

Monetary Policy and Asset Price Bubbles: A Laboratory Experiment

Jordi Galí* Giovanni Giusti† Charles N. Noussair‡

May 2021

Abstract

Leaning-against-the-wind (LAW) policies, whereby interest rates are raised in the face of a growing asset price bubble, are often advocated as a means of dampening such bubbles. On the other hand, there are theoretical arguments suggesting that such a policy could have the opposite effect (Galí, 2014). We study the effect of monetary policy on asset price bubbles in a laboratory experiment with an overlapping generations structure. Participants in the role of the young generation allocate their endowment between two investments: a risky asset and a one-period riskless bond. The risky asset pays no dividend and thus the possibility of selling it to the next generation is its only source of value. Consequently, its price is a pure bubble. We study how variations in the interest rate affect the evolution of the bubble in an experiment with three treatments. One treatment has a fixed low interest rate, another a fixed high interest rate, and the third has a LAW interest rate policy in place. We observe that the bubble increases (decreases) when interest rates are lower (higher) in the period of a policy change. However, the opposite effect is observed in the following period, when higher (lower) interest rates are associated with greater (smaller) bubble growth. Direct measurement of expectations reveals that traders expect prices to follow previous trends and tend to correct for prior errors in their predictions.

*Centre de Recerca en Economia Internacional (CREI), Universitat Pompeu Fabra and Barcelona GSE. E-mail: jgali@crei.cat. Galí acknowledges the European Research Council for financial support under the European Union's Seventh Framework Programme (FP7/2007-2013, ERC Grant agreement number 339656), as well as generic financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S), and from the Generalitat de Catalunya, through CERCA and SGR Programme (2017-SGR-1393)

†Escola Superior de Ciències Socials i de l'Empresa de Tecnocampus-Universitat Pompeu Fabra and Economic Science Laboratory at the University of Arizona. Email: ggiusti@tecnocampus.cat

‡Department of Economics, University of Arizona. Email: cnoussair@email.arizona.edu

1 Introduction

The Global Financial Crisis of 2007-2008 and the subsequent Great Recession, commonly attributed to the bursting of housing bubbles in a number of countries, has shown how damaging the effects of a collapse in asset prices can be to the real economy. This episode has renewed the debate regarding the stance that central banks should take in response to a growing asset bubble. The common view among policy makers before the crisis was that central banks should restrict their mandate to the stabilization of inflation and the output gap. This view, however, has not gone unchallenged in the aftermath of the crisis, with many authors and policymakers arguing for a more active role for central banks in preventing overinflated asset prices by means of "Leaning Against the Wind" (henceforth, LAW) monetary policies. A policy of this type specifies that the interest rate be raised in response to asset price increases that are viewed as purely speculative, i.e. not justified by fundamentals, in order to attenuate the growth of a bubble. The rationale for LAW is that higher interest rates increase the opportunity cost of holding a bubble asset, reduce its demand, and lower the size of the bubble.

The use of LAW policies in response to asset price bubbles has been criticized on several grounds. Firstly, it is argued that in practice it is difficult, if not impossible, to establish whether or not asset prices reflect their fundamental values, given the difficulty in measuring the latter. Secondly, such a policy is viewed as potentially having undesired consequences on sectors of the economy not affected by the bubble. Thirdly, Galí (2014) calls into question, on theoretical grounds, the notion that a higher interest rate would reduce the size of the bubble. He argues that, if agents are rational, the bubble component must grow in expectation at the rate of interest. As a result, a higher interest rate could end up increasing, rather than decreasing, the size of the bubble.¹ On the other hand, Miao et al. (2019) point to the existence, under certain conditions, of equilibria in which a bubble decreases in size in response to an increase in the interest rate.

In this paper, we study the relationship between interest rate policy and bubble dynamics in the laboratory. The use of laboratory experimental methods provides a controlled environment, in which we can isolate the impact of different interest rate policies on the size and evolution of an asset price bubble. In the experiment, the bubble component of the asset price can be observed with precision: it corresponds to the price itself, since the fundamental component is zero. Different interest rate policies can be studied and their effects compared, while keeping all other aspects of the environment constant.

Our experimental environment has an overlapping generations structure. Each participant plays the role of *young* and *old* consecutively over two periods. When young, participants decide how to allocate their cash endowment between a single-period riskless bond, yielding a known interest rate, and a long-lived asset that pays no dividends. The only incentive to purchase the long-lived

¹Galí and Gambetti (2015) provide some empirical evidence, based on U.S stock market data, in support of that hypothesis.

asset comes from the possibility of reselling it to a "young" participant in the following period, hopefully with a capital gain. A positive price for that asset indicates a pure bubble. When they are old, participants collect the principal plus interest from their bonds, as well as the proceeds from the sale of the asset to the current young generation.

Within this environment, we implement three different monetary policies, each specifying a particular rule determining the interest rate on the riskless bond. Each of these policies represents a treatment in our experiment. In the first two treatments, the interest rate remains constant at either a "low" (3%) or a "high" (15%) level, over the entire experimental session. In the third treatment, the interest rate varies as a function of the change in the price of the long-lived asset in the previous period. The economy begins with a 9% interest rate, which we raise (lower) by 3% each time that the asset price increases (declines) by more than 10% from one period to the next. We interpret this third treatment as a LAW policy.

For an asset that pays a dividend and is priced at its fundamental value, a higher interest rate lowers the asset's present discounted dividend stream and thus its price. A LAW policy thus has the effect of lowering asset prices when interest rates are raised and increasing prices when rates are lowered. Higher rates increase the opportunity cost of committing capital to an asset. The attraction of using LAW policies to combat asset prices that seem too high appears to be based on this logic. However, as argued by Galí (2014), if agents have rational expectations, an asset whose price has a bubble component may behave in the opposite manner. While the initial effect of a change of interest rate on price is indeterminate, the bubble component would subsequently grow at the rate of interest (and at a greater rate if traders are risk averse and the asset carries some risk), meaning that higher interest rates may lead to larger bubbles, all else equal. In our experiment, we study the behavior of an asset that has *only* a bubble component, a favorable scenario for a LAW policy to be counterproductive. If we observe that the policy is effective in stabilizing asset prices in our study, it would suggest that the theoretical arguments against the use of a leaning against the wind policy rely on assumptions that are not satisfied in our environment, and that may not be satisfied in the outside world.

Our experiment is designed to address the following questions. (1) Do bubbles grow faster under high or low interest rates? (2) What is the initial effect of an interest rate change on the size of a bubble? (3) What is the effect of an interest rate change on the size of the bubble in the period subsequent to the change?

Our main findings can be summarized as follows. First, the average price increase for the risky asset is close to the riskless rate when the latter is low and constant, but significantly smaller when the riskless rate is high. Second, a LAW policy of increasing (decreasing) interest rates in response to asset price increases (decreases) has two effects. There is an immediate effect in the current period of decreasing (increasing) the size of the asset price bubble. However, in the subsequent period, the result is reversed and the higher (lower) interest rates tend to exacerbate (mitigate) asset price bubbles. Our results thus suggest a

means of reconciling the apparent success of LAW policies in guiding asset price bubbles in the short run with the theoretical claims arguing that they would be counterproductive afterward. Increasing interest rates reduces a bubble in the short-run, but increases it thereafter. The latter pattern is consistent with the model of Galí (2014).

Some of the asset pricing patterns suggest that the assumption of rational expectations may not be appropriate. To study this possibility, we conduct a follow-up experiment. This experiment assesses how participants form their expectations about future asset prices. We find that expectations are backward-looking, with adaptive and trend extrapolating elements, rather than rational, suggesting that the source of the departures from the theoretical framework lies in the manner that agents form expectations. We then argue that assuming *trend-following* or *adaptive* expectations can account for a number of patterns in our data.

The paper is structured as follows. In Section 2, we review the related literature and in Section 3 we propose a benchmark theoretical model. In Section 4, we describe the experimental design and state our hypotheses. The results are reported in Section 5. We discuss the role of expectations in Section 6 and we conclude with a discussion in Section 7.

2 Related literature

The four most closely related lines of experimental literature are (1) the work on asset market bubbles in experimental finance, (2) experimental studies of the effect of monetary policies on asset market behavior, (3) the experimental literature studying the behavior of economies with an overlapping generations structure, and (4) recent studies of the effects of interest rates in Learning-to-Forecast experimental environments.

