

Rational and Irrational Bubbles: An Experiment¹

Sophie Moinas and Sebastien Pouget ²

July 23, 2010

¹We would like to thank Andrea Attar, Bruno Biais, Snehal Banerjee, Christophe Bisière, Georgy Chabakauri, John Conlon, Tony Doblal-Madrid, Xavier Gabaix, Alex Guembel, Guo Kai, Charles Noussair, Alessandro Pavan, Jean Tirole, Reinhard Selten, Paul Woolley, and especially James Dow and Thomas Mariotti, as well as seminar participants in Bonn University, Luxembourg University, Lyon University (GATE), Toulouse University, Paris Dauphine University, the Paris School of Economics, the London Business School, the Midwest Macroeconomics Meetings, the CSIO-IDEI workshop in Northwestern University for helpful comments. This research was conducted within and supported by the Paul Woolley Research Initiative on Capital Market Dysfunctionalities at IDEI-R, Toulouse.

²University of Toulouse (IAE and Toulouse School of Economics), Place Anatole France, 31000 Toulouse, France, sophie.moinas@univ-tlse1.fr and sebastien.pouget@univ-tlse1.fr

Abstract

Rational bubbles form when traders buy an overvalued asset, rationally expecting to resell at a higher price. We design a novel experimental setting with finite trading opportunities that precludes backward induction from unraveling this process. This is the case if no player is ever sure to be last in a trading sequence. Otherwise, only irrational bubbles can form. In our experiment, both rational and irrational bubbles are observed, and are not eliminated by replication. Players speculate more when more steps of reasoning are needed to conclude their followers will be last in the trading sequence. Maximum likelihood estimations show that the observed behaviors are consistent with boundedly rational models of choice.

Keywords : Rational bubbles, irrational bubbles, experiment, cognitive hierarchy model, quantal response equilibrium.

1 Introduction

This paper presents an experimental investigation of rational and irrational bubbles in asset markets. Historical and recent economic developments such as the South Sea, Mississippi, and dot com price run-up episodes suggest that financial markets are prone to bubbles and crashes. However, to the extent that fundamental values cannot be directly observed in the field, it is very difficult to empirically demonstrate that these episodes actually correspond to mispricings.¹

To overcome this difficulty and study bubble phenomena, economists have relied on the experimental methodology: in the laboratory, fundamental values are induced by the researchers who can then compare them to asset prices. Starting with Smith, Suchanek and Williams (1988), many researchers document the existence of irrational bubbles in experimental financial markets. These bubbles are irrational in the sense that they would be ruled out by backward induction. The design created by Smith, Suchanek and Williams (1988) features a double auction market for an asset that pays random dividends in several successive periods. The subsequent literature shows that irrational bubbles also tend to arise in call markets (Van Boening, Williams, LaMaster, 1993), with a constant fundamental value (Noussair, Robin, Ruffieux, 2001), with lottery-like assets (Ackert, Charupat, Deaves, and Kluger, 2006), and tend to disappear when some traders are experienced (Dufwenberg, Lindqvist, and Moore, 2005), when there are futures markets (Porter and Smith, 1995), and when short-sells are allowed (Ackert, Charupat, Church, and Deaves, 2005). Lei, Noussair and Plott (2001) further show that, even when they cannot resell and realize capital gains, some participants still buy the asset at a price which exceeds the sum of the expected dividends, a behavior consistent with risk-loving preferences or judgmental errors.

Extending the experimental analysis to rational bubbles raises two difficulties. On the one hand, existing theories of rational bubbles rely on infinite trading opportunities.² This is a challenge because it is not feasible to organize an experiment in which there are infinite trading opportunities. Indeed,

¹In this paper, we define the fundamental value of an asset as the price at which agents would be ready to buy the asset given that they cannot resell it later. See Camerer (1989) and Brunnermeier (2009) for surveys on bubbles.

²Infinite trading opportunities may derive from an infinite horizon models (see, for example, Tirole (1985) for deterministic bubbles, Blanchard (1979) and Weil (1987) for stochastic bubbles, Abreu and Brunnermeier (2004) and Doblus-Madrid (2010) for clock games), or via continuous trading models (see Allen and Gorton, 1993).

an experiment cannot last for ever (even with a small probability) nor can it offer an infinite trading speed. On the other hand, rational bubbles generate potentially infinite gains and losses for market participants because of the price explosiveness. This is again a challenge because it is not acceptable to impose potentially huge losses neither on experimental subjects nor on the experimenter.³

The contribution of this paper is to propose an experimental design to analyze rational bubbles, while also allowing for the study of irrational bubbles. Another advantage of our experimental design is to allow for the study, at the individual level, of three different types of speculative behavior that play a role in bubble formation: irrational speculation (due to mistakes or erroneous beliefs), speculation on others' irrationality (betting on the fact that others are going to do mistakes), and speculation on others' rationality (betting on the fact that others are going to act rationally).

To circumvent the two difficulties discussed above, we propose an economic setting with finite trading opportunities and symmetric information on the asset payoff in which bubbles can be rational. This setting features a sequential market for an asset which is worth nothing (and this is common knowledge).⁴ We show that bubbles can be rational if no trader is ever sure to be last in the market sequence.⁵ This is the case when the price proposed to the first trader in the market sequence is random and potentially infinite. We further prove that a bubble can be rational even if traders are endowed with limited liability and are financed by outside financiers. We design the game such that the potential infinity in gains and losses is concentrated in the hands of outside financiers. In order to avoid potentially infinite losses for experimental subjects (and for the experimenter), outside financiers are not part of the experiment.

To study irrational bubbles, we introduce a price cap. In this case, upon

³The theoretical analyses of Allen, Morris, and Postlewaite (1993), and Conlon (2004) show that rational bubbles can occur with a finite number of trading opportunities and without exposing participants to potentially infinite losses. These analyses however involve asymmetric information regarding the asset cash flows. In order to be in line with the literature on experimental bubbles, we design an experiment in which there is no asymmetric information on the asset payoff. In this setting, asset prices as well as potential gains and losses have to grow without bounds for a bubble to be sustained at equilibrium (see Tirole, 1982).

⁴Our model of financial market can be viewed as a generalization of a centipede game in which players do not know at which node of the game they play.

⁵This idea is in line with Allen and Gorton (1993) and Abreu and Brunnermeier (2004). We complement these analyses by theoretically showing that continuous trading or infinite horizon are not required for a bubble to emerge at equilibrium.

receiving the highest potential price, a trader realizes that he or she is last in the market sequence and, if rational, refuses to buy. Even if not sure to be last in the market sequence, the previous trader, if rational, also refuses to buy because he anticipates that the next trader will know he is last and will refuse to trade. This backward induction argument rules out the existence of bubbles when there is a price cap, if all traders are rational and this is common knowledge. Observing irrational bubbles would mean that one of these assumptions does not hold. By increasing the level of the cap, one increases the number of steps of reasoning needed for participants to realize that the agent who receives the highest price refuses to trade. As a result, varying the level of the cap allows the experimenter to understand how bounded rationality or lack of higher-order knowledge of rationality affect bubble formation.

Our experimental protocol is as follows. There are three traders on the market. The experimenter proposes a price to a trader, who can accept or refuse to buy at this price. If he buys, he will resell at a higher price if the next trader accepts to buy. The price proposed to the first trader in the market sequence is determined following a geometric distribution of parameter $1/2$ defined on the set of positive powers of 10. This first price can thus be 1 with probability $1/2$, 10 with probability $1/4$, 100 with probability $1/8$,... After each transaction, the price is multiplied by 10. As a result, upon being proposed to buy at a price of 100, for example, a trader has a conditional probability $1/7$ to be first in the market sequence, $2/7$ to be second, and $4/7$ to be last in the market sequence. The baseline experiment features different treatments. In one treatment, there is no cap on the first price and there thus exists a rational bubble. In other treatments, there is a cap K on the first price, with $K = 10, 100, 1000$, or 1. In these alternative treatments, there is no rational bubble.⁶ Subjects participate in only one treatment and in a one-shot game. In a second experiment, the same four treatments are used but the game is now repeated five times. In a third experiment, a one-shot game experiment with $K = 10, 100$ is organized with executive MBA students.

Our experimental results are as follows. First, when there is a cap on the first price, bubbles frequently appear. This result suggests that in our simple design, we are able to replicate the results found in the previous experimental literature that irrational bubbles can be observed. Second, we

⁶When there is a cap on the first price, this price is also determined randomly with the truncated probability being added to the highest possible realization of the price. For example, in the case of a cap at 100, the first price equals 1 with probability $1/2$, and 10 or 100 with the same probability $1/4$.

find that, when there is no price cap, bubbles are also observed but not always. This may be due to strategic uncertainty, because of the existence of two equilibria, with or without a bubble. Third, a regression analysis shows that the propensity for a subject to enter a bubble is negatively related to the conditional probability to be last in the market sequence, and to risk aversion. These results suggest that subjects' decision to enter bubbles demonstrate some level of rationality. The analysis also shows that the probability that subjects enter a bubble increases with the number of reasoning steps needed to realize that a participant might know that he or she is last. We refer to this phenomenon as a snow-ball effect. Our results are robust in the sense that they also hold when the game is repeated five times or when it is played by experienced business professionals. Experience with the game or in business reduces but does not eliminate the propensity to enter into bubbles.

To better understand subjects' individual behavior, we estimate two models of bounded rationality that depart from the Nash equilibrium in two different aspects: the cognitive hierarchy model (hereafter CH) developed by Camerer, Ho and Chong (2004), and the quantal response equilibrium (hereafter QR) of Mc Kelvey and Palfrey (1995). They are able to account for the snowball effect that is observed in irrational bubbles. At the same time, they capture well the fact that not all traders are willing to enter into rational bubbles.

The CH model states that agents best-respond to mutually inconsistent beliefs. Maximum likelihood estimation indicates an average level of sophistication equal to 0.67 which is in line with the estimations of Camerer, Ho and Chong (2004) on other types of games. In the CH model, bubbles form even even if there is a price cap because agents tend to overestimate the probability that the next trader buys the asset: unsophisticated traders make mistakes with some probability, which induces more sophisticated traders to buy the asset. The CH model therefore features a snowball effect, which explains why it fits the data better than the Nash equilibrium.

The QR equilibrium also takes into account the fact that players make mistakes but it retains beliefs' consistency. At equilibrium, players are responsive to payoffs to the extent that more profitable actions are chosen more often. Maximum likelihood estimation indicates an average responsiveness of 2.54 in line with previous estimations of Mc Kelvey and Palfrey (1995) on other types of games. In the QR equilibrium, bubbles form even if there is a price cap because, if some agents mistakenly buy the asset, it becomes less costly for others to also buy: the agent who is sure to be last enters the bubble with a small probability thereby increasing the incentives

of the previous traders to enter. The QR model also features a snowball effect, and fits the data slightly better than the CH model. This is because it captures the increase in the expected payoff from being at the beginning of the market process when the traders are sure not to be last.

