

Volatility After-Effects: Evidence from the Field

Elise Payzan-LeNestour ^a, Lionnel Pradier ^a, and Tālis J. Putniņš ^{b*}

^a *University of New South Wales*

^b *University of Technology Sydney and Stockholm School of Economics in Riga*

May 7, 2015

Abstract

We propose and test the idea that investor perceptions exhibit volatility ‘after-effects’ whereby perceived volatility is distorted after prolonged exposure to extreme volatility levels. Using VIX to measure perceived volatility in S&P 500 stocks, we find evidence of significant perceptual distortions in the aftermath of volatility regimes, consistent with the after-effect theory and recent experimental evidence. These distortions are larger after both stronger and longer volatility regimes, and are absent after volatility changes that are not preceded by extreme volatility levels, consistent with the after-effect theory and inconsistent with alternative explanations. Our study shows that perceptual biases can have a significant distortionary effect on asset prices, even in very actively traded financial securities.

JEL classification: D83, D87, G02, G14, G17

Keywords: after-effect, perception bias, volatility, VIX, neuroeconomics, neurofinance

* Payzan-LeNestour: UNSW Australia Business School, University of New South Wales, NSW 2052, Australia; email: elise@unsw.edu.au; phone: +61 2 9385 4273. Pradier: UNSW Australia Business School, University of New South Wales, NSW 2052, Australia; email: lionnel.pradier@wanadoo.fr. Putniņš: UTS Business School, University of Technology Sydney, PO Box 123 Broadway, NSW 2007, Australia; email: talis.putnins@uts.edu.au; phone: +61 2 95143088.

The Internet Appendix that accompanies this paper can be found at <http://ow.ly/He7v0>.

We thank the Securities Industry Research Centre of Asia-Pacific, and Thomson Reuters for providing access to data used in this study. We also thank seminar participants at the Stockholm School of Economics in Riga and the Baltic International Centre for Economic Policy Studies for helpful comments and suggestions.

1. Introduction

Conventional economic theory assumes no distortions in the way agents perceive realized asset returns. Yet, McFadden (1999) points out that perception errors are important and should be accounted for, as they explain many behavioral anomalies. Here we seek to follow this lead, by postulating and testing the presence of important distortions in investor perceptions of asset return volatility.

A large body of literature in neurophysiology has documented that after prolonged exposure to a stimulus, a perception bias subsequently emerges which creates the illusion of an opposite stimulus. This bias is called *after-effect*. For instance, after viewing a red square a gray square appears greenish (Hurvich and Jameson, 1957); after a few moments looking at the downward flow of a waterfall, the static rocks to the side appear to ooze upward (Barlow and Hill, 1963); and prolonged viewing of a male face makes subsequently seen androgyne faces appear more feminine than they normally would (Webster, Kaping, Mizokami, and Duhamel, 2004; Rutherford, Chattha, and Krysko, 2008). After-effects appear to be ubiquitous. They occur for stimuli of all stripes, running the gamut from simple stimuli to highly abstract properties such as the perceived numerosity of dots in patches (Burr and Ross, 2008). They also occur across different time horizons—some after-effects occur in the order of a few seconds whereas others have a daily or monthly horizons (Delahunt, Webster, Ma, and Werner, 2004; Webster, McDermott, and Bebis, 2007).

On the theoretical side, Woodford (2012) proposes that the after-effect phenomenon is only one instantiation of *neuronal adaptation*, the principle by which the brain maximizes accuracy of perceptions, subject to a limit on information-processing capacity.. Neuronal adaptation has two central properties: (1) diminishing sensitivity to value contrasts that are far away from the prior mean stimulus (the stimulus level that is expected to be encountered most often); (2) reset to the mean or reference-dependence, i.e., the brain perceives a given stimulus level with respect to the prior mean level, not the stimulus level itself. As such, neuronal adaptation has two major implications for decision-making under uncertainty: (1) predicts the shape of the *value function* featured by *Prospect Theory* (Kahneman and Tversky, 1979), which has received much emphasis, and (2) predicts the after-effect phenomenon, which is the novel focus of this study.

Inasmuch as after-effects appear to be not only ubiquitous but also necessary given our neurobiological constraints, it seems natural to postulate that they affect investor perceptions of asset return volatility. We therefore propose that investors perceive volatility to be lower than actual after prolonged exposure to high volatility levels, and higher than actual after prolonged exposure to low volatility levels. We further conjecture that this perception bias affects asset prices. In this paper we provide strong evidence for this conjecture.

Recent experimental work documents the presence of strong volatility after-effects in the laboratory. Payzan-LeNestour, Balleine, Berrada, and Pearson (2014) design a computer task that is a stylized version of what a trader experiences on a Bloomberg terminal. Task participants are shown a time-series representing trajectories of a stock market index over a year at a daily frequency. They are asked to report how volatile they perceive each trajectory. By design, the volatility of the test trajectory is always 10%. However, the task participants' perceptions differ from 10% in a systematic way. Perceived volatility is 32% higher after prolonged (50 seconds) exposure to low volatility (2%) trajectories than after prolonged exposure to high volatility (45%) trajectories. Hence after-effects appear to distort perceptions of volatility in the laboratory.

What about in financial markets? Do such after-effects distort the VIX, which reflects investor forecasts of volatility? Our empirical evidence indicates the answer is a definite yes. After-effects significantly influence the VIX and thus underlying asset prices. This finding is not a foregone conclusion because although the *average* individual's perception may be distorted, asset prices are determined by the *marginal* trader, who may well be sufficiently sophisticated so as to not suffer systematic perceptual distortions.

We focus on the change in VIX when transitioning from a state of either very low or very high volatility to a neutral volatility state (neither high nor low). We report that the part of the change in VIX that cannot be attributed to changes in either fundamentals or risk aversion levels, can however be attributed to the after-effect. To establish this, we construct a variable that equals +1 (-1) on the day that a prolonged high (low) volatility state reverts to a neutral level, and 0 at all other times. We find that this variable is a

significant determinant of changes in VIX. The change in VIX in the aftermath of a low volatility regime is higher than the corresponding change in the aftermath of a high volatility regime. The impact of a change of regime on VIX is as large as 3.5% or 76 bps (the same impact on VIX as a 1% change in the S&P 500, which is the most important predictor of change in VIX—more on this below). This finding is consistent with our conjecture that investors' perception of volatility is higher in the aftermath of a prolonged period of very low volatility than in the aftermath of a prolonged period of very high volatility, all other things being equal.

Furthermore, the significance of our indicator variable increases linearly with the strength of the regimes, as it should if after-effects drive our results. It is maximal for regimes featuring extremely high or low levels of volatility, and nil for regimes in which the volatility levels do not depart markedly from the levels observed during the neutral states (around 13.5% on average in our data). That the magnitude of the effect increases linearly with the strength of the regimes conforms to what the after-effect theory predicts.

Additionally, we find that the significance of our indicator variable increases with the duration of the regimes (the exposure time to very high or low volatility levels). This again conforms to the after-effect theory. Experiments by psychologists have indeed documented that the magnitude of the after-effect builds up logarithmically with the duration of exposure to a given stimulus (Magnussen and Johnsen, 1986; Hershenson, 1989; Leopold, Rhodes, Muller, and Jeffery, 2005).

Together these findings constitute strong evidence for our conjecture that perceptual after-effects bias the VIX. To our knowledge, no competing theory can explain the collection of empirical findings.

Importantly, our results are robust to assuming that the agents have adaptive expectations about the S&P500 volatility level. Our benchmark model assumes that the agents have *rational* expectations, which seems at odds with the growing body of evidence that investors have *adaptive* or *extrapolative* expectations (i.e., return forecasts are positively correlated with recent returns) and that these forecasts have implications for expected returns (see Greenwood and Shleifer, 2014; Barberis, Greenwood, Jin and Shleifer, 2015; Choi and Mertens, 2013). In fact, modifying our model to account for

potential extrapolative biases in the VIX strengthens the evidence for volatility after-effects that we document here. We elaborate in Section 4.6 (Robustness Checks).

Finally, we provide evidence that the VIX distortions we document here are asymmetric. While the VIX exhibits an abnormal decrease in the aftermath of a high volatility regime, the corresponding VIX increase in the aftermath of a low volatility regime is not apparent in our data. We investigate this finding by revisiting the experimental findings of Payzan-LeNestour, Balleine, Berrada, and Pearson (2014). We run follow-up experimental sessions in the laboratory in which the volatility parameters are similar to those that we observe in the field (low volatility: 7% versus it was 2% in the original experimental sessions; neutral: 13.5% versus 10% in the original sessions; high volatility: 40% versus 45% in the original sessions). Quite strikingly, with those parameter values, the asymmetry that we observe in the field emerges in the laboratory as well.

The absence of after-effects in the aftermath of low volatility regimes suggests that from a perceptual viewpoint, the levels of volatility that prevail during low volatility states (on average 7-9% in our data) do not markedly contrast with the intermediate levels that prevail during the transition states (13.5%).

The current study adds to the growing literature in behavioral finance. Prior behavioral finance studies have documented a number of behavioral biases such as limitations in the number of variables that agents can keep track of or pay attention to.¹ Here, we document a novel behavioral bias, which relates to how investor perceptions of volatility are distorted in the aftermath of volatility regimes. Notably, we do not simply document that this perception bias exists among some individuals, but rather, we show that it has a meaningful impact on asset prices. As such, the current study adds to the literature that has shown that the presence of irrational noise traders can significantly affect stock prices.² One distinctive trait of our study is that the bias we focus on here does not arise from a lack of intelligence in some agents; rather, it is a direct implication

¹ See, among others, Simon (1955), Kahneman (1973), Huberman and Regev (2001), DellaVigna and Pollet (2009).

² See, among others, De Long, Shleifer, Summers, and Waldmann (1990), Lee, Shleifer and Thaler (1991), Shleifer and Vishny (1997), Froot and Dabora (1999), Barberis and Shleifer (2002), Mitchell, Pulvino and Stafford (2002), and Lamont and Thaler (2003).

of the way our perceptual system works and consequently it potentially affects all agents, including very sophisticated arbitrageurs.

