Three Essays in Experimental Economics

Giovanni Giusti

TESI DOCTORAL UPF / ANY 2014

DIRECTORS DE LA TESI

Prof. Robin M. Hogarth, Departament Economia i Empresa

Prof. Charles N. Noussair, Tilburg University

Universitat Pompeu Fabra Barcelona

ACKNOWLEDGMENTS

I would like to thank my two advisors, Robin Hogarth and Charles Noussair.
Robin Hogarth for his support and patience. He has always been a source of motivation and encouragement while writing this thesis.
Charles Noussair made me discover the beauty of doing academic research. He taught me the importance of looking very deeply into results and the importance of exploring many alternative hypotheses before coming to a conclusion.

I also owe a special thanks to Hans-Joachim Voth for introducing me to the elegance of research in economic history. I learnt a lot from his classes and from co-authoring a project with him.
I am especially grateful to Janet Jiang for her co-authorship in the second essay of this thesis and to Rosemarie Nagel for helpful advices and insightful comments during my graduate studies.

I have particularly benefitted by working and interacting with Manel Baucells, Fabrizio Colella, Gert Cornelissen, Jordi Galí, Marc Le Menestrel and Gaël Le Mens.
Many thanks to Laura Agustí, Marta Araque and Mariona Novoa for their superb administrative support.

I am also very grateful to Pablo López-Aguilar for spending many hours teaching me how to program in Java and z-Tree.

I would like to express my gratitude to the Spanish Ministry of Education for the financial support, to the department of economics at Tilburg University for hosting me several times and especially to the department of economics and business at Universitat Pompeu Fabra.

I thank my wife and my parents for their encouragement and unconditional support. I dedicate this thesis to them.

Giovanni Giusti, September 2014

SUMMARY

This thesis is composed of three essays. In the first essay (joint with Janet Jiang and Xiping Xu) we study asset price bubbles in a laboratory experiment. By introducing interest payments on cash we separate the effect of trading opportunity cost from the role of asset fundamental value trend on bubble formation. Results show that the fundamental value trend plays a more critical role. In the second essay (joint with Charles Noussair and Hans-Joachim Voth) we study in a laboratory setting the importance of several historical institutional features that characterized the South Sea bubble. Our main finding is that the debt-equity swap was the single biggest contributor for the stock price explosion. In the third essay we study in an experiment how different dynamics of piece rate monetary incentives affect participants' effort provision. Our main finding shows that a decrease in piece rate following an increase has detrimental effects for participants' effort provision.

RESUM

Aquesta tesi conté tres assaigs. En el primer assaig (conjunt amb Janet Jiang i Xiping Xu) estudiem la bombolla de preus d'actius en un experiment de laboratori. Introduint els pagaments d'interessos en efectiu separem l'efecte del cost d'oportunitat de comerciar de la trajectòria del valor fonamental de l'actiu. Els resultats mostren que la trajectòria del valor fonamental juga un paper molt crític. En el segon assaig (conjunt amb Charles Noussair i Hans-Joachim Voth) estudiem en un laboratori la importància de diverses característiques institucionals històriques que van caracteritzar la "bombolla dels mars del sud". El nostre principal descobriment és que el "swap" de deute per accions és l'únic gran contribuïdor per l'explosió del preu de les accions. En el tercer assaig estudiem en un experiment com diferents dinàmiques de preu fet com a incentiu monetari afecten la prestació d'esforç dels participants. El nostre principal descobriment indica que una disminució del preu fet després d'un increment té efectes perjudicials per a la prestació d'esforç dels participants.

PREFACE

This thesis consists of three self-contained chapters. The first chapter contributes to a large literature studying bubble formation in experimental asset markets following the seminal paper by Smith, Suchanek and Williams (1988).

In Smith, Suchanek and Williams (1988), subjects trade a finitely lived asset that pays a random dividend with known distribution in each period. The dividend payment is the only source of intrinsic value of the asset.

Theory assuming full rationality predicts that bubbles should not occur in this setting. However, the authors find that the trading price frequently exceeds the fundamental value, which provides strong evidence against full rationality.

In our experimental asset market we deviate from Smith, Suchanek and Williams (1988) by the introduction of interest on cash: interest payments interpreted as interest earnings of a savings account, or interest charges interpreted as banking fees.

By exploiting the flexibility created by the presence of interest payments we investigate two aspects that could be responsible for the “irrational trading” in the Smith, Suchanek and Williams (1988) setting: low opportunity cost of speculation and confusion about the fundamental value process.

Our results suggest two main conclusions. The first is that paying positive interest on cash and consequently increasing the opportunity cost for treading is ineffective in reducing bubbles. The second is that the fundamental value generating process plays a critical role in the formation of asset bubbles in the laboratory. In particular, bubbles tend to occur whenever there is a conflict between the sign of the time trend of the fundamental value and the sign of the expected dividend payment.

In the second chapter we use a laboratory experiment to study one of the major asset bubble episodes in the history, the South Sea bubble.

In order to examine what factors might have caused the asset price bubble to become very large, we reproduce some of the specific institutional features investors in the South Sea Company faced in 1720. Several factors have been proposed as potentially contributing to one of the greatest periods of asset overvaluation in history: an intricate debt-for-equity swap, deferred payment for these shares, and the possibility of default on the deferred payments. The results of the experiment suggest that the company's attempt to exchange its shares for government debt was the single biggest contributor to the stock price explosion, because of the manner in which the swap affected fundamental value. Issuing new shares with only partial payments required, in conjunction with the debt-equity swap, also had a significant effect on the size of the bubble. Limited contract enforcement, on the other hand, does not appear to have contributed significantly.

The third chapter contributes to the extensive literature studying the effect of changes in monetary incentives on workers' effort provision.

Specifically, we examine how effort provision of participants varies in a three-period laboratory setting in which different performance based compensation schemes are implemented for two differentiated tasks.

The first task, consisting of downward counting from a big number while each time subtracting a fixed quantity, is considered relatively more challenging, while the second task consisting of counting the number of A's in a paragraph is considered relatively less challenging.

Furthermore, tasks differed on subjects' individual evaluation on an interest/enjoyment scale. Our findings are as follows. First, performance contingent incentives affect participants' effort provision. Second, the effect of monetary incentives is task dependent and it is much stronger for the less challenging task. Third, a unique increase in the amount of piece rate leads to an increase in performance only in the less challenging task. Fourth, a decrease in piece rate incentive negatively affects subjects' performance on both tasks, but only provided that the decrease follows a previous increase.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTEREST ON CASH, FUNDAMENTAL VALUE PROCESS AND BUBBLE FORMATION: AN EXPERIMENTAL STUDY

(With Janet Jiang and Yiping Xu)

1.1 Introduction

Bubbles refer to the phenomena associated with dramatic increases in asset prices exceeding the asset's fundamental value. Bubbles occur when asset owners believe that they can resell the asset at an even higher price in the future. There are, broadly speaking, two approaches to modelling bubbles (see Brunnermeier 2009, and Scherbina 2013, for two surveys of the literature on asset bubbles). The first approach assumes full rationality. One result of this approach is that if all information is common knowledge, bubbles must be infinitely-lived and grow at the same rate as the discount rate[1].
The second approach deviates from perfect rationality and assumes that at least some traders are behavioral.

There is a large literature of experimental studies on asset bubbles following the seminal paper by Smith, Suchanek and Williams (1988, hereafter SSW). The experimental approach constitutes a good complement to research using field data. The advantages include cleaner control of trading environments and clearer definition of the fundamental value. In SSW, subjects trade a single asset in a simple experimental asset market environment. The asset has a finite lifetime and pays a random dividend in each period. The dividend payment and 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. Theory assuming full rationality predicts that bubbles should

[1] The result that bubbles cannot occur in a finite horizon model is derived by backward induction. Since a bubble cannot grow from the last trading period, there cannot be a bubble of this size in the second-last period, and so on.

not occur in the SSW setting. However, SSW find that the trading price frequently exceeds the fundamental value, which provides strong evidence against full rationality.[2] SSW conjecture that bubbles on the experimental asset market are caused by the lack of common knowledge of rationality leading to speculation. Lei, Noussair and Plott (2001) reject the conjecture after observing bubbles even when resale opportunities are removed, and suggest real irrationality as the source of bubbles. In addition, Lei, Noussair and Plott (2001) find that removing speculative opportunities and simultaneously adding a commodity market greatly reduce trading volumes and price bubbles on the asset market. As a result, they raise the active-participation hypothesis that much of the trading activity that accompanies bubble formation is due to the lack of an alternative activity during the experiment. More generally, Lei, Noussair and Plott (2001) suggest that bubble formation in experimental asset markets could have origins in aspects of the methodology of the experiment.

In this paper, we respond to Lei, Noussair and Plott's (2001) call to explore the aspects of the SSW design that may have contributed to irrational trading on experimental asset markets. The framework that we choose deviates from the SSW design by the introduction of interest on cash: interest payments interpreted as interest earnings of a savings account, or interest charges interpreted as banking fees.[3]
The design allows us to investigate two possible factors contributing to irrational trading in the SSW experimental asset markets. First, the absence of interest payments on cash implies a low opportunity cost of speculation on the asset market, which may have boosted active trading and bubble formation. This is related to the point raised in Lei, Noussair and Plott (2001), that subjects tend to trade (irrationally) due to boredom and the lack of alternative activities during the experiment. We study the case where cash earns positive interest. Although it does not provide alternative trading activities as in Lei, Noussair and Plott (2001), it raises the opportunity cost of speculation on the asset market and may induce more prudent trading and a lower trading volume.

[2] In the experimental literature, the fundamental value is usually calculated under the assumption of risk neutrality. If agents are risk averse, the fundamental value should be lower than that implied by risk neutrality, which makes the observation of pricing bubbles (relative to risk-neutral fundamental values) even more striking.
[3] Bostian and Holt (2009) and Fischbacher, Hens and Zeisberger (2013) also feature interest on cash. We will discuss the differences between their work and ours in section 1.2, where we review related literature.

Second, researchers have suspected that the fundamental value generating process featured in the SSW design can induce to confusion and irrational trading. The introduction of interest on cash allows us to alter the fundamental value process in more flexible ways. Exploring a richer set of fundamental value processes helps us to identify the source of confusion in the SSW design. Note that, without interest payments, two aspects of the fundamental value generating process -- the sign of the expected dividend payment and the time trend of the fundamental value -- are tied together: positive, zero, or negative expected dividend payments give rise to decreasing, flat, or increasing fundamental values, respectively. The introduction of interest breaks the connection and allows for greater flexibility in terms of more possible combinations of the sign of the dividend payment and the sign of the time slope of the fundamental value.

To investigate how the two aspects of the SSW design -- low opportunity cost of speculation and confusion about the fundamental value -- may have contributed to irrational trading on experimental asset markets, we design three treatments with interest on cash characterized by different combinations of the sign of dividend payments and the sign of the time slope of the fundamental value. In the first two treatments (treatments F and R), cash earns positive interest, and stocks pay positive expected dividends. The difference lies in the dynamics of the fundamental value: the fundamental value decreases in treatment F and increases in treatment R. In the third treatment (treatment N), banking fees are charged on cash holdings, the expected dividend is negative (interpreted as carrying costs), and the fundamental value decreases over time.

There is substantial overpricing in treatment F, very weak mispricing in treatment R, and, on average, slight underpricing in treatment N. There is no significant difference in terms of trading volume among the three treatments.

Treatment F shares similarities with a standard SSW design in that both designs have positive dividend payments and decreasing fundamental values; the main difference is that cash earns positive (no) interest in treatment F (SSW). Since treatment F involves significant overpricing, as with the SSW design, we can conclude that paying interest on cash is not sufficient to suppress price bubbles through the reducing-active-participation channel. In addition, a comparison between treatments F and R suggests that active participation is not the reason why bubbles appear in treatment F, because the two treatments share similar levels of trading intensity. A more likely reason for bubble formation is confusion about the fundamental value of the asset. In particular, the results from our study, together with those from other papers (which we will discuss in more detail in the next section), suggest that mispricing tends to occur whenever there is a conflict between the sign of dividend payments and the sign of the time slope of the fundamental value: overpricing tends to occur with positive dividend payments and decreasing fundamental values, and underpricing with negative dividends and increasing fundamental values.

The rest of the chapter is organized as follows. We discuss related literature in section 1.2, describe the experimental design in section 1.3, analyze the experimental results in section 1.4, and conclude in section 1.5.

1.2 Related literature

Following the seminal paper by SSW, the experimental literature on asset bubbles has largely followed two directions (see Palan 2013 for a detailed survey of the literature).

The first keeps the same fundamental value specification as in SSW and tries to find measures to reduce bubbles. King et al. (1993) examine the effect of allowing for short sales, using non-student subjects, transaction fees, equal endowment and price-change limits; these measures are ineffective in eliminating bubbles.[4]

Fisher and Kelly (2000) introduce two simultaneous asset markets and find that bubbles exist in both markets. Fischbacher, Hens and Zeisberger (2013) and our paper study the case with interest payments on cash; both studies suggest that paying interest on cash in itself is ineffective in eliminating bubbles. Rigid measures to curb trading are more effective in reducing or eliminating bubbles; such measures include removing speculative opportunities and diverting subjects' attention from the asset market to a commodity market (Lei, Noussair and Plott 2001), reducing liquidity or controlling for the cash/asset ratio (Caginalp, Porter and Smith 1998, 2001; Fischbacher, Hens and Zeisberger 2013; Kirchler, Huber and Stöckl 2012), and imposing holding caps (Lugovskyy et al. 2012). Different trading mechanisms have also been studied. Van Boening, Williams and LaMaster (1993) find that bubbles continue to occur in call

[4] Ackert et al. (2006) and Haruvy and Noussair (2006) also study the effect of short sales.

markets. Lugovskyy, Puzello and Tucker (2011) show that the Tâtonnement mechanism is effective in reducing bubbles. Cheung and Palan (2012) find that trading teams reduce the bubble phenomenon compared to individual traders.

The second direction of the experimental research on asset bubbles contends that the fundamental value generating process in the SSW is conducive to misunderstanding and bubble formation. Most studies use a framework without interest on cash. Smith, van Boening and Wellford (2000), Noussair, Robin and Ruffieux (2001), Huber, Kirchler and Stöckl (2012), and Kirchler, Huber and Stöckl (2012) study the case with flat fundamental values assuming a zero expected dividend payment, and find that overpricing is greatly reduced. Huber, Kirchler and Stöckl (2012) also study the case with increasing fundamental values by assuming negative dividend payments and find underpricing in this situation.[5]

Bostian and Holt (2009) conduct a classroom experiment using a framework with interest payments on cash. They study a single regime which features a flat fundamental value induced by equating the buyout value to the ratio of the expected dividend payment over the interest rate, and observe frequent occurrences of bubbles. Our study exploits the flexibility created by the introduction of interest on cash and includes three different fundamental value processes.[6]

There are two explanations of how the fundamental value process affects trading behavior, both focusing on the effect of the time trend of the fundamental value. First, as pointed out by Smith (2010) and Oechssler (2010), subjects may find it hard to comprehend the decreasing fundamental value in the SSW design, because asset prices tend to increase or stay constant in the long run in real life. Huber, Kirchler and Stöckl

[5] Noussair and Powell (2010) and Breaban and Noussair (2014) examine environments where the fundamental values experience different time trends during the trading game. Noussair and Powell (2010) conduct two sets of experiments. In the "peak" treatment, fundamentals first rise and then fall, while in the "valley" treatment fundamentals first fall and then recover. They find that bubbles still occur in both treatments, but in smaller magnitudes in the peak treatment. Breaban and Noussair (2014) study markets in which a trend in fundamentals sets in after an interval of constant value. They find that prices tend to track fundamentals more closely when the trend is decreasing than when it is increasing. Breaban and Noussair (2014) conclude that the contrast between their results and those from previous studies indicate that the timing of the onset of a trend in fundamentals is an important feature influencing how the trend affects the price discovery process.

[6] Fischbacher, Hens and Zeisberger (2013) also have interest on cash. In their design, the interest rate changes in response to prices, making it difficult to define the fundamental value. They do not focus on the effect of fundamental value dynamics.

(2012) provide a second explanation, proposing that anchoring on information generated by the trading process drives under-reaction and, in turn, mispricing on the experimental asset market. According to this explanation, decreasing fundamental values would give rise to overpricing, increasing fundamentals would lead to under-pricing, and flat fundamentals would involve little mispricing. Neither of the two explanations is fully compatible with the existing studies. In particular, the results in the two studies with interest on cash -- Bostian and Holt (2009) and our study -- provide evidence against either explanation. According to Smith (2010) and Oechssler (2010), there should be great overpricing in our treatment N with decreasing fundamentals and no mispricing in Bostian and Holt (2009) with a flat fundamental. However, there is, on average, slight under-pricing in our treatment N and frequent bubbles in Bostian and Holt (2009). According to the explanation suggested by Huber, Kirchler and Stöckl (2012), there should be under-pricing or negative bubbles in our treatment R with increasing fundamentals, positive bubbles in our third treatment with decreasing fundamentals and no mispricing in Bostian and Holt (2009) with a constant fundamental value; the prediction is inconsistent with the experimental results.

We propose a third explanation that suggests both the time trend of the fundamental value and the sign of the expected dividend payment affect mispricing. In particular, mispricing tends to occur whenever there is a conflict between the sign of dividend payments and the sign of the time slope of the fundamental value: overpricing tends to occur with positive dividend payments and decreasing fundamental values, and under-pricing tends to occur with negative dividends and increasing fundamental values. This new explanation is consistent with all existing studies, including our own.

There is also a strand of studies that keeps the same fundamental value generating process as in SSW, but adopts various measures to help subjects understand the process. One measure is to allow subjects to repeat the experiment. Many studies, including SSW, King et al. (1993), Van Boening, Williams and LaMaster (1993), Dufwenberg, Lindqvist and Moore (2005), and Hussam, Porter and Smith (2008), find that past experience with the same game significantly reduces the magnitude of bubbles.[7]

[7] The effect of experience only applies if subjects repeat the same game. The effect will disappear if there is a large shock to the environment, such as liquidity and dividend uncertainty as in Hussam, Porter and Smith (2008), and reshuffling of subjects and admission of new subjects as in Xie and Zhang (2012).

Lei and Vesely (2009) include a pre-trading phase where subjects hold the asset and experience dividend flows. Kirchler, Huber and Stöckl (2012) change the context from "stocks" to "stocks of a depletable gold mine." Huber, Kirchler and Stöckl (2012) display the fundamental value on the trading screen. These salient measures are found to be effective in eradicating bubbles. Porter and Smith (1995) find that the existence of a futures market, in which contracts are realized at the halfway point of the trading horizon, helps to reduce but does not remove asset bubbles. A follow-up study by Noussair and Tucker (2006) finds that the futures markets can eliminate bubbles if the contracts are such that there is one maturing in each period of the life of the asset.

Finally, there are also a few papers that study asset bubbles with indefinite horizons. For example, Camerer and Weigelt (1993) study asset markets where subjects trade stochastically lived assets that pay a dividend each period and live from period to period with a known probability; they find that asset prices converge slowly to the fundamental value. On the other hand, in a classroom experiment, in which an asset with a constant fundamental value was traded in an indefinite horizon, Ball and Holt (1998) find systematic overpricing. In Hens and Steude (2009), the dividend process follows a random walk with a positive drift, which implies that the fundamental value depends on the dividend realization in the current period (the paper does not investigate the problem of mispricing). Kose (2013) finds that concerns about bankruptcy risk (the asset becomes worthless once the game ends randomly) cause under-pricing irrespective of the time trend of the fundamental value.

Crockett and Duffy (2013) and Fenig, Mileva and Petersen (2014) consider general-equilibrium economies. Crockett and Duffy (2013) find that in an environment where subjects are induced to adjust share holding in order to smooth consumption, assets trade at a discount relative to the risk-neutral fundamental price. In Fenig, Mileva and Petersen (2014), subjects are induced to maximize their utility from consumption and leisure, and at the same time engage in speculative activities in the asset market; they find that asset prices consistently grow above the fundamental value and do not decline significantly with learning.

1.3 Experimental design

The major departure from the SSW design is the introduction of interest payments or charges on cash holdings. Within this framework, we can study whether positive interest payments on cash, which increase the opportunity cost of asset-market speculation, will reduce speculation and bubbles on the asset market. We can also investigate the effect of different fundamental value generating processes, taking advantage of the flexibility created by the introduction of interest on cash.

Shares have a finite life of T periods. Each share pays a random dividend at the end of each period from time *1* to T, plus a fixed buyout value, K, at the end of period T. The distribution of the dividend is *iid* over time. The expected value of the dividend is fixed at d. If $d < 0$, we interpret it as carrying costs. In each period, cash earns interest or bears charges at the net rate of r. When $r > 0$, we interpret that cash is parked in an interest-bearing savings account. When $r < 0$, we say cash is placed in a banking account that charges banking fees. Subjects can use money from their savings/banking account to purchase shares. Revenues from share sales, carrying costs of shares and banking fees are automatically deposited into or deducted from the savings/banking account. Following the usual practice in the literature, we define the fundamental value as the holding value for a risk-neutral agent (the fundamental value for a risk-averse agent is lower). The fundamental value of the asset at the beginning of period t is calculated as the net present value of all remaining dividend payments and the buyout value at the end of T, i.e.,

$$FV_t = d\left[\sum_{\tau=1}^{T-t+1}(1+r)^{-\tau}\right] + K(1+r)^{-(T-t+1)}$$

$$= \begin{cases} d(T-t+1) + K \; if \; r = 0, \\ \frac{d}{r} + \left(K - \frac{d}{r}\right)(1+r)^{-(T-t+1)} \; if \; r \neq 0. \end{cases}$$

The time trend of the fundamental value is therefore given by:

$$\frac{d(FV_t)}{dt} = \begin{cases} -d \; if \; r = 0, \\ \left(K - \frac{d}{r}\right)[ln(1 + r)](1 + r)^{-(T-t+1)} \; if \; r \neq 0. \end{cases}$$

Note that in the absence of r (i.e. $r = 0$), the time trend of the fundamental value is fully determined by the sign of the dividend payment (negative dividend payments can be interpreted as carrying costs). If $d > 0$ as in the SSW design, the fundamental value must decrease over time. To generate a flat fundamental value, d must be equal to 0 as in Smith, van Boening and Wellford (2000), Noussair, Robin and Ruffieux (2001), and Kirchler, Huber and Stöckl (2012). To induce an increasing fundamental value, d must be negative as in Huber, Kirchler and Stöckl (2012). The introduction of r allows for more flexibility. In particular, it is possible to have increasing fundamental values with $d > 0$ (for example, by setting $r > 0$ and $K > d/r$) and decreasing fundamental values with $d < 0$ (for example, by setting $r < 0$ and $K > d/r$).

Using the above framework, we investigate two possible factors that may contribute to irrational trading. The first is the "active-participation hypothesis" raised by Lei, Noussair and Plott (2001). We would like to study whether paying positive interest on cash holdings, which increases the opportunity cost of speculation on the asset market, will reduce trading and overpricing of the asset. The second is the fundamental value generating process. As discussed earlier, existing studies (except our paper) focus on the effect of the time slope of the fundamental value, but the time slope itself cannot explain the results in all studies. We conjecture that the way to generate the time slope is also important. Taking advantage of the flexibility created by the introduction of interest on cash, we run three treatments with different combinations of the interest rate, dividend payment, the buyout value, and the time slope of the fundamental value. We run a total of 19 sessions. Detailed information for each session is listed in Table 1.

Treatment	Session	Subjects	Trading periods	Dividend	Initial shares	Initial cash	Interest rate (r)	Buyout (K)	FV_1	FV_{15}	FV_{15}/FV_1	CA_1
F	F1	10	15	(0,8,28,60)	4	2000	0.1	72	200	87	0.44	2.50
	F2	10	15	(0,8,28,60)	4	2000	0.1	72	200	87	0.44	2.50
	F3	9	15	(0,8,28,60)	4	1022	0.1	0	183	22	0.12	1.40
	F4	10	15	(0,8,28,60)	4	1054	0.1	24	188	44	0.23	1.40
	F5	10	15	(0,8,28,60)	4	827	0.15	60	148	73	0.49	1.40
	F6	10	15	(0,8,28,60)	4	591	0.15	60	148	73	0.49	1.00
R	R1	10	15	(0,8,28,60)	4	2000	0.1	720	355	676	1.91	1.41
	R2	10	15	(0,8,28,60)	4	2000	0.1	720	355	676	1.91	1.41
	R3	10	15	(0,8,28,60)	4	2000	0.1	720	355	676	1.91	1.41
	R4	10	15	(0,8,28,60)	4	4259	0.1	720	355	676	1.91	3.00
	R5	10	15	(0,8,28,60)	4	2000	0.15	720	229	647	2.83	2.19
	R6	10	15	(0,8,28,60)	4	2000	0.15	720	229	647	2.83	2.19
	R7	10	12	(0,8,28,60)	4	745	0.15	300	186	282	1.52	1.00
N	N1	9	15	(0,-8,-28,-60)	4	50000	- 0.1	500	1503	529	0.35	8.32
	N2	10	15	(0,-8,-28,-60)	4	50000	- 0.1	500	1503	529	0.35	8.32
	N3	9	15	(0,-8,-28,-60)	4	50000	- 0.1	500	1503	529	0.35	8.32
	N4	10	15	(0,-8,-28,-60)	4	40000	- 0.1	500	1503	529	0.35	6.65
	N5	9	15	(0,-8,-28,-60)	4	40000	- 0.1	500	1503	529	0.35	6.65
	N6	10	15	(0,-8,-28,-60)	4	40000	- 0.1	500	1503	529	0.35	6.65

Table 1*:* Parameters used in the experiment

In the first treatment, cash earns a positive interest payment with $r = 10\%$ or 15%[8]

The dividend payment has four possible realizations, *0, 8, 28,* and *60*, with equal probabilities, which implies a positive expected dividend payment with $d = 24$.[9]

The buyout value K is set to be less than d/r to induce a decreasing fundamental value. This treatment is labeled "F" to reflect the falling fundamental value. We run six experimental sessions (sessions F1-6) of this treatment.

The second treatment is similar to the first treatment in terms of interest and dividend payments, but with $K > d/r$, which implies increasing fundamentals. We label the second treatment "R" to capture the rising fundamental values. We run seven treatment-R sessions (R1-7).

The third treatment has interest charges on cash holdings with $r = -10\%$ and negative dividend payments interpreted as carrying costs with $d = -24$ (the dividend is equal to *0, -8 , -28,* or *-60* with equal probabilities); the buyout value, *K,* is set at 500; and the

[8] The interest rate of 10% or 15% seems to be unrealistically high. Because subjects play with small stakes in the experiment, we set the interest rate at conspicuously high levels to induce meaningful responses from subjects.

[9] Some treatments in SSW feature the same dividend distribution.

fundamental value decreases over time. We label the treatment "N" to reflect the negative dividend and interest payment. There are six sessions (N1-6) of this treatment.

Treatment F differs from a standard SSW design in that cash earns positive interest (both designs have positive dividend payments and decreasing fundamental values). Therefore, we can use the result from treatment F to evaluate whether increasing the opportunity cost of asset trading by paying positive interest on cash is effective in eliminating bubbles through the reducing-active-participation channel. In addition, if bubbles are observed only in treatment F, but not in treatment R, we can infer that the main reason for the different results from the two treatments is due to the different fundamental value process.[10]

To examine the effect of the fundamental value generating process, we compare the results from all three treatments, together with the results from other papers that study the effect of alternative fundamental dynamics, including Smith, van Boening and Wellford (2000), Noussair, Robin and Ruffieux (2001), Bostian and Holt (2009), Huber, Kirchler and Stöckl (2012), and Kirchler, Huber and Stöckl (2012). Our purpose is to find an explanation to reconcile all existing studies on the effect of fundamental value generating processes.