Our experimental analysis is related to Brunnermeier and Morgan (2010) who study clock games both from a theoretical and an experimental standpoint. These clock games can indeed be viewed as metaphors of “bubble fighting” by speculators, gradually and privately informed of the fact that an asset is overvalued. Speculators do not know if others are already aware of the bubble. They have to decide when to sell the asset knowing that such a move is profitable only if enough speculators have also decided to sell. Their experimental investigation and ours share two common features. First, the potential payoffs are exogenously fixed, that is, there is a pre-determined price path. Second, there is a lack of common knowledge over a fundamental variable of the environment. In Brunnermeier and Morgan (2006), the existence of a bubble is not common knowledge. In our setting, the existence of the bubble is common knowledge but traders’ position in the market sequence is not. There are several differences between our approach and theirs. A first difference is the time dimension. The theoretical results tested by Brunnermeier and Morgan (2006) depend on the existence of an infinite time horizon. They implement this feature in the laboratory by randomly determining the end of the session. On the contrary, we design an economic setting in which there could be bubbles in finite time with finite trading opportunities, even if traders act rationally. A second difference is that we study the formation of irrational bubbles. In particular, we analyze the various types of speculation that can arise in irrational bubbles. A third difference is that we rationalize the formation of rational and irrational bubbles by showing that bounded rationality models can explain observed behavior.

The rest of the paper is organized as follows. Section 2 presents a theory of rational bubbles with finite trading opportunities and symmetric information on the asset’s cash flow. Section 3 presents an experimental design based on this theory. Results are in Section 4. Section 5 concludes and provides potential extensions.

2 A theory of rational bubbles

The objective of this section is to show that bubbles can emerge in a financial market with perfectly rational traders and finite trading opportunities,

without asymmetric information on asset payoffs. Consider a financial market in which trading proceeds sequentially. There are T agents, referred to as traders. Traders' position in the market sequence is random with each potential ordering being equally likely. Traders can trade an asset that generates no cash flow and this is common knowledge.⁷ This enables us to unambiguously define the fundamental value of the asset: it is zero in our case because, if the asset cannot be resold, an agent would not pay more than zero to buy it.

The asset is issued by agent 0, referred to as the issuer.⁸ The first trader in the sequence is offered to buy the asset at a price P_1 . If he does so, he proposes to resell at price P_2 to the second trader. More generally, the t -th trader in the sequence, $t \in \{1, \dots, T - 1\}$, is offered to buy the asset at price P_t and resell at price P_{t+1} to the $t + 1$ -th trader. Traders take the price path as given, with $P_t > 0$ for $t \in \{1, \dots, T\}$. Finally, the last trader in the sequence is offered to buy the asset at price P_T but cannot resell it. If the t -th trader buys the asset and is able to resell it, his payoff is $P_{t+1} - P_t$. If he is unable to resell the asset, his payoff is $-P_t$. For simplicity, we consider that if a trader refuses to buy the asset, the market process stops.

We consider that traders are risk neutral. We show in Appendix A that our results hold with risk averse traders. Individual i has an initial wealth denoted by W_i , $i \in \{1, \dots, T\}$.⁹ As a benchmark, consider the case in which traders have perfect information, that is, each trader i knows that his position in the sequence is t and this is common knowledge. In this perfect information benchmark, it is straightforward to show that no trader will accept to buy the asset except at a price of 0 which corresponds to the fundamental value of the asset. Indeed, the last trader in the queue, if he buys, ends up with $W_T - P_T$ which is lower than W_T . Since he knows that he is the last trader in the queue, he prefers not to trade. By backward induction, this translates into a no-bubble equilibrium. This result is summarized in the next proposition.

Proposition 1 *When traders know their position in the market sequence, the unique perfect Nash equilibrium involves no trade.*

Let's now consider what happens when traders do not initially know their position in the market sequence, and this is common knowledge. We

⁷The asset cash flow could be positive and risky without changing our results.

⁸The potential bubbles that may arise in our environment can be interpreted as Ponzy schemes, and the issuer of the asset as the scheme organizer.

⁹In our model, traders might end up with negative wealth.

model this situation as a Bayesian game. The set of players is $\{1, \dots, T\}$. The set of states of the world is Ω which includes the $T!$ potential orderings. ω refers to a particular ordering. The set of actions is identical for each player i and each position t and is denoted by $A = \{B, \emptyset\}$ in which B stands for buy and \emptyset for refusal to buy. Denote by $\omega_t^i \subset \Omega$ the set of orderings in which trader i 's position in the market sequence is t . The set of signals that may be observed by player i is the set of potential prices denoted by P . The signal function of player i is $\tau(i) : \omega_t^i \rightarrow P_t$, in which P_t refers to the price that is proposed to the t -th trader in the market sequence. The price path P_t is defined as follows. The price P_1 proposed to the first trader in the sequence is random and is distributed according to the probability distribution $g(\cdot)$ on P .¹⁰ Other prices are determined as $P_{t+1} = f(P_t)$, with $f(\cdot) : P \rightarrow P$ being a strictly increasing function that controls for the explosiveness of the price path. A strategy for player i is a mapping $S_i : P \rightarrow A$ in which $S_i(p)$ indicates what action is chosen by player i after observing a price p . Conditional on observing $p = P_t$, player i understands that the next player j in the market sequence observes $f(P_t)$, and that he chooses $S_j(f(P_t))$. Using the signal function, players may learn about their position in the market sequence. A strategy profile $\{S_1^*, \dots, S_T^*\}$ is a Bayesian Nash equilibrium if the following individual rationality (IR) conditions are satisfied:

$$\mathbb{E} [\pi [S_i^*(P_t), S_j^*(f(P_t))] | P_t] \geq \mathbb{E} [\pi [S_i(P_t), S_j^*(f(P_t))] | P_t],$$

$$\forall (i, j) \in \{1, \dots, T\} \times \{1, \dots, T\} \text{ with } j \neq i, \text{ and } \forall P_t \in P.$$

$\pi [S_i(P_t), S_j^*(f(P_t))]$ represents the payoff received by the risk-neutral player i given that he chooses action $S_i(P_t)$ and that other players choose actions $S_j^*(f(P_t))$. Remark that agents' payoff not only depend on others' actions but also on the state of nature because it is possible that they are last in the market sequence.

We now study under what conditions there exists a bubble equilibrium $\{S_1^* = B, \dots, S_T^* = B\}$. The crucial parameter a player i has to worry about in order to decide whether to enter a bubble is the conditional probability to be last in the market sequence, $\mathbb{P}(\omega \in \omega_T^i | P_t)$. The IR condition can be rewritten as:

$$(1 - \mathbb{P}[\omega \in \omega_T^i | P_t]) \times (W_i + f(P_t) - P_t) + \mathbb{P}[\omega \in \omega_T^i | P_t] \times (W_i - P_t) \geq W_i,$$

¹⁰One can consider that this first price P_1 is chosen by Nature or by the issuer according to a mixed strategy characterized by $g(\cdot)$.

$$\begin{aligned}
& \forall i \in \{1, \dots, T\}, \text{ and } \forall P_t \in P \\
& \Leftrightarrow (1 - \mathbb{P}[\omega \in \omega_T^i | P_t]) \times f(P_t) \geq P_t, \\
& \forall i \in \{1, \dots, T\}, \text{ and } \forall P_t \in P.
\end{aligned}$$

If $\mathbb{P}[\omega \in \omega_T^i | P_t] = 1$ for some i and some P_t , the IR condition is not satisfied and the bubble equilibrium does not exist. This is for example the case when the support of the distribution $g(\cdot)$ is bounded above by a threshold K . Indeed, a trader upon observing $P_t = f^{T-1}(K)$ knows that he is last and refuses to trade. Backward induction then prevents the existence of the bubble equilibrium. The IR function is also not satisfied if the signal function $\tau(i)$ is injective. Indeed, by inverting the signal function, players, including the one who is last in the sequence, learn what their position is. These results are summarized in the following proposition.

Proposition 2 *The no-bubble equilibrium is the unique Bayesian Nash equilibrium if i) the signal function is injective, ii) the first price is randomly distributed on a support that is bounded above, iii) the price path is not explosive enough, or iv) the probability to be last in the market sequence is too high.*

We now propose an environment where the IR condition derived above is satisfied. Consider that the set of potential prices is defined as $P = \{m^n \text{ for } m > 1 \text{ and } n \in \mathbb{N}\}$, that is, prices are positive powers of a constant $m > 1$. Also, assume that $g(P_1 = m^n) = (1 - q)q^n$, that is, the power n follows a geometric distribution of parameter $q \in (0, 1)$. Finally, we set $f(P_t) = m \times P_t$. If there are T players on the market, the probability that a player i is last in the sequence, conditional on the price P_t that he is proposed, is computed by Bayes' rule:

$$\begin{aligned}
\mathbb{P}[\omega \in \omega_T^i | P_t = m^n] &= \frac{\mathbb{P}[P_t = m^n | \omega \in \omega_T^i] \times \mathbb{P}[\omega \in \omega_T^i]}{\mathbb{P}[P_t = m^n]} \\
&= \frac{(1 - q)q^{n-(T-1)} \times \frac{1}{T}}{\sum_{j=n-(T-1)}^{j=n} (1 - q)q^j \times \frac{1}{T}} = \frac{1 - q}{1 - q^T} \text{ if } n \geq T - 1,
\end{aligned}$$

and

$$\mathbb{P}[\omega \in \omega_T^i | P_t = m^n] = 0 \text{ if } n < T - 1.$$

Under our assumptions, Bayes' rule implies that the conditional probability to be last in the market sequence is 0 if the proposed price is strictly smaller than m^{T-1} , and $\frac{1-q}{1-q^T}$ if the proposed price is equal to or higher than m^{T-1} .

This conditional probability thus does not depend on the level of the price that is proposed to the players.¹¹ The IR condition can be rewritten:

$$\left(\frac{q - q^T}{1 - q^T}\right) \times m \geq 1.$$

This condition is less restrictive when there are more traders present on the market.

There thus exists an infinity of price paths characterized by $m \geq \frac{1-q^T}{q-q^T}$ that sustain the existence of a bubble equilibrium. Obviously, there always exists a no-bubble equilibrium.¹² Indeed, if players anticipate that other players do not enter the bubble, then they are better off refusing to trade. These results are summarized in the next proposition.

