COMPLEXITY, CONFUSION AND BUBBLES IN EXPERIMENTAL ASSET MARKETS [*]

Stephen L. Cheung [†]
Discipline of Economics, The University of Sydney
Merewether Building H04, Sydney NSW 2006, AUSTRALIA
S.Cheung@econ.usyd.edu.au

Morten Hedegaard
Department of Economics, University of Copenhagen
Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K, DENMARK
Morten.Hedegaard@econ.ku.dk

Stefan Palan
Institute of Banking and Finance, Karl-Franzens University Graz
Universitätsstraße 15/F2, 8010 Graz, AUSTRIA
stefan.palan@uni-graz.at

Preliminary draft: 26 February 2010.

Abstract:

Experimental asset markets of the type introduced by Smith, Suchanek and Williams (1988) are known to produce price bubbles and crashes with inexperienced participants. We examine the effects of reducing the complexity of the decision environment and controlling for confusion upon the performance of such markets. We first show that a 'sledgehammer' treatment combining dividend certainty, elimination of speculation and a battery of control questions essentially eliminates price bubbles with inexperienced participants. More surprisingly, we obtain results that are statistically indistinguishable from this when applying just the control questions alone. In follow-up experiments, we hope to determine whether it is the direct effect of control questions upon participants' understanding – or the common knowledge thereof – that drives our results.

Keywords: asset market experiment, price bubbles, bounded rationality.

JEL codes: C92, D84, G12.

[*] We thank the Centre for Experimental Economics at the University of Copenhagen for access to their laboratory and subject pool, and our respective universities for financial support. We thank Robert Slonim and Jean-Robert Tyran for valuable discussions, seminar audiences at the University of Copenhagen and The University of Sydney, and participants at the EURO Working Group Workshop in Graz in November 2009, the DGPE Workshop in Ebeltoft in November 2009, the 4th Australian Workshop on Experimental Economics in Christchurch in December 2009, and the Asia-Pacific meetings of the Economic Science Association in Melbourne in February 2010.

[†] To whom correspondence should be addressed.

1. INTRODUCTION

In asset market experiments of the type pioneered by Smith, Suchanek and Williams (1988), hereinafter SSW, it is well-known that price bubbles – 'trade in high volume at prices that are considerably at variance from intrinsic value' (King et al 1993, p. 183) – followed by crashes are commonplace when participants are inexperienced. In the two decades since SSW first documented this pattern, an extensive body of literature – much of which is summarised in Porter and Smith (2008) – has sought to extinguish these bubbles by applying a very wide range of manipulations to the market environment. These efforts have largely proven fruitless, such that the conventional wisdom remains that the only condition known to eliminate these bubbles is repeated experience in a stationary market environment, as part of the same group of participants.

In this paper, we build upon two aspects of this previous research. Firstly, motivated by somewhat different concerns to our own, previous research has shown that when applied separately, neither dividend certainty (Porter and Smith 1995) nor the elimination of speculation (Lei, Noussair and Plott 2001) is sufficient to extinguish price bubbles. We conjecture that since the effect of each of these treatments is to reduce the complexity of the decision environment facing a trader, then applying them jointly should yield a more pronounced effect. Secondly, we observe that unlike some other branches of experimental research, the literature on asset markets does not consistently make use of control questions. We thus investigate the effect of a particularly robust set of such questions in controlling for confusion or misunderstanding on the part of participants.

Our findings are provocative. We first show that a 'sledgehammer' treatment combining dividend certainty, no speculation and our control questions suffices to essentially eliminate bubbles even in inexperienced markets. Given the rather heavy-handed nature of the intervention, this result is perhaps not unexpected. More surprisingly, we obtain essentially the same result using just the control questions alone – the marginal effect of applying dividend certainty and no speculation over that of the control questions alone is statistically insignificant for three out of four measures of the magnitude of a bubble.

In the next section, we briefly review the most pertinent previous research. We then describe our design and procedures, and document our results. In follow-up experiments, we plan to identify whether it is the direct effect of control questions on understanding, or the common knowledge thereof, that explains our results. We close with some brief concluding remarks, including methodological implications for research on experimental markets.

2. REVIEW OF RELATED LITERATURE

The canonical design of an experimental asset market is due to SSW.[1] In this design each market typically has between eight and twelve traders, each of whom is given an initial endowment of experimental money and shares which they can trade in a computerised double auction. The market operates over fifteen trading periods, with each period typically running for four minutes.

At the end of each period, each share pays a dividend to its current owner. This dividend is the same for all shares, and the probability distribution from which it is drawn is common information. In particular, the dividend takes values of 0, 8, 28 or 60 currency units, each with equal probability, such that the expected dividend in each period is 24. After the fifteenth dividend has been paid, shares have no further value. The intrinsic value of each share is thus 24 times the number of remaining dividend payments, so it is 360 in the first period and declines by 24 after each dividend has been paid. Not only are these facts common information, they are also clearly presented to participants in the form of an 'Average Holding Value Table'. Table 1 shows a typical example of such a table, taken from Porter and Smith (1995).

[Table 1 about here.]

By designing this relatively simple market environment, SSW originally set out to create:

> a "transparent" asset trading market where shares had a well-defined intrinsic value based on common trader information concerning share expected, or average, dividend value. Using these experimental results as a baseline, the research program originally was expected to inquire if price bubbles – trading away from intrinsic value – could be created by controlling information or other elements. (Plott and Smith 2008, p. 12.)

[1] The following description relates to Design 4 in SSW, which is the standard and most extensively-studied set of parameters.

Contrary to expectation, this intentionally simple design has instead been found to consistently generate price bubbles and crashes when participants are inexperienced. By way of illustration, Figure 1 depicts the median transaction prices in each trading period for six baseline sessions reported in Haruvy, Noussair and Powell (2009), hereinafter HNP.

[Figure 1 about here.]

In this figure the lower, dark, stepped line represents the time path of intrinsic value while the upper, light, stepped line represents the *maximum* possible holding value of a share in the event that it pays the highest possible dividend value of 60 in each of the remaining periods. Figure 1 shows that prices tend to start out below intrinsic value in the first few periods before rising steeply above it – frequently to a level in excess of the maximum value – before crashing back toward intrinsic value as the end of the market approaches.