The bulk of experimental research on asset bubbles builds on the seminal paper of Smith, Suchanek and Williams (1988). In the markets that they study, participants trade units of a single asset. The asset has a finite lifetime and pays a random dividend in each period. The dividend payment and (in some cases) a fixed terminal buyout value are the only sources of intrinsic value of the asset. The distribution of the dividend process is common knowledge to all traders. This means that the fundamental value is unambiguous, but only when traders are risk neutral. All cash not invested in the asset yields an interest rate of zero. Smith et al. (1988) find, however, that a price bubble tends to emerge and prices become decoupled from fundamental values. Much of the subsequent work has implemented changes to the Smith et al. (1988) environment in order to consider the robustness of the bubble phenomenon and to search for ways to prevent or eliminate it. Some approaches that have been taken include the introduction of short selling (Haruvy and Noussair (2006)), the addition of futures markets (Noussair and Tucker (2006), Porter and Smith (1995)) and the inclusion of experienced traders (Dufwenberg, Lindqvist and Moore (2005)). A number of studies have considered traders' beliefs about

future prices (Smith et al., 1988; Haruvy et al., 2007; Carle et al., 2019), and have found that they tend to extrapolate prior trends. For recent overviews of this literature see Palan (2013), Powell and Shestakova (2016) or Nuzzo and Morone (2017).

We are aware of four papers that study the effects of monetary policies on asset markets (Fischbacher, Hens and Zeisberger (2013), Giusti, Jiang and Xiu (2016) Bostian and Holt (2009), and Fenig, Mileva and Petersen (2018)). The environments in these studies have the common feature that agents have access to an alternative asset (bonds or deposit accounts) paying an interest rate each period. All of these studies find that the presence of the alternative investment is not sufficient to completely eliminate asset bubbles. Fischbacher et al. (2013) implement a policy where the interest rate on cash is varied to try to push prices of the type of asset studied by Smith et al. (1988) toward fundamentals, in effect implementing a LAW interest rate policy. They find that doing so does not have an appreciable effect on bubbles. Fenig et al. (2018) observe, in a production economy, that raising the interest rate to curb inflation has the additional effect of causing asset prices to decrease.²

Our paper relates to a literature that studies the behavior of overlapping generations economies in the laboratory. Aliprantis and Plott (1992) find that prices converge to a stationary competitive equilibrium. Marimon and Sunder (1993) observe convergence to a low-inflation steady-state in a setting with multiple equilibria, and Marimon and Sunder (1994) observe convergence to expectationally-driven cycles. Bernasconi and Kirchkamp (2000) investigate the effect of different monetary policies on inflation volatility and expectations using an overlapping generation framework. They find a tendency for oversaving and that monetary policies affect outcomes.

Finally, our work is related to a pair of recent studies of interest rate policies in Learning-to-Forecast (LTF) experimental environments. These are environments in which the individuals in the economy submit forecasts. The forecasts then determine the outcome of the economy in the current period under the assumption of optimal individual decisions given the forecasts and the underlying structure of the economy, which is typically unknown to participants. Hennequin and Hommes (2019) study how interest rate policies affect price bubbles in a LTF economy. They find that a strong response of interest rates to price levels reduces bubbles, while a weak response aggravates them. Bao and Zong (2019) also consider the effect of an LAW policy in an LTF environment. Their policy is strongly responsive to differences between asset price and fundamental value. They find that LAW has the effect of moderating price bubbles.

²Other experimental paradigms have been used to investigate the effects of monetary policies in general equilibrium economies. Assenza et al. (2013) and Pfajfar and Zakelj (2016) study the interaction between the formation of inflation expectations and the use of monetary policies within a New Keynesian framework. Assenza et al. (2013) find that an interest rate rule that reacts more than one-for-one to inflation has a stabilizing effect on prices. Pfajfar and Zakelj (2016) find that a forward-looking Taylor rule with a high reaction coefficient contributes to the reduction of inflation variability.

3 A Benchmark Theoretical Model

In this section we develop a stylized theoretical model of bubbles that we use as the basis for the design of the experiment described in Section 4. Consider an economy that consists of a sequence of overlapping generations. The population is constant, i.e. each outgoing generation is replaced with a new one of identical size. Each individual lives for two periods. When "young," an individual born in period t receives an endowment Y_t of the single good in the economy. This endowment is assumed to grow from one generation to the next at a rate $\gamma \geq 0$, i.e. $Y_t = Y_0(1+\gamma)^t$. Consumption takes place only when the individual is "old."

In order to provide for that consumption, each young individual invests her endowment in two assets: a riskless one-period bond yielding an interest rate r_t and a long-lived risky asset. The supply of the riskless bond is assumed to be perfectly elastic at the given interest rate (i.e. it can be thought of as being traded in a large "international" market). The risky asset is in positive net supply (normalized to unity) and can be traded only among domestic savers. It yields no dividends (currently or in the future), so it can be thought of as a pure bubble asset. Furthermore, there is a constant probability $\theta \in (0, 1)$, known to all agents, that the world comes to an end right after any given period. Agents learn whether the current period is the last, after trading and consumption have taken place for the period.

When old, an individual born in period t consumes an amount C_{t+1} of the single good, which equals the payoff from the investments made when young. Formally:

$$C_{t+1} = (Y_t - P_t S_t)(1 + r_t) + P_{t+1} S_t \quad (1)$$

where P_t is the price of the risky asset and S_t denotes the number of units of that asset purchased by the young. Free disposal guarantees that $P_t \geq 0$ for all t .

The young in period t make their portfolio decision in order to maximize $\mathbb{E}_t\{C_{t+1}\}$, subject to (1). When making that decision, the price of the bubble asset in the following period, P_{t+1} , is uncertain. The possibility that the bubble bursts anytime (so that P_{t+1} becomes zero) and the non-negativity of consumption justify the assumption that borrowing is not allowed, i.e. $P_t S_t \leq Y_t$. The optimality condition for the young (assuming an interior solution) is given by:

$$\mathbb{E}_t \left\{ \frac{P_{t+1}}{P_t} \right\} = 1 + r_t \quad (2)$$

i.e., the expected growth of the bubble (and, hence, its expected return) must be equal to the interest rate on the riskless bond. Letting $R_{t+1} \equiv (P_{t+1} - P_t)/P_t$ be the ex-post return on the risky asset (which corresponds to the net growth in its price), we must have

$$R_{t+1} = r_t + \xi_{t+1}$$

where $\mathbb{E}_t \{\xi_{t+1}\} = 0$. In equilibrium, $S_t = 1$ for all t , so the borrowing constraint will not be binding as long as $P_t < Y_t$, which we assume. Note that the

previous condition requires that the bubble does not grow (persistently) faster than the aggregate endowment. In the deterministic case, a sufficient condition to guarantee this is $r_t \leq \gamma$ for all t .³

When rational expectations are assumed, equilibrium condition (2) has several implications that will be tested in the context of our experiments. Firstly, it implies that when comparing two economies with different interest rates, the *average* growth of the bubble should be higher in the economy with the higher interest rate. Secondly, when looking at the evolution of the bubble over time in an economy with time-varying interest rates, it implies that in response to an interest rate increase (decrease) the *subsequent* growth of the bubble should be higher (lower).

As discussed in Galí (2014), however, (2) doesn't pin down the *level* of P_t , i.e. the size of the bubble itself, which remains indeterminate.⁴ A corollary of this indeterminacy result is that the size of the bubble may respond systematically to unanticipated changes in the interest rate (or to news of any other nature, for that matter) without violating the equilibrium condition. In particular, an equilibrium such that $\xi_t = \alpha(r_t - \mathbb{E}_{t-1}\{r_t\}) + \xi_t^*$ (where ξ_t^* satisfies $\mathbb{E}_{t-1}\{\xi_t^*\} = 0$ while being uncorrelated with interest rate innovations) is consistent with a rational expectations equilibrium, with the sign (as well as the size) of α being indeterminate. In the experiments described below, changes over time in the interest rate (associated with LAW policies) are always anticipated (since they are a function of the lagged bubble growth), thus the previous indeterminacy result does not apply. Accordingly, changes in the size of the bubble should be related only to the lagged interest rate (in addition to having a noisy surprise component) in a rational expectations equilibrium.

4 Experimental Design

We conducted a total of nine sessions for our main experiment. Sixteen subjects participated in each session for a total of 144 subjects.⁵ The sessions were conducted at the LEEX laboratory at Universitat Pompeu Fabra, in Barcelona, Catalonia, Spain. Subjects were undergraduate students in business and economics. The Ztree platform (Fischbacher, 2007) was used to program the computerized environment. A session took approximately two hours and the average payment was 23€ per subject.⁶ The earnings received by each subject were proportional to her payoff in the experimental economy, in keeping with conventional procedures in experimental economics.

³In the stochastic case, one must assume in addition that $P_0 \ll Y_0$ and that positive innovations in the size of the bubble are small enough so that $P_t < Y_t$ for all t .

⁴In particular, the bubble may burst at any time, with $P_t = 0$ subsequently until the end of trading.

⁵These totals of sessions and participants do not include the sessions of the follow-up experiment described in section five.

