

DYNAMIC WARP ANALYSIS: A NEW APPROACH FOR DETECTING AND TIMING BUBBLES

DECEMBER 12, 2023

Mark Kritzman, Huili Song, and David Turkington

Mark Kritzman is the chief executive officer at Windham Capital Management in Boston, MA, and a senior lecturer at the MIT Sloan School of Management in Cambridge, MA.

kritzman@mit.edu

100 Main Street, Cambridge MA, 02142

617-901-4455

Huili Song is a vice president at State Street Associates in Cambridge, MA.

hsong@statestreet.com

140 Mt Auburn Street, Cambridge MA, 02138

857-472-2022

David Turkington is senior managing director and head of State Street Associates in Cambridge, MA.

dturkington@statestreet.com

140 Mt Auburn Street, Cambridge MA, 02138

617-955-7679

Key Takeaways

- Dynamic warp analysis enables investors to rescale individual stock price bubbles that progress along different calendar paths into synchronized steps.
- This synchronization enables investors to observe stock characteristics that coincide with different phases of a bubble as well as periods when a stock is not experiencing a bubble.
- By warping nearly 7 million bubble pairs, the authors offer compelling evidence that investors may be able to profit by detecting bubbles and recognizing how far they have progressed.

Abstract

The authors apply a technique called dynamic warp analysis to rescale the unique cadences of 2,638 bubbles into synchronized steps. They then observe the distributions of chosen stock characteristics for each step across all the bubbles. They also observe these stock characteristics during periods when a stock is not experiencing a bubble. The authors use this information to detect when a bubble is underway and how far it has progressed. They test several trading rules to assess the potential to profit from this information.

DYNAMIC WARP ANALYSIS: A NEW APPROACH FOR DETECTING AND TIMING BUBBLES

Investors have long been challenged to detect when a stock price bubble has begun and if so, whether it is in its early, middle, or late stage.¹ This task has proven to be daunting because bubbles progress at different paces. Some bubbles fully evolve from inception to conclusion in just a few days whereas others proceed over several years. Moreover, though we tend to visualize a bubble as a smooth and symmetric concave progression of prices, bubbles ascend and descend non-monotonically. It would be much easier to detect a bubble and where it is along its path if we could rescale calendar time into synchronized units.

We, therefore, apply a rescaling technique called dynamic warp analysis to analyze 2,638 individual stock price bubbles that occurred between January 1, 1973 and May 16, 2023. We show that when converted to synchronized warped units, bubbles conform more closely to our stylized visualization of them, and they exhibit common characteristics that enable us to predict with considerable success the emergence of a bubble and how advanced it is along its journey.

We proceed as follows. First, we illustrate our warping algorithm with a numerical example. We then give an example of two bubbles that proceeded at dramatically different paces in calendar time, but when warped proceeded along remarkably similar paths. Next, we describe our data including the rules we use to define bubbles and the stock characteristics we

use to indicate the phases of a bubble. We then describe our methodology for estimating whether a bubble has begun and, if so, how far it has progressed. Finally, we provide evidence of the efficacy of our bubble detection system by testing trading rules designed to exploit bubble dynamics.

Dynamic Warp Analysis

Dynamic warp analysis is a technique for synchronizing series that proceed at different cadences. It was introduced in the 1970s to aid with speech recognition and later shown to be closely related to Hidden Markov Models.² We apply this technique to synchronize the evolution of stock price bubbles that evolved disparately when observed in calendar time.

Warping Algorithm

Consider two series, A and B.

Exhibit 1: Pre-warped Series

Step	Series A	Series B
1	0.01	-1.04
2	0.07	-1.12
3	-0.55	-0.49
4	-0.33	-0.55
5	-0.15	-0.52
6	0.21	-0.11
7	-0.09	-0.35
8	0.81	-0.50
9	0.37	-0.45
10	0.63	-0.54

We begin by constructing a cumulative distance matrix as shown in Exhibit 2.

Exhibit 2: Cumulative Distance Matrix

	A	0.01	0.07	-0.55	-0.33	-0.15	0.21	-0.09	0.81	0.37	0.63
B											
-1.04	1.10	2.32	2.56	3.06	3.84	5.40	6.29	9.70	11.68	14.46	
-1.12	2.38	2.52	2.65	3.19	4.00	5.62	6.46	10.03	11.93	14.75	
-0.49	2.63	2.69	2.53	2.55	2.67	3.15	3.31	5.00	5.73	6.98	
-0.55	2.95	3.02	2.53	2.58	2.71	3.25	3.37	5.17	5.85	7.13	
-0.52	3.23	3.30	2.53	2.57	2.71	3.25	3.44	5.15	5.95	7.18	
-0.11	3.25	3.27	2.73	2.58	2.57	2.67	2.67	3.52	3.75	4.30	
-0.35	3.38	3.43	2.77	2.58	2.61	2.88	2.74	4.01	4.04	4.71	
-0.50	3.64	3.71	2.77	2.61	2.70	3.12	2.91	4.46	4.78	5.32	
-0.45	3.86	3.92	2.78	2.62	2.70	3.14	3.04	4.51	5.14	5.95	
-0.54	4.16	4.23	2.78	2.67	2.77	3.27	3.24	4.87	5.34	6.51	

The top row and the left most column of the cumulative distance matrix are the two series we wish to warp. The interior cells of the matrix are the squared Euclidean distances between the values of Series A and Series B plus the minimum value of the adjacent cells that precede it horizontally, diagonally, and vertically, as given by Equation 1.

$$d_{i,j} = (A_i - B_j)^2 + \min(d_{i-1,j-1}, d_{i,j-1}, d_{i-1,j}) \quad (1)$$

To calculate these distances, we begin with the cell in the top row and first column. We calculate its value as $1.10 = (0.01 - (-1.04))^2 + \min(0, 0, 0)$, because there are no preceding values. We then proceed horizontally, diagonally, and vertically to fill out the matrix. Consider, for example, the cell in seventh row and sixth column. We calculate its value as $2.88 = (0.21 - (-$

$0.35)^2 + \min(2.61, 2.57, 2.67)$. Now consider the cell in the last row and last column. We calculate its value as $6.51 = (0.63 - (-0.34))^2 + \min(5.14, 5.34, 5.95)$.

To find the warped series that are most closely aligned, we proceed in reverse from the cell in the last row and last column to the cell in the first row and first column, always moving to the prior adjacent cell with the minimum distance. The path that best aligns the two series, which is shown in red, reveals whether the series proceed in lockstep or if one series proceeds faster than the other at some of the steps.

Although we work backwards to identify the most closely aligned warped series, our goal is to synchronize their forward progression. If the two series proceed at the same pace, the synchronized path advances from the upper left to the lower right along the cells in the diagonal. If Series A proceeds at a faster pace than Series B, the synchronized path moves to the vertical adjacent cell to give Series B time to catch up, as in the case with the cell in the fourth row and third column. If Series B proceeds at a faster pace than Series A, the synchronized path moves to the horizontal adjacent cell as in the case with the cell in the sixth row and sixth column.