One explanation for why initial trades typically occur at prices below intrinsic value is that they could be motivated by risk aversion on the part of sellers. On the other hand, the only *rational* explanation for purchases at prices exceeding the maximum dividend value is that these trades are motivated by speculation.[2] That is, the only reason why a rational trader would be willing to pay such a price would be if she believed there was a second trader to whom she could resell at an even higher price – and for this to be the case, this second trader must himself either be irrational or else believe there is some third trader to whom he can resell at a yet higher price, and so on.

Building upon such reasoning, it can be argued that even though the *dividend process is common information*, this is insufficient to induce *common knowledge of rationality*. For this reason, a price bubble can occur even when all traders are indeed rational and correctly understand the dividend structure of the asset – but there are at least some traders who believe that some others may not. Alternatively, and more directly, it could simply be the case that overpriced transactions are the result of actual irrationality (or confusion, or decision errors) on the part of some traders.[3]

To examine the hypothesis that low initial transaction prices might be due to risk aversion, Porter and Smith (1995) study markets in which dividends are certain and equal to the

[2] Of course, this is not to say that trades at less extreme prices might not also be motivated by speculation.

[3] We follow Palan (2009) in defining an overpriced transaction as one that occurs at a price in excess of the maximum dividend holding value.

expected dividend value of 24 in the original SSW design. They nonetheless observe the familiar bubble-and-crash pattern in their inexperienced markets, indicating that risk attitudes are not the major cause of price deviations from intrinsic value.

The hypothesis that price bubbles are due to failure of common knowledge of rationality is consistent with the observation that prices track more closely to intrinsic value as participants gain repeated experience in the same market environment as part of the same group of traders (SSW; van Boening, Williams and LaMaster 1993). With experience, it appears that traders become increasingly confident that their counterparts will not behave irrationally, and so perceive less opportunity for profitable speculation. Nonetheless, it is unclear from this result to what extent it is actual rationality, as opposed to the common knowledge thereof, that is improved with experience.

Conclusive evidence of *actual irrationality* in inexperienced markets is provided by Lei et al (2001), who control for speculative motives by assigning each participant to a role either as a buyer (with no opportunity for resale) or seller (with no opportunity to repurchase). They find that 38 percent of trades in such markets occur at prices in excess of the *maximum* dividend value where, given that it is not possible to resell, the buyer is sure to incur a net loss from such a trade.[4]

Consistent with the results of both Porter and Smith (1995) and Lei et al (2001), the bubble-and-crash pattern has proven highly robust to a very wide range of treatment manipulations involving various aspects of the market environment when participants are inexperienced. Much of this work is surveyed in Porter and Smith (2008).

One conspicuous exception is provided by Noussair and Tucker (2006) who find that when a complete set of futures markets is provided, price bubbles in the spot market are eliminated. These futures markets are opened sequentially in reverse order – starting with the futures market for period fifteen – *before* the spot market is opened. One implication of this procedure is that fully half of the session time is taken up with this sequential opening of futures markets before the first period of spot trade commences. Noussair and Tucker (2006) explicitly acknowledge that this is done to 'facilitate the solution of the backward induction

[4] In the Lei et al no-speculation sessions the dividend is either 20 or 40 with equal probability, such that the maximum dividend value is 4/3 times intrinsic value. By contrast, in the SSW design the maximum dividend value is 60/24 = 2.5 times intrinsic value. This difference contributes to the relatively high incidence of overpriced transactions in the Lei et al data.

problem and the comprehension of the decision environment' (p. 169) and 'to facilitate the backward reasoning that is required for agents to realize that the expected future dividend stream corresponds to a limit price for a rational trader' (p. 172).

Noussair and Tucker report one of the very few treatments to eliminate the bubble-and-crash phenomenon for inexperienced participants in the SSW environment. However, it remains unresolved whether their result holds because they provide a complete set of markets (as implied by theory) or, as they themselves suggest, because their treatment hammers home the logic of backward induction to participants.[5] Accordingly, one aspect of our approach in this paper is to ask whether price bubbles can be eliminated by highlighting the logic of backward induction, without providing a complete set of markets.

3. DESIGN AND PROCEDURES

Motivated by the findings reviewed above, our research in this paper aims to determine whether price bubbles in SSW-type markets can be eliminated in inexperienced markets by means of two sets of treatment interventions, either individually or jointly:

Reducing the complexity of the decision environment. It is known that neither the dividend certainty treatment of Porter and Smith (1995) nor the no-speculation treatment of Lei et al (2001) is sufficient on its own to prevent bubbles from forming in inexperienced markets.[6] Nonetheless, we observe that the effect of each of these treatments is to simplify the decision environment confronting a trader. Under dividend certainty, considerations of risk preference are removed, while the Lei et al treatment eliminates strategic considerations associated with speculation. We therefore conjecture that combining these into a *Certainty-No-speculation* treatment should yield an even greater reduction in complexity relative to the baseline case of *Uncertainty-Speculation*.

---

5 Subsequent to commencing the work reported here, we also became aware of the results of Lei and Vesely (2009). In their experiment there is a pre-market phase during which participants passively experience the realisation and accrual of a stream of dividends at periodic intervals corresponding to the length of a trading period. We observe that this procedure is similar to that of Noussair and Tucker insofar as participants first gain experience in a related activity – in the case of Noussair and Tucker, trading in futures markets – for a length of time equal to the life of the experimental asset itself, prior to the commencement of trade in the asset market proper. Both results may thus also be related to the robust finding, noted above, that price bubbles tend to be diminished with repeated experience as part of the same group of participants.

6 However, we observe that Lei et al depart from standard SSW parameters both in that their dividend has only two as opposed to four possible realisations, and their asset has a life of twelve periods as opposed to fifteen. We are not aware of any attempts to replicate their result under classic SSW parameters; however we think it unlikely that their treatment would be *more* effective in the more complex SSW environment.

Controlling for confusion on the part of participants. Although it is standard procedure in SSW-type markets to carefully explain the dividend process and provide participants with an Average Holding Value Table similar to Table 1, little is known regarding the extent to which participants either understand the information or make use of the table. In particular, in contrast to some other branches of experimental economics research, the literature on asset market experiments following SSW does not consistently make use of control questions to check on participants' understanding of the decision environment. However, given the evidence of confusion or decision errors identified by Lei et al, we sought to control for such misunderstandings by implementing a more robust structure of control questions than is typical in this literature.