The novel contribution of this study is to propose and test the idea that the presence of volatility regimes in itself may contribute to distorting asset prices as per the after-effects channel that we postulate. This idea builds on a large body of data from psychophysics and neurology on human perception in many sensory domains. As such, the current study complements the neurologically grounded economics literature that proposes to augment conventional economic theory with consideration of the fundamental constraints imposed by our brains' hardware (Glimcher, 2011; Woodford, 2012).

As emphasized earlier, the foregoing neuroeconomics work has established that the after-effect phenomenon and the shape of the value function proposed by Prospect Theory are two different instantiations of the same neural principle—namely, neuronal adaptation. While the importance of Prospect Theory in our understanding of decision-making under uncertainty has long been recognized, the current findings compellingly suggest that after-effects are of equal importance.

The evidence of volatility after-effects in the laboratory leads us to test the presence of such perceptual after-effects in asset prices in the field. The results that emerge from the field study then lead us to run follow-up investigations in the laboratory. To our best knowledge, this approach, which involves going from laboratory data to field data and back to the laboratory, is novel in experimental finance.

The rest of the paper is organized as follows. Section 2 explains the after-effect theory. Section 3 details the data and empirical strategy. Section 4 documents the main findings as well as robustness tests. Section 5 documents asymmetry in the after-effects that we identify and reports the main results of the follow-up laboratory experiment. Section 6 concludes.

2. Theory

To explain the phenomenon of after-effects, *opponent-process theory* (see, e.g., Hurvich and Jameson, 1957; Hering, 1964; Griggs, 2009) invokes antagonistic connectivity between pairs of neurons coding for alternative stimulus representations; for

example, pairs of motion-selective neurons coding for upward versus downward motions, pairs of face-selective neurons coding for happy versus sad face expressions or male versus female traits, pairs of color-selective neurons coding for red versus green, and so on. The after-effect follows from an imbalance among the pair of feature-selective neurons. Take for instance the two color representations red versus green. When viewing a red square the neurons coding for ‘red’ are strongly stimulated, while the competing neurons coding for ‘green’ are weakly stimulated (the square looks red). After a few moments of stimulation the neurons coding for ‘red’ show diminished responses, owing to a mechanism of synaptic depression named *neuronal adaptation*. When subsequently viewing an ambiguous (grey) square, the neurons signaling the red color are less stimulated than those signaling the green color, whereby the green perception spontaneously emerges to win the competition (the grey square looks greenish).

Neuronal adaptation reflects how neurons adjust to the mean stimulus level and perceive contrasts around the mean (Kandel, Schwartz, and Jessell, 2000). Such adjustment confers a number of functional advantages to the observer (e.g., Webster, McDermott, and Bebis, 2007; Woodford, 2012). Woodford (2012) shows the mechanism is actually optimal under neurobiological constraints on the degree of precision of people’s awareness of their environment. Benefits of adaptation include maximizing the limited dynamic range available for visual coding and improving visual discrimination. For instance, visual sensitivity adjusts to the mean light level so that the exposure level remains at an appropriate level for perceiving the variations in light around the mean (Barlow, 1972). Adjusting to the mean stimulus level also allows differences around the mean to be more easily distinguished.³

Applying the after-effect theory to a financial decision context, we postulate that investors’ perception of variability in a broad sense—volatility of a time-series as well as variance of a sequence of numbers—involves a pair of variability-selective neurons. After prolonged exposure to low volatility levels, the neurons signaling low volatility would show diminished baseline activity relative to the competing neurons coding for

³ This process might underlie high-level perceptual judgments such as *the other race effect* (Eysenck and Keane, 2013) in face perception, in which we can readily discriminate differences between faces within the ethnic group we are exposed to while faces drawn from novel groups appear similar (Webster, McDermott, and Bebis, 2007).

high volatility. This imbalance would result in subsequent neutral volatility levels (the counterpart of the grey square in the previous example) looking more volatile than they truly are. Likewise, the neurons signaling high volatility would be relatively depressed after being overstimulated in a high volatility regime, resulting in investors perceiving neutral volatility levels as less volatile than they are. The theory therefore predicts that perceived volatility is biased downward (upward) in the aftermath of prolonged exposure to high volatility (low volatility). Figure 1 illustrates this prediction. Notably, standard behavioral theories predict the opposite perception bias. In particular, under *adaptive expectations* (expectations are adjusted by a fraction of the prediction error—the difference between the predicted and realized volatility) and *anchoring* (making insufficient adjustments from a reference point, which could be the previous volatility level), VIX is distorted upward (resp. downward) in the aftermath of a high (resp. low) volatility regime.

< Figure 1 here >

As stressed in the Introduction, several studies document that the magnitude of the after-effect increases with both the intensity of the stimulus to which the agent is exposed to during the adaptation phase as well as the duration of this stimulus. In light of this, we predict that volatility after-effects depend on both the strength and duration of the volatility regimes. The more extreme and the longer the regime, the stronger the neuronal adaptation and hence the larger the after-effect. Finally, after-effects theory does *not* predict a perception bias when transitioning from a neutral state to a state of very high or very low volatility. Thus, we do not expect to see any perception bias when volatility jumps to a very high or very low level after having been at a neutral level for a prolonged period of time.

3. Empirical methods

3.1 Data

To test our theory we use data on the S&P 500 cash index and VIX index values for the period January 2, 1996 to May 31, 2014. Our key variable of interest is VIX

squared (the one-month variance forecast for the cash index),⁴ or more precisely, changes in VIX squared, i.e., VIX squared first difference. The logic is that contrary to VIX squared, which reflects not only the variance currently perceived by the agents but many other variables such as forecast errors and variance risk premium, changes in VIX squared are mainly driven by contemporaneous changes in the variance perceived by the agents. (We show this formally below). So, if perceived variance is biased following a period of prolonged extreme variance, as per the foregoing after-effect, VIX squared first difference should directly exhibit this bias. The Internet Appendix contains further details on the raw data and various cleaning procedures.⁵

3.2 Estimation of realized volatility

To estimate volatility, we use the Zhang, Mykland, and Aït-Sahalia (2005) multi-grid estimator, which provides a good compromise between accuracy and simplicity. It is more accurate than the Andersen, Bollerslev, Diebold, and Labys (2000) low frequency estimator, which is commonly used in the literature, yet its implementation is relatively simple. Its higher accuracy stems from the fact that it utilizes multiple sampling grids, effectively averaging out much of the measurement error contained in estimates derived from a single grid. Denote the log S&P 500 index value by p . A daily interval $[t - 1, t]$ consists of N tick-by-tick observations $\{t_0, t_1, \dots, t_N\}$. The multi-grid estimator of daily realized variance with K grids results from the summation of squared K -period-returns:

$$RV_{t-1,t}^2 = \frac{1}{K} \sum_{i=0}^{N-K} [p(t_{i+K}) - p(t_i)]^2 . \quad (1)$$

We select K , the sampling frequency of returns, using variance signature plots following Andersen, Bollerslev, Diebold, and Labys (2000).⁶ The optimal K depends on the degree of trading activity, among other factors, which changes substantially through

⁴ VIX, formally the Chicago Board Options Exchange Market Volatility Index, is an estimate of the implied volatility of the S&P 500 index over the next 30 days. As inputs to the calculation, VIX takes the market prices of the all out-of-the-money call and put options for the front and second-to-front expiration months. VIX is computed as the option price implied par variance swap rate for a 30-day variance swap (using a kernel-smoothed estimator), and expressed as an annualized standard deviation (volatility) in percentage points by taking the square root of the variance swap rate.

⁵ The Internet Appendix can be found at <http://ow.ly/He7v0> .

⁶ To determine the optimal sampling frequency we use the volatility signature tool instead of other common techniques (e.g., Zhang, Mykland, and Aït-Sahalia, 2005; Bandi and Russell, 2006). This is because the common techniques assume negative first order autocorrelation of returns, whereas the S&P 500 cash index returns exhibit positive autocorrelation, as we document in the Internet Appendix.

time from the start to the end of our 18-year sample. We therefore use three different sampling frequencies (ten, five and three minutes) in three different time periods, increasing the frequency in line with trading activity. See the Internet Appendix for details.

After computing $RV_{t-1,t}^2$ at the optimal frequency K^* , we calculate realized volatility for the daily interval $[t-1, t]$ by taking the square root of $RV_{t-1,t}^2$ and annualizing using a year of 252 business days:

$$RV_{t-1,t} = \sqrt{RV_{t-1,t}^2 \times 252} . \quad (2)$$

To simplify notation, we refer to realized variance and realized volatility for the daily interval $[t-1, t]$ with a single time subscript corresponding to the end of the daily interval, RV_t^2 and RV_t . In the Internet Appendix, we provide detailed descriptive statistics for realized volatility RV_t , its log $\ln RV_t$, and its first difference, ΔRV_t . We also provide descriptive statistics for the VIX and VIX first difference series. Among other things, we document that the log realized volatility appears to be close to normally distributed, a result discussed by Andersen, Bollerslev Diebold, and Labys (2000, 2001, 2003). We also find that the structure of realized volatility autocorrelation is typical of a long-memory process. Fitting a HAR model to the daily realized variance series we find coefficients that are very close to those found in prior studies (e.g., Corsi, 2009). Similar to realized volatility, the autocorrelation structure of the daily VIX series is typical of a long-memory process. By contrast, VIX and realized volatility first differences are not persistent. Their autocorrelation is not significant beyond the first lag. This result is consistent with earlier studies (Fleming, Ostdiek, and Whaley, 1995; Carr and Wu, 2006; Ahoniemi, 2008). See the Internet Appendix for details.

3.3 Identification of the volatility regimes that induce after-effects

To identify episodes in which after-effects are likely to be triggered, we must first identify very high, very low and neutral volatility states. To do this, we start by computing the mean and standard deviation of the distribution of the daily log realized volatility, during a rolling three-month (63 business days) window. Our motivation for using the log realized volatility (rather than realized volatility) is that it appears to be

approximately normally distributed, as pointed out above. Furthermore, the volatility of the log realized volatility shows little persistence (Corsi, Mittnik, Pigorsch, and Pigorsch, 2008).

