NBER WORKING PAPER SERIES

BUBBLES, BOOMS AND CRASHES IN THE US STOCK MARKET 1792-2024

William N. Goetzmann
Otto Manninen
James Tyler

Working Paper 34903
http://www.nber.org/papers/w34903

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2026

We thank Laurence B. Siegel and Roger G. Ibbotson for helpful comments. We gratefully acknowledge research support from the International Center for Finance at the Yale School of Management. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w34903

Bubbles, Booms and Crashes in the US Stock Market 1792-2024
William N. Goetzmann, Otto Manninen, and James Tyler
NBER Working Paper No. 34903
February 2026
JEL No. G1, G10, G12, G4

ABSTRACT

We examine the historical frequency of stock market booms, crashes, and bubbles in the United States from 1792 to 2024 using aggregate market data and industry-level portfolios. We define a bubble as a large boom followed by a crash that reverses the market's prior gains. Bubbles are extremely rare. We extend the industry-level analysis of Greenwood, Shleifer, and You (2019) through 2024 and replicate their findings out of sample using Cowles Commission industry data from 1871 to 1938. Booms do not reliably predict crashes, but they do predict higher subsequent volatility, increasing the likelihood of both large gains and large losses.

William N. Goetzmann
Yale University
School of Management
and NBER
william.goetzmann@yale.edu

James Tyler
Yale University
james.tyler@yale.edu

Otto Manninen
Yale University
otto.manninen@yale.edu

Introduction

Bubbles are of great interest to investors because they are periods in which stock prices rise dramatically and drop quickly, generating large speculative profits for some and large losses for others. There are many theories about them, ranging from market constraints, to irrational behavior or to the fact that bubbles are only identifiable ex post: an extreme temporary rise will not be called a bubble unless it is followed by a crash.

A theoretical definition of an asset bubble is a large, positive, temporary deviation from fundamental value. Testing for a bubble under this definition requires a valuation model. Such models must incorporate not only fundamentals like expected future cash flows but also the complex interaction of differing investor beliefs and behavior. In this chapter we settle on a simple, empirical definition which most would agree upon. A bubble is a large boom in stock prices followed by a crash. If, for example, a market doubles in value over three years it would qualify in most people's minds as stock market boom. If prices then halved in the following years, this would certainly qualify as a crash. Putting these two together defines an empirical bubble. One could choose different thresholds or compounding intervals but the principle of defining a bubble in terms of price dynamics provides an intuitive framework for enumerating their historical frequencies.

A practical issue for investors is whether a rapid increase in prices predicts a rapid future price drop. The historical rarity of bubbles, at least in the US capital market, makes this question difficult to test empirically. One approach is to expand the sample to the full population of global equity markets using all available time periods, the approach taken in Goetzmann (2015).[4] Another approach is to follow Greenwood *et al.* (2019) and test whether industry-level price booms are more likely to be followed by crashes. This is clearly relevant to investors concerned with a boom in a single industry.

We begin by using the available history of aggregate US stock market returns to document the historical frequency of booms, crashes, and bubbles. We then use industry portfolios to test for bubble dynamics. We extend the Greenwood *et al.* (2019) analysis through 2024 and

[4] See also Goetzmann and Kim (2018), who use a global sample to study negative bubbles.

perform an out-of-sample analysis using industry returns from an earlier era in US capital market history, 1871 to 1938. Finally, we compare the conditional distribution of returns following 12-month industry booms to the unconditional distribution of returns.

We find that bubbles are extremely rare in the aggregate index as well as in industry indexes. Negative bubbles, in which a large decline is followed by a large gain, are slightly more frequent than positive bubbles. Our replication of Greenwood *et al.* (2019), using earlier data and a sample from 1926 through 2024, closely matches their original findings. Booms do not predict crashes, but crashes are more frequent following booms. Comparison of conditional to unconditional distributions shows why this is so. Booming markets are more likely to be volatile markets in the subsequent period, with a higher likelihood of both large gains and losses.

Booms and busts in the US stock market, 1792-2024

In this section we examine the frequency of booms, crashes, and bubbles over the whole course of US stock market history from 1792 to the present. We use a monthly price index of 100 large cap stocks, compiled by Finaeon, a data provider specializing in historical financial data. For simplicity we ignore the substantial return provided by dividends and their re-investment – we are interested primarily in counting big price rises followed by declines.

We define a bubble as a boom followed by a crash that reverses all of the market's prior gains. For various boom thresholds—ranging from a 50% gain over three years to a doubling—we count the number of episodes in our overlapping monthly sample and then count how many of those episodes were followed, within one, three, or five years, by either (a) a further boom of the same magnitude, or (b) a full reversal that gave back all of the earlier gain. We present these counts graphically rather than as tables of conditional probabilities, because the raw counts make two facts immediately visible: the rarity of extreme events, and the very small number of observations on which any inference about extreme outcomes must rest.

36-month booms. Figure 1 plots episode counts for booms measured over trailing 36-month windows. Red bars in the figure count the number of bubbles according to the simple definition of a price drop that fully reverses a large prior gain. Consider the 50% threshold: of the 394 overlapping episodes in which the market rose by at least 50% over three years (an annualized price increase of roughly 14.5%), 144 were followed by another gain of at least 50% within five years [dark green bar], while only 50 were followed by a full reversal [dark red bar]. At the three-year horizon the contrast is even starker: 111 further booms [medium green bar] versus 27 reversals [medium red bar]. Even within one year, the count of further booms [light green bar] (14) is comparable to the count of crashes [light red bar] (23). The visual pattern is clear: green bars dominate red bars at most horizons. For big price runups – 75% and 100% over three years, the number of bubbles is a tiny fraction of the sample size T.