⁶At the beginning of each experimental session, we provided written instructions to all subjects and read them aloud. This took approximately 20 minutes.

The experimental environment has an overlapping generation structure. Each participant is active for two consecutive periods, in which he plays the role of a member of the “young” and “old” cohorts sequentially. Each period, a cohort of four young participants interacts with an old cohort with four members.⁷

When entering the economy, each young participant receives an endowment that she has to allocate between two assets: (i) a riskless one-period bond in perfectly inelastic supply, yielding a pre-announced interest rate and maturing in the next period, and (ii) an infinitely-lived risky asset. This asset does not yield a dividend at any time, but it can be sold in the following period, thereby possibly leading to capital gains or losses. The fundamental value of the risky asset is zero, and therefore we can think of it as a “bubble asset”. If the asset trades at a positive price, it is a pure bubble. The participant allocates endowment to the bubble asset by purchasing it in a market that is operating in all periods. More precisely, “young” participants make bids on the existing units of the bubble asset held by the “old,” with the cash endowment not used to purchase the asset automatically allocated to the purchase of the bond at the end of the young period.

At the end of the period in which a participant is young, she earns interest on her holdings of the bond, and her holdings of the bubble asset are carried over into her old period. In her old period, she can sell units of the bubble asset to young agents. The lifetime payoff of each subject is realized at the end of her old period, and is made up of the principal plus interest on her bonds, plus the proceeds from the sale of her bubble asset holdings.⁸

Each experimental session consists of several horizons, where each horizon refers to an entire multi-period economy. At the beginning of each horizon, all endowments are reinitialized at the same starting level. Each horizon, in turn, consists of a random number of periods, with a constant probability (equal to 0.1) that the horizon would end after any given period.⁹ This is a standard

⁷The implementation of the overlapping generations environment follows Marimon and Sunder (1993), but with some differences. In each period, a subset of the individuals in the laboratory enters the market as “young” participants and remains in the market for two consecutive periods, consisting of a young, followed by an old, period. In our study, an agent is then reborn as a new young agent in the period immediately following her old period, while in Marimon and Sunder (1993), there was always at least one period between an individual’s old period and the next time they were reborn as a young agent.

⁸The aggregate supply of the long-lived asset (i.e. the number of units available for purchase every period) always remains unchanged. In order to guarantee this condition, units that are unsold by current old agents at the end of a given period through the auction mechanism are automatically and randomly allocated to current young buyers who are forced to pay 1.5 times the average transaction price in the period. For such transactions, the price received by the old sellers is 0.5 times the average transaction price. In this way, we maintain a constant supply of shares while incentivizing trade. In the event of no trade in a given period, the average transaction price is set at 0. Thus, a zero price equilibrium, which does exist in our laboratory environment, is a feasible outcome in the experiment, though it is never observed.

⁹As discussed in Duffy (2016) this is theoretically equivalent to an infinite horizon with a discount factor $\beta = .9$. Including a constant probability of ending serves to create the same incentives that exist under an infinite horizon, and also serves to make all periods alike in terms of the expected number of periods remaining. It induces no trend in pricing over time, and does not affect the prediction that the price of the asset would grow at the rate of interest.

technique in experiments evaluating infinite horizon models, and was first used in experimental economics by Roth and Murningham (1978), and in experiments for long-lived assets by Camerer and Weigelt (1993). For recent examples of its use in asset market experiments, see Crockett et al. (2019) or Carbone et al. (2021).¹⁰ Each experimental session includes as many horizons as we could conduct in two hours and 30 minutes. We did not start a new horizon unless there remained at least 20 minutes left until the scheduled end of the session.

In each of the nine sessions, there are two horizons and thus two markets operating concurrently, each with 8 participants. At the end of each horizon, before starting a new market, participants in both markets are randomly re-grouped. This is common knowledge.

There are three treatments, LOW, HIGH, and LAW, with the treatments differing only in the interest rate policy that is in effect. The policy defines the interest rate to be paid on the bond in each period.

LOW Treatment: A constant interest rate of 3% prevails throughout the entire experimental session.

HIGH Treatment : A constant interest rate of 15% is in force throughout the entire session.

LAW Treatment: A Leaning-Against-the-Wind policy is in effect, whereby the interest rate is adjusted as a function of the percentage change in the price of the bubble asset in the previous period. Specifically, the interest rate starts at 9% at the beginning of each horizon. It is increased (reduced) by 3% whenever the price of the bubble asset rises (declines) by more than 10% from one period to the next. Otherwise, the interest rate is left unchanged. Furthermore, we set upper and lower bounds for the nominal interest rate of 3% and 15%, respectively. The resulting interest rate rule can be written as:

$$r_t = \begin{cases} \min\{r_{t-1} + 0.03, 0.15\} & \text{if } P_{t-1}/P_{t-2} \geq 1.1 \\ \max\{r_{t-1} - 0.03, 0.03\} & \text{if } P_{t-1}/P_{t-2} \leq 0.9 \\ r_{t-1} & \text{otherwise} \end{cases}$$

Only one treatment is in effect in a given session. In the LAW treatment, participants are not explicitly informed about the policy rule nor given any reason for the changes in the interest rate. In all cases, the interest rate prevailing in any given period is communicated to participants at the beginning of the period.

¹⁰There are a number of other ways that the horizon could be terminated while preserving the continuation incentives in the model. For example, in their OLG economies, Marimon and Sunder (1993) asked some participants currently not participating in the economy, on a rotating basis, to submit incentivized price forecasts for the upcoming period. When the experimenters ended the economy, they used the prices forecasted by the observers to calculate the payoffs that those who were young when the economy terminated would have earned in their old period. These payoffs are then paid out to the young as if the horizon had continued into their old period.

The trading mechanism is a continuous double auction with an open order book (Smith, 1962). In each trading period, subjects can initiate a potential transaction by posting offers to buy or to sell the bubble asset. Each offer is for a single unit of asset, but subjects can post multiple offers to buy or sell. Active buy and sell orders are ranked and displayed in two separate columns, with the best available offers listed at the bottom. Subjects execute a trade by selecting the best order available at any point in time and clicking on the "buy" or "sell" button located at the bottom of the order book. Short sales or buying on margin are not permitted.

Each trading period lasts for 150 seconds. At the end of a trading period, a summary screen appears, showing relevant information, such as the participant's current quantity of risky asset and bonds held, and the interest they receive in the period.

Young participants are endowed with 5000 ECU in cash at the beginning of period 1 of each horizon. The size of the endowment of each new young generation increases at a 15% rate per period. The reason for this increase is the theoretical requirement of an interest rate at or below the growth rate of the economy's resources for a rational bubble to exist in equilibrium. Since our highest possible interest rate is equal to 15% per period, our strategy guarantees that this requirement is satisfied. In period 1, each current "old" participant is endowed with 3 units of the bubble asset in order to initialize the market. Therefore, in any period, 12 units are exchanged between the old and young generations.

The payoff for each subject over each two-period lifetime is calculated at the end of the old period. It is given by the initial endowment, plus interest on bonds plus the capital gains (or minus the losses) from the purchase and sale of the bubble asset. Since each subject typically plays several "lives" in each horizon and there are several horizons in each experimental session, the program randomly selects one of the "lives" to compute as a participant's earnings from the experiment. The experimental currency is converted into Euros at a fixed, pre-specified exchange rate.

5 Results

5.1 Comparison between treatments

Figure 1 shows the average price in each of the three treatments. In the figure, the horizontal axis indicates the market period. All horizons are included in the figure. Periods beyond seven are not shown because they have fewer observations since relatively few horizons reach them. The vertical axis is the average price in a period. The figure shows that the average price in the HIGH treatment, 278, is slightly greater than under the LOW treatment, 251, though the difference is not significant ($t = -0.78$, $p=0.44$; Mann-Whitney rank-sum test $z=-1.16$, $p=0.24$). Thus, higher interest rates exert neither a dampening nor an exacerbating effect on average asset prices. The average price in the LAW treatment is equal to

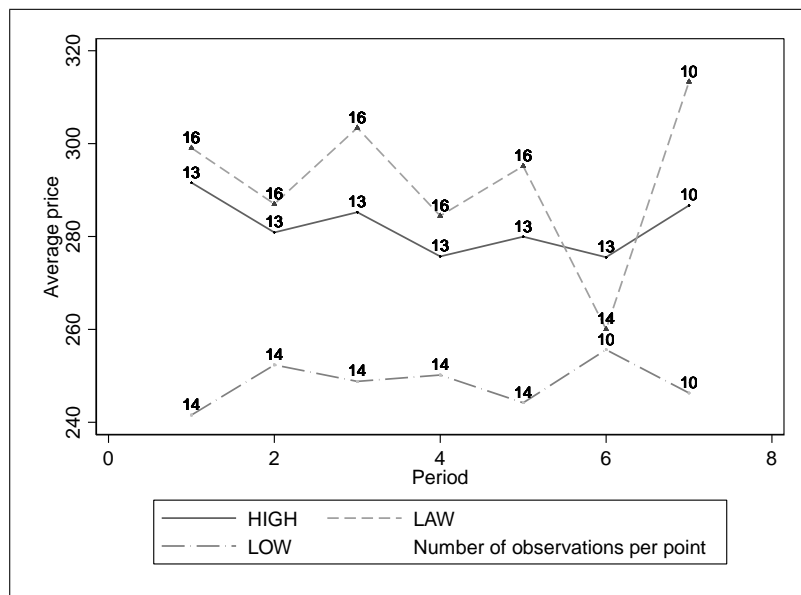


Figure 1: Asset prices in treatment by period

289, and not statistically different neither from the HIGH treatment ($t=0.28$, $p=0.78$) nor from the LOW treatment ($t=1.13$, $p=0.27$). At the period 1 prices, on average young agents are holding on average 16.3%, 17.3% and 18.0% of their initial wealth of 5000 Eurux in the risky asset, in the LOW, HIGH and LAW treatments, respectively. This share tends to decline over time as the cash available increases more rapidly than asset prices on average.