Exhibit 3 shows the warped series. Notice that the warped series require 11 steps whereas the original series has only 10 steps. This additional step is required because each series moves without the other one time.

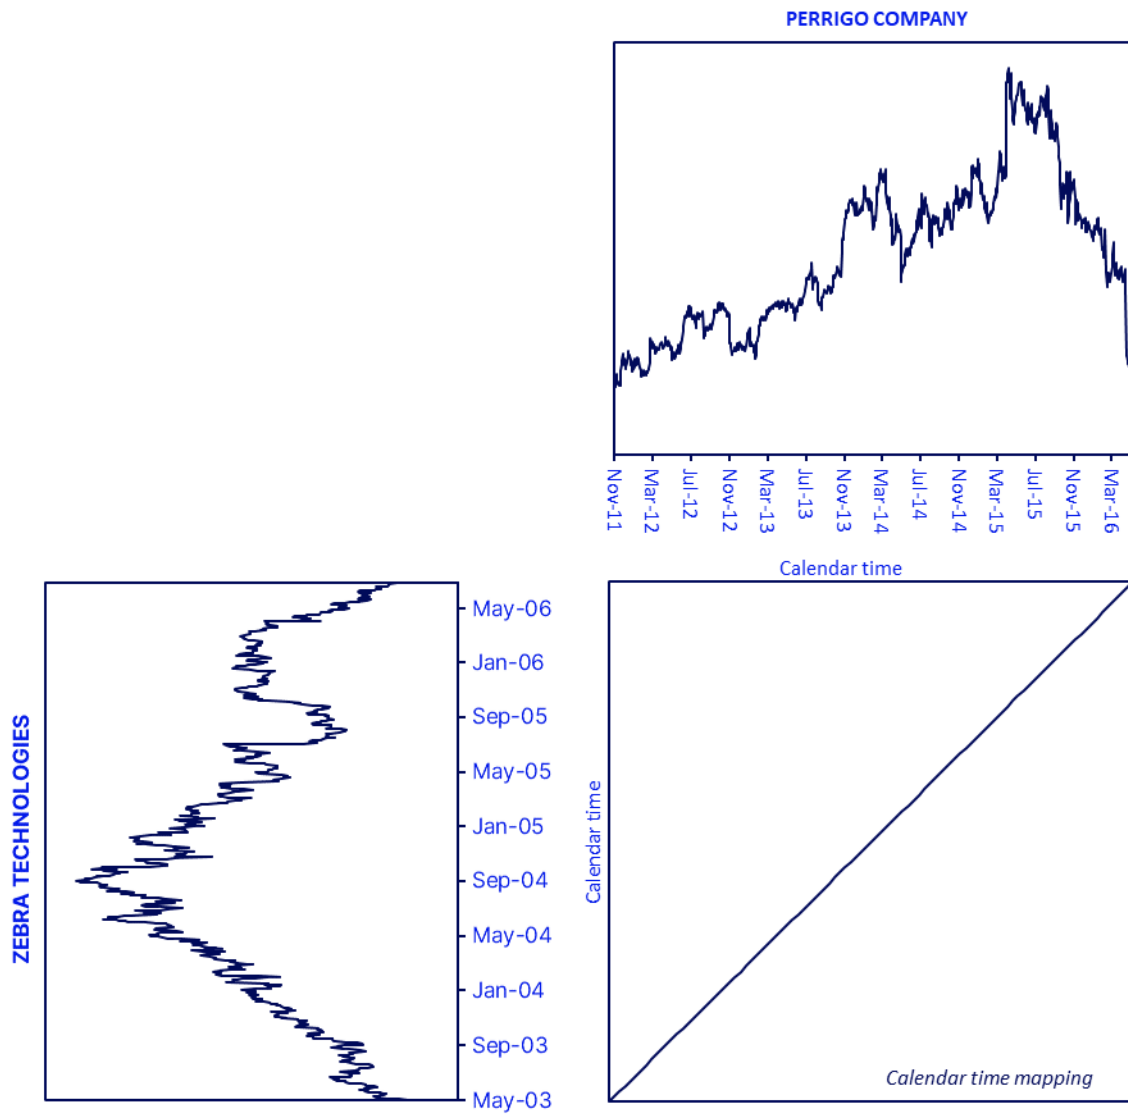
Exhibit 3: Warped Series

Step	Warped Series A	Warped Series B
1	0.01	-1.04
2	0.07	-1.12
3	-0.55	-0.49
4	-0.55	-0.55
5	-0.33	-0.52
6	-0.15	-0.11
7	0.21	-0.11
8	-0.09	-0.35
9	0.81	-0.50
10	0.37	-0.45
11	0.63	-0.54

Exhibits 4, 5, and 6 illustrate the transformation effect of warping based on two bubbles that occurred in our historical sample according to the bubble definition we describe in the next section: Zebra Technologies, which occurred over three years and two months from May 26, 2003 through July 21, 2006; and Perrico Company, which occurred over four and a half years from November 9, 2011 through May 13, 2016.³