The program used to conduct the experiment is written in z-Tree (Fischbacher, 2007). See the Appendix for the experimental instructions. There are 9 or 10 subjects participating in each session, trading a single asset called "shares." Communication among subjects is prohibited during the experiment. The number of trading periods, T, is 15, except for session R7, which has 12 trading periods. Each trading period lasts for 150 seconds. Subjects are given the opportunity to practice with the trading interface. There is also a training period during which subjects familiarize themselves with the task that they will perform. Each subject starts the first formal trading period with the same endowment of shares and cash. The share endowment for each subject is 4. The amount of cash endowment is chosen to control for the initial cash/asset ratio, which

[10] The result that substantial overpricing occurs only in treatment R, but not in treatment F, is unlikely to occur.

ranges from 1 to 2.5 in treatment F, 1 to 3 in treatment R, and 6.65 to 8.32 in treatment N.[11]

Following the usual practice in the literature, we provide subjects with a table to list the holding value of a share in terms of cash. The trading mechanism is a continuous double auction with open order books. Subjects initiate a transaction by posting offers to buy (bids) and offers to sell (offers). Each offer is for the transaction of one share, but subjects can post multiple offers to buy or sell. Active orders to buy and orders to sell are ranked in two separate columns, with the best available offers at the bottom of the lists. Subjects execute a trade by selecting the best order and press the "buy" or "sell" button located at the bottom of the order book. To facilitate the comparison between our results and those from other papers that study the effect of interest payments and the fundamental value dynamics, we adopt the same design to ban short sales of shares and borrowing money to buy shares. The sessions were conducted from October 2011 to March 2012 at two universities: Universitat Pompeu Fabra (UPF), Barcelona, and University of International Business and Economics (UIBE), Beijing. Each session of experiment lasts for about 90 minutes. The average earning is 13 euros at UPF and 100 RMB at UIBE.

1.4 Experimental results

We document the experimental results in Figures 1-3 and Tables 2-7. Before discussing the experimental results, we first describe the information in these figures and tables.

Figures 1-3 plot the time series of the fundamental value (FV_t), the median trading price (P_t) and the trading volume (N_t) for each of the experimental session.[12]

[11] Caginalp, Porter and Smith (1998, 2001) suggest that the effect of the cash/asset ratio is stronger in early periods of the experiment.

[12] We use the median trading price instead of the average trading price because the former is less affected by errors made by subjects while posting offers.

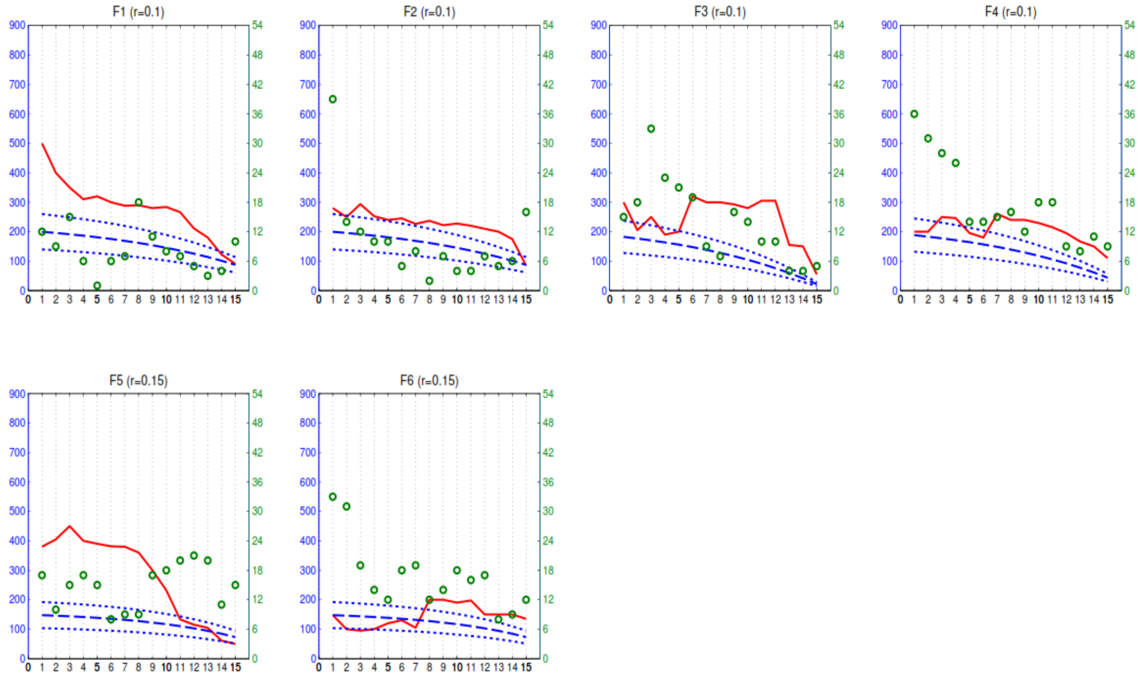

Figure 1: Experimental results - treatment F

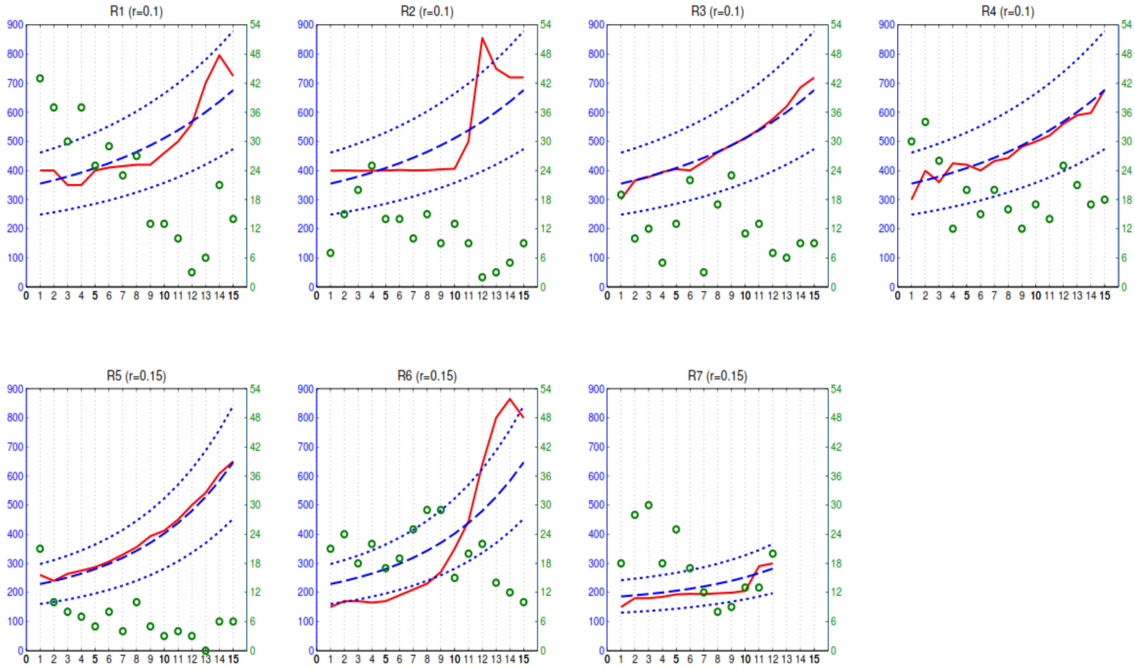

Figure 2: Experimental results - treatment R

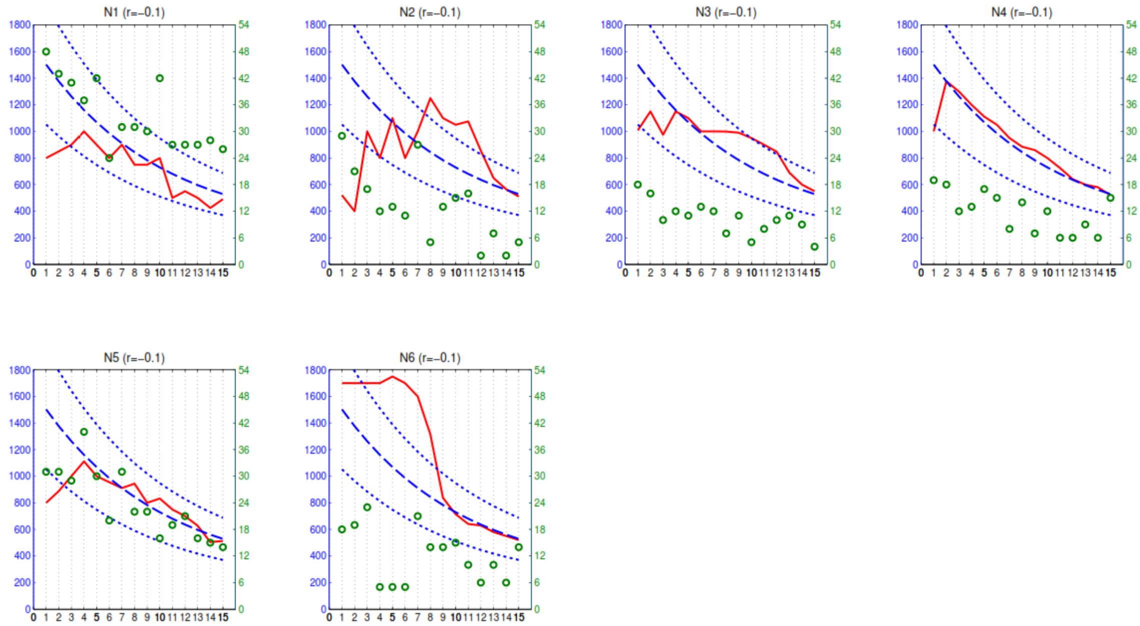

Figure 3: Experimental results - treatment N

The six sessions with treatment F are graphed in Figure 1 (F1 –F6). Figure 2 reports the seven treatment-R sessions (R1 – R7). The six treatment-N sessions are represented in Figure 3 (N1 – N6). The horizontal axis indicates the trading period running from 1 to 15. Prices are depicted along the left vertical axis: the solid line is the path of (P_t), the dashed line represents (FV_t), the upper dotted line indicates $(1 + 30\%)\ x\ FV_t$, and the lower dotted line represents $(1 - 30\%)\ x\ FV_t$. The two dashed lines serve as reference lines to visualize the extent of mispricing. The trading volume is graphed against the right vertical axis in circles.

Session	RAD	RD	PA	ST
R1	0.10	0.01	0.47	8.28
R2	0.14	0.04	0.81	4.25
R3	0.03	0.00	0.21	4.48
R4	0.04	-0.02	0.18	7.43
R5	0.04	0.04	0.08	2.50
R6	0.32	-0.01	1.04	7.43
R7	0.11	-0.08	0.32	5.28
Treatment R Average	**0.11**	**0.00**	**0.44**	**5.66**
F1	0.81	0.81	1.92	3.05
F2	0.45	0.45	0.67	3.73
F3	1.03	1.03	1.71	5.78
F4	0.57	0.57	0.77	6.63
F5	1.33	1.28	2.72	5.55
F6	0.40	0.20	1.12	6.30
Treatment F Average	**0.77**	**0.72**	**1.49**	**5.17**
N1	0.21	-0.20	0.85	14.00
N2	0.34	-0.07	1.53	4.88
N3	0.17	0.02	0.79	4.36
N4	0.07	0.00	0.64	4.43
N5	0.15	-0.09	0.89	9.92
N6	0.31	0.29	0.83	4.63
Treatment N Average	**0.21**	**-0.01**	**0.92**	**7.03**

Table 2: Statistics for treading behavior

Table 2 provides four statistics to quantify the trading behavior. We provide the statistics for each individual session, and the treatment statistics (in bold face) averaged across sessions of the same treatment. There are three statistics to measure price deviations: relative absolute deviation (RAD), relative deviation (RD) and price amplitude (PA). The fourth statistic, share turnover (ST), measures trading intensity. The four statistics are calculated as follows.

Let $\overline{FV} = (\sum_{t=1}^{T} FV_t)/T$ be the average lifetime fundamental value. Denote the number of outstanding shares as N_O, which is equal to 40 in sessions with 10 subjects and 36 in sessions with 9 subjects. The relative absolute deviation

$RAD = \left(\sum_{t=1}^{T} \frac{|P_t - FV_t|}{\overline{FV}}\right)/T$ measures the average level of mispricing relative to the average lifetime fundamental value of the asset. The relative deviation $RD = \left(\sum_{t=1}^{T} \frac{P_t - FV_t}{\overline{FV}}\right)/T$ measures the extent of over or under-valuation. The price amplitude $PA = \max_{1\leq t\leq T}[(P_t - FV_t)/\overline{FV}] - \min_{1\leq t\leq T}[(P_t - FV_t)/\overline{FV}]$ measures the overall size of mispricing. The share turnover is calculated as $ST = \sum_{t=1}^{T} \frac{N_t}{N_o}$.

Note that we use the average lifetime fundamental value, $\overline{FV}$, to calculate the three measures of price deviation. As discussed in Stöckl, Huber and Kirchler (2010), it is more appropriate to use $\overline{FV}$ (than FV_1, as in many studies) for comparison among different experimental settings, especially among treatments with different time paths of the fundamental value.

		RAD	RD	PA	ST
F		0.766	0.723	1.486	5.171
R		0.111	-0.003	0.445	5.661
	Z-value	3.000	3.000	2.429	-0.429
	p-value	0%	0%	1%	67%
	Sample size	13	13	13	13
F		0.776	0.723	1.486	5.171
N		0.209	-0.009	0.920	7.034
	Z-value	2.882	2.722	1.121	-0.320
	p-value	0%	0%	26%	75%
	Sample size	12	12	12	12
R		0.111	-0.003	0.445	5.661
N		0.209	-0.009	0.920	7.034
	Z-value	-1.857	0.714	-2.000	-0.429
	p-value	6%	48%	5%	67%
	Sample size	13	13	13	13

Table 3: Treatment statistics

Table 3 reports the results from two-tailed Mann-Whitney tests that compare the trading patterns of the different treatments. Each session is counted as one observation of the treatment: there are six observations for treatment F, seven observations for treatment R, and six observations for treatment N. The test is performed for each of the four trading statistics and for each pair of the three treatments. A total of 12 tests (four statistics and three pairs) are conducted. For each test, we list the average treatment statistic for the pair of treatments being compared, the Z-statistic, the p-value and the combined sample size of the pair of treatments.

Treatment	Number of Sessions	Number of Sessions with Bubbles		
		30% rule	40% rule	50% rule
F	6	6	6	5
R	7	0	0	0
N	6	2	1	0

Table 4: Incidences of bubbles

In Table 4, we identify the incidence of bubbles. We use a commonly adopted rule in the literature: we say that a bubble occurs in a session if the median transaction price exceeds the fundamental value by at least *x%* for more than five consecutive periods (see, for example, Noussair, Robin and Ruffieux 2001). Table 4 lists the number of bubbly sessions for each of the three treatments for x = 30, 40, and 50.

Session	Transactions	Offers
F1	57	74
F2	66	90
F3	62	82
F4	52	84
F5	66	80
F6	64	94
Treatment F average	**61**	**84**
R1	92	100
R1	92	100
R2	100	100
R3	100	100
R4	100	100
R5	100	100
R6	72	94
R7	98	100
Treatment R average	**94**	**99**
N1	85	92
N2	71	98
N3	91	98
N4	98	100
N5	88	95
N6	86	87
Treatment N average	**87**	**95**

Table 5: Percentage of fundamental transactions and offers -- session and treatment

		Median (% fundamental transactions)	Median (% fundamental offers)
F		61	84
R		94	99
	Z-value	-3.046	-3.085
	p -value	0%	0%
	Sample size	13	13
F		61	84
N		87	95
	Z-value	-2.887	-2.406
	p -value	0%	2%
	Sample size	12	12
R		94	99
N		87	95
	Z-value	2.103	2.178
	p -value	4%	3%
	Sample size	13	13

Table 6: Mann-Whitney test of fundamental transactions and offers

Finally, Tables 5 and 6 provide information about individual trading behavior. In particular, we check the extent of "fundamental trading" for each subject. We say a transaction or offer is "fundamental" if the price is ≤ (1+30%) *x FV* for share purchases, and ≥ (1-30%) *x FV* for share sales. We then measure the extent to which an individual is a fundamental trader by the percentage of fundamental transactions and offers that the individual engages in across the session. We provide session-level and treatment-level statistics in Table 5. The session statistic is calculated as the median of the individual statistics for subjects who participated in that session. The treatment-level statistic (in bold face) is the average of session statistics for sessions belonging to the same treatment. Table 6 shows the results from the Mann-Whitney test of the three treatments using the session-level statistics as observations.

1.5 Description of experimental results

Now we describe the experimental results. It is clear from Figure 1 that the median trading price frequently exceeds the fundamental value (and the (1+30%) *x FV* line) in sessions with treatment F. In contrast, in sessions with treatment R, the median trading price closely tracks the fundamental value most of the time. For treatment N, except for two sessions N2 and N6, the median trading price does not deviate substantially from the fundamental value: the line for the median trading price lies within the (1±30%) *x FV* band most of the time.

From Table 2, which provides the trading statistics, one can see that treatments R and N have, on average, mild underpricing (relative to the risk-neutral fundamental value): the average treatment RD is -0.3% for treatment R and -0.9% for treatment N. In contrast, treatment F exhibits substantial overpricing, with the average treatment RD being high at 72.3%. In terms of general mispricing, treatment R involves the smallest deviation from the fundamental value (the treatment average RAD is 11.1%), followed by treatment N (the treatment average RAD is 20.9%). Treatment F exhibits substantial mispricing: the treatment average RAD is 76.6%. The price amplitude is very high in treatment F averaged at 148.6%, compared with 44.5% for treatment R and 92.0% for treatment N. The three treatments have a comparable trading intensity: the share turnover is 5.171 for treatment F, 5.661 for treatment R and somewhat higher at 7.034 for treatment N.

Table 3 shows the results from Mann-Whitney tests on RAD, RD, PA and ST. The tests suggest that treatment F generates statistically higher mispricing, mainly in the form of overpricing, than the other two treatments. In terms of RAD, which measures overall mispricing, the test between treatments F and R has a Z-value of 3, and the test between treatments F and N gives a Z-value of 2.882. In terms of RD, which measures overpricing, the test between treatments F and R has a Z-value of 3, and the test between treatments F and N has a Z-value of 2.722. All four tests have a p-value of 0%[13].

The tests on PA suggest that treatment R involves much smaller price fluctuations than the other two treatments. The comparison between treatments F and N is not statistically different, with a p-value of 26%. Finally, in terms of trading intensity, measured by share turnover, the three treatments are not statistically different: the tests between treatments F and R, F and N and R and N, have a p-value of 67%, 75% and 67%, respectively.

In terms of the number of bubbles (see Table 4), we find that all treatment-F sessions have bubbles if we define a bubble as the situation where the median transaction price

[13] Note that the RADs for all seven treatment-R sessions are universally lower than those for all six treatment-F sessions. The same pattern applies to RD (see Table 2)

exceeds the fundamental value by at least 30% or 40% for more than five consecutive periods (i.e., x = 30% or 40%). Even if we increase x to 50% (which tends to give a low bubble count), there are still five bubbly sessions out of the six treatment-F sessions. There are no bubbles in treatment R using all three rules. For treatment N, two sessions (N2 and N6) have bubbles if we set x = 30% (which tends to give a high bubble count), one session (N6) has a bubble if x = 40%, and there are no bubbles if x = 50%.

From Tables 5-6, we can see that there are more fundamental traders in treatments R and N than in treatment F. For example, among the 70 subjects who participate in the seven sessions of treatment R, the percentage of fundamental transactions is 80% for 80% of the subjects (remember that in a fundamental transaction, the purchasing price is $\leq$ (1+30%) *x FV* and the sale price is $\geq$ 30%) *x FV*). In contrast, among the 59 participants in the six sessions of treatment F, only 14% of subjects have more than 80% of fundamental transactions. For treatment N, among the 57 participants, 72% have more than 80% fundamental transactions, which is much higher than the 14% in treatment F, but somewhat lower than the 80% for treatment R. The treatment average percentage of fundamental transactions is 94% for treatment R, 87% for treatment N and 61% for treatment F. The treatment average percentage of fundamental offers is 99% for treatment R, slightly lower at 95% for treatment N and much lower at 84% for treatment F.[14]

The Mann-Whitney test shows that the difference between the three treatments in terms of the percentage of fundamental trading and posting is statistically significant.

To summarize, treatment F involves substantial mispricing in the form of overpricing, while treatments R and N involve, on average, much lower mispricing and very mild under-pricing. The three treatments are not significantly different in terms of trading intensity.

1.6 Effect of interest payments on cash

We first check whether paying positive interest on cash helps to reduce bubbles by increasing the opportunity cost of speculation on the asset market.

[14] The data in Tables 5 and 6 show with regularity that posted offers follow the fundamental value more closely than trading prices.

Note that treatment F differs from a standard SSW design in that cash earns positive interest (both treatments have positive dividend payments and decreasing fundamental values). Since treatment F involves significant overpricing (as with the SSW design), we can conclude that paying interest on cash is not sufficient to suppress overpricing through the reducing-active-participation channel. The result is consistent with the findings in previous studies that have interest payments on cash. Bostian and Holt (2009) observe that bubbles frequently occur in an environment where cash earns positive interest and the fundamental value is constant. Fischbacher, Hens and Zeisberger (2013) investigate the effect of monetary policy in correcting mispricing by raising (cutting) the interest rate when the trading price is above (below) the fundamental value throughout the whole trading session. They find that raising the interest rate cannot eliminate bubbles.

In addition, a comparison between treatments F and R suggests that active participation is not the reason why bubbles (do not) appear in treatment F (R), because the two treatments share similar trading intensity. A more likely reason for bubble formation is confusion about the fundamental value process.

1.7 Effect of the fundamental value generating process

Given that the fundamental value generating process plays a critical role in the formation of bubbles, the next step is to identify features of the fundamental generating process that are responsible for the occurrence of bubbles. One explanation, as formulated in Smith (2010) and Oechssler (2010), is that, since asset prices tend to increase or stay constant in the long run in real life, subjects may find it difficult to comprehend that the fundamental price of the asset could decrease over time. According to this explanation, we should observe little overpricing in treatments with flat or increasing fundamental values. This explanation is consistent with some studies, but conflicts with others. For example, Smith, van Boening and Wellford (2000), Noussair, Robin and Ruffieux (2001), Huber, Kirchler and Stöckl (2012), and Kirchler, Huber and Stöckl (2012) find that bubbles are greatly reduced or disappear with a flat fundamental value, and our treatment R shows that there are no bubbles with increasing fundamental values; these results support the explanation. On the other hand, Bostian and Holt (2009) find positive bubbles with a flat fundamental value, Huber, Kirchler and Stöckl

(2012) find under-pricing or negative bubbles with increasing fundamental values, and our treatment N has decreasing fundamental values but no substantial overpricing; these results are inconsistent with the explanation.

Huber, Kirchler and Stöckl (2012) provide another explanation. They conduct an experiment with increasing, decreasing and flat fundamental values (in the absence of interest), induced by a positive, negative and zero expected dividend payment, respectively. They find no bubbles with a flat fundamental value, positive bubbles with decreasing fundamental values and negative bubbles with increasing fundamental values. As a result, Huber, Kirchler and Stöckl (2012) propose that anchoring on information generated by the trading process drives under-reaction and, in turn, mispricing on the experimental asset market. Again, this explanation is consistent with some studies, but incompatible with others. According to the explanation, there should be negative bubbles in our treatment R, no bubbles in Bostian and Holt (2009), and positive bubbles in our treatment N. However, there is minimal mispricing in our treatment R, significant overvaluation in Bostian and Holt (2009) and no substantial mispricing in our treatment N. To identify the source of confusion, we investigate in detail the fundamental value generating process, particularly the way in which each paper controls the time trend of the fundamental value of the traded asset. In papers that feature environments without interest payments on cash, the time trend of the fundamental value is determined by the sign of the expected dividend payment, d. A constant fundamental value requires $d = 0$, as in Smith, van Boening and Wellford (2000), Noussair, Robin and Ruffieux (2001), Huber, Kirchler and Stöckl (2012), and Kirchler, Huber and Stöckl (2012). More specifically, in Smith, van Boening and Wellford (2000), shares do not pay dividends. In Noussair, Robin and Ruffieux (2001), the positive dividend is offset by a carrying cost, which implies a zero net expected dividend. In Huber, Kirchler and Stöckl (2012) and Kirchler, Huber and Stöckl (2012), the dividend follows a random process with a zero expected value. To achieve increasing fundamental values, d has to be negative as in Huber, Kirchler and Stöckl (2012), who find undervaluation in this treatment. If $d > 0$, as in a standard SSW design, the time series of the fundamental value has a decreasing trend and bubbles frequently appear in this setting. Bostian and Holt (2009) and our paper study the formation of bubbles in environments where cash earns interest. The time trend of the fundamental value is determined by three parameters: K, d and r. Bostian and Holt (2009) achieve a flat fundamental value by setting $K = d/r$ with $K,\ d\ r\ >\ 0$. The fundamental value increases over time in our treatment R with $K, d, r > 0$ and $K > d/r$. In treatment F, we have $K, d, r > 0$ and $K < d/r$. In treatment N, we have $K > d/r > 0$ with $d < 0$ and $r < 0$.

Based on the above observations, we propose a third explanation: mispricing tends to occur whenever there is a conflict between the sign of the time trend of the fundamental value and the sign of the expected dividend. Subjects are more likely to perceive that the value of the share should increase if it pays a positive dividend, decrease if the dividend is negative and remain flat if the dividend is zero. The hypothesis is compatible with the results from all existing studies. In Smith, van Boening and Wellford (2000), Noussair, Robin and Ruffieux (2001), Huber, Kirchler and Stöckl (2012), and Kirchler, Huber and Stöckl (2012), the fundamental value is flat and d=0; the trading price tends to follow the fundamental value well. In the SSW design, Bostian and Holt (2009) and our treatment F, the time slope of the fundamental value is negative or zero, but the dividend is positive; this conflict induces overvaluation or positive bubbles. In our treatment R (N), both the time slope of the fundamental value and the expected dividend payment are positive (negative); the extent of overpricing is, in general, small. In the increasing treatment in Huber, Kirchler and Stöckl (2012), the fundamental value increases over time but the dividend is negative, a conflict that results in undervaluation or negative bubbles.

1.8 Conclusion

In this paper, we investigate the formation of price bubbles in an experimental asset market with interest payments or charges on cash holdings. We investigate two aspects of SSW design that may have contributed to irrational trading on experimental asset markets: low opportunity cost of speculation and confusion about the fundamental value.

We have run three treatments. In the first two treatments (treatments F and R), cash earns a positive interest payment, and stocks pay positive expected dividends. The difference lies in the dynamics of the fundamental value: the fundamental value decreases over time in treatment F and increases in treatment R. In the third treatment (treatment N), banking fees are charged on cash holdings, the expected dividend is negative (interpreted as carrying costs), and the fundamental value decreases over time. We find little mispricing in treatment R, substantial overpricing in treatment F and mild under-pricing in treatment N. There is no significant difference in terms of trading volume among all three treatments.