The return of the asset at time t is given by:

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}} \quad (3)$$

The figure gives the impression that there is at most a modest tendency for the price of the bubble asset to rise over time, suggesting that returns average close to zero. Statistical tests confirm this impression. The average return of the asset in the LOW treatment, 1.98%, is not significantly different from the risk free rate of 3% ($t=-0.39$, $p=0.70$; $z=-1.791$, $p=0.073$), which is consistent with a rational bubble. However, it is also not significantly different from 0% ($t=0.76$, $p=0.46$; $z=0.47$, $p=0.637$). The average return for the asset in the HIGH treatment, -0.7%, is significantly lower than the risk free rate of 15% ($t=-12.24$, $p < 0.01$; $z=-3.18$, $p=0.0$), though not different from zero ($t = -0.53$, $p = 0.60$, $z = -0.45$, $p = 0.64$). The returns are also not significantly different between the LOW and HIGH treatments at conventional levels of significance ($t=0.89$, $p = 0.37$; $z = 0.72$, $p = 0.46$).

Thus, we do not observe significantly different bubble size or growth between the LOW and HIGH treatments. While the rate of bubble growth is close to

the interest rate under LOW, it is significantly lower than the interest rate under HIGH.¹¹

We also compare the volatility of returns among the three treatments. To compute a measure of volatility, we first compute the absolute value of the difference in return between one period and the next, $|R_t - R_{t-1}|$, and then we calculate the average of these differences for each market. Using each market as a unit of observation, we perform a t-test to check for differences between treatments. The average volatility for the HIGH treatment is equal to 0.23, which is not significantly different from the average volatility of 0.24 in the LAW treatment ($t = -0.26, p = 0.79$), nor from the average of 0.16 of the LOW treatment ($t = 1.14, p = 0.26$). The comparison between LOW and LAW shows a marginally significant difference at the 10% level between the two treatments ($t = 1.81, p = 0.08$) with greater volatility under the LAW treatment.¹² Thus, we do not find that a LAW policy reduces asset price volatility.

5.2 Effect of leaning against the wind

In this subsection, we first consider the association between interest rates and prices under the LAW treatment. Figure 1 gives the impression that LAW treatment exhibits large oscillations. The claim of advocates of Leaning Against the Wind policies is that an interest rate increase (decrease) in period t would reduce (increase) the price in period t . In our experiment, a LAW policy takes the form of an interest rate change in period t being enacted in response to a price change between periods $t-1$ and t . That is, $r_t > r_{t-1} + .03$ if $P_{t-1} > 1.1 * P_{t-2}$ and $r_t < r_{t-1} + .03$ if $P_{t-1} < .9 * P_{t-2}$. This change in r_t would affect the return from holding bonds between periods t and $t + 1$, changing the demand for the bubble asset. A higher interest rate would increase the opportunity cost

¹¹To conduct a test for a rational bubble, we estimate the following regression for the pooled data for all treatments.

$$R_{t+1} = \alpha + \beta r_t + \varepsilon_{t+1}, \quad (4)$$

and test the hypothesis that $\beta = 1$. A rational bubble requires the price of the asset to grow at the riskless rate. We also test whether $\alpha > 0$, which would indicate that the asset carries a risk premium. We estimate $\alpha = 0.01$, which is not significantly different from 0. However, the estimated $\beta = -0.07$, which is highly significantly different from 1. The lack of a positive overall relationship between the return of the asset and the riskless rate calls the assumption of rational expectations into question. In section 5, we investigate agents' expectations in detail.

¹²We can also consider other measures of dispersion. We distinguish between *within-horizon* dispersion of returns from period to period within a horizon, and *between-horizon* dispersion, a measure of the variance in average return across different horizons. We define the within-horizon dispersion as $\sum_{t=1}^T (R_t - R_{t-1})^2$, where R_t is the return of the asset in period t and R_{t-1} is the return in the preceding period. This measure averages 0.09 for the LOW, and 0.12 for the HIGH treatment, respectively. The difference between the two treatments is not significant ($t = -0.37, p = 0.71; z = -1.26, p = 0.21$). We define between-session dispersion as $\sum_{j=1}^J (\bar{R}_j - \bar{R})^2$, where \bar{R}_j represents the average return in a horizon of a given treatment and \bar{R} is the average return in the treatment. The values of this measure are 0.0088 in LOW, and 0.0019 in the HIGH, treatment, suggesting more between-session heterogeneity in the LOW treatment.

of holding the asset, leading to lower demand and consequently a lower price. A lower interest rate would have the opposite effect.

We consider this relation with two measures. The first is a test of whether interest rate increases (decreases) correlate with a contemporaneous decrease (increase) in asset price. That is, we test whether

$$(P_t - P_{t-1}) * (r_t - r_{t-1}) < 0 | r_t \neq r_{t-1} \quad (5)$$

in a significant majority of instances. We also evaluate the effect of interest rate changes on the current price trend. Since the price trend is $P_t - P_{t-1}$, the test is:

$$(P_t - P_{t-1} - (P_{t-1} - P_{t-2})) * (r_t - r_{t-1}) < 0 | r_t \neq r_{t-1} \quad (6)$$

However, as noted by Galí (2014), there may be an offsetting effect of interest rate changes on subsequent price trends. The interest rate in period t affects the return of a bond held from period t to $t + 1$. Therefore, the return of the bubble asset between periods t and $t + 1$ should be affected in a similar direction as the interest rate change. In our experiment, we measure this effect as a positive relationship between a change in interest rate in period t and a change in the price of the bubble asset between periods t and $t + 1$. That is, we test whether:

$$(P_{t+1} - P_t) * (r_t - r_{t-1}) > 0 | r_t \neq r_{t-1} \quad (7)$$

after a significant majority of interest rate changes. We also consider how interest rate changes affect the change in the return of the asset compared to the prior trend. That is, we test whether:

$$(P_{t+1} - P_t - (P_t - P_{t-1})) * (r_t - r_{t-1}) > 0 | r_t \neq r_{t-1} \quad (8)$$

Table 1 shows the percentage of observations that are consistent with each of the measures. The first column of the table is the measure being considered, the second column indicates the total number of periods in which there were interest rate changes, and for which the test can be evaluated. The third column contains the number of periods in which the relationship in the first column was observed. The table shows that the large majority of observations are consistent with all four measures. In the last column, we report the p-values from binomial tests of the hypothesis that the probability that each criterion is satisfied is 0.5, and in all cases, $p < .001$.

These results indicate that the immediate effect of a LAW policy is to push asset prices in the intended direction. Interest rate increases have the effect of lowering prices, while rate reductions increase them in the current period. However, the effect one period ahead is very different. In the overwhelming majority of cases, the price bubble in the asset is increased by a rate hike and lowered by a rate decrease.

Measure	Obs. with $r_t \neq r_{t-1}$	Cons. with measure	p-value of binomial test
$(P_t - P_{t-1}) * (r_t - r_{t-1}) < 0$	54	51	<0.001
$(P_t - P_{t-1} - (P_{t-1} - P_{t-2})) * (r_t - r_{t-1}) < 0$	48	47	<0.001
$(P_{t+1} - P_t) * (r_t - r_{t-1}) > 0$	45	39	<0.001
$(P_{t+1} - P_t - (P_t - P_{t-1})) * (r_t - r_{t-1}) > 0$	41	38	<0.001

Table 1: Effects of Leaning Against the Wind Policy on Current and Future Prices of the Risky Asset

6 The Role of Expectations

In section 5.1, we observed that the price of the asset under LOW tends to increase modestly, but at a rate that was not significantly different from zero. The price of the asset in the HIGH treatment also did not change at a rate significantly different from zero. It was negative in sign, despite the positive interest rate. The price trend in HIGH is significantly different than that predicted by a no arbitrage condition between the bond and the bubble asset. In this section, we explore what might be behind these patterns. We consider the potential connection between the price trend and the expectations individuals hold.