Proposition 3 *If i) the T traders are equally likely to be last in the market sequence, ii) the price P_1 proposed to the first trader in the sequence is randomly chosen in powers of m according to a geometric distribution with parameter q , and iii) $P_t = m \times P_{t-1}$, $\forall t \in \{2, \dots, T\}$, there exists a bubble Bayesian Nash equilibrium if and only if $m \geq \frac{1-q^T}{q-q^T}$. There always exists a no-bubble equilibrium.*

Our results hold even if one introduces randomness in the underlying asset payoff, and (potentially random) payments at interim dates. In Appendix A, we show that our results hold if traders are risk averse. One could be tempted to interpret our results as an inverse-Hirshleifer effect: going from perfect to imperfect information seems to imply a creation of gains from trade in our setting even with risk-neutral agents. However, note that it is not possible to compute the ex-ante welfare created by the game of imperfect information. Indeed, the expected payoffs of the players are infinite. To see this, remark that these expected payoffs are equal to:

$$\lim_{x \rightarrow +\infty} \left[\frac{m-1}{2} + \frac{m(m-1)}{4} + \left(\frac{q-q^T}{1-q^T} (m-1) - \frac{1-q}{1-q^T} \right) \sum_{n=2}^{n=x} q^{n+1} m^n \right].$$

¹¹We implicitly assume here that players cannot observe if transactions occurred before they trade. However, we do not need such a strong assumption. For example, if each transaction was publicly announced with a probability strictly smaller than one, our results would still hold. This probability should be small enough so that the likelihood of being last in the sequence is not too high.

¹²When there exists a bubble-equilibrium in pure strategies, there can also exist mixed-strategies equilibria in which traders enter the bubble with a positive probability that is lower than 1. We have characterized these equilibria for the two-player case. They involve peculiar evolutions of the probability to enter the bubble depending on the price level that is observed. We thus do not use these mixed-strategy equilibria in our analysis.

This limit converges if and only if $qm < 1$. This inequality is in conflict with the IR condition according to which $m \geq \frac{1-q^T}{q-q^T}$. This implies that the only games in which the ex-ante welfare is well-defined are the games where only the no-bubble equilibrium exists. This makes it hard to conclude that the imperfect information game is actually creating welfare even if interim (that is, knowing the proposed price), all traders are strictly better off entering the bubble if they anticipate that other traders are also going to do so.¹³

3 Experimental design

This section proposes a simple experimental design in which bubbles can arise at equilibrium. Based on the previous theoretical analysis, this design features a sequential market for an asset whose fundamental value is 0. There are three traders on the market.¹⁴ Trading proceeds sequentially. Each trader is assigned a position in the market sequence and can be first, second or third with the same probability $\frac{1}{3}$. Traders are not told their position in the market sequence but get some information when observing the price at which they can buy the asset. Prices are exogenously given and are powers of 10. For simplicity, in this experimental design, we do not include the issuer of the asset. The first trader is offered to buy at a price $P_1 = 10^n$. In the baseline experiment, the power n follows a geometric distribution of parameter $\frac{1}{2}$, that is $\mathbb{P}(n = j) = \frac{1}{2}^{j+1}$, with $j \in \mathbb{N}$. The geometric distribution is useful from an experimental point of view because it is simple to explain and implies that the probability to be last in the market sequence conditional on the proposed price is constant. This probability is equal to 0 if the proposed price is 1 or 10, and is equal to $\frac{4}{7}$ otherwise.¹⁵ If he decides to buy the asset, trader i in the market sequence proposes the asset to the next

¹³The St. Petersburg Paradox can be solved by considering that players are risk averse. In contrast, the supplementary appendix to this paper (Moinas and Pouget, 2010) shows that the expected welfare is not defined even for risk-averse players, for games in which bubbles are rational. This is because of the conflict between the IR condition that requires the utility function to be not too concave, and the convergence requirement that requires the utility function to be concave enough. We show that this conflict cannot be resolved in our context.

¹⁴We could have designed an experiment with only two traders per market. However, this would have required higher payments for bubbles to be rational. Indeed, the conditional probability to be last would be higher. We could also have chosen to include more than three traders per market. We decided not to do so in order to have a sufficiently high frequency of the different price levels.

¹⁵The probabilities to be first, second or third conditional on the prices, which are computed using Bayes' rule, are given to the participants in the Instructions.

trader at a price $P_{i+1} = 10P_i$. In order to avoid participants from discovering their position in the market sequence by hearing other subjects making their choices, subjects play simultaneously: once P_1 has been randomly determined, the first, second and third traders are simultaneously offered prices of P_1 , P_2 , and P_3 , respectively, at which they can choose to buy the asset.¹⁶ If they decide to buy, they automatically try and resell the asset.

If we stopped here, participants' net payoffs (that is, their gains and losses relative to their initial wealth) would be as illustrated in Figure 1. Payoffs depend on the various traders' decisions. For example, if the first and the second traders buy while the third one refuses to buy, payoffs are $P_2 - P_1$, $-P_2$, and 0, respectively. Except for the case in which the first trader refuses to buy (so that the bubble does not start), each potential market outcome of the game translates into a loss for one of the market participants (the one who is not able to resell the asset).¹⁷ Since the first price is distributed on an unbounded support, this loss can be very large. This feature is unappealing because experimental subjects cannot be asked to pay large amounts of money. We thus introduce limited liability in a way that does not affect subjects incentive to enter into bubbles. To do so, we rely on a delegated portfolio situation that we refer to as the trader/financier game (as opposed to the self-financed trader game that we considered up to now).

We now present the trader/financier game and show that a bubble equilibrium can also exist in this case, i.e., that both traders and financiers find it beneficial to ride/finance the bubble. Each trader is endowed with 1 euro. If additional capital is required in order to buy the asset at price P_t , this additional capital (that is, $P_t - 1$) is provided by an outside financier. We assume that each trader is financed by a different financier. Payoffs (potential gains and losses) are then divided between the trader and the financier in proportion of the capital initially invested: a fraction $\frac{1}{P_t}$ for the trader, and a fraction $\frac{P_t-1}{P_t}$ for the financier. When a trader decides to buy the asset at price P_t , if the next trader refuses to trade is, his final wealth is 0 (he loses the euro he was initially endowed with). The outside financier also ends up with 0 (he loses the $P_t - 1$ euros he has invested). This enables

¹⁶When a trader does not accept to buy the asset, subsequent traders end up with their initial wealth whatever their decisions.

¹⁷At each outcome of the game (except if the first trader refuses to buy), the total payoff is equal to $-P_1$. This aggregate loss corresponds to the gain of the issuer of the asset. Taking the issuer's payoff into account, the game is thus a zero-sum game. As explained before, the issuer is not part of the experiment. This role is performed by the experimenter.

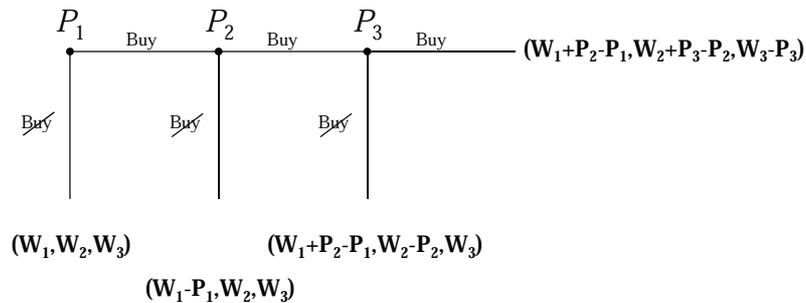


Figure 1: Timing of the market in the trader game
Market proceeds sequentially. Traders are equally likely to be first, second or third. The first, second and third traders are offered to buy at prices P_1 , P_2 , and P_3 respectively. This figure displays traders net payoff, that is their gains or losses relative to their initial wealth. Traders can end up losing very large amount of money.

us to implement limited liability in the experiment: the maximum potential loss of a trader is 1. The potentially infinite gains and losses are incurred by the financiers. As a result, we do not include them in the experiment.¹⁸

When the trader is able to resell the asset, he gets $\frac{1}{P_t}$ percent of the proceed $P_{t+1} = 10P_t$ and thus ends up with a final wealth of 10 which corresponds to 1 euro of initial wealth plus a gain of 9 euros. Overall, whatever the price at which the trader buys, he can lose 1 euro or win 9 euros. If he anticipates that other traders buy the asset, it is in a trader's best interest to also buy the asset if the following individual rationality (IRi) condition is satisfied:

$$\frac{3}{7}U_i(10) + \frac{4}{7}U_i(0) \geq U_i(1),$$

where $U_i(\cdot)$ is the trader's utility function. For a bubble equilibrium to exist, the IRi condition has to be satisfied for all traders on the market. It is straightforward to show that there exists functions $U_i(\cdot)$ for which the IRi condition holds.

¹⁸The experimenter is implicitly playing the role of the financier for all traders in addition to playing the role of the issuer. As a result, the experimenter's maximum payment to the subjects is 20 euros. This maximum payment occurs when all subjects decide to enter the bubble. The experimenter is thus not subject to a bankruptcy risk.

This IRi condition in the trader/financier game mirrors the IR condition that prevails for the trader game presented in the previous section (see appendix A, for an analysis of the trader game with risk averse agents). Given a proposed price, the strategic incentives that agents face in both games are similar. In fact they are exactly the same if the proposed price is 1, and they are the same modulo a factor of P_t if the proposed price is P_t .

In order to show that the trader/financier game is meaningful, we now check that financiers have an interest in providing capital to traders. The individual rationality of an outside financier (IRf) is written as:

$$\frac{3}{7}U_f(f_t(P_t) - P_t - 10) + \frac{4}{7}U_f(-P_t + 1) \geq U_f(0), \forall t \in \{1, \dots, T\}, \text{ and } \forall P_t \in P,$$

where $U_f(\cdot)$ is a financier's utility function. For a bubble equilibrium to exist, the IRf condition has to be satisfied for all financiers on the market. It is straightforward to show that there exist functions $U_f(\cdot)$ for which the IRf condition holds. Again, in order to limit potential losses in our experiment, subjects do not play the role of outside financiers.

The timing of the trader/financier game in which subjects are involved is illustrated in Figure 2 (the payoffs of the financiers who are not part of the experiment are also indicated for completeness).¹⁹ Notice that the sum of traders' and financiers' payoffs in the trader/financier game equals the payoffs of the traders in the trader game. If one considers the payoffs of the issuer, it is also a zero-sum game (not counting the three euros initially provided to subjects by the experimenter).

In order to study how formation of bubbles is influenced by their rationality, we also focus on an experimental setting where there is a cap K on the first price. This will allow us to relate our work to the previous experimental literature on bubbles initiated by Smith et al. (1988) that focuses on irrational bubbles. Indeed, as shown in the previous section, in this setting, bubbles are irrational in the sense that they would be ruled out by backward induction. A cap of K on the first price translates into a cap of $100K$ on the highest potential price in the experiment. Upon being proposed this price, a subject should understand that he is last in the market sequence and, consequently, should refuse to buy. Anticipating this refusal, subjects who are proposed lower prices should also refuse to buy. At equilibrium, the bubble never starts. However, given the experimental results on the failures of backward induction (for results on the centipede game, see Mc Kelvey and

¹⁹This timing does not correspond to the extensive form of the game. Indeed, it leaves aside the issue of which player is first, second, or third. The extensive form game is provided in Appendix B for the two-player case. When there is no cap on the first price, it includes an infinite number of nodes.