3.1 Initial experiments

Our initial experiments were conducted at the Laboratory for Experimental Economics at the University of Copenhagen in October and November 2009. We adopt the classic SSW parameters of fifteen double auction trading periods each lasting four minutes, with a dividend that takes values of 0, 8, 28 or 60 currency units, each with equal probability, at the end of each period.[7] We oversubscribed sessions to ensure that there would be exactly ten participants in each market, such that in our no-speculation markets there were five traders assigned the role of buyers and five sellers. Finally we operated two completely independent markets in each session, for a total of twenty participants. No participant had ever taken part in any previous asset market experiment. Each session lasted up to 2.5 hours, and the average earnings were 230 Danish kroner (approximately 46 US dollars). The experiments were conducted in English, and the computerised market was programmed using the z-Tree environment (Fischbacher 2007).

At the start of each session, we first distributed and read aloud the first three pages of the instructions which dealt with the mechanics of using the computer interface to make price offers and to buy and sell shares.[8] This was followed by a ten-minute practice period, which did not count toward participants' earnings. Thus note that participants completed this practice task before they had been told about the dividend structure of the asset or how their earnings would be determined. We next circulated and read aloud the remainder of the

7 Details of technical parameters, such as endowments and exchange rates, are summarised in the Appendix.

8 The full text of the instructions is available as an Annexe to our working paper.

instructions, dealing with the buyer and seller roles (for the no-speculation sessions only), dividend structure of the asset, holding value table and calculation of earnings. Following this, we required participants to complete the control questions, which we describe in detail below. After the conclusion of the experiment, participants were also asked to complete a questionnaire.

In the first instance, our experiments focused on two treatments. Firstly, *CN-Q* is an extremely heavy-handed treatment that combines dividend certainty, no speculation and our battery of control questions. Secondly, *US-Q* is a standard SSW market in which dividends are uncertain and speculation is permitted, but with the imposition of our control questions. Thus *US-Q* enables us to identify what part of the effect observed under *CN-Q* is attributable to our control questions, and what part is due to the combination of dividend certainty and no speculation. Note that in both treatments, it is common knowledge that all participants must complete the questions, and that the experiment does not begin until all participants (in both markets) have answered all of them correctly. Our initial set of experiments yielded six independent observations in each of these treatments, where each market is treated as the unit of observation.[9]

3.2 Control questions

Prior to constructing our control questions, we conducted a thorough search of the literature for appropriate precedents, and could identify very few. Given that control questions are typically only included in working papers and do not find their way into final publications, we cannot claim on these grounds that they are seldom used, but it is nonetheless clear to us that their use is not universal. Moreover, many of the examples we identified relate to aspects of the market that are novel to a specific paper (for example, the futures market treatment of Noussair and Tucker 2006), as opposed to the standard SSW environment itself. In short, the existing literature provides little clear guidance as to what to include in an appropriate set of control questions.

[9] In addition, we were able to recover partial data from a seventh *US-Q* market in which we suffered a fatal server crash which we could not restore while the session was in progress. However, we omit the data from this crashed market in our statistical analysis.

Our main intervention involved two sets of control questions – one framed from the perspective of buying a share, and the second framed from the perspective of selling.[10] Since one aspect of our interest in control questions was to train participants in the logic of backward induction without providing a complete set of futures markets, we included fifteen questions in each frame, ordered from period fifteen to period one.[11] For example, the first buyer control question asked:

> Suppose that you buy one share in period 15 and that you keep it until the end of the market (i.e. until period 15). What is the {average} total dividend that you will receive from this share?[12]

Similarly, the second seller control question asked:

> Suppose that you sell one share in period 14 and that you do not buy it back. What is the {average} total dividend that you give up on this share?

Our control questions thus effectively require each participant to enter the values from the final column of the holding value table twice, from the bottom up. The vast majority of participants learned to do this relatively quickly, and without requiring any assistance from the experimenters. However, in each session there were also up to 20 percent of participants who took much longer – in some cases over twenty minutes – requiring further instruction from an experimenter in the process.

4. RESULTS

4.1 Results of initial experiments

Figure 2 shows the time path of median transaction prices for each trading period in each of our six *CN-Q* markets, which combine dividend certainty, no speculation and our battery of control questions. This shows that under our rather heavy-handed 'sledgehammer' treatment, prices track exceptionally closely to intrinsic value for the most part in each of the markets. However, the final few trades in market 071 provide a conspicuous exception to this finding.

---

[10] Thus note that in our no-speculation sessions every participant answers both sets of questions, even though their role will be constrained to either be a buyer or a seller. Moreover, they do not learn which role they have been assigned until after the control questions have been completed.

[11] In addition, all participants were also required to answer a set of four general control questions.

[12] The curly braces indicate that the word 'average' was omitted from this question in the *CN* sessions.

To understand this aberration, note that in a no-speculation market it is only possible for each unit of stock to transact at most one time. In recognition of this, we increased the number of shares on issue in our *CN*-*Q* markets in an effort to ensure the liquidity of the market.[13] Nonetheless, in two *CN*-*Q* markets the maximum turnover of 100 percent was achieved. Thus although there is no rational explanation for the very high prices paid in the final trades in market 071, they do reflect market power on the part of the sellers of the last remaining units.[14]

[Figure 2 about here.]

The corresponding median transaction price paths for our six *US*-*Q* markets are shown in Figure 3. These are standard SSW-type markets in which dividends are uncertain and speculation is allowed, except that all participants must first complete our battery of control questions. Consistent with previous research on SSW-type markets, we observe a tendency for prices to start out below intrinsic value in the first period – an effect that was not observed to nearly the same extent in the *CN*-*Q* data. However from period two onward, prices in four of the six markets track very closely in line with intrinsic value. (In a seventh market, the prices tracked intrinsic value from period two through to twelve, at which time we experienced a fatal server crash; the data from this crashed market are not included in the analysis that we report below.) In the remaining two markets we observe a residual tendency for shares to trade at prices above intrinsic value through the second half of the life of the asset, but even this is mild compared to the mispricing typically observed in inexperienced SSW markets such as the HNP baseline data in Figure 1. Thus overall, although our *US*-*Q* markets do not track intrinsic value as precisely as our *CN*-*Q* markets, they also clearly do not display the pronounced tendency to bubble and crash that is the norm for inexperienced participants.

