) the method requires a relatively small series of training observations (at minimum, enough to reliably calculate a standard deviation); and (vii) the method can be expressed as a "probability of a bubble" rather than a discrete bubble detection test. For (vii), simply calculate $A = |\Delta P_t - \hat{\mu}_t| \hat{\sigma}_t^{-1}$, and then (from the bound given by Cheybychev's inequality) the probability of a bubble is (weakly) greater than $1 - \min\{A^{-2}, 1\}$. I use (vii) in the results below.

The limitations of my proposed approach are that the training series must be bubble-free, it is not immediately obvious how to translate the test to other products and the identifying assumption of two otherwise equal products is ultimately qualitative in nature.[13] Further, the method does not distinguish the nature of the bubble (i.e. rational bubble versus sunspot), will have low power if the bubble evolves slowly over time, and (as with all bubble

[13]When the proposed method is applied to wine, I could provide evidence on the identifying assumption (the relative quality of the products) by testing the equality of expert ratings across products. This is left for future work.

detection methods) depends on the frequency of observations, which is often determined by the availability of the data.

3. Data

I use data collected by the London International Vinters Exchange (Liv-Ex) from July 2005 to January 2017. Liv-Ex collects daily transaction records on the price of fine wines exchanged through global secondary markets. Price updates are collected through over 400 member distributors and auction houses around the world, accounting for approximately 35,000 transactions worth $30 million daily. I use the monthly average of daily mid-point prices (i.e. simple average of the highest bid and lowest offer on a given day) for standard 750mL bottles. Liv-Ex uses these prices to construct indices which aim to reflect current and past conditions in the fine wine secondary market and are reported on professional data vendor services including Bloomberg and Thomson Reuters.

Liv-Ex provided monthly mid-point prices for the component wines and vintages used to construct their *Fine Wine 50* (FW50) and *Bordeaux Fine Wine 500* (FW500) indices. An exhaustive list of the component wines and vintages (used in the Jan. 2017 index) is provided in Tables 1 and 2. The FW50 is composed of the most recently released ten vintages of the five first growth Bordeaux wines. The FW500 is ten vintages of all of the wines in the FW50 as well as five 'super brands' from the Right Bank, five second growth brands, five Sauternes brands, and 30 well-regarded brands from the greater left (20 brands) and right (10 brands) bank regions. All records in the FW50 and FW500 are complete except for three wines, which do not enter the FW500 until 2016 (Pape Clement, Smith Haut Lafitte, and Clinet) and are hence dropped. Hereafter I exclude the ten vintages of my superstar wine, Lafite, from the FW50 and thus refer to it as the FW40. Similarly, I exclude the FW50 from the FW500 and refer to it as the FW450. The synthetic control considers all wines in the FW500 except for Lafite. Due to sample availability, the period of consideration is July 2005 to January 2017 ($T = 140$).

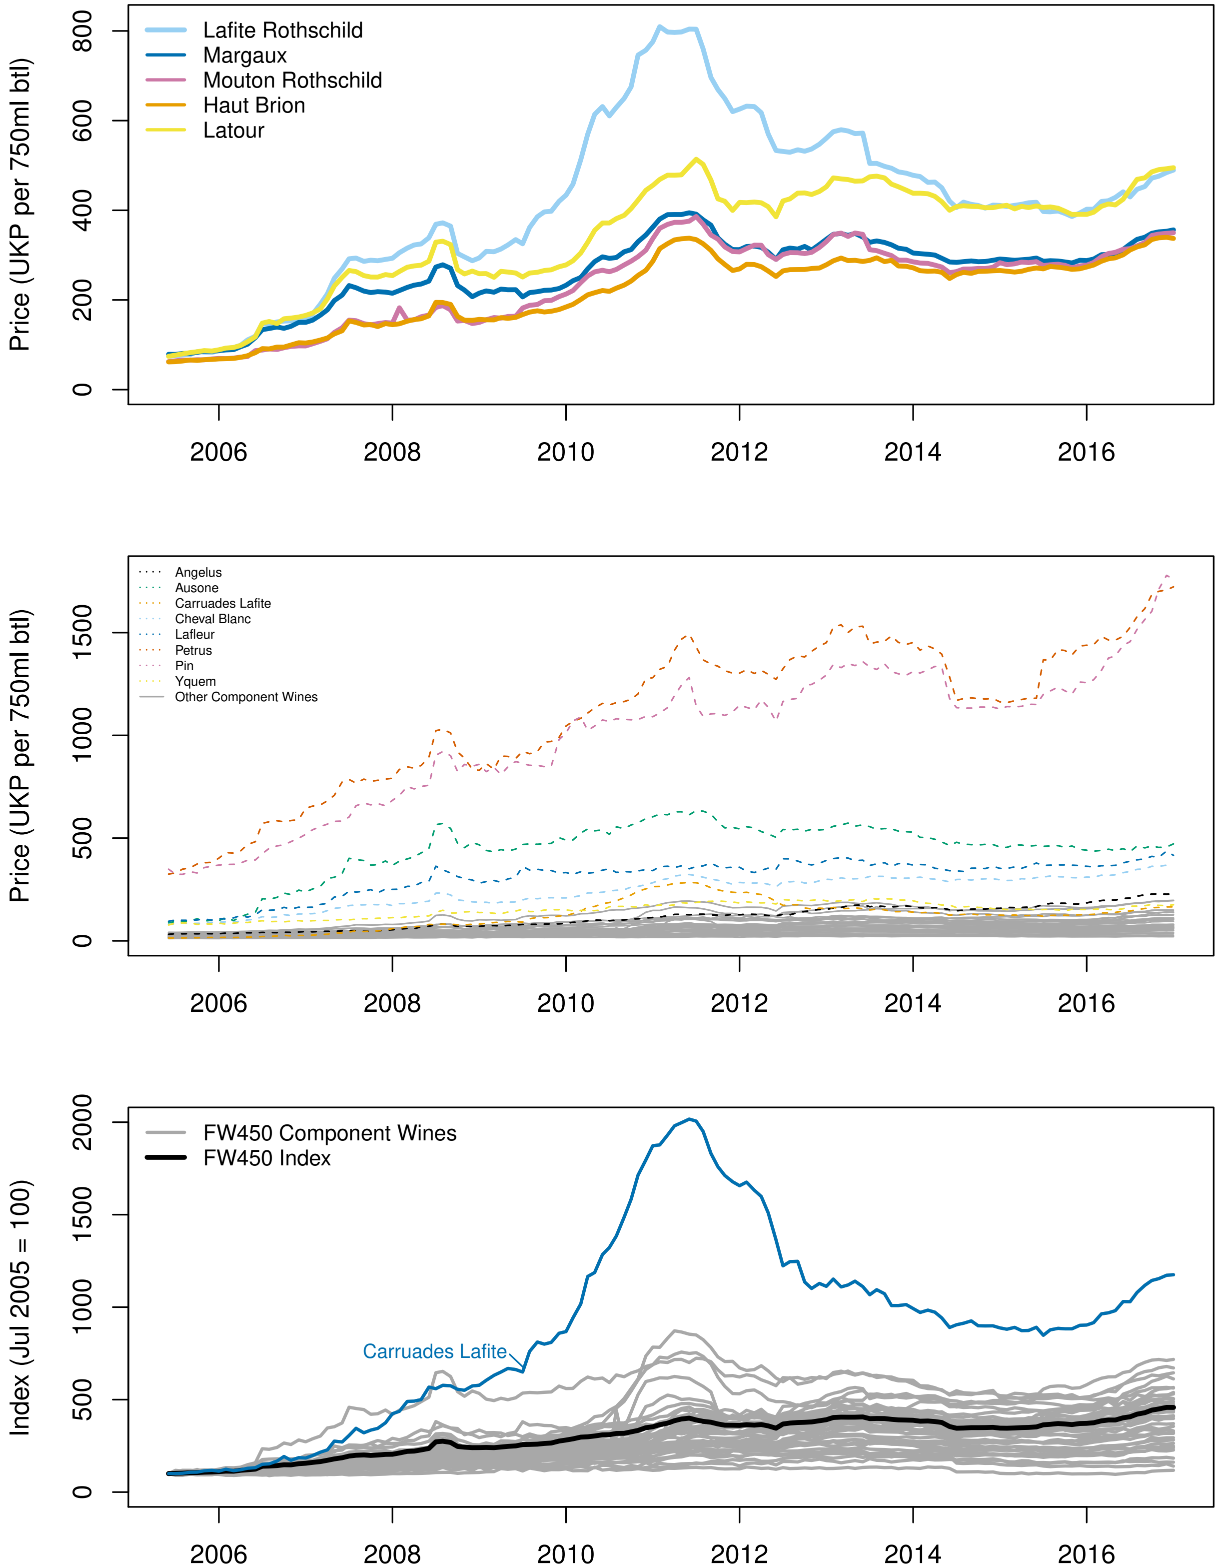

FIGURE 3. Top: prices of the FW50 components. Middle: prices of the FW450 (FW500 exc. FW50) components. Bottom: index values for the FW450 components (Base: July 2005 = 100).

Figure 3 plots the prices (£/750mL) of the individual component wines over time in a number of different ways. The top panel plots Lafite against the four other premier cru wines. Lafite is clearly the most expensive of the five wines, though after 2014 Latour begins to catch up. The middle panel plots the prices of the FW450 component wines, with the eight most expensive component wines highlighted with colors and all other wines in gray. Three wines in the FW450—Petrus and Le Pin, and at times Ausone—are priced at or above the magnitude of the premier cru wines. The bottom panel plots the individual components of the FW450 in index terms (standardized to a value of 100 for July 2005) where the potential "Lafite bubble" is evident even in Lafite Rothschild's second wine, Carruades Lafite. All values are in real terms (2015 base), deflated using the monthly UK consumer price index.