The results suggest that paying interest on cash is not likely to reduce asset bubbles through the reducing-active-participation channel, and the occurrence of bubbles is mainly due to confusion about the fundamental value of the asset. In order to identify the source of confusion, we investigate in detail the designs of existing studies, including our own three treatments, with a particular focus on how each study controls the time trend of the fundamental value. We offer a new explanation that bubbles tend to occur whenever there is a conflict between the sign of the time trend of the fundamental value and the sign of expected dividend payments. This new explanation is consistent with all existing studies.

1.9 Appendix

The experiment consists of a sequence of trading periods, each one lasting for 150 seconds. During each period, you will make decisions to invest your money between two forms of investment: shares of stocks of a fictitious Company, and a savings account. The currency used in the market is called EURUX, which will be converted into RMB at the end of the experiment. The conversion rate is 1000 EURUX for 1 RMB.

TRADING INTERFACE

In each trading period, you start with some money invested in two forms of investment: savings account and shares. Money in the savings account earns interest. Shares earn dividends (dividends and interests will be described later).

During each trading period, you make investment decisions to allocate money between the two forms of investment: you can use money in the savings account to buy shares, or sell shares and deposit the revenue in your savings account. Here is a sample trading screen.

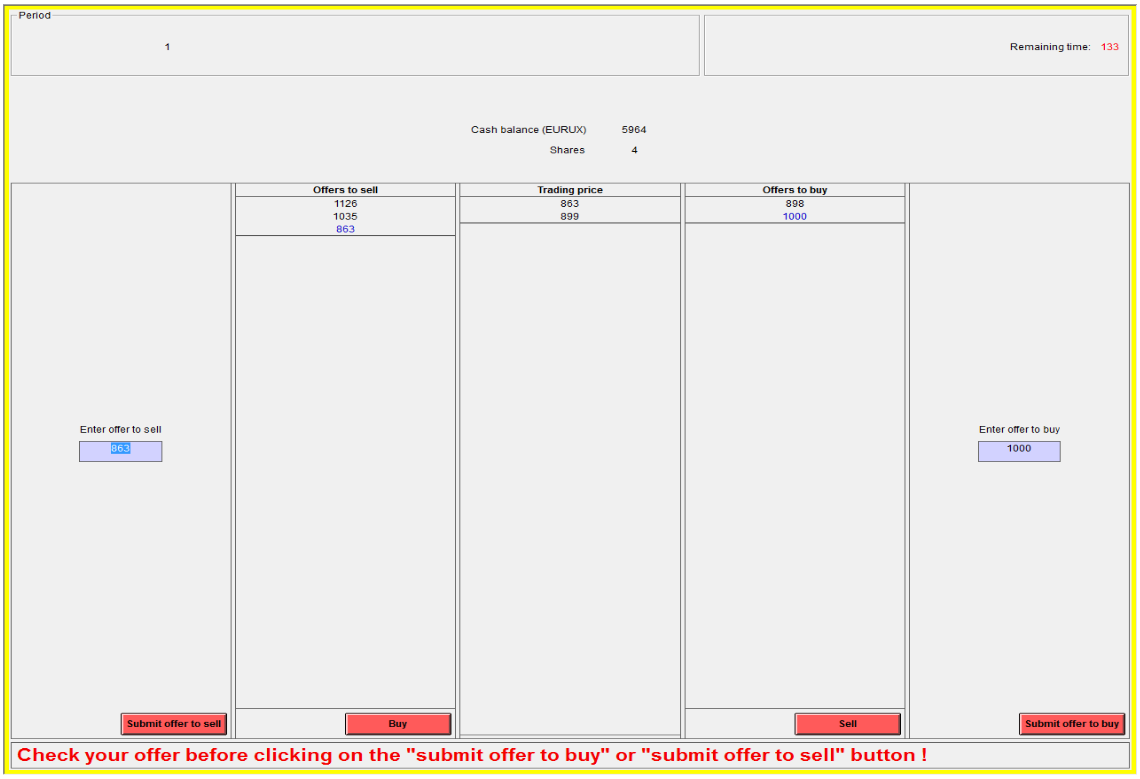

The top left corner shows the current trading period, and the top right corner shows how much time (in seconds) is left in the current period. Your investment portfolio – money in your saving account and the number of shares you own – are shown in the middle of the screen. On this screen you can buy or sell shares in four ways.

<u>First</u>, you can initiate a sale of shares by submitting an offer to sell.

If you have shares, you may choose to sell them. You can initiate a sale in the text area below "Enter offer to sell" in the first column. Here you can enter the price at which you are offering to sell a share. To send the offer, you have to click the "Submit offer to sell" button. After that, your offer to sell will appear in the second column labelled "Offers to sell". Each offer introduced corresponds to one single share. If you want to sell more shares, repeat this process.

Note that by submitting an offer to sell, you initiate a sale, but the sale will not be executed until someone accepts it.

Try offering to sell a share now. Write a number (integer) in the text area labelled "Enter offer to sell" and then click on the button "Submit offer to sell". You can see that a set of numbers will appear in the column labelled "Offers to sell". Each number corresponds to an offer from one of the participants. Your own offers are shown in blue; others' offers are shown in black. The offers to sell are ranked from high to low, so that the cheapest (best) price is displayed at the bottom of the list.

<u>Second</u>, you can realize a purchase of shares by accepting an offer to sell.

If you have enough money in your savings account, you can buy a share at one of the prices in the "Offers to sell" column (which also contains your previously submitted offer to sell). You buy a share by selecting one of the others' offers (shown in black) and then clicking on the red button "Buy". Note that you are not allowed to accept your own offers, which are shown in blue. Remember that the cheapest (best) price is displayed at the bottom of the list.

It may happen that when you select the best price and press the "Buy" button, someone else is doing the same thing but acting slightly faster than you. In that case, a message "someone has been faster than you" will show up.

Try buying a share now. Choose a price in the column "Offer to sell" and then click on the "Buy" button; or directly click on the "Buy" button and buy at the cheapest price listed in the column "Offers to sell".

Whenever an offer is accepted, a transaction is executed. Immediately when you accept an offer to sell, you realize a purchase and the number of EURUX in your savings account goes down by the trading price; at the same time, your trading partner realizes a sale and the balance in his/her savings account increases by the trading price. In contrast, when your offer to sell is accepted, you realize a sale, your trading partner realizes a purchase, and money is transferred from your trading partner's savings account to your savings account by the amount of the trading price.

Given that you all submitted one offer to sell and accepted one offer to sell, you all realized one purchase and one sale so you have the same number of shares as you started out with.

Third, you can initiate a purchase of a share by submitting an offer to buy.

If you have money in your savings account and would like to buy a share, you can initiate the purchase by submitting an offer to buy. Enter a number in the text box under "Enter offer to buy" situated on the right side of the screen and then click on the "Submit offer to buy" button.

Try submitting an offer to buy a share now. Write a number in the text area "Enter offer to buy." Then press the red button labelled "Submit offer to buy". Immediately in the column labelled "Offers to buy" you will see a list of numbers ranked from low to high, so that the highest (best) price is displayed at the bottom of the list. If you want to sell more shares, repeat this process. Again, your own offers are shown in blue; others' offers are shown in black.

<u>Fourth</u>, you can realize a sale of a share by accepting an offer to buy.

You can sell a share at one of the prices offered in the "Offers to buy" column (which also contains your previously submitted offer to buy). Select one of the offers and then click on the red button "Sell". Again, note you are not allowed to accept your own offers (shown in blue). Remember that the highest (best) price is displayed at the bottom of the list.

Try selling a share now. Choose a price in the column "Offer to buy" and then click on the "Sell" button.

Again, a transaction is executed whenever an offer to buy is accepted. If you accept an offer to buy posted by others, you realize a sale and as a result, the amount of EURUX in your savings account increases by the trading price. In contrast, when your offer to buy is accepted by someone else, you realize a purchase and the number of EURUX in your savings account decreases by the trading price. The reverse happens to your trading partner.

You can see that the these four trading methods are complementary: you can initiate a trade by offering a price to sell or buy and wait for the offer to be accepted by others; you can execute/realize a trade by accepting an offer to buy or sell submitted by other participants.

In the column situated in the middle of the screen and labelled “Trading price”, you can see the prices at which shares have been traded during the trading period by all participants present in the market.

SHARE and SAVINGS ACCOUNT

Shares

At the end of the trading period, you receive dividends for the shares you hold. Dividends are automatically added to your savings account.

The amount of dividend per share is determined by a random device (the Company’s business may go well or bad, which will affect how much dividend you get) and takes one of four values with the same probability:

1/4 probability you get 0 EURUX per share,

1/4 probability you get 8 EURUX per share,

1/4 probability you get 28 EURUX per share, and

1/4 probability you get 60 EURUX per share

Each participant gets the same dividend per share. There is a new random dividend draw for each new trading period.

Since all four outcomes are equally likely, we can calculate the average dividend as (0 + 8 + 28 + 60)/4=24 EURUX.

At the end of the game, the Company will purchase your shares at a buyout value of 72 EURUX per share.

Savings Account

The money in your savings account earns interest rate at 10% per period.

An Example

Here is an example to illustrate how dividends and interest are paid. Suppose after trading, you have 2 shares and 1000 EURUX in your savings account. The random device shows that each share receives a dividend of 28 EURUX. At the end of the period, you will receive 28x2=56 EURUX of dividend and 1000x10%=100 EURUX of interest. As a result, the balance in your savings account at the end of the period will be 1000+56+100=1156.

END-OF-PERIOD INFORMATION SCREEN

At the end of the trading period, after dividends and are paid and deposited in your savings account, you will be shown an "information screen". The screen shows you the dividend payment, and also the information about your end-of-period inventory of shares and the balance in your saving account.

The "information screen" contains the following information:

1. *Period*: the period just finished
2. *Your shares*: number of shares you own after trading in the period
3. *Savings account balance before dividend and interest*: amount of *EURUX* you have in your savings account right after trading and before dividend and interest payment
4. *Dividend per share:* the amount of dividend in *EURUX* you receive for each share you own.
5. *Total dividend:* calculated as *Your shares* x *dividend per share.*
6. *Interest:* net amount of interest you receive in the period for money in your savings account, which is calculated as *Savings account balance before dividend and interest* x *10%.*

7. *Savings account balance after dividend and interest:* money in your savings account after dividend and interest have been paid and deposited, which is calculated as *Savings account balance before dividend and interest + Total dividend + Interest*

The experiment consists of 15 consecutive trading periods. Each period will last for 150 seconds. You start period 1 with a certain investment portfolio of shares and money in your savings account. In each of the 15 trading periods, you trade among yourselves using the interface you just practiced with. At the end of each trading period, you see the "information screen" which shows your end-of-period portfolio position after dividend payment.

Your inventory of shares and savings account balance carry over from one period to the next. For example, if at the end of period 4 you have 2 shares and 1000 EURUX. You start period 5 with the same portfolio of 2 shares and 1000 EURUX before trading.

The game ends after 15 periods. If you own some shares at the end of period 15, the Company will purchase your shares at a buyout value of 72 EURUX per share.

For example, suppose after trading in period 15, you own 3 shares and 2000 EURUX. At the end of period 15, after dividend and interest payment, you can sell your shares to the Company at the buyout value. If the dividend payment is 8 EURUX per share, you receive 3x8=24 EURUX as dividends. The interest payment is 2000x10%=200 EURUX. Your 3 shares are sold to the Company for 3x72=216 EURUX. Your total earnings in this game are calculated to be 216+24+200+2000=2440 EURUX, which will be converted into RMB.

HOLDING VALUE TABLE

The objective of your investment decisions is to maximize your end-of-game total earnings. In each trading period, you decide how to allocate your money between the two forms of investment: shares and savings account.

To facilitate your decision-making, we provide you a table called "Holding value table" (See next Page), which can be used through the entire experiment. The table calculates the average amount of money you earn if you buy a share in the current period and hold it until the end of the game. Of course, you may choose not to hold the share until the end of the game, if, for example, you can sell it at a good price before the end of the game. The holding value table is just for your reference.

The table has 6 columns, which we will go through one-by-one.

1. *Current period*: The current trading period.
2. *Average dividend*: The average amount of dividend per share per period. This, as explained earlier, is equal to 24 EURUX.
3. *Average remaining dividends*: If you hold 1 share of stock until the end of the game, you will be entitled to a dividend payment at the end of each of the remaining periods. The remaining dividend is calculated as the total amount of money you will accumulate **at the end of the game** if you deposit all dividend payments into your savings account which earns 10% interest per period. For example, for each share you hold in period 14, there are two remaining dividend payments: one at the end of period 14, and one at the end of period 15. You deposit the period 14 dividends in your savings account, which will increase your money balance at the end of the game by 24x1.1=26.4 EURUX. The period 15 dividend is paid at the end of the game (so will not earn interest) and will increase your end-of-game money balance by 24 EURUX. The average remaining dividends is calculated as the sum of the two amounts = 26.4+24=50.4 EURUX.
4. *Buyout value.* At the end of game, each share you own will be purchased by the Company at 72 EURUX.
5. *End average holding value*. The average amount of EURUX you will receive at the end of the game if you hold one share for the remainder of the experiment. It is calculated as the sum of *average remaining dividend* (column 3) and the *buyout value* (column 4). For example, the average holding value I for a share in period 14 is calculated as 50.4+72=122.4 EURUX.
6. *Current average holding value*. To buy a share in the current period, you have to use money currently in your savings account. When you make the buying decision, you may want to know the average holding value of a share measured in terms of money in the current savings account. Call this the *current average holding value*. Let us illustrate how to calculate the value by an example. Suppose you are trading in period 14. One EURUX in the current saving account will generate 1.1^2 (there are two remaining interest payments) units of EURUX at the end of the game. Holding one share generates (on average) 122.40 EURUX at the end of the game. Holding

one share is thus (on average) equivalent to holding 122.40/1.21=101 EURUX in the current savings account.

Holding Value Table

1 Current period	2 Average dividend	3 Average remaining dividends	4 Buyout value	5 End average holding value	6 Current average holding value
1	24	762.54		834.54	200
2	24	671.40		743.40	196
3	24	588.55		660.55	191
4	24	513.22		585.22	186
5	24	444.75		516.75	181
6	24	382.50		454.50	175
7	24	325.91		397.91	169
8	24	274.46		346.46	162
9	24	227.69		299.69	154
10	24	185.17		257.17	145
11	24	146.52		218.52	136
12	24	111.38		183.38	125
13	24	79.44		151.44	114
14	24	**50.40**		**122.40**	**101**
15	24	24.00	72	96.00	87

Quiz

Please read carefully the Holding Value Table and make sure that you understand it. Raise your hand whenever you have any questions. When you think you understood the table, please answer the following questions:

1. Suppose you are in period 5. How much is the *average dividend* you should expect at the end of this period? ________
2. Which is the maximum and minimum dividend you can get in any period? _________
3. Suppose you are in period 5 and a share pays the average dividend in each of the remaining periods. The *current holding value of* one share in terms of money in the current savings account is ______.
4. Please explain on one sentence or two what the *current holding value* is.

CHAPTER 2

RECREATING THE SOUTH SEA BUBBLE: LESSON FROM AN EXPERIMENT IN FINANCIAL HISTORY

(Joint with Charles N. Noussair and Hans-Joachim Voth)

2.1 Introduction

From the Dutch Tulipmania of the 17th century to the NASDAQ bubble, the rise and fall of speculative bubbles has produced massive gains and losses for investors. As the financial crisis of 2007-08 has demonstrated, bubbles can also be a major source of economic instability. Large swings in asset prices over short periods of time have often been considered as a sign of inefficiency and a telling testament to the "irrational exuberance" of investors. One important strand in the literature denies the existence of bubbles altogether (Garber 2001; Fama 1965). Other scholars have sought to explain their emergence as a result of risk shifting, investor inexperience, and limitations of market micro-structure such as an inability to short over-valued stocks (Allen and Gale 2000; Greenwood and Nagel 2009; Hong, Scheinkman, and Xiong 2006; Lintner 1971). Empirical work has demonstrated that sophisticated investors – instead of attacking mispricing – often "ride" bubbles, aggravating price swings (Brunnermeier and Nagel 2004; Temin and Voth 2004). The experimental evidence shows that the size and duration of bubbles are sensitive to the market parameters, institutions, and incentives in place (see Palan 2013 for a review of this literature).

While the theoretical and empirical literature offers explanations for the continuation of mispricing, the origins of bubbles are not well-understood. In particular, there is no convincing explanation for why bubbles appear at certain times in particular markets, but not at other times or in other markets. Experimental work suggests that bubbles emerge readily in laboratory settings (Smith, Suchanek, and Williams 1988). This makes it all the more puzzling that major bubbles have only erupted on a handful of occasions over the last 400 years, such as during the Tulipmania, the South Sea bubble, and the NASDAQ episodes. Milton Friedman (2001) concluded that "…the start and end of a bubble just cannot be explained rationally."

In this paper, we report the results of a laboratory experiment designed to study which specific institutional features played a role in igniting the South Sea mania. We recreate many of the incentives faced by investors in 1720 in our laboratory experiments, and then isolate their effects by "switching them off" one-by-one. Of course, many historical features cannot be recreated. Nevertheless, if an institution exerts a systematic effect on mispricing in our experiments, some 400 years after the event, we argue that it is likely to have played an important role in the original episode. Because our interest is in understanding a specific historical episode, we do not adhere closely to any previously studied paradigm, but rather develop a new experimental design that is tailored to the purpose of our study.

Together with the Tulipmania and the Mississippi Bubble, the South Sea bubble is one of the three famous, early bubbles that occurred during the 17th and 18th centuries (Carswell 1960; Dale 2004) . It is among both the best-documented and the least well-understood episodes in the history of financial markets.[15] Originally created to trade with Spanish America, the South Sea Company's main source of revenue was interest payments from its holding of UK government bonds. In late 1719, it proposed to swap all outstanding government debt for its own equity. After the contract was awarded by Parliament, it began to issue new stock through subscriptions. Eventually, it exchanged government debt for equity. The South Sea Company's stock rose in value from a little more than £120 at the start of the year 1720 to nearly £1,000 in June, before crashing by some 80 percent before year-end (figure 4). While other stocks also saw their prices surge, the sheer scale and speed of the South Sea Company's price explosion and decline are without parallel.[16] The South Sea bubble did not only matter for investors at the time. As a result of lobbying by the company, when the stock price was near its peak, England effectively prohibited the issuance of shares in new companies. This

[15] One example demonstrating the extent of disagreement is the debate between Shea (2007) and R. S. Dale, Johnson, and Tang (2005).
[16] Temin and Voth (2004) compare the magnitude of the price run-up and decline in the South Sea and the NASDAQ bubble, and find the former to be markedly larger.

closed the stock market to firms for over a century, resulting in markedly greater difficulties in raising funds for new ventures (Harris 1994; Temin and Voth 2013).

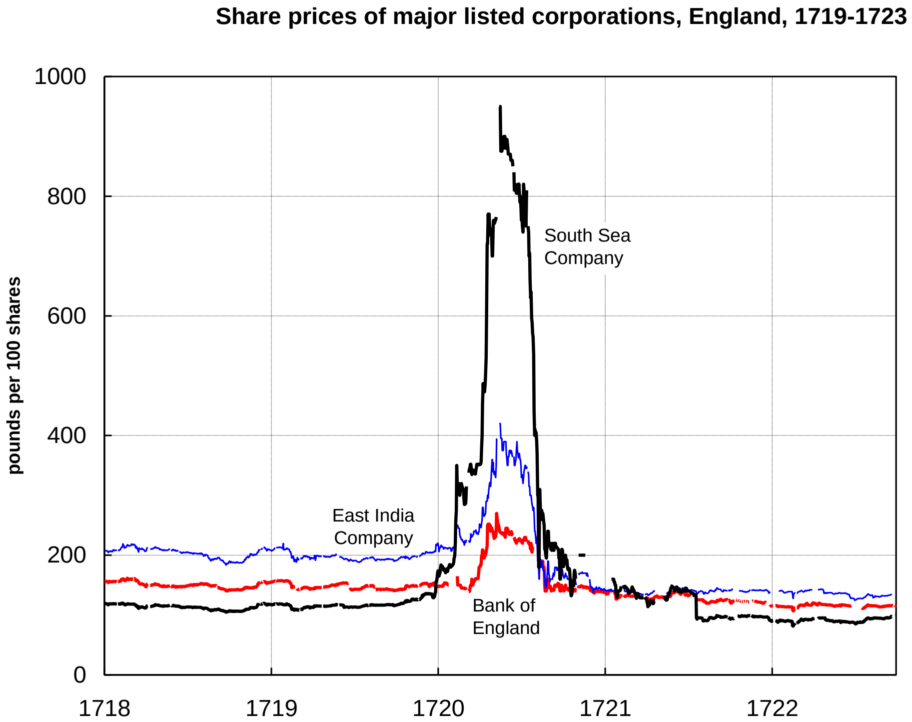

Figure 4: Share prices of major listed corporations, England, 1719-1723

The institutional features of the South Sea scheme inform our choice of experimental treatments. The explosion in the South Sea Company's stock price took place when it (1) was seeking to swap government debt for equity, (2) allowed investor to purchase shares while deferring payment of most of the purchase price, and (3) had limited ability to enforce collection of the deferred payments. We examine all these specific features in isolation. The baseline treatment of our design includes all of these features and generates a large bubble. In the other treatments, we remove the institutional features (1) – (3), one at a time. The magnitude of the bubble in these treatments is compared to that in the baseline treatment. In this manner, we are able to identify which aspects likely contributed to one of the greatest and most famous bubbles in history.[17]

[17] By focusing on the relative magnitude of bubbles, we also sidestep the issues raised about bubbles in experimental settings by Kirchler et al. (2012).

Our paper makes two main contributions. The first is substantive: Our key finding is in providing evidence about what factors might have been responsible for the eruption of one of history's greatest bubbles. Our findings suggest that the attempt to swap government debt for company equity was the single most important contributing factor to the South Sea bubble. The plan to take over all of England's government debt was the most unusual aspect of the South Sea Company's finances; its large role in igniting the bubble offers one explanation why similar episodes have been rare. In addition, the ability to defer payment of shares helped to increase the likelihood of a bubble forming, and increased its size. Limited contract enforcement played a markedly smaller role. The second contribution is methodological: We use the laboratory to try to understand a specific episode in economic history, recreating incentives agents faced nearly 400 years ago. We exploit the possibility, offered by laboratory experimentation, to observe the fundamental value and thus measure the magnitude of a bubble, to introduce and remove institutional features keeping all else constant, and to generate new data designed to reproduce conditions that ceased to exist hundreds of years ago. We view this last feature as particularly beneficial to the study of economic history.

The paper proceeds as follows. Section 2.2 summarizes the historical context and background of the South Sea bubble. Section 2.3 presents our experimental design and describes how we capture the essential historical details of trading "in the South Seas" (as contemporaries would have said). In Section 2.4, we report our results, and Section 2.5 concludes.

2.2 Historical context and background

The South Sea Company was founded to trade with South America. The Peace Treaty of Utrecht in 1713, which brought the War of the Spanish Succession to an end, granted Britain the right to send trading ships periodically to Spain's possessions in the Americas. The company took over some of the government's debt in exchange for the trading privileges. Its mercantile operations never amounted to much. By the late 1710s, the South Sea Company amounted to little more than a shell company distributing interest payments on government debt to its shareholders.

In 1719, the Company took over another part of the national debt, referred to as the "lottery loan". While paying a high rate of interest, the loan was highly illiquid. Bonds could not readily be transferred; price discounts were substantial. The operation that swapped these government bonds for equity in the South Sea Company was widely considered a success – the investors gained a more liquid asset, the government lowered the interest charges on its debt, and the company made a profit.

The 1720 scheme was vastly more ambitious – and it contained one crucial difference with the 1719 operation. The South Sea Company proposed to take over the entire remaining national debt (except for the parts held by the Bank of England and the East India Company). Instead of swapping debt for equity at a pre-established price, the company remained vague as to the exchange ratio. This implied that as the stock price appreciated, more debt could be bought for each share.

Bidding against the Bank of England for the right to do the debt conversion, the South Sea Company finally won the contract in a parliamentary vote in 1720. Massive bribery preceded the award of the contract. By this stage, the stock price had more than doubled. After the award of the contract, the South Sea Company began to issue new shares in repeated rounds of offerings. As shown in table 7, it did so at steadily rising prices – for £300 in early April, £400 in late April, £1,000 in June, and £1,000 in August. These were known as "subscriptions", and were bought on installment plans. Actual down payments amounted to only £40-200 (10-20% of the total cost). Subscribers did not become owners of shares until all payments had been made. Subscription receipts could themselves be traded. Their prices moved in parallel with the price of the underlying stock, but in relative terms, price changes were magnified – as they are with options.[18]

[18] For the exact details of the analogy with options pricing, cf. Shea (2007)

Subscription round	**1**	**2**	**3**	**4**
Date	14 April	29 April	16 June	24 August
Issue Price	300	400	1000	1000
Final Payment Due*	14 August, 1721	24 April, 1723	2 January, 1725	24 August 1722
Premium**	-1 percent	9.7 percent	21.3 percent	27.6 percent
Gain of Subscription Value within One Week***	10 percent	10 percent	150 percent	53.5 percent

* according to the original issuance schedule

** NPV of subscription payments relative to the market price at time of issuance

*** calculated as the difference between the subscription price and the price of South Sea stock one week after the subscription closed

Table 7: South Sea Company issues of new shares, 1720

Throughout the spring and summer of 1720, the stock price moved up, reaching nearly £1,000 by June. Many other stock schemes sprang up during the same time, luring investors. Also, many inexperienced investors entered the market, often in the expectation of a quick profit. The company initially did not use the proceeds from share issues to actually buy back government bonds, as the original scheme had envisaged. Instead, it lent generously against its own shares. Many of these loans were later not repaid.

The actual exchange of government debt for equity took place in May and in August. Eventually, in the spring, the company began to offer bond holders a chance to exchange bonds for South Sea stock. Terms were not overly generous, but a significant share of bondholders nonetheless accepted the deal. As the stock price increased, existing bonds could be bought in exchange for fewer share certificates. This increased the intrinsic value of the stock. In other words, the rise in the stock price, the issuance of new shares, and the possibility of buying out debt holders created winners and losers (Carswell 1993 : 120). Contemporaries were keenly aware of this fact. As a matter of fact, a Member of Parliament – Archibald Hutcheson – published several pamphlets in the course of 1720, pointing out the losses and gains to different subscribers and the original bondholders, as a function of when they had bought.

One of the tables published by Hutcheson is reproduced as Table 8[19]. We see that in the fall of 1720 – after the second subscription of shares at £1,000 per 100 shares had closed – there were clear winners and losers. Column 1 gives the proportion of the stock held by different groups – the old proprietors, who bought South Sea stock at £100, and the subscribers, who had bought at increasing prices. The company had issued shares for a nominal value of £42 million. In the aggregate, it had sold them to the public for £234 million, or £557 per 100 shares. At the time, South Sea stock was worth close to £600. Thus, all the subscribers who paid less than this had made money; and those who had bought for £1,000 had lost. Column 2 gives, for each group, the market value of stock held; column 3 summarizes what investors paid for it. The magnitude of gains and losses is summarized in column 4. Some £90 million of losses accrued to the new subscribers at £1,000. Most of these ended up in the pockets of the old proprietors (£ 57 million) and, in much smaller quantities, in those of the early subscribers.