Much prior experimental evidence supports the notion that people have backward-looking expectations of some form. Indeed, in experimental tasks that involve predicting market prices, prior studies consistently find that the majority of subjects use *trend extrapolation*, which makes expectations depend only on previous price data, or *adaptive expectations*, which involves basing current predictions on past prices and past own predictions (see e.g. Haruvy et al. (2007) or Hommes et al. (2008)). Indeed, one should only hold rational expectations if it is common knowledge that others do so as well. Therefore, the failure to observe rational expectations prices does not necessarily mean that there is any irrational behavior, only that it is believed to be possible. That is, a lack of common knowledge of rationality can make expectations of prices, and in turn prices themselves, depart from fundamentals, even when all traders are rational.

6.1 A follow-up experiment to elicit expectations

To establish what might be the appropriate assumptions on expectations to describe behavior in our experiment, we conducted additional experimental sessions with a different cohort of participants than those in the prior treatments. The new participants were randomly selected from the same subject pool, students in business and economics from Universitat Pompeu Fabra. A total of 30 subjects, divided among four laboratory sessions, participated in this experiment.

In these sessions, subjects received information about prices in sessions of the

three main treatments of the experiment. Specifically, participants were given the following information: (1) the average asset price in the current period P_t , (2) the prior interest rate on the bond r_{t-1} , and (3) the interest rate that would prevail between the current period and the next, r_t . They were asked to forecast the average price of the asset in the next period, P_{t+1} . After that, they observed the actual data from the period that they had forecast, and were asked to forecast the average price of the following period. The sequences of realized prices and interest rates for the forecasting task were randomly taken from the market data of the three treatments. A total of 22 sequences (out of a total of 43) were used for this experiment, and each of the sequences was used once.

Participants received a monetary payment based on the accuracy of their forecasted prices compared to the actual realizations. Following Haruvy et al. (2007) the incentive scheme for correct predictions was established through a simple payment schedule (shown as Table 1B in the instructions in online Appendix B). The experiment took on average 70 minutes, including instruction, and the average payment was 25€. In total, we obtained 2642 price predictions from these sessions.

6.2 What expectations do participants have?

We begin our analysis of the data by considering how well predictions conform to rational expectations. There are two senses in which we can view expectations as rational. The first notion of rationality is that expectations are unbiased predictors of the subsequent price. To test this, we estimate

$$R_{t+1} = \alpha + \beta R_{t+1}^e + \varepsilon_{t+1}, \quad (9)$$

where $R_{t+1}^e = (P_{t+1}^e - P_t)/P_t$. P_{t+1}^e is an individual's prediction of the price in period $t + 1$ and R_{t+1}^e is the predicted return for that period. If expectations are unbiased, we would observe that $\alpha = 0$ and $\beta = 1$.

However, as shown in the first column of Table 2, the estimated β coefficient is equal to 0.139, significantly different from 1 (p -value < 0.01). The estimate of α is significantly different from 0 (p -value < 0.01) and equal to 0.029. A F-test of the hypothesis that the two coefficients are as predicted yields a statistic of $F = 261.59$ ($p < .0001$). The reported R^2 is equal to 0.0057, so that the model explains a very small amount of variation. The data are inconsistent with the hypothesis that expectations are unbiased predictors of upcoming prices. The estimates in Columns 3 and 5 in the table, which include fixed effects for markets and for individuals, yield very similar estimates.

Another notion of rational expectations is that individuals apply a theoretical model that assumes common knowledge of rationality to the data that they observe to form their expectations. Under this assumption, an agent expects the return on the asset to equal the rate of interest on bonds. To test this, we can estimate the following model:

$$R_{t+1}^e = \alpha + \gamma r_t + \varepsilon_{t+1} \quad (10)$$

Under the null hypothesis, $\alpha = 0$ and $\gamma = 1$. The estimates are shown in the second column of Table 2. The estimated γ coefficient is equal to -0.043, significantly different from 1 (p -value < 0.01). The estimate of α is 0.009, not significantly different from 0. A F-test of the hypothesis that the two coefficients are equal to the predicted values yields $F = 529.34$ ($p < .0001$). The reported R^2 is essentially 0. The estimates in Columns 4 and 6, which allow for market and individual fixed effects, lead to similar conclusions. Thus, the data are inconsistent with the hypothesis that expectations reflect the application of a fully rational, forward-looking, theoretical model to the available data.

	(1)	(2)	(3)	(4)	(5)	(6)
	R_{t+1}	R_{t+1}^e	R_{t+1}	R_{t+1}^e	R_{t+1}	R_{t+1}^e
R_{t+1}^e	0.139*** (0.038)		0.129*** (0.038)		0.136*** (0.036)	
r_t		-0.043 (0.059)		0.049 (0.077)		-0.028 (0.162)
Constant	0.029*** (0.006)	0.009 (0.007)	0.207*** (0.052)	0.007 (0.016)	0.022 (0.029)	-0.003 (0.01)
N	2323	2323	2323	2323	2323	2323
R^2	0.006	0.000	0.065	0.014	0.006	0.008
Market FE			✓	✓		
Individual FE					✓	✓

Table 2: OLS estimation of Models of Rational Expectations.

Beta coefficients reported and robust standard errors in parentheses. Columns (3) and (4) include fixed effects for each market and Columns (5) and (6) include fixed effects for each individual. *** significant at 1%, **significant at 5%, * significant at 10%

In view of the fact that our incentivized observers do not base their forecasts on a model with rational expectations, we ask if they employ rules that have been widely observed in other experimental paradigms. We consider three alternative models in the regressions reported in Table 3. In the first and fourth columns of the table, we report estimated parameters of a simple Trend Extrapolation rule, given by:

$$P_{t+1}^e - P_t = \beta_0 + \beta_1(P_t - P_{t-1}) + \varepsilon_{t+1}. \quad (11)$$

We also estimate an adaptive rule:

$$P_{t+1}^e - P_t^e = \beta_0 + \beta_1(P_t - P_t^e) + \varepsilon_{t+1} \quad (12)$$

and the estimates are given in columns 2 and 5. Both the trend following and the adaptive models are estimated with and without the interest rate r_t as a control.

The simple trend extrapolation rule that we consider assumes that individuals use only prior price data and assume that the most recent trend will

continue. There has been support for models with this feature in the asset market experiments of Haruvy et al. (2007). Adaptive rules have been proposed and supported in Learning-to-Forecast experiments, often in combination with trend following rules (Anufriev and Hommes (2012)). The simple adaptive rule we evaluate compares one’s prediction in the preceding period with the actual outcome, and updates the prediction for the subsequent period in the direction of the outcome.

There is also the possibility that expectations are governed by a mix of the two types of rule. We also estimate a hybrid rule that includes trend following and adaptive expectation components:

$$P_{t+1}^e - P_t^e = \beta_0 + \beta_1(P_t - P_{t-1}) + \beta_2(P_t - P_t^e) + \varepsilon_{t+1} \quad (13)$$

Table 3 shows that the trend following and adaptive terms are typically highly significant, and the R^2 of regressions that include the adaptive term are between .45 and .53. In Appendix A, we also report the same analysis separating the LAW treatment, in which the interest rate is subject to frequent changes, from the LOW and HIGH treatments. In Appendix A, we also present the estimates for the regression models of Table 3 including both market and individual fixed effects. The fixed effect specifications yield almost identical results to those reported in Table 3.

Interestingly, the estimated coefficients on the recent trend, $P_t - P_{t-1}$, in equations (11) and (13) are negative. This means that individuals expect prices to exhibit a form of negative autocorrelation or mean reversion, in that the larger the change in a given direction, the less in the same direction the next change is expected to be. However, the estimates in the tables in Appendix A show that this expected mean reversion under the Hybrid model is confined to the LAW treatment, and does not appear in the data for the LOW and HIGH treatments. This indicates that under the HIGH and LOW treatments, observers expect a continuation of previous trends, while in the LAW treatment they (correctly) anticipate a reversal in the price trend, presumably resulting from the LAW policy response.

The coefficients on $P_t - P_t^e$ in equations (12) and (13) are also positive, in agreement with previous work. This means that the more prices have exceeded (been lower than) an individual’s expectations in period t , the more she adjusts her predictions upward (downward) in the next period. This is a consistent with the idea that individuals try to correct for prior errors in their predictions by adjusting their next prediction in the direction of the previously observed market price. This pattern holds in all three treatments.