By most standards the wines are expensive: in January 2017 the average price for wines in the FW500 was £203.16, while the median wine cost £93.78. At its maximum, the price of Lafite increases more than ninefold relative to July 2005, compared to nearly fivefold for FW40 and over threefold for FW450. In terms of downside, no prices drop below their value in July 2005. The coefficient of variation in prices for Lafite (0.425) is considerably higher than that of FW40 (0.309) and FW450 (0.267). Returns for these wines were also quite high: holding Lafite provided, on average, an annualized return of over 15 percent ($1.012^{12} \approx 1.154$). The average returns were 12.7 percent for FW40 and 11.4 percent for FW450. Holding a rolling portfolio of Lafite whilest shorting FW40 or FW450 would result in monthly returns of 2.0 or 5.3 percent (26.8 and 85.8 percent annualized), respectively—large in magnitude by any standard.

4. Bounded Fundamental Value Model-Free Bubble-Test Results

I present results comparing Lafite to two benchmarks: an index of the four other premier cru wines and an index of wines selected using data-driven synthetic control methods. Intuitively, the synthetic control estimator selects weights on other Bordeaux wines that minimize the mean squared prediction error in the July 2005–Jan. 2009 pre-period. The procedure selects nonzero weights on two wines—19.5 percent Ausone and 80.5 percent Latour—which

Figure 2 shows provides an excellent representation of Lafite in the pre-period (indeed, the average monthly prediction error is 2.9 units relative to an average index value of 281.1). The results are divided into three subsections as follows. In the first subsection, the concentration inequality is calculated based on a fixed bubble-free pre-period of July 2005–December 2008 (inclusive). In the second subsection, the concentration inequality is calculated using a 24-month rolling window, allowing the relationship between two products reflected in ΔF_t and its bounds to change over time. The third subsection shows the results are robust to different lengths of the bubble-free period and rolling window, as well assuming FW450 is the benchmark wine and re-estimating the synthetic control with a shorter and longer pre-period.

4.1. **Model-Free Test with Time-Invariant Bounds.** The first set of results uses the pre-2009 bubble-free pre-period to calculate the parameters of the concentration inequality and assumes those parameters are constant throughout the sample. The panels in Figure 4 summarize the bubble-test at pre-specified levels for Lafite−FW40 (top) and Lafite−Synthetic Lafite (bottom). For Lafite−FW40, the test detects a bubble beginning in December 2009 at the ten percent level and March 2010 at the five percent level. The bubble is detected for 42 months at ten percent and 26 months at five percent. Using synthetic Lafite, the results are even stronger: The anomaly in Lafite prices is first detected in August 2009 and lasts until May 2013, a total of 46 months, all at the five percent level (there is also a detection in June 2009 at the ten percent level).

The results summarized in Figure 4 are important for at least three reasons. First, the early periods when an anomoly is detected correctly anticipate a large run-up in the relative price of Lafite, followed by a subsequent crash—exactly what would be expected in a bubble. Second, Lafite is compared to two peer groups: four other premier cru wines (the FW40), which have long been considered of equivalent prestige and quality to Lafite, and a synthetic Lafite based on data-driven synthetic control methods using other fine Bordeaux wines as possible controls. Third, I am finding strongly significant, model-free direct evidence of a bubble. I find this evidence despite using the higher than conventional statistical thresholds given by Cheybychev's inequality.

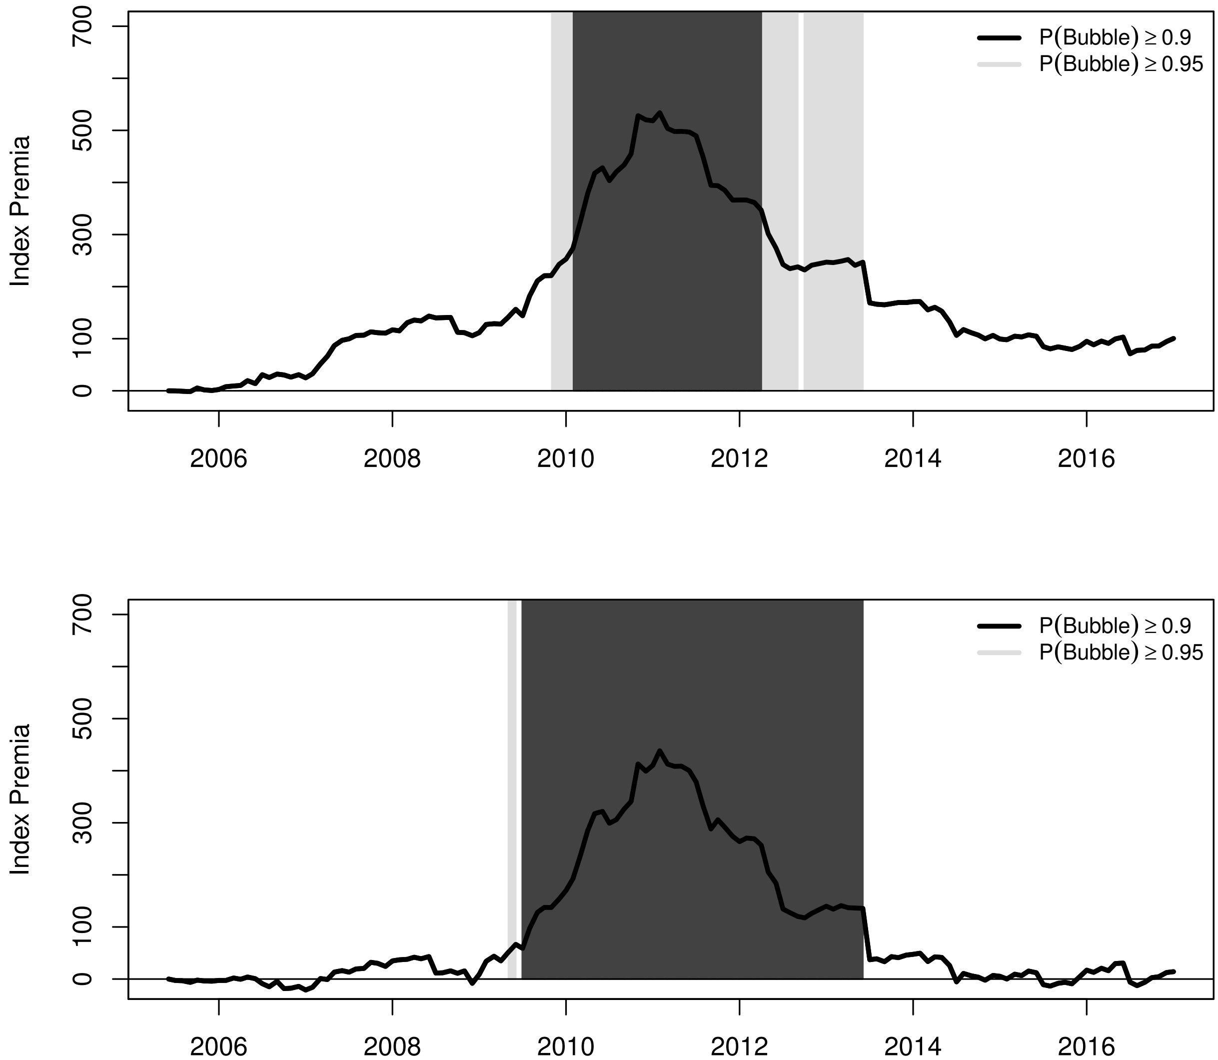

FIGURE 4. Bubble detection for Lafite with time invariant concentration inequality. Top panel: Lafite versus the four other premier cru wines (FW40). Bottom panel: Lafite versus synthetic Lafite.

4.2. **Model-Free Test with Time-Varying Bounds.** Now I allow the fundamental value relationship between the two products to be time-varying through the use of a rolling window. Recall a rolling window approach is common to the literature, especially for predictive purposes (e.g. Phillips, Wu, and Yu, 2011) but necessarily involves a trade-off: by allowing the parameters used in the concentration inequality to update, B_t may be greater than zero but if it is not sufficiently large to exceed the π_t^* threshold, it will be difficult to detect a bubble in the subsequent period. Thus, the probability of failing to detect a bubble when a bubble is present is higher with the rolling window approach. As a baseline, I use a 24-month rolling window. The results for a given comparison are best summarized in figures, each of which has three panels: the top panel plots ΔP_t over time along with the calculated values of μ_t and π_t^*, where π_t^* is calculated using both $\theta = \sqrt{20}$ and $\theta = \sqrt{10}$; the middle panel illustrates the results of the proposed method where areas shaded in gray indicate bubble detection at the five (dark gray) and ten (light gray) percent levels; and the bottom panel plots the inferred probability of a bubble at each point in time, with horizontal lines indicating the levels indicative of a 90 and 95 percent probability of a bubble and black diamond indicating the (lower bound) probability of a bubble.

The top panel of Figure 5 illustrates how the rolling window works. It is clear in this figure the π_t^* threshold change over time as a result of both changes in the mean and dispersion of the ΔF_t random variable. The threshold changes relatively slowly with changes in the mean, but quickly with changes in the dispersion—not surprising given that the standard deviation is updated with squared deviations, which results in larger changes from one period to the next relative to those generated by changes in the rolling mean. This is quite clear at the points immediately before and after 2009: prior to this point, the threshold and mean are both increasing, but after, as the standard deviation in the rolling window decreases, the threshold decreases despite the mean still going upwards. Note that if the rolling mean were used for π_t^* as in Baur and Glover (2014), the method would detect a bubble for more than 24 months, which illustrates how the proposed method is relatively conservative.