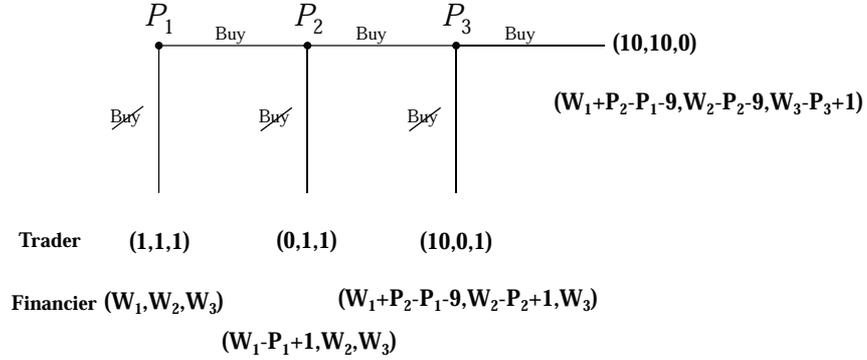


Figure 2: Timing of the market in the trader/financier game. Market proceeds sequentially. Traders are equally likely to be first, second or third. The first, second and third traders are offered to buy at prices P_1 , P_2 , and P_3 respectively. This figure displays traders net payoff, that is their gains or losses relative to their initial wealth. Traders can end up losing very large amount of money.

Palfrey (1992) or Kawagoe and Takizawa (2009)), bubbles might develop in our experiment.

Varying the level of the cap K offers potentially interesting comparative statics because it affects the number of iterated steps of reasoning that are needed in order to reach the Nash equilibrium. When the proposed price is $P = 100K$, a subject knows that he is last and there is no iterated step of reasoning needed. When the proposed price is $P = 10K$, a subject knows that he is not first in the market sequence (he can be second or third). At equilibrium, he has to anticipate that the next trader in the market sequence (if any) would not accept to buy the asset. One step of iterated reasoning is thus needed to derive the equilibrium strategy. More generally, when the proposed price is $1 \leq P \leq 100K$, the required number of iterated steps of reasoning is $\log_{10} \left(\frac{100K}{P} \right)$. In order to study whether this required number of iterated steps of reasoning could affect bubble formation, we have chosen to experimentally study cases in which K equals 1, 100, and 10,000.

The experimental protocol is as follows. Our baseline experiment includes a total of 93 subjects. Subjects are in the last year of the Bachelor in Business Administration at the University of Toulouse. We have four sessions with 21 or 24 subjects per session. Each subject participates in only one session and receives a 5-euro show-up fee. Each session includes only

one replication of the trading game.²⁰ Subjects' risk aversion is measured thanks to a procedure which is inspired from Holt and Laury (2002). We adjust their questionnaire in order to match the set of possible decisions to the decisions subjects actually face in our experiment.²¹ The instructions indicate the conditional probabilities to be first, second, and third given the price subjects are proposed. The minimum, median, maximum, and average gains in the experiment are respectively 0, 1, 10, and 3.62 euros (not including the show-up fee). The instructions for the case where $K = 10,000$ are in Appendix C.

Our experimental design is summarized in Table I.

Session	# Replications	# Subjects	cap on initial price, K	Equilibrium
1	1	24	1	no-bubble
2	1	21	100	no-bubble
3	1	24	10,000	no-bubble
4	1	24	$+\infty$	no-bubble or bubble

Table 1: Experimental design

4 Baseline Results

4.1 Market behavior

We first start by analyzing overall market behavior. At this aggregate level, we can measure the frequency as well as the magnitude of bubbles. The frequency of bubbles is defined as the proportion of replications where the first trader accepted to buy the asset. The magnitude of bubbles is referred to as large if all three traders accepted to buy the asset, medium if the first two traders accepted, and small if only the first trader accepted. Figure 3 presents the results per session.

²⁰The influence of learning on bubble formation is analyzed in the robustness section.

²¹The questionnaire is composed by a table with 14 decisions. For each decision i , subjects may choose between the riskless option A, which is to receive 1 euro for sure, or the risky option B, which is to receive 10 euros with probability $\frac{i}{14}$, or 0 euro with probability $\frac{14-i}{14}$. This questionnaire features what Harrison, List and Towe (2007) refer to as a higher frame: a risk-neutral agent switches to the risky option B in the upper part of the table. Such questionnaire gives us a fine estimation of the willingness to accept the bets at stake in our bubble game.

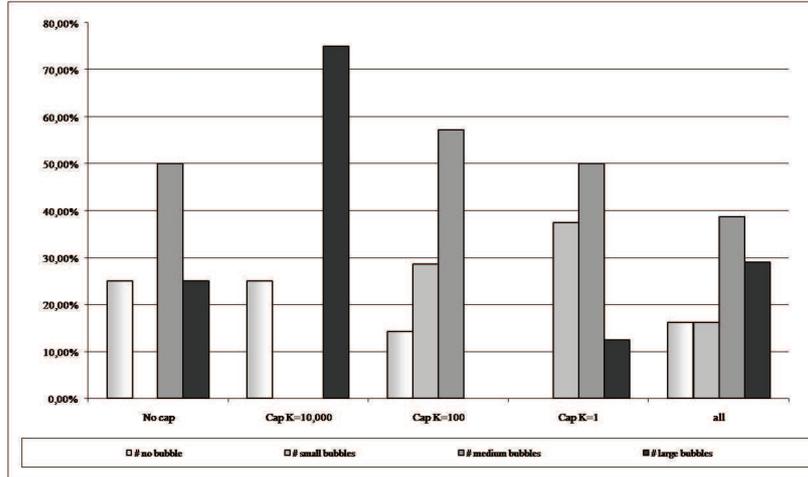


Figure 3: Probability to observe bubbles, depending on the cap on the initial price

Figure 3 shows that there are bubbles in an environment where backward induction is supposed to shut down speculation, namely when there exists a price cap. This is in line with the previous experimental literature cited in the introduction. We observe large bubbles even in situations where the existence of a cap enables some subjects to perfectly infer that they are last in the market sequence. A potential explanation is related to bounded rationality.²² It is indeed possible that some traders make mistakes and buy, in particular (but not only) when proposed a price of $100K$. We also observe bubbles when there is no price cap, that is, when there exists a bubble equilibrium. However, bubbles are not more likely when there is no price cap than when there is a price cap. This indicates that traders fail to perfectly coordinate on the bubble equilibrium. Proposition 3 indicates that this could be due to strategic uncertainty because there always exists a no-bubble equilibrium, or to risk aversion because if traders are sufficiently risk averse, it is not beneficial to enter the bubble. An additional interpretation is that the possible existence of irrational traders, who may not buy when it would be rational, or the failure of the common knowledge of rationality, increase the risk of entering the bubble for rational traders and may prevent

²²An alternative explanation could be related to social preferences. However, extreme altruism would be required in order for a subject to be willing to lose in order to let other subjects gain. We therefore do not focus on this interpretation.

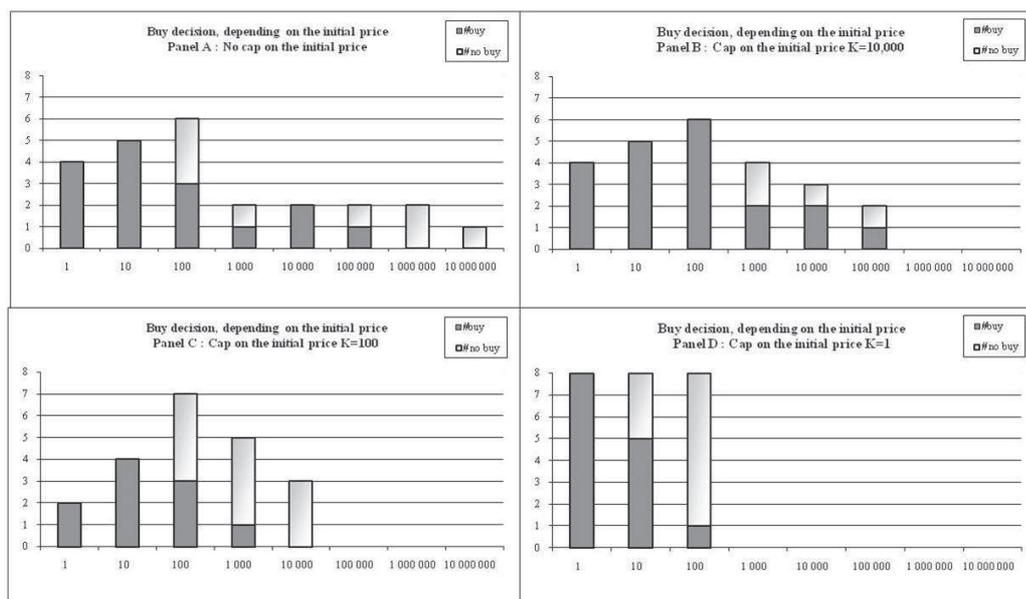


Figure 4: Buy decision, depending on the initial price

them from doing so.

4.2 Individual behavior

To gain more insights on bubbles formation, we now study individual decisions to buy the overvalued asset. Figure 4 plots our entire data set. For each session, each bar represents the number of times a given price has been proposed. Within each bar, the dark gray part corresponds to buy decisions while the white part corresponds to refusals to trade. Figure 4 complements the results found in the previous subsection. Let's first focus on the case where there is no cap on the first price (panel A). In this treatment, participants who are proposed prices of 1 or 10 always buy the asset while those who are proposed higher prices buy approximately half of the time. This pattern is consistent with Nash equilibrium if one takes into account subjects' risk aversion: when they are proposed a price of 1 or 10, subjects are sure not to be last, whereas when they are proposed higher prices, they have 4 chances out of 7 to be last. When prices are 100 or above, if participants coordinate on the same equilibrium, their decisions should be the same for all price levels. In line with this hypothesis, using a Wilcoxon rank sum test, we cannot reject the fact that the probability to buy is the same after observing prices of 100, 1,000, and 10,000, and after observing higher prices.