	(1)	(2)	(3)	(4)	(5)	(6)
	Trend Follow	Adapt. Exp.	Hybrid	Trend Follow	Adapt. Exp.	Hybrid
$P_t - P_{t-1}$	-0.34*** (0.02)		-0.14*** (0.03)	-0.34*** (0.02)		-0.21*** (0.03)
$P_t - P_t^e$		0.61*** (0.02)	0.71*** (0.03)		0.66*** (0.02)	0.82*** (0.03)
r_t				-224.20*** (9.30)	-167.64*** (8.64)	-192.94*** (9.22)
<i>Cons</i>	-11.21*** (0.62)	-7.20*** (0.61)	-8.27*** (0.59)	10.97*** (0.78)	8.90*** (0.81)	9.76*** (0.79)
N	2323	2323	2323	2323	2323	2323
R^2	0.20	0.45	0.46	0.33	0.50	0.53

Table 3: OLS estimates for Backward-Looking Expectational Rules. Coefficients reported and robust standard errors in parentheses. *** significant at 1%, **significant at 5%, * significant at 10%

To consider whether expectation formation varies between treatments, we estimate the following version of the hybrid model:

$$\begin{aligned}
P_{t+1}^e - P_t^e = & \beta_0 + \beta_1(P_t - P_{t-1}) + \beta_2 HIGH + \beta_3 LAW + \beta_4 HIGH * (P_t - P_{t-1}) \\
& + \beta_5 LAW * (P_t - P_{t-1}) + \beta_6 (P_t - P_t^e) + \beta_7 HIGH * (P_t - P_t^e) + \beta_8 LAW * (P_t - P_t^e) + \varepsilon_{t+1}
\end{aligned}
\tag{14}$$

The model includes interaction terms between treatments and the trend-following and adaptive terms. This allows us to consider whether the coefficients of these terms differ among the three treatments. The estimates are reported in Table 4. An F-test of the restriction that $\beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_7 = \beta_8 = 0$ yields a highly significant value of $F = 137.3$ for the specification in Column 1. For Column 2, an F-test of the restriction that all of the interaction terms = 0 results in $F = 60.5$, which is also highly significant. This shows that the expectational rules significantly differ depending on the interest rate environment.

	(1)	(2)
$P_t - P_{t-1}$	-0.05 (0.09)	-0.05 (0.09)
<i>HIGH</i>	-14.37*** (1.42)	-3.98 (24.33)
<i>LAW</i>	-0.33 (1.42)	7.33 (6.56)
<i>HIGH</i> * ($P_t - P_{t-1}$)	0.23** (0.09)	0.23** (0.09)
<i>LAW</i> * ($P_t - P_{t-1}$)	-0.47*** (0.09)	-0.46*** (0.09)
$P_t - P_t^e$	0.98*** (0.08)	0.98*** (0.07)
<i>HIGH</i> * ($P_t - P_t^e$)	-0.40*** (0.08)	-0.40*** (0.08)
<i>LAW</i> * ($P_t - P_t^e$)	-0.03 (0.08)	-0.02 (0.08)
r_t		-86.79 (202.9)
<i>LAW</i> * r_t		-57.36 (204.52)
<i>Cons</i>	-0.36 (0.56)	2.27 (6.24)
N	2323	2323
R^2	0.61	0.61

Table 4: OLS estimation for Hybrid Expectational Rule, Allowing Coefficients to Differ by Treatment.

Coefficients reported and standard errors in parentheses. *** significant at 1%, **significant at 5%, * significant at 10%

The estimates reveal that the trend-following component, $P_t - P_{t-1}$, is stronger in the HIGH condition than under LOW, and that it is negative in the LAW treatment. This is consistent with the notion that a reversal of the trend is expected as a result of the LAW policy. In the LOW treatment, the adaptive term $P_t - P_t^e$ is very close to 1, meaning that predictions fully correct for the forecast error of the previous period. Under LAW, the correction is very similar to LOW, while under HIGH, the adjustment is about 60 percent of the prior prediction error. Overall, under the HIGH riskfree rate, there is more weight placed on price trends and less on previous forecasting errors.

6.3 Expectations and price dynamics

As described in section five, the behavior of the LOW and HIGH treatments shows some unexpected patterns. The first is that the price of the asset does

not exhibit a significant increase over time. Indeed, in the HIGH treatment, the price actually decreases over time on average, though the effect is not significant. The second is that the price changes between one period to the next, $P_t - P_{t-1}$, in the HIGH and LOW treatments are not significantly different from each other. Indeed, the trend is actually lower under HIGH than under LOW. While these patterns are inconsistent with rational expectations, we argue here in this subsection that they are consistent with expectations of the form described in equations (11) or (13).

We show how under appropriate assumptions on expectations, the two following counterintuitive patterns can arise: (1) a downward trend in prices despite a positive riskfree interest rate, and (2) smaller (larger) price appreciation (depreciation) under higher riskless rates. While these patterns are extreme caricatures of what we have observed, the arguments can be easily refined to capture the more modest effects that we have obtained. Consider a trader with expectations as in equation (13), and for simplicity assume that $\beta_0 = 0$. Denote the expectation the agent holds at the beginning of period s about the price of the asset in period t as P_t^s . Before period 1, there is no basis to form expectations, because there is no past data. Thus, we assume that a young agent's initial expectations for the price level in periods 1 - 3, P_1^1, P_2^1, P_3^1 are arbitrary. However, for a given expected price for period 1, we assume that $P_1^1 = P_2^1 = P_3^1$.¹³ That is the individual expects that the asset price will remain constant for the first three periods.

If the return on a riskless bond between periods 1 and 2 is the interest rate r , then the willingness-to-pay of traders for the long-lived asset in period 1 is $\frac{P_2^1}{1+r}$. This is because the return on bonds is r , and the return on the asset and the bond must be equal. Therefore, the period 1 price must be $\frac{P_2^1}{1+r}$. It is clear that if all traders approach the valuation of the asset in this manner, the price in period 1 will be below the level traders predicted, so that $P_1 = P_1^1/(1+r) < P_1^1$, where P_t is the observed market price in period t . This means that prior period 1 predictions are incorrect, in line with the relaxation of the assumption of rational expectations. The actual period 1 price is lower than predicted.

Now consider period 2. An agent with Hybrid Expectations will now update her prediction for period 2 downward, in view of the fact that the price in period 1 was lower than their prediction, so that $P_2^2 < P_2^1$. Because there is not yet enough data to establish a trend, they will also have the expectation that $P_3^2 = P_2^2$. Thus, in period 2, the current young will value the asset at $P_2 = \frac{P_3^2}{1+r}$. Notice that this price is lower than $1+r$ times the observed price in period 1, so that $P_2 < (1+r) * P_1$, since $P_2 = \frac{P_3^2}{1+r} < P_2^2 < P_2^1 = (1+r) * P_1$. This means that the asset appreciates at a rate less than r , and may even depreciate.

In period 3, agents will lower their expectations of the current price for pe-

¹³The assumption that the expectations of individuals who are just entering a new asset market are that the price trajectory will remain constant is strongly supported by the results of Haruvy et al. (2007), who observed a strong tendency for flat expectations among inexperienced agents in their markets, even though all agents knew that the fundamental value of the asset was decreasing over time.

riod 3 in light of the fact that $P_2 < P_2^2$, so that $P_3^3 < P_3^2$. A similar effect will occur in subsequent periods $t > 3$ and any increase in price will be strictly less than the risk-free rate and may even be negative. The trend-following effect will lead to anticipation of a continuation of the price trend, and the adaptive expectations component will continually push agents to update their beliefs in the direction of lower prices.

Example: Suppose that $P_1^1 = P_2^1 = P_3^1 = 250$, $r = 0.09$, and agents have the expectations given in (13) with $\beta_1 = \beta_2 = 1/2$. In period 1, the price would equal $P_1 = \frac{P_1^1}{1+r} = 250/1.09 = 229.36$. Then, a young agent's expectation in period 2 of the upcoming period 2 price would be $P_2^2 = P_1 + \beta_2(P_1 - P_1^{e1}) = 229.36 + 0.5*(229.36 - 250) = 219.03$. She would have the same expectation for period 3, that $P_3^2 = P_2^2 = 219.03$. Therefore, the price that a young agent would be willing to pay in period 2, $P_2 = 219.03/1.09 = 200.95$.