Figure 1: Counts of Booms and Crashes After a 36-month Price Boom

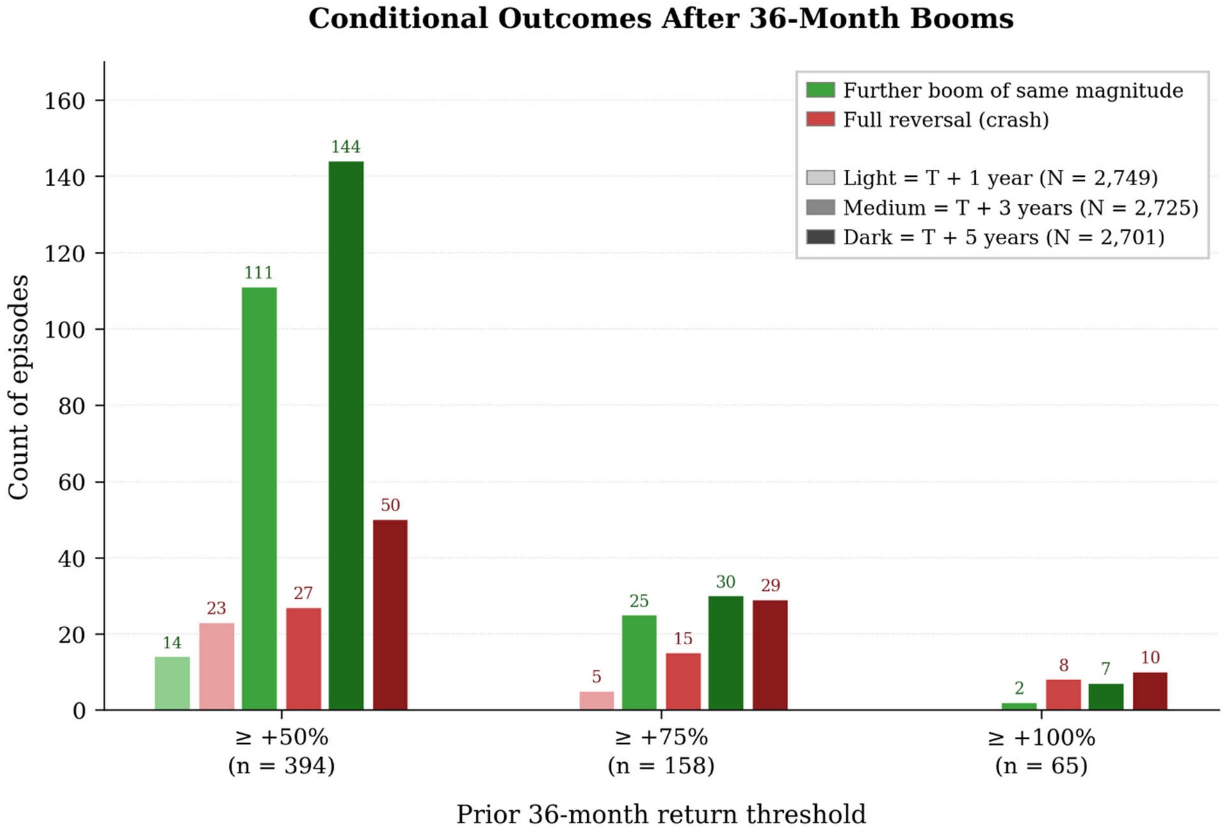

Note: Raw counts of episodes in which a subsequent gain (further boom) or loss (full reversal) of equal magnitude occurred, at forward horizons of 1, 3, and 5 years. Overlapping monthly observations; US stock market, 1792–2024.

At higher thresholds the same qualitative pattern holds, but the counts become very small. Among the 158 episodes of a 75% three-year gain, 30 were followed by a further 75% gain within five years and 29 by a full reversal—nearly a coin flip, but one based on overlapping observations of a handful of historical events. At the 100% threshold there are only 65 episodes in 232 years. Within five years, 7 experienced a further doubling and 10 a full reversal. These single-digit counts caution against computing precise odds ratios; the figure makes the sparseness of the data immediately apparent.