The high probabilities to buy in the sessions where there is a price cap (panels B, C and D) are inconsistent with Nash equilibrium. Still, we keep observing the same pattern as in the situation where there is no cap on the

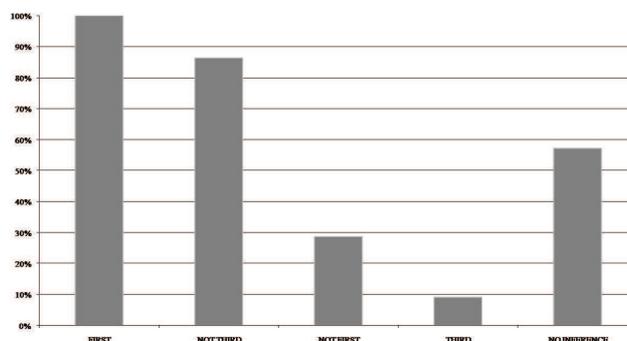


Figure 5: Probability to buy conditional on subjects' inferences

initial price. In sessions with a cap though, higher prices are informative on a trader being last in the sequence. To see whether the probability to be third matters, Figure 5 reports the probability to buy, conditional on participants' inference on their position. On the one hand, the probability to buy decreases with the likelihood that a trader is last in the market sequence. Traders buy very often when they are sure to be first and sure not to be last, while they buy half of the time when they cannot infer their position. This is consistent both with the fact that they face a more risky decision and with the fact that traders are reluctant to coordinate on the bubble equilibrium. Also, this indicates that there are some elements of rationality in subjects' decisions. On the other hand, this result holds whether bubbles are irrational or not. In particular, the propensity of subjects to enter a bubble is extremely large when they know that they are first or second, even in a situation in which there exists a price cap. This result contradicts the predictions of Nash equilibrium, so that rationality does not appear to be perfect. Our results also suggests that some traders anticipate this and speculate on others' mistakes or on the fact that others speculate on others' mistakes.

Figure 6 reports the probability to buy for various level of risk aversion. It shows that risk aversion decreases the propensity to enter into bubbles.

Figure 7 reports the probability to buy as a function of the number of steps of reasoning needed for a trader to realize that someone knows he is last. This figure shows that the propensity to enter into bubbles increases with the number of steps. When there is no cap, that is, when the number of steps is infinite, it seems that the probability to buy is lower than when

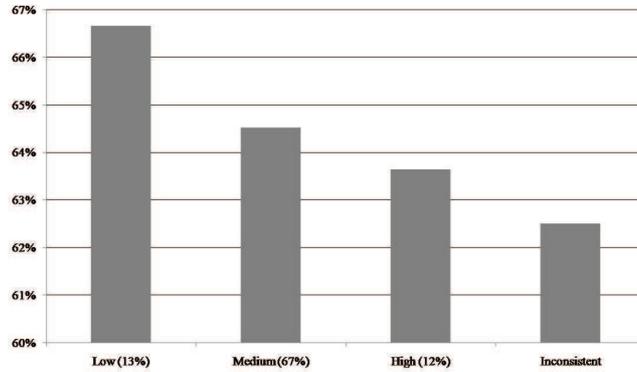


Figure 6: Probability to buy depending on the level of risk aversion

there are 4, 5 or 6 steps but this result disappears when controlling for other variables.

In order to statistically establish these results, we run a logit regression.²³ The propensity to buy the overvalued asset is explained by several variables: a dummy indicating that a subject is one step from the maximum price and cannot be last (that is, has observed a price of 1 or 10), a dummy indicating that a subject is one step from the maximum price and may be last (that is, has observed a price strictly higher than 10), a dummy indicating that a subject is two or more steps from the maximum price and cannot be last, a dummy indicating that a subject is two or more steps from the maximum price and may be last, a dummy indicating that there is no price cap and that the subject cannot be last, a dummy indicating that there is no price cap and that the subject may be last. We also include as explanatory variables the proposed price, and the degree of risk aversion.²⁴ The constant reflects the propensity to enter of subjects who are proposed to buy at the maximum price (100K). This is useful because, since these subjects are expected not to enter, their probability to enter can be viewed as the incompressible level of noise in our data. The regression then enables us to study whether the subjects who are further away from the maximum price enter the bubble significantly more.

The results are in column 1 of Table II. We first focus on the subjects who

²³The results are the same if we run a probit regression.

²⁴The coefficient of risk aversion is computed assuming a constant relative risk aversion utility function as in Holt and Laury (2002).

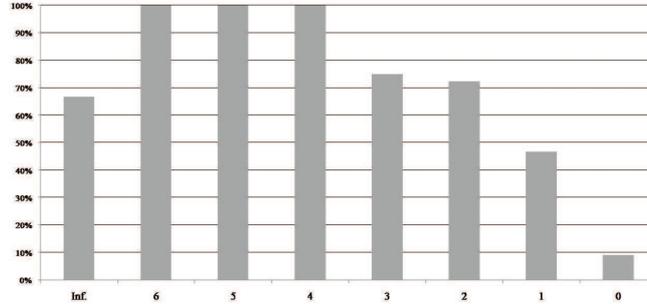


Figure 7: Probability to buy depending on the number of steps of iterated reasoning

know they are last in the market sequence. We can reject the fact that these subjects enter the bubble with probability zero. Indeed, we observe that, out of the eleven subjects who knew they were last, one subject entered into the bubble. This number is low but it is not zero. This result is in line with the findings of Lei, Noussair and Plott (2001) that subjects were buying an overvalued asset even when prohibited to resell. These agents can be viewed in our framework as “step 0” subjects. In page 853, they report that 6 out of 36 subjects made at least one dominated transaction.²⁵ This proportion is slightly higher than our proportion of dominated choices (1/11), maybe reflecting the more complicated framework used in their experiment.

We then complement their results by studying the behavior of subjects who are further away from the maximum price, that is, “step1” and “step 2 and more” subjects. This is interesting because these subjects might be tempted to rationally speculate on others’ irrationality. We uncover a snowball effect: the propensity to enter bubbles increases with the distance from the maximum price. The coefficients of the “step 1” dummy (1.61) is positive even if not significant (p-value=0.25). The coefficient of the “step 2 and more” dummy (3.05) is also positive and it is statistically significant (p-value=0.01). The propensity of subjects to enter bubbles when there is no cap (the estimated coefficient is 2.66 and the p-value is 0.03) is the same as when subjects are two or more steps from the maximum price (a Wald

²⁵We focus here on the experiment of Lei, Noussair and Plott (2001) during which subjects could participate in several markets, the so-called TwoMarket/NoSpec treatment. This is because, in this treatment, subjects could participate actively in the experiment without being forced to participate in the bubble. This provides a lower bound for the number of subjects who make mistakes in the market.

	1. Baseline (one-shot)		2. Baseline (one-shot) and first replication		3. Baseline (one-shot), first and fifth replication	
	Coefficient	p-value	Coefficient	p-value	Coefficient	p-value
Constant	-1.28	0.267	-1.98	0.066	-2.06	0.053
One step from maximal price and not last	3.12	0.022	3.74	0.002	3.68	0.003
One step from maximal price and maybe last	1.61	0.246	1.98	0.141	1.93	0.148
Two or more steps from maximal price and not last			7.05	0.000	6.92	0.000
Two or more steps from maximal price and maybe last	3.05	0.011	3.42	0.002	3.39	0.002
No cap and maybe last	2.66	0.034	3.15	0.006	3.12	0.006
Zero step from maximal price, and fifth replication					0.91	0.551
One step from maximal price and not last, and fifth replication					-0.85	0.455
Two or more steps from maximal price and not last, and fifth replication					-1.20	0.411
Two or more steps from maximal price and maybe last, and fifth replication					-0.63	0.336
No cap and maybe last, and fifth replication					-0.41	0.624
Price	-0.00	0.225	-0.00	0.235	-0.00	0.257
Risk aversion	-2.08	0.043	-1.77	0.019	-1.52	0.018
Log likelihood	-31.81		-57.18		-82.32	
Number of observations	61		144		197	

Table 2: Logit Regression on the Buy Decision

test cannot reject the equality of the two coefficients, p-value=0.63).

We find a high propensity to enter bubbles for subjects who know they are not last. For subjects who are at one step from the maximum price, when they know they cannot be last, their propensity to enter is significantly higher than "step 0" subjects (the coefficient is 3.12 and the p-value is 0.02). When there are two or more steps from the maximum price or when there is no price cap, the propensity to enter is extremely high: in fact, all the subjects who received a price of 1 or 10 and who knew they were not last decided to enter. This is the reason why we drop these observations from our sample (the sample size for this regression is thus 61 instead of 93; we drop 23 observations that correspond to the "step 2 and more" subjects who were sure not to be last, and 9 observations that correspond to the "no cap" subjects who were sure not to be last).

Overall, these results indicate that, in our experiment, the formation of irrational bubbles is related to a snowball effect. Indeed, since there is only a $\frac{1}{11}$ probability that "step 0" subjects enter the bubble, if "step 1" subjects correctly anticipate this, they should not enter the bubble. This is because the expected payoff from entering the bubble is $\frac{10}{11}$ which is smaller than the payoff of 1 from not buying the asset. Only "step 1" subjects who make a mistake would enter, that is, a proportion $\frac{1}{11}$ of agents if we assume that the error rate is the same for "step 0" and "step 1" agents. "step 2" subjects would thus not have an incentive to enter the bubble either and the reasoning goes on for higher steps. As a result, the propensity to enter the bubble would be expected to be very small whatever the distance from the maximum price. A snowball effect has to be incorporated into a model in order to capture the observations from the laboratory. The study of this snowball effect is done in the next section.

4.3 Fitting models of bounded rationality

Our results so far suggest that some players have bounded rationality and that the formation of bubbles is related to a snowball effect. To account for these phenomena, we estimate models that explicitly incorporate bounded rationality: Camerer, Ho and Chong (2004)'s cognitive hierarchy model and McKelvey and Palfrey (1995)'s quantal response equilibrium.