The middle panel of Figure 5 conveniently communicates the bubble detection method at pre-specified levels of θ. In the Lafite$-$FW40 comparison, the method detects a bubble five times at the ten percent level and one time at the five percent level, in two groups: August 2009–October 2009 and March 2010–May 2010. Both sets seem to accurately predict the bubble, though the crash does not occur until a few months afterwards. The concentration inequality detects large anomolies, but in subsequent periods these anomolies are incorporated into the rolling window, so subsequent detection becomes less likely (even if a bubble is present). Given the possibility of false detection (analogous to type I error),

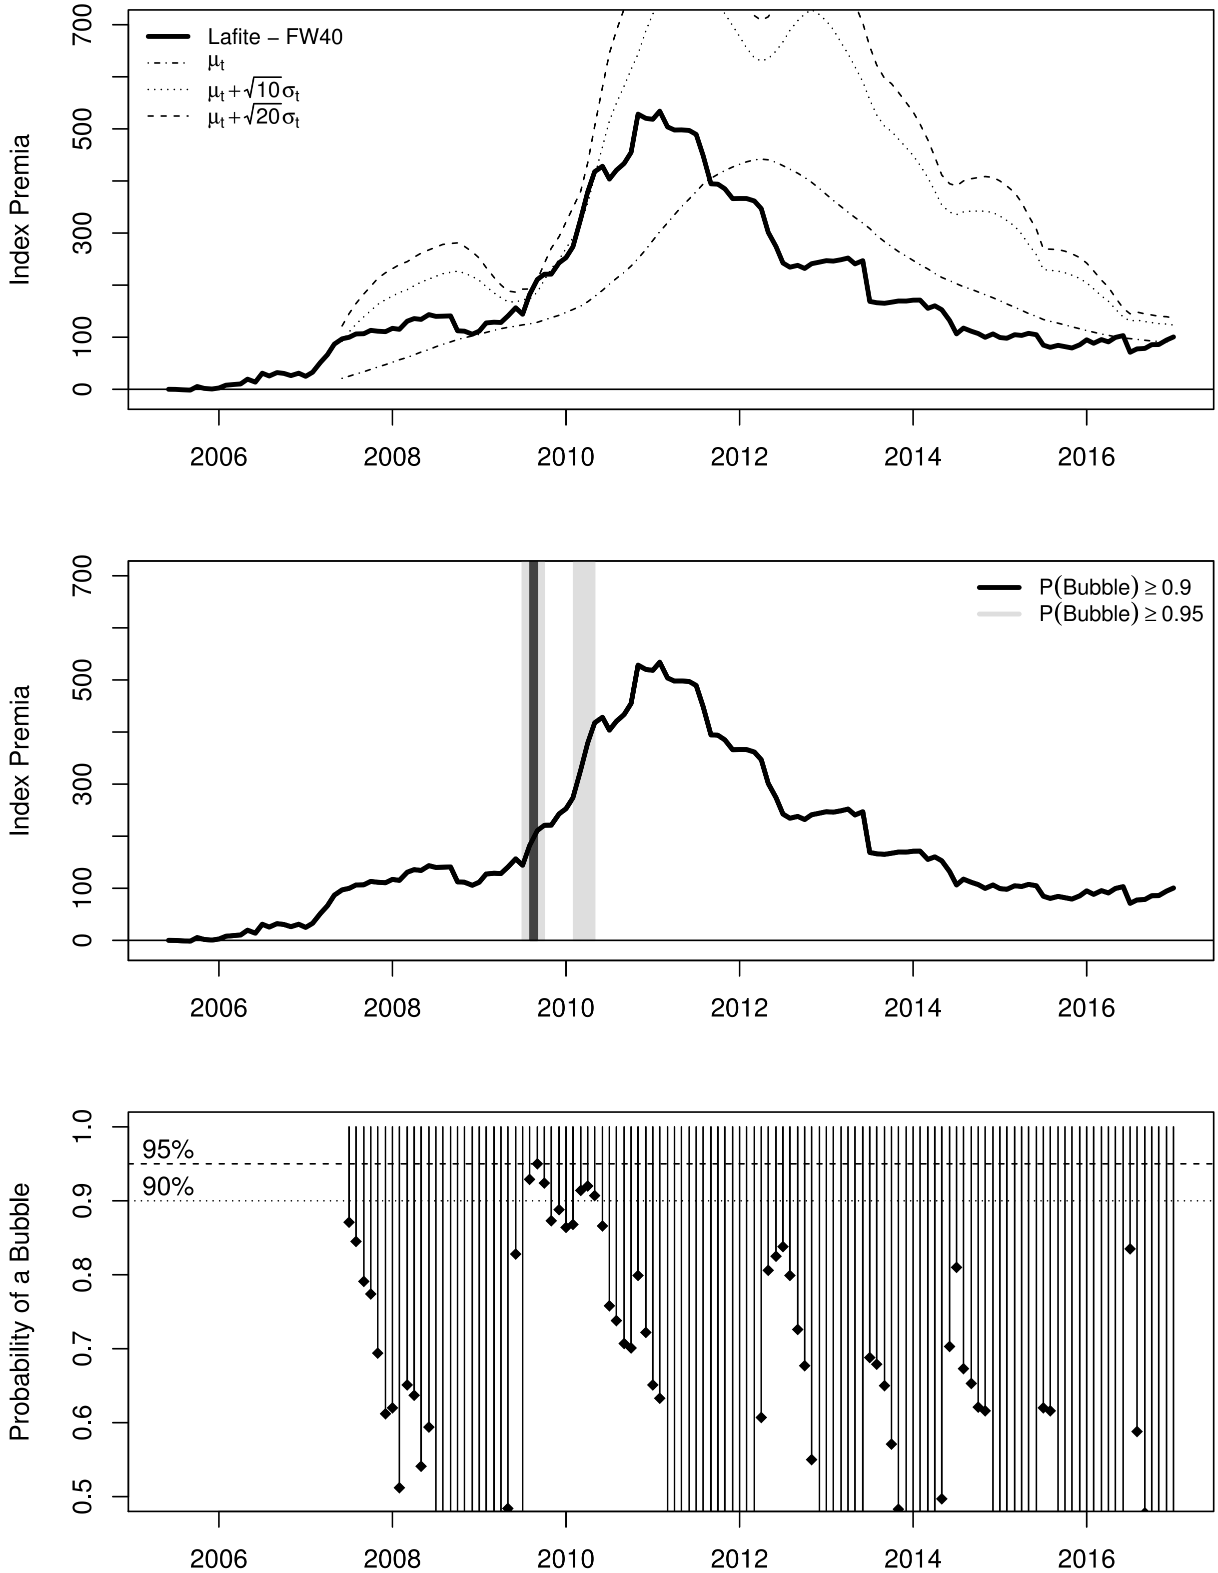

FIGURE 5. Bubble detection for Lafite−FW40 with time varying concentration inequality. Top: monthly index premia with ΔF_t upper-bound. Middle: Summary of bubble detection. Bottom: implied probability of a bubble.