In period 3, an agent's expectation for the upcoming period $P_3^3 = P_2 + \beta_1(P_2 - P_1) + \beta_2(P_2 - P_2^2) = 200.95 + 0.5*(200.95 - 229.36) + 0.5*(200.95 - 219.03) = 200.95 - 14.21 - 9.04 = 177.7$. If agents hold the same expectation for period 4, so that $P_4^3 = P_3^3$, then $P_3 = 177.7/1.09 = 163.03$.¹⁴ Since $P_3 < P_2 < P_3$, the example illustrates how a negative trend in the valuation of a risky asset can arise, even in an environment with a positive riskfree interest rate, when the assumption of rational expectations is replaced by those beliefs that are held by our participant pool.

In Appendix B, we provide additional examples that illustrate that (a) the price decline may be steeper when the interest rate is higher, (b) the same general pattern of decreasing asset prices under a positive riskless rate can be observed if $\beta_1, \beta_2 < 0$, and (c) a constant asset price trajectory is also possible.

7 Discussion

In this paper, we have reported on an experiment to study the effects of monetary policy on the size of an asset price bubble. The asset has no intrinsic value, and can only be held for the purpose of speculation. It is widely thought, within the central banking community, that a tight interest rate policy can be used to mitigate asset market bubbles, while an accomodative policy can help alleviate asset price declines. However, as pointed out by Galí (2014), higher interest rates have the effect of raising the return on all assets, so that an asset in a bubble can experience a more rapid run-up in price, the higher is the interest rate. Indeed, this is an unavoidable consequence of assuming rational expectations.

¹⁴Note that the downward trend remains, though is less pronounced, if there is no trend-following component in agents's expectations. Suppose that $\beta_2 = 0$. Then $P_3^3 = 186.75$ and $P_3 = 171.33$.

Our experimental data show the way to a reconciliation of the two viewpoints. We find that there is an immediate, short-term decrease in the price of the risky asset following an interest rate hike and similar increases following an interest rate cut. This effect appears to be behavioral in origin and is not anticipated by received theoretical models,¹⁵ which allow for the effect only when the asset pays dividends in the future. The response to the interest rate change occurs presumably because the return of the alternative, safe asset is directly affected by the policy, resulting in changes in demand for the bubble asset. The increase in return in the bond resulting from a rate hike draws investment out of the risky asset, and the decrease in return induced by a rate cut has the opposite effect. We also observe that the effect predicted by Galí (2014) appears strongly in the data. When a higher interest rate is in effect, the bubble asset appreciates more than under a lower interest rate. Thus, the typical effect of a rate hike is to decrease a bubble in the current period, and to magnify the bubble in the next period. It appears to us that a central bank policy of leaning against the wind must have very frequent adjustments, before the second effect can appear, in order to be effective in moderating asset prices.

We also conducted two treatments, HIGH and LOW, in which the interest rates were constant. Under Rational Expectations, the bubble asset would appreciate at a greater rate in the HIGH than in the LOW treatment, since the interest rate is greater under HIGH. In our data, we observe that the two assets exhibit similar average price changes. The changes average close to zero, but are not significantly different from the riskfree rate in the LOW treatment. It may be the case that individuals make better predictions and market decisions when interest rates are low, since participants typically only have experience with rates in the 0 - 5 percent range outside of the laboratory. In Section 6, we have argued that such a pattern can arise if individuals have adaptive or trend-following, rather than rational, expectations. The experiment demonstrates a clear pattern of responses to interest rate policy among inexperienced market participants under relatively low stakes. Follow-up work can investigate whether similar results are to be observed in markets populated with sophisticated investors and markets with high stakes.

References

- Aliprantis, C. D. & Plott, C. R. (1992), ‘Competitive equilibria in overlapping generations experiments’, *Economic Theory* **2**(3), 389–426.
- Anufriev, M. & Hommes, C. (2012), ‘Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments’, *American Economic Journal: Microeconomics* **4**(4), 35–64.

¹⁵See Miao et al. (2019) for an environment in which that outcome may arise under certain conditions.

- Assenza, T., Heemeijer, P., Hommes, C. H. & Massaro, D. (2013), ‘Individual expectations and aggregate macro behavior’, *Tinbergen Institute Discussion Paper*.
- Bao, T. & Zong, J. (2019), ‘The impact of interest rate policy on individual expectations and asset bubbles in experimental markets’, *Journal of Economic Dynamics and Control* **107**, 103735.
- Bernasconi, M. & Kirchkamp, O. (2000), ‘Why do monetary policies matter? An experimental study of saving and inflation in an overlapping generations model’, *Journal of Monetary Economics* **46**(2), 315–343.
- Bostian, A. J. A. & Holt, C. A. (2009), ‘Price bubbles with discounting: A Web-based classroom experiment’, *The Journal of Economic Education* **40**(1), 27–37.
- Camerer, C. F. & Weigelt, K. (1993), ‘Convergence in experimental double auctions for stochastically lived assets’, *The Double Auction Market: Institutions, Theories, and Evidence. Proceedings Volume in the Santa Fe Institute Studies in the Sciences of Complexity* **14**, 355–396.
- Carbone, E., Hey, J. & Neugebauer, T. (2021), ‘An experimental comparison of two exchange economies: Long-lived asset vs. short-lived asset’, *Management Science* **forthcoming**.
- Carlé, T. A., Lahav, Y., Neugebauer, T. & Noussair, C. N. (2019), ‘Heterogeneity of beliefs and trade in experimental asset markets’, *Journal of Financial and Quantitative Analysis* **54**(1), 215–245.
- Crockett, S., Duffy, J. & Izhakian, Y. (2019), ‘An experimental test of the Lucas asset pricing model’, *The Review of Economic Studies* **86**(2), 627–667.
- Duffy, J. (2016), ‘Macroeconomics: a survey of laboratory research’, *Handbook of Experimental Economics* **2**, 1–90.
- Dufwenberg, M., Lindqvist, T. & Moore, E. (2005), ‘Bubbles and experience: An experiment’, *American Economic Review* **95**(5), 1731–1737.
- Fenig, G., Mileva, M. & Petersen, L. (2018), ‘Deflating asset price bubbles with leverage constraints and monetary policy’, *Journal of Economic Behavior & Organization* **155**, 1–27.
- Fischbacher, U. (2007), ‘z-Tree: Zurich toolbox for ready-made economic experiments’, *Experimental Economics* **10**(2), 171–178.
- Fischbacher, U., Hens, T. & Zeisberger, S. (2013), ‘The impact of monetary policy on stock market bubbles and trading behavior: Evidence from the lab’, *Journal of Economic Dynamics and Control* **37**(10), 2104–2122.
- Galí, J. (2014), ‘Monetary policy and rational asset price bubbles’, *American Economic Review* **104**(3), 721–752.

- Galí, J. & Gambetti, L. (2015), ‘The effects of monetary policy on stock market bubbles: Some evidence’, *American Economic Journal: Macroeconomics* **7**(1), 233.
- Giusti, G., Jiang, J. H. & Xu, Y. (2016), ‘Interest on cash, fundamental value process and bubble formation: An experimental study’, *Journal of Behavioral and Experimental Finance* **11**, 44–51.
- Haruvy, E., Lahav, Y. & Noussair, C. N. (2007), ‘Traders’ expectations in asset markets: experimental evidence’, *American Economic Review* **97**(5), 1901–1920.
- Haruvy, E. & Noussair, C. N. (2006), ‘The effect of short selling on bubbles and crashes in experimental spot asset markets’, *The Journal of Finance* **61**(3), 1119–1157.
- Hennequin, M. & Hommes, C. (2019), ‘Managing bubbles in experimental asset markets with monetary policy’, *working paper, University of Amsterdam*.
- Hommes, C., Sonnemans, J., Tuinstra, J. & Van de Velden, H. (2008), ‘Expectations and bubbles in asset pricing experiments’, *Journal of Economic Behavior & Organization* **67**(1), 116–133.
- Marimon, R. & Sunder, S. (1993), ‘Indeterminacy of equilibria in a hyperinflationary world: experimental evidence’, *Econometrica* **61**(5), 1073–1107.
- Marimon, R. & Sunder, S. (1994), ‘Expectations and learning under alternative monetary regimes: an experimental approach’, *Economic Theory* **4**(1), 131–162.
- Miao, J., Shen, Z. & Wang, P. (2019), ‘Monetary policy and rational asset price bubbles: Comment’, *American Economic Review* **109**(5), 1969–1990.
- Noussair, C. & Tucker, S. (2006), ‘Futures markets and bubble formation in experimental asset markets’, *Pacific Economic Review* **11**(2), 167–184.
- Nuzzo, S. & Morone, A. (2017), ‘Asset markets in the lab: A literature review’, *Journal of Behavioral and Experimental Finance* **13**, 42–50.
- Palan, S. (2013), ‘A Review of bubbles and crashes in experimental asset markets’, *Journal of Economic Surveys* **27**(3), 570–588.
- Pfajfar, D. & Žakelj, B. (2018), ‘Inflation expectations and monetary policy design: Evidence from the laboratory’, *Macroeconomic Dynamics* **22**(4), 1035–1075.
- Porter, D. P. & Smith, V. L. (1995), ‘Futures contracting and dividend uncertainty in experimental asset markets’, *Journal of Business* **68**(4), 509–541.