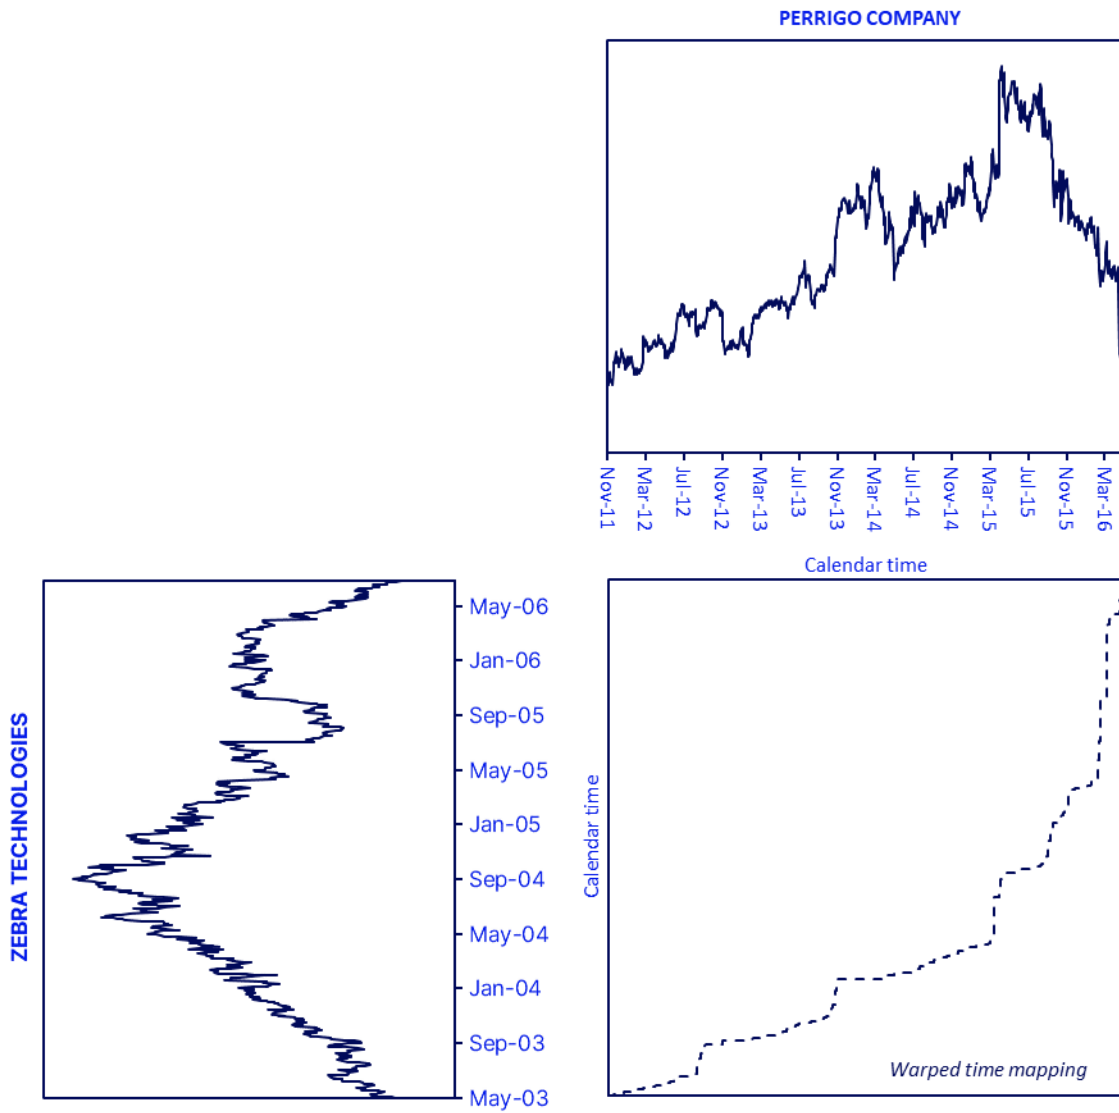
The lower left and upper right graphs in Exhibit 4 show the calendar time progression of these two bubbles. The lower right graph links the fraction of each bubble's duration that has elapsed since their inceptions. It reveals that when we express the phases of bubbles as percentages of elapsed calendar time they are perfectly aligned when mapped in calendar units.

Exhibit 4: Calendar Time Bubbles Linked in Elapsed Time



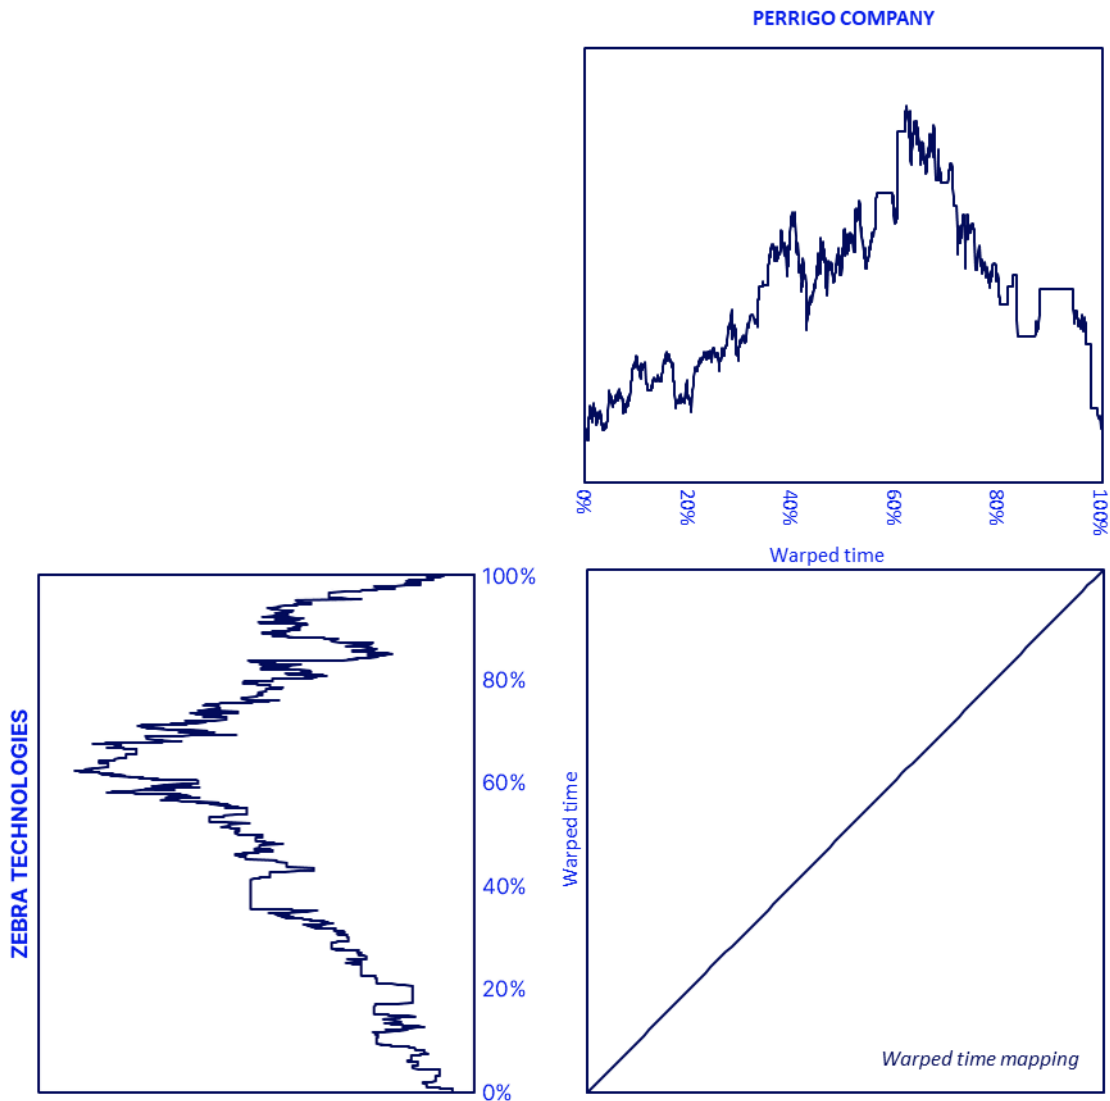
The lower left and upper right graphs in Exhibit 5 show the calendar time progression of the bubbles, just like the graphs in Exhibit 4. However, the lower right graph now links them in warped units. It shows that we must bend calendar time to align the bubbles when we express their joint progression in warped units.

Exhibit 5: Calendar Time Bubbles Linked in Warped Time



The lower left and upper right graphs in Exhibit 6 are expressed in warped units. The lower right graph shows that the bubbles are synchronously aligned when we link them in warped units. Moreover, the shapes of the two bubbles are nearly identical when plotted in warped units, unlike their shapes when plotted in calendar units.