[Figure 3 about here.]

To formalise these initial impressions, we follow the literature on asset market experiments in computing a range of standard measures of the severity of a price bubble. Before reporting

13 Under *CN*-*Q* each of the five sellers was endowed with 30 shares, for a total float of 150. Under *US*-*Q* the median trader's endowment included four shares, giving a total float of 40. Details of technical parameters are summarised in the Appendix.

14 In this market, there were two trades in period eleven, two in period twelve, and one in period fifteen. All five of these transactions, as well as the nine trades that took place in period ten, involved only two buyers.

this analysis, we make two preliminary observations. Firstly, our initial experiments did not include any baseline SSW markets without control questions, although we plan to collect our own baseline data prior to finalising our results. Accordingly, for the purpose of this preliminary analysis we use the HNP baseline data to provide an indicative comparison to our two main treatments. Although the HNP baseline sessions employ very similar parameters and procedures to our own, such comparisons can only be suggestive since they also involve a different subject population.[15]

Secondly, because our *Certainty-No-speculation* treatment imposes additional restrictions not present in standard SSW markets, some care must be exercised in the selection of bubble measures to ensure a fair basis for comparison. In particular, in a no-speculation market each share can transact at most once, whereas in a market in which speculation is possible each share can be resold and repurchased multiple times. For this reason, the standard measures of *Turnover* and *Normalised Absolute Price Deviation* (King et al 1993) should not be used to compare markets in which speculation is permitted to ones in which it is not.[16] Bearing this in mind we selected a set of four bubble measures, which we define below, as the basis for our comparisons. For each of these measures, a larger value indicates a more severe price bubble.

Amplitude (King et al 1993) measures the overall magnitude of peak-to-trough deviations in mean period transaction prices from intrinsic value, normalised by the initial intrinsic value:

$$Amplitude = \max_t \left[\left(\bar{P}_t - f_t \right) / f_1 \right] - \min_t \left[\left(\bar{P}_t - f_t \right) / f_1 \right]$$

where $\bar{P}_t$ is the mean transaction price in period t and f_t is intrinsic value in period t.

Duration (Porter and Smith 1995) is the length of the longest sequence of periods over which the difference between the mean transaction price and intrinsic value increases continually:

$$Duration = \max\left(m : \bar{P}_t - f_t < \bar{P}_{t+1} - f_{t+1} < \ldots < \bar{P}_{t+m} - f_{t+m} \right)$$

15 Specifically, we employ the same initial endowments as HNP, although their markets have nine traders instead of ten. We also derive our software and procedures from those of HNP, although the exact wording of our instructions differs substantially. The HNP baseline sessions were conducted in Tilburg and Dallas.

16 Similarly, measures of price volatility and extreme overpricing should be interpreted with caution when comparing markets in which dividends are certain to ones in which they are uncertain.

Total Dispersion (Haruvy and Noussair, 2006) measures the aggregate *absolute* discrepancy between the median transaction price in each period from intrinsic value in that period:

$$Total\ Dispersion = \sum_{t=1}^{T} \left| \tilde{P}_t - f_t \right|$$

where $\tilde{P}_t$ is the median transaction price in period t.

Average Proximity is a new measure of absolute mispricing at the level of individual transactions, which we have defined specifically to be comparable between treatments in which speculation is permitted and ones in which it is not. It is defined as the average absolute deviation of each individual transaction price from intrinsic value:

$$Average\ Proximity = \sum_{t=1}^{T} \sum_{i=1}^{q_t} \left| P_{it} - f_t \right| \Big/ \sum_{t=1}^{T} q_t$$

where q_t is the number of transactions, and P_{it} is the price of the ith transaction, in period t.[17]

Our analysis of these bubble measures is reported in Table 2. The first section of this table reports values of each of the measures for each of the six HNP baseline sessions, along with the treatment mean. The next two sections report the corresponding values for our *CN-Q* and *US-Q* markets respectively. The final three rows of the table report *p*-values for Wilcoxon rank-sum tests of the null hypothesis that the distribution of each of the bubble measures is the same in the HNP and *CN-Q*, HNP and *US-Q*, and *CN-Q* and *US-Q* markets respectively.

The treatment averages in Table 2 confirm the visual impressions obtained from a comparison of Figures 1 through 3. On average, we obtain smaller values of all four bubble measures under our *CN-Q* treatment than in the HNP baseline data. Moreover, the formal tests reported in the third from bottom row of Table 2 confirm that these differences are statistically significant at a level of between $p = 0.0547$ (*Amplitude*) and $p = 0.0027$ (*Duration*). Recall that *CN-Q* combines dividend certainty, no speculation and our battery of control questions. Given this exceedingly heavy-handed intervention, it should come as little surprise that price bubbles are significantly diminished as a result.

17 *Average Proximity* is thus related to the standard measure of *Normalised Average Price Deviation* (King et al 1993), in which the denominator is replaced by the total number of shares in the market, to create a measure that penalises both absolute price deviations and the volume of trade.

Treatment *US-Q* enables us to identify what part of the effect observed under *CN-Q* is due to the combination of dividend certainty and no speculation, and what part is attributable to our control questions. This is a standard SSW market, except that participants must first complete the control questions. Again, we observe on average smaller values of all four bubble measures under *US-Q* than in the HNP baseline. For three of the four measures, the treatment mean under *US-Q* lies between those observed under *CN-Q* and the HNP baseline – the exception to this is that we observe on average slightly lower *Amplitude* under *US-Q* than *CN-Q*. The tests reported in the second-last row of Table 2 confirm that all four measures are significantly smaller under *US-Q* than in the HNP baseline, at a significance level of between $p = 0.0887$ (*Duration*) and $p = 0.0065$ (*Total Dispersion* and *Average Proximity*). These tests thus confirm that simply requiring all participants to complete our control questions – without the additional simplifications imposed under our *CN-Q* treatment – is alone sufficient to significantly diminish the magnitude of price bubbles.