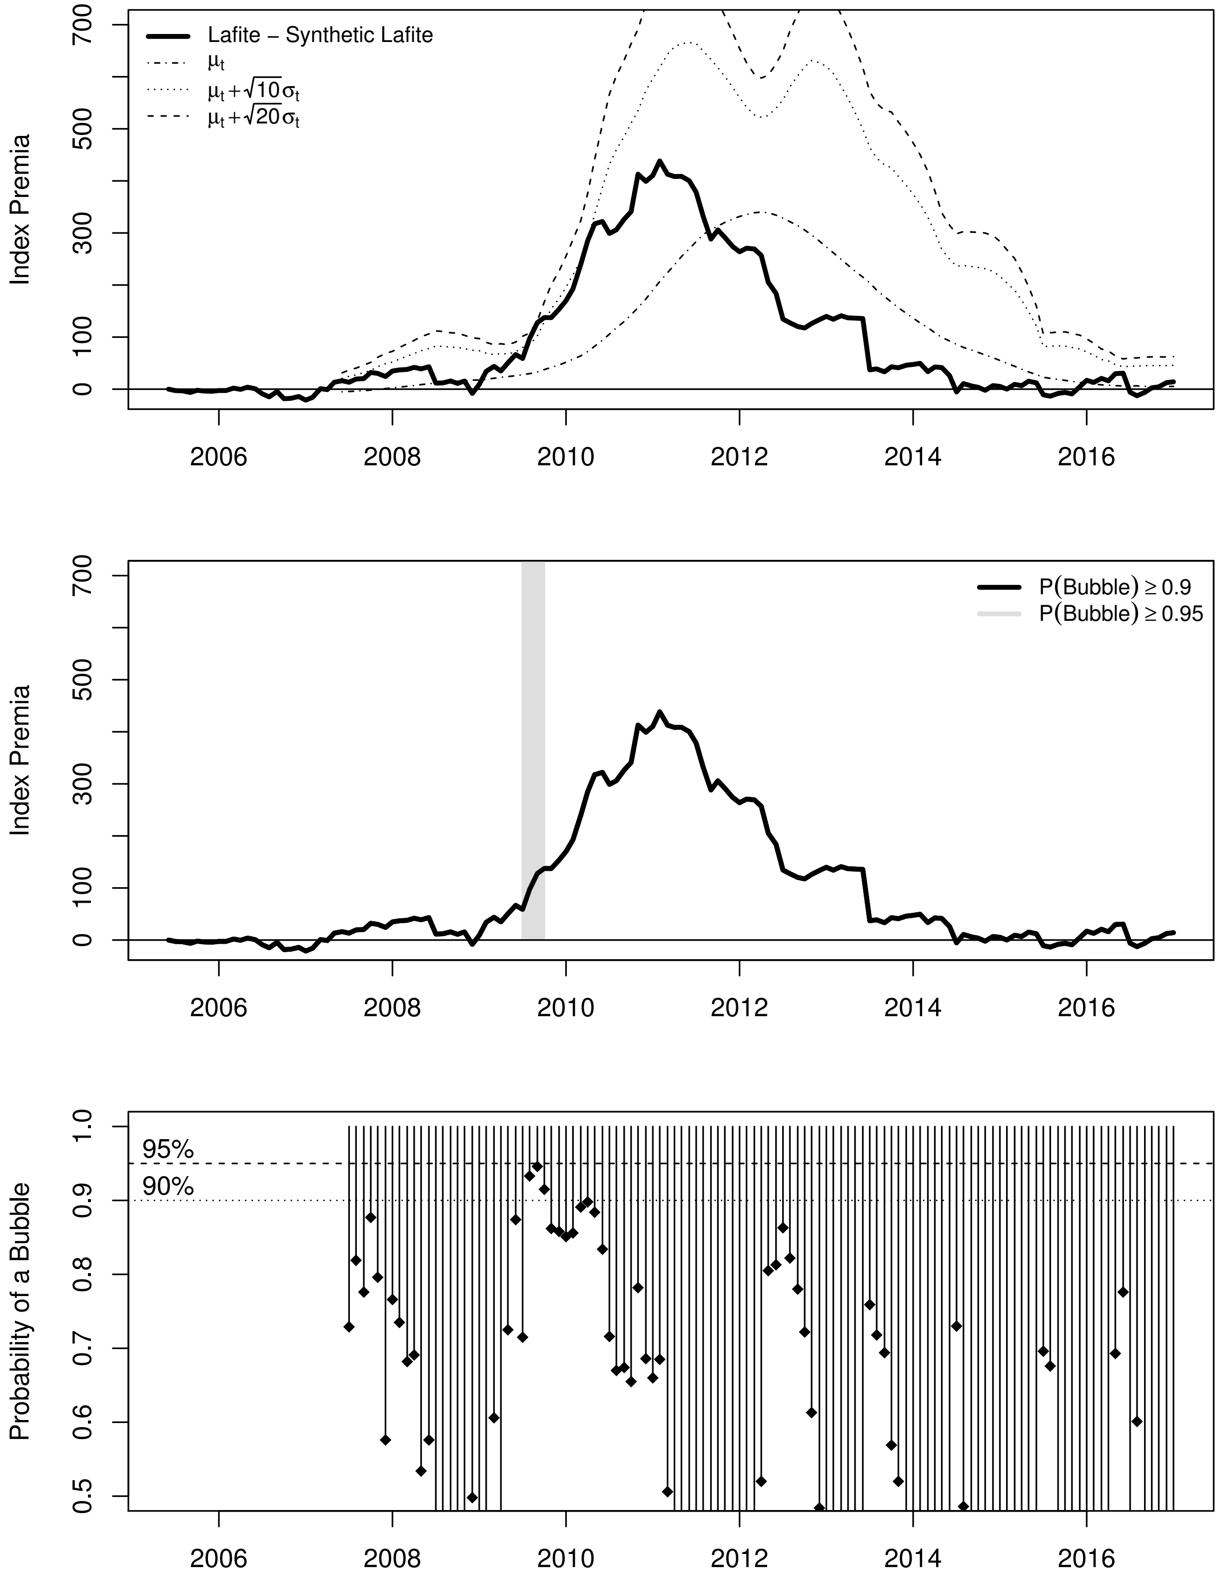

Figure 6. Bubble detection for Lafite−SYN with time varying concentration inequality. Top: monthly index premia with ΔF_t upper-bound. Middle: Summary of bubble detection. Bottom: implied probability of a bubble.

it is important to note both detection sets are three sequential periods long and that each sequential period of detection implies the possibility of a random false detection is less likely. Roughly-speaking, the probability of randomly finding three sequential periods of detection at the ten percent level is less than $(0.1)^3 = 0.001$ and two at the ten percent with one at the five percent level is less than $(.1)^2 \times 0.05 = 0.0005$. Thus, we again have compelling evidence of a bubble in Lafite−FW40 during 2009–2010, even with time-varying parameters which make bubble detection more difficult.

While the middle panel of Figure 5 is convenient, it does use a discrete threshold for detection, which by definition does not address the overall behavior of the test. The bottom panel addresses this by plotting the implied probability of a bubble. Of particular note in this panel is that the observations before and after September 2009 (when a bubble is detected at the five percent level) are fairly close to meeting the threshold, whereas the other points detected at the ten percent level are further from the five percent threshold.

Figure 6 provides the same three panels using the synthetic Lafite as the benchmark wine. We once again find evidence consistent with the presence of a bubble, though the evidence is somewhat weaker. In the middle panel we see the method detects a bubble in three consecutive months from August 2009 to October 2009 all at the ten percent level. While detection occurs only at the ten percent level, recall that the probability of random (false) detection of three consecutive observations at this level is less than $(0.1)^3 = 0.001$. Interestingly, the bubbly periods detected with the synthetic Lafite correspond exactly to the first detection set in Lafite−FW40, but the probability of a bubble with synthetic Lafite floats just below 90 percent when the strongest bubble is detected in Lafite−FW40.[14]

Taking the results of Figures 5–6 together, even with a time varying concentration inequality which works against detecting a bubble, I again find evidence of a bubble in Lafite using a model-free direct bubble test. The detection of three consecutive bubbly periods

[14] Bubbly and non-bubbly observations receive the same weight in the rolling window calculations, which makes it even more difficult for the test to detect a bubble in sequential periods or when prices rise quickly, but just below the threshold. Future work could address by down-weighting (or trimming) observations in proportion to their probability of a bubble.

implies the probability of randomly detecting the anomolies observed in the data is less than 0.1 percent.

4.3. **Robustness Checks.** Figures 7–9 summarize the robustness checks, which examine different assumptions on the: length of the time-invariant bubble-free period (36, 40, and 48 months); length of the time-varying rolling window bubble-free period (18, 36, and 48 months); and wines forming the benchmark comparison wine. Figures 7 and 8 plot results for Lafite−FW40. Lafite−Synthetic Lafite is qualitatively similar.

Figure 7 summarizes the results of the model-free test using time-invariant bounds based on 36 (top), 40 (middle), and 48 (bottom) months, whereas the main results use all observations prior to 2009 (43 months). A bubble is detected at the five percent level from March 2010 to May 2012 with 36 months and March 2010 to June 2012 for both 40 and 48 months. At the ten percent level, the test first detects a bubble in October 2009 with 36 months, December 2009 with 40 months, and January 2010 with 48 months. With longer pre-periods, the detection thresholds are (slightly) higher: evident in the period after May/June 2012 to June 2013, where 36 month pre-period detects a bubble continuously up to June 2013, 40 months for all but one observation, and 48 months for about half of the observations as ΔP_t lies very close to the 10 percent detection threshold. The time-invariant results with different bubble-free periods are consistent with those in Figure 4.

Figure 8 summarizes the results of the model-free test using time-varying bounds based on rolling windows of 18 (top), 36 (middle), and 48 (bottom) months, whereas the main results use 24 months. Bubbles are detected after 2008 using rolling windows of all three window lengths: August 2009 to October 2009 with 18 months and March 2010 to June 2010 with both 36 months and 48 months. Again, these robustness checks are broadly consistent with the main results.

All tests in Figure 9 use the 24-month rolling window approach. The top panel of Figure 9 uses the FW450 Bordeaux wines instead of the FW40. A bubble is detected at the ten percent level for three consecutive months from March 2010 to May 2010. At first glance,