Exhibit 6: Warped Time Bubbles Linked in Warped Time



These graphs illustrate why warping is necessary to synchronize bubbles. When we map bubbles in percentages of elapsed time, we only synchronize their durations. When we map them in warped time, we synchronize their entire shapes, which allows us to account for the pace at which they proceed as well as irregularities in their patterns.

Data and Methodology

Data

Our data comprises the total returns of all the stocks in the S&P 500 Index as well as three types of stock characteristic data: investor behavior data, price-based data, and fundamental data, as shown in Exhibit 7. The investor behavior category comprises flows and holdings indicators produced by State Street Associates and pertain to the overall GICS industry to which a given stock belongs. All other data are specific to the individual stock.

Exhibit 7: Stock Characteristics

Indicators	Source	Description
Investor Behavior Indicators		
Sentiment	State Street Associates	Stock-level sentiment (rolling 30 day average)
Disagreement	State Street Associates	Stock-level disagreement (rolling 30 day average)
Industry flows	State Street Associates	Industry-level flows (rolling 20 day average)
Industry holdings	State Street Associates	Industry-level holdings
Price-Based Indicators		
Momentum	QAD DataSteam	Total return - past 1 year
Reversal	QAD DataSteam	Total return - past 20 days
Volatility 60 day	QAD DataSteam	Standard deviation - past 60 days
Market beta	QAD DataSteam, Fama/French 3 Factors	Stock beta relative to market factor
Value beta	QAD DataSteam, Fama/French 3 Factors	Stock beta relative to value factor
Size beta	QAD DataSteam, Fama/French 3 Factors	Stock beta relative to size factor
Fundamental Indicators		
Dividend yield	QAD WorldScope PIT	Annual dividend per share/price per share
P/E multiple	QAD WorldScope PIT	Price per share/quarterly earnings per share
EPS 1 year growth	QAD WorldScope PIT	Percentage change in EPS - last 12 months
Net margin	QAD WorldScope PIT	Net income/total revenue (x 100)
Cash and equivalents/of total assets	QAD WorldScope PIT	Cash and marketable securities/total current assets
Long-term debt/common equity	QAD WorldScope PIT	Long term debt/shareholder equity
Fixed charge coverage ratio	QAD WorldScope PIT	Earnings before interest and taxes/fixed charges
Cash earnings return on equity	QAD WorldScope PIT	Operating cash flow/equity
Sales estimate	QAD IBES	Mean and standard deviation across analysts
Pretax profit estimate	QAD IBES	Mean and standard deviation across analysts
Earnings per share estimate	QAD IBES	Mean and standard deviation across analysts
Cash flow per share estimate	QAD IBES	Mean and standard deviation across analysts

All the data is for the period beginning January 1, 1973 and ending May 16, 2023. Additionally, we standardize the data by converting the observations to cross-sectional percentile ranks.

Bubble Definition

We define a bubble as an event in which the total return index of a stock increased 50% or more from its previous low point and then declined by 50% or more from its previous peak. The bubble is deemed to have ended when the return index reached a new low point prior to recovering to 30% below its prior peak. We identify the start date of the bubble as the most recent time the index value was as low as the value at the conclusion of the bubble.

We identified 2,638 bubbles for 866 stocks from January 1, 1973 through May 16, 2023.

Training and Prediction

Training Process

Our training process proceeds as follows.

1. We select randomly without replacement 10% of the bubbles from the full sample to use as the holdout sample, and we use the 90% complement as the training sample.
2. We warp bubble 1 and bubble 2 from the training sample into 21 time steps of 5% intervals from 0% to 100% of warp time.

3. For time step 1, which is 0% of warp time and therefore the inception of the bubble, we record a vector of the stock characteristics, previously shown in Exhibit 7 and expressed as cross-sectional percentile ranks.
4. We repeat this process for bubble 1 with every other bubble, bubble 2 with every other bubble, bubble 3 with every bubble, and so on until we have warped every bubble pair in our training sample, recording the vectors of stock characteristics along the way.
5. We then repeat this entire process for time steps 2 through 21, producing distributions of stock characteristics for every time step of the warped bubbles.
6. We repeat this process 10 times, thereby evaluating all the bubbles in the full sample.

It is important to note that warp time is not universal. It is unique to each bubble pair.

Holdout Sample Prediction

Next, we use a statistic called the Mahalanobis distance to estimate the time step of a bubble from the holdout sample by comparing its stock characteristics to the distributions of the stock characteristics we observed during each time step of the training sample bubbles. The Mahalanobis distance is given by Equation 2.

$$d_s(x_{i,t}) = (x_{i,t} - \mu_s)\Sigma_s^{-1}(x_{i,t} - \mu_s)' \quad (2)$$

In Equation 2, $d_s(x_{i,t})$ is the Mahalanobis distance of the stock characteristics of bubble i observed at time t in the holdout sample from the stock characteristics of warp step s from the training sample; $x_{i,t}$ is a vector of the bubble stock characteristics at the time it is observed in the holdout sample; μ_s is a vector of the average of the stock characteristics at warp step s

from the training sample; and Σ_s is the covariance matrix of the stock characteristics at warp step s from the training sample.

The vector $(x_{i,t} - \mu_s)$ measures how different the stock characteristics of the currently observed bubble are from the average characteristics of a time step from the training sample. By multiplying this vector by the inverse of the covariance matrix, we capture the interaction of the characteristics associated with the training sample time step. By multiplying this product by the transpose of the vector we consolidate the outcome into a single number, which represents the covariance-adjusted distance between the stock characteristics of the holdout sample bubble and the average characteristics of the various time steps from the training sample. Based on information about a stock experiencing a bubble at time t , we estimate its current warp time step as the step s to which its characteristics are least distant.

Once we identify the most likely time step of the currently observed bubble, we calculate the percentage price appreciation remaining to its peak, assuming it has not yet reached its peak, as:

$$\text{Remaining price percentage to peak} = \frac{\text{peak date price} - \text{prediction date price}}{\text{peak date price} - \text{bubble start date price}} \quad (3)$$

If the bubble is in its selloff phase, we calculate the remaining percentage price depreciation to the conclusion of the bubble as:

$$\text{Remaining price percentage to conclusion} = \frac{\text{conclusion date price} - \text{prediction date price}}{\text{conclusion date price} - \text{peak date price}} \quad (4)$$

We calculate these remaining price changes in both warped time and calendar time to compare the efficacy of these dimensions for assessing a bubble's progression.