- Powell, O. & Shestakova, N. (2016), ‘Experimental asset markets: A survey of recent developments’, *Journal of Behavioral and Experimental Finance* **12**, 14–22.
- Roth, A. E. & Murnighan, J. K. (1978), ‘Equilibrium behavior and repeated play of the prisoner’s dilemma’, *Journal of Mathematical Psychology* **17**(2), 189–198.
- Smith, V. L. (1962), ‘An experimental study of competitive market behavior’, *Journal of Political Economy* **70**(2), 111–137.
- Smith, V. L., Suchanek, G. L. & Williams, A. W. (1988), ‘Bubbles, crashes, and endogenous expectations in experimental spot asset markets’, *Econometrica* **56**(5), 1119–1151.

We include three appendices. The first, Appendix A, contains a number of additional regressions that investigate the determinants of trader expectations. Appendix B consists a number of additional worked out examples to supplement the analysis in subsection 6.3. Appendix C provides an English translation of the instructions to participants, which were originally in Spanish.

A Additional Regression Results

	(1)	(2)	(3)	(4)	(5)	(6)
	Trend Follow	Adaptive Exp.	Hybrid	Trend Follow	Adaptive Exp.	Hybrid
$P_t - P_{t-1}$	-0.37*** (0.02)		-0.26*** (0.03)	-0.36*** (0.02)		-0.29*** (0.03)
$P_t - P_t^e$		0.64*** (0.02)	0.83*** (0.03)		0.68*** (0.02)	0.90*** (0.03)
r_t				-187.26*** (10.86)	-149.54*** (10.82)	-173.54*** (10.90)
<i>Cons</i>	0.89 (1.74)	0.15 (1.63)	0.51 (1.76)	16.34*** (1.82)	12.58*** (1.72)	14.99*** (1.83)
N	2323	2323	2323	2323	2323	2323
R^2	0.35	0.51	0.54	0.43	0.54	0.59
Market FE	✓	✓	✓	✓	✓	✓

Table A1: OLS estimation for Backward-Looking Expectational Rules, including Market Fixed Effects

Coefficients reported and robust standard errors in parentheses. *** significant at 1%, **significant at 5%, * significant at 10%

	(1)	(2)	(3)	(4)	(5)	(6)
	Trend Follow	Adaptive Exp.	Hybrid	Trend Follow	Adaptive Exp.	Hybrid
$P_t - P_{t-1}$	-0.35*** (0.02)		-0.17*** (0.03)	-0.35*** (0.02)		-0.24*** (0.03)
$P_t - P_t^e$		0.62*** (0.02)	0.74*** (0.03)		0.66*** (0.02)	0.84*** (0.03)
r_t				-216.53*** (9.12)	-162.70*** (8.62)	-190.49*** (9.22)
<i>Cons</i>	-16.89*** (3.10)	-11.40*** (2.50)	-13.38*** (2.74)	5.28** (2.70)	4.74** (2.32)	4.75* (2.55)
N	2323	2323	2323	2323	2323	2323
R^2	0.25	0.47	0.49	0.37	0.52	0.55
Individual FE	✓	✓	✓	✓	✓	✓

Table A2: OLS estimation of Backward-Looking Expectational Rules, including Individual Fixed Effects

Coefficients reported and robust standard errors in parentheses. *** significant at 1%, **significant at 5%, * significant at 10%

	Trend Following	Adaptive Exp.	Hybrid	Trend Following	Adaptive Exp.	Hybrid
$P_t - P_{t-1}$	-0.15*** (0.04)		0.25*** (0.04)	-0.17*** (0.03)		0.17*** (0.05)
$P_t - P_t^e$		0.69*** (0.03)	0.54*** (0.04)		0.73*** (0.03)	0.62*** (0.04)
r_t				-195.29*** (9.83)	-147.63*** (8.81)	-124.35*** (10.62)
<i>Cons</i>	-16.08*** (0.74)	-11.84*** (0.63)	-9.58** (0.71)	5.79*** (0.80)	4.20*** (0.82)	3.19*** (0.86)
N	1567	1567	1567	1567	1567	1567
R^2	0.04	0.56	0.59	0.17	0.60	0.61

Table A3: OLS estimation for Backward-Looking Expectational Rules - LOW and HIGH treatments only

Coefficients reported and robust standard errors in parentheses. *** significant at 1%, **significant at 5%, * significant at 10%

	Trend Following	Adaptive Exp.	Hybrid	Trend Following	Adaptive Exp.	Hybrid
$P_t - P_{t-1}$	-0.543*** (0.03)		-0.52*** (0.03)	-0.52*** (0.02)		-0.51*** (0.03)
$P_t - P_t^e$		0.46*** (0.03)	0.95*** (0.04)		0.49*** (0.04)	0.97*** (0.04)
r_t				-146.54*** (23.77)	-190.63** (30.16)	-144.16*** (23.63)
<i>Cons</i>	-0.82 (0.86)	0.52 (1.12)	-0.69 (0.86)	9.67*** (1.54)	14.08*** (1.85)	9.59*** (1.53)
N	756	756	756	756	756	756
R^2	0.60	0.27	0.56	0.62	0.30	0.58

Table A4: OLS estimation for Backward-Looking Expectational Rules - LAW Treatment only

Coefficients reported and robust standard errors in parentheses. *** significant at 1%, **significant at 5%, * significant at 10%

B Additional examples of price trajectories to accompany section 6.3

Example B1: This example demonstrates, when contrasted with the example provided in Section 6.3, that the price decline is steeper, the higher the interest rate. Consider the behavior of a similar market when the interest rate $r = .03$ instead of $r = .09$ as in the example described in Section 6.3. Then, $P_1^1 = P_2^1 = P_3^1 = 250$, and $P_1 = 250/1.03 = 242.72$. Consequently, $P_2^2 = 239.08$. and $P_2 = 232.12$. In period 3, $P_3^3 = 232.12 - 5.3 - 3.48 = 223.34$ and $P_3 = 216.83$. The downward trend is more modest than under the higher interest rate.

Example B2: This example shows that the same general pattern of decreasing asset prices over time is observed if $\beta_1, \beta_2 < 0$. Assume that $P_1^1 = P_2^1 = P_3^1 = 250$, $r = 0.09$, and agents have the expectations given in (13), but with $\beta_1 = \beta_2 = -0.5$. The β coefficients are negative and initial expectations are of a flat trajectory. In period 1, the price would equal $P_1 = \frac{P_2^1}{1+r} = 250/1.09 = 229.36$, and $P_2^2 = P_1 + \beta_2(P_1 - P_1^1) = 229.36 - 0.5*(229.36 - 250) = 239.68$. Since $P_3^3 = P_2^2 = 239.68$, $P_2 = 239.68/1.09 = 219.89$. In period 3, $P_3^3 = P_2 + \beta_1(P_2 - P_1) + \beta_2(P_2 - P_2^2) = 219.89 - 0.5*(219.89 - 229.36) - 0.5*(219.89 - 239.68) = 219.89 + 4.73 + 9.90 = 233.91$. If $P_4^3 = P_3^3$, then $P_3 = 233.91/1.09 = 214.59$. Prices decline over time.

Example B3: This example illustrates how, with appropriate long-term expectations, the price trend can be flat regardless of the signs of the coefficients β_1 and β_2 . Suppose, for example, that short term expectations for the upcoming

period t are those in (13), while the price is expected to appreciate by the rate of interest in period $t+1$. That is, assume that $P_1^1 = 250$, $P_2^1 = (1+r)*P_1^1$, and $P_3^1 = (1+r)*P_2^1$, and $r = 0.09$. Then, we have that $P_1 = P_2^1/(1+r) = 250$. As a consequence, $P_2^2 = P_1 + \beta_2(P_1 - P_1^1) = 250 + 0.5*(250 - 250) = 250$, and $P_3^2 = P_2^2 * (1+r) = 250*1.09 = 272.5$. Then the price in period 2, $P_2 = P_3^2/(1+r) = 250$, and $P_3^3 = P_2 + \beta_1(P_2 - P_1) + \beta_2(P_2 - P_2^2) = 250 + 0.5*(250 - 250) + 0.5*(250 - 250) = 250$, and so on. The price stays constant at 250 despite the positive interest rate. The prices would remain constant for any values of β_1 and β_2 under the assumed expectations.

C Experimental Instructions

In this Appendix, we provide the instructions for the main experiment, described in Section 3, followed by those in the sessions in which participants made price predictions, described in Section 5.