In cognitive hierarchy (CH) theories, traders best-respond to mutually inconsistent beliefs. Traders differ in their level of sophistication l , and each player believes he understands the game better than the other players. Specifically, level-0 traders play randomly. Level-1 traders believe that other traders are level-0, and level- l traders believe that other traders are a mixture

		Session				
		No Cap	Cap K=10,000	Cap K=100	Cap K=1	All
Data						
	Sample size	24	24	21	24	93
	Av. probability buy	67%	83%	48%	58%	65%
Nash Equilibrium						
	Av. probability buy	100%	0%	0%	0%	26%
	Log L	-73,68	-184,21	-92,10	-128,95	-478,94
Cognitive Hierarchy						
	Tau	0,41	1,37	0,27	2,65	0,67
	Av. probability buy	67%	86%	58%	56%	69%
	Log L	-15,28	-10,36	-13,96	-9,25	-53,04
	90% CI	[0 - 1.10]	[0.67 - 3.91]	[0 - 1.79]	[0.87 - 4]	[0.48 - 1.12]
Quantal Response						
	Lambda	0,43	2,69	2,49	1,93	2,54
	Av. probability buy	80%	79%	48%	58%	70%
	Log L	-5,18	-7,99	-7,53	-8,31	-25,90
	90% CI	[0.56 - 2.52]	[1.77 - 2.86]	[0.54 - 2.83]	[0.46 - 1.16]	[0.51 - 2.74]

Table 3: Comparison of fits of Nash, Cognitive Hierarchy and Quantal Response Equilibrium

of level- $(l-1)$, level- $(l-2)$, ..., and level-0. Traders then best-respond using this belief. As in Camerer, Ho and Chong (2004), we assume that traders' types are distributed according to a Poisson distribution. Let τ denote the average level of sophistication. For each parameter τ and each level of price P , we find the best-response of a risk-neutral level-1 trader who considers that the next trader is a level-0 player observing $P \times 10$. After one iteration, we find the best-response of a risk-neutral level-2 player who considers that the next trader, observing $P \times 10$, is a level-0 player with probability $\frac{\exp(-\tau) \times \tau^0}{0!}$, and a level-1 player with probability $\left(1 - \frac{\exp(-\tau) \times \tau^0}{0!}\right)$. We use a similar iterative process to find the best-response of a level- l player who considers that the next trader, observing $P \times 10$, is a level- j player with probability $\frac{\exp(-\tau) \times \tau^j}{j!}$ for $j < l - 1$, and a level- $(l - 1)$ player with the complementary probability. This process is fully described in the supplementary appendix of this paper (Moinas and Pouget, 2010) for the case where the price cap is $K = 1$. This enables us to determine the likelihood function under the assumption that subject are risk-neutral. We then estimate the parameter τ by maximum likelihood in each session as well as for the entire dataset.

Results are reported in Table III. The fit of the CH-Poisson model is compared with the fit of Nash equilibrium, by session and overall. The

first two lines describe our data, namely, the group size, and the observed average probability to buy. The middle three lines show the predictions of the Nash equilibrium under a similar assumption of risk-neutrality, the log-likelihood of this model, and its mean squared deviation.²⁶ The mean choices are generally far off from the Nash equilibrium; the probability to buy is too low when there exists a bubble-equilibrium, and too high when it does not exist. The five next lines report the estimate of the parameter τ , the predictions of the cognitive hierarchy model for this value of the estimate, and the corresponding log-likelihood and mean squared deviation. We further compute the 90 percent confidence interval for τ estimated from a randomized resampling (bootstrap) procedure using 10,000 simulations. We estimate an overall average level of sophistication of 0.67, which is consistent with the estimates reported in Camerer, Ho and Chong (2004). This suggests a high proportion of level-0 players, that is, almost 50%. Interestingly, what drives this result is not really the fact that traders enter too much into bubbles when they should not. Indeed the fact that there is only around 10% of subjects who buy when they know they are last in the market sequence suggests a proportion of level-0 players equal to 20%. What explains the high estimated proportion of level-0 players is rather the fact that subjects do not buy as much as expected by the cognitive hierarchy model (with a higher average sophistication level) when there is no cap on the initial price or when the cap is large. The reason why subjects do not buy as much as expected by the risk-neutral cognitive hierarchy model can be related to risk aversion. This suggests that the average level of sophistication may be underestimated due to the risk-neutrality assumption.

The Poisson CH model retains best-response (except for level-0 players), but it weakens equilibrium (that is, belief-choice consistency). McKelvey and Palfrey (1995) propose an alternative approach, which retains equilibrium but weakens best-response. In their Quantal Response Equilibrium, players may make mistakes. However, the likelihood of these errors depends on the impact of such errors on players' expected utility. More specifically, the following logit specification of the error structure is assumed, so that, if the buy decision yields an expected profit of u_B while the no buy decision yields an expected profit of u_\emptyset , the probability to buy is: $\Pr(B) = \frac{e^{\lambda u_B}}{e^{\lambda u_B} + e^{\lambda u_\emptyset}}$.

This enables us to determine the likelihood function under the assump-

²⁶We consider that traders coordinate on the bubble equilibrium when there is no cap on the initial price. The no-bubble Nash equilibrium has a lower log-likelihood. In order to compute these likelihoods, we assume that players choose non-equilibrium strategies with a probability of 0.0001.

tion that subject are risk-neutral (The supplementary appendix describes how this probability to buy is computed conditionally on the proposed price for the case where the price cap is $K = 1$). We then estimate the parameter λ , which we refer to as responsiveness to expected payoffs, by maximum likelihood in each session and for the entire dataset. Responsiveness is inversely related to the level of errors made by subjects. The results are reported in the last five lines of Table III. We estimate an overall average responsiveness of 2.54, which is consistent with the results of McKelvey and Palfrey (1995). The QRE seems to fit better our data than the cognitive hierarchy model. This result seems to contradict those of Kawagoe and Takizawa (2008), who compare the goodness of fit of both models in laboratory experiments of the centipede game. To further investigate this issue, we compare in Figure 8 the probability to buy conditional on the proposed price, in the CH model and in the QRE, with our observations. What the QRE seems to better capture in our data is the drop in the probability to buy for prices larger than $P = 100$. In the CH model, this pattern either does not characterize the expected outcome (see the case in which there is no cap on the initial price), or captures it less intensively and with a lag (see the cases in which there is a cap at $K = 100$ or $K = 10,000$). In the QRE however, costlier mistakes are less likely. This model is thus able to capture the drop in players' expected utility from buying: when they are proposed a price $P \geq 100$, the conditional probability to be third is greater than or equal to $\frac{4}{7}$, whereas, when they are proposed a price of 1 or 10, the conditional probability to be third is zero. This feature is present in our design and is different from the centipede games analyzed by Kawagoe and Takizawa (2008).

Both the CH model and the QR equilibrium framework capture the snowball effect that underlies the formation of bubbles in our experiment but for different reasons. In the CH model, agents tend to overestimate the probability that the agent at the next step will enter. Indeed, at step 0, only level-0 players enter with probability one half. These players also enter at step 1 but, at this step, level-1 players also enter because they wrongly believe that all the other players are level-0. These wrong beliefs imply that when the distance from the maximal price is higher, the level of the players that are willing to enter the bubble is also higher. In the QR equilibrium, at step 0, entering in the bubble is pretty costly, since it generates a loss for sure, so that agents enter pretty rarely. At step 1, given that an agent expects that some players mistakenly enter the bubble at step 0, entering the bubble is less costly than at step 0, which translates into a higher probability to enter. This phenomenon is repeated at each step which generates the snowball effect.

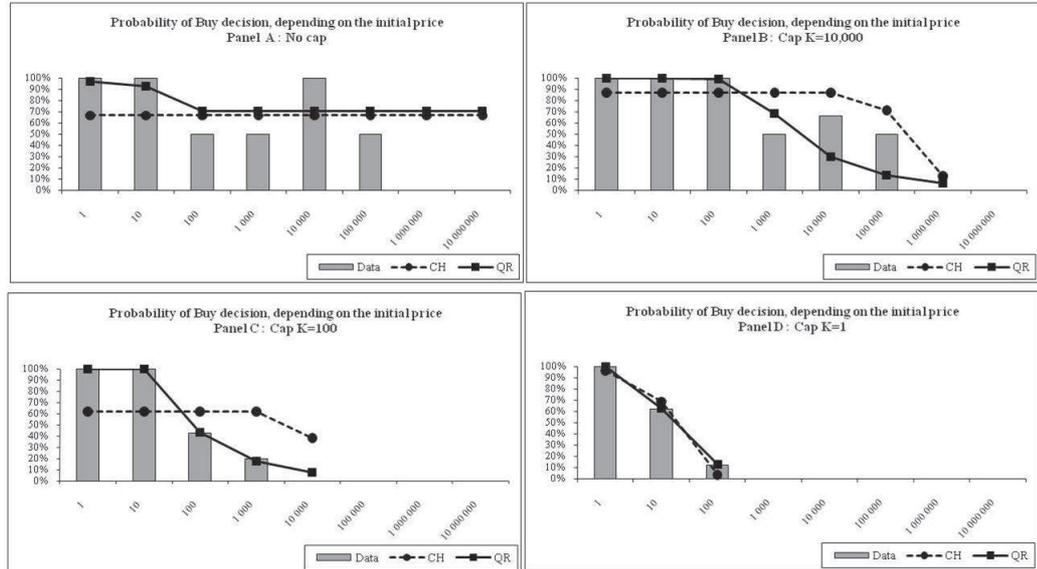


Figure 8: Probability of Buy decision, depending on the initial price

5 Robustness: learning and professional experience

This section extends our analysis to study the effect of learning and professional experience. The first subsection reports an experiment where subjects play five replications of the game. The second subsection reports an experiment run with Executive MBA students at the London Business School.

5.1 Learning

In order to study how learning affects bubble formation, we run exactly the same experiment as before except for the number of replications. Subjects are now playing five replications in a stranger design (and this is common knowledge): subjects do not know with whom they are playing and it is very unlikely that they will be playing again with the same subjects. The experiment includes 66 subjects from the first year of Master in Finance at the University of Toulouse. This pool of students is very similar to the baseline experiment pool. There are four sessions with 15 or 18 subjects. Each subject participates in only one session and receives a 5-euro show-up fee. The minimum, median, maximum, and average gains in this experiment

are respectively 1, 13, 41, and 16 euros (not including the show-up fee).

This experimental design is summarized in Table IV.

Session	# Replications	# Subjects	cap on initial price, K	Equilibrium
5	5	15	1	no-bubble
6	5	18	100	no-bubble
7	5	15	10,000	no-bubble
8	5	18	∞	no-bubble or bubble

Table 4: Experimental design of the 5-period experiments

We start by constructing a data set that includes the baseline (one-shot) experiment and the first replication of the learning experiment. We thus have 159 observations. We run the same logit regression as in the baseline experiment. The results are in column 2 of Table II. The coefficient estimates and significance levels are very similar to the baseline case. Very few “step 0” enter the bubble but the proportion is not zero: 1/17. The propensity to enter bubbles increases with the distance from the maximum price. This propensity increases when subjects know they are not last. In particular, when there is no price cap, the propensity to enter bubbles is very high: in this case, 100% of the subjects buy the asset after receiving a price of 1 or 10. We drop 15 out of our 159 observations corresponding to these subjects who know they are not last.

A first look at the effect of learning in our experiment is offered by adding to the data set the fifth replication of the learning experiment. We then include in the experiment a dummy variable indicating that the observation corresponds to the fifth session and we interact this dummy with the other explanatory variables of interest. The results are in column 3 of Table II. Overall, the coefficients of the fifth replication dummy variable and its interactions appear negative but insignificant. This seems to indicate that the propensity to enter bubbles is not really lower during the fifth period.