Finally, by comparing *US-Q* to *CN-Q* we can quantify the additional effect attributable to dividend certainty and no speculation, over and above that of our control questions alone. These tests are reported in the final row of Table 2. For three out of four measures, there is no significant difference between our *US-Q* and *CN-Q* markets ($p \geq 0.1093$). The one measure for which we obtain a significant effect is *Duration*: we find that price bubbles are shorter under *CN-Q* compared to *US-Q* at a significance level of $p = 0.0495$. Thus with the exception of lower *Duration*, we find no significant effect of reducing complexity in the decision environment beyond that achieved by controlling for confusion on the part of participants.

We summarise the main results of our initial experiments as follows: we find that price bubbles are essentially eliminated in SSW-style markets with inexperienced traders under a rather heavy-handed combination of dividend certainty, no speculation and control questions. However, we are also able to replicate this result using our control questions alone. In a between-groups comparison we find that with the exception of *Duration* there is no significant difference between our full 'sledgehammer' and a treatment in which we apply the control questions alone. Thus our control questions alone are largely sufficient to eliminate the bubble-and-crash phenomenon that has puzzled researchers for over twenty years.

5. CONCLUSION

When SSW first devised the design of their experiment, they intended it be a particularly simple and transparent bubble-free environment that would serve as a baseline for research into other factors that might contribute to the formation of bubbles. Nonetheless, SSW were initially sanguine about their observation of price bubbles with inexperienced participants, as they observed that rational expectations theory was not necessarily violated even if all traders were indeed rational but simply lacked common knowledge of this fact. This interpretation was consistent with the observation that bubbles were diminished with repeated experience as part of the same group of participants. However, the results of Lei et al (2001) provided conclusive evidence of actual irrationality in inexperienced participants.

In this paper, we demonstrate that price bubbles can be essentially eliminated, even in inexperienced markets, by requiring all participants to complete an extensive battery of control questions prior to the commencement of trade. Thus it appears that the incidence of confusion and decision error, as first documented by Lei et al, is substantially reduced (albeit not completely eliminated) by these control questions. Further, we find that additional heavy-handed interventions to reduce complexity in the decision environment – in the form of dividend certainty and the elimination of speculation – have no significant effect beyond that of our control questions.

There is one further aspect of our findings that we plan to explore in our follow-up experiments. In all of our experiments to date, it has been common knowledge that all participants are required to answer the control questions, and that the experiment does not begin until all participants have answered all of the questions correctly. Thus our finding that price bubbles are essentially eliminated may not be attributable solely to the control questions themselves, but to the combination of the control questions and this common knowledge. The results of our initial experiments point to the possibility of disaggregating these effects by implementing our control questions in an environment without common knowledge. By so doing, we hope to at last resolve whether – in an environment in which confusion and decision error have been substantially mitigated by the control questions – an absence of common knowledge is alone sufficient to ignite a price bubble, as first conjectured by SSW.

Our results also have broader methodological implications for research on experimental markets. Whereas SSW specifically set out to design a particularly simple asset market experiment, policy makers are today increasingly turning to applied experimental economics research to provide 'wind tunnel' tests of quite nuanced aspects of institutional design in settings such as markets for electricity and carbon abatement. It is sometimes the case that these experiments challenge the cognitive limitations of standard experimental participants. The persistent finding of price bubbles in relatively simple environments – one that has been consistently replicated for over two decades – is troubling for experiments in more complex settings. Our results thus reinforce the importance of implementing thorough checks on participants' understanding of the decision environment.

REFERENCES

Caginalp, G., Porter, D. P. and Smith, V. L. (2001). 'Financial bubbles: Excess cash, momentum, and incomplete information', *Journal of Psychology and Financial Markets*, vol. 2(2), pp. 80–99.

Fischbacher, U. (2007). 'z-Tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics*, vol. 10(2), pp. 171–178.

Haruvy, E. E. and Noussair, C. N. (2006). 'The effect of short selling on bubbles and crashes in experimental spot asset markets', *Journal of Finance*, vol. 61(3), pp. 1119–1157.

Haruvy, E. E., Noussair, C. N. and Powell, O. (2009). 'The impact of asset repurchases and issues in an experimental market', Tilburg University.

King, R. L., Smith, V. L., Williams, A. W. and van Boening, M. V. (1993). 'The robustness of bubbles and crashes in experimental markets', in (R. H. Day and P. Chen Eds.), *Nonlinear Dynamics and Evolutionary Economics,* pp. 183–200: Oxford University Press.

Lei, V., Noussair, C. N. and Plott, C. A. (2001). 'Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality', *Econometrica*, vol. 69(4), pp. 831–859.

Lei, V. and Vesely, F. (2009). 'Market efficiency: Evidence from a no-bubble asset market experiment', *Pacific Economic Review*, vol. 14(2), pp. 246–258.

Noussair, C. N. and Tucker, S. J. (2006). 'Futures markets and bubble formation in experimental asset markets', *Pacific Economic Review*, vol. 11(2), pp. 167–184.

Palan, S. (2009). *Bubbles and Crashes in Financial Asset Markets*, Lecture Notes in Economics and Mathematical Systems, vol. 626: Springer.

Plott, C. A. and Smith, V. L. (2008). 'Markets', in (C. A. Plott and V. L. Smith Eds.), *Handbook of Experimental Economics Results,* pp. 3–13: North Holland.

Porter, D. P. and Smith, V. L. (1995). 'Futures contracting and dividend uncertainty in experimental asset markets', *Journal of Business*, vol. 68(4), pp. 509–541.

Porter, D. P. and Smith, V. L. (2008). 'Price bubbles', in (C. A. Plott and V. L. Smith Eds.), *Handbook of Experimental Economics Results,* pp. 247–255: North Holland.

Smith, V. L., Suchanek, G. L. and Williams, A. W. (1988). 'Bubbles, crashes, and endogenous expectations in experimental spot asset markets', *Econometrica*, vol. 56(5), pp. 1119–1151.

van Boening, M. V., Williams, A. W. and LaMaster, S. (1993). 'Price bubbles and crashes in experimental call markets', *Economics Letters*, vol. 41(2), pp. 179–185.