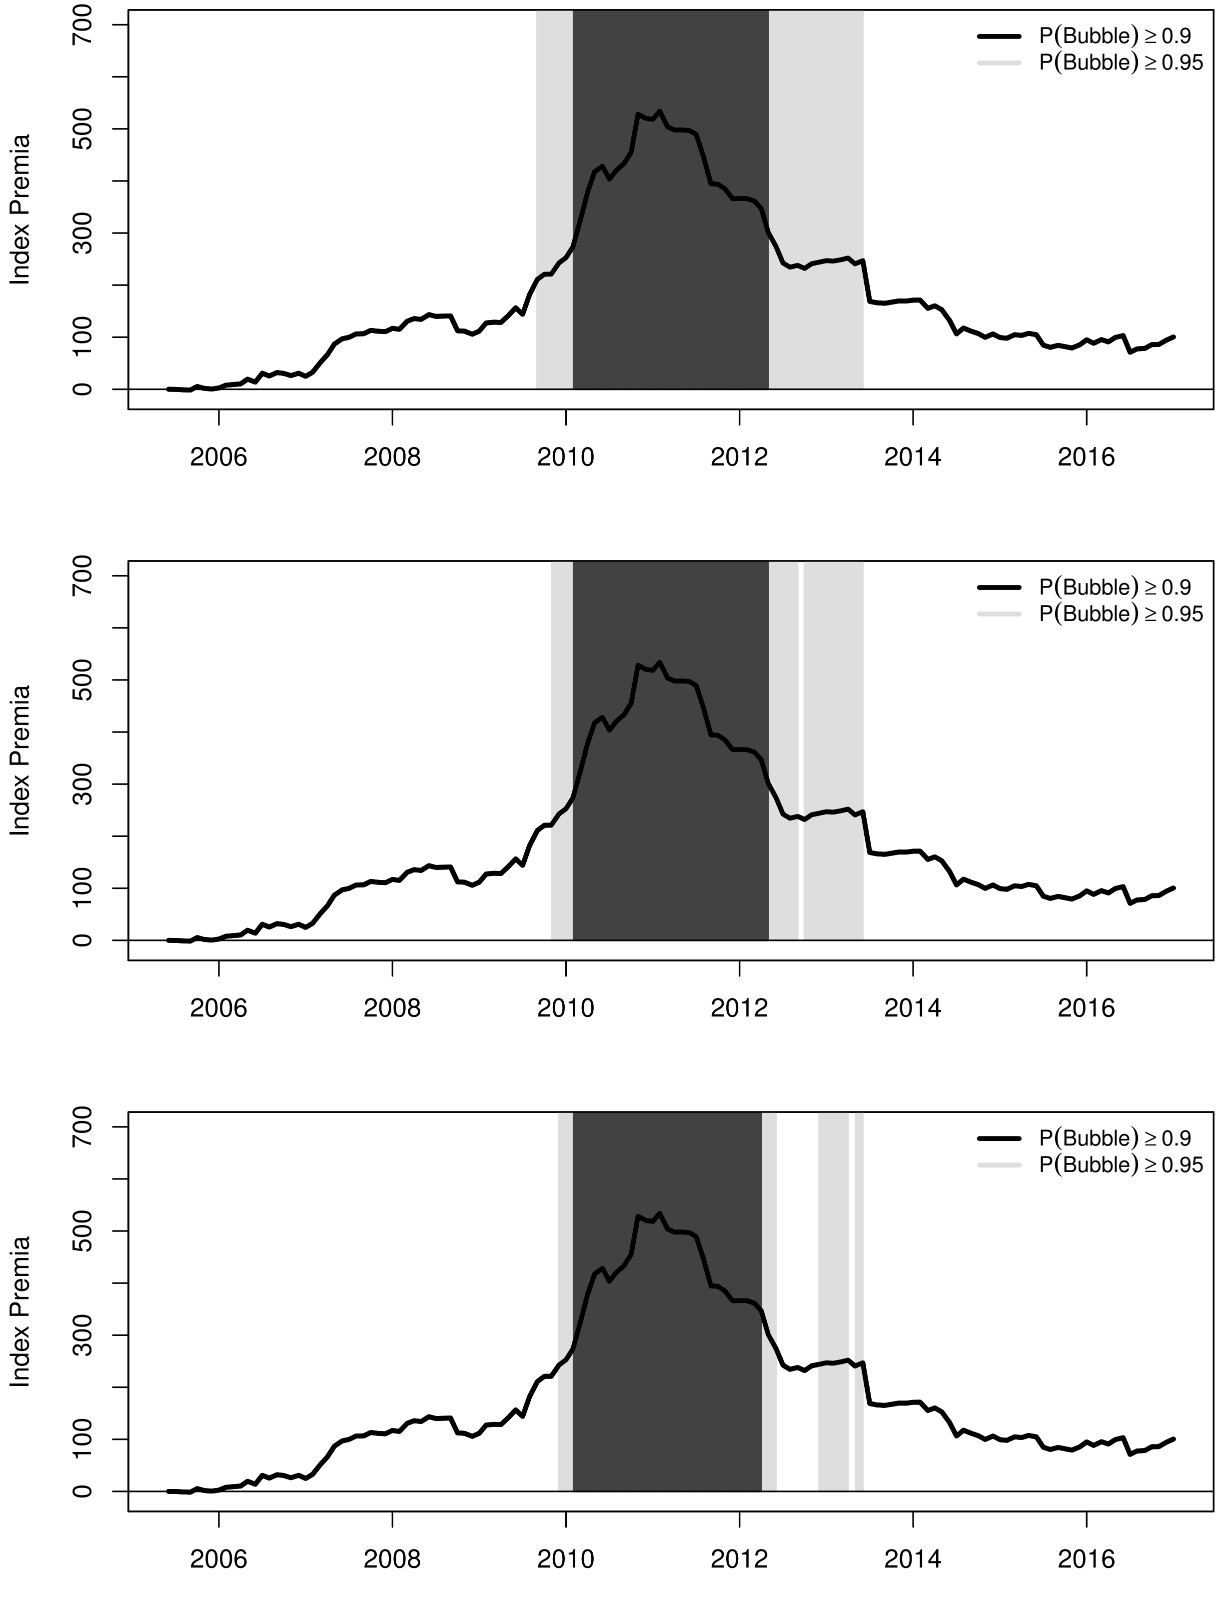

FIGURE 7. Robustness check: Lafite−FW40, time-invariant concentration inequality based on bubble-free pre-period of 36 (top), 40 (middle), and 48 (bottom) months.

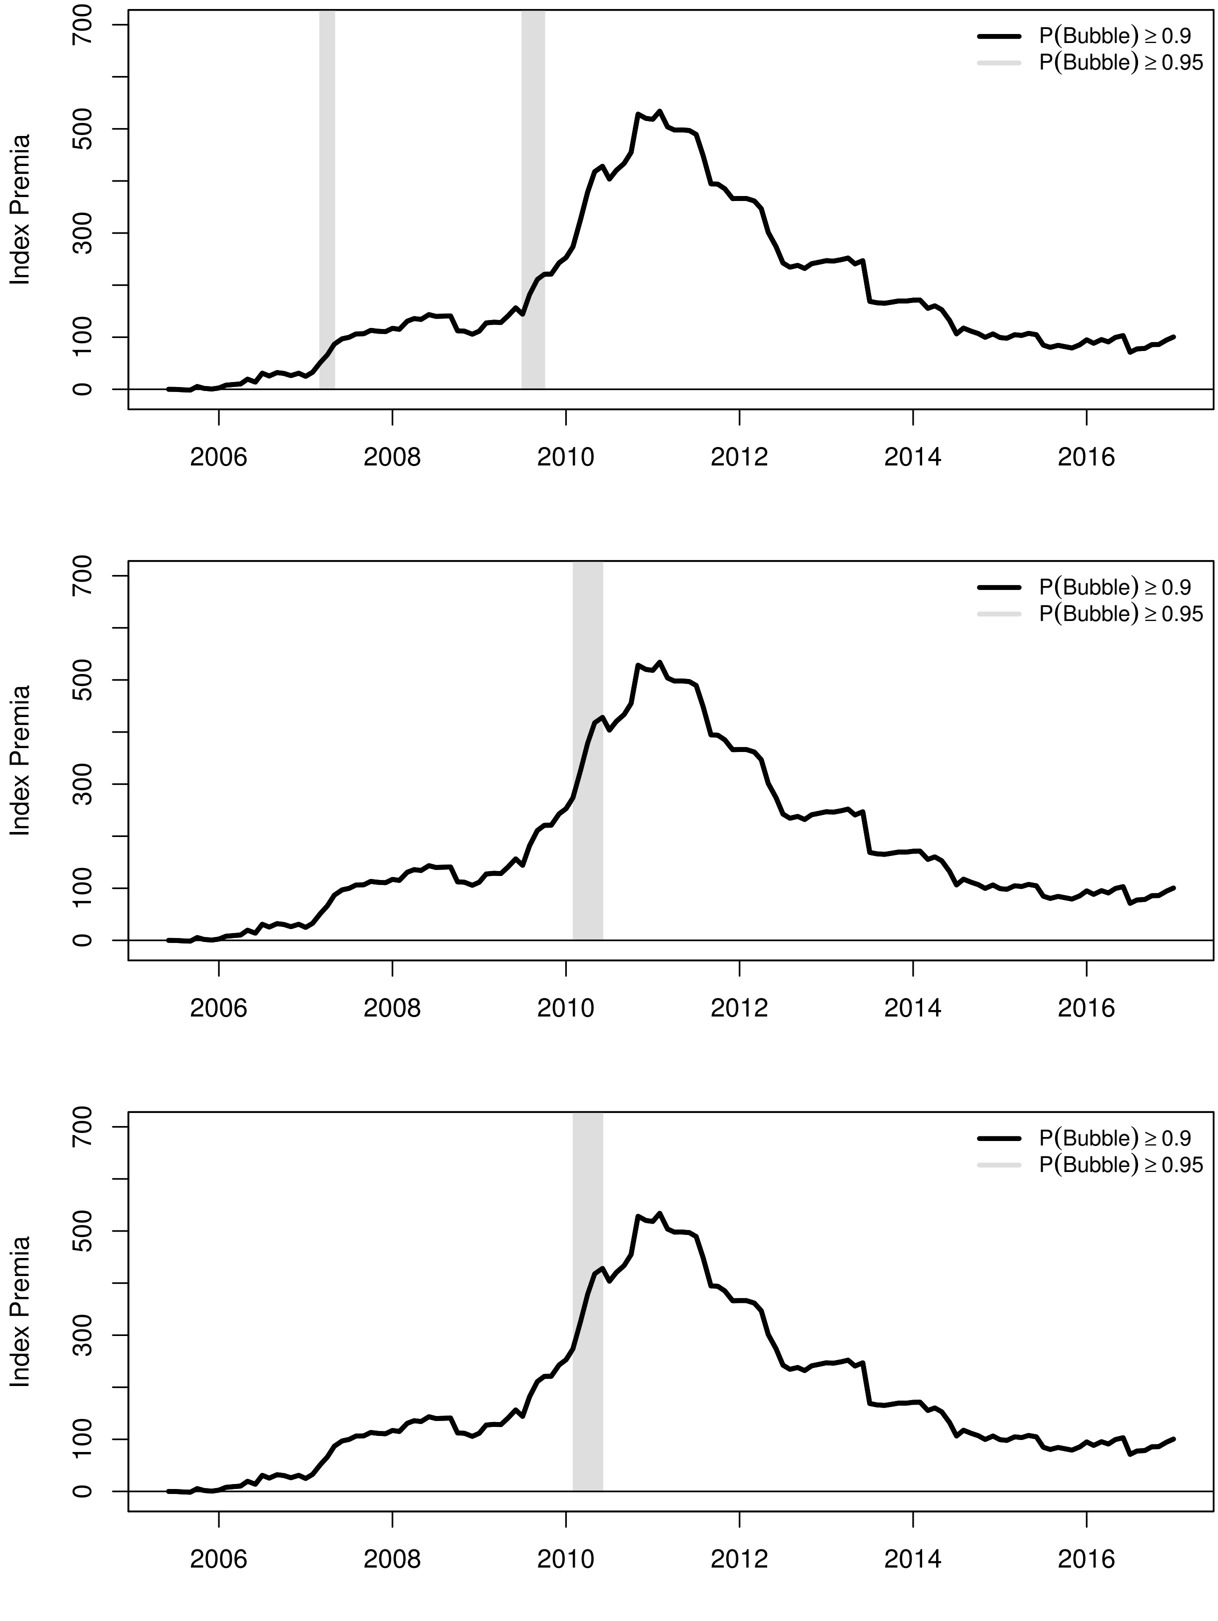

FIGURE 8. Robustness check: Lafite−FW40, time-varying concentration inequality based on bubble-free pre-period of 12 (top), 36 (middle), and 48 (bottom) months.