Figure 2: Counts of Booms and Crashes After a 12-month Price Boom

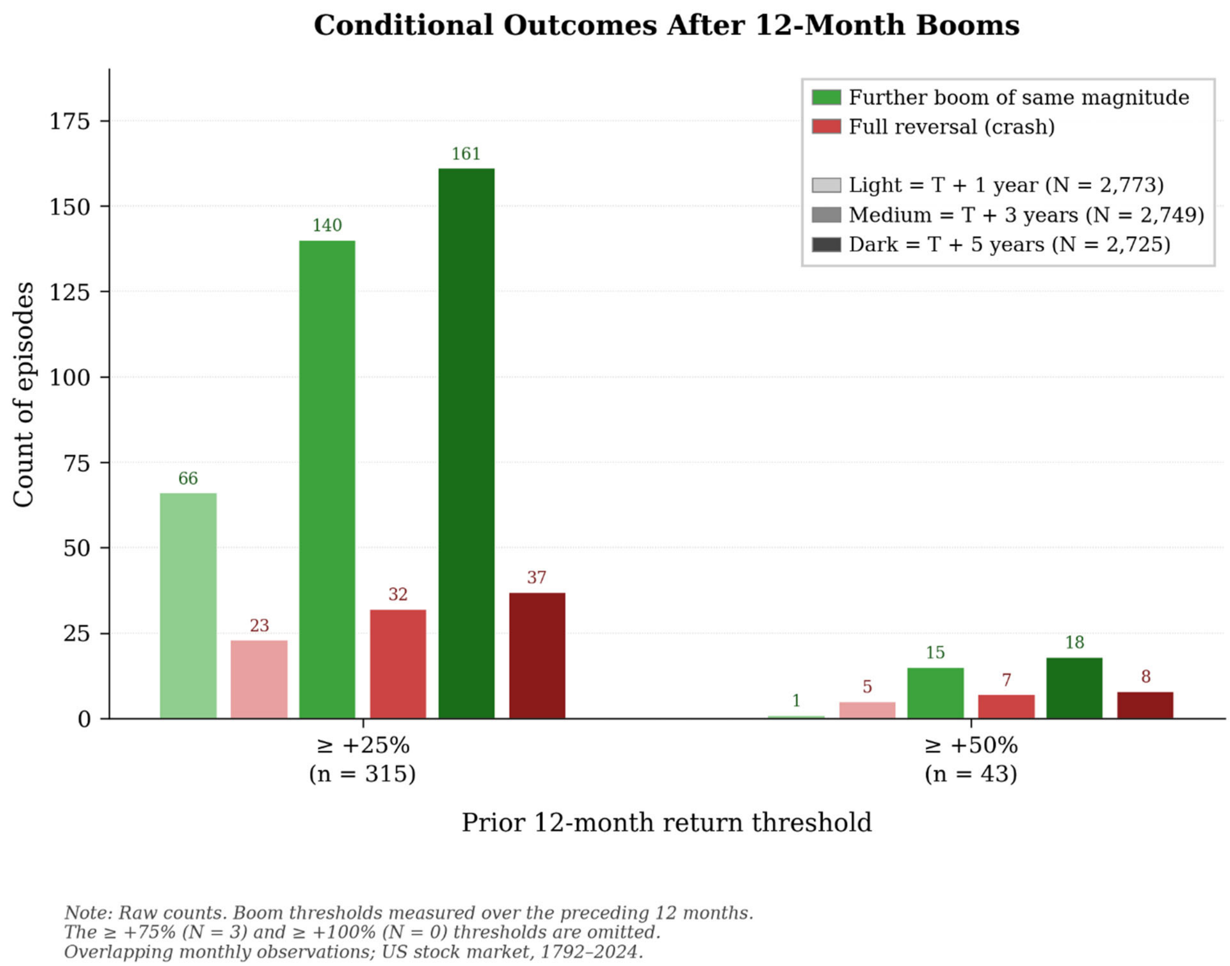

12-month booms. Figure 2 repeats the exercise for booms measured over 12-month windows. Rapid price increases of this magnitude are rarer still. A gain of 25% or more in a single year occurred 315 times in the sample; 161 of those were followed by another 25% gain within five years, versus 37 full reversals. A gain of 50% or more in twelve months occurred only 43 times. Of those, 18 subsequently gained another 50% within five years, while 8 fully reversed. At higher thresholds the data are too thin to report: a 75% twelve-month gain occurred just 3 times, and a doubling in twelve months never occurred. We omit these thresholds from the figure. Taken together, Figures

Figure 3: Counts of Booms and Crashes After a 36-month Crash

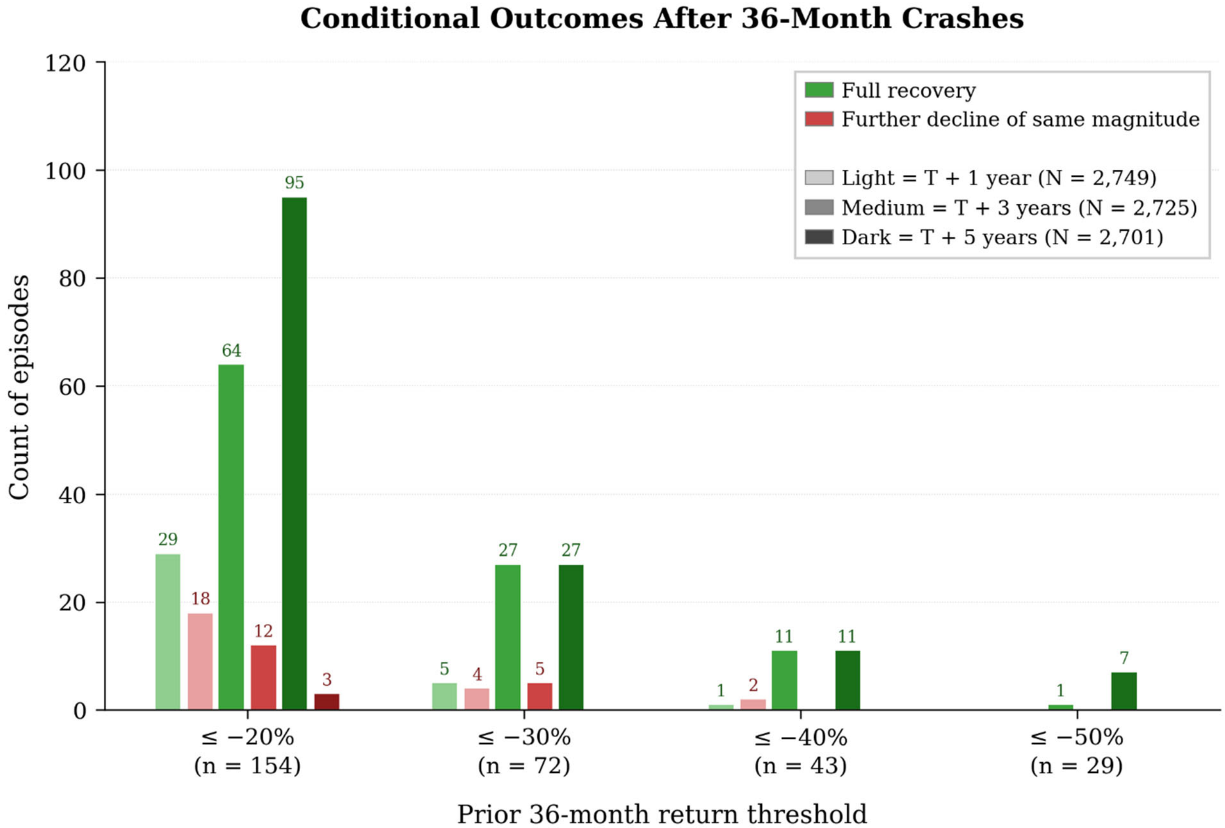

Note: Raw counts of episodes in which a subsequent recovery (reversing the prior decline) or further decline of equal magnitude occurred, at forward horizons of 1, 3, and 5 years. Overlapping monthly observations; US stock market, 1792–2024.