	Proport^s. in this Stock.	Value thereof.		Paid for the ſame.	Loſs thereon.			
	l.	*l.*	*Decim.*	*l.*	*l.*	*Decim.*	*l.*	*Decim.*
The Firſt Subſcribers at 1000 *l.*	6,600,000	36,771,428	571428	60,000,000	23,228,571	428572		
The Second Subſcribers at 1000 *l.*	15,400,000	85,800,000	- - - -	154,000,000	68,200,000	- - - -	91,428,571	428572
					Gained thereon.			
The Old Proprietors at 100 *l.*	12,320,000	68,640,000	- - - -	11,200,000	57,440,000	- - - -		
Subſcribers at 300 *l. per Cent.*	2,475,000	13,789,285	714286	6,750,000	7,039,285	714286		
Ditto at 375 *l.*	3,665,000	20,419,285	714286	12,494,318	7,924,967	714286		
Ditto at 400 *l.*	1,540,000	8,580,000	- - - -	5,600,000	2,980,000	- - - -	75,384,253	428572
	42,000,000	234,000,000	- - - -	250,044,318	- - - - - - - -		16,044,318	

Table 8: Winners and loser from the South Sea conversion scheme

19 Source: Archibald Hutchenson (1720)

The basic principle that ensured that investors put their money into the South Sea scheme was simple enough. The expected losses and gains for various groups, as set out in detail by Hutcheson, could not have remained a secret. One key question is then why the new owners bought shares which gave them cash flow rights that were lower than their market price? Put simply, an investor in 1720 could acquire the rights to future interest payments by buying a government bond, or by buying South Sea stock. Why pay more for the latter? One logical possibility is that some commercial venture might produce vast profits. This was highly unlikely.

Our answer is that new subscribers could possibly benefit from rising inherent values of their shares as a result of additional stock issuance in the future – again, if bondholders could be bought out more cheaply, the intrinsic value of shares would increase. Note that for this mechanism to work no actual purchase of bonds is necessary – it is enough that it is planned. New investors are willing to buy because they hope that they will gain if prices continue on an upward trend (which then translates into a self-fulfilling prophecy). The prospect of future issuance can turn the loss in the present into a purchase that, at least in expectation, can turn out to be profitable at some point in the future.

The structure we just described is, of course, that of a classic Ponzi scheme. The secret to success is to join early (and to get out before things fall apart). As long as there is a good chance that another wave of investors will enter, it is a good idea to participate. The details of the South Sea operations were complex, and there is some evidence that investors and the general public did not find it easy to see through it all. The *Flying Post*, a newspaper at the time, argued on April 9, 1720 that the intrinsic value of the South Sea Company stock would be £448 if the share price went to £300. At £600, it would be £880. We do not know how these numbers were calculated; it is clear that they cannot be correct. The basic structure is right, with higher share prices justifying a higher fundamental. At the same time, the relative prices are wrong – the intrinsic value in this operation can never catch up with the price at which the last issue was undertaken.

2.3 Hypotheses

Based on the historical background described above, it is possible to distill a number of hypotheses about the factors that might have contributed to one of the biggest bubbles in history:

Hypothesis 1: Redistributing revenue from new issues to existing shareholders increases prices and bubble magnitude.

As indicated earlier, the redistribution of new revenues to existing shareholders, in the form of interest payments on the bonds purchased with these revenues, increased the return to existing shareholders. This increase in fundamental values could launch a bubble, in which prices depart from intrinsic values. This could occur, for example, if heterogeneous beliefs about the likelihood of future issues lead to those with the most optimistic beliefs about future issues to bid up the price. Furthermore, the increases in price could attract the attention of momentum traders who bid up prices, merely because they expect past price trends to continue. Speculators may attempt take advantage of the presence of momentum traders, bid up prices further, and magnify a bubble.

Hypothesis 2: When new shares issues can be paid for in installments, bubbles are greater than when they must be fully paid for up front.

The possibility of payment by installment would, in principle, relax cash constraints on those individuals who are speculating on future new share issues. This would allow them to take larger long positions, possibly inflating a bubble. Furthermore, deferment of payment for new shares allows greater leverage for investors who are speculating on a rising market. Consider an individual who is interested in a purchase for later resale, and who expects the price to appreciate by 20% between period t and t + 1. If deferment of payment is not feasible and he purchases and resells, his return is simply 20%. However, if he can pay only 20% in period t, sell for 120% of the original contract price and pay back the remaining 80% at that time, his return is 100%. This leveraging effect could magnify a bubble once prices have begun to rise.

Hypothesis 3: If outstanding installment payments do not have to be paid back at the end of the life of the asset, bubbles are greater than if they do have to be paid.

The ability to purchase in installments increases the capacity to speculate on future price increases. If the resulting debt does not have to be paid back, it lowers the downside risk of holding the asset. If prices fall to a level below the amount the individual owes, she might find it more profitable to default on her remaining installment payments and give up the share. Thus, if prices are volatile, the default option gives the asset the properties of a call option, limiting downside risk in a similar manner. In our experiment, we implement the possibility of default by nullifying any outstanding debts owed at the end of the life of the asset. If traders take this effect into account, it could increase demand for the asset when future prices are unpredictable, and thus exacerbate any bubbles that occur.[20]

2.4 Experimental design

2.4.1 Procedures common to all treatments

The experiment consisted of four treatments, called Baseline, NoInstall, NoDefault, and NoSwap. We conducted eight sessions of the baseline treatment and four sessions of each of the other three treatments, for a total of 20 sessions. The sessions were conducted at University Pompeu Fabra and Tilburg University.[21] Each session consisted of two consecutive horizons, with each horizon made up of a sequence of between ten and twenty 150-second periods. At the beginning of the second horizon, endowments were reinitialized to the same levels as the beginning of the first.[22] This means that the

[20] Here, the intuition is similar to the one in recent work that explores the analogy between options pricing and stock prices under high uncertainty (Pastor and Veronesi 2006).

[21] The sessions were conducted at University Pompeu Fabra (17 sessions) and at Tilburg University (3 sessions) between July 2012 and December 2012. Each session took approximately 2 hours and 40 minutes. The average payment was 9 Euros per hour at UPF and 13 Euros per hour at Tilburg.

[22] We use the term session in this paper in the manner in which it is typically employed in experimental economics, which differs from common usage in finance. In experimental economics, a session refers to a continuous time interval in which a particular cohort is present in the laboratory. The sessions of our experiment are divided up into two trading horizons. The horizons are independent of each other, in that at the beginning of the second horizon all parameters are reinitialized, and the only link between the two horizons is the possible effect of subjects' prior experience in the first horizon on behavior in the second. Each horizon is divided up into periods. A trader's cash and asset position at the end of one period carry over to the next period within the same horizon. Thus, according to one common notion in finance, a

two horizons can be viewed as two distinct economies, linked only by the experience participants accumulate in during the first horizon.

When subjects arrived at the laboratory, they received approximately 45 minutes of instructions. They were instructed in both the use of the computer software and the specific conditions of the treatment in effect during the session. All subjects had the role of traders, with the ability to both purchase and sell the asset. After instruction, they familiarized themselves with the game by trading during a practice phase in which there was one auction of new shares and three periods of trading of shares. Thereafter, the first trading horizon began.[23]

In each period, up to two markets could operate simultaneously. In Market A, the shares originally issued at the beginning of the trading horizon could be exchanged. That is, market A served for trading those shares in circulation at the outset of the first period of a horizon. Market B enabled trading of newly-issued shares, those shares issued after the horizon began. New shares could be issued in any period, beginning in period four. Shares trading in market A and market B were identical in terms of the dividends paid and their expected lifetime, and therefore also identical in terms of their fundamental values.[24]

The markets were organized using continuous double auction rules (Smith 1962), and implemented with the z-tree computer software (Fischbacher 2007). Trade took place in terms of an experimental currency, called “ducats”, which was converted into Euros at the end of the experiment at a conversion rate that was common knowledge.

2.4.2 The asset

At the end of each period, each unit of the asset paid a dividend that took on one of four values, 0, 8, 28, or 60 Ducats, each with equal probability. Therefore the expected dividend per period was equal to 24 Ducats. Dividend realizations were unknown until the time they were determined.

---

period in our experiment corresponds to a trading session or a trading day. However, in this paper, we maintain the usage of the term conventional in experimental economics.

[23] Participants were graduate and undergraduate students in various majors, but most were majoring in business and economics. No subjects had previously participated in any asset market experiment. The number of subjects participating in each session ranged from 6 to 10.

[24] The reason for keeping asset markets separate is that during the South Sea bubble, shares and “scrib” (subscription certificates) were trading in parallel.

The asset had a life of multiple periods. The maximum possible lifetime was 20 periods and the minimum was 10 periods. For each period from 10 to 19, whether the trading session would continue to the next period or not was determined randomly at the end of the period. The probability that the current period was the last one was equal to 1/6 in each of the periods from 10 to 19.

The fundamental value of both assets can be calculated at any time from the dividend and the probability of termination. This fundamental value is the sum of the expected dividends to be received from the current period *t* until the end of the life of the asset. This is given by the expression:

$$FV_t = \sum_{s=t}^{T} d_s * (1 + \pi)^{(T-s)}$$

where d_s denotes the expected dividend in each period and π is the probability that the horizon ends after the current period. For example, consider a fixed d_t of 24 for each period until end of the horizon. The fundamental value in period 16 would be:
$24 * (5/6)^4 + 24 * (5/6)^3 + 24 * (5/6)^2 + 24 * (5/6)^1 + 24 = 85.54$ Ducats.
Subjects were provided with a table that indicated the fundamental value in each period and how it was calculated. The same table was also displayed on each subject's computer screen before the start of each trading period.

2.4.3 Initial endowments and new issues

Before period 1, each subject was endowed with 5 shares and 30,000 ducats. These shares of asset could then be traded in market A at any time and the cash endowment could be used for new asset purchases. New issues of asset could occur in any period beginning in period 4. Issues occurred at the beginning of the period, just after the dividend in the previous period was paid and before the market opened for trade in the current period. The criterion for whether or not a new issue would take place in period t was the following: If the average transaction price in market B was greater than the fundamental value of the preceding period, a new issue would occur.[25]

New shares were issued with sealed bid auctions that took place at the beginning of some periods. In each period in which there was an issue, a subject could bid for up to two shares of the K new shares offered for sale. The bids were ordered from highest to lowest and the K highest bidders were awarded units. Winning bidders paid a per-unit price equal to the lowest of the accepted bids, that is, the K^{th} highest bid. Thus the auction format was a uniform-price sealed bid auction with lowest accepted bid pricing.[26] In the first period in which there was a subscription, eight shares were auctioned. In subsequent subscriptions, five shares were auctioned. The criterion for issuing shares was unknown to participants.

2.4.4 Treatments

The treatments are designed to assess the impact of (a) swapping debt for equity, (b) deferring payment for newly-issued shares, and (c) the possibility of default on these payments, for bubble formation. The baseline treatment is characterized by the presence of all institutional features (a) – (c). Each of the other three treatments eliminates one of the features (see Table 9). Consequently, we can isolate the marginal contribution of each factor when all of the others are present.[27] Table 9 displays the features characterizing each of the treatments, and distinguishing it from the other treatments.

	Baseline	NoInstall	NoDefault	NoSwap
Swapping	Yes	Yes	Yes	No
Installments	Yes	No	Yes	Yes
Default	Yes	No	No	Yes

Table 9: Treatment features

[25] Before starting trading in period 4, there is a subscription auction if the price in period 3 for asset A is greater than the fundamental value.

[26] While not demand-revealing, this auction format tends to generate highly efficient allocations when individuals demand is for a single unit or for two units (Alsemgeest, Noussair, and Olson 1998).

[27] Notice that the absence of payment in installments automatically implies that default was impossible. However this does not cause problems of inference, because we are still able to disentangle the effect of each single factor.

- **The *NoInstall* treatment: No payment through installments**

In three of the treatments, an individual who purchases newly-issued shares is not required to pay the full price of the share he has bought at the time of purchase. Only 20% of the price is subtracted immediately from his current cash balance. The remaining 80% is registered as debt that the individual owes. This debt is repaid as follows: In each of four periods immediately following the period of the share issue, 20% of the price originally paid is subtracted from his cash balance. In the *NoInstall* treatment, newly issued shares must be paid in full at the time of purchase.

- **The *NoDefault* treatment: Limited contract enforcement**

As explained above, the number of periods that the trading horizon continues is not known to traders. The horizon terminates with a 1/6 probability at the end of each period from 10 to 19, and with probability 1 at the end of period 20. The horizon may end before the debt of a trader from a purchase of new shares in the auction has been fully repaid. Subjects were informed that they would not have to repay any debts outstanding at the time a horizon ends. In other words, the debt owed is not subtracted from her cash holdings at the end of the trading horizon. The exception to this is the NoDefault treatment, in which the amount due is subtracted in full from his cash account when a horizon ends.

In the market for the South Sea shares, the default option was typically exercised during the crash because the market value of the shares fell below the level of the installments investors owed. In the experiment, we did not allow investors to voluntarily default at any time, as this would have been an additional demand on subjects who were already in a relatively complicated experiment. We instead implemented the default automatically at the end of the horizon, when all individuals would clearly prefer to default. This setup retains scope for the possibility of defaulting to increase bubble magnitude, since it can be expected to increase demand considerably near the end of the horizon when the default probability is relatively high.

- **The *NoSwap* treatment: Distribution of revenue from new share issues to existing shareholders.**

There was no direct swapping of shares for equity. Instead, we created a similar effect on incentives by supplementing the dividend payment to shareholders with revenue from new shares. Our setup makes the Ponzi-scheme nature of the South Sea scheme clearer, but it leaves the key payoff features intact. While the mechanics "in the background" are different, the direct recycling of payments creates an analogous effect to using the revenue from new shares to purchase government bonds and transferring the interest payments to shareholders. In three of the treatments, each time that new shares were issued, 15% of the new revenue from the issue was distributed to existing shareholders in proportion to their total share holdings. For example, if there are 20 shares outstanding and 5 more are issued for 1000 each, then 750 is given out in total in the current period, 37.5 to the holder of each outstanding share. This supplement of 15% of the revenue from the issue continues to be paid in each period for the remainder of the horizon. This new dividend is added directly to the random dividend originally paid on the share. Because the fundamental value of a share is the discounted stream of future dividends, the fundamental value increases after each new subscription.[28] In the NoSwap treatment, the revenue from the issue of new shares is not paid out to shareholders. This means that new issues do not contribute to the future expected dividend stream and thus do not affect fundamental values.

One feature that is not included in our design is short selling. The design here focuses on and isolates features that distinguished the South Sea share market from other asset markets. During the South Sea episode, short-sellers existed, though short selling was highly risky activity due to counterparty risk.[29] The role and extent of short selling was similar in the market for South Sea shares to that in other asset markets throughout history, in which bubbles did not occur. In experimental markets, short-selling tends to reduce prices. Prices are lower the larger the short selling capacity of traders is, and a sufficiently large short sale capacity pushes prices to levels below fundamental values (Haruvy and Noussair, 2006). There is no reason to suppose that the effect of introducing short-selling would be any different across our four treatments. It would have the effect of pushing prices lower, and would exert a larger effect in this direction, the larger the short sale capacity.

[28] At the beginning of each market period, the table displayed on subjects' computer screens that indicates the future expected dividend stream, is updated accordingly.
[29] There were also difficulties enforcing contracts in law. Cf. Neal 2012.

The market for South Sea shares was characterized by heterogeneity in experience among investors. It was a setting in which there was scope for relatively sophisticated investors to try to exploit poor decisions of less sophisticated ones, such as new entrants. This is also the case for other asset markets throughout history, including those that did not see bubbles or crashes. Thus, we do not include a flow of new entrants in our design. Nevertheless, our laboratory markets, like most asset markets, are characterized by differences in trader sophistication. Traders are sampled from the student population of diverse public universities. In experimental asset markets with such populations, individuals with relatively high cognitive ability substantially outperform those with lower ability (see for example Corgnet et al., 2012, or Breaban and Noussair, 2014). In this regard, our study has the potential to reproduce the heterogeneity present in the South Sea share market. While we employ a student population obviously differing in many ways from the investors in the South Sea company, the bulk of the evidence available indicates that the behavior of students and professionals in experimental paradigms does not differ systematically (Frechette, 2011). In our experiment, in which there are two consecutive trading horizons, we are able to gather data at two different experience levels for each cohort. As can be seen from figures 1 and 2 in the next section, the results do not differ appreciably between horizons, particularly with regard to which treatment condition generates a greater bubble.

2.5 Results

2.5.1 Market outcomes

The time series of transaction prices, as well as of the difference between transaction price and fundamentals, are shown in figure 5. Absolute transaction price time series are shown in the left panel and the excess of prices over fundamentals in the right panel. The upper half of the figure corresponds to the first horizon, and the lower half to the second horizon. Each of the series of data is the average over all four sessions of a given treatment. The data from markets A and B are pooled.[30]

[30] In principle, even though both assets have the same fundamental value at any point in time, it is possible that their trading prices differ. In order to check for this possibility, we run a set of t-test comparisons on the pooled data from all treatments. First, we test for differences in median prices

Figure 5: Treatment result - Market outcomes

between the two markets from period 4 on. A total of 464 prices were compared (232 prices for each asset type).The differences in prices between the two markets are very small. In market A, the mean price is 797, with a standard deviation of 584. In market B the mean price is = 792, with a standard deviation of 590. A t-test of whether the difference in prices is different from zero, using the pair of prices within each period as the unit of observation, yields t = 0.0987, with p = 0.9214. Second, we run 20 independent-sample t-tests, one for each session. Only in one out of 20 sessions we do find significantly different prices between the two markets, and in that case the difference was only present during five periods. The result that identical assets trade at identical prices is consistent with previous experimental work that shows that two assets with identical fundamental valuetrading simultaneously do not generate significantly different price patterns (Fisher and Kelly, 2000).

Some consistent patterns are evident in the figures.[31] The first is that the Baseline treatment exhibits prices much greater than fundamental values, as well as the highest prices among any of the treatments. This suggests that each of the three factors contributes to a bubble. The second pattern is that prices are lowest and closest to fundamental values in the NoSwap treatment. Indeed, in the second horizon, prices in the NoSwap treatment closely track fundamentals. This indicates that the paying out of revenues from new issues to existing shareholders is the most substantial factor in creating the bubble we have observed in our experiment. This suggests that the swapping of equity for government bonds was the most important contributor to the bubble during the South Sea episode.

The differences between treatments are similar in the two horizons. The Baseline treatment generates the largest bubble, followed by the NoDefault, the NoInstall, and lastly the NoSwap treatments. The differences tend to be more pronounced and of generally greater magnitude in the first trading horizon. The price dynamics, relative to fundamentals, follow a similar pattern in three of the treatments, with a price boom in early periods and steep price declines late in the session the typical dynamic. The exception is the NoSwap treatment, in which prices tend to track fundamental values closely.

Figure 6 shows the time series of transaction prices for the second horizon of each session, as well as the time path of fundamental values.[32]The fundamental value increases at the time of the first issue in all treatments except for NoSwap, since swapping debt for equity redistributes some of the new funds to existing shareholders as dividends. There is considerable heterogeneity between sessions within each treatment but, broadly speaking, they follow the basic overall patterns illustrated in figure 5.

---

[31] In the absence of a swap, the fundamental value declines over time due to the decreasing future expected lifetime of the asset. Swapping debt for equity increases the fundamental value of the asset, offsetting this effect to some extent.

[32] Figure 6: 1) First two rows: Baseline. 2) Third row: NoInstall. 3) Fourth row: NoDefault. 4) Fifth row: NoSwap. Graph legend: Left vertical axis: Price. Right vertical axis and green circles: number of trades. Horizontal axis: trading period. Red line: Average median trading price. Bold-dotted blue line: Fundamental value.

In most sessions, after several new share issues, the price begins to decline. This decline typically begins at approximately periods 6 – 8. In the three treatments other than NoSwap, prices exceed fundamentals by a considerable margin for an extended interval of time and in some sessions rapidly crash. In some sessions, the crashes are so severe that prices are below fundamentals at the time the market ends. This eventual drop in prices might be triggered by the impending end of the market by period 20 as well as the ever greater supply of shares on the market. The increasing supply lowers the quantity of cash available for purchases relative to the quantity of asset available to purchase. Reductions in this ratio tend to lead to lower prices (Caginalp et al., 2001; Haruvy and Noussair, 2006).

To measure bubble magnitudes and compare them across treatments, we use the RD and RAD measures, as proposed in Stöckl et al. (2010). These are particularly well-suited for situations in which different markets in the dataset are characterized by differing numbers of periods and/or by different fundamental values, as they are in our study. This is because they normalize for average fundamental value and number of periods.

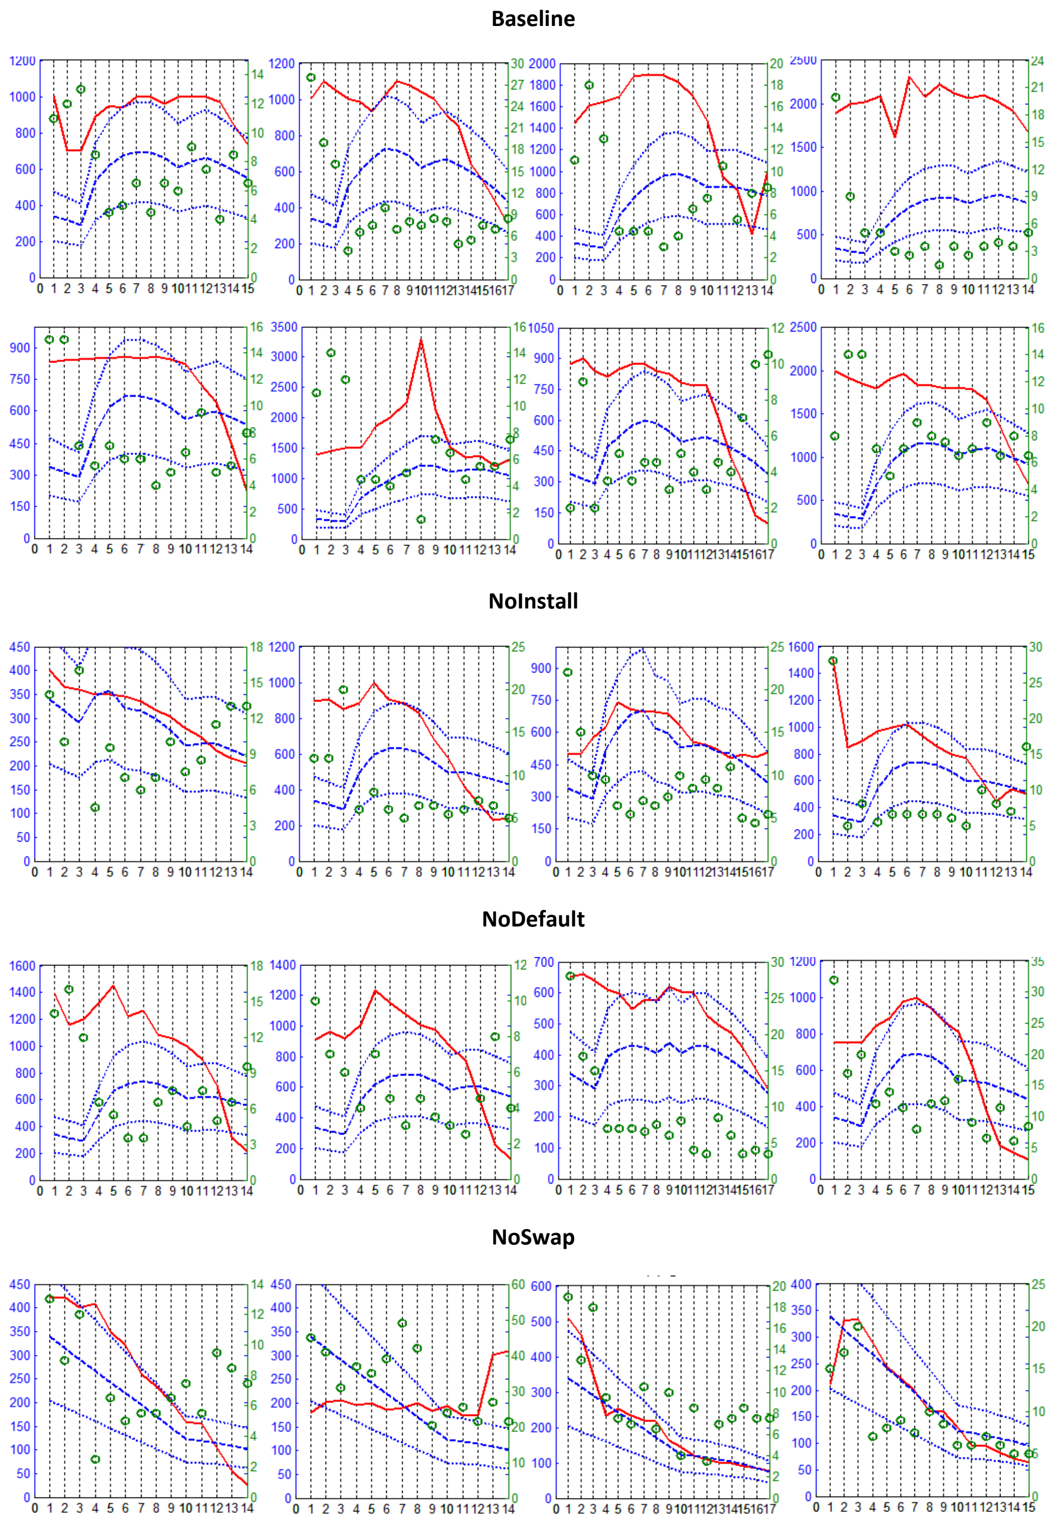

Figure 6: Results of each experimental session

The relative absolute deviation, RAD $= \left(\sum_{t=1}^{T} \frac{|P_t - FV_t|}{\overline{FV}}\right)/T$, measures the average level of mispricing compared to the average fundamental value of the market. The relative deviation, $RD = \left(\sum_{t=1}^{T} \frac{P_t - FV_t}{\overline{FV}}\right)/T$ measures the extent of over or under-valuation. The values of RD and RAD observed in the second horizon of each session are indicated in table 10, along with the treatment averages.

Treatment	**Baseline**	**NoInstall**	**NoDefault**	**NoSwap**
	0.610	0.093	0.952	0.379
	0.608	0.602	0.758	0.422
	1.069	0.170	0.425	0.181
RAD	1.744	0.527	0.580	0.135
	0.511			
	0.917			
	0.627			
	0.939			
Average	0.878	0.348	0.679	0.279
Treatment	**Baseline**	**NoInstall**	**NoDefault**	**NoSwap**
	0.610	0.066	0.797	0.277
	0.568	0.410	0.533	0.049
	0.984	0.162	0.425	0.149
RD	1.744	0.483	0.291	-0.057
	0.399			
	0.917			
	0.459			
	0.904			
Average	0.823	0.280	0.512	0.105

Table 10: RD and RAD by session - 2nd horizon

To investigate whether these differences are statistically significant, we run four (2-tailed) Mann-Whitney U-tests to test for differences in RD and RAD between the Baseline and each of the other treatments. Table 11 reports the z-scores and the resulting significance levels, for each pairing of treatments. The significant difference between the Baseline and NoSwap treatments provides the basis for our first result.