We define a very high (*VH*) volatility level as one that is more than x standard deviations above the mean, and a very low (*VL*) volatility level as one that is more than x standard deviations below the mean. A medium or neutral volatility level (*M*) is one that falls within y standard deviations of the mean.⁷

We refer to episodes that, according to theory, are likely to trigger after-effects as volatility ‘regimes’. Our volatility regime indicator variable, $VolReg_t$, is defined over a four-day period. Specifically, $VolReg_t$ takes a value of 1 if we observe very high volatility levels in the three preceding days (‘high volatility state’) and day t has a neutral volatility level. $VolReg_t$ takes the value -1 if we observe very low volatility levels in the three preceding days (‘low volatility state’) and day t has a neutral volatility level. $VolReg_t$ is 0 in all other instances. Formally:

$$VolReg_t = \begin{cases} +1 & \text{if } \{LnRV_{t-3}, LnRV_{t-2}, LnRV_{t-1}, LnRV_t\} = \{VH, VH, VH, M\} \\ -1 & \text{if } \{LnRV_{t-3}, LnRV_{t-2}, LnRV_{t-1}, LnRV_t\} = \{VL, VL, VL, M\} \\ 0 & \text{otherwise} \end{cases} , \quad (3)$$

where $LnRV_t$ is log realized volatility on day t . The identification of volatility states (very high, very low and neutral) and regimes (transitions from very high or very low to neutral volatility) is illustrated in Figure 2.

< Figure 2 here >

Note that volatility regimes ($VolReg_t = \pm 1$) can involve very high or very low realized volatility that persists for more than three days before transitioning to the neutral level. We calibrate x and y to ensure that in the analysis there is a sufficiently large number of regimes (from a statistical viewpoint), while significantly separating the

⁷ We choose to use symmetrical intervals because the distribution is approximately symmetrical (see the Internet Appendix). We use five buckets to avoid threshold effects happening when volatility has been in the highest or lowest bucket and a small change brings it into the adjacent middle bucket. Setting $x = y$ collapses the five buckets into three adjacent ones.

volatility states in the sense that the level of volatility in states VH and VL is sufficiently different from that in the neutral state, M .

After removing the first three months of the sample used in the rolling window that determines high/low/neutral levels, we are left with 4,539 daily observations. When $x = y = 1$, there are about 100 volatility regimes (transitions from very high or very low volatility states to the neutral state) over the whole sample 1996-2014. Figure 3 illustrates the temporal distribution of volatility regimes for $x = y = 1$. Regimes occur regularly throughout the sample with some evidence of clustering, for example, during the second semester of 2009. Table 1 reports the number of regimes for a range of x and y between 1.00 and 1.75 standard deviations.

< Figure 3 here >

< Table 1 here >

In the main analysis we use volatility states that are relatively close to each other because when $x - y$ is very large (i.e., when the difference between the volatility levels in the very high/low states and the neutral state is large—the ideal scenario to detect the after-effect if any), the number of regimes is too low from a statistical viewpoint.⁸

Table 2 reports the absolute difference between the average log realized volatility at the onset of a neutral state and the average log realized volatility over the previous three days, for different values of x and y . The difference increases with the strength of the stimulus (x) and it decreases with the threshold that defines a neutral volatility level (y). The average jump from either a very high or very low volatility state to a neutral volatility state is 0.42 in log terms (about 40% in realized volatility terms).

< Table 2 here >

⁸ This is not surprising: our regime indicator is defined over four days and we have 4,539 days in the sample. So we have a maximum of 1,135 non-zero values. To get a non-zero value, realized volatility has to stay in the tail of the distribution for three days in a row before jumping. This is an unlikely path (albeit it is possible given the persistent nature of realized volatility and the presence of jumps).

3.4 Structural model

The after-effect theory predicts that the VIX, a model-free measure of volatility expectations is impacted by a systematic perception error in the aftermath of a volatility regime. Changes in the VIX are mainly driven by changes in expected future volatility, which in turn are driven by the volatility perceived by the agent. Normally, those changes in perceived volatility mirror corresponding changes in the level of realized volatility except in the aftermath of prolonged exposure to very high or very low volatility. To see this formally, we write a structural model that assumes rational expectations but allows for a perceptual bias due to after-effects. In the Internet Appendix we investigate an alternative assumption of adaptive expectations and show that assuming adaptive expectations would merely strengthen our evidence for the existence of an after-effect in volatility perception.

VIX squared is the price of a synthetic variance swap⁹, which is the sum of expected realized variance and a variance risk premium as in Carr and Wu (2006):

$$VIX_t^2 = E_t[RV_{t,t+30}^2] + VRP_t, \quad (4)$$

By differencing, we obtain:

$$\begin{aligned} \Delta VIX_t^2 &= VIX_t^2 - VIX_{t-1}^2 \\ &= E_t[RV_{t,t+30}^2] - E_{t-1}[RV_{t-1,t+29}^2] + \Delta VRP_t. \end{aligned} \quad (5)$$

In forming expectations about future variance, a rational agent is likely to use (either explicitly or implicitly) a variance forecasting model that has the highest possible forecasting accuracy. A model that fits that criterion is the popular HAR model used in Andersen, Bollerslev and Diebold (2007).¹¹ While being quite simple, this model together with related models such as ARFIMA has excellent forecasting performance, beating competing models such as the ARCH family of models (e.g., Corsi, 2009). One-week and one-month realized variances are given by the following:

$$RV_{t,t+7}^2 = \frac{1}{7}(RV_{t,t+1}^2 + RV_{t+1,t+2}^2 + \dots + RV_{t+6,t+7}^2). \quad (6)$$

⁹ VIX is quoted as an annualized standard deviation, so VIX squared is annualized risk-neutral expected variance.

¹¹ Specifically, we use equation 11 of Andersen, Bollerslev and Diebold (2007). The reader will note that here we use variance instead of volatility and calendar instead of business days. Despite their simplicity, HAR models capture the two most important empirical characteristics of volatility: volatility clustering and long memory (e.g., Corsi, 2009).

$$RV_{t,t+30}^2 = \frac{1}{30} (RV_{t,t+1}^2 + RV_{t+1,t+2}^2 + \cdots + RV_{t+29,t+30}^2). \quad (7)$$

The HAR model for monthly variance is:

$$RV_{t,t+30}^2 = \beta_0 + \beta_D RV_{t-1,t}^2 + \frac{\beta_W}{7} RV_{t-7,t}^2 + \frac{\beta_M}{30} RV_{t-30,t}^2 + \epsilon_{t,t+30}. \quad (8)$$

This gives (rewriting the difference of expectations in (5)):

$$\begin{aligned} E_t^P [RV_{t,t+30}^2] - E_{t-1}^P [RV_{t-1,t+29}^2] &= \beta_D RV_{t-1,t}^2 + \frac{\beta_W}{7} RV_{t-7,t}^2 + \frac{\beta_M}{30} RV_{t-30,t}^2 \\ &\quad - (\beta_D RV_{t-2,t-1}^2 + \frac{\beta_W}{7} RV_{t-8,t-1}^2 + \frac{\beta_M}{30} RV_{t-31,t-1}^2) \\ &= \beta_D (RV_{t-1,t}^2 - RV_{t-2,t-1}^2) \\ &\quad + \frac{\beta_W}{7} (RV_{t-1,t}^2 - RV_{t-8,t-7}^2) \\ &\quad + \frac{\beta_M}{30} (RV_{t-1,t}^2 - RV_{t-31,t-30}^2). \end{aligned} \quad (9)$$

From the estimated HAR coefficients for S&P 500 realized variance (Table IIB of Andersen, Bollerslev and Diebold (2007) and our own estimated coefficients in Table 8 of the Internet Appendix), one can see that the coefficients β_W and β_M are both quite small (around 0.3), which means that the difference in expectation is mainly driven by the first term, the change in realized variance over $[t-1, t]$:

$$E_t [RV_{t,t+30}^2] - E_{t-1} [RV_{t-1,t+29}^2] \approx \beta_D (RV_{t-1,t}^2 - RV_{t-2,t-1}^2) = \beta_D \Delta RV_t^2. \quad (10)$$

Substituting (10) into (5), we thus have:

$$\Delta VIX_t^2 \approx \beta_D \Delta RV_t^2 + \Delta VRP_t. \quad (11)$$

According to the after-effect theory described above, after a very high or very low volatility regime, perceived variance will differ from actual by a perception error. Using subscript π for perceived variance and PE for the perception error, we have:

$$RV_{\pi,t}^2 = \begin{cases} RV_t^2 - PE, & \text{after a high volatility regime} \\ RV_t^2 + PE, & \text{after a low volatility regime} \\ RV_t^2, & \text{otherwise} \end{cases} \quad (12)$$

More formally, using the definitions of volatility regimes and $VolReg_t$ variable, we have:

$$\{RV_{\pi,t-3}^2, RV_{\pi,t-2}^2, RV_{\pi,t-1}^2, RV_{\pi,t}^2\} = \begin{cases} \{RV_{t-3}^2, RV_{t-2}^2, RV_{t-1}^2, RV_t^2 - PE\} & \text{if } VolReg_t = +1 \\ \{RV_{t-3}^2, RV_{t-2}^2, RV_{t-1}^2, RV_t^2 + PE\} & \text{if } VolReg_t = -1 \\ \{RV_{t-3}^2, RV_{t-2}^2, RV_{t-1}^2, RV_t^2\} & \text{if } VolReg_t = 0 \end{cases}. \quad (13)$$

This leads to:

$$RV_{\pi,t}^2 = RV_t^2 - PE.VolReg_t, \quad (14)$$

and

$$\Delta RV_{\pi,t}^2 = \Delta RV_t^2 - PE.VolReg_t. \quad (15)$$

We therefore incorporate the after-effect bias into equation (11) by rewriting it using equation (15):

$$\begin{aligned} \Delta VIX_t^2 &\approx \beta_D \Delta RV_{\pi,t}^2 + \Delta VRP_t \\ &\approx \beta_D \Delta RV_t^2 - \beta_D PE.VolReg_t + \Delta VRP_t. \end{aligned} \quad (16)$$

Our strategy is thus to regress changes in VIX squared onto changes in realized variance and the $VolReg_t$ variable. Following Merton's (1980) arguments, the variance risk premium is slow-moving; Bollerslev, Gibson, and Zhou (2011) find that it varies with the business cycle. For this reason, the third term in (16), ΔVRP_t , should be negligible. One may deem the Mertonian assumption too strong though, in which case it is advisable to augment the regression with control variables for ΔVRP_t . We test both specifications of the model and the results appear to be robust to the inclusion of a variety of control variables. We also conduct 'placebo' tests (described in Section 4.4), which address the concern that large jumps in volatility (such as those captured by $VolReg_t$) may impact VIX via changes in the variance-risk premium.