36-month crashes. Crashes tell a different story. Figure 3 plots episode counts for 36-month declines. Here the green bars represent full recoveries—cases in which the market regained all of its prior losses—while the red bars represent further declines of the same magnitude. The visual asymmetry is striking: as the horizon extends, the green recovery bars grow while the red further-decline bars shrink toward zero.

After a decline of 20% or more over three years (154 episodes), 95 experienced a full recovery within five years, while only 3 suffered a further 20% loss. After a 30% decline (72 episodes), 27 recovered within five years and none declined further by 30%. At the 40% and 50% thresholds the pattern is the same: recoveries occur, further declines of equal magnitude do not. Not a single episode of a 50% three-year decline was followed by an additional 50% decline at any horizon.

Figure 4: Counts of Booms and Crashes After a 12-month Crash

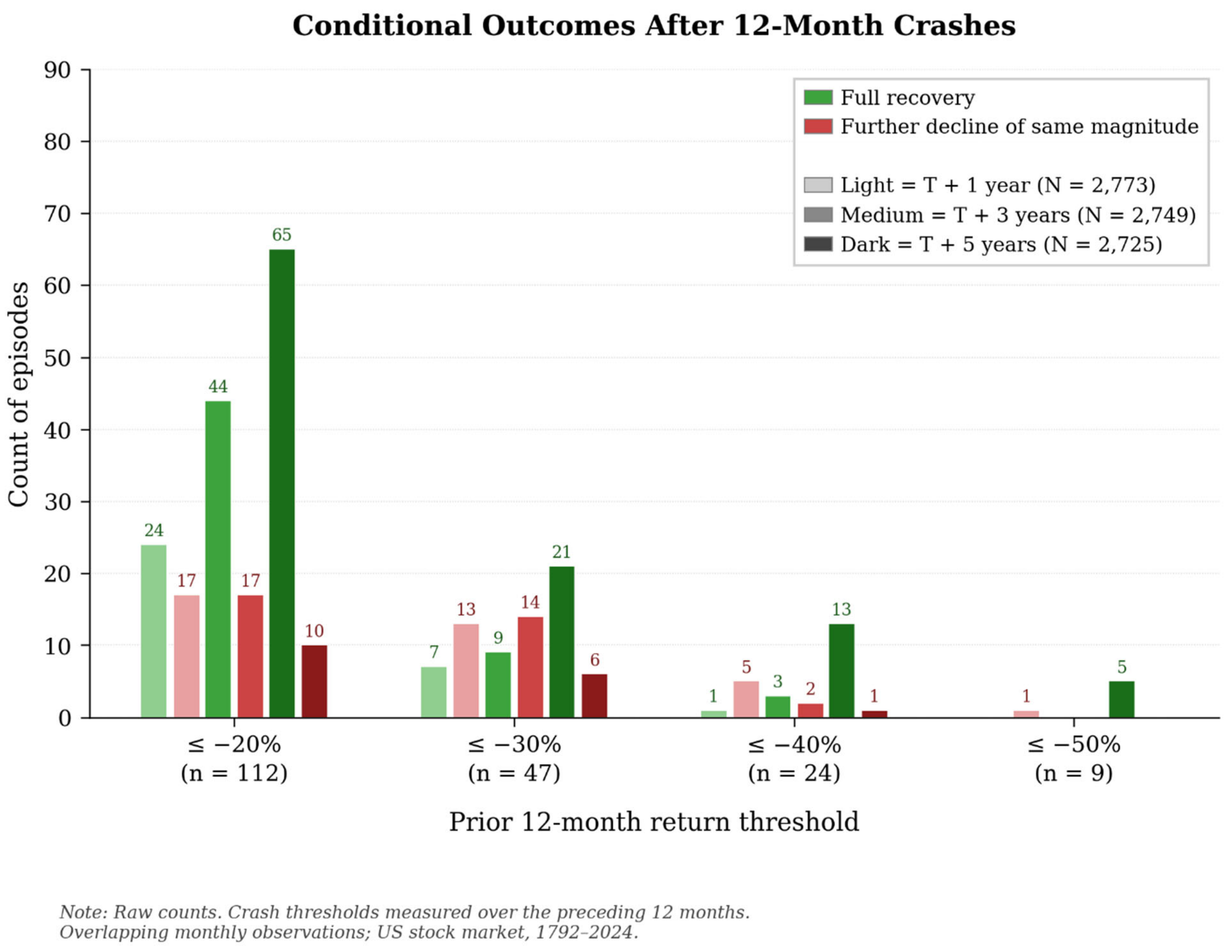

12-month crashes. Figure 4 examines rapid, twelve-month crashes—the events most frightening to investors. There were only 9 episodes in which the market fell by 50% or more in a single year, less than a third of a percent of the sample. None of these was followed by a further 50% decline at any horizon, but 5 of the 9 had fully recovered within five years.

At the 30% threshold (47 episodes), the near-term picture is more sobering: within one year, 13 episodes suffered a further 30% decline versus only 7 recoveries. But by five years the pattern reverses: 21 full recoveries versus 6 further declines. This suggests that rapid crashes carry genuine short-horizon risk of further loss, but that over longer periods the dominant outcome is recovery.

Discussion. Several features of these figures deserve comment. First, booms are more likely to be followed by further booms than by full reversals. This is true at every threshold and nearly every horizon for the 36-month measurement window, and at the 25% and 50% thresholds for the 12-month window. The popular intuition that a large run-up in stock prices is a reliable signal of an impending crash is not supported by 232 years of US data.

Second, crashes are even more likely to be followed by recoveries than booms are to be followed by further booms. At the five-year horizon, recoveries dominate further declines at every crash threshold. An investor experiencing an extended decline in the market and considering whether to sell would find little historical support for doing so.