	RAD	RD
Baseline	0.878	0.823
NoInstall	0.348	0.28
Z-Value	2.378	2.208
Significance	2%	3%
N	12	12
Baseline	0.878	0.823
NoDefault	0.679	0.512
Z-Value	0.85	1.53
Significance	40%	13%
N	12	12
Baseline	0.878	0.823
NoSwap	0.279	0.105
Z-Value	2.717	2.717
Significance	< 1%	< 1%
N	12	12

Table 11: Mann-Whitney test for RAD and RD

Result 1: Redistributing revenues from new issues to existing shareholders increases bubble magnitude.

Support for result 1: Comparison of the Baseline and NoSwap treatments in table 11 indicates a strong and significant difference between the treatments. In all four sessions of the Baseline treatment, both RD and RAD are greater than in any session of the NoSwap treatment. The hypothesis that RD is equal in the two treatments, as well as the analogous hypothesis for RAD, is rejected at $p < .01$ ($z – 2.717$).

Figures 5 and 6 also indicate that prices are considerably lower in NoInstall than in the Baseline treatment. Result 2 is a statement that the difference is significant, and indicates that paying in installments does increase prices and bubble magnitude.

Result 2: Payment for new shares in installments increases prices and bubble magnitude.

Support for result 2: In every session of the Baseline treatment, both RD and RAD are greater than in any session of the NoInstall treatment. Mann-Whitney tests indicate that the Baseline treatment has higher RAD (z = 2.378, p = 2%) and RD (z = 2.208, p < 3%) than the NoInstall treatment.

Finally, the possibility of defaulting on debt if the life of the asset ends does not have a significant effect on both bubble measures.

Result 3: The ability to default on debts owed does not increase bubble magnitude.

Support for result 3: When the Baseline treatment is compared with NoDefault, there is no significant difference in RD (z = 1.53, p = 13%) or in RAD (z = 0.85, p = 40%).

In addition to comparing average prices, we can study the likelihood that a bubble forms. There is no consensus definition of a bubble, but the purposes of our experiment, we define a bubble as five consecutive periods during which the asset price remains more than a certain percentage above its fundamental value. That is, we say that a bubble occurs in a session if the session contains five consecutive periods in which median transaction prices exceed the fundamental by a threshold value. In table 12 we report the number of sessions with a bubble for threshold values of 30%, 40%, and 50% of fundamentals.

%	Baseline	NoInstall	NoDefault	NoSwap
30%	8	2	4	2
40%	7	2	4	1
50%	7	1	2	0
N	8	4	4	4

Table 12: Number of sessions with bubble in each treatment

Table 12 shows that the Baseline and NoDefault treatments produce bubbles in all sessions if the 30% criterion is used, and in at least 50% of sessions under the 50% criterion. NoInstall and NoSwap always generate fewer bubbles, independent of the cut-off used. Under the 50% criterion, the difference is particularly marked, with 7 sessions in the baseline showing bubbles – but only 1/2 in the NoInstall/NoDefault treatments. NoSwap generates no bubbles at all under the 50% criterion.

2.5.2 Individual behavior

In this section we explore how the institutional factors that we have highlighted interact with trading strategies of individuals to either magnify or dampen bubbles. We employ the profile of trader types first proposed by Delong et al. (1990), and applied to experimental data by Haruvy and Noussair (2006) and Haruvy et al. (2013).

In this structure, there are three types of trader. (1) *Fundamental Value Traders* use the fundamental value as a limit price. They increase (decrease) share holdings when the current price is below (above) fundamental value. (2) *Momentum Traders* increase (decrease) share holdings in response to an upward (downward) price trend in the recent past. (3) *Rational Speculators* correctly anticipate the next period's price movement. If the price move is upward (downward), they increase (decrease) current holdings of shares.

We define an individual's behavior as consistent with the Fundamental Value Trader type in period t if one of two conditions is met: Either (1) if ($p_t > FV_t$), then ($s_{it} < s_{i,t-1}$), or (2) if ($p_t < FV_t$), then ($s_{it} > s_{i,t-1}$), holds. In conditions (1) and (2), p_t is the average price in period *t*, FV_t is the fundamental value in period *t*, and s_{it} is the number of units of asset that individual *i* holds in period *t*. That is, the fundamental value trader purchases more shares than she sells in period t if the price is below fundamentals, and sells more than she buys if the price is above fundamentals. A market populated predominantly with such types will tend to have prices that broadly track fundamental values.

Her behavior is consistent with the Momentum Trader type in case either condition (3), if ($p_{t-1} < p_{t-2}$), then ($s_{it} < s_{i,,t-1}$) or (4), if ($p_{t-1} > p_{t-2}$), then ($s_{it} > s_{i,t-1}$), holds. The momentum trader accumulates units in period t if there has been an increasing price trend between periods $t – 2$ and $t – 1$, or sells units if there has been a decreasing trend. These types generally do not contribute to pricing close to fundamentals. Indeed, they tend to generate momentum – they cause a continuation of previous trends, which can aid in inflating a bubble once it has begun.

A trader's behaviour is consistent with the Rational Speculator type if her behaviour in period t satisfies one of the following conditions: (5) if ($p_{t+1} < p_t$), then ($s_{it} < s_{i,t-1}$), or

(6) if ($p_{t+1} > p_t$), then ($s_{it} > s_{i,t-1}$). This type of agent anticipates the price in the next period in an unbiased manner. She makes positive net purchases if the price is about to increase between the current and the next period. She makes net sales if the price is about to decrease. These traders anticipate future trends and can initiate and sustain bubbles. When they expect prices to increase, they demand more shares, and thus their prediction becomes self-fulfilling.

We first assess each person's trading record in each period t, using the prices from periods $t - 2$ and $t + 1$ and the change in her holdings of asset, $s_{it} - s_{i,t-1}$. We then classify her type by the trading style with which she is consistent for the greatest number of periods. If there is a tie between two types, we classify the trader as belonging to each type with proportion 0.5, and if there is a tie between all three types, he is assigned each type with proportion 0.33.

As a measure of how much influence an individual exerts in her market, we use two variables. The first is a measure called Market Portfolio Influence. The *Market Influence* of subject i in period t is defined as:

$$MI_{i,t} = P_t * s_{it} + c_{it}$$

Where s_{it} and c_{it} indicate shares and cash owned by subject i in period t, respectively while P_t indicates the median transaction price for period t. The second measure of an individual's impact on the market that we use is simply the share of the total outstanding shares that the individual holds, which is equal to s_{it}/s_t.[33]

Table 13 shows the average number of individuals in a market, classified as belonging to each of the three types, by treatment. The table shows that the percentage of traders classified as each type is comparable in the four treatments.

[33] Haruvy and Noussair (2006) and Haruvy et al. (2013) use another measure of an individual's weight in the market, which they call *Market Power*. This is a weighted average of the percentage of the total cash, and the total stock of asset, in the market that an individual holds. The results reported in this section concerning Market Influence are very similar if Market Power is used instead as a measure of an individual's weight in the market. The market influence measure is an index of the current market value of a trader's position.

	Baseline	NoInstall	NoDefault	NoSwap
Momentum	0.4	0.4	0.4	0.4
Fundamental value	0.3	0.4	0.3	0.5
Rational speculator	0.2	0.2	0.2	0.1

Table 13: Percentage of traders that are of each type, by treatment

The market influence of traders of each type is shown in figure 7. It reveals that the NoSwap treatment is the only one characterized by having Fundamental Value traders with more market influence on average than the other two types at all times during the trading horizon.[34] In each of the other treatments, Momentum traders have more market influence than the other two types for most of the trading session. In the Baseline treatment, momentum traders are dominant. NoDefault is similar in that momentum traders have the most influence throughout the trading session. In NoInstall, momentum traders have the most influence for most of the session, but fundamental value traders surpass them at the end. Speculators have similar market influence in all of the treatments, except for NoSwap, where they have less.

Does looking at individual-level trading behavior shed light on the origins of the South Sea bubble? The debt-for-equity swap/new issuance of shares raises the fundamental value of shares. This increases the price Fundamental Value Traders are willing to pay, by a corresponding amount. Momentum traders, on the other hand, when observing any increase in price – whether justified or not –increase their demand during subsequent periods. This tends to create and enhance bubbles. Speculators correctly anticipate this behavior of momentum traders (and in the case of new issuance, Fundamental Value traders). They buy in advance of these investors, during the early stages of the price boom. The role of the debt swap is important according to the laboratory evidence because it causes an initial boost in fundamentals. The (appropriately) higher price then leads to reactions by momentum traders and speculating investors that magnifies the run-up in equity values, increasing the size of the bubble.

[34] Periods 15,16, and -17 contain data from only one session per each treatment, because only one of the randomly-generated termination sequences lasted that long. This accounts for the abrupt changes in period 15 in some of the panels of figure 7.

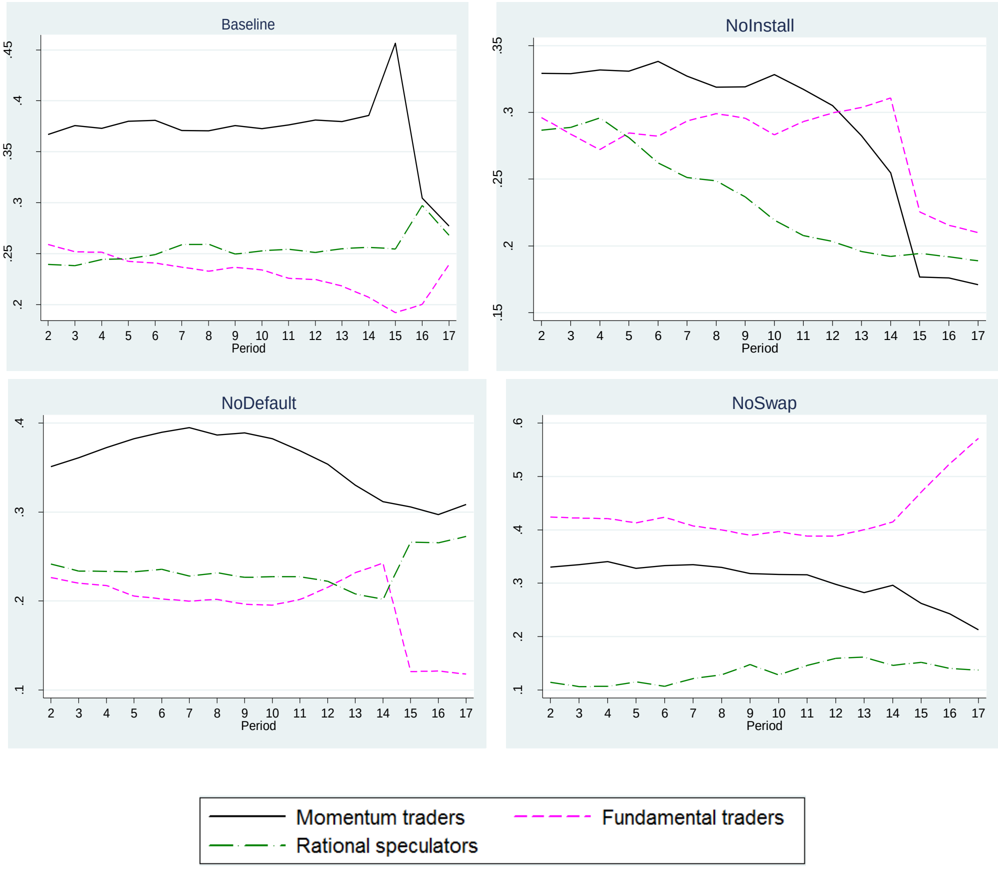

Figure 7: The market influence for an individual of each trader type, averaged over all sessions of a treatment

Table 14 illustrates how the proportion of units, as well as the market influence of each type, depends on the asset's return in the preceding period. The independent variables are the return of the asset from the previous period to the current period. This is defined as:

$$Return_t = \frac{P_t - P_{t-1} + d_{t-1}}{P_{t-1}}$$

The second independent variable is the order of the new issue. This variable is equal to the number of issues that have occurred up to and including the current period. In the first two columns, the market influence and percentage of total shares held by rational speculators are the dependent variables. In the last four columns, the analogous variables for fundamental value and momentum traders are the dependent variables.

	Rational Speculators		Fundamental Value Traders		Momentum Traders	
	market portfolio influence	% of shares owned	market portfolio influence	% of shares owned	market portfolio influence	% of shares owned
Return	0.142***	0.235***	-0.245***	-0.209***	0.1085**	-0.0116
	-0.0565	-0.0833	-0.0741	-0.0643	0.0559	-0.0044
Order	0.00466	-0.0039	-0.0069	-0.0082***	0.0006	-0.0044
	-0.00267	-0.00387	-0.0069	-0.00299	0.0026	0.0039
Constant	0.186***	0.253***	0.348***	0.249***	0.33***	0.407***
	-0.0161	-0.0238	-0.348	-0.0183	0.0159	0.0244
Observations	182					

Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1

Table 14: Market Influence and quantity of shares held by each type of trader as a function of return and number of share issues

The estimates reveal several interesting patterns. High returns in the previous period attract purchases from rational speculators and momentum traders; they induce sales by fundamental value traders. These purchases increase the market influence of momentum and speculators and decrease those of fundamental value types. Low returns have the opposite effect, and induce the rational speculators to sell to fundamental value traders. The results for market influence imply that rational speculators on the whole are right in their prediction that past returns during our bubble experiment predict future returns. Fundamental traders, on the other hand, lose out – their influence decreases as they "lean against the bubble".

We also find an additional effect from share issuance. As more new issues come to the market, all else equal, fundamental value traders hold fewer units and exercise less market influence. This suggests that the degree of return predictability induced by the debt-for-equity swap weakens the influence of fundamental traders, creating opportunities for trend-chasers and for rational speculators.

The data in table 13 and 14 are consistent with the following account of how new issues lead to a bubble. Momentum traders are typically the modal type of agent in our markets. If a new issue occurs, there is an immediate increase in fundamental value, and thus in the return on the asset. Fundamental-driven investors bid up the price of the asset. In a market with many momentum traders, this price increase will continue for a while since momentum traders are attracted to the increasing trend. The anticipation of an upward trend leads rational speculators to also make purchases, further driving up the price. Once prices get to be greater than fundamentals, it is the fundamental value traders who sell to speculators and to momentum traders. The price increase accelerates over time as more and more shares are issued since these are purchased by momentum traders and speculators. As Haruvy et al. (2013) argue, increases in the weight in the market of fundamental values are closely related to greater adherence of prices to fundamentals. The new issues serve to lower the market influence of fundamental value traders in the South Sea experiment because they create return predictability, drawing in momentum traders and rational speculators; this is ultimately the cause for the breakdown in the link between prices and intrinsic value. Under NoSwap, in which new issues have no effect on fundamental value, the process described here does not begin and prices remain close to fundamentals.

2.6 Discussion

The South Sea bubble and crash was one of the most spectacular episodes in financial history. Many idiosyncratic features of stock trading in the 1720s have been proposed as contributing to the bubble and crash, including corrupt allocation of shares to influential individuals, misleading information issued by the company about its prospects in the New World, and contagion from concurrent bubbles in Amsterdam and Paris. In this paper, we focus on specific institutional features of the asset market itself and analyse if these play a role in promoting bubble formation more generally. The use of laboratory experiments allows us to examine the impact of each aspect individually and in combination, allowing a precise determination of which market features promote instability the most.

Several historical features contribute to the formation of large bubbles in our experiments. At its core, the South Sea bubble is about a debt-for-equity swap, wherein the company offered to exchange government debt for its own shares (with the support of the UK government). At its peak, the South Sea company held 23% of Britain's entire stock of public debt. The key manner in which this changed incentives to shareholders is that each issue of new shares meant that more interest payments were paid out to shareholders. In our NoSwap treatment, in which these payments did not occur, we observe no bubble. When no swapping of debt was possible, the other features that could enhance the formation of bubbles, risk shifting in the form of delayed payments for new shares, or the chance to default on these payments, did not generate a bubble.

In our experiments, we classify investors, according to their predominant behavior into momentum traders, rational speculators, and fundamental-driven investors. Crucially, the swapping of equity for debt raised the fundamental value of shares *for old owners*. This meant that new subscriptions of shares raised the expected value of future dividends, and thus justified an increase in prices. An increasing price in turn attracted momentum investors and speculators. Because these buy in parallel with the fundamental investors, prices increase and can exceed fundamentals for some time. The increase in fundamentals induced by particular contractual features of the 1720 debt swap is therefore crucial for generating the rapid rise in prices. In particular, the fact that the South Sea Company did not fix the "exchange rate" between stock and debt ex ante allowed it to buy out bond holders ever more cheaply (in terms of stock) as long as the stock price kept increasing. Interestingly, a fixed exchange rate was present in 1719, when the South Sea Company swapped its shares for so-called "lottery tickets", another kind of government debt (Carswell, 1960). At that time, there was also no explosion in share prices.

In the experiment, when there was swapping and new issuance took place, the ability to delay payment on newly-issued shares further increased the likelihood and magnitude of bubbles. The possibility of leverage raised the return to speculation in treatments with swapping because the swapping induced an increasing trend in prices. Without the establishment of an initial increasing price trend, higher leverage would not enhance the returns to speculation.

More generally, a large literature has argued that bubbles are caused by initially fundamentals-driven, rational increases in share prices that become exaggerated through feedback loops in the market (Shiller, 2000) – such as in the case of electrification, the automobile, and aviation in the 1920s, or the internet in the 1990s. Our evidence about the importance of the debt-equity swap, which boosted the intrinsic value of South Sea stock, appears to be in line with this interpretation. Note also that the laboratory findings coincide precisely with Charles Kindleberger's (1987) classic definition of the origins of a bubble:

> "A bubble may be defined loosely as a sharp rise in price of an asset …, with the initial rise generating expectations of further rises and attracting new buyers—generally speculators interested in profits from trading rather than in its use or earning capacity."

Overall, the experimental evidence lends support to classes of bubble models emphasizing investor heterogeneity. For example, in Hong, Scheinkman, and Xiong (2006), investors receive different signals about the fundamentals of an asset. Bubbles can form because the resale (option) value to the other group induces even those who are pessimistic about fundamentals to buy. This is similar to the interaction between fundamental-driven investors, rational speculators, and momentum traders in our experiment, which requires an increase in fundamental value, resulting from the debt-equity swap, to ignite a major speculative frenzy.

2.7 Appendix

Instruction

Baseline treatment instruction

The experiment consists of a sequence of trading periods, each one lasting 150 seconds. During each period, you will have the opportunity to buy and sell units of two assets of an imaginary Company named "Blue River" in a market. The first asset is called "Share A" while the second is called "Share B". During the instructions we will explain differences between them. The currency used in this experiment is called *ducats*, which will be converted into Euros at the end of the experiment at an exchange rate of xx DUCATS for 1 Euro. You will have to participate in two separate markets. At the end of the experiment one market is randomly selected and your payment will correspond to this selection.

How to buy and sell shares

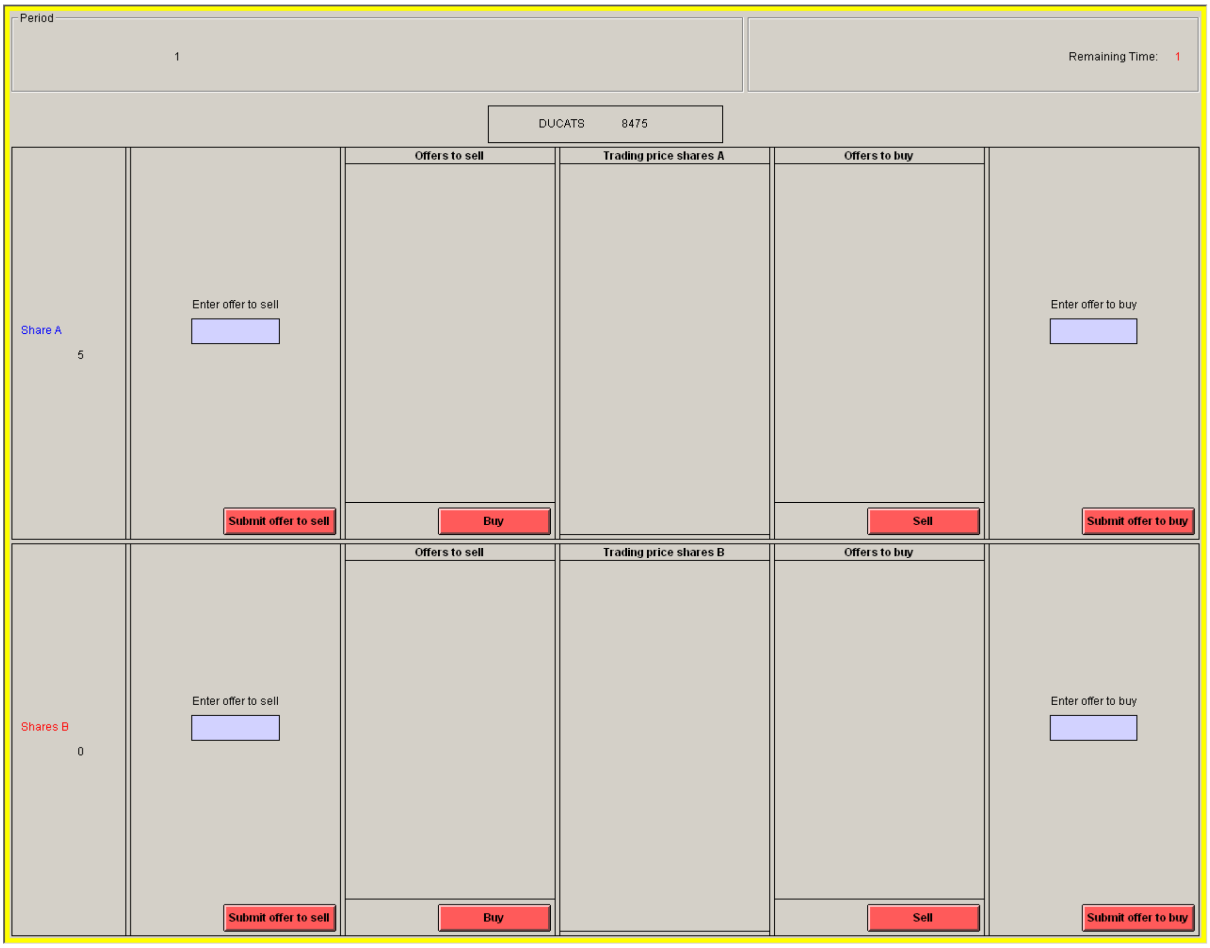

This is the trading screen you will use during the experiment.

You start the experiment with a quantity of money called "*DUCATS*" and a number of shares "A". You will not start with any share "B".

We will explain how trade occurs in the market for "Share A", which occurs exactly in the same way as for "Share B". The top left corner shows the current trading period, and the top right corner shows how much time (in seconds) is left in the current period. Your ducat balance is shown in the middle of the screen. Using this screen, you can buy or sell shares in four ways. First, you can initiate a sale of shares by submitting an offer to sell. If you have shares, you may choose to sell them. You can initiate a sale in the text area below "Enter offer to sell" in the first column. Here you can enter the price at which you are offering to sell a share. To send the offer, you have to click the "Submit offer to sell" button. After that, your offer to sell will appear in the second column labelled "Offers to sell". Each offer introduced corresponds to one single share. If you want to sell more shares, repeat this process.

Note that by submitting an offer to sell, you initiate a sale, but the sale will not be executed until someone accepts it.

Try offering to sell a share now. Enter a number in the text area labelled "Enter offer to sell" in the market for "Shares A" and then click on the button "Submit offer to sell". You can see that a set of numbers will appear in the column labelled "Offers to sell". Each number corresponds to an offer from one of the participants. Your own offers are shown in blue; others' offers are shown in black. The offers to sell are ranked from high to low, so that the cheapest (best) price is displayed at the bottom of the list.Second, you can buy shares by accepting an offer to sell. If you have enough *ducats*, you can buy a share at one of the prices in the "Offers to sell" column (which also contains your previously submitted offer to sell). You buy a share by selecting one of the others' offers (shown in black) and then clicking on the red button "Buy". Note that you are not allowed to accept your own offers, which are shown in blue. Remember that the cheapest (best) price is displayed at the bottom of the list.It may happen that when you select the best price and press the "Buy" button, someone else is doing the same thing but acting faster than you. In that case, a message "someone has been faster than you" or "you have to select a price" will appear.

Try buying a share now. Choose a price in the “Offer to sell” column and then click on the “Buy” button, or click directly on the “Buy” button and buy at the cheapest price listed in the “Offers to sell” column. Whenever an offer is accepted, a trade occurs. When you accept an offer to sell, you realize a purchase and the number of *ducats* in your ducat balance goes down by the transaction price; at the same time, your trading partner makes a sale and his/her ducat balance increases by the trading price. In contrast, when your offer to sell is accepted, you make a sale, your trading partner makes a purchase, and an amount of money is transferred from your trading partner to you that is equal to the amount of the trading price. Because you have each submitted one offer to sell and accepted one offer to sell, you have all realized one purchase and one sale so that you have the same number of shares as you started out with.

Third, you can initiate a purchase of a share by submitting an offer to buy.

If you have *ducats* and would like to buy a share, you can initiate the purchase by submitting an offer to buy. Enter a number in the text box under “Enter offer to buy” situated on the right side of the screen and then click on the “Submit offer to buy” button.

Try submitting an offer to buy a share now. Enter a number in the text area “Enter offer to buy” in the market for “Shares A” Then press the red button labelled “Submit offer to buy”. Immediately, in the column labelled “Offers to buy”, you will see a list of numbers ranked from low to high, so that the highest (best) price is displayed at the bottom of the list. If you want to sell more shares, repeat this process. Again, your own offers are shown in blue; others’ offers are shown in black.

Fourth, you can sell a share by accepting an offer to buy. You can sell a share at one of the prices offered in the “Offers to buy” column (which also contains your previously submitted offer to buy). Select one of the offers and then click on the red button “Sell”. Again, note you are not allowed to accept your own offers (shown in blue). Remember that the highest (best) price is displayed at the bottom of the list. *Try selling a share now*. Choose a price in the column “Offer to buy” and then click on the “Sell” button.

A transaction occurs whenever an offer to buy is accepted. If you accept an offer to buy posted by another person, you make a sale and as a result, the amount of *ducats* you have increases by the trading price. In contrast, when your offer to buy is accepted by someone else, you realize a purchase and the number of *ducats* you have decreases by the trading price. The opposite happens to your trading partner. You can see that these four trading methods are complementary: you can initiate a trade by offering a price to sell or buy and wait for the offer to be accepted by others; you can execute a trade by accepting an offer to buy or sell submitted by another participant. In the column situated in the middle of the screen and labelled "**Trading price**", you can see the prices at which shares have been traded during the trading period by all participants playing in the market.

Information display at the end of each period

At the end of each trading period a screen appears that summarizes your situation.