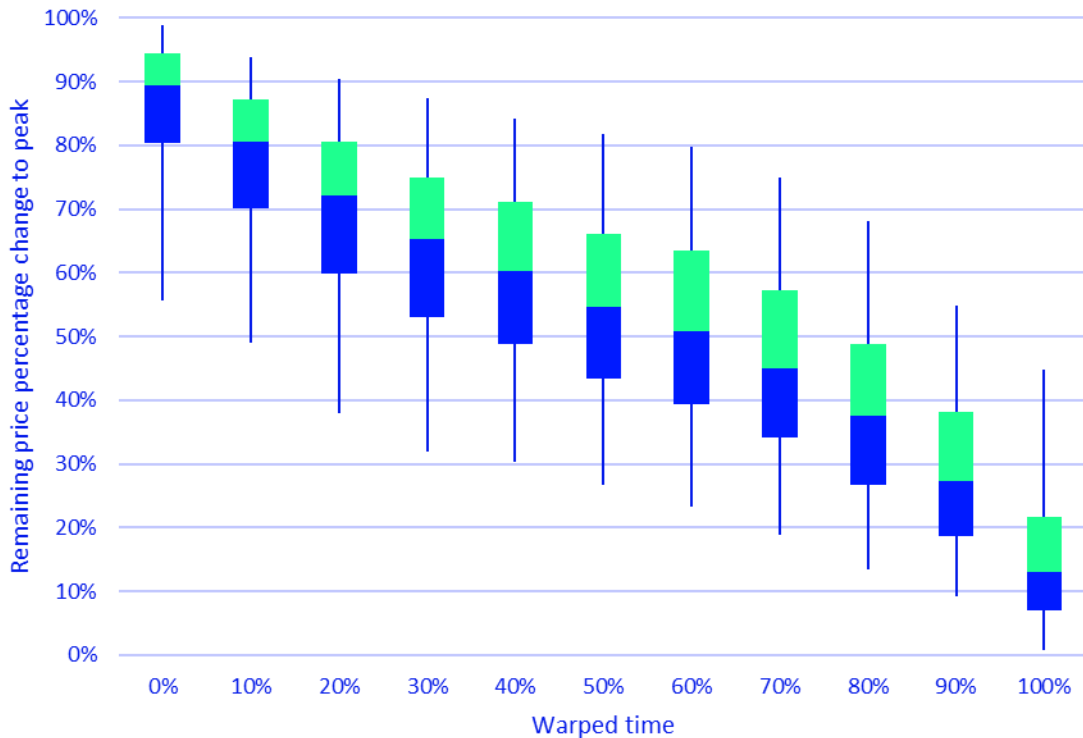
Results

As we mentioned previously, we repeat the training process and holdout estimation 10 separate times to produce our results. Each time we randomly select a 10% holdout sample without replacement. For each training sample, we consider 6,260,765 bubble pairs ($2,638 \times 2,637 \times 0.90$), and we evaluate 264 bubbles ($2,638 \times 0.10$) in the holdout sample. Given that we repeat the process 10 times, in total we consider all 2,638 bubbles in our sample.

Bubble Phases: Warped Time versus Calendar Time

Exhibit 8 shows the composite distribution of subsequent price appreciation to the bubble peak for bubbles in the holdout samples that were estimated to be in each stage of the run-up phase. The horizontal axis reflects the percentage of elapsed warped time for each time step from inception to peak. The vertical axis represents the remaining percentage of price appreciation from inception to peak that subsequently occurred. The box plots show the 25th, 50th, and 75th percentile values, with lines extending to the 5th and 95th percentiles. These distributions comprise the bubble step estimates made for every day of every bubble in the composite holdout sample.

Exhibit 8: Elapsed Warp Time versus Realized Percentage Change from Inception to Peak

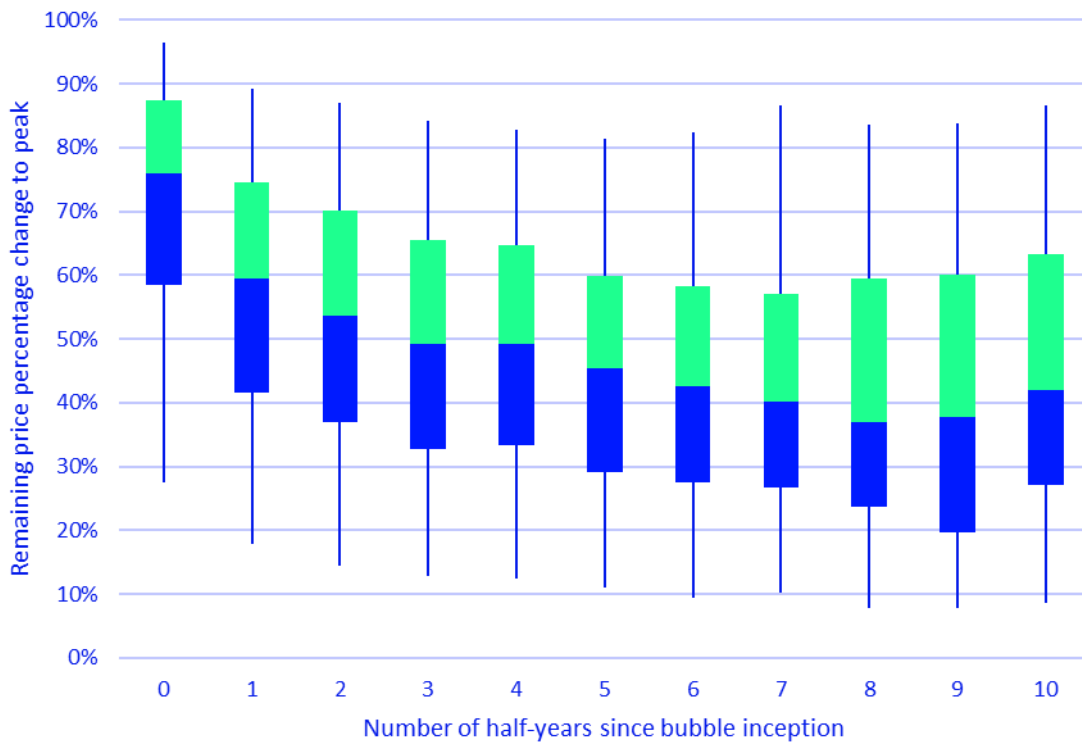


Consider, for example, 0% warp time to peak. This warp time represents holdout sample observations for which the stock's attributes at the time suggested it is most likely at the inception of a bubble. In 50% of these cases, the bubble had more than 89% of its total price appreciation remaining (the median of the distribution is 89%). When warp time was estimated to be at 50% of the bubble appreciation phase, the realized median percentage to peak was 54%. And when the bubbles were estimated to be at their peaks, the median realized percentage change to peak was 13% across all bubbles.

Exhibit 9 presents the distribution of remaining price appreciation if, instead of using warp time, we estimate a bubble's stage of progression based on the elapsed calendar time since inception. Because the calendar duration of the bubbles can vary dramatically, we

calibrate the horizontal axis as the number of six-month periods from the bubbles' inceptions up to three years, which captures the total duration, or at least a large fraction of it, for most bubbles in our sample.

Exhibit 9: Elapsed Calendar Time versus Realized Percentage Change from Inception to Peak



In contrast to the relationship between elapsed warp time and remaining percentage price change to peak, which shows a pronounced downward and relatively steady slope, the relationship between elapsed calendar time and remaining percentage change to peak has a much shallower and less monotonic slope, and the dispersion around the median outcomes is far wider.