Third, the unconditional probability of a bubble—defined as a boom followed by a full reversal—is very small. At the 50% three-year threshold, there were 50 reversals out of roughly 2,700 overlapping observations, an unconditional frequency below 2%. At the 100% threshold there were 10 reversals out of 2,700 observations, roughly one-third of one percent. Reverse bubbles—crashes followed by full recoveries—are somewhat more frequent, consistent with the general tendency of markets to rise over long horizons.

Fourth, the counts at extreme thresholds are very small, and all are based on overlapping monthly observations. A single historical event like October 1929 appears in every 36-month window that overlaps it. The seven further doublings and ten reversals at the 100% threshold are not independent observations; they may reflect one or two episodes. We present counts rather than conditional probabilities precisely to make this limitation transparent. Any precise odds ratio computed from a handful of overlapping observations of extreme events should be treated as anecdote, not evidence.

Caveats discussed above similarly apply. We only have one historical time series to study. The fact that it is the US market should also make us cautious. Over the period since 1792, the American capital market rose from obscurity to prominence, partly through survival in its early years as an emerging market, recovery after a massive civil war in the 19th century, success in two major wars in the 20th century, and successful management of the global financial crisis of the 21st century. Drawing inferences about market recovery from disasters would likely differ if we were using data from pre-revolutionary Russia, pre-World War 2 Eastern European markets, or markets that survived large shocks but faded in importance with secular changes in the world economic order.

Documenting extreme events in financial history is important for forecasting the distribution of future investment returns. The few large-scale booms and crashes experienced by the US stock market over its history help assess the probability of future tail events, but are likely to be conservative estimates of their potential magnitude, if only because they are a very small sample. Bubbles are even more rare than extreme booms and crashes because they require two rare things to occur in succession, although if one is the cause of the other that changes the probabilities. As we see below when we examine industry bubbles, both booms and busts could be correlated due to higher volatility and thus are more likely to occur concurrently.

Industry bubbles

Thus far we have examined booms, crashes, and bubbles in the historical record of the US market as a whole. Investors are equally interested in the possibility that a single industry is in a bubble. In this section we extend our boom, crash, and bubble analysis to industry portfolios comprising the US index. Historical data provide a rich cross-section of industries extending back to 1871. We use the Cowles (1938) data to test the findings of Greenwood *et al.* (2019) out of sample, and Fama and French industry portfolio data to extend their analysis to the current time. We did not chain the two periods together because definitions of industries slightly differ. The Cowles data includes more industries with presumably finer gradations among them.

As we did with the aggregate market index analysis above, we define a boom as a large return over a given time period, either 12 months or 36 months. Because there are many industry indexes, we can construct a large sample of extreme booms and extreme crashes. We thus define a boom as a 100% price increase over a 12-month window, a crash as a 50% price decline, and a bubble as a boom followed by a crash in the subsequent 12 months. To avoid double-counting, once a month is tagged as a boom, the following 12 months cannot initiate a new boom episode. This methodology combines both the approach of Greenwood *et al.* (2019) and Goetzmann (2015) and provides a simple, transparent way to identify extreme events.

We first conduct our analysis using the 49 Fama-French total return industry portfolios that span the period 1926 to present, and then use the 68 Cowles Commission monthly total return industry portfolios[5] that span the period 1871 to 1938. The benefit of using these two samples is that, by analyzing sub-sectors of the broad index, we can control for overall market movements, defining a boom as an increase over the market trend as opposed to an absolute return. Greenwood *et al.* (2019) show that this control is potentially important, because it avoids the results being driven by a handful of market-wide boom and crash periods. Finally, the industry-level analysis addresses an important question for investors. What happens when one industry suddenly booms? Is it likely that a boom in the sub-sector of the market continues, reverses, or reverts to average industry behavior?

Following Greenwood *et al.* (2019) with slight modification, a boom is identified when an industry earns both a raw and market-adjusted return of 100% or more over the past two years, and a raw return of at least 50% over the past five years. Figure 1 reports two event studies conditional on this boom definition.

The booms in the sample had an average growth of 300% over two years. However, conditional on this huge gain, their average subsequent price path was flat. The figure on the

[5] Cowles Commission indices use monthly average prices versus the typical end of month prices.

left repeats this analysis for the Cowles industries. The figures look nearly identical. They support the basic finding by Greenwood *et al.* (2019) that booms on average are not followed by crashes. The number of booms in each sample was about the same: 52 for the Cowles data, 48 for the Fama-French data.

Figure 2 takes the less-stringent boom definition, used in Goetzmann (2015), of an industry doubling or more in 12 months, and a crash definition of a decline by half or more. Normalizing the return paths to 1 at the end of their booms or crashes allows us to better see the post-trend cumulative returns. The two graphs on the left-hand side show the conditional average path of raw returns with confidence bands. Both paths drift up over the subsequent 36 months following their rise or decline, consistent with positive average equity returns for the US market over the sample periods. Confidence bands indicate that, for most horizons, the null of post event equality of boom paths and crash paths cannot be rejected, except for the first few months following the event-date.

The two right-hand-side figures define booms and crashes as doubling or halving in excess returns over the market. This significantly reduces the sample size; however, it removes the market trend from the conditional return paths. For both the early sample and the more recent sample, the average paths following booms and crashes are indistinguishable from the market. Notice a subtle difference between the raw versus excess figures. In the Cowles sample in the top row, the raw return paths exhibit a lot of time-series variation, even though the sample size is larger than the excess return sample. This is likely due to the influence of a handful of major market events that made a number of industries boom or crash at the same time. Defining booms and crashes in terms of excess returns over the market mitigates some of this sample selection bias.

A possible survivor bias is introduced by dropping industries when they disappear from the sample, resulting in averaging only across surviving industries in event time. To address this, we make the conservative assumption that capital in a discontinued industry was moved to the market portfolio at the end of its last monthly industry observation. The results are little changed. The terminal date of 1938 also censors industries from the event study. To test the effect of this, we analyzed a subset of booms and crashes prior to 1936, allowing us a full three-year potential history. The results do not differ significantly from those reported. To further investigate the source of the slight positive paths in the excess return figures, we plotted the median of the cumulative excess return indices in each period. This resulted in a terminal median value close to zero.