"End-of- period screen"

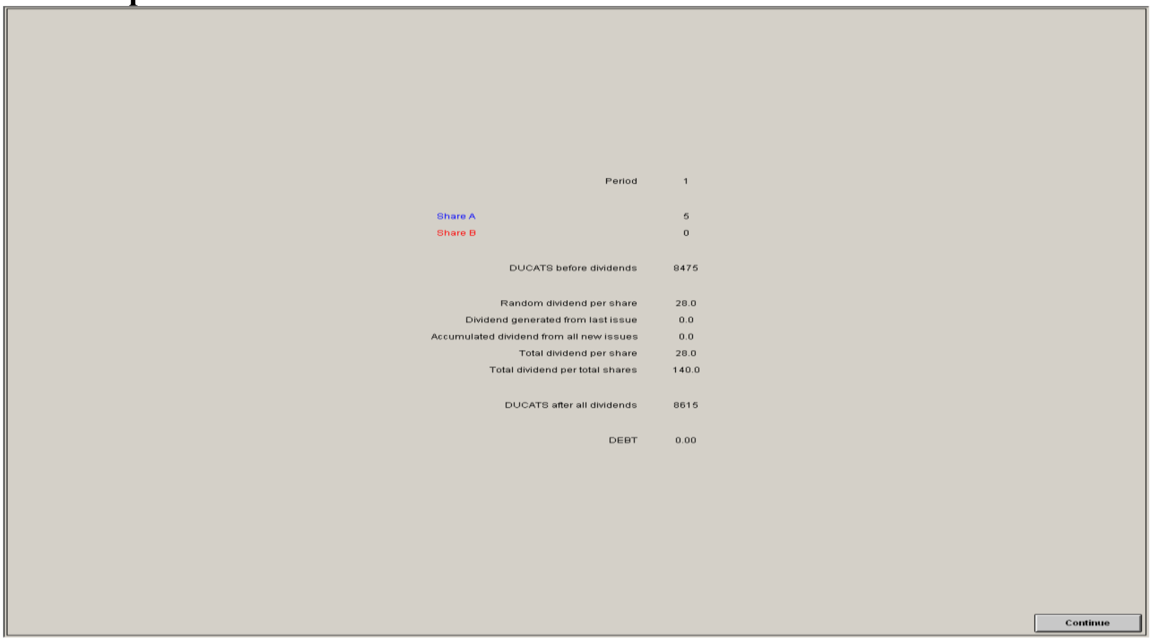

In this screen you have the following information:

Period: This is the trading period just finished.

SHARE A: The number of shares of type "A" you own at the end of this period.

SHARE B: The number of shares of type "B" you own at the end of this period.

DUCATS before dividends (dividends will be explained later): This is the money you have before you receive the current period's dividend payment.

Random dividend per share: As explained later, each share you own may pay a dividend at the end of each trading period. The random dividend per share is the amount in *ducats* you get in this trading period for each share you own.

Dividend generated from last issue: Share issues, in which the experimenter makes new B shares available, will be described later in the instructions. Each time that new shares are issued on the market, the dividend you receive on both shares A and B increases. This *dividend generated from last issue* is the increase in dividend that results, from the last issue of shares only. If no shares are issued it is equal to 0.

Accumulated dividend from all new issues: As will be explained later, during the experiment new shares may be issued. Each time that new shares arrive on the market, if you decide to buy them, the amount of dividend you receive on the shares that you own increases. The Accumulated dividend from all new issues is the total of all dividends derived from all new issues of shares.

Total dividend per share: This is the sum of E + G.

Total dividend x total share: This is the total dividend per share [H] times the number of shares (shares A + shares B) you own since each share "A" and "B" receives the same dividend per period.

Ducats after all dividends: This is the sum of total dividend x total shares plus your Ducats before dividend. (D + I).

Debt: This will be explained later.

Random Dividend

For each share you own of both type of shares "A" and "B" you receive a dividend at the end of each trading period. A random device determines the dividend you get for the period for each share you own.

The amount and the chance of each possible dividend are the following:

25% chance you get 0 *ducats* per share you own

25% chance you get 8 *ducats* per share you own

25% chance you get 28 *ducats* per share you own

25% chance you get 60 *ducats* per share you own

Each participant gets the same dividend per share at the end of the period. The average dividend in a period is equal to 24 *ducats*. This is calculated as following: (0 + 8 + 28 + 60)/4=24.

How long the market lasts

The market lasts for a minimum of 10 and a maximum of 20 trading periods, each period lasting 150 seconds. In each of the periods between 10 and 20 the experiment could end. Whether it ends or not is determined with a random device. Beginning in period 10, at the end of every period there is:

an 83% chance that the market continues to the next period, and

a 17% chance that the market ends immediately.

This means that dividends from period 10 onward are not guaranteed because the market might end.

Your average holding value table

Your "Average holding value" table indicates the total of the dividends you would get on average if you held a share from any period until the market ends, <u>if there are no new shares issued (</u>i.e. subscribed) you can use this table to help you make decisions.

1 Current period	2 Average random dividend per share **this period**	3 Average **Remaining** Dividends per share	4 Chance next period occurs
1	24	339	100%
2	24	315	100%
3	24	291	100%
4	24	267	100%
5	24	243	100%
6	24	219	100%
7	24	195	100%
8	24	171	100%
9	24	147	100%
10	24	123	83%
11	24	119	83%
12	24	115	83%
13	24	109	83%
14	24	103	83%
15	24	95	83%
16	24	86	83%
17	24	74	83%
18	24	60	83%
19	**24**	**44**	**83%**
20	24	24	0

The columns in the table refer to the following:

Current period: The period corresponding to the number in the row. For instance for period 2 you would receive on average total future dividends of 315 for each share you held until the end of the market in case there are NO issues of new shares.

Average random dividend per share this period: The average amount of dividend per period for each share you own. This is equal to 24 in each period as explained above.

Average remaining dividends per share: The average amount of total dividend you will receive for each share that you own from now until the end of the experiment. It is the average amount of dividend (24 *ducats*) multiplied by the average remaining number of periods. From period 11 on, there is an 83% chance that the market will continue for another period. The calculation of the average remaining dividends takes this into account. For instance if you are currently in period 5, the average expected dividend you would get from owning one share until the end of the experiment is 243.

Chance next period occurs: The probability that the following period exists.

Issues of type "B" Shares

Beginning in period 3 you might have the opportunity to buy additional shares of type "B" from the experimenter. Type B shares pay the same dividends as type A shares and can be bought and sold in the same way as A shares. B shares are issued in some rounds for a price determined in an auction, which is described below. However, you will not be informed in advance about the total number of times or in which periods B shares will be issued.

If an issuing round takes place (i.e. a screen like this appears), you have to participate to an auction.

You can make bids on up to 2 B shares.

You must bid on two shares. If you do not want to buy a share, you can type in a bid of 0 in both spaces. If you are sure that you don't want to buy any more than 1 unit, you can type 0 in the box labeled "Bid 2" and enter the bid for the one unit you are interested in bidding for in the box labeled "Bid 1"

Once you and the other participants have submitted their bids, a new screen appears.

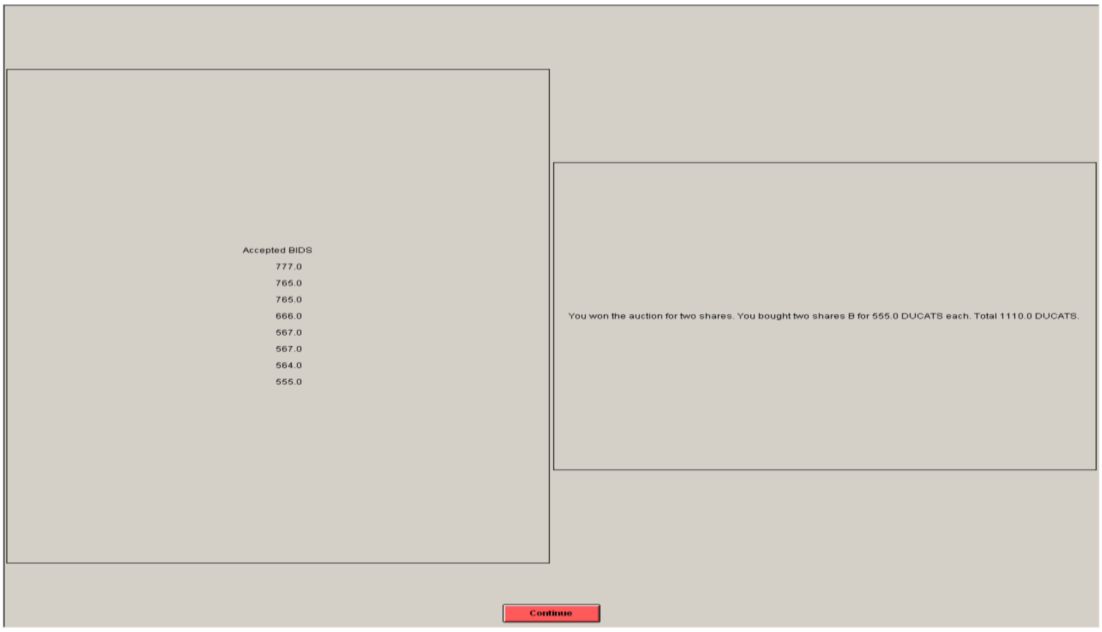

On this screen you can see the accepted bids from all participants in the market, ranked in order of how large they are. The people who sent in these accepted bids receive a unit for each bid they have accepted. If you have one bid on the list, you get 1 unit, and if you have 2 bids on the list you get two units.
The last price on this list, which is the lowest accepted bid, is the amount of Ducats that you will have to pay for each of your bids on the list.
On the right corner of the screen you can see whether you purchased one or two newly issued shares B, and the price you pay for each share.
If you purchase newly issued shares “B”, the experimenter **lends you some of the money to buy them.** This means that you do not have to have enough ducats on hand at the moment to buy the shares, because you pay most of what you owe for the shares later on.

In the period where you buy the newly-issued B shares from the experimenter, you pay 20% of the auction price. The remaining 80% are debt that you owe to the experimenter. In each of the following 4 periods, 20% of the price will be subtracted from your ducat balance. Your debt is therefore paid off over 5 periods. The amount of debt that you have at any time will be reported on your “After period screen” after each period in the field labeled DEBT.
If the market ends before you pay back your debt, you do not have to pay back your remaining debt.
For example. Suppose that there are two issues of shares “B”, one in period 4 and one in period 12. Suppose that the game ends after period 13, and that the issued price of each share was of 300. You purchased one B share in period 4 and one in period 12.
In period 8, you will have fully repaid your debt on the first share (20% in period 4 and 20% in each of the following four periods).
In period 13, when the market ends, you will have paid back 120 ducats on the share you purchased in period 12. The 120 is equals to 20% of the 300 plus the 20% of the 300 of period 12 and 13 respectively. **If the market ends and you still owe debt (as it is the case in this hypothetical example), you do not have to repay it.**
Notice that after the unit has been purchased in the auction, <u>A and B shares are the same from the point of view of any new buyer</u>. The two types of share pay the same dividends

in each period, and the two types of shares will continue to exist until the game ends. Because the two types of share are the same, an individual can make a profit if she can sell one at a higher price and buy the other at a lower price. This is because the asset she bought and the asset she sold will always pay the same dividends and she has more cash than he did before. For example, suppose that you sell an A share at 300 Ducats and you purchase a B share at 200 Ducats five seconds later. You then have the same number of shares as before, but have 100 Ducats more than before, so you have made a profit of 100 Ducats.

Increasing in the amount of dividend received.

The number of issues is unknown. Each time that you subscribe to new shares B, "Blue River Company" is using the collected money to buy investment certificates in a secondary market. The higher the price at which Blue River shares are issued, and the more buyers there are for these new shares, the more of these investment certificates can be bought. The dividends from these certificates will be added to the ones paid out by Blue River on share A and B.

For example: suppose that in period 10 the Company offers the possibility of buying new shares B. Therefore you participate in an auction. Imagine that the result of the auction is that the price for one share B is 320 ducats each. If all of you in total have bought 8 shares, Blue River will gather 2560 ducats (320 *8).

With this money the Company buys 25.6 certificates (2560 ducats/100). Each certificate pays a dividend of 15 ducats. Therefore a total of 384 ducats will be divided for the total number of shares outstanding in the market and paid to each of you as dividends after each market.

This means that if in the market there is a total of 30 shares (22 initials and 8 new issued), you receive additionally to the random dividend (explained in the instructions) a new increment of 12.8 ducats for each share you own (384/30).

This means that the value of owning both A and B shares increases as the total number of B shares in the market increases.

If B shares are issued, the numbers "Average remaining dividend per share" column will change. For example, after a single issuing round of shares "B", in period 4, for a price of 320 ducats (the same example as above), the "Average remaining dividend per share" column of the row 4 will change from 267 ducats to 351 ducats. The entire column will change automatically, and the average remaining dividends will be increased again each time new shares are bought.

A table called “Average remaining dividends per share” will be displayed at the end of each period (see your computer screen). This table is the current updated version of the “Average holding value table” described earlier and should be read exactly in the same way. You might want to use the information in the table to see how “Average remaining dividends” are updated after each issuing round.

In the screenshot shown here you can see how “average remaining dividends per share” has been affected after a single share issue and compare it with the “Average holding value table” above. In the example reported below, at the end of period 3 there was an issue of shares and all values in the “Average remaining dividends” column for future periods increased substantially compared to the original “Average holding table”

CHAPTER 3

THE EFFECT OF INTRINSIC MOTIVATION AND REFERENCE POINT UPDATING ON EFFORT PROVISION: A LABORATORY EXPERIMENT

3.1 Introduction

Consider the situation in which the amount of piece rate payment a worker receives for his/her job changes through time in a non-performance contingent way.

This can be quite common, especially for some types of low level seasonal jobs including fruit pickers or tree-planters.

In these types of jobs, a worker is often paid for example based either on the number of boxes he/she manages to fill each day or the number of trees he/she plants. Thus, he/she receives a salary that depends on the level of effort he/she exerts.

Often, the amount of piece rate that the employer pays changes substantially from season to season, on a monthly base, or from one crop to another.

There can be several reasons for these changes. For example, the market price at which a specific good is sold is lower this season compared to the previous one, or because planting conditions vary substantially among soils (Paarsch and Shearer, 2009).

Considering the high workforce mobility in these types of jobs, it could also be a case of the worker being aware that in the previous seasons a specific plantation was paying a different piece rate amount.

Thus, even if that worker was not employed there previously, he/she might arrive with some expectations formed which differ from the reality.

We believe that an important issue is understanding if these changes in piece rate from a previous level can affect worker's motivation and, consequently, the level of effort he/she is willing to exert.

The answer to such questions should be of particular interest to an employer who is looking for the best available contract that maximizes firm's profit.

Despite the ease of finding data regarding employees' piece rate amounts received and their performance, establishing a causal relationship is very hard. This is because there are many unobservable factors that may affect this relationship.

The best way to control for these confounding factors and to eliminate endogeneity problems is to study the question in an environment in which changes in piece rate are completely exogenous.

For this reason, we decided to study how changes in the amount of piece rate incentive affect task performance by using a laboratory experiment.

In a three period laboratory experiment we implemented four different incentive schemes across three distinct periods.

(1) Fixed payment: participants received a constant and fixed amount of money in each period that did not depend on the number of correct or incorrect answers.

(2) Constant piece rate[35]*:* participants received a piece rate per correct answer and received a penalty for incorrect answers. The amount remained unchanged over the three periods.

(3) Peak piece rate: The piece rate received (and penalty) is doubled in the second period compared to the first and third periods.

(4) Valley piece rate: The piece rate received (and penalty) is halved in the second period compared to the first and third periods.

Each treatment is implemented in two different tasks. The first task was designed to be relatively more challenging while the second task was designed to be less challenging. This task difference was confirmed by subjects' average evaluations on a post-treatment interest/enjoyment scale.

We find that, consistent with standard economic theory, performance contingent monetary incentives affect subjects' effort provision. At the same time, as suggested by the literature on intrinsic motivation, we found that the effect of monetary incentives varies between tasks and is much stronger in the less challenging task.

We also obtained that a unique, positive change in piece rate leads to an increase in subjects' performance, but only in the less challenging task.

[35] p.r. is often used in the future as abbreviation for the words "piece rate".

Finally, in line with literature based on reference-dependent outcomes, we discovered that a decrease in piece rate that follows an increase negatively affects performance on both tasks, while an increase that follows a decrease does not significantly affect performance.

The main message that we have derived from our results is that employers of especially boring and repetitive types of jobs should not increase the amount of piece rate if they cannot maintain the increase over time.

3.2 Literature review

In this section we will refer to four different but related literatures from which we derived our set of hypotheses.

3.2.1 Effect of monetary incentive on performance

In economic theory, incentives serve as an instrument to align principal-agent objectives (Holmstrom, 1979; Kale et al., 2009; Gibbons, 1992).

For an employer, it is therefore a common practice to use financial incentives to improve employees' performance and thus increase the firm's profit.

Several papers support this view.

Lazear (2000), for instance, found that the average level of output per worker increases between 20% and 36% when switching from hourly wage rate to performance based payment.

Banker et al. (1996) analyzed a panel of 15 retail outlets and found that sales increased when performance-based compensation was introduced.

In a field experiment, Hossain and List (2009) found that incentives lead to higher productivity, especially for teams.

In laboratory experiments, Libby et al. (1992) found that subjects worked harder and were more accurate on a recall and recognition task if rewarded for performance. Fehr, Goette and Lienhard (2008) showed that effort provision of participants incremented by 15% when piece rate substituted fixed wages.

Despite much evidence, several other papers found that under some conditions, or in some specific tasks, monetary incentives do not significantly affect workers' performance (see Camerer and Hogarth, 1999 and Bonner and Sprinkle, 2002 for a comprehensive review of the effect of incentive in performance).

3.2.2 Incentive and intrinsic motivation

One situation in which the presence of monetary incentives might affect performance differently from what economic theory predicts is, for instance, when workers are intrinsically motivated by doing the task (Frey, 1997).

According to Deci (1971) "One is said to be intrinsically motivated to perform an activity when he receives no apparent reward except the activity itself" (p.105).

In a famous experiment, Deci, Koestner, & Ryan, (1999) showed that adding monetary incentives (i.e. extrinsic reward) for an intrinsically motivated task can "crowd out" motivation and consequently undermine individuals' effort.

But most of the time extrinsic incentives co-exist with intrinsic motivation, given that a person could enjoy doing a task and at the same time be paid for doing it.

Therefore, recent research has tried to assess how the presence of intrinsic motivation interacts with the presence of monetary incentives and which is the most efficient combination between them (Cerasoli et al, 2014).

Results suggest that monetary (i.e. extrinsic) incentives are generally more powerful for very repetitive and boring tasks since they are characterized by low levels of intrinsic motivation (Cerasoli et al, 2014).

3.2.3 Incentive magnitude and incentive variations

There are two other questions that have received attention in the incentives literature. The first investigates how *absolute magnitude* of monetary reward affects performance, and the second investigates what people's responses are to *relative changes in the magnitude of incentives*.

Regarding *absolute magnitude,* laboratory evidence showed mixed results.

For instance, Gneezy and Rustichini (2000) found that the effect of monetary compensation on performance does not monotonically increase as one might think.

In fact, they obtained that for an IQ test task, subjects who were paid a small amount of money for a correct answer performed poorer than subjects that were not paid at all. In order for monetary incentives to matter they observed that the magnitude had to be sufficiently high.

On the other hand, Pokorny (2008) and Takahashi et al. (2014) found the opposite phenomenon. For example, Pokorny (2008) found that in a laboratory experiment subjects exerted more effort when they received very low incentives compared to when they were not paid or when they were paid a lot.

With respect to the effect of *relative changes in the magnitude of incentives*, results are less mixed and seem to converge in one direction.

For instance, in a field experiment Paarsch and Shearer (2009) analyzed how positive changes in piece rate amount affect the performance of tree-planters. They observed that an increase in piece rate leads to a significant increase in performance.

Dikinson (1999) conducted a laboratory experiment and was able to disentangle income from substitution effects when varying wages. For half of the subjects he substantially increased the piece rate amount from the first to the second period, while doing the opposite for the other half. He found that when subjects could only vary their amount of effort and not the duration of work, a compensated wage increase (decrease) made subjects work harder (less hard).

3.2.4 Prospect theory and reference point updating

The results obtained by Dikinson (1999) and Paarsch and Shearer (1999) are very much in line with what standard economic theory predicts.

At the same time, starting with the work of Camerer et al. (1997), several empirical studies found evidence of negative temporal substitution in response to temporary wage shocks. For instance, Camerer et al. (1997) found that taxi drivers reduced the number of working hours following an increase in their per-hour salary. These results clearly go against standard economic theory and could be explained with behavioral models based on prospect theory (Kahneman and Tversky, 1979).

According to prospect theory, individuals' preferences depend not only on the achieved income level but also on some relative reference level (or reference point).

In the case of taxi drivers for instance, as pointed out by Camerer et al. (1997), the reference level was provided by the targeted income a taxi driver made on an average day. As a result, he/she would diminish his effort by reducing the number of hours worked once that level was attained.

Both prospect theory and prospect theory based reference dependence models (e.g. Koszegi and Rabin, 2006) make clear predictions regarding individuals' outcome evaluation from a reference point.

They furthermore suggest plausible candidates as reference points such as individuals' status quo or individuals' expectations.

Despite that, they do not say much regarding how reference points are updated, shift and evolve over time.

Recent research attempts to fill this gap (Baucells et al. 2011; Arkes et al. 2008, 2010; Chen and Rao 2002).

For instance, Chen and Rao (2002) proposed a theoretical model and a set of laboratory experiments to understand how people's reference points are sequentially updated.

In their work they show that the order in which reference points changes can critically affect their final evaluation.

More specifically, they studied the question of how individuals feel after being exposed to a sequence of two events of equal magnitude but opposite direction (i.e. gain followed by a loss vs. a loss followed by a gain).

They show, both theoretically and empirically, that the order in which two equivalent events occurs matters. In fact, they found that a gain followed by a loss will be evaluated less favorably (i.e. generate a higher loss in individual utility) compared to a loss followed by a gain[36].

According to them, the reason for the difference in the evaluation is due to the way people update their reference points.

In particular, they point out that a first outcome moves the reference point in a particular direction. Thus, a sequence of events, in which an initial event is surprisingly reversed, will be more favorable if the first event is a loss than if it is a gain.

[36] This result finds additional support in Baucells et al. 2011.

3.3 Description of the tasks

As we have seen in the literature reviewed in section 3.2.2, the effect of monetary incentives critically interacts with the degree of intrinsic motivation that a task generates. For this reason it becomes particularly important for our study to control whether intrinsic motivation also affects subjects' responses to different dynamics of incentives.

Therefore, we implemented the same set of four treatments on two differentiated tasks.

The first task is assumed to be more challenging. Subjects should derive relatively more pleasure from the task itself and should, therefore, be relatively more intrinsically motivated by doing it.

The second task is assumed to be relatively less challenging and more boring. Therefore, subjects should derive relatively less intrinsic motivation from doing it.

3.3.1 The more challenging task

From the experimental literature in social sciences, we observe that when the experimenter intended to assign subjects a challenging and effortful task he/she often proposed exercises having subjects perform non-trivial mathematical operations.

This is, for instance, the case of Brüggen and Strobel (2007) who had subjects multiply sequences of two digit numbers; or, for example, in Niederle and Vesterlund (2007) where subjects had to add sets of 4 two-digit numbers.

In our experiment, we construct a mathematical task that consists of downward counting from a very large number (e.g. 1,500,000) while continually subtracting a fixed quantity from the previous number.[37]

In particular, the numbers to be subtracted were 13, 17 and 27 respectively for the first, second and third periods. [38]

[37] We could also have used the same task as Brüggen and Strobel (2007) or as Niederle and Vesterlund (2007). At the same time our task is conceived to be relatively more entertaining and for this reason might perhaps induce a relatively higher level of intrinsic motivation (see footnote 38).

[38]In a previous pilot session (with 15 paid subjects) we individually asked (with open type of questions) to the subjects their perception about the task. Their answers suggested that they perceived the task as challenging and entertaining.

3.3.2 The less challenging task

In order for a task to be effortful and yet not constitute a challenge for subjects, it needs to be sufficiently boring, meaningless and repetitive.

From the experimental literature in social science we find that typical tasks with these characteristics involve, for instance, counting quantities of given numbers or letters within some sets. For example Abeler et al. (2009) made subjects count the number of zeros contained in a matrix consisting of 1s and zeros, or for instance Rey-Biel et al. (2013) made subjects count a specific letter within a fixed sequence of sentences.

Our task is similar to the one proposed by Rey-Biel et al (2013). Subjects had to count the number of A's contained in paragraphs where, in addition, the font was particularly small[39]

In our experiment, the exact same task was carried out in the three periods, each period lasting 4 minutes[40].

3.3.3 Post –treatment task evaluation

In order to insure that tasks differed on the level of intrinsic motivation they generated, we made subjects rate them by answering a post treatment questionnaire.

Specifically, we used the interest/enjoyment subscale taken from the Intrinsic Motivation Inventory scale (Ryan, 1982). This subscale is considered to be the self-reported measure of intrinsic motivation and has been extensively used, especially in the cognitive psychology literature (e.g. Ryan, 1982; Ryan, Mims & Koestner, 1983).

The interest/enjoyment subscale is composed of seven questions below. Each question was rated from 1 to 5. 1 = Not at all true, 2 = Not true, 3 = Somewhat true, 4 = True and 5 = Very true.

The questions were as follows[41]:

1. I enjoyed doing this activity very much
2. This activity was fun to do
3. I thought this was a boring activity

[39] A screenshot of both tasks is reported in the appendix with its description.
[40] Experimental sessions for the boring task were always run at the end of another experiment of an unrelated subject.
[41] Note that at the moment of calculating the average, the scale for questions 3 and 4 were reversed.

4. This activity did not hold my attention at all
5. I would describe this activity as very interesting
6. I thought this activity was quite enjoyable
7. While I was doing this activity, I was thinking about how much I enjoyed it

3.4 Experimental design

The experiment took place at the Leex (Laboratory of Experimental Economics) at the University Pompeu Fabra between spring 2010 and spring 2011. Participants were recruited via the Leex list of participants which is mostly composed of undergraduate students from all discipline. Participants' average age was of 21 and the number of males was approximately the same as the number of females.

The experiment consisted of eight treatments. The same four treatments were implemented in two different tasks in a 2x4 between subjects design. Tasks were programmed with Ztree (Fischbacher, 2007).

Each session was composed of 20 participants. Subjects recruited in each session were randomly assigned to one of the eight treatments (four treatments for each task).

Before starting each session the experimenter read aloud the general part of the instructions, this was common to all groups for the same task. After general instructions, subjects could familiarize themselves with the task through a trial period of thirty seconds.

Before beginning each period, in each subject's screen an individual text with instructions appeared. In this screen the critical information that changes between groups was the amount of piece rate participants received for each correct answer (and penalty for incorrect ones) for the period that was about to start.

Importantly, the change in piece rate was made particularly salient to subjects. For instance, in addition to reporting the amount of incentive in points, we used sentences such as: "Be careful! For this period the amount you earn per correct answer is doubled (or halved) compared to the previous period".

Furthermore participants also had to write the amount of piece rate per correct answer before starting each period in a piece of paper to return at the end of the experiment[42].
In each treatment, participants did not know that the amount of piece rate would change in following periods. They simply received the new information at the beginning of each period.

Treatment	Period 1	Period 2	Period 3
(1) Fixed payment	Fixed payment	Fixed payment	Fixed payment
(2) Constant p.r.	P* (correct answers) – (0.25*P* incorrect answers)	P* (correct answers) – (0.25*P* incorrect answers)	P* (number of correct answers) – (0.25*P* incorrect answers)
(3) Peak p.r.	P* (correct answers) – (0.25*P* incorrect answers)	**2P*** (number of correct answers) – (**0.5***P* incorrect answers)	P* (number of correct answers) – (0.25*P* incorrect answers))
(4)Valley p.r.	P* (correct answers) – (0.25*P* incorrect answers)	½ **P*** (number of correct answers) – (**0.125***P* incorrect answers)	P* (number of correct answers) – (0.25*P* incorrect answers)

Table 15: Experimental structure

Table 15 reports the characteristics of each experimental treatment for each of the two tasks.
It is important to notice that in each period of each treatment the penalty for an incorrect answer was valued only ¼ of the value of each correct answer. Given this difference we decided to carry out separately the analysis of correct and incorrect answers.[43]

In the *Fixed payment* subjects received 3 Euros in each period of the more challenging task and 2 Euros in each period of the less challenging task independently from their performance. This difference was because a period in the more challenging task lasted 8 minutes while in the less challenging task it lasted 5 minutes.
In the other treatments, payoffs depended on performance.