Equation (16) makes it clear that changes in VIX_t^2 mainly depend on the changes of perceived variance over the period $[t-1, t]$ and not on the higher order lags. That the VIX_t^2 first difference is mainly driven by contemporaneous changes in the variance level perceived by the agents justifies that we use it in the analysis instead of using the level of VIX. Using the latter would not work given our purpose here; if we were to use it, monthly variations in important VIX factors such as the degree of risk aversion would potentially cloud the perceptual after-effect that we are looking for.

It should also be noted that the VIX first difference time series features low persistence (as documented in prior empirical work—and we verify that this is true in our data as well). As such it is a suitable dependent variable in the analysis contrary to VIX, which is highly persistent.¹³

¹³ Granger and Joyeux (1980) show that one cannot infer much from a regression in which the dependent variable is highly persistent.

3.5 Regression strategy

We use a log specification of (16) as the logs of VIX and logs of realized volatility are approximately normally distributed, as pointed out earlier.¹⁴ In robustness tests we find similar results, in some cases even stronger, under alternative specifications including not logged series. The benchmark form of our regression is thus:¹⁵

$$\Delta \ln VIX_t = \alpha + \beta VolReg_t + \gamma \Delta \ln RV_t + \varepsilon_t. \quad (17)$$

We estimate this benchmark regression and then augment it by adding control variables one at a time. Among the control variables, we include the market return during the transition period (r_t) as a proxy for changes in the variance risk premium (ΔVRP_t in equation (16)). We also include negative market returns ($r_t^- = \min(r_t, 0)$) to account for possible leverage effects. We include first lags of realized volatility and VIX first differences since the two series display significant auto-correlation (as described earlier). Finally, we include dummy variables for the well-known day-of-the-week effect in VIX (Fleming, Ostdiek, and Whaley, 1995). The regression with the complete set of control variables is thus:

$$\begin{aligned} \Delta \ln VIX_t = & \alpha + \beta VolReg_t + \gamma_0 \Delta \ln RV_t + \gamma_1 \Delta \ln RV_{t-1} + \delta r_t + \delta^- r_t^- \\ & + \rho_1 \Delta \ln VIX_{t-1} + \sum_{i=2}^5 \theta_i D_{it} + \varepsilon_t, \end{aligned} \quad (18)$$

where $\{D_{it}\}_{i=2,3,4,5}$ are dummy variables for Tuesday to Friday (Monday is base case).

4. Results

4.1 Impact of a regime change on VIX

We find that the impact of a change of regime on VIX is significant, as predicted by after-effects theory. Table 3 column (1) reports estimates from the baseline regression (17) using threshold parameters $x = 1.75$ and $y = 1.50$ (a compromise that ensures a sufficiently large number of regimes and sufficiently large difference between the volatility states). The impact of $VolReg_t$ on changes in VIX has a sign that is consistent with volatility after-effects, and is statistically significant. The economic impact of

¹⁴ We multiply the log difference by 100 to make it consistent with the definition of S&P 500 returns.

¹⁵ In logging the series, the squares become linear terms, with the factor of two being absorbed into the corresponding coefficients and regression intercept.

$VolReg_t$ is large: a transition from a very high or very low volatility state to neutral volatility changes VIX by about 2.73%.

We augment the baseline regression with a number of control variables. The coefficient of the key variable, $VolReg_t$, is hardly affected by the additional control variables. Column (2) reports the results of the regression with S&P 500 returns and negative returns added as control variables that correlate with changes in the variance risk premium. Both return variables are highly significant both statistically and economically: a +1% S&P 500 return is associated with a 3.31% drop in VIX. There is asymmetry in the impact of returns confirming results found in previous literature: a 1% increase in S&P 500 decreases VIX by 2.41% while a 1% decrease increases VIX by 4.20%. Including the S&P 500 returns significantly increases the R^2 of the regression: realized volatility differences and S&P 500 returns explain close to 60% of the variation in VIX first differences. Most importantly, the coefficient of our main variable of interest, $VolReg_t$, hardly changes compared to its estimate in the baseline regression.

In column (3) the regression includes lagged VIX difference and lagged realized volatility differences. The coefficient of lagged VIX differences is significant and large. Finally, column (4) reports the results of the regression that includes day-of-the-week dummies. All of the day-of-the-week dummies are significant, which indicates daily seasonality in VIX, consistent with existing literature. The coefficient of $VolReg_t$ is remarkably stable across the four regressions. It always has the sign that is consistent with predictions based on after-effects theory and is statistically significant.

< Table 3 here >

4.2 Impact of a regime change on VIX as a function of regime strength

We now turn to examining some of the more nuanced predictions of after-effects theory. We find that the more extreme the volatility levels during the adaptation phase (the three days preceding a transition to neutral volatility), the more significant our volatility regime variable, suggesting a stronger after-effect. Table 4 reports the estimated coefficients and significance of $VolReg_t$ in regression (18) for different values of x (which determines the volatility levels during the stimulus phase) and y (which

determines the volatility levels in the neutral state). In all cases the estimated coefficient of $VolReg_t$ is negative and statistically significant, consistent with the presence of volatility after-effects.¹⁶ The largest coefficient is 3.572, which implies an effect size that is of the same order of magnitude as the impact of S&P 500 returns in the regression. Put differently, depending on the strength of the volatility stimulus, the impact of the after-effect on VIX can be about the same as the impact of a 1% change in S&P 500.

< Table 4 here >

Figure 4 displays the coefficient of $VolReg_t$ as a function of the level of (log) realized volatility in the adaptation phase along the diagonal of Table 4 (i.e., when $x = y$). We see a steady increase in the size of the after-effect as the average level of realized volatility in the adaptation phase increases. This finding is consistent with our theoretical prediction that the more extreme the volatility during the adaptation period, the stronger the neuronal adaptation and hence the larger the after-effect.

< Figure 4 here >

4.3 Impact of regime change on VIX as a function of stimulus duration

After-effects theory also predicts that the longer the exposure to the stimulus, the stronger the after-effect. Our tests in this subsection support this prediction: the longer the stimulus phase, the more significant the perception bias. To establish this we modify $VolReg_t$ so that the adaptation period spans two days, three days or five days.¹⁷

The after-effect increases with the number of days in the adaptation window. For the threshold values $x = 1.50$ and $y = 1.25$ for instance, the estimated coefficients on the $VolReg_t$ variable are: 1.723, 2.439 and 3.592 for two-day, three-day and five-day adaptation windows. This finding is consistent with the prediction that the longer the

¹⁶ The coefficients of the control variables are virtually unchanged for the different values of x and y .

¹⁷ In the Internet Appendix, we document that there are fewer ‘transitions’ using a three-day adaptation window than a two-day window and even fewer when using a five-day window (as one would expect). The absolute differences between the volatility level in the very high / very low state compared to the neutral state are almost identical when using the two-day and three-day adaptation windows. The absolute differences are slightly larger with the five-day windows.

adaptation period (the time spent in a very high or very low volatility state) the stronger the neuronal adaptation and hence the larger the after-effect in the neutral state.

< Figure 5 here >

4.4 Transition from neutral to very high or very low volatility states

While our results so far are consistent with the predictions of perceptual after-effects, a competing explanation is related to the fact that the $VolReg_t$ variable measures large jumps in realized volatility. It could be that agents have adaptive expectations about volatility *changes*: after seeing an increase (resp. decrease) in realized volatility, they expect a further increase (resp. decrease). According to that theory, immediately after transitioning from a high (resp. low) volatility state to a neutral state, the agent expects volatility to further decrease (resp. increase). Consequently, the agent revises his expectation of 30-day future volatility downward (resp. upward), causing a negative (resp. positive) change in VIX. The negative (resp. positive) changes in VIX coincide with $VolReg_t = +1$ (resp. $VolReg_t = -1$) and therefore if agents have adaptive expectations of volatility *changes* we would expect $\beta < 0$ in our main regression, which is consistent with our results.

To tease apart the after-effect and adaptive expectations theories, we construct a ‘placebo’ test in which we modify our volatility regime indicator variable so that similar to the original definition it measures jumps between adjacent volatility states after a period of stability in volatility levels, but unlike the original definition the jumps are not predicted to cause perceptual after-effects. According to the after-effect theory, there should be no after-effect when realized volatility jumps from a neutral state to a very high or very low volatility state. In contrast, the adaptive expectations theory predicts a bias in VIX when realized volatility jumps from a neutral state to a very high or very low volatility state.

To perform our placebo test, we modify our volatility regime indicator variable as follows:

$$ModVolReg_t = \begin{cases} +1 & \text{if } \{LnRV_{t-3}, LnRV_{t-2}, LnRV_{t-1}, LnRV_t\} = \{M, M, M, VH\} \\ -1 & \text{if } \{LnRV_{t-3}, LnRV_{t-2}, LnRV_{t-1}, LnRV_t\} = \{M, M, M, VL\} \\ 0 & \text{otherwise} \end{cases} \quad (19)$$

Table 5 shows the number of transitions from the neutral volatility state to the very high or very low state. There are a lot more transitions than before. This is because it is more likely that volatility will stay in a neutral state three days in a row and jump to a very high or very low state than it is that it will stay in a very high or very low state three days in a row and jump to a neutral state. Table 6 displays the absolute difference between the average log realized volatility in the neutral state and the average log realized volatility in the very high or very low volatility states. The values are similar to those found in Table 2.

< Table 5 here >

< Table 6 here >

Table 7 reports estimated coefficients of the modified volatility regime variable in regression (18) (replacing $VolReg_t$ with $ModVolReg_t$ in the regression). As the after-effects theory predicts, there is no effect for transitions from a neutral state to a very high or very low volatility state: the coefficients are small and none of them are statistically significant. This finding suggests that it is not large changes in volatility per se that drive changes or bias in VIX, but only those changes that take a specific form, namely a transition from prolonged very high/low volatility to neutral volatility. Such evidence rules out the aforementioned adaptive expectations explanation for the underlying mechanism that drives our results.