TABLE 1: AVERAGE HOLDING VALUE TABLE
(FROM PORTER AND SMITH 1995)

End Period	Begin Period	Periods Held	×	Average per Period Dividend Value	=	Average per Unit Inventory Value
15	1	15	×	24	=	360
15	2	14	×	24	=	336
15	3	13	×	24	=	312
15	4	12	×	24	=	288
15	5	11	×	24	=	264
15	6	10	×	24	=	240
15	7	9	×	24	=	216
15	8	8	×	24	=	192
15	9	7	×	24	=	168
15	10	6	×	24	=	144
15	11	5	×	24	=	120
15	12	4	×	24	=	96
15	13	3	×	24	=	72
15	14	2	×	24	=	48
15	15	1	×	24	=	24

TABLE 2: BUBBLE MEASURE ANALYSIS

	Amplitude	*Duration*	*Total Dispersion*	*Average Proximity*
HNP Baseline Sessions				
Session 56	0.83	10	923.0	119.40
Session 57	2.56	7	4229.0	276.42
Session 58	0.72	13	1904.5	169.15
Session 62	0.93	7	1265.0	136.88
Session 63	1.65	9	2557.5	188.47
Session 64	0.81	9	1310.5	99.86
Treatment Average	1.25	9.17	2031.58	165.03
***CN-Q* Markets**				
Market 061	0.51	2	263.5	94.05
Market 062	0.15	2	80.0	27.76
Market 071	2.77	3	1749.5	97.91
Market 072	0.19	2	43.5	29.27
Market 081	0.13	2	77.0	13.92
Market 082	0.03	2	11.0	2.86
Treatment Average	0.63	2.17	370.75	44.29
***US-Q* Markets**				
Market 092	0.27	5	68.0	20.79
Market 101	0.58	2	347.5	39.95
Market 102	0.73	4	472.0	60.25
Market 111	0.84	9	1104.0	106.61
Market 121	0.41	10	664.5	42.90
Market 122	0.55	2	342.0	46.73
Treatment Average	0.56	5.33	499.67	52.87
Wilcoxon rank-sum test (*p*-value)				
HNP vs. *CN-Q*	0.0547*	0.0027***	0.0163**	0.0039***
HNP vs. *US-Q*	0.0250**	0.0887*	0.0065***	0.0065***
CN-Q vs. *US-Q*	0.1093	0.0495**	0.1495	0.3367

For the rank-sum tests, *, ** and *** denote significance at the 0.1, 0.05 and 0.01 levels respectively.