Next, we consider the remaining percentage change to conclusion for bubbles that have already reached their peaks. Exhibit 10 compares estimates of the bubbles' elapsed times as percentage changes from peak to conclusion in warped units, shown along the horizontal axis, with the distributions of the realized percentages changes from peak to conclusion that remained, expressed as boxplots, shown on the vertical axis.

Exhibit 10: Elapsed Warp Time versus Realized Percentage Change from Peak to Conclusion

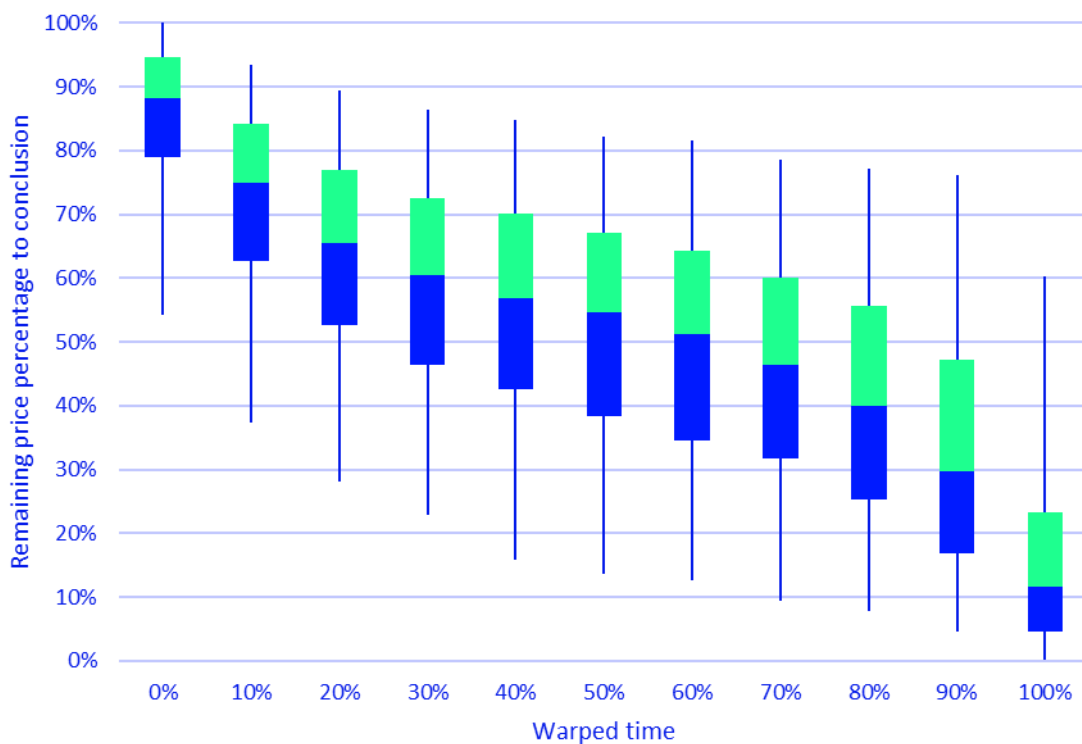


Exhibit 10 shows a similar steady downward slope for peak to conclusion as we observed for inception to peak. For example, when warp time estimated that the bubbles were at their peaks (0%), the median percentage change to conclusion was 89%. When warp time indicated that 50% of the bubbles' times to conclusion had transpired, the median percentage change to conclusion was 54%. And when warp time indicated that the bubbles had reached

their conclusions, the median percentage change to conclusion was 11%. However, the bands around the median estimates are wider than they are for bubble runups, indicating that the warp predictions for bubble selloffs are less reliable than they are for bubble runups, except near their conclusion.

Exhibit 11 presents the same selloff comparison but based on calendar time instead of warped time.

Exhibit 11: Elapsed Calendar Time versus Realized Percentage Change from Peak to Conclusion