< Table 7 here >

4.5 No transition in volatility

The neuronal adaptation mechanism described in Section 2 suggests that perceived volatility (and thus the level of VIX) may slightly decrease (increase) toward the end of a prolonged high (low) volatility state as agents get used to the extreme volatility level. Practitioners often refer to this effect as ‘reset of the mean’. If present, the phenomenon may potentially cloud our identification of the after-effects that are the focus of this study. We expect reset-of-the-mean, if present, to be a second-order phenomenon

compared to the strength of after-effects (which is why we do not account for it in the outline of the theory illustrated in Figure 1).

To test for the reset-of-the-mean effect, we define a new variable ($NonVolReg_t$) that, in contrast to $VolReg_t$ takes non-zero values when realized volatility stays in the same very high or very low state *without* transitioning to the neutral state. We vary the duration of the stimulus window from two to four days. For a four-day window, $NonVolReg_t$ is defined as:

$$NonVolReg_t = \begin{cases} +1 & \text{if } \{LnRV_{t-3}, LnRV_{t-2}, LnRV_{t-1}, LnRV_t\} = \{VH, VH, VH, VH\} \\ -1 & \text{if } \{LnRV_{t-3}, LnRV_{t-2}, LnRV_{t-1}, LnRV_t\} = \{VL, VL, VL, VL\} \\ 0 & \text{otherwise} \end{cases} . \quad (20)$$

Table 8 reports the coefficients of $NonVolReg_t$ when it replaces $VolReg_t$ in regression (18). The reset-of-the-mean effect predicts that the coefficient of $NonVolReg_t$ should be negative, i.e., prolonged exposure to very high volatility should cause perceived volatility to be lower than actual, and vice versa for very low volatility. The table reports the coefficient estimates of $NonVolReg_t$ for different lengths of the adaptation window and different values of the threshold that defines the very high/low volatility states, x . The coefficient estimates are generally small negative values, which could indicate some adaptation to extreme volatility states, but they are neither statistically nor economically significant. This supports our modeling choice to neglect the reset-of-the-mean perceptual bias during extreme volatility states.

< Table 8 here >

4.6 Robustness checks

The foregoing findings constitute strong evidence that investor perceptions of volatility are biased by after-effects. We run many robustness checks which we report in detail in the Internet Appendix A.5.

First, we add additional control variables to the regressions. The key coefficient estimates barely change. Second, we use alternative specifications for the regression. For example, we use VIX and realized volatility in levels rather than in logs. Our main results are robust to this alternative specification.

Third, we check that use of a three-month window to calculate the volatility of realized volatility is not pivotal for our main results. We repeat the analysis using a six-month window and find that even though the choice of the window has an impact on the volatility regime identification, the regression results are robust to the use of a different window. We further find that our results still hold when classifying volatility states (very high, very low, neutral) using percentiles of the realized volatility distribution, rather than fractions of standard deviations of the realized volatility distribution.

Finally, we find that our results still hold when assuming that agents have adaptive expectations about volatility levels (rather than assuming rational expectations as in Section 3.4). In fact, the results are strengthened, i.e., the after-effect is actually stronger than the one reported in Section 4. Intuitively, this is because adaptive expectations and the after-effect phenomenon work as antagonistic forces, inasmuch as adaptive expectations about volatility levels push the VIX in the opposite direction to the after-effect. See Internet Appendix A.5.5 for a formal proof.

5. Follow-up investigations

5.1 Asymmetry of the after-effect in the field

We investigate whether the after-effect is as strong when transitioning from a very high volatility state to a neutral state (henceforth, ‘post high’) as it is when transitioning from a very low volatility state to a neutral state (‘post low’). To that goal, we decompose our volatility regime variable $VolReg_t$ into the following two variables:

$$VolReg_t^+ = \begin{cases} +1 & \text{if } \{LnRV_{t-3}, LnRV_{t-2}, LnRV_{t-1}, LnRV_t\} = \{VH, VH, VH, M\} \\ 0 & \text{otherwise} \end{cases}, \quad (21)$$

$$VolReg_t^- = \begin{cases} -1 & \text{if } \{LnRV_{t-3}, LnRV_{t-2}, LnRV_{t-1}, LnRV_t\} = \{VL, VL, VL, M\} \\ 0 & \text{otherwise} \end{cases}. \quad (22)$$

We separately compute the number of transitions from very high volatility states to neutral and from very low volatility states to neutral for a range of thresholds, x and y . Table 9 summarizes the number of transitions. The number of transitions is fairly symmetrical. Table 10 reports the difference between the average volatility in the very high/low states and the volatility in the neutral state. The difference is around 10% smaller for transitions from the very low volatility state to the neutral state than it is for transitions from the very high volatility state to neutral.

< Table 9 here >

< Table 10 here >

Table 11 reports the estimated coefficients of $VolReg_t^+$ and $VolReg_t^-$ when they are used as replacements for $VolReg_t$ in regression (18). The after-effect seems to occur only when transitioning from prolonged exposure to very high volatility, not when transitioning from very low volatility. The coefficients of $VolReg_t^+$ are even larger than the coefficients of $VolReg_t$ in Table 4. In contrast, the coefficient of $VolReg_t^-$ is not statistically distinguishable from zero. Thus, the after-effect appears to be asymmetric.

< Table 11 here >

To understand this asymmetry, it is useful to look at the average levels of realized volatility in the very high/low states (see Table 12). Whereas the average volatility level in the very high states is between the 90th and 95th percentile of the volatility distribution, the average volatility level in the very low states is never below the 25th percentile.¹⁸ So it seems that in the field, prolonged periods of truly low volatility are absent, which explains the asymmetry found in our tests of after-effects. We further explore this idea in the next subsection.

< Table 12 here >

5.2 Asymmetry of the after-effect in the laboratory

We revisit the data of Payzan-LeNestour, Balleine, Berrada, and Pearson (2014) to investigate whether the foregoing asymmetry in the after-effect post high vs. post low prevails in the laboratory as well. The data reveal that the after-effect post high is stronger than the after-effect post low, which suggests some potential asymmetry in the after-effect, albeit not as large as the one we see in the field (see Table 13, “Original

¹⁸ Key percentiles of the realized volatility distribution are: 5th = 5.87%, 10th = 6.64%, 25th = 8.32%, 75th = 15.99%, 90th = 22.53%, and 95th = 28.20%.

Laboratory Experiment”). We attribute this difference between the two studies to the fact that in the laboratory, the low volatility parameter was set to a very low value (2%) relative to the values typically observed in the field during the low volatility states (7-9%).

To test this conjecture, we run follow-up experimental sessions that replicate the original laboratory experiment except that we set the volatility parameters to values that are in line with the values that we observe in the field. The experimental design is described in detail in Payzan-LeNestour, Balleine, Berrada, and Pearson (2014). We repeat the essentials here for ease of reference. In each trial, the subjects (N=31) are shown for 50 seconds a time-series representing the trajectory of a stock market index over a year at a daily frequency. Then, during a 15-second test phase, they are asked how volatile they perceive a second (unrelated) time-series to be, on a scale of 1-5 (1: very flat; 5: very volatile). By design, the trajectory in the test phase always has a volatility level of 13.5% (the average volatility level observed during the neutral state in our field study). The volatility levels in the first phase of the experimental trials alternate between 40% (to mimic the very high volatility state in the field) and 7% (to mimic the very low volatility state). Control and diversion trials are also randomly interspersed, to get a benchmark on subjects’ reports in the absence of any after-effect. In the control trials, the subjects see the neutral stimulus in both phases. In the diversion trials, the subjects see the neutral stimulus in the first phase and the very high / very low volatility stimulus (volatility levels alternate between 40% and 7%) in the second phase.

For generalization purposes, we also run a variant of the task in which instead of assessing the *volatility* of a time-series, subjects (N=57) are asked to assess the *variance* of a sequence of balls that are sequentially drawn from a bucket. The balls are drawn from normal distributions with varying means and standard deviations. Subjects are shown a test sequence (standard deviation: 13.5%) for less than 20 seconds after being exposed for 50 seconds to either a low variance stimulus (standard deviation: 5%) or a high variance stimulus (standard deviation: 40%). They are asked to report how unstable

they perceive the test sequence to be on a scale of 1-5 (1: very stable; 5: very unstable).¹⁹ The results obtained in both variants of the task are very similar so we merge them here for simplicity, but the results hold in each subsample.

< Table 13 here >

Table 13 summarizes the main results. The after-effect appears to be asymmetric in the follow-up experiment. Like in the field, the after-effect post high is very strong, whereas the after-effect post low is essentially absent. In the original experimental sessions, the difference between the mean perceived variability post high (post low) and the mean perceived variability in the control trials is -0.44 (0.34). So the magnitude of the after-effect post low is smaller than the corresponding magnitude post high, but the after-effect post low is still significant (p-value: 0.000). In contrast, in the follow-up experimental sessions, the after-effect post low vanishes (p-value: 0.650) whereas the after-effect post high is very significant. The only difference between the original and follow-up experiments is that in the latter, the variability parameters are set to match the values observed in the field.

These results are consistent with our conjecture that the absence of after-effect after exposure to low volatility in the field reflects the lack of truly low volatility regimes in the stock market, rather than an asymmetry of the after-effect.

6. Conclusions

We examine whether the after-effect phenomenon, which has been documented in a large number of settings outside of economics and finance, affects investor perceptions of asset return volatility and consequently impacts on asset prices. Using VIX for S&P 500 stocks, we provide strong evidence that investors perceive volatility to be lower than actual volatility after prolonged exposure to very high volatility levels. The magnitude of this perception bias is highly economically meaningful as noted earlier.

¹⁹ In the volatility version of the experiment (resp. variance version), subjects are explained intuitively by means of exemplar stimuli what volatile (resp. unstable) means. Demonstrations of the task and the task instructions are available at elisepayzan.com\na.

Our empirical analysis further finds support for a series of more nuanced facts predicted by our after-effect theory. For example, the perception bias becomes stronger, the longer the exposure to the very high / very low volatility ‘stimulus’. The perception bias is also stronger when the stimulus is more intense, i.e., following more extreme levels of volatility. Furthermore, consistent with after-effects being the driver of our results, we find no perception bias following jumps in volatility that are not preceded by prolonged exposure to very high / very low volatility levels—even though these jumps are comparable in magnitude to those that induce after-effects. These additional results rule out the most plausible alternative explanations. To our knowledge, no competing theory can explain the collection of empirical facts.