FIGURE 1: MEDIAN TRANSACTION PRICES IN HNP BASELINE DATA

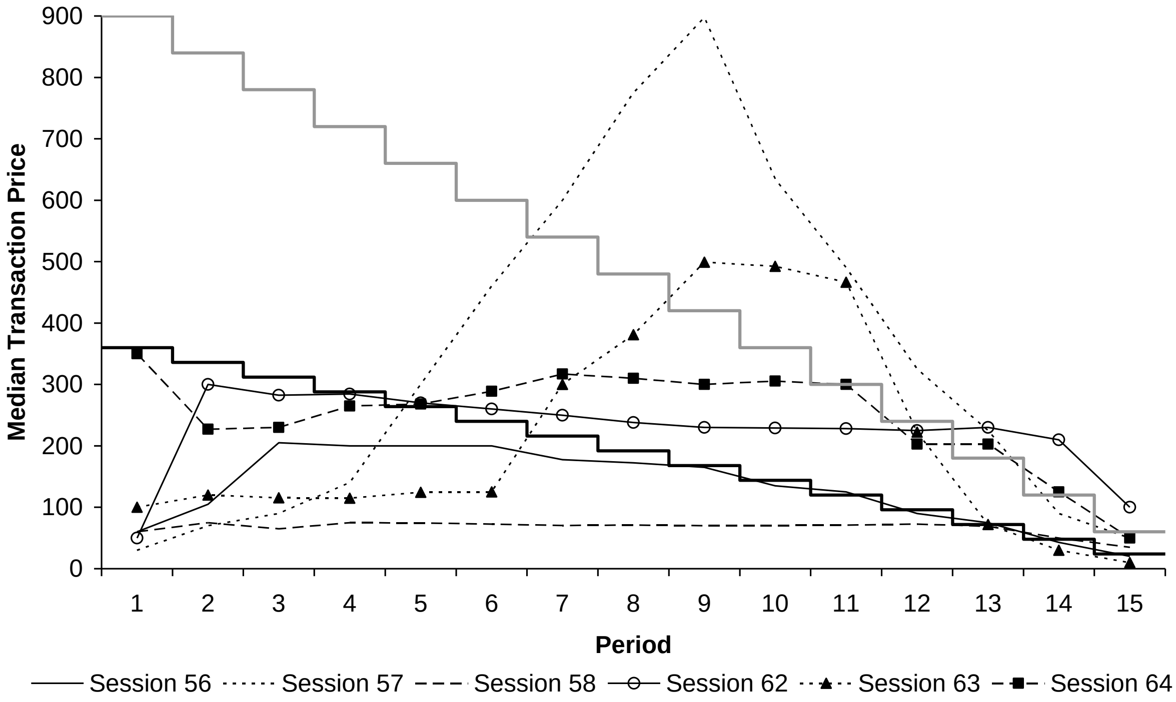

FIGURE 2: MEDIAN TRANSACTION PRICES IN CN-Q MARKETS

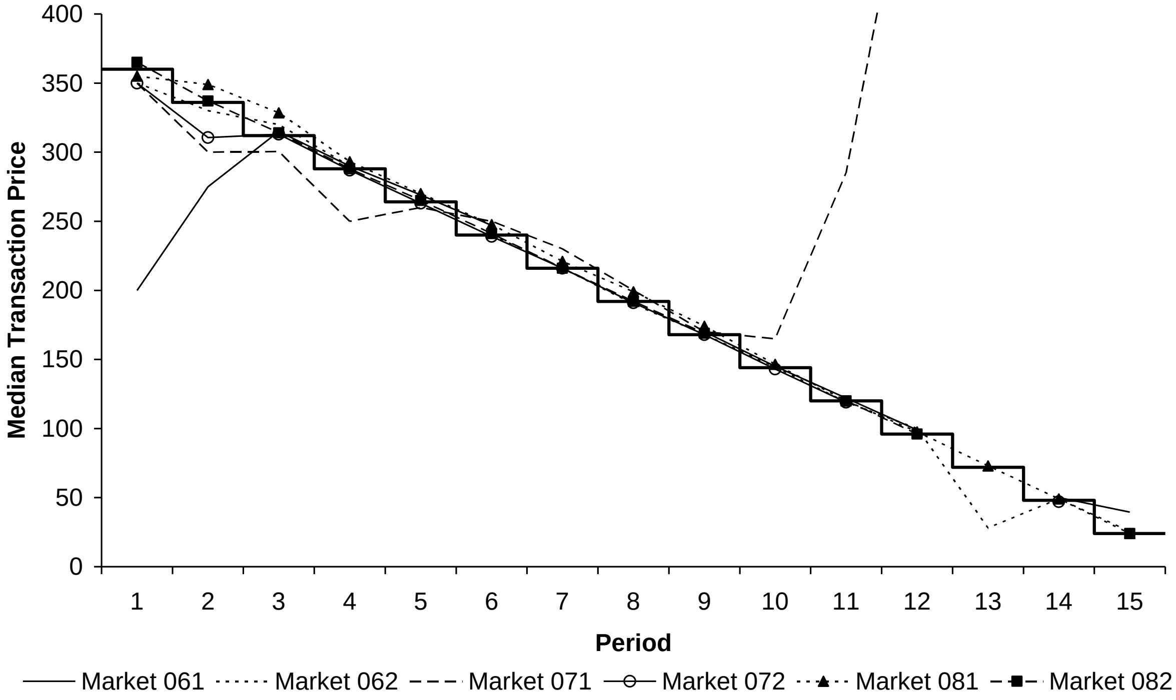

FIGURE 3: MEDIAN TRANSACTION PRICES IN US-Q MARKETS

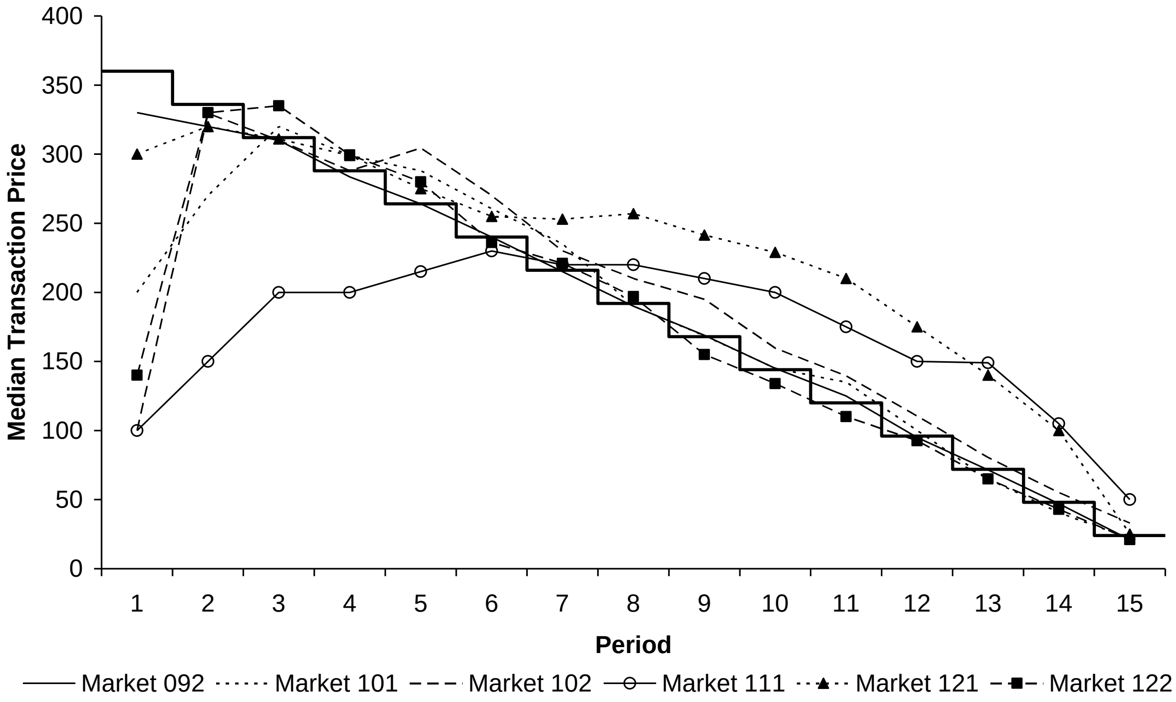

APPENDIX: TECHNICAL PARAMETERS OF EXPERIMENT DESIGN

Table A1 summarises the initial endowments and exchange rates for our *US-Q* and *CN-Q* markets. The *US-Q* endowments are derived by doubling the original endowments in Design 4 of SSW, and are thus identical to those used in the baseline sessions from HNP. In a no-speculation market, endowments must differ quite radically from one in which speculation is allowed, since each unit of stock may transact at most once. Whereas Lei et al (2001) endowed each of their sellers with 20 shares, we increased this number to 30 to account for the fact that our market has a life of 15 periods as opposed to 12 in their design.

[Insert Table A1 about here]

Another consideration in the construction of our endowments was the findings of Caginalp, Porter and Smith (2001) and Haruvy and Noussair (2006) that a high ratio of cash to stock increases the severity of price bubbles. To take account of this, we took care to ensure that the ratio of cash to stock in our *CN-Q* markets matched those of our *US-Q* design – and thus also of both SSW and HNP – thereby providing a clean basis for comparison.

Specifically, for the median (Type II) trader in our *US-Q* markets, the ratio of her initial cash endowment to the initial intrinsic value of her stock endowment is 0.8125. From this, the initial endowments of the high cash (Type I) and high stock (Type III) traders are derived by respectively deducting / adding two units of stock and adding / deducting 720 units of cash, thereby ensuring that the initial endowments of all three types have identical intrinsic value. Since there are equal numbers of Type I and III traders, the overall initial ratio of cash to stock for the market as a whole is thus identical to that of the median trader.

Similarly, in our *CN-Q* design, the ratio of the initial total cash to the initial intrinsic value of the total stock in the market is 0.8182. However in these markets, all cash is initially in the hands of the buyers, and most stock is initially in the hands of the sellers. Given that the initial market cash-to-stock ratio is less than one, we also give the buyers a small initial stock holding. This ensures that the initial endowments of both buyers and sellers are of equal intrinsic value.