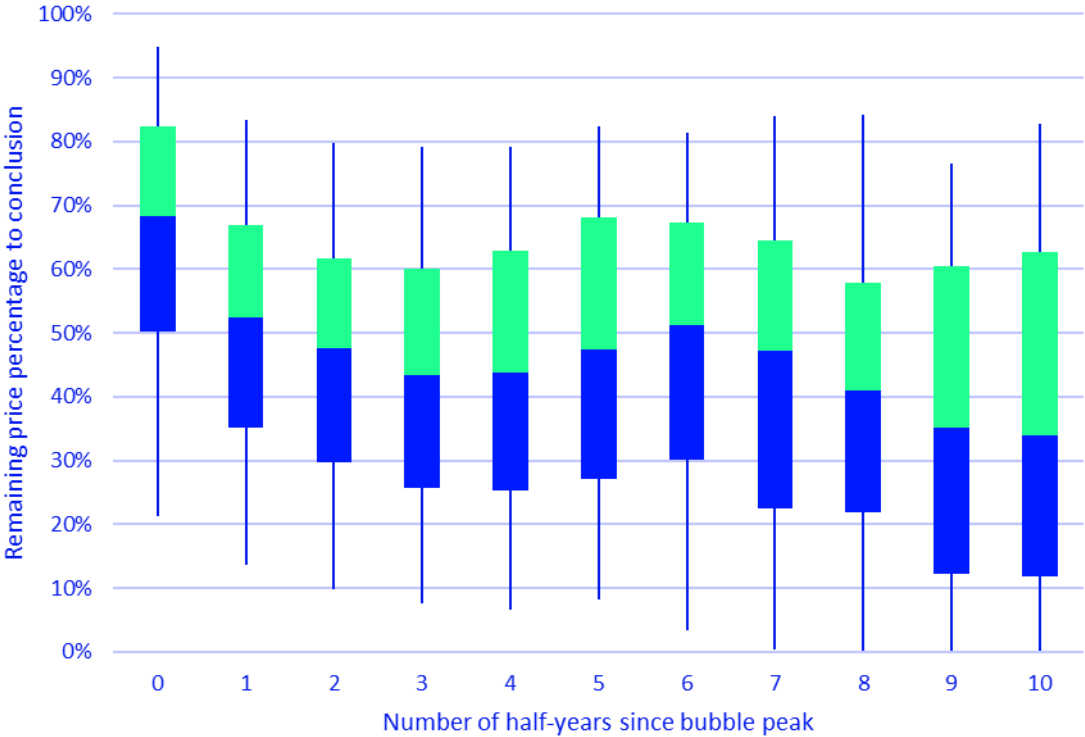


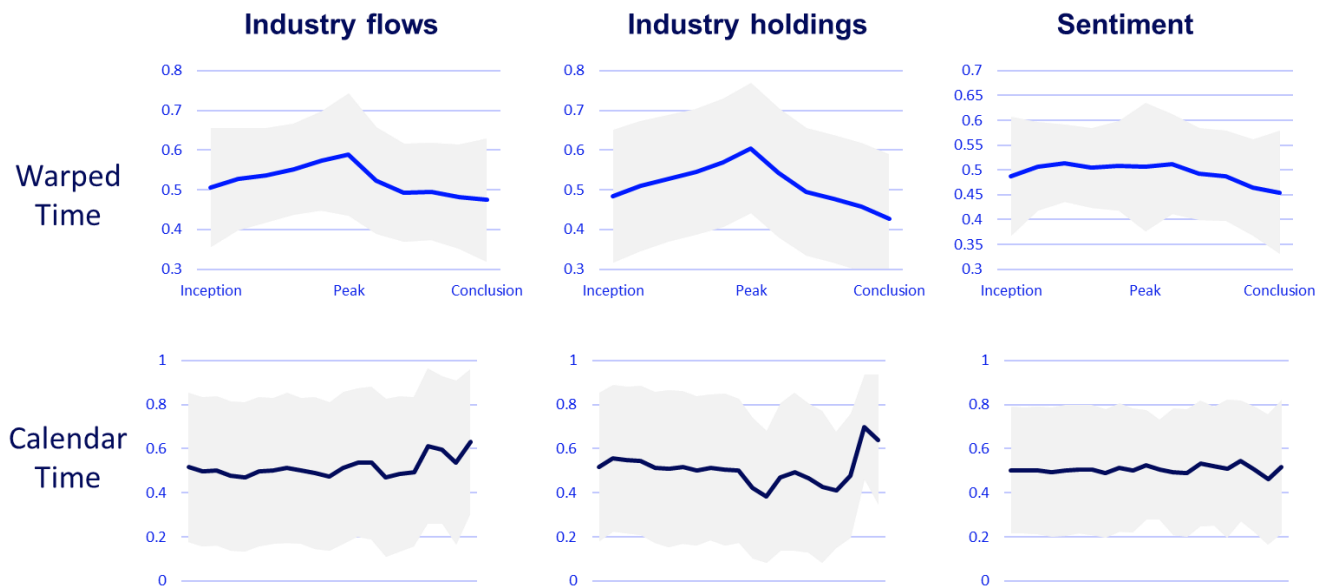
Exhibit 11 reveals that calendar time gives a much less reliable estimate of the percentage change remaining from a bubble’s peak to conclusion, as indicated by the wider

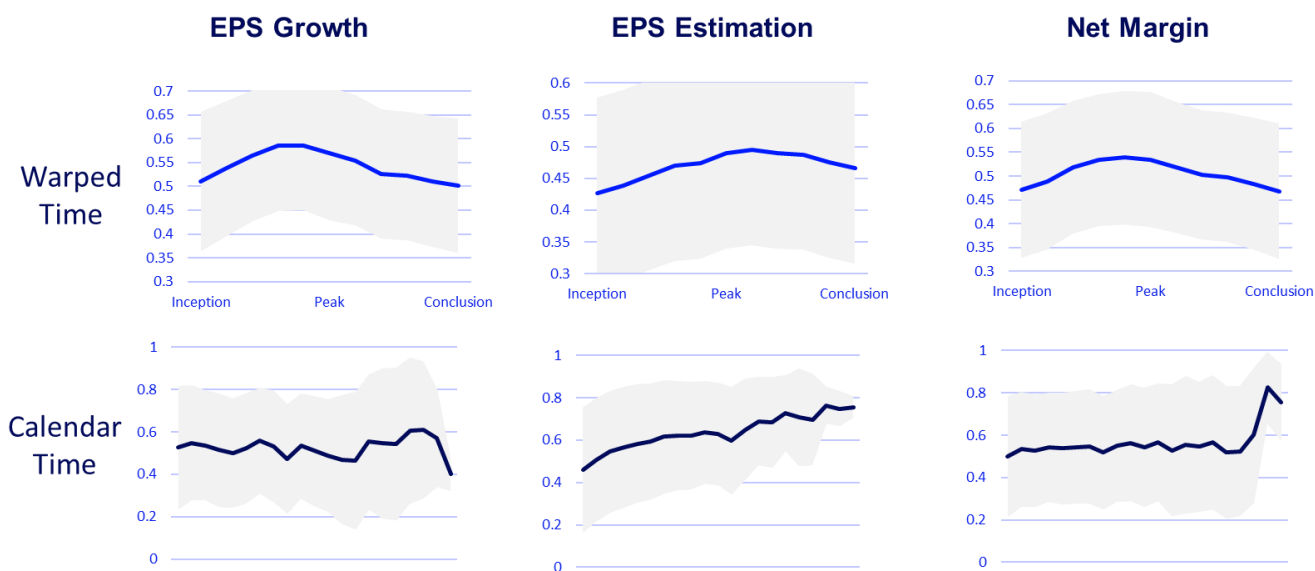
bands around the median estimates and the significantly more shallow and less monotonic slope.

Stock Characteristics: Warped Time versus Calendar Time

As further evidence that warping enables us to evaluate bubbles more effectively, we show in Exhibit 12 how certain stock characteristics progressed in warped time (top panels) versus calendar time (bottom panels).

Exhibit 12: Progression of Selected Stock Characteristics in Warped Time and Calendar Time





The results we have presented thus far assume that we knew the bubbles in the holdout samples were underway, but that we did not know at what time step they were along their journeys; hence, these results are not fully out of sample. Next, we show how to detect the inception of bubbles as well as how far they have progressed.

Out-of-Sample Testing

We test three market neutral trading rules to determine if observing bubbles in warped time enables investors to detect their arrival and the phase of their progression. For these tests we include an additional time step to represent times when a stock was not experiencing a bubble. Just as we use the Mahalanobis distance to detect a bubble's time step, we similarly use it to detect "no bubble" periods. These tests are therefore fully out of sample. They presume no foreknowledge of a bubble's existence nor how far a bubble has progressed.

We begin by learning bubble characteristics from data beginning January 1, 1973 through December 31, 1999, and we expand this window each month as we move forward in

our testing sample. We rebalance the positions monthly, and we weight them according to their capitalizations. We test three trading rules: one in which we seek to participate in bubble run-ups and exit before sacrificing accumulated gains; one in which we seek to exploit over-reaction near the end of bubble selloffs; and one in which we combine these trading rules. Our rebalancing occurs at each month end from January 31, 2000 through May 31, 2023. Hence, our measurement period runs from February 2000 through June 2023.

1. Run-up Trading Rule

- Purchase bubble stocks weighted by their capitalization that are estimated to be between 20% and 80% of the elapsed warp time from the bubble inception to the bubble peak.
- Sell S&P 500 in equal amount to create market neutral exposure.

2. Over-reaction Trading Rule

- Purchase bubble stocks weighted by their capitalization that are estimated to be between 80% and 100% of the elapsed warp time from the bubble peak to the bubble conclusion.
- Sell S&P 500 in equal amount to create market neutral exposure.

3. Run-up and Over-reaction Trading Rule

- Purchase bubble stocks weighted by their capitalization that are estimated to be between 20% and 80% of the elapsed warp time from the bubble inception to the bubble peak and bubble stocks that are estimated to be between 80% and 100% of the elapsed warp time from the bubble peak to the bubble conclusion.
- Sell S&P 500 in equal amount to create market neutral exposure.