The psychology, physiology and behavioral economics/finance literatures document a number of different cognitive limitations and biases that affect individuals’ perceptions and behavior in ways that depart from traditional financial economics assumptions about rationality. Although there is little doubt that a large number of such biases exist in individuals, there is an active and unresolved debate in the literature about which of them, if any, affect aggregate market outcomes such as equilibrium asset prices, and in what settings. Much of the debate revolves around whether highly capitalized sophisticated and relatively unbiased arbitrageurs/speculators steer markets to outcomes consistent with rational behavior, or whether limits to arbitrage, frictions or the sheer mass of biased individuals cause biases to impact equilibrium prices. What is remarkable about the current study is that we do not simply document that a perception bias exists among some individuals, but rather, we show that the bias has a meaningful impact on asset prices. What is more, we show this not in a small and illiquid market where frictions and limits to arbitrage may be large, but in one of the most actively traded markets in the world.

Our evidence that asset prices in even a very actively traded market can be substantially impacted by a neurologically based perception bias naturally raises a number of further questions. First, how pervasive are after-effects in financial markets – do they play a role in different contexts such as perceptions of trading activity or liquidity, do they affect the valuations of companies in addition to stock option prices, are their effects stronger in smaller and less active markets? Does the influence of after-

effects diminish with increased use of computer algorithms to make trading decisions? Are there other deep-rooted neurological processes that systematically bias investors' perceptions and influence asset prices? These are all important directions for future research, which will help reconcile the behavioral and classical paradigms.

References

- Ahoniemi, K., 2008, Modeling and forecasting the VIX index, Working paper.
- Andersen, T., T. Bollerslev, and F.X. Diebold, 2007, Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility, *Review of Economics and Statistics* 89, 701-720.
- Andersen, T., T. Bollerslev, F.X. Diebold, and P. Labys, 2000, Great realisations, *Risk* (March 2000), 105-108.
- Andersen, T., T. Bollerslev, F.X. Diebold, and P. Labys, 2001, The distribution of realized exchange rate volatility, *Journal of the American Statistical Association* 96, 42-55.
- Andersen, T., T. Bollerslev, F.X. Diebold, and P. Labys, 2003, Modeling and forecasting realized volatility, *Econometrica* 71, 529-626.
- Bandi, F., and J. Russell, 2006, Separating microstructure noise from volatility, *Journal of Financial Economics* 79, 655-692.
- Barberis, N., and A. Shleifer, 2002, Style investing, *Journal of Financial Economics* 68, 161-199.
- Barberis, N., R. Greenwood, L. Jin, and A. Shleifer, 2015, X-CAPM: An extrapolative capital asset pricing model, *Journal of Financial Economics* 115, 1-24.
- Barlow, H.B., and R.M. Hill, 1963, Evidence for a physiological explanation of the waterfall phenomenon and figural after-effects, *Nature* 200, 1345-1347.
- Barlow, H.B., 1972, Dark and light adaptation: Psychophysics, in *Handbook of Sensory Physiology*, Jameson, D., and L. M. Hurvich, Editors, Springer-Verlag: New York, 1-28.
- Bollerslev, T., M. Gibson, and H. Zhou, 2011, Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realised volatilities, *Journal of Econometrics* 160, 235-245.
- Burr, D., and J. Ross, 2008, A visual sense of number, *Current Biology* 18, 425-428.
- Carr, P., and L. Wu, 2006, A tale of two Indices, *Journal of Derivatives* 13, 13-29.
- Choi, J.J., and T.M. Mertens, 2013, Extrapolative expectations and the equity premium, Working paper.

- Corsi, F., S. Mittnik, C. Pigorsch, and U. Pigorsch, 2008, The volatility of realized volatility, *Econometric Reviews* 27, 46-78.
- Corsi, F., 2009, A simple approximate long-memory model of realized volatility, *Journal of Financial Econometrics* 7, 174-196.
- Delahunt, P.B., M.A. Webster, L. Ma, and J.S. Werner, 2004, Long-term renormalisation of chromatic mechanisms following cataract surgery, *Visual Neuroscience* 21, 301-307.
- DellaVigna, S., and J.M. Pollet, 2009, Investor inattention and Friday earnings announcements, *Journal of Finance* 64, 709-749.
- De Long, J.B., A. Shleifer, L.H. Summers, and R.J. Waldmann, 1990, Noise trader risk in financial markets, *Journal of Political Economy* 98, 703-738.
- Eysenck, M. and M.T. Keane, 2013, *Cognitive Psychology* (6th ed.), Psychology Press: Hove/New York.
- Fleming, J., B. Ostdiek, and R.E. Whaley, 1995, Predicting stock market volatility: A new measure, *Journal of Futures Markets* 15, 265-302.
- Froot, K.A., and E.M. Dabora, 1999, How are stock prices affected by the location of trade?, *Journal of Financial Economics* 53, 189-216.
- Gabaix, X., 2014, A sparsity-based model of bounded rationality, *Quarterly Journal of Economics* 129, 1661-1710.
- Glimcher, P.W., 2011, *Foundations of Neuroeconomic Analysis*, Oxford University Press: Oxford.
- Granger, C.W.J., and R. Joyeux, 1980, An introduction to long memory time series models and fractional differencing, *Journal of Time Series Analysis* 1, 15-39.
- Greenwood, R., and A. Shleifer, 2014, Expectations of returns and expected returns, *Review of Financial Studies* (forthcoming).
- Griggs, R.A., 2009, *Psychology: A Concise Introduction* (2nd ed.), Worth Publishers: New York.
- Hering, E., 1964, *Outlines of a Theory of the Light Sense*, Harvard University Press: Cambridge, Massachusetts.

- Hershenson, M.B., 1989, Duration, time constant, and decay of the linear motion aftereffect as a function of inspection duration, *Perception and Psychophysics* 45, 251–257.
- Huberman, G., and T. Regev, 2001, Contagious speculation and a cure for cancer: A nonevent that made stock prices soar, *Journal of Finance* 56, 387-396.
- Hurvich, L.M., and D. Jameson, 1957, An opponent-process theory of color vision, *Psychological Review* 64, 384-404.
- Kahneman, D., 1973, *Attention and Effort*, Prentice-Hall: Englewood Cliffs, New Jersey.
- Kahneman, D., and A. Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263-291.
- Kandel, E.R., J.H. Schwartz, and T.M. Jessell, 2000, *Principles of Neural Science* (4th ed.), McGraw–Hill: New York.
- Koszegi, B., and A. Szeidl, 2013, A model of focusing in economic choice, *Quarterly Journal of Economics* 128, 53–104.
- Lamont, O.A., and R.H. Thaler, 2003, Can the market add and subtract? Mispricing in tech stock carve-outs, *Journal of Political Economy* 111, 227–268.
- Lee, C.M.C., A. Shleifer, and R.H. Thaler, 1991, Investor sentiment and the closed-end fund puzzle, *Journal of Finance* 46, 75–109.
- Leopold, D.A., G. Rhodes, K.-M. Müller, and L. Jeffery, 2005, The dynamics of visual adaptation to faces, *Royal Society Proceedings B* 272, 897–904.
- McFadden, D., 1999, Rationality for economists?, *Journal of Risk and Uncertainty* 19, 73-105.
- Magnussen, S., and T. Johnsen, 1986, Temporal aspects of spatial adaptation. A study of the tilt aftereffect, *Vision Research* 26, 661–672.
- Merton, R.C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323-361.
- Mitchell, M., T. Pulvino, and E. Stafford, 2002, Limited arbitrage in equity markets, *Journal of Finance* 57, 551–584.
- Payzan-LeNestour, E., B. Balleine, T. Berrada, and J. Pearson, 2014, Variance after-effects bias risk perception in humans, Working paper.

- Rutherford, M.D., H.M. Chattha, and K.M. Krysko, 2008, The use of aftereffects in the study of relationships among emotion categories, *Journal of Experimental Psychology* 34, 27-40.
- Shleifer, A., and R.W. Vishny, 1997, The limits of arbitrage, *Journal of Finance* 52, 35-55.
- Simon, H.A., 1955, Behavioral model of rational choice, *Quarterly Journal of Economics* 49, 99–118.
- Woodford, M., 2012, Inattentive valuation and reference-dependent choice, Working paper.
- Webster, M.A., D. Kaping, Y. Mizokami, and P. Duhamel, 2004, Adaptation to natural facial categories, *Nature* 428, 558-561.
- Webster, M.A., K. McDermott, and G. Bebis, 2007, Fitting the world to the mind: Transforming images to mimic perceptual adaptation, 3rd International Symposium on Visual Computing (ISVC07), *Lecture Notes of Computer Science* 4842, 757-768.
- Zhang, L., P.A. Mykland, and Y. Ait-Sahalia, 2005, A tale of two time scales: Determining integrated volatility with noisy high-frequency data, *Journal of the American Statistical Association* 100, 1394-1411.




Figure 1. How after-effects bias investor perceptions of realized variance.

This figure illustrates the perception bias that is predicted by the after-effects theory. After prolonged exposure to high (low) realized variance, perceived variance is lower (higher) than actual realized variance.

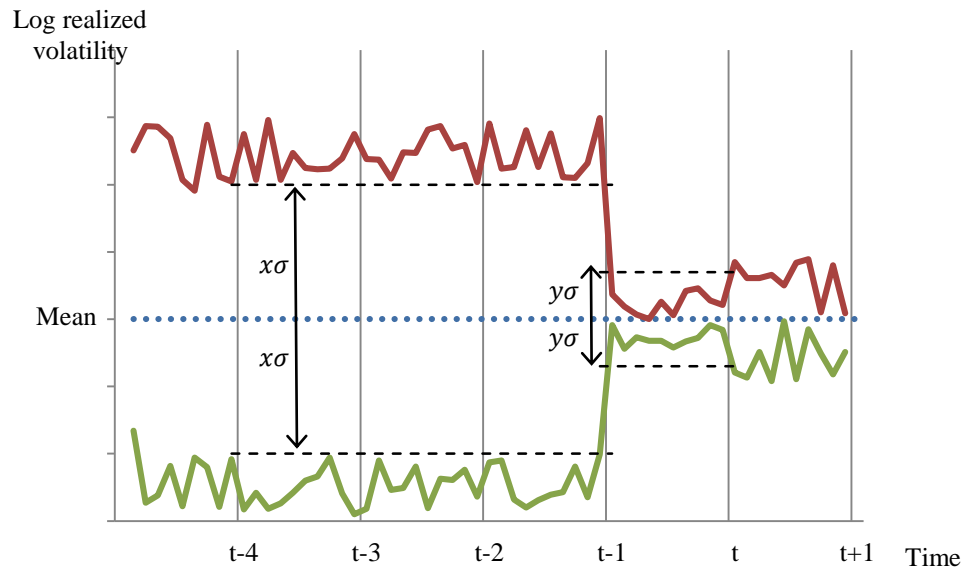


Figure 2. Methodology to identify the volatility regimes.