The results are strikingly similar across the two datasets: on average, prices rise into the boom and then on average level off, with mean paths roughly in line with the market thereafter. In the Cowles data, prices typically go flat after a sharp increase; in the more modern Fama–French data, the price evolution is closer to a smooth transition from the initial run-up.

To get another view, we examine both positive and negative bubble episodes. Figure 4 plots average paths for booms (12-month doublings) and "negative booms" (12-month halvings), using raw and market-adjusted returns for Cowles and Fama-French (hereafter FF49) industries.[6] Again, the patterns look similar across datasets. When we focus on raw returns, average outcomes after both sharp run-ups and sharp declines are generally positive over the following three years. When we look at returns in excess of the market, average paths after either type of event are close to flat.

[6] FF49 refers to the 49 industries defined by Fama and French at https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/det_49_ind_port.html.

Figure 5: Post-boom outcomes for Cowles and FF49 industries, average return for each boom-episode from t-12 to t+24.

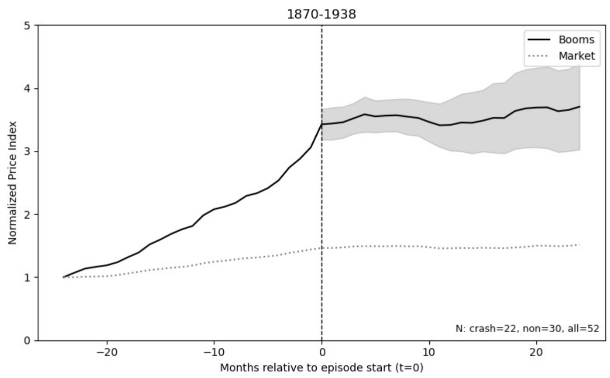

(a) Post-boom Cowles

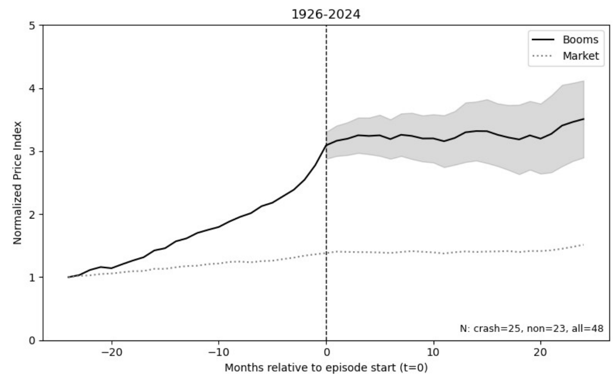

(b) Post-boom FF49

Figure 6: Post-boom and -bust outcomes based on raw and excess of market returns, Cowles and FF49.