[42] The reason for these two additional measures was to raise subjects' awareness that there was a change in the piece rate they received. We decided to adopt these measures after running the pilot session given that several subjects at the end of the pilot could not remember the piece rate they received in each period.
[43] Analysis with incorrect answers is reported in section 3.6.2.

The piece rate for a correct answer is represented in table 1 by P which is always 100 points. What changes is the factor that multiplies P which characterized the dynamic of the incentives studied.

The exchange rate was 100 points = 4 cents for the more challenging task and 100 points = 10 cents for the less challenging task. In this way we keep the same structure of points for both tasks[44]. Average payment per subject at the end of the experiment was 8 Euros[45] for the more challenging task and 5 Euros for the less challenging task.

3.5 Hypotheses

Based on the literature review we developed in section 2, we present the following set of hypotheses.

In accord with standard economic theory and based on the literature we reported in section 3.2.1 we form the first hypothesis.

Hypothesis I: If subjects are paid based on performance, they will exert more effort compared to the case in which their payment is fixed and independent of performance.

We should therefore find that for period 1 subjects in the fixed payment treatment obtain less correct answers than subjects in other treatments[46]:

$$Fixed\ payment_1 < Average(Constant\ p.r._1 + Peak\ p.r._1 + Valley\ p.r._1)$$

Where $p.r._1$ means piece rate in period 1.

Based on the literature on the effect of relative changes in piece rate incentive magnitude described in section 3.2.3 we form the second hypothesis.

[44] Piece rate was set in a way that subjects' average payment with respect to the time spent in the laboratory would not differ substantially between the two tasks.
[45] In each treatment subjects received as well an additional payment of 2€ for showing up.
[46] Given that the payment per correct answer in the first period is the same for each of the piece rate treatment (i.e. 100 points), we can pool together subjects' number of correct answers.

Hypothesis II: Given the impossibility of intertemporal substitution in our experiment, a change in the amount of piece rate paid will lead to a change in the level of effort provided in the same direction of the change.

Therefore we should find that the difference between number of correct answers in period 2 and number of correct answers in period 1 are as follow[47]:

$$Peak\ p.r._{2-1} > Constant\ p.r._{2-1}$$
$$Valley\ p.r._{2-1} < Constant\ p.r._{2-1}$$

Based on the literature on asymmetric reference point updating described in section 3.2.4 and assuming that changes in piece rate incentive do shift reference points we form the following hypothesis for the number of correct answer in period 3:

Hypothesis III: A sequence of two consecutive changes in reference points will have a negative and stronger effect on performance when the first change is positive and the second change is negative, than the other way around.

This implies that we should find an asymmetric response to a second change in the magnitude of incentives.
Specifically we should obtain that the negative change in period 3 for the peak piece rate treatment will have a stronger effect on subjects' performance than the change in period 3 for the valley piece rate treatment.

$$Effect(Peak\ p.r._{3-2} - Constant\ p.r._{3-2})$$
$$>$$
$$Effect(Valley\ p.r._{3-2} - Constant\ p.r._{3-2})$$

[47] We use the differences in correct answers from one period to the other since we are interested in measuring performance's changes. The sub-index denotes for which periods we are calculating the difference of correct answers. For example, if the sub-index is 2-1 it means correct answers in period 2 minus correct answers in period 1.

Based on the literature on the interaction of the type of task with performance based monetary incentives reported in section 3.2.2 we form the fourth hypothesis.

Hypothesis IV: The effect of performance based monetary incentive on subjects' effort level will be stronger for the less challenging task.

This because in the less challenging task the level of intrinsic motivation should be lower compared to the more challenging task and consequently the overall response to extrinsic monetary incentive should be stronger.

$$\begin{gathered}\textit{Piece rate incentive } Effect\ (Less\ challenging\ task) \\ > \\ \textit{Piece rate incentive } Effect\ (More\ challenging\ task)\end{gathered}$$

3.6 Experimental results

In table 16 and table 17 we report the mean and standard deviation for correct and incorrect answers for each of the three periods of the four treatments for the more challenging task and for the less challenging task, respectively. In figure 8 and 9 we report graphically the evolution of the number of correct answers for each task.

More Challenging Task

Treatment		Mean Correct	S.d. Correct	Mean Incorrect	S.d. Incorrect	N
Period 1	Fix payment	42	17	3	4	19
	Constant piece rate	47	18	2	2	20
	Peak piece rate	38	26	4	3	19
	Valley piece rate	48	14	2	3	20
Period 2	Fix payment	37	15	5	6	19
	Constant piece rate	44	16	3	3	20
	Peak piece rate	38	19	7	5	19
	Valley piece rate	43	13	3	5	20
Period 3	Fix payment	30	13	4	6	19
	Constant piece rate	38	17	3	4	20
	Peak piece rate	29	17	4	2	19
	Valley piece rate	35	15	3	3	20

Table 16: Mean and standard deviation of correct and incorrect answers - more challenging task

Less Challenging Task						
Treatment		**Mean Correct**	**S.d. Correct**	**Mean Incorrect**	**S.d. Incorrect**	N
Period 1	Fix payment	3	3	3	1	20
	Constant piece rate	11	2	2	1	30
	Peak piece rate	9	3	2	2	30
	Valley piece rate	9	2	2	1	30
Period 2	Fix payment	3	2	4	2	20
	Constant piece rate	8	2	3	1	30
	Peak piece rate	12	3	3	1	30
	Valley piece rate	6	3	2	1	30
Period 3	Fix payment	5	2	4	2	20
	Constant piece rate	8	2	3	1	30
	Peak piece rate	5	3	4	1	30
	Valley piece rate	8	2	3	1	30

Table 17: Mean and standard deviation of correct and incorrect answers - less challenging task

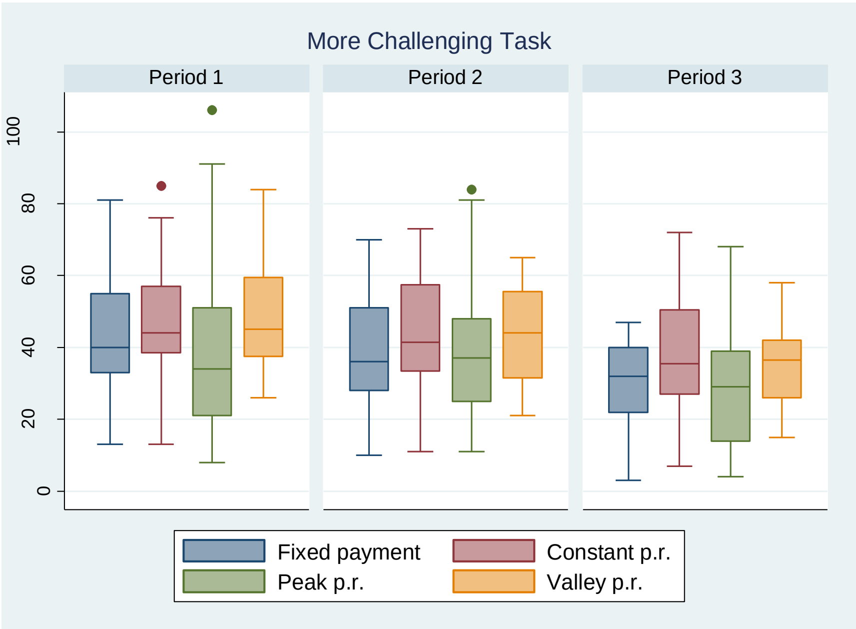

Figure 8: Number of correct answers for each period in each treatment - more challenging task

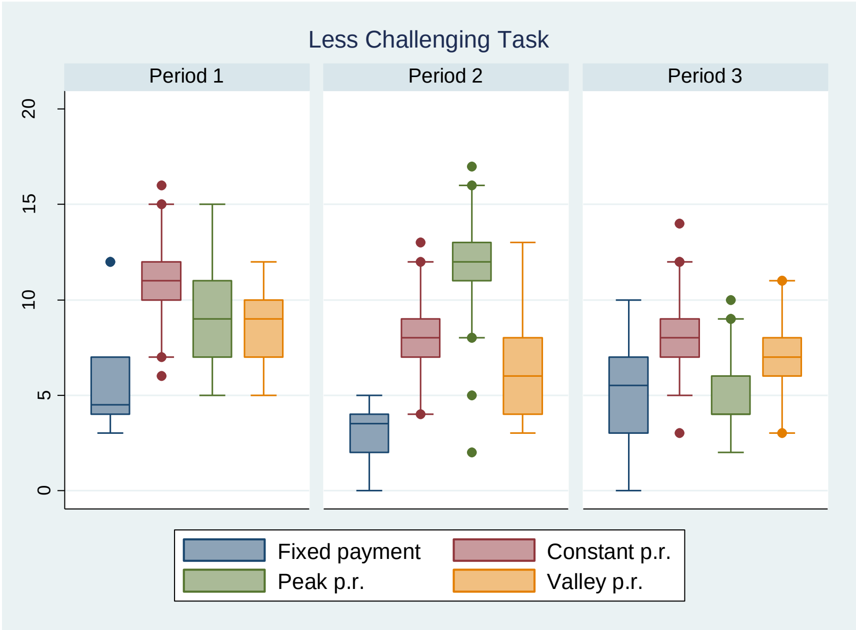

Figure 9: Number of correct answers for each period in each treatment - less challenging task

Figure 8 and figure 9 report a visual representation (i.e. a box plot) of the distribution of the number of correct answers for each of the treatments for the more challenging task and for the less challenging task respectively.

The box plot splits the distribution of correct answers into quartiles. The horizontal line contained in each box represents the median of the distribution of correct answers for each treatment for the corresponding period. The two horizontal lines outside each box represent lowest and highest observations. Points represents outliers.

From the distribution of correct answers for the more challenging task (figure 8) we observe a general tendency of the median of correct answers to decrease from period 1 to period 3 in each treatment. This could be caused both because of subjects' fatigue and because the number to be subtracted differed from period to period, perhaps increasing task difficulty in late periods.

In the same figure we also observe that within each period, medians of correct answers between treatments do not differ remarkably.

It is particularly interesting to notice that for the fixed payment treatment, median of correct answers is very similar to the median of the other three performance contingent treatments.[48]

From the distribution of correct answers for the less challenging task reported in figure 9 we observe a more marked difference between treatments.

In particular the fixed payment treatment is characterized by a much lower number of correct answers compared to each of the other treatments.

Particularly low is as well the number of correct answers in period 3 for the peak p.r. treatment and the visually substantial decrease that the peak p.r. treatment experienced from period 2 to period 3.

3.6.1 Hypothesis testing

The statistical inference that we present in next section will use only the number of correct answers as measure for performance[49].

The analysis of the evolution of the incorrect answers is reported separately in section 3.6.2.

Testing Hypothesis I

The first hypothesis suggests that for period 1 the level of effort exerted by subjects will be higher when subjects are remunerated in a performance contingent way. That is:

$$Fixed\ payment_1 < Average(Constant\ p.r._1 + Peak\ p.r._1 + Valley\ p.r._1)$$

[48] This could perhaps be explained by the fact that subjects feel particularly intrinsically motivated with this task or can also be explain by the incompleteness of contract between subjects and experimenter as speculated by Gneezy and Rustichini (2000).

[49] In order to test our hypotheses we adopt a statistical approach that is widely used in psychology and begin to be commonly used as well in other social sciences. The approach is part of what is named "*The new statistics"* (Cumming, 2013) and it suggests that instead of using the null – hypothesis significant testing approach, we should estimate our results in terms of effect sizes and confidence intervals.

The reason we favor this approach is that it allows us not only to establish whether there are differences among treated groups but instead it tells us the size of the differences and as consequence it allows us to easily compare results across different design or, as in our case across different tasks.

Furthermore, in addition to reporting effect sizes (estimated in our case by using *Cohen's d*) we also report the 95% confidence interval of each effect size. In this way we are able to capture better the extent of uncertainty in our inference.

More Challenging Task: Correct answers in period 1

Treatment	Statistic			
	Mean	S.d.	Average p.r.	Fixed Payment
Average p.r.	44.86	17.7	------	0.13 [-0.39; 0.64]
Fixed Payment	42.36	20.24	------	------

Table 18: Testing hypothesis I - more challenging task

Less Challenging Task: Correct answers in period 1

Treatment	Statistic			
	Mean	S.d.	Average p.r.	Fixed Payment
Average p.r.	9.72	4.26	------	1.68 [1.14; 2.21]
Fixed Payment	5.5	9.2	------	------

Table 19: Testing hypothesis I - less challenging task

Table 18 and 19 report the difference in means for the number of correct answers in period 1 in the fixed payment treatment compared with the average of the other three groups for the more challenging task and for the less challenging task respectively. Numbers reported in the right part of the table represent Cohen's d measuring the effect size of groups' differences and its 95% confidence interval.

Result I: The number of correct answers for the less challenging task is higher if subjects are paid for performance compared to when they received a fixed payment.

Support for result I:

From table 19 we observe a large effect size for the average number of correct answers in period 1 for the three performance contingent treatments pooled. In particular Cohen's d is 1.68, 95% CI [1.14; 2.21] compared to the fixed payment treatment.

The same analysis for the more challenging task (table 19) shows only a very small effect size of 0.13 with a 95% CI of [-0.39; 0.64].

Testing Hypothesis II

The second hypothesis suggests that a change in the amount of piece rate paid to subjects will lead to a change in the level of effort provided in the same direction of the change.

$$Peak\ p.r._{2-1} > Constant\ p.r._{2-1}$$

$$Valley\ p.r._{2-1} < Constant\ p.r._{2-1}$$

More Challenging Task: Correct answers in period 2 - Correct answers in period 1

Treatment	Statistic		Pairwise comparisons			
	Mean	S.d.	Fixed payment	Constant p.r.	Peak p.r.	Valley p.r.
Fixed payment	-4.9	7.9	------	-0.21	-0.55	0.01
Constant p.r.	-2.9	10.26	------	------	-0.29	0.20
Peak p.r.	0.11	9.97	------	------	------	0.50
Valley p.r.	-5.00	10.29	------	------	------	------

Table 20: Testing hypothesis II - more challenging task

Less Challenging Task: Correct answers in period 2 - Correct answers in period 1

Treatment	Statistic		Pairwise comparisons			
	Mean	S.d.	Fixed payment	Constant p.r.	Peak p.r.	Valley p.r.
Fixed payment	-2.40	3.50	------	-0.23	-1.21	0.02
Constant p.r.	-3.10	2.50	------	------	-1.59	0.23
Peak p.r.	2.13	3.90	------	------	------	1.34
Valley p.r.	-2.46	2.89	------	------	------	------

Table 21: Testing hypothesis II - less challenging task

.

Tables 20 and 21 report the difference in means of correct answers given in period 2 minus the correct answers given in period 1 for the more challenging task and for the less challenging task respectively. Numbers reported in the right part of the table represent Cohen's d and measure the effect size of groups' differences.

Result II: A positive change in the magnitude of piece rate per correct answer leads to a positive change in subjects' performance. Furthermore a negative change in the magnitude of piece rate per correct answer does not affect subjects' performance.

Support for result II:

We compare mean differences between the number of correct answers in period 2 minus the correct answers in period 1for both tasks.[50]

[50] Numbers in the right part of table 20 – 25 reports *Cohen's d*, a measure for the effect size of the difference in means between treatments. In each cell (of each tables in this paper), a negative (positive) number means that the treatment reported in the column is characterized by a highest (lowest) average compared to the treatment reported in the row of the table.

The effect size is particularly marked once comparing the *constant p.r.* treatment (mean = -3.10) with the *peak p.r.* treatment (mean = 2.13) for the less challenging task.
The estimated differences between the means is -1.59, 95% CI [-2.17; -1.00].
The same comparison for the more challenging task also suggest a similar result although the effect size is much smaller and equal to -0.55, 95% CI [-1.20; 0.09].
In both tasks, when comparing valley p.r. and constant p.r. we find very small effect sizes. This result suggests that a decrease in piece rate per correct answer does not have a substantial effect on subjects' performance.

Testing Hypothesis III

The third hypothesis suggests that the order of the sequence in which the piece rate incentive changes matters for subjects' performance. In particular it suggests that the negative effect will be much stronger when the first change is positive and the second change is negative.

$$Effect(Peak\ p.r._{3-2} - Constant\ p.r._{3-2}) > Effect(Valley\ p.r._{3-2} - Constant\ p.r._{3-2})$$

More Challenging Task: Correct answers in period 3 - Correct answers in period 2

Treatment	Statistic		Pairwise comparisons			
	Mean	S.d.	Fixed payment	Constant p.r.	Peak p.r.	Valley p.r.
Fixed payment	-7.42	7.09	------	-0.18	0.38	0.15
Constant p.r.	-6.15	6.88	------	------	0.57	0.33
Peak p.r.	-10.05	6.79	------	------	------	-0.22
Valley p.r.	-8.50	7.15	------	------	------	------

Table 22: Testing hypothesis III - more challenging task

Less Challenging Task : Correct answers in period 3 - Correct answers in period 2

Treatment	Statistic		Pairwise comparisons			
	Mean	S.d.	Fixed payment	Constant p.r.	Peak p.r.	Valley p.r.
Fixed payment	1.80	2.37	------	0.50	2.49	0.24
Constant p.r.	0.33	3.18	------	------	1.95	-0.26
Peak p.r.	-6.66	3.93	------	------	------	-2.24
Valley p.r.	1.13	2.95	------	------	------	------

Table 23: Testing hypothesis III - less challenging task

Table 22 and 23 report the difference in means of correct answers given in period 3 minus correct answers given in period 2 for the more challenging task and for the less challenging task respectively. Numbers reported in the right part of the table represent Cohen's d and measure the effect size of groups' differences.

Result 3: The order of the sequence in which a decrease in piece rate follows a previous increase made subjects to reduce substantially the number of correct answers compared to the case in which piece rate remains constant.

Oppositely, the effect of an increase in piece rate that followed a decrease does not have a strong effect compare to the case in which piece rate remains constant.

Support for result 3:

We compare mean differences between the number of correct answers in period 2 minus the correct answers in period 1for both tasks.

Mean comparison for the more challenging task (table 8) comparing the *constant p.r.* treatment with the *peak p.r.* treatment shows an estimated difference of 0.57, 95% CI [-0.07; -1.21].

Mean comparison for the less challenging task (table 9) comparing the *constant p.r.* treatment with the *peak p.r.* treatment indicates a very large effect size of 1.95, 95% CI [1.33; 2.57].

Once we compare mean difference for valley p.r. and constant p.r we observed that for both tasks the effect size reported is very small. [51]

Testing Hypothesis IV

The fourth hypothesis suggests that the effect of performance contingent monetary incentive is stronger for the less challenging task.

$$Piece\ rate\ incentive\ \ Effect\ (Less\ challenging\ task)$$
$$>$$
$$Piece\ rate\ incentive\ Effect\ (More\ challenging\ task)$$

[51] At the same time it is interesting to notice the large Cohen's d once we compare Peak p.r. treatment with Valley p.r. treatment for the less challenging task, d = -2.24, 95% CI [-2.88; -1.58]. This perhaps further indicating a detrimental effect caused by a decrease in piece rate which follows a previous increase.

More Challenging Task: Correct answers in period 1			
Treatment	Constant p.r	Peak p.r.	Valley p.r.
Fixed Payment	-0.26	0.17	-0.39

Table 24: Testing hypothesis IV - more challenging task

Less Challenging Task: Correct answers in period 1			
Treatment	Constant p.r	Peak p.r.	Valley p.r.
Fixed Payment	-2.30	-1.46	-1.42

Table 25: Testing hypothesis IV - less challenging task

Tables 24 and 25 report difference in means of number of correct answers in period 1 for the fixed payment treatment compared with each of the other treatment for the more challenging task and for the less challenging task respectively. Numbers reported in the right part of the table represent Cohen's d measuring the effect size of groups' differences.

Result 4:

The positive effect of performance contingent monetary incentives on effort provision is larger for the less challenging task.

Support for result 4:

We compare mean differences on the number of correct answers in period 1 for both tasks.

Table 24 compares for the more challenging task the average number of correct answers for the *fixed payment* treatment against each piece rate treatment individually.

We observe that the estimated difference between the mean of the fixed payment treatment is small compared to the other three performance contingent treatments.

Differences of the fixed payment treatment with respect to each of the other treatments are:

-0.26, 95% CI [-0.88; 0.37] compared to the constant p.r. treatment;

0.17, 95% CI [-0.47; 0.80] compared to the peak p.r. treatment;

-0.39, 95% CI [-1.03; 0.24] compared to the valley p.r. treatment;

Table 25 reports the same analysis for the less challenging task.

The estimated difference between the mean of the fixed payment treatment is much larger and in the direction predicted compared to each of the other three performance contingent treatments.

Differences of the fixed payment treatment with respect to each of the other treatments are:

-2.30, 95% CI [-3.03; -1.57] compared to the constant p.r. treatment;

-1.46, 95% CI [-2.08; -0.81] compared to the peak p.r. treatment;

-1.42, 95% CI [-2.05; -0.79] compared to the valley p.r. treatment;

3.6.2 Analysis with number of incorrect answers

In figure 10 and 11 we report graphically the evolution of the number of incorrect answers for each task.

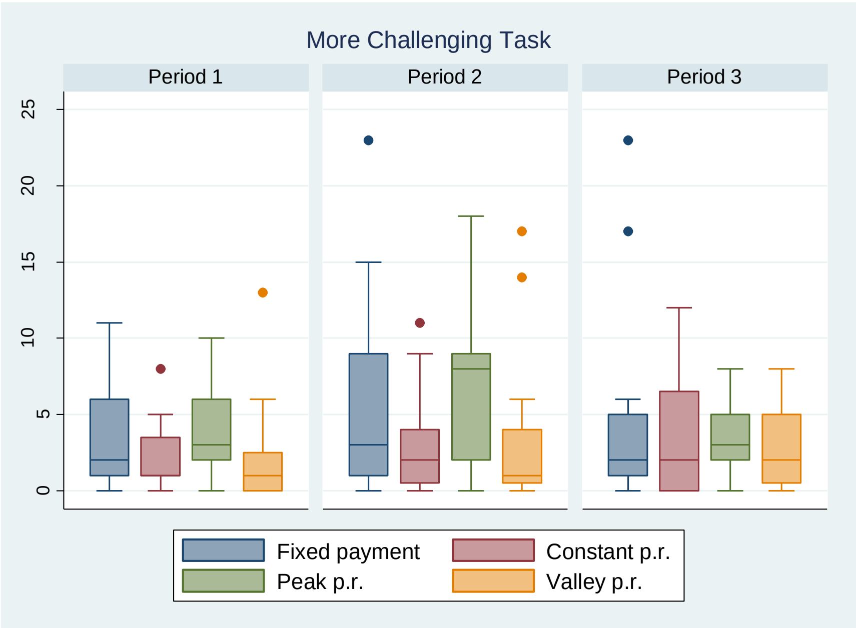

Figure 10: Number of incorrect answers for each period in each treatment - more challenging task

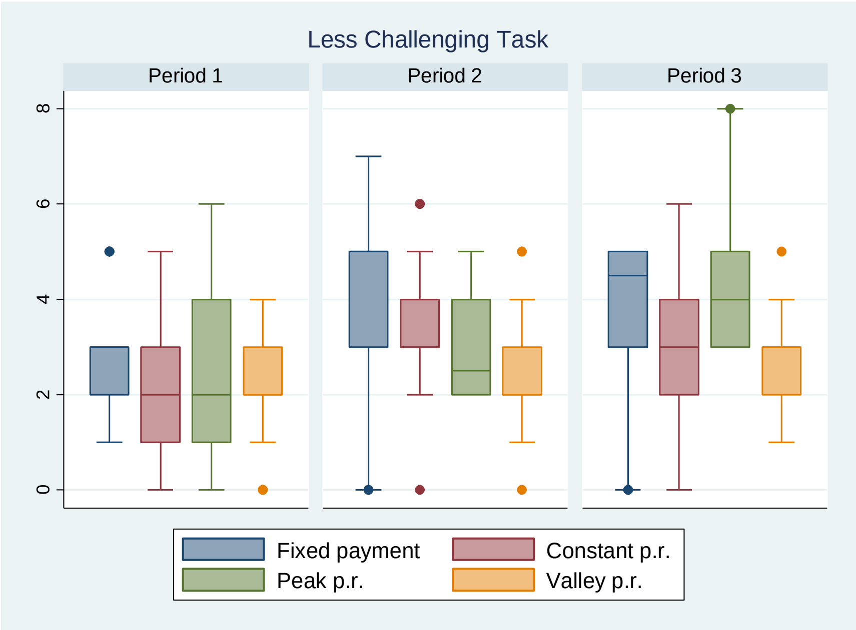

Figure 11: Number of incorrect answers for each period in each treatment - more challenging task

In figure 10 we observe the distribution of the number of incorrect answers for the more challenging task.[52]

Generally, the number of incorrect answers does not appear to be very different between treatments in period 1 and period 2. While in period 3, for the peak p.r. treatment we notice an important increase in the number of incorrect answers.

This might be the result of subjects attempting to introduce as much numbers as possible moved by the relatively higher piece rate.

In figure 11 we observe the distribution of the number of incorrect answers for the less challenging task.

A general pattern that we observe is that in late periods the number of incorrect answers increases in each treatment, probably as a result of the boredom of the task.

It is particularly interesting the way in which the number of incorrect answer increases for the peak treatment in the third period compared to the constant p.r. and to the valley p.r.

[52] In section 3.6 we explained how to interpret a box plot.

The estimated effect size for the comparison of incorrect answers given in period 3 minus incorrect answers given in period 2 for the peak p.r. against the constant p.r is quite large and equal to 1.01, 95% CI [0.46; 1.53].

The same comparison for the peak p.r. against the valley p.r. returned an estimate of 0.93, 95% CI [0.39; 1.46].

Both results confirmed that the increase the number of incorrect answers in the peak p.r. treatment compared with the constant p.r. and the valley p.r. (in the third period) was quite large.

3.6.3 Additional results

Post treatment task evaluation

At the end of the experiment each subject answered the Intrinsic Motivation Inventory scale (Ryan, 1982) reported in this paper in section 3.3.

For each subject we calculated the average evaluation of the task and we compare mean differences.

The estimated effect size for the difference between the more challenging task (M= 3.18; SD=1.29) and the less challenging task (M=1.68, SD= 0.87) is 1.42, 95% CI [1.23; 1.60].

This result highlights an important difference in evaluation between the two tasks. In particular it says that participants were on average relatively more intrinsically motivated by the more challenging task.

Gender differences

We investigate whether there are differences in ability as well as differences in response to monetary incentives between males and females.

To test for differences in ability we compare total number of correct answers given during the entire experiment between males and females across all treatments.

For the more challenging task the average of total correct was of 124 for males and 112 for females.

We find a negligible effect size for gender difference of 0.23, 95% CI [-0.21; 0.67].

For the less challenging task the average of total correct answers was of 22 for males and 24 for females with very small effect size on their differences.

In order to investigate if there were gender differences in response to monetary incentives we replicate the same analysis we carried out when testing hypothesis I while controlling for gender.
For both tasks we do not observe any gender difference.