To investigate further this result, we focus on the 66 subjects who participated in the five replications and we run a panel logit regression that controls for period and individual fixed-effects. We drop 25 observations corresponding to 5 subjects who either always buy or always do not buy. Our regression uses the set of explanatory variables detailed above, aggregating “step-0” and “step-1” observations due to the low number of “step-0” observations. In addition, we include the following variables: a dummy that indicates that a subject bought and lost at least once in a previous replica-

	Coefficient	Statistic	p-value
Constant	-9.44	-4.55	0.000
Two or more steps from maximal price and not last	50.24	14.65	0.000
Two or more steps from maximal price and maybe last	45.30	13.28	0.000
No cap and not last	28.29	9.18	0.000
No cap and maybe last	22.65	8.36	0.000
The subject bought and lost at least once in a previous replication and when he or she is not sure to be last	-3.15	-3.38	0.001
The subject bought and won at least once in a previous replication and when he or she is not sure to be last	2.24	1.59	0.113
The subject has been last and knew it at least once in a previous replication	2.54	1.63	0.103
Accumulated gains	-0.42	-2.86	0.004
Risk aversion	-24.57	-11.32	0.000
Dummy which takes value 1 in the 2nd period	0.11	0.12	0.903
Dummy which takes value 1 in the 3rd period	1.63	1.53	0.126
Dummy which takes value 1 in the 4th period	2.27	1.68	0.093
Dummy which takes value 1 in the 5th period	4.94	3.17	0.002
<hr/>			
Log likelihood	-65.29		
Number of observations	305		

Table 5: Panel Logit Regression on the Buy Decision

tion and that he or she is not sure to be last, a dummy that indicates that a subject bought and won at least once in a previous replication and that he or she is not sure to be last, and a dummy that indicates that a subject has been last and knew it at least once in a previous replication. The first two dummies are designed to capture reinforcement or belief-based learning.²⁷ The last dummy captures the behavior of subjects that have experienced what it means to receive the highest potential price. To capture potential wealth effects, we also include an additional control variable, namely the accumulated gains of a subject.

The results are in Table V. As before, a subject's propensity to buy the overvalued asset significantly increases with the number of steps of iterated

²⁷See Camerer and Ho (1999) for a theoretical and experimental analysis of learning in games.

reasoning needed to derive the equilibrium strategy, and significantly decreases with his probability to be last and with his risk aversion. Our estimation further shows that learning has an ambiguous effect on the propensity to enter a bubble: subjects tend to speculate more after good experiences and less after bad experiences. Overall, it is thus not clear that learning leads, at least rapidly, to the no-bubble equilibrium. Finally, it seems that those who have been confronted with the highest price may be more likely to buy when they are subsequently not sure to be last. This might be due to the fact that they realize the complexity of the game and are more ready to bet on other subjects' mistakes.

5.2 Professional experience

In order to study how experienced business people behave as far as bubble formation is concerned, we run exactly the same experiment as in the baseline case except for the origin of the subjects and for experimental incentives. Subjects are now students from the Executive MBA program at the London Business School. Instead of playing for euros, they played for fine chocolate boxes (worth 5 euros each). There is thus a five-time increase in the power of incentives. If a subject buys the asset, he ends up with 10 chocolate boxes if he is able to resell and 0 box if he is not. If he decides not to buy, he keeps the chocolate box. The rest of the design is exactly the same as in the baseline case (subjects played only once).²⁸ This experiment includes 54 subjects. There is only one session with a cap of 10.000 on the first price. The minimum, median, maximum, and average gains in this experiment are respectively 0, 1, 10, and 3.08 chocolate boxes. This experimental design is summarized in Table VI.

Session	# Replications	# Subjects	cap on initial price, K	Equilibrium
9	1	54	10.000	no-bubble

Table 6: Experimental design of the experiment with LBS students

Our results are obtained thanks to a logit regression of the probability to buy the asset. We pool the 54 observations corresponding to LBS executive students with the 24 subjects from Toulouse University who played the one-shot game with a cap at 10.000. Overall we thus have 78 observations. The

²⁸In the interest of time, we did not measure the level of risk aversion of the Executive MBA students.

	Coefficient	Statistic	p-value
Constant	1.12	1.91	0.056
Dummy which equals 1 for LBS subjects	-1.64	-2.43	0.015
Dummy which equals 1 when the subject knows that he is not last	2.87	3.53	0.000
Price	-3.03e-06	-0.55	0.585
<hr/>			
Log likelihood	-36.75		
Number of observations	78		

Table 7: Logit Regression on the Buy Decision, LBS students

explanatory variables are a dummy indicating that a subject is an executive from LBS, and a dummy indicating that a subject knows he is not last. We also include as control variable the proposed price.

The results are in Table VII. Again, a subject’s propensity to buy the overvalued asset significantly increases when he knows he is not last. We also find that executive students at LBS buy the asset less often. This indicates that their behavior is slightly closer to the unique no-bubble equilibrium. Only slightly closer because more than 50% percent of LBS subjects enter in the bubble.

6 Conclusion

This paper studies speculative behavior in a laboratory experiment. The objective is to better understand the various types of speculation that may be observed during bubble episodes: irrational speculation (due to mistakes or erroneous beliefs), rational speculation on others’ irrationality (betting on the fact that others are going to do mistakes), and rational speculation on others’ rationality (betting on the fact that others are going to act rationally). In order to design an experimental environment (that necessarily induces a finite number of trading opportunities) in which bubbles can occur at equilibrium, a theory is developed that extends the insights of Tirole (1982), Allen and Gorton (1993), and Abreu and Brunnermeier (2003). The

idea is to model a financial market as a game in which agents trade an asset sequentially. When agents do not know at which position they are in the market sequence, it can be in their interest to enter the bubble. We design an experimental setting based on this insight. Our design includes several treatments that differ by only one parameter, namely the level of a cap on prices. When there is no cap (or an infinite cap), there exists a bubble equilibrium. When there is a finite cap, there is only a no-bubble equilibrium. However, the higher the cap is, the higher is the number of iterated steps of reasoning needed to reach equilibrium.

Our results show that bubbles are frequently observed when there is a price cap. This is in line with the previous literature on bubbles initiated by Smith, Suchanek and Williams (1988) that shows that bubbles arise even when theory predicts they are irrational. We complement this literature by showing that, when bubbles can be expected in theory, they do materialize but not always. We also show that the decision to speculate and enter into a bubble is positively related to the number of iterated steps of reasoning required to rule out bubbles. This decision is negatively related to risk aversion and to the likelihood to be last in the market sequence.

Maximum likelihood estimations show that both the cognitive hierarchy model of Camerer, Ho, and Chong (2004) and the quantal response equilibrium of McKelvey and Palfrey (1995) capture two of our main experimental observations: we observe bubbles when there is a price cap and we do not always observe bubbles when there is no price cap. This is because both models incorporate the following two effects. First, irrational traders may miss profitable speculation opportunities, thereby reducing the occurrence of rational bubbles. Second, the presence of irrational traders, who buy when it is not beneficial, creates a rationale motive for speculation, even if speculation would be ruled out by backward induction if all traders were rational.

Robustness experiments show that the propensity to enter bubbles does not significantly diminish when the game is repeated five times. Subjects appear to speculate more (respectively, less) when past speculation have been successful (respectively, unsuccessful). When the one-shot game is played by experienced professional businessmen, the propensity to enter into irrational bubbles is lower than with inexperienced students but is still at pretty high levels.

The experimental setting proposed in the present paper opens several avenues of research. It could be interesting to study whether the occurrence of bubbles (rational and irrational) vary with the number of traders, the introduction of risk in the underlying asset payoff, the level of transparency

(one could proxy for transparency by setting a non-null probability that a trade is publicly announced). It would also be interesting to extend the experimental setting to cases in which the price path and the timing are left at the discretion of traders. This would allow testing whether traders are able to coordinate on a price path and a timing that sustains rational bubbles.

7 Appendix

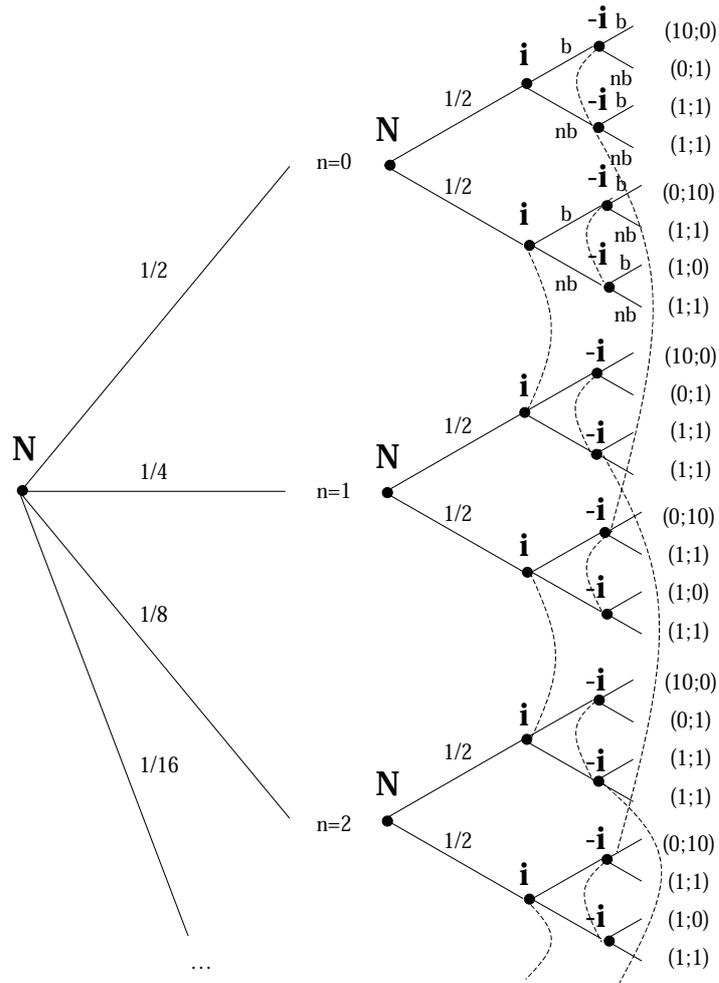
Appendix A: Bubble equilibrium with risk aversion