This figure illustrates how volatility regimes are defined. A very-high-to-neutral transition ($VolReg_t = +1$) occurs when realized volatility is very high (greater than x standard deviations above the mean) for at least three consecutive days and then neutral (within y standard deviations from the mean) the next day. Similarly, a very-low-to-neutral transition ($VolReg_t = -1$) occurs when realized volatility is very low (more than x standard deviations below the mean) for at least three consecutive days and then neutral (within y standard deviations from the mean) the next day.

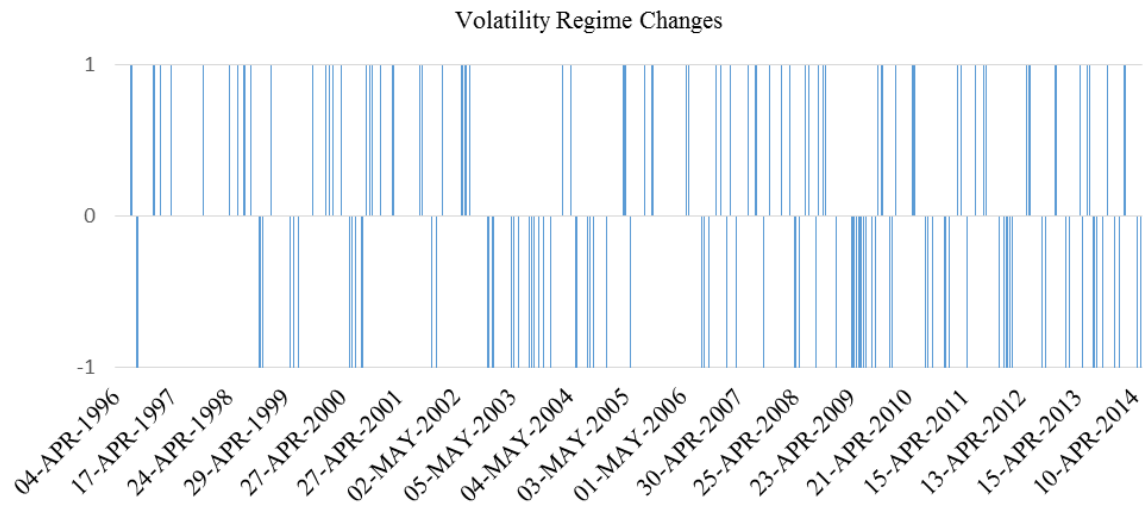


Figure 3. Distribution of realized volatility regime changes through time.

The horizontal axis measures time from the start of our sample (4 April 1996) until the end (31 May 2014). Vertical lines indicate very-high-to-neutral ($VolReg_t = +1$) and very-low-to-neutral ($VolReg_t = -1$) transitions in realized volatility.

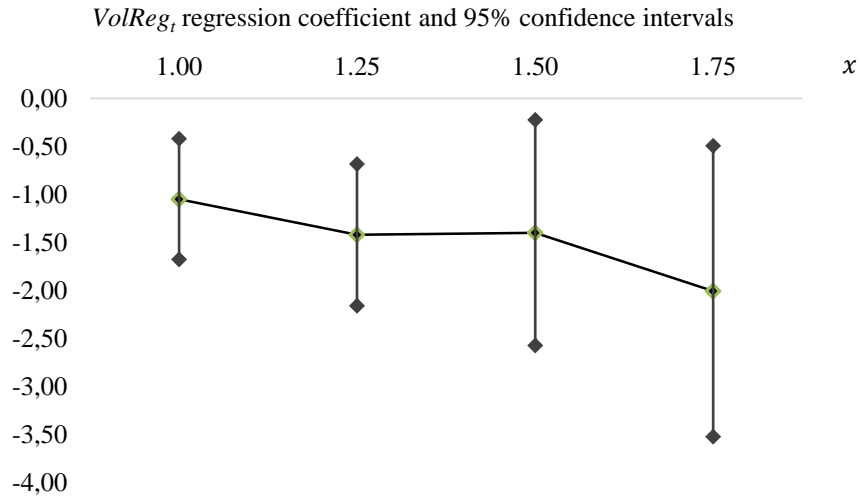


Figure 4. Stimulus strength and magnitude of the after-effect.

This figure plots the coefficient of $VolReg_t$ (vertical axis) in regression (18). Negative values of the coefficient of $VolReg_t$ are consistent with a perception bias due to after-effects. The figure plots the coefficient estimates (strength of the after-effect) and 95% confidence intervals, for four different values of x , the threshold that defines very high and very low volatility states (horizontal axis).

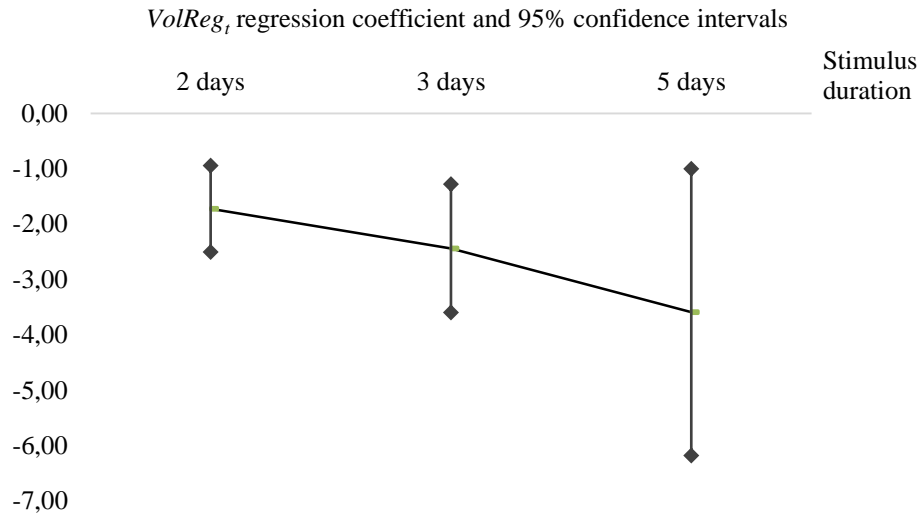


Figure 5. Stimulus duration and magnitude of the after-effect.

This figure plots the coefficient of $VolReg_t$ (vertical axis) in regression (18). Negative values of the coefficient of $VolReg_t$ are consistent with a perception bias due to after-effects. The figure plots the coefficient estimates (strength of the after-effect) and 95% confidence intervals, for three different values of the stimulus duration (the period of very high or very low volatility, horizontal axis).

Table 1**Number of transitions from high/low to neutral volatility for different threshold values**

This table reports the number of realized volatility regimes (including both very-high-to-neutral ($VolReg_t = +1$) and very-low-to-neutral ($VolReg_t = -1$) transitions) for different threshold values. Columns report different values of the threshold that defines very high and very low volatility states (volatility that is greater than x standard deviation from the mean). Rows report different values of the threshold that defines the neutral volatility state (volatility that is within y standard deviations from the mean).

y	x			
	1.00	1.25	1.50	1.75
1.00	166	80	38	14
1.25	.	114	58	26
1.50	.	.	78	36
1.75	.	.	.	44

Table 2**Differences in volatility in high/low states versus the neutral state**

This table reports absolute difference between the average log realized volatility in the high or low volatility state relative to the log realized volatility in the neutral state for different threshold values. Columns report different values of the threshold that defines very high and very low volatility states (volatility that is greater than x standard deviation from the mean). Rows report different values of the threshold that defines the neutral volatility state (volatility that is within y standard deviations from the mean).

y	x			
	1.00	1.25	1.50	1.75
1.00	0.38	0.41	0.44	0.55
1.25	.	0.37	0.40	0.46
1.50	.	.	0.35	0.41
1.75	.	.	.	0.38

Table 3
Regressions testing for a perception bias

This table reports coefficient estimates from regression (18). Negative values of the coefficient of $VolReg_t$ are consistent with a perception bias due to after-effects. The other coefficients are for control variables that are defined in the text. The regression uses the threshold parameters $x = 1.75$ and $y = 1.50$. t-statistics (in parenthesis) are calculated with heteroskedasticity robust standard deviations. ***, ** and * denote coefficients significant at the 1%, 5% and 10% level respectively.