Cost-Benefit tradeoff

One of the main objectives of this research was to inform an employer about the most efficient way for his/her firm to remunerate his/her employees with performance based types of incentive.
In table 26 and table 27 we report a cost-benefit analysis for both tasks.
We report the average of the total number of correct and incorrect answers divided by treatment. In the last column we computed the average cost incurred by the experimenter in each treatment.

More Challenging Task: Total cost for the experimenter by treatment

	Correct answers	Incorrect answers	Show up fee	Payment
Fixed payment	109	12	2.0 €	11.0 €
Constant p.r.	129	8	2.0 €	7.1 €
Peak p.r.	105	15	2.0 €	7.5 €
Valley p.r.	126	8	2.0 €	6.1 €

Table 26: Cost-Benefit tradeoff - more challenging task

Less Challenging Task: Total cost for the experimenter by treatment

	Correct answers	Incorrect answers	Show up fee	Payment
Fixed payment	14	11	2.0 €	8.0 €
Constant p.r.	27	9	2.0 €	4.5 €
Peak p.r.	26	10	2.0 €	5.5 €
Valley p.r.	23	7	2.0 €	3.1 €

Table 27: Cost-Benefit tradeoff - less challenging task

The payment column was computed by multiplying the period's correspondent piece rate (and penalty) for the number of correct (incorrect) answers while adding the show-up fee.

From table 26 we observe that among the piece rate treatments, the peak p.r. results to be the less efficient one given that it is characterized by the lowest number of correct answers and by the highest cost for the experimenter.

In the analysis of the less challenging task reported in table 27, we obtain a similar result. In particular we observe that the constant p.r. treatment provides almost the same average of correct answers than the constant p.r. but it is 20% more expensive.

3.7 Discussion

3.7.1 Summary of results

In this paper we examined the impact of different payment schemes on performance on two effortful tasks by means of a laboratory experiment.
Tasks differed since the first was designed to be relatively more challenging compared to the second. Furthermore, the first task scored significantly higher on an interest/enjoyment scale aimed at measuring subjects' intrinsic motivation toward the task.
In each task we implemented four different payment schemes separately in a three period environment. The first schema rewarded subjects with a fixed and known amount in each period that did not depend on subjects' performance.
The other three schemes were performance contingent and only differed since, in the second period, the piece rate per correct answers was doubled, halved or remained unchanged compared to period one and period three.
This simple structure allowed us to investigate subjects' response to a variety of changes in incentives.

We obtained the following results:
First we observed that in line with standard economic theory, performance contingent monetary incentives affect subjects' effort provision.
At the same time we observed that:

(1) The effect of monetary incentive is task dependent. Consistent with the literature on intrinsic motivation, we infer that monetary incentives are more effective for a less challenging task in which the level of intrinsic motivation is lower.
(2) For a unique change in incentive magnitude we found that subjects' response is stronger when the change is positive compared to when the change is negative.
(3) For two consecutive changes in incentive magnitude we observed an asymmetric response of the level of effort provided by subjects. In particular we realized that, consistent with the literature on reference point shifting, the effect of a decrease in piece rate which follows a previous increase is negative and much stronger than the effect of an increase following a decrease.

3.7.2 Managerial implications

Taking into account the usual concern of external validity that characterizes laboratory experiments, the study we proposed has mainly three potentially important implications for an employer looking for the best way to motivate his employees.
First, it suggests that for boring and repetitive tasks an employer should enhance employees' effort by targeting their salary on performance instead of paying them on a fixed based wage. This might allow the employer to motivate his employees more effectively.

Second, it suggests that an increase in an employee's performance based salary might positively affect employees' willingness to exert effort. Perhaps this result supports the use of positive rewards and bonuses in firms to incentivize workers.

Third, and in our view the most important suggestion, if an increase in employees' remuneration is followed by a decrease, the result can be detrimental for employees' motivation. An employer should therefore avoid substantially increasing an employee's salary if the increase cannot be sustained over time.
This result might also suggest that the effect of bonus and prizes might not have long lasting benefits.

3.7.3 Limits and future research

In this paper we started to investigate how different incentives' dynamics can affect subjects' willingness to exert effort.

Clearly in our experiment changes in piece rate were completely exogenous and furthermore, subjects did not receive any explanation regarding the reason for these changes.

At the same time we acknowledge that often in real life these changes come with an explanation. Think for example of a case where an employer has to cut employees' wages because he is facing financial problems.

An interesting question would therefore be to investigate if, in presence of a plausible justification, subjects will react to a change in incentive differently. More specifically, if social preferences such as altruism, reciprocity or guilt aversion could affect subjects' effort provision.

Another interesting question could be to investigate how long a damage caused by an increase followed by a decrease in the amount of piece rate lasts. This could be studied by constructing longer sequences of changes in piece rate amount in which more complete incentives' dynamics could be implemented.

3.8 Appendix

Description of the tasks.

Here we report a screenshot for the less challenging task. In this task, subjects have to count the number of A's contained in each paragraph. Once they insert the number, they have to press the button "Enter" followed by the button "Next Paragraph".

The program will sum the number of correct and incorrect answers without communicate it to the subjects.

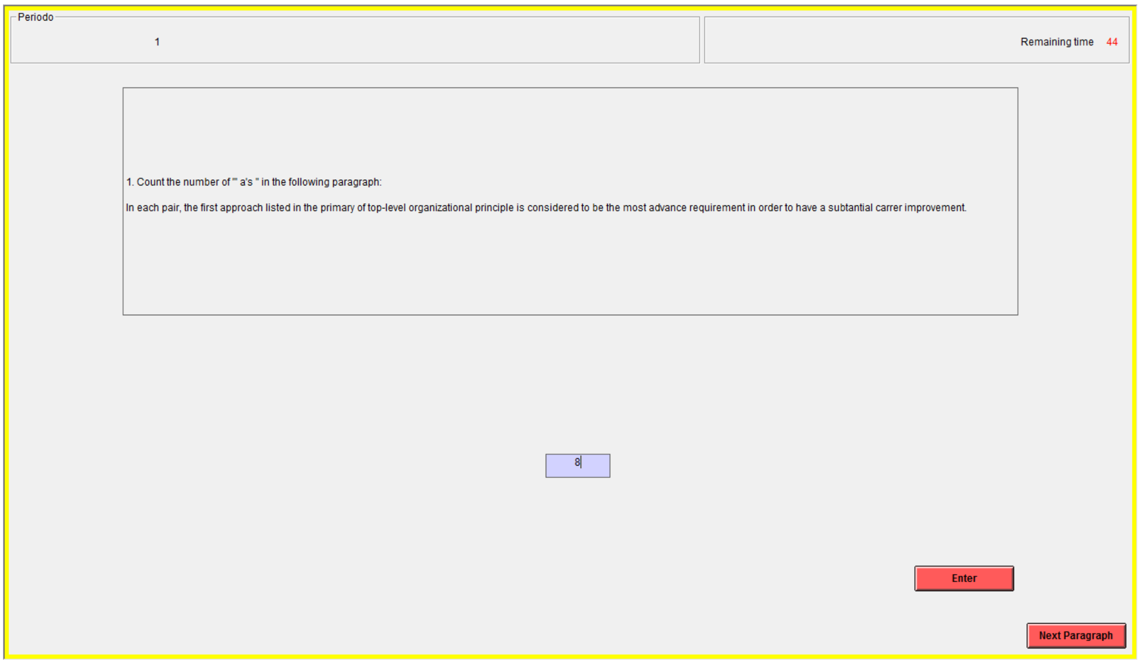

Here we report a screenshot for the more challenging task (in Spanish). Subjects have to downward counting from the number displayed under the field "Número" each time subtracting a fixed number announced at the beginning of the period. Once inserted a number in the field "Tu respuesta" they have to press the "OK" button. The program will sum the number of correct and incorrect answers without communicate it to the subjects.

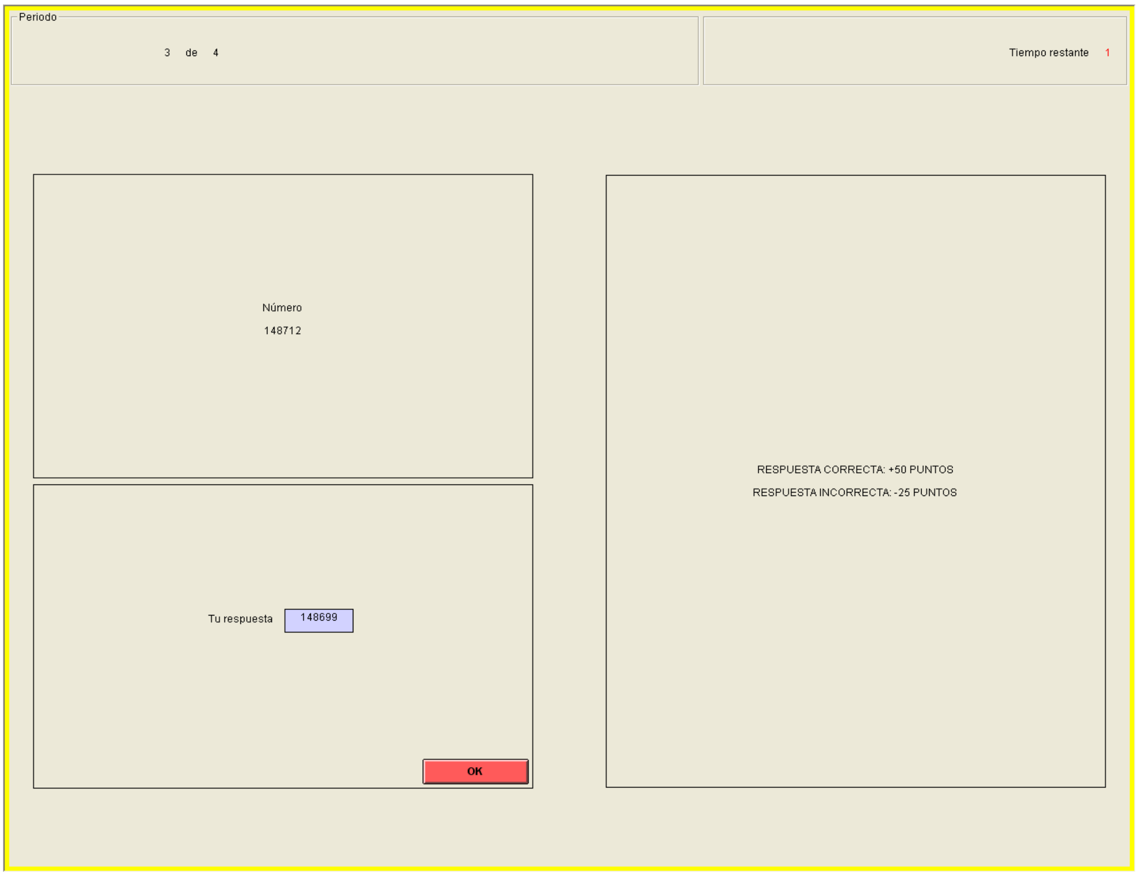
Periodo
3 de 4
Tiempo restante 1
Número
148712
Tu respuesta 148699
OK
RESPUESTA CORRECTA: +50 PUNTOS
RESPUESTA INCORRECTA: -25 PUNTOS

BIBILIOGRAPHY

Abeler, J., Falk, A., Götte, L., and Huffman, D. (2011), "Reference points and effort provision." *American Economic Review*, 101, 470–492.

Ackert, L., Charupat, N., Church, B. and Deaves, R. (2006) Margin, short selling, and lotteries in experimental asset markets, *Southern Economic Journal* 73(2), 419-436.

Allen, Franklin, and Douglas Gale. (2000). "Bubbles and Crises." *The Economic Journal* 110 (460): 236–55.

Alsemgeest Paul, Noussair Charles and Mark Olson. (1998). "Experimental Comparisons of Auctions Under Single-and Multi-Unit Demand." *Economic Inquiry* 36 (1): 87–97.

Arkes, H., Hirshleifer, D., Jiang, D., Lim, S. (2008), "Reference point adaptation: Tests in the domain of security trading." *Organ. Behav. Human Decision Process* 105(1) 67–81

Arkes, H., Hirshleifer, D., Jiang, D., Lim, S. (2010), "A cross-cultural study of reference point adaptation: Evidence from China, Korea and the US." *Organ. Behav. Human Decision Process* 112(2) 99–111

Ball, S.B. and Holt, C.A. (1998) Classroom games: speculation and bubbles in an asset market, *Journal of Economic Perspectives* 12(1), 207-218.

Baker, G., Jensen, M., and Murphy, K. J. (1988), "Compensation and Incentives: Practice vs. Theory". *Journal of Finance*, 43, pp. 593-616.

Banker, Rajiv; Seok-Young Lee, and Gordon Potter. (1996). "A Field Study of the Impact of a Performance-Based Incentive Plan". *J. Acc. Econ.*, 21, pp. 195-226.

Baucells, Manel, Martin Weber, & Frank Welfens (2011). Reference Point Formation and Updating. *Management Science* 57(3), 506–519

Bonner, S.E; and G. B. Sprinkle. (2002). The effects of monetary incentives on effort and task performance: Theories, evidence, and a framework for research. *Accounting, Organizations and Society* 27 (4–5): 303–345.

Bostian, A.J. and Holt, C.A. (2009) Price bubbles with discounting: a web-based classroom experiment, *Journal of Economic Education* 40, 27-37.

Breaban, A. and Noussair, C. (2014) Fundamental value trajectories and trader characteristics in an asset market experiment, CentER Discussion Paper Series No. 2014-010.

Bruggen, Alexander, Strobel, Martin, (2007). Real effort versus chosen effort in experiments. *Economics Letters* 96 (2), 232–236.

Brunnermeier, M. (2009) Bubbles, in S. Durlauf and L. Blume (eds), The New Palgrave Dictionary of Economics (2nd edition), Palgrave Macmillan.

Brunnermeier, Markus K., and Stefan Nagel. (2004). “Hedge Funds and the Technology Bubble.” *The Journal of Finance* 59 (5).

Camerer, Colin and Hogarth Robin. (1999). "The Effects of Financial Incentives in Experiments: A Review and Capital-Labor Framework," *J. Risk Uncertainty* 19, 42

Camerer, C., Babcock, L., Loewenstein, G., Thaler, R., (1997). Labor supply of New York city cab drivers: One day at a time. *Quarterly Journal of Economics* 111, 408–441.

Camerer, C.F. and Weigelt, K. (1993) Convergence in experimental double auctions for stochastically lived assets, in: The Double Auction Market: Institutions, Theories, and Evidence, Proceedings Volume in the Santa Fe Institute Studies in the Sciences of Complexity (14). Addison-Wesley, Reading, MA, pp. 355-396.

Caginalp, G., Porter, D. and Smith, V. (1998) Initial cash/asset ratio and asset prices: an experimental study. Proceedings of the National Academy of Sciences of the United States of America 95, 756--761.

Caginalp, G., Porter, D. and Smith, V. (2001) Financial bubbles: excess cash, momentum, and incomplete information. *Journal of Psychology and Financial Markets* 2(2), 80-99.

Carswell, John. 1960a. *The South Sea Bubble*. Cresset Press.
———. 1960b. *The South Sea Bubble*. Stanford University Press.
———. 1993. *The South Sea Bubble*. Alan Sutton.

Cerasoli, C. P., Nicklin, J. M., & Ford, M. T. (2014). Intrinsic motivation and extrinsic incentives jointly predict performance: A 40-year meta-analysis. *Psychological Bulletin*

Chen, H. A., A. R. Rao. 2002. “Close encounters of two kinds: False alarms and dashed hopes”. *Marketing Science*. 21(2) 178–196

Cheung, S. and Palan, S. (2012) Two heads are less bubbly than one: team decision-making in an experimental asset market, *Experimental Economics* 15(3), 373--397.

Corgnet, B., Hernán, R., Kujal, P., & Porter, D. (2012). The effect of earned vs. house money on price bubble formation in experimental asset markets. Working paper.

Crockett, S. and Duffy, J. (2013) An experimental test of the Lucas asset pricing model, working paper, University of Pittsburgh.

Cumming, G. (2013). The New Statistics: Why and How. *Psychological Science* 25,7-29

Dale, Richard. (2004a). *The First Crash: Lessons from the South Sea Bubble*. Princeton University Press.

———. (2004b). *The First Crash: Lessons from the South Sea Bubble*. Princeton University Press.

Dale, Richard S, Johnnie EV Johnson, and Leilei Tang. (2005). “Financial Markets Can Go Mad: Evidence of Irrational Behaviour during the South Sea Bubble1.” *The Economic History Review* 58 (2): 233–71.

Deci, Edward L. (1971). “Effects of Externally Mediated Rewards on Intrinsic Motivation*.” Journal of Personality and Social Psychology*, 18(1): 105–115

Deci, E., Koestner, R. and Ryan, R. (1999), "A Meta-Analytic Review of Experiments Examining the Effects of ExtrinsicR ewardso n IntrinsicM otivation" *Psychological Bulletin*, 125 (6), 627-668

Dickinson, D.L. (1999). "An Experimental Examination of Labor Supply and Work Intensities." *Journal of Labor Economics*. 17(4) Part 1,638-670.

Dufwenberg, M., Lindqvist, T. and Moore, E. (2005) Bubbles and experience: an experiment, *American Economic Review* 95, 1731-1737.

Fama, Eugene. (1965). “The Behavior of Stock-Market Prices.” *Journal of Business*: 34–105.

Fenig, G., Mileva, M. and Petersen, L. (2014) Asset trading and monetary policy in production economies, working paper, Simon Fraser University.

Fehr, E., Goette, L., & Lienhard, M. (2008). Loss Aversion and Effort: Evidence from a Field Experiment. Working Paper. (University of Zurich)

Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics*, 10 (2), 171-178

Fischbacher, U., Hens, T. and 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.

Fisher, E. and Kelly, F. (2000) Experimental foreign exchange markets, *Pacific Economic Review* 5, 365-387.

Frey, B.S. (1997). On the relationship between intrinsic and extrinsic work motivation. *International Journal of Industrial Organi*zation, 15, 427-439

Fréchette G. (2011), "Laboratory Experiments: Professionals versus Students," in *The Methods of Modern Experimental Economics*, G. Fréchette and A. Schotter (eds.), Oxford University Press.

Garber, Peter M. (2001). *Famous First Bubbles: The Fundamentals of Early Manias*. Reprint. The MIT Press.

Gibbons, R. and Murphy, K. (1992). "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence", *Journal of Political Economy*, 100, 468-505

Gneezy, U. and Rustichini, A. (2000), "Pay Enough or Don't Pay at All", *Quarterly Journal of Economics*, 115(3), 791-810

Greenwood, Robin, and Stefan Nagel. (2009). "Inexperienced Investors and Bubbles." *Journal of Financial Economics* 93 (2): 239–58.

Harris, Ron. (1994). "The Bubble Act: Its Passage and Its Effects on Business Organization." *The Journal of Economic History* 54 (3): 610–27.

Haruvy, Ernan, and Charles N. Noussair. (2006). "The Effect of Short Selling on Bubbles and Crashes in Experimental Spot Asset Markets." *Journal of Finance*.

Haruvy, Ernan, Charles N. Noussair, and Owen Powell. (2013). "The Impact of Asset Repurchases and Issues in an Experimental Market." *Review of Finance*.

Hens, T. and Steude, S.C.(2009). The leverage effect without leverage, *Finance Research Letters* 6, 83-94.

Holmstrom, Bengt. "Moral Hazard and Observability." *Bell J. Econ*. 10 (Spring 1979): 74-91.

Hong, Harrison, José Scheinkman, and Wei Xiong. (2006). “Asset Float and Speculative Bubbles.” *The Journal of Finance* 61 (3): 1073–1117.

Hossain, T., List, J.A. (2009). “The Behavioralist Visits the Factory: Increasing Productivity Using Simple Framing Manipulations”, *Management Science*, vol.58 (12), 2151 -2167

Huber, J., Kirchler, M. and Stöckl, T. (2012) Underreaction and anchoring in experimental asset markets with distinct fundamental value regimes, working paper, University of Innsbruck.

Hussam, R., Porter, D. and Smith, V. (2008) Thar she blows: can bubbles be rekindled with experienced subjects? *American Economic Review* 98, 924-937.

Hutcheson, A. (1720). *Collection of Calculations and Remarks Relating to the South Sea Scheme*. London.

Kahneman, D., A. Tversky. (1979). “Prospect Theory: An analysis of decision under risk”. *Econometrica* 47(2) 263–292

Kale, J., Reis, E., Venkateswaran, A., (2009). Rank-order tournaments and incentive alignment: the effect on firm performance. *Journal of Finance* 64(3), 1479-1512

King, R., Smith, V., Williams, A. and Van Boening, M. (1993) The robustness of bubbles and crashes in experimental stock markets, in: I. Prigogine, R. Day, and P. Chen (eds). *Nonlinear Dynamics and Evolutionary Economics*, Oxford University Press, Oxford.

Kirchler, M., Huber, J. and Stöckl, T. (2012) Thar she bursts -- reducing confusion reduces bubbles, *American Economic Review* 102, 865-83.

Kose, T. (2013) Price convergence and fundamentals in asset markets with bankruptcy risk: an experiment, working paper, University of Pittsburgh.

Koszegi, B., & Rabin, M. (2006). A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics*, 121 (4), 1133-1165.

Lazear, Edward. (2000). "Performance Pay and Productivity." *American Economic Review*. December, 90:5, pp. 1346-361.

Lei, V., Noussair, C. and Plott, C. (2001) Nonspeculative bubbles in experimental asset markets: lack of common knowledge of rationality vs. actual irrationality, *Econometrica* 69, 831-859.

Lei, V. and Vesely, F. (2009) Market efficiency, evidence from a no-bubble asset market experiment, *Pacific Economic Review* 14, 246-258.

Libby, R., & Lipe, M. G. (1992). "Incentive effects and the cognitive processes involved in accounting judgments". *Journal of Accounting Research*, 30, 249–273.

Lintner, John. (1971). "The Effect of Short Selling and Margin Requirements in Perfect Capital Markets." *The Journal of Financial and Quantitative Analysis* 6 (5): 1173–95.

Long, J. Bradford de, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. (1990). "Positive Feedback Investment Strategies and Destabilizing Rational Speculation." *The Journal of Finance* 45 (2): 379–95.

Lugovskyy, V., Puzzelo, D. and Tucker, S. (2011) An experimental study of bubble formation in asset markets using the Tâtonnement trading institution, working paper, Indiana University.

Lugovskyy, V., Puzzelo, D., Tucker, S. and Williams, A. (2012) Can concentration control policies eliminate bubbles?, working paper, Indiana University.

Mohd-Sanusi, Z., Mohd-Iskandar, T. *(*2007*) "*Audit judgement performance: Assessing the effect of performance incentives, effort and task complexity". *Managerial Auditing Journal, Vol.* 22*, pp.* 34–52

Niederle, Muriel, and Lise Vesterlund. (2007). "Do Women Shy Away from Competition? Do Men Compete Too Much?" *Quarterly Journal of Economics*, 122(3): 1067-1101.

Noussair, C. and Powell, O. (2010) Peaks and valleys: price discovery in experimental asset markets with non-monotonic fundamentals, *Journal of Economic Studies* 37(2), 152-180.

Noussair, C., Robin, S. and Ruffieux, B. (2001) Price bubbles in laboratory asset markets with constant fundamental values, *Experimental Economics* 4, 87-105.

Noussair, C. and Tucker, S. (2006) Futures markets and bubble formation in experimental asset markets, *Pacific Economic Review* 11, 167-184.

Oechssler, J. (2010) Searching beyond the lamppost: let's focus on economically relevant questions, *Journal of Economic Behavior and Organization* 73, 65-67.

Paarsch, H.J., Shearer, B.S. (2009). "The Response to Incentives and Contractual Efficiency: Evidence from a Field Experiment". *European Economic Review* (53), 481-49

Palan, S. (2013) A review of bubbles and crashes in experimental asset markets, *Journal of Economic Surveys* 27, 570-588.

Pastor, L., and P. Veronesi. (2006). Was There A Nasdaq Bubble in the Late 1990s? *Journal of Financial Economics* 81:61–100

Prendergast, Canice. (1999). “The Provision of Incentives in Firms.” *Journal of Economic Literature*, 37(1): 7–63.

Pritchard, R. D., Jones, S. D., Roth, P. L., Stuebing, K. K., & Ekeberg, S. E. (1988). “Effects of group feedback, goal setting, and incentives on organizational productivity”. *Journal of Applied Psychology*, 73, 337–358

Pokorny, K. (2008). Pay – but do not pay too much: An experimental study on the impact of incentives. *Journal of Economic Behavior & Organization*. 66, 251-264
Porter, D. and Smith, V. (1995) Futures contracting and dividend uncertainty in experimental asset markets, *Journal of Business* 68, 509-541.

Rey-Biel Pedro, Roman Sheremeta and Neslihan Uler (2013).The effects of information and culture on giving. Working paper

Riedel, J. A., Nebeker, D. M., & Cooper, B. L. (1988). “The influence of financial incentives on goal choice, goal commitment, and task performance”. *Organizational Behavior and Human Decision Processes*, 42, 155–180.

Ryan, R.M. (1982). “Control and information in the intrapersonal sphere: An extension of cognitive evaluation theory. *Journal of Personality and Social Psychology*,43, 450-461

Ryan, R. M., Mims, V., & Koestner, R. (1983). “Relation of reward contingency and interpersonal context to intrinsic motivation: A review and test using cognitive evaluation theory”. *Journal of Personality and Social Psychology*, 45, 736–750

Scherbina, A. (2013) Asset price bubbles: a selective survey, working paper, IMF Working Paper WP/13/45.

Shea, Gary S. (2007). “Financial Market Analysis Can Go Mad (in the Search for Irrational Behaviour during the South Sea Bubble) -Super-1.” *Economic History Review* 60 (4): 742–65.

Shi, Lan (2007) "Productivity Effect of Piece Rate Contracts: Evidence from Two Small Field Experiments," Mimeo, Department of Economics, University of Washington, Seattle

Smith, Vernon L. (1962). “An Experimental Study of Competitive Market Behavior.” *Journal of Political Economy*.

Smith, V. (2010) Theory and experiment: what are the questions? *Journal of Economic Behavior and Organization* 73, 3-15.

Smith, V., Suchanek, G. and Williams, A. (1988) Bubbles, crashes and endogenous expectations in experimental spot asset markets, *Econometrica* 56, 1119-1151.

Smith, V., van Boening, M. and Wellford, C. (2000) Dividend timing and behavior in laboratory asset markets, *Economic Theory* 16, 567-583.

Stöckl, T., Huber, J. and Kirchler, M. (2010) Bubble measures in experimental asset markets, *Experimental Economics* 13, 284-298.

Takahashi, H., Shen, J., Ogawa, K. (2014). “An experimental Examination of Compensation Schemes and Level of Effort in Differentiated Tasks”. Working paper

Temin, Peter, and Hans-Joachim Voth. (2004). “Riding the South Sea Bubble.” *American Economic Review* 94 (5): 1654–68.

———. (2013). *Prometheus Shackled: Goldsmith Banks and England’s Financial Revolution after 1700*. Oxford University Press, USA

Van Boening, M., Williams, A. and LaMaster, M. (1993) Price bubbles and crashes in experimental call markets, Economics Letters 41, 179-185.

Xie, H. and Zhang, J. (2012) Bubbles and experience: an experiment with a steady inflow of new traders, CIRANO Scientific Series, CIRANO, Montreal.