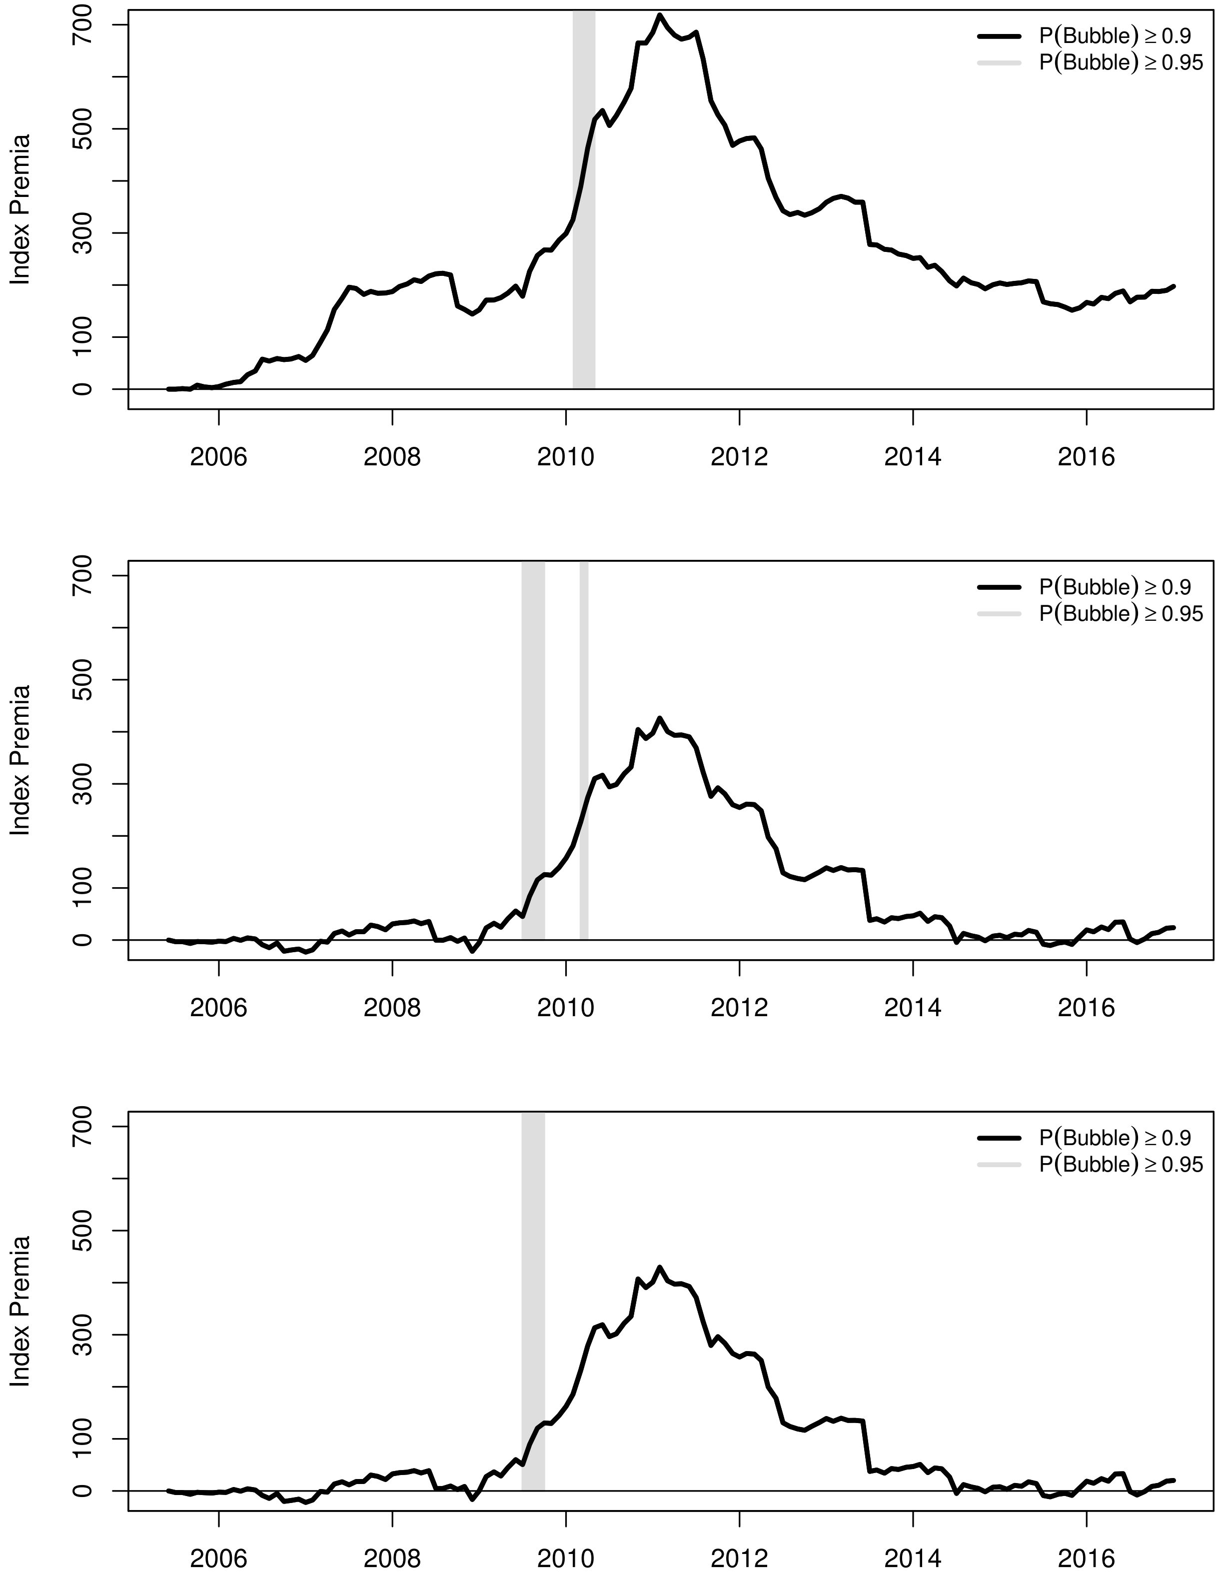

Figure 9. Robustness check: alternative benchmark wines using time-varying test. Top: 45 other fine Bordeaux wines (FW450). Middle: (bottom:) synthetic control using four less (more) months in pre-period.

detection at the five percent level in Lafite−FW45 but not in Lafite−FW450 is counterintuitive; after all, intuition suggests it should be easier to detect a Lafite bubble in *less* similar products. However, this finding is likely the result of the fact that being less related leads to Lafite−FW450 having a higher variance early in the sample, which in turn results in a higher detection threshold and the method finding the subsequent price changes "less anomalous." That being said, a bubble is detected for three consecutive periods at the ten percent level whether FW40 is the benchmark wine or the broader set of wines in the FW450.

The middle and bottom panel of Figure 9 use four fewer and four more months of bubble-free period, respectively, to estimate the synthetic controls. With four fewer months the synthetic Lafite is formed with 27.3 percent Ausone, 4.1 percent Forts Latour, and 68.6 percent Latour, while with four more months, the synthetic Lafite is formed with 24.4 percent Ausone, 3.3 percent Forts Latour, and 72.3 percent Latour. In both cases, a bubble is detected from August 2009 to October 2009. The main qualitative result—direct, model-free evidence of a bubble in the superstar Bordeaux wine Lafite Rothschild—is robust to changes in any of these assumptions.

5. Alternative Explanations and Future Work

I address some other possible reasons why asset prices quickly increase and address their applicability to Lafite.

5.1. **Supply-Side.** Sudden, unexpected shortages of a product can lead to rapid, large price increases. In the context of wine, such shortages of present (and perhaps future) inventory may be due to production shocks, while shortages of past inventory may be due to stock depletion. Note that for a shortage to induce the price anomalies observed in the data, the shortage would have to occur in Lafite *relative* to other wines.

Production shocks reducing the expected present or future inventory of Lafite are unlikely to have caused the observed price anomalies. When one thinks of a production shock in the context of wine, a "bad vintage" immediately comes to mind—typically referring to a specific region or appellation. The five premier cru chateau are *all within* 30 miles of each other: the furthest distance being between Lafite and Haut-Brion, the second furthest chateau from Lafite is Margaux at 13.7 miles. Close geographic proximity implies similar soil type, drainage capacity, weather, pest pressures, and hail incidence. Latour, Lafite, and Mouton are all within the same appellation of Bordeaux, Pauillac, and Margaux is in an adjacent appellation. The five premier cru wines are all based on the same combination of grape varietals, and appellations are delineated based on similarities in important production factors such as soil type. Indeed, the case for a bubble is actually stronger with Lafite−FW40 than in the robustness check Lafite−FW450. The latter includes more diverse production conditions (white wines and red wines from the Right Bank) and wineries that tend to be further away from Lafite (for example, Petrus is 30 miles, Angelus 31 miles, and Yquem 60 miles). Thus, a production shock would have to be limited to Lafite to cause the price spike, for example a pest outbreak. However, such a shock in one vintage would have to be very large to lead to a spike in the ten-vintage average price. Perhaps the most serious weakness with this argument is that new vintages are introduced into the ten-vintage price in June while the bubble is detected in late 2009. The 2006 vintage that came onto the market in June 2009 is widely considered to be "above average to excellent."[15] Information on vintage quality is marketed long before a vintage is released through the *en premier* futures market and expert reviews, so wine- and vintage-specific price shocks are unlikely shortly after a release. Overall, a production shock is not a very compelling explanation for the (extended) anomaly in Lafite prices observed in the data.