Exhibit 13 shows that all three trading rules generated profits on average over the testing period from January 2000 through June 2023. The run-up strategy underperformed during the Dot Com Bubble, most likely due to systematic influences.⁴ The over-reaction trading rule, by contrast, performed extremely well during this period, suggesting systematic negative over-reaction by investors. On average, the over-reaction trading rule produced a higher return than the run-up trading rule but with considerably more volatility. The combined trading rules generated a better return than the run-up trading rule, but they did not generate as large a profit as the over-reaction trading rule. The performance of the combined trading rules is more similar to that of the run-up trading rule than the over-reaction trading rule because run-up periods tend to last longer than over-reaction periods; therefore, they tend to dominate the combined sample.

Exhibit 13: Cumulative Excess Returns

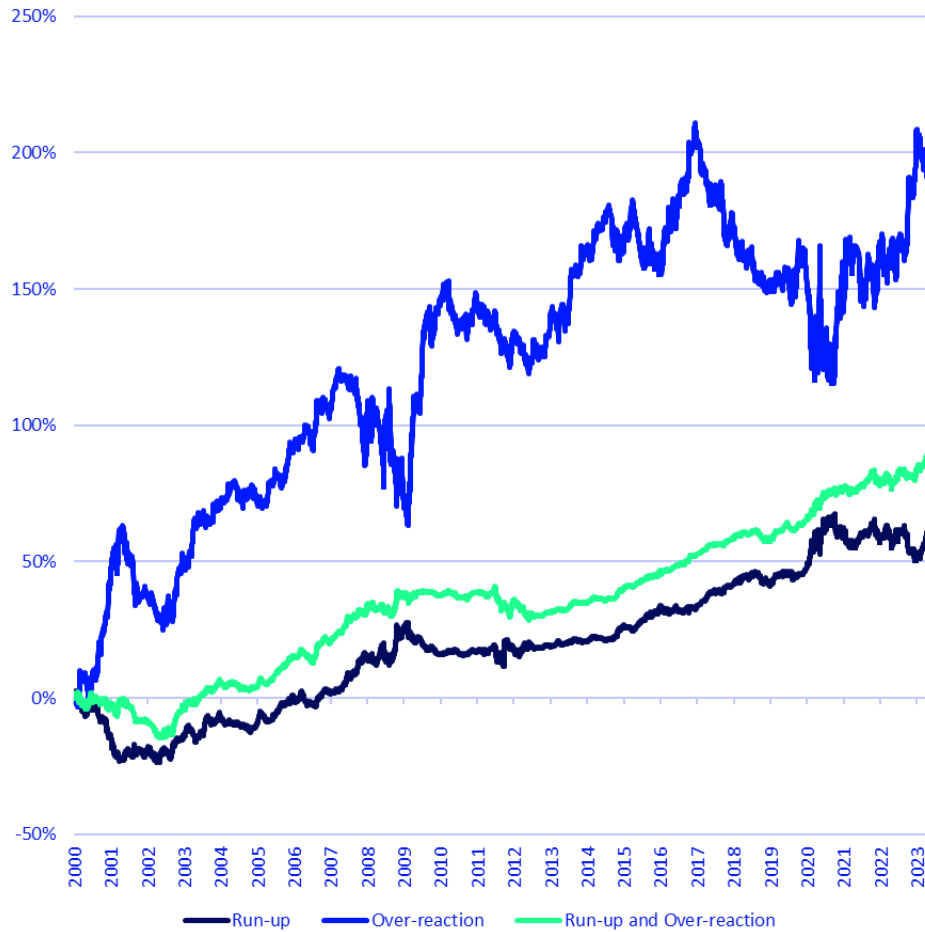


Exhibit 14 shows the annualized details of these three trading rules. It shows that the combined trading rules produced the best risk-adjusted outcome by a significant margin.

Exhibit 14: Trading Rules Annualized Return and Risk

	Run-up	Over-reaction	Combined
Return (annualized)	2.2%	5.4%	2.9%
Risk (annualized)	5.7%	10.1%	3.9%
Information Ratio	0.39	0.53	0.74

Conclusion

We employed a technique called dynamic warp analysis to convert the calendar progression of bubble pairs into synchronized warped time steps. For each warped time step we recorded a variety of stock characteristics across all the warped bubble pairs in a training sample. We then used the Mahalanobis distance to measure the relative proximity of an out-of-sample bubble observed at an unknown time step to the distribution of stock characteristics of each time step from the training sample. Next, we showed that observing bubbles in warped time gives much more reliable estimates of the realized remaining percentage changes from inception to peak and from peak to conclusion than observing bubbles in calendar time. We also showed that the differences across stock characteristics conforms more closely to the stylized image of a bubble when observed in warped time as opposed to calendar time. We cautioned that, although our results presumed we had no foreknowledge of how far a bubble had progressed along its journey, they did presume we had foreknowledge that a bubble was underway.

We then tested three trading rules fully out of sample to determine if observing bubbles in warped units has the potential to generate profits. We considered a rule to determine if one could participate in bubble run-ups and exit sufficiently early to preserve accumulated gains, as

well as a rule that exploited investor over-reaction during the final phase of bubble selloffs. We tested these two trading rules independently and in combination. These tests presumed we had no foreknowledge of whether a bubble was underway, and if it was underway, how far along it was in its journey. Our tests offer compelling evidence that dynamic warp analysis has the potential to enable investors to profit by detecting new bubbles and by revealing how far they have progressed.

References

Juang, B.-H. 1984. "On the Hidden Markov Model and Dynamic Time Warping for Speech Recognition—A Unified View." AT&T Bell Laboratories Technical Journal, September.

[Eugene Fama on Finance](#), an EconTalk podcast, January 30, 2012.

Greenwood, R, A Shleifer, and Y. You. 2019. "[Bubbles for Fama](#)." Journal of Financial Economics 131 (1): 20-43.

Notes

This material is for informational purposes only. The views expressed in this material are the views of the authors, are provided "as-is" at the time of first publication, are not intended for distribution to any person or entity in any jurisdiction where such distribution or use would be contrary to applicable law and are not an offer or solicitation to buy or sell securities or any product. The views expressed do not necessarily represent the views of Windham Capital Management, State Street Global Markets®, or State Street Corporation® and its affiliates.

¹ Nobel Laureate Eugene Fama famously asserted in a 2010 interview "It's easy to say prices went down, it must have been a bubble, after the fact. I think most bubbles are twenty-twenty hindsight. . . . People are always saying that prices are too high. When they turn out to be right, we anoint them. When they turn out to be wrong, we

ignore them.” Greenwood, Shleifer, and Who (2019) argue that certain features of stocks correspond to a heightened probability of bubbles.

² See, for example, Juang (1984).

³ We chose these two historical examples arbitrarily for the purpose of illustration. This illustration does not imply any views about these stocks.

⁴ We could neutralize the market effect associated with bubble identification by considering returns in excess of the market’s return, but we believe that absolute bubbles resonate more with investors than relative bubbles.