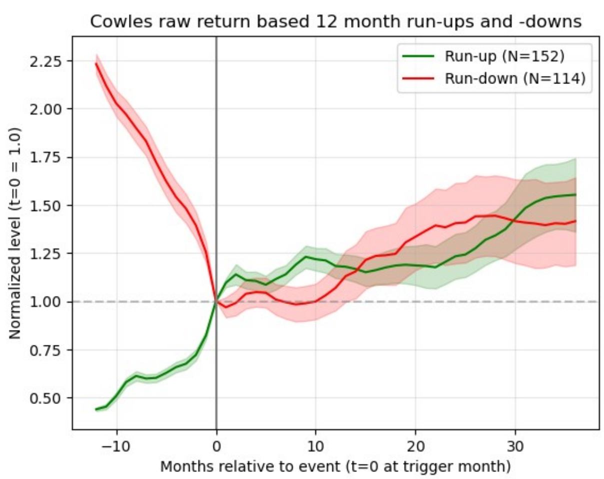

(a) Cowles — raw

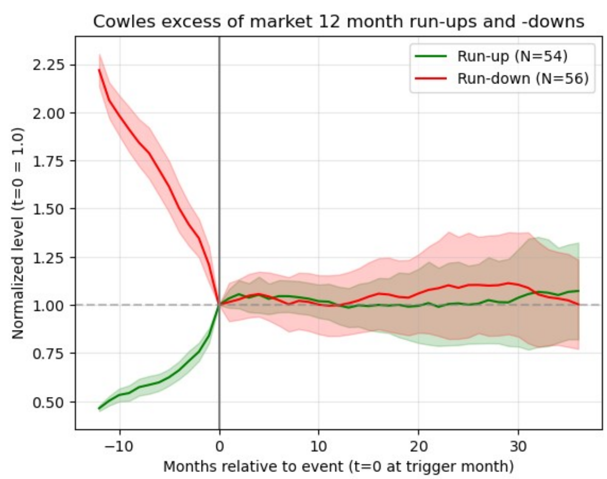

(b) Cowles — excess

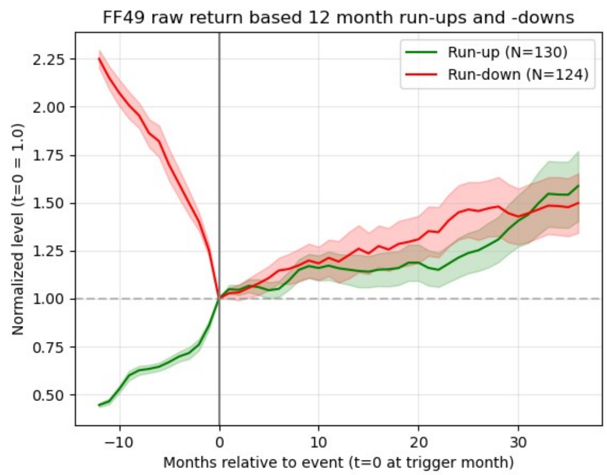

(c) FF49 — raw

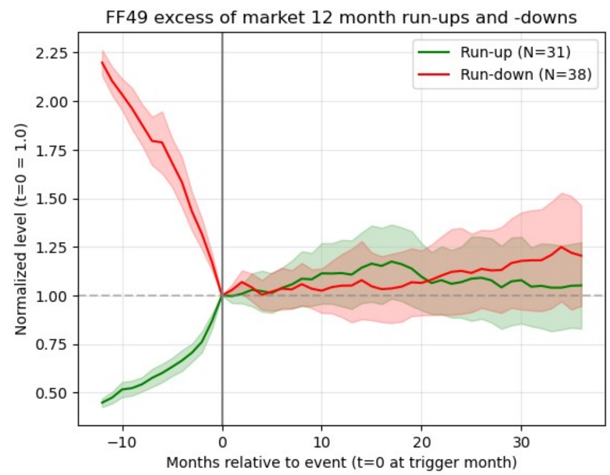

(d) FF49 — excess

Conditional distributions

In the analysis to this point we have studied boom-conditional and crash-conditional return distributions. In this section we explicitly test whether crashes are more likely if there is a preceding boom. We do this by comparing conditional distributions to unconditional distributions, i.e. to the entire historical industry return paths, not just those that condition on a previous boom or crash. We test the null hypothesis that post-boom return paths are not unusual when compared to the large, unconditional sample.

Greenwood *et al.* (2019) point out that, even though booms do not on average predict crashes, they are associated with a higher subsequent probability of a crash. This apparent contradiction is due to the fact that, by picking industries that had dramatic run-ups and declines, we are conditioning on high volatility. In this section we investigate the shape of the conditional distributions compared to the unconditional distribution of industry returns, and test for conditional volatility and skewness using the unconditional distribution as a benchmark.

Figure 3 plots the distribution of industry returns conditional on a prior 12-month boom in pink, and the entire distribution of unconditional 12 month returns in blue. As noted above, the sample of unconditional returns includes multiple overlapping periods, while conditional 12-month returns do not overlap.

The top row uses the Cowles sample and the bottom row uses the Fama-French sample. The left column shows results for raw returns; the right column shows results for excess returns. The means of each distribution are shown by vertical lines.

Figure 3: Post-boom outcomes based on raw and excess-of-market returns, clipped to (-100%, +200%), Cowles and FF49.

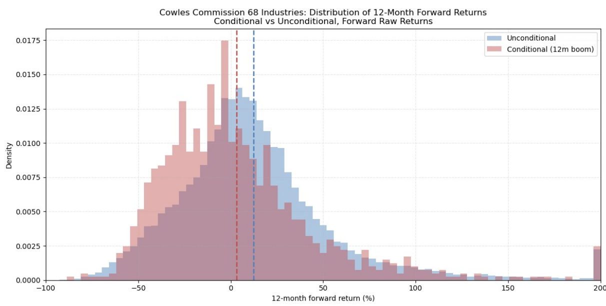

(a) 1871-1938 period — raw returns

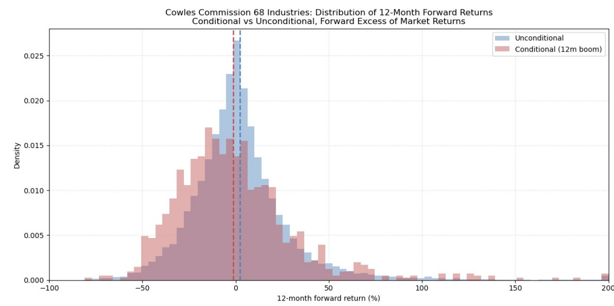

(b) 1871-1938 period — excess returns

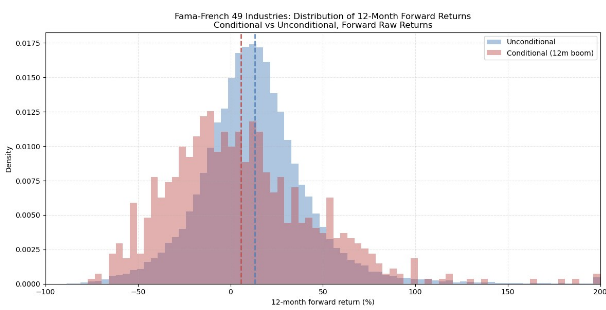

(c) 1926-2024 period — raw returns

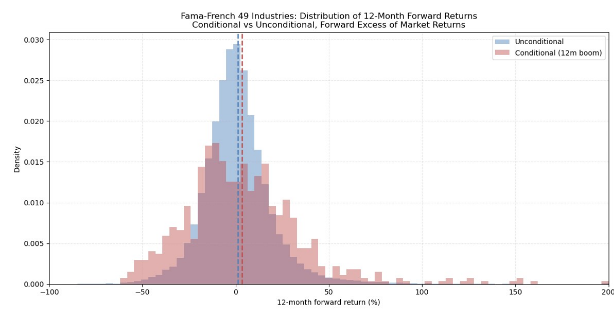

(d) 1926-2024 period — excess returns

Notice that the distribution of conditional raw returns for the 1926-2024 period is wider and shifted to the left compared to the unconditional distribution. The difference in means is large and significant. The difference in modes is also dramatic. The figure would suggest that, at the industry level, 12-month booms are a strong predictor of low excess returns in the following 12 months. However, the distribution of subsequent excess returns shown to the right suggests this is not necessarily the case. The mean of the subsequent excess returns is slightly higher than the unconditional mean. The wide spread remains: the probability of a crash after an industry boom is much higher than the unconditional probability, but so is the probability of another boom.

Conditioning on a dramatic boom selects either for high volatility industries, or for industries that are in a transitory but persistent phase of high volatility.

Figure 4: Post-bust and outcomes based on raw and excess-of-market returns, clipped to (-100%, +200%), Cowles and FF49.