	(1)	(2)	(3)	(4)
$VolReg_t$	-2.726** (-2.34)	-2.522*** (-2.77)	-2.606*** (-2.87)	-2.530*** (-2.79)
$\Delta LnRV_t$	0.062*** (16.18)	0.024*** (9.46)	0.034*** (9.54)	0.037*** (9.72)
$\Delta LnRV_{t-1}$			0.017*** (4.62)	0.019*** (5.09)
r_t		-3.309*** (-18.36)	-3.204*** (-17.88)	-3.209*** (-18.11)
r_t^-		0.894*** (3.02)	0.895*** (3.13)	0.847*** (3.00)
$\Delta LnVIX_{t-1}$			-0.149*** (-4.33)	-0.149*** (-4.34)
D_{2t}				-1.252*** (-5.51)
D_{3t}				-1.569*** (-6.62)
D_{4t}				-1.228*** (-5.25)
D_{5t}				-1.867*** (-8.46)
α	0.004 (0.04)	-0.292** (-2.48)	-0.296*** (-2.58)	0.923*** (4.75)
R^2	8.75%	57.40%	59.33%	60.20%

Table 4**Strength for the after-effect for different volatility state thresholds**

This table reports coefficient of $VolReg_t$ variable in regression (18). Negative values of the coefficient of $VolReg_t$ are consistent with a perception bias due to after-effects and larger magnitude coefficients suggest a stronger after-effect. Columns report different values of the threshold that defines very high and very low volatility states (volatility that is greater than x standard deviation from the mean). Rows report different values of the threshold that defines the neutral volatility state (volatility that is within y standard deviations from the mean). t-statistics (in parenthesis) are calculated with heteroskedasticity-robust standard errors.***, ** and * denote coefficients significant at the 1%, 5% and 10% level respectively.

y	x			
	1.00	1.25	1.50	1.75
1.00	-1.050*** (-3.27)	-1.155*** (-2.60)	-2.180*** (-2.87)	-3.049*** (-2.71)
1.25	.	-1.422*** (-3.78)	-2.439*** (-4.13)	-3.572*** (-4.46)
1.50	.	.	-1.400** (-2.33)	-2.530*** (-2.79)
1.75	.	.	.	-2.008*** (-2.60)

Table 5

Number of transitions from neutral to high/low volatility for different threshold values

This table reports the number of realized volatility transitions from the neutral volatility state to very high and very low volatility states. Columns report different values of the threshold that defines very high and very low volatility states (volatility that is greater than x standard deviation from the mean). Rows report different values of the threshold that defines the neutral volatility state (volatility that is within y standard deviations from the mean).

y	x			
	1.00	1.25	1.50	1.75
1.00	355	211	130	81
1.25	.	367	231	140
1.50	.	.	323	202
1.75	.	.	.	275

Table 6
Differences in volatility in high/low states versus the neutral state

This table reports absolute difference between the log realized volatility in the high or low volatility state relative to the average log realized volatility in the neutral state for different threshold values. These values (in contrast to those in Table 2) are computed for transitions from the neutral volatility state to very high and very low volatility states (in Table 2 the transitions are from very high or very low to neutral volatility). Columns report different values of the threshold that defines very high and very low volatility states (volatility that is greater than x standard deviation from the mean). Rows report different values of the threshold that defines the neutral volatility state (volatility that is within y standard deviations from the mean).

y	x			
	1.00	1.25	1.50	1.75
1.00	0.38	0.45	0.51	0.57
1.25	.	0.42	0.48	0.55
1.50	.	.	0.45	0.52
1.75	.	.	.	0.49

Table 7**Coefficients of modified volatility regime variable for different volatility state thresholds**

This table reports coefficient estimates for the modified measure of volatility regimes, $ModVolReg_t$ (which measures changes from the neutral volatility state to very high and very low volatility states). The coefficient estimates are obtained from regression (18), replacing $VolReg_t$ with $ModVolReg_t$. Unlike $VolReg_t$ (for which after-effect theory predicts a negative coefficient), after-effects theory does not predict a significant coefficient for $ModVolReg_t$. Columns report different values of the threshold that defines very high and very low volatility states (volatility that is greater than x standard deviation from the mean). Rows report different values of the threshold that defines the neutral volatility state (volatility that is within y standard deviations from the mean). t-statistics (in parenthesis) are calculated with heteroskedasticity-robust standard errors.***, ** and * denote coefficients significant at the 1%, 5% and 10% level respectively.

y	x			
	1.00	1.25	1.50	1.75
1.00	0.033 (0.09)	-0.039 (-0.10)	0.160 (0.32)	-0.369 (-0.53)
1.25	.	0.240 (0.91)	0.411 (1.18)	-0.216 (-0.46)
1.50	.	.	0.397 (1.38)	-0.191 (-0.52)
1.75	.	.	.	-0.184 (-0.60)

Table 8
Tests of the reset-of-the-mean effect

This table reports coefficient estimates for the variable $NonVolReg_t$, which is non-zero if the volatility stays in the very high or very low state (in contrast to $VolReg_t$, which is non-zero after transitioning to the neutral state). The coefficient estimates are obtained from regression (18), replacing $VolReg_t$ with $NonVolReg_t$. The ‘reset-of-the-mean’ effect, if present, should lead to significant negative coefficients for $NonVolReg_t$. Columns report different length in days of the stimulus window. Rows report different values of the threshold that defines very high and very low volatility states (volatility that is greater than x standard deviation from the mean). t-statistics (in parenthesis) are calculated with heteroskedasticity-robust standard errors. ***, ** and * denote coefficients significant at the 1%, 5% and 10% level respectively.

x	Length of stimulus window		
	2 days	3 days	4 days
1.00	-0.315*	-0.295	-0.412
	(-1.94)	(-1.47)	(-1.59)
1.25	-0.117	-0.050	-0.035
	(-0.53)	(-0.17)	(-0.09)
1.50	-0.289	-0.126	-0.303
	(-0.98)	(-0.29)	(-0.48)
1.75	-0.228	-0.254	-0.136
	(-0.52)	(-0.37)	(-0.13)

Table 9**Number of transitions from very high and very low to neutral volatility**

This table reports the number of realized volatility transitions from the very high (*VH*) state to the neutral state (Panel A) and from the very low (*VL*) state to the neutral state (Panel B) for different threshold values. Columns report different values of the threshold that defines very high and very low volatility states (volatility that is greater than x standard deviation from the mean). Rows report different values of the threshold that defines the neutral volatility state (volatility that is within y standard deviations from the mean).

y	x			
	1.00	1.25	1.50	1.75
Panel A: Transitions from <i>VH</i> to <i>M</i>				
1.00	80	40	21	8
1.25	.	58	35	18
1.50	.	.	47	27
1.75	.	.	.	29
Panel B: Transitions from <i>VL</i> to <i>M</i>				
1.00	86	37	14	3
1.25	.	53	20	5
1.50	.	.	28	6
1.75	.	.	.	12

Table 10**Differences in volatility in high/low states versus the neutral state**

This table reports the absolute difference between the log realized volatility in the very high (*VH*) volatility state relative to the neutral state (*M*) for very-high-to-neutral volatility transitions (Panel A), and the absolute difference between the log realized volatility in the very low (*VL*) volatility state relative to the neutral state (*M*) for very-low-to-neutral volatility transitions (Panel B), for different threshold values. Columns report different values of the threshold that defines very high and very low volatility states (volatility that is greater than x standard deviation from the mean). Rows report different values of the threshold that defines the neutral volatility state (volatility that is within y standard deviations from the mean).

y	x			
	1.00	1.25	1.50	1.75
Panel A: Transitions from <i>VH</i> to <i>M</i>				
1.00	0.41	0.44	0.46	0.55
1.25	.	0.41	0.44	0.48
1.50	.	.	0.40	0.42
1.75	.	.	.	0.41
Panel B: Transitions from <i>VL</i> to <i>M</i>				
1.00	0.34	0.38	0.41	0.54
1.25	.	0.33	0.34	0.42
1.50	.	.	0.27	0.39
1.75	.	.	.	0.32

Table 11

This table reports coefficients for the $VolReg_t^+$ (Panel A) and $VolReg_t^-$ (Panel B) variables, which measure very-high-to-neutral and very-low-to-neutral volatility transitions, respectively. The coefficient estimates are obtained from regression (18), replacing $VolReg_t$ with $VolReg_t^+$ (Panel A) and $VolReg_t^-$ (Panel B). Columns report different values of the threshold that defines very high and very low volatility states (volatility that is greater than x standard deviation from the mean). Rows report different values of the threshold that defines the neutral volatility state (volatility that is within y standard deviations from the mean). t-statistics (in parenthesis) are calculated with heteroskedasticity-robust standard errors. ***, ** and * denote coefficients significant at the 1%, 5% and 10% level respectively.

y	x			
	1.00	1.25	1.50	1.75
Panel A: Transitions from VH to M				
1.00	-2.410*** (-4.86)	-2.708*** (-4.39)	-3.784*** (-3.86)	.
1.25	.	-3.018*** (-5.76)	-3.924*** (-5.39)	-3.849*** (-3.93)
1.50	.	.	-2.624*** (-3.20)	-2.597** (-2.38)
1.75	.	.	.	-2.607** (-2.56)
Panel B: Transitions from VL to M				
1.00	0.226 (0.62)	0.530 (1.03)	0.225 (0.25)	.
1.25	.	0.328 (0.75)	0.152 (0.21)	.
1.50	.	.	0.617 (0.92)	.
1.75	.	.	.	-0.578 (-0.72)

Table 12**Average realized volatility in the very high and very low volatility states**

This table reports the average realized volatility in the very high (Panel A) and very low (Panel B) volatility states for different threshold values. Columns report different values of the threshold that defines very high and very low volatility states (volatility that is greater than x standard deviation from the mean). Rows report different values of the threshold that defines the neutral volatility state (volatility that is within y standard deviations from the mean).

	x			
y	1.00	1.25	1.50	1.75
Panel A: Transitions from <i>VH</i> to <i>M</i>				
1.00	22.52	22.32	25.55	.
1.25	.	23.98	26.80	29.40
1.50	.	.	27.95	30.36
1.75	.	.	.	30.55
Panel B: Transitions from <i>VL</i> to <i>M</i>				
1.00	9.30	8.49	7.32	.
1.25	.	8.91	8.51	.
1.50	.	.	9.22	.
1.75	.	.	.	8.30

Table 13**Asymmetry of the after-effect in the laboratory**

This table reports the magnitude of the after-effects post low and post high. The after-effect post low (resp. post high) is measured by the difference between the subjects' mean report of the volatility of the test stimulus after exposure to low (resp. high) volatility and their mean report of the volatility of the test stimulus after exposure to neutral volatility in the control trials. The mean difference across subjects is reported, as well as standard deviation, t-statistic and p-value, both for the original laboratory experiment (first two columns on the left) and for the follow-up experiment (second two columns).

	Original Laboratory Experiment (N=57) [Low: 2% Neutral: 10% High: 45%]		Follow-up Laboratory Experiment (N=88) [Low: 5-7%; Neutral: 13.5%; High: 40%]	
	Post low	Post high	Post low	Post high
Mean	0.34	-0.44	0.02	-0.30
Standard error	0.07	0.06	0.04	0.06
t-statistic	5.16	-7.38	0.46	-5.38
p-value	0.000	0.000	0.650	0.000