Dramatic price increases can occur when stocks of a non-renewable resource approach exhaustion (e.g. Jovanovic, 2007) ; however, stock depletion is similarly not a compelling argument for the price anomaly. If stock depletion of one vintage were to explain the run-up—and explain the subsequent decrease in prices—it would have to apply to one of the wines in the ten-vintage portfolio during 2009 and exit the portfolio during the price fall in late 2011. Even then, stock depletion of particular vintages tends to unfold over many

[15]For example, see Robert Parker's vintage chart: (`https://www.robertparker.com/resources/vintage-chart`. Other prominent critics have expressed similar sentiments including Jane Anson (Decanter) and Jancis Robinson.

years and is largely foreseeable by market participants. A number of the weaknesses of the production shock explanation also apply: the stock depletion would have to occur in Lafite relative to the other wines and would have to be sufficiently dramatic so as to impact the ten-vintage price. Further, this would require the resource to be exhaustible, and while in theory that is true of Lafite for any given vintage, Lafite produces a large volume of its first and second wines: while its annual production is confidential, it is the largest winery in Bordeaux and produces roughly 210,000 standard 750mL bottles of its first wine each year and more of Carruades de Lafite (Johnson and Robinson, 2007). In practice, the past ten vintages of Lafite are readily available and there is no reason to expect that to change in the foreseeable future. Thus, I would argue stock depletion is not a compelling explanation for the observed price behavior.

5.2. **Demand-Side.** Convenience yields, changes in preferences, conspicuous consumption, and increases in global demand are possible demand-side explanations for the price increase. Convenience yields can be immediately ruled out: it is hard to imagine a scenario under which many people would "need" a Lafite—and only Lafite—from any of the past ten vintages sufficiently to create a large price spike. While it could be that aggregate preferences changed suddenly to strongly favor Lafite over other fine Bordeaux wines, Stigler and Becker (1977) argue persuasively that preferences do not change instantaneously, especially at a market level. Conspicuous consumption runs into the same issue because it would imply a change in the preference for conspicuous consumption. Perhaps the most compelling argument is an increase in global demand. For example, higher wealth (in, say, China) might explain the increase in global demand. However, even this explanation presents difficulties for the same reasons that motivated Shiller (1981) to look at stock prices: it does not explain *how much* prices change from one period to the next, just the direction. In this regard, correlating fluctuations in Lafite prices with increases in wealth would likely be relatively unfruitful, as the latter changes (relatively) gradually. Note that this question, and others where fundamentals may be changed during to observable exogenous information, could be addressed with regression techniques and estimated residuals to conduct the direct, model-free bubble test *conditional* on observable information. Useful observable information might include the availability of wines (at auction or elsewhere), the quality of wines as captured through expert scores, the prices of other wines or other explanatory variables such as changes in wealth, value of investments, or the demand for other luxury goods. Proper identification using this approach may require additional assumptions.[16]

6. Conclusion

I extend the direct, model-free bubble-test of Giglio, Maggiori, and Stroebel (2016) to examine the possibility of a bubble in the Bordeaux wine Lafite. While Giglio, Maggiori, and Stroebel fails to find statistically significant evidence of a bubble in two housing markets widely considered to have a bubble, I find strong evidence of a bubble in Lafite under a more general approach than Giglio, Maggiori, and Stroebel, most importantly that the fundamental value relationship between two related products is assumed to be bounded rather than equivalent. Using assumptions that make it difficult to detect a bubble, namely a rolling window that allows the relationship between fundamental values to change over time, three consecutive observations are statistically anomalous at the ten percent level under nonparametric assumptions. The probability of randomly detecting three consecutive anomalies when no bubble is present is less than 0.1 percent. The evidence is robust to different assumptions on how the test is conducted. Future work should explore incorporating exogenous information into the test and documenting how the bubble grew and popped. One

[16]For example: assume time-invariant coefficients and that for products $j = 0, 1$ the fundamental value is given by the equation $F_{jt} = \alpha_j N_{jt} + \beta_j Q_{jt} + \rho_j P_{(-j)t-1} + \delta_j Q_{(-j)t} + \zeta_j Z_t + u_{jt}$ where the number available and quality of wine i are N_i and Q_i; prices and qualities of other products currently existing on the market are denoted P_{-i} and Q_{-i}; and Z is an m-length vector of exogenous shifters. We would require three further assumptions to get an expression that can be estimated: (i), the lagged reaction coefficient, ρ, is constant across products, i.e. $\rho_0 = \rho_1 = \rho$; (ii) scarcity premiums are equivalent, i.e. $\alpha_0 = \alpha_1$, and (iii) differences in quality are time invariant, i.e. $\eta := (\beta_1 - \delta_0)Q_{1t} - (\beta_0 + \delta_1)Q_{0t}$. Note that assumptions (i)–(iii) are strong and untestable. Under (i)—(iii), I have: $\Delta F_t = \alpha(N_{1t} - N_{0t}) + \eta - \rho(P_{1t-1} - P_{0t-1}) + \zeta Z_t + v_t$ where $\zeta := \zeta_1 - \zeta_0$ and $v_t := B_{1t} + u_{1t} - u_{0t}$, which can be estimated given data on the relative scarcity of wine one and an exogenous shifter Z. In economic terms, (i) implies changes in the price level of either product are symmetrically transmitted to the other product after a one period lag, (ii) implies both products react the same to changes in their absolute level of scarcity, and (iii) any difference in quality between the products is constant throughout time.

promising avenue for the latter is the recent literature on diverse perceptions and learning in financial markets (David and Veronesi, 2013; Eyster and Piccione, 2013; Adam, Marcet, and Nicolini, 2016).

Table 1. Component wines of the Fine Wine 500 for Jan. 2017, Wines 1–250.

Liv-Ex Group	Components	Price (£/750mL)		
		Jun. 2005	Mar. 2011	Jan. 2017
First Growth 50 (FW50)	Lafite Rothschild	74.73	797.54	489.65
	Margaux	78.72	390.28	356.30
	Mouton Rothschild	62.42	368.49	349.88
	Haut Brion	61.82	325.47	337.20
	Latour	74.34	478.27	494.96
Left Bank 200	Beychevelle	14.47	54.14	63.54
	Calon Segur	15.95	34.08	56.73
	Cos d'Estournel	21.66	74.36	95.09
	Ducru Beaucaillou	22.77	67.97	102.16
	Duhart Milon	10.98	79.63	49.42
	Grand Puy Lacoste	15.22	28.55	37.81
	Gruaud Larose	16.53	30.21	39.59
	Haut Bailly	12.65	33.24	59.01
	Leoville Barton	19.01	42.99	50.15
	Leoville Las Cases	35.24	104.14	111.87
	Leoville Poyferre	15.00	46.51	59.71
	Lynch Bages	21.61	69.67	80.00
	Mission Haut Brion	40.38	185.96	197.22
	Montrose	19.56	61.63	80.87
	Palmer	29.85	98.10	149.78
	Pape Clement	–	–	77.98
	Pichon Baron	17.32	57.35	76.14
	Pichon Lalande	26.76	69.40	75.50
	Pontet Canet	12.51	42.60	70.36
	Smith Haut Lafitte	–	–	67.45

Note: Data from the London International Vinters Exchange. Reported price is the average price of the ten most recently released vintages at that point in time (e.g. for June 2016–July 2017, the ten most recently released vintages are 2005 to 2014 for most wines). Prices are real in 2015 terms deflated with UK CPI. All wines are red except Sauternes, which are dessert wines.

Table 2. Component wines of the Fine Wine 500 for Jan. 2017, Wines 251–500.

		Price (£/750mL)		
Liv-Ex Group	Components	Jun. 2005	Mar. 2011	Jan. 2017
Right Bank 50	Petrus	324.81	1336.21	1722.72
	Ausone	87.56	613.69	471.94
	Cheval Blanc	79.49	302.58	368.86
	Pin	348.86	1134.73	1759.87
	Lafleur	95.26	363.04	416.02
Right Bank 100	Angelus	31.87	121.90	228.65
	Pavie	30.36	116.59	195.86
	Clinet	–	–	67.58
	Fleur Petrus	26.57	102.91	126.34
	Evangile	35.18	82.47	103.39
	Conseillante	24.31	65.22	78.34
	Vieux Chateau Certan	26.80	64.59	109.50
	Clos Fourtet	13.65	34.74	69.00
	Troplong Mondot	15.15	46.36	65.58
	Eglise Clinet	41.55	107.51	137.46
Sauternes 50	Yquem	77.20	171.19	175.68
	Climens	26.82	59.82	37.63
	Coutet (Barsac)	12.57	16.86	20.32
	Suduiraut	15.28	22.84	28.02
	Rieussec	20.38	26.28	24.01
Second Wines 50	Carruades Lafite	14.10	272.14	165.67
	Petit Mouton	24.82	93.74	140.46
	Forts Latour	18.89	157.20	126.83
	Pavillon Rouge	17.47	107.71	106.95
	Bahans/Clarence Haut Brion	14.57	76.46	65.29

Note: Data from the London International Vinters Exchange. Price is the average price of the ten most recently released vintages at that point in time (e.g. for June 2016–July 2017, the ten most recently released vintages are 2005 to 2014). Prices are real in 2015 terms deflated with UK CPI. All wines are red except Sauternes, which are dessert wines.

References

Abadie, A., A. Diamond, and J. Hainmueller. 2015. "Comparative politics and the synthetic control method." *American Journal of Political Science* 59:495–510.

—. 2010. "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program." *Journal of the American Statistical Association* 105:493–505.

Abadie, A., and J. Gardeazabal. 2003. "The economic costs of conflict: A case study of the Basque Country." *American Economic Review* 93:113–132.

Adam, K., A. Marcet, and J.P. Nicolini. 2016. "Stock market volatility and learning." *Journal of Finance* 71:33–82.

Al-Anaswah, N., and B. Wilfling. 2011. "Identification of speculative bubbles using state-space models with Markov-switching." *Journal of Banking & Finance* 35:1073–1086.