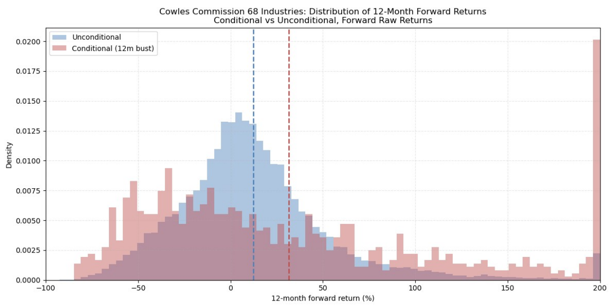

(a) 1871-1938 period — raw returns

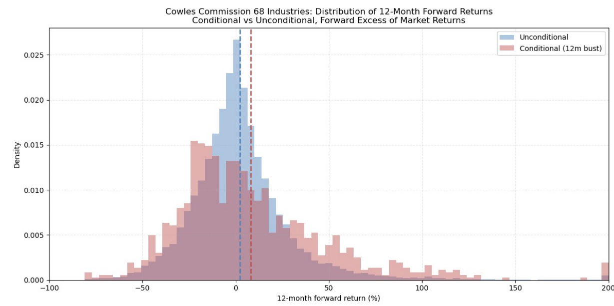

(b) 1871-1938 period — excess returns

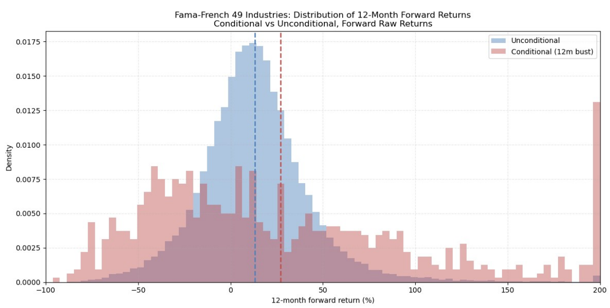

(c) 1926-2024 period — raw returns

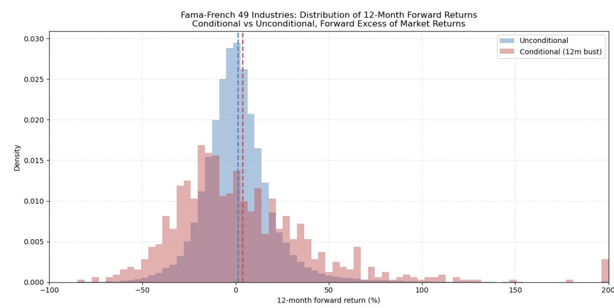

(d) 1926-2024 period — excess returns

Discussion

This chapter uses an empirical approach to define stock market booms, crashes, and bubbles. Large-scale booms and crashes have been relatively rare in US market history. Bubbles, defined as a boom followed by a crash, are by definition even rarer. Most important for investors is that the bubbles and crashes the US market experienced over more than two centuries of data have not been fatal to either the existence of the market, or to the fortunes of long-term investors.

There are many potential explanations for this continuity. Our analysis of the aggregate US index is a study of one market that grew to global dominance by the end of the 20th century. Indeed, it is the first market for which high-quality, long-term equity return data was collected and studied by financial economists (cf. Cowles, 1938). Scholars studying the cross-sectional history of global equity markets (e.g., Jorion and Goetzmann [1999], Dimson, Marsh, and Staunton [2002], Siegel [1994 *et seq.*]) have had to deal with breaks caused by war, political disruption, and rejection of the market economy. Ultimately, analysis of international equity data showed that the US was an unusually good performer, but that the equity premium is widespread and robust, even taking into account the major shocks and disruptions experienced over the past two centuries. The problem with analyzing a single market is the rarity of extreme events, which makes it difficult to test predictive models. In practical terms, the conclusion we can draw about booms and crashes based on a single, albeit multi-century return series is that they are infrequent and hard to predict with prior price trends.

The industry-level analysis in this chapter provides a richer test-bed for boom and crash prediction and the study of bubbles. The rise and fall of industrial sectors is a fundamental prediction of Schumpeter (1942). For much of its span the US stock market has channeled investment into innovative technologies that rise, succeed, and decline through time, replaced by successive waves of competition. Pástor and Veronesi (2009) model this

Schumpeterian process and find that technologies characterized by high uncertainty and fast adoption are prone to bubbles.

The industry-level analysis shows that booming (and crashing) industries are associated with subsequent high volatility, making both crashes and repeat booms more likely. However, once one controls for market trends, post-boom performance is not significantly different from unconditional returns. For crashing industries, rebounds are slightly more likely than further declines. This is potentially useful information for investors considering tactical industry allocations in real time.

References

[1] Cowles, Alfred. 1938. *Common Stock Indexes, 1871–1937.* Bloomington, IN: Principia Press.

[2] Dimson, Elroy, Paul Marsh, and Mike Staunton. 2002. Long-run global capital market returns and risk premia. *London Business School Accounting Research Paper Series.* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=299335

[3] Goetzmann, William N. 2015. Bubble investing: Learning from history. *NBER Working Paper 21693.* http://www.nber.org/papers/w21693

[4] Goetzmann, William N., and Dasol Kim. 2018. Negative bubbles: What happens after a crash. *European Financial Management* 24(2):171–191.

[5] Greenwood, Robin, Andrei Shleifer, and Yang You. 2019. Bubbles for Fama. *Journal of Financial Economics* 131(1):20–43.

[6] Jorion, Philippe, and William N. Goetzmann. 1999. Global stock markets in the twentieth century. *The Journal of Finance* 54(3):953–980.

[7] Pástor, Ľuboš, and Pietro Veronesi. 2009. Technological revolutions and stock prices. *American Economic Review* 99(4):1451–1483.

[8] Schumpeter, Joseph A. 1942. *Capitalism, Socialism and Democracy.* London: Routledge.

[9] Siegel, Jeremy J. 2022. *Stocks for the Long Run: The Definitive Guide to Financial Market Returns & Long-Term Investment Strategies,* Sixth Edition. New York: McGraw Hill.