Consider the environment in which a bubble-equilibrium exists when players are risk-neutral. We now show that a bubble equilibrium can still exist if players are risk averse. The environment is as follows. There are T players. The set of potential prices is defined as $P = \{P_n = m^n \text{ for } m > 1 \text{ and } n \in \mathbb{N}\}$. P_1 is randomly determined following a geometric distribution: $g(P_1 = m^n) = (1 - q)q^n$ with $q \in (0, 1)$. Finally, the price path is defined as $P_{t+1} = m \times P_t$ for $t \in [1, \dots, T - 1]$. For simplicity, we assume that utility functions are piecewise linear with a kink at agents' initial wealth, that is player i 's utility function is: $U_i(x) = x \mathbb{1}_{x \leq W_i} + [W_i + (1 - \gamma_i)(x - W_i)] \mathbb{1}_{x > W_i}$, where $\gamma_i \in]0, 1]$ is a measure of player i 's risk aversion. The IR condition is now written as:

$$\begin{aligned} & \left(\begin{array}{l} (1 - \mathbb{P}[\omega = \omega_T^t | P_t]) \times U_i(W_i + f(P_t) - P_t) \\ + \mathbb{P}[\omega = \omega_T^t | P_t] \times U_i(W_i - P_t) \end{array} \right) \geq U_i(W_i), \forall i \in \{1, \dots, T\}, \text{ and } \forall P_t \in P \\ \Leftrightarrow & \left(\begin{array}{l} (1 - \mathbb{P}[\omega = \omega_T^t | P_t]) \times [W_i + (1 - \gamma_i)(f(P_t) - P_t)] \\ + \mathbb{P}[\omega = \omega_T^t | P_t] \times (W_i - P_t) \end{array} \right) \geq W_i, \forall i \in \{1, \dots, T\}, \text{ and } \forall P_t \in P. \\ \Leftrightarrow & \gamma_i \leq 1 - \frac{\mathbb{P}[\omega = \omega_T^t | P_t] \times P_t}{(1 - \mathbb{P}[\omega = \omega_T^t | P_t]) (f_t(P_t) - P_t)}, \forall i \in \{1, \dots, T\}, \text{ and } \forall P_t \in P. \\ \Leftrightarrow & \gamma_i \leq 1 - \frac{(1 - q)}{(q - q^T)(m - 1)}, \forall i \in \{1, \dots, T\}. \end{aligned}$$

This inequality indicates that, if players are not too risk averse, there exists a bubble equilibrium. It also shows that, when m gets larger, the range of risk aversion for which a bubble equilibrium exists is larger.

Appendix B: Extensive form of the game with two players



At each node, Nature (**N**), player **i** or player **-i** choose an action. $(x;y)$ represents the payoffs; x for player **i**, and y for player **-i**. Dotted lines relate nodes that are observationally equivalent. **b** refers to the buy decision, **nb** to the refusal decision.

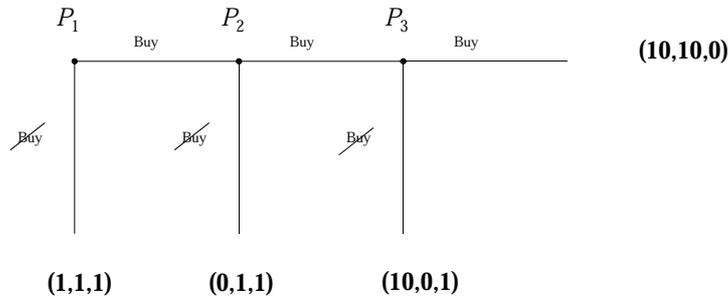
Appendix C: Instructions for the case where $K = 10,000$

Welcome to this market game. Please read carefully the following instructions. They are identical for all participants. Please do not communicate with the other participants, stay quiet, and turn off your mobile phone during the game. If you have questions, please raise your hand. An instructor will come and answer.

As an appreciation for your presence today, you receive a participation fee of 5 euros. In addition to this amount, you can earn money during the game. The game will last approximately half an hour, including the reading of the instructions.

Exchange process

To play this game, we form groups of three players. Each player is endowed with one euro which can be used to buy an asset. Your task during the game is thus to choose whether you want to buy or not the asset. This asset does not generate any dividend. If the asset price exceeds one euro, you can still buy the asset. We indeed consider that a financial partner (who is not part of the game) provides you with the additional capital and shares profits with you according to the respective capital invested. The market proceeds sequentially. The first player is proposed to buy at a price P_1 . If he buys, he proposes to sell the asset to the second player at a price which is ten times higher, $P_2 = 10 \times P_1$. If the second player accepts to buy, the first player ends up the game with 10 euros. The second player then proposes to sell the asset to the third trader at a price $P_3 = 10 \times P_2 = 100 \times P_1$. If the third player buys the asset, the second player ends up the game with 10 euros. The third player does not find anybody to whom he can sell the asset. Since this asset does not generate any dividend, he ends up the game with 0 euro. This game is summarized in the following game tree.



At the beginning of the game, players do not know their position in the market sequence. Positions are randomly determined with one chance out of three for each player to be first, second or third.

Proposed prices

The price P_1 that is proposed to the first player is random. This price is a power of 10 and is determined as follows:

Price P_1	Probability that this price is realized
1	1/2 (50%)
10	1/4 (25%)
100	1/8 (12.5%)
1,000	1/16 (6.3%)
10,000	1/16 (6.3%)

Players decisions are made simultaneously and privately. For example, if the first price $P_1 = 1$ has been drawn, the prices that are simultaneously proposed to the three players are: $P_1 = 1$ for the first player, $P_2 = 10$ for the second player, and $P_3 = 100$ for the third player. Identically, if the first price $P_1 = 10,000$ has been drawn, the prices that are simultaneously proposed to the three players are: $P_1 = 10,000$ for the first player, $P_2 = 100,000$ for the second player, and $P_3 = 1,000,000$ for the third player.

The prices that you are been proposed can give you the following information regarding your position in the market sequence:

- if you are proposed to buy at a price of 1, you are sure to be first in the market sequence;
- if you are proposed to buy at a price of 10, you have one chance out of three to be first and two chances out of three to be second in the market sequence. You are sure not to be last;
- if you are proposed to buy at a price of 100 or 1,000, you have one chance out of seven to be first, two chances out of seven to be second, and four chances out seven to be last in the market sequence;
- if you are proposed to buy at a price of 10,000, you have one chance out of four to be first, one chance out of four to be second, and two chances out four to be last in the market sequence.
- if you are proposed to buy at a price of 100,000, you have one chance out of two to be second, and one chance out of two to be third. In this case, you you are sure not to be first in the market sequence.
- if you are proposed to buy at a price of 1,000,000, you are sure to be last in the market sequence.

In order to preserve anonymity, a number will be assigned to each player. Once decision will be made, we will tell you (anonymously) the group to which you belong, your position in the market sequence, if you are proposed to buy, and your final gain.

Do you have any question?

References

- [1] Abreu D., and M. K. Brunnermeier, 2003, Bubbles and Crashes, *Econometrica*, 71(1), 173-204.
- [2] Ackert L. F., N. Charupat, B. K. Church, and R. Deaves, 2005, Bubbles in Experimental Asset Markets: Irrational Exuberance No More, *mimeo*.
- [3] Ackert L. F., N. Charupat, R. Deaves, and B. Kluger, 2006, The Origins of Bubbles in Laboratory Asset Markets, *mimeo*.
- [4] Allen, F., and G. Gorton, 1993, Churning Bubbles, *Review of Economic Studies*, 60, 813-836.
- [5] Allen, F., S. Morris, and A. Postlewaite, 1993, Finite Bubbles with Short-Sale Constraints and Asymmetric Information, *Journal of Economic Theory*, 61, 206-229.
- [6] Blanchard O., 1979, Speculative Bubbles, Crashes, and Rational Expectations, *Economic Letters*, 3, 387-389.
- [7] Brunnermeier M. K., 2009, Bubbles, Entry in *The New Palgrave Dictionary of Economics*, edited by Steven Durlauf and Lawrence Blume, 2nd edition.
- [8] Brunnermeier M. K., and J. Morgan, 2010, Clock Games: Theory and Experiments, *Games and Economic Behavior*, forthcoming.
- [9] Camerer C., 1989, Bubbles and Fads in Asset Prices: A Review of Theory and Evidence, *Journal of Economic Surveys*, vol. 3, 3-41.
- [10] Camerer C., T.-H. Ho, and J.-K. Chong, 2004, A cognitive hierarchy model of one-shot games, *Quarterly Journal of Economics*, vol. 119, 861-898.
- [11] Conlon J., 2004, Simple Finite Horizon Bubbles Robust to Higher Order Knowledge, *Econometrica*, vol. 72, 927-936.
- [12] Doblaz-Madrid A., 2010, A Robust Model of Bubbles with Multidimensional Uncertainty, Working Paper, Michigan State University.
- [13] Dufwenberg M., T. Lindqvist, and E. Moore, 2005, Bubbles and Experience: An Experiment, *American Economic Review*, vol. 95, 1731-1737.
- [14] Harrison G., J. List and C. Towe, 2007, Naturally Occurring Preferences and Exogenous Laboratory Experiments: A Case Study of Risk Aversion, *Econometrica*, 75, 433-458.
- [15] Holt, C. A., and S. K. Laury, 2002, Risk Aversion and Incentive Effects, *American Economic Review*, 92, 1644-1655
- [16] Kawagoe, T. and H. Takizawa, 2009, Level-K Analysis of Experimental Centipede Games, *mimeo*.
- [17] Lei V., C. Noussair and C. Plott (2001), Nonspeculative Bubbles in Experimental Asset Markets: Lack of Common Knowledge of Rationality vs. Actual Irrationality, *Econometrica*, 69(4), 831-859.
- [18] Mc Kelvey, R. D., and T. R. Palfrey, 1992, An Experimental Study of the Centipede Game, *Econometrica*, vol. 60, 803-836.
- [19] McKelvey, R. D. and T. R. Palfrey, 1995, Quantal Response Equilibrium for Normal Form Games, *Games and Economic Behavior*, 10, 6-38.

- [20] Moinas, S. and S. Pouget, 2010, Supplementary Appendix to “Rational and Irrational Bubbles: An Experiment”, *mimeo*.
- [21] Noussair, C., S. Robin, and B. Ruffieux, 2001, Price Bubbles in Laboratory Asset Markets with Constant Fundamental Values, *Experimental Economics*, 4, 87-105.
- [22] Porter, D. P & Smith, V. L, 1995, Futures Contracting and Dividend Uncertainty in Experimental Asset, University of Chicago Press, vol. 68, 509-41.
- [23] Smith V., G. Suchanek, and A. Williams, Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets, *Econometrica*, 1988, 56, 1119-1151.
- [24] Tirole J., On the possibility of speculation under rational expectations, *Econometrica*, 1982, 50, 1163-1181.
- [25] Tirole J., Asset Bubbles and Overlapping Generations, *Econometrica*, 1985, 53, 1499-1528.
- [26] Van Boening, M.V., Williams, A.W., and LaMaster S. 1993. Price bubbles and crashes in experimental call. markets. *Economic Letters*, 41, 179-185.
- [27] Weil P., 1987, Confidence and the Real Value of Money in an Overlapping Generations Economy, *Quarterly Journal of Economics*, 102, 1-22.