Andrade, E.B., T. Odean, and S. Lin. 2015. "Bubbling with excitement: An experiment." *Review of Finance* 20:447–466.

Authers, J. 2012. "Vintage strategy tempts palata of jaded investors." In *Financial Times*. The Long View, June 9, p. 18.

Baulch, B. 1997. "Transfer costs, spatial arbitrage, and testing for food market integration." *American Journal of Agricultural Economics* 79:477–487.

Baur, D.G., and K.J. Glover. 2014. "Heterogeneous expectations in the gold market: Specification and estimation." *Journal of Economic Dynamics and Control* 40:116–133.

Bhattacharya, U., and X. Yu. 2008. "The causes and consequences of recent financial market bubbles: An introduction." *The Review of Financial Studies* 21:3–10.

Bikhchandani, S., D. Hirshleifer, and I. Welch. 1992. "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of Political Economy* 100:992–1026.

Blanchard, O.J. 1979. "Speculative bubbles, crashes and rational expectations." *Economics Letters* 3:387–389.

Blanchard, O.J., and M.W. Watson. 1982. "Bubbles, rational expectations and financial markets." In P. Wachtel, ed. *Crisis in the Economic and Financial System*. Lexington, MA: Lexington Books, pp. 293–315.

Camerer, C. 1989. "Bubbles and fads in asset prices." *Journal of Economic Surveys* 3:3–41.

Campbell, J.Y., and R.J. Shiller. 1987. "Cointegration and tests of present value models." *Journal of Political Economy* 95:1062–1088.

—. 1988a. "The dividend-price ratio and expectations of future dividends and discount factors." *The Review of Financial Studies* 1:195–228.

—. 1988b. "Stock prices, earnings, and expected dividends." *Journal of Finance* 43:661–676.

Cardebat, J.M., and J.M. Figuet. 2010. "Is Bordeaux wine an alternative financial asset?" In *The 5th International Academy of Wine Business Research Conference*. pp. 8–10.

Carvalho, V.M., A. Martin, and J. Ventura. 2012. "Understanding bubbly episodes." *American Economic Review* 102:95–100.

Cheung, A., E. Roca, and J.J. Su. 2015. "Crypto-currency bubbles: an application of the Phillips–Shi–Yu (2013) methodology on Mt. Gox bitcoin prices." *Applied Economics* 47:2348–2358.

Conlon, J.R. 2004. "Simple finite horizon bubbles robust to higher order knowledge." *Econometrica* 72:927–936.

David, A., and P. Veronesi. 2013. "What ties return volatilities to price valuations and fundamentals?" *Journal of Political Economy* 121:682–746.

Diamond, P.A. 1965. "National debt in a neoclassical growth model." *American Economic Review* 55:1126–1150.

Diba, B.T., and H.I. Grossman. 1988a. "Explosive rational bubbles in stock prices?" *American Economic Review* 78:520–530.

—. 1987. "On the inception of rational bubbles." *Quarterly Journal of Economics* 102:697–700.

—. 1988b. "The theory of rational bubbles in stock prices." *The Economic Journal* 98:746–754.

Doblas-Madrid, A. 2012. "A robust model of bubbles with multidimensional uncertainty." *Econometrica* 80:1845–1893.

Dufwenberg, M., T. Lindqvist, and E. Moore. 2005. "Bubbles and experience: An experiment." *American Economic Review* 95:1731–1737.

Eckel, C.C., and S.C. Füllbrunn. 2015. "Thar she blows? Gender, competition, and bubbles in experimental asset markets." *American Economic Review* 105:906–20.

Eyster, E., and M. Piccione. 2013. "An approach to asset pricing under incomplete and diverse perceptions." *Econometrica* 81:1483–1506.

Faye, B., E. Le Fur, and S. Prat. 2015. "Dynamics of fine wine and asset prices: evidence from short-and long-run co-movements." *Applied Economics* 47:3059–3077.

Flood, R.P., and P.M. Garber. 1980. "Market fundamentals versus price-level bubbles: The first tests." *Journal of Political Economy* 88:745–770.

Flood, R.P., and R.J. Hodrick. 1990. "On testing for speculative bubbles." *Journal of Economic Perspectives* 4(2):85–101.

Froot, K.A., and M. Obstfeld. 1991. "Intrinsic bubbles: The case of stock prices." *American Economic Review* 81:1189–1214.

Giglio, S., M. Maggiori, and J. Stroebel. 2016. "No-bubble condition: Model-free tests in housing markets." *Econometrica* 84:1047–1091.

Gürkaynak, R.S. 2008. "Econometric tests of asset price bubbles: Taking stock." *Journal of Economic Surveys* 22:166–186.

Hall, S.G., Z. Psaradakis, and M. Sola. 1999. "Detecting periodically collapsing bubbles: A Markov-switching unit root test." *Journal of Applied Econometrics* 14:143–154.

Hamilton, J.D., and C.H. Whiteman. 1985. "The observable implications of self-fulfilling expectations." *Journal of Monetary Economics* 16:353–373.

Harrison, J.M., and D.M. Kreps. 1978. "Speculative investor behavior in a stock market with heterogeneous expectations." *Quarterly Journal of Economics* 92:323–336.

Hussam, R.N., D. Porter, and V.L. Smith. 2008. "Thar she blows: Can bubbles be rekindled with experienced subjects?" *American Economic Review* 98:924–37.

Johnson, H., and J. Robinson. 2007. *The World Atlas of Wine*, sixth ed., G. Pitts, ed. Mitchell Beazley.

Jovanovic, B. 2007. "Bubbles in prices of exhaustible resources." Working paper No. 13320, National Bureau of Economic Research, Cambridge, MA.

Kirchler, M., J. Huber, and T. Stöckl. 2012. "Thar she bursts: Reducing confusion reduces bubbles." *American Economic Review* 102:865–83.

Kleidon, A.W. 1986. "Variance bounds tests and stock price valuation models." *Journal of Political Economy* 94:953–1001.

Le Fur, E., H.B. Ameur, and B. Faye. 2016. "Time-varying risk premiums in the framework of wine investment." *Journal of Wine Economics* 11:355–378.

Le Roy, S.F. 2004. "Rational exuberance." *Journal of Economic Literature* 42:783–804.

Le Roy, S.F., and R.D. Porter. 1981. "The present-value relation: Tests based on implied variance bounds." *Econometrica* 49:555–574.

Lei, V., C.N. Noussair, and C.R. Plott. 2001. "Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality." *Econometrica* 69:831–859.

Lister, E. 2011. "New ways to tap into fast-growing asset class." In *Financial Times*. FT Report - Buying and Investing in Wine, June 18, p. 3.

Mackay, C. 1852. *Extraordinary Popular Delusions and the Madness of Crowds*. Bentley (reprint ed., New York: Harmony Books 1980).

Miao, J., and P. Wang. 2018. "Asset bubbles and credit constraints." *American Economic Review* 108:2590–2628.

Moinas, S., and S. Pouget. 2013. "The bubble game: An experimental study of speculation." *Econometrica* 81:1507–1539.

O'Hara, M. 2008. "Bubbles: Some perspectives (and loose talk) from history." *The Review of Financial Studies* 21:11–17.

Palan, S. 2013. "A review of bubbles and crashes in experimental asset markets." *Journal of Economic Surveys* 27:570–588.

Palfrey, T.R., and S.W. Wang. 2012. "Speculative overpricing in asset markets with information flows." *Econometrica* 80:1937–1976.

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

Phillips, P.C., S. Shi, and J. Yu. 2014. "Specification sensitivity in right-tailed unit root testing for explosive behaviour." *Oxford Bulletin of Economics and Statistics* 76:315–333.

—. 2015a. "Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500." *International Economic Review* 56:1043–1078.

—. 2015b. "Testing for multiple bubbles: Limit theory of real-time detectors." *International Economic Review* 56:1079–1134.

Phillips, P.C., Y. Wu, and J. Yu. 2011. "Explosive behavior in the 1990s Nasdaq: When did exuberance escalate asset values?" *International Economic Review* 52:201–226.

Phillips, P.C., and J. Yu. 2011. "Dating the timeline of financial bubbles during the subprime crisis." *Quantitative Economics* 2:455–491.

Rabinovitch, S. 2010. "Chinese investors raise glasses to first wine fund." In *Financial Times*. World News, August 24, p. 6.

Samuelson, P.A. 1958. "An exact consumption-loan model of interest with or without the social contrivance of money." *Journal of Political Economy* 66:467–482.

Shiller, R.J. 1981. "Do stock prices move too much to be justified by subsequent changes in dividends?" *American Economic Review* 71:421–436.

—. 2015. *Irrational Exuberance*. Princeton, NJ: Princeton University Press.

Smith, V.L., G.L. Suchanek, and A.W. Williams. 1988. "Bubbles, crashes, and endogenous expectations in experimental spot asset markets." *Econometrica* 56:1119–1151.

Stigler, G.J., and G.S. Becker. 1977. "De gustibus non est disputandum." *American Economic Review* 67:76–90.

Stimpfig, J. 2010. "Notes of recovery, hints of bubble." In *Financial Times*. FT Report - Buying and Investing in Wine, June 19, p. 1.

Temperton, P. 2011. "Try a useful asset bubble signal - the buy-to-not." In *Financial Times*. Fund Management, August 22, p. 8.

Tirole, J. 1985. "Asset bubbles and overlapping generations." *Econometrica* 53:1499–1528.

—. 1982. “On the possibility of speculation under rational expectations.” *Econometrica* 50:1163–1181.

West, K.D. 1988. “Bubbles, fads and stock price volatility tests: A partial evaluation.” *Journal of Finance* 43:639–656.

—. 1987. “A specification test for speculative bubbles.” *Quarterly Journal of Economics* 102:553–580.

Xiong, W. 2013. “Bubbles, crises, and heterogeneous beliefs.” Working paper No. 18905, National Bureau of Economic Research, Cambridge, MA.