TABLE A1: ENDOWMENTS AND EXCHANGE RATES

	US-Q			*CN-Q*	
Endowment type	US I	US II	US III	Seller	Buyer
Number of traders of this type	3	4	3	5	5
Initial stock	2	4	6	30	3
Initial cash	1,890	1,170	450	0	9,720
Endowment value (ECU)	2,610			10,800	
Exchange rate (DKK/ECU)	1/11			1/45	
Endowment value (DKK)	237.27			240	
Total Stock of Units (TSU)	40			150	

ANNEXE (NOT FOR PUBLICATION): INSTRUCTIONS FOR THE EXPERIMENT ‡

General Instructions

This is an experiment on decision making in a market. The instructions are simple and if you follow them carefully and make good decisions, you may earn a considerable amount of money which will be paid to you in cash at the end of the experiment.

Please do not communicate with other participants during the experiment. If you have a question please raise your hand, and an experimenter will assist you.

In this experiment, you have the opportunity to buy or sell in a market. The money used in this market is 'Experimental Currency Units' (ECU). All trading will be done in terms of ECU. The cash payment to you at the end of the experiment will be in Danish kroner.

[CN ONLY] The conversion rate will be **45 ECU to 1 krone**.

[US ONLY] The conversion rate will be **11 ECU to 1 krone**.

You will then be asked to complete a questionnaire, after which you will receive your payment. The entire experiment will last approximately two-and-a-half hours, including half an hour for instructions and practice.

How to use the Computerised Market

On the top right of the screen you will see how much time is left in the current trading period. The items you can buy and sell in the market are called shares. In the centre of your screen you will see the number of shares and the amount of money you currently have.

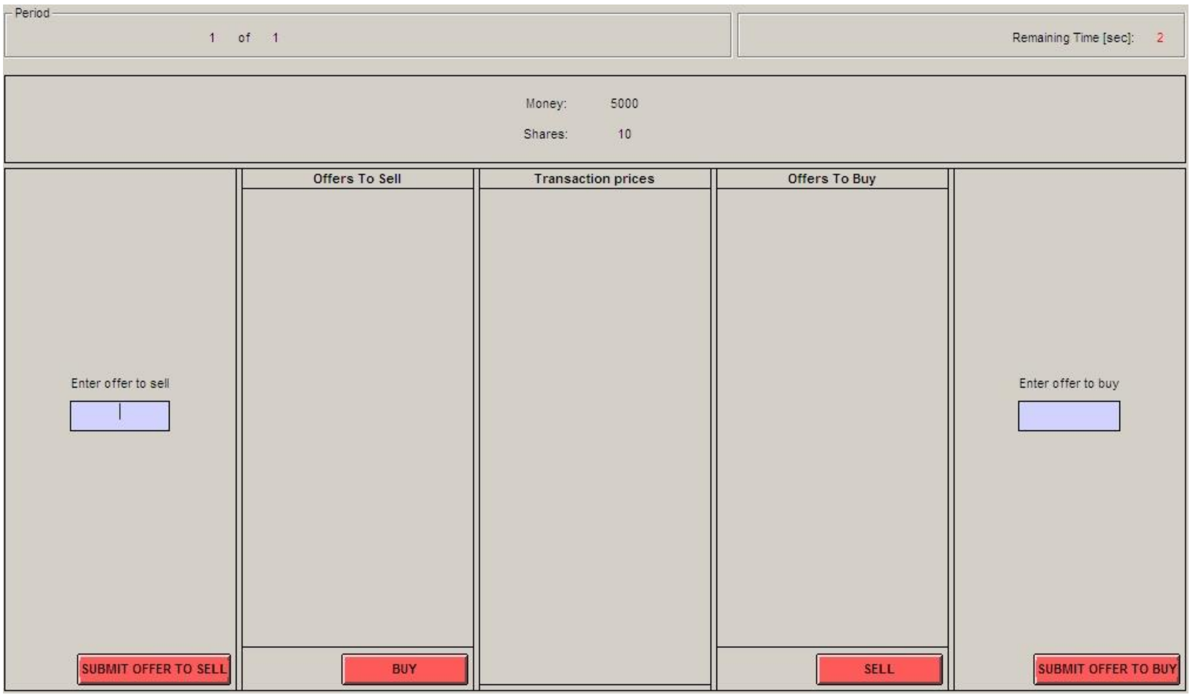

‡ Horizontal rules denote the positions of the page breaks in the original instructions.

The screen can be used to participate in the market in one of four ways.

Making an offer to sell a share, by entering the price at which you would like to sell:

To offer to sell a share, enter the price at which you would like to sell in the box labelled 'Enter offer to sell' on the left of the screen, then click on the button 'Submit offer to sell'.

The second column from left will show a list of offers to sell, each submitted by a different participant. The lowest offer-to-sell price will always be on the bottom of the list. Your own offer will appear in blue. Submitting a new offer will replace your previous one.

Making an offer to buy a share, by entering the price at which you would like to buy:

To offer to buy a share, enter the price at which you would like to buy in the box labelled 'Enter offer to buy' on the right of the screen, then click on the button 'Submit offer to buy'.

The second column from right will show a list of offers to buy, each submitted by a different participant. The highest offer-to-buy price will always be on the bottom of the list. Your own offer will appear in blue. Submitting a new offer will replace your previous one.

Buying a share, by accepting an offer to sell:

You can select an offer to sell in the second column from left by clicking on it. If you click the 'Buy' button at the bottom of this column, you will buy one share at the selected price. However you are not allowed to buy a share from yourself.

When you accept an offer to sell, it will disappear from the list. If you had also placed an offer to buy, it will disappear from the offers to buy list because you have just bought a share.

Selling a share, by accepting an offer to buy:

You can select an offer to buy in the second column from right by clicking on it. If you click the 'Sell' button at the bottom of this column, you will sell one share at the selected price. However you are not allowed to sell a share to yourself.

When you accept an offer to buy, it will disappear from the list. If you had also placed an offer to sell, it will disappear from the offers to sell list because you have just sold a share.

Transaction prices

When you buy a share your money decreases by the price of the purchase. You can only buy a share if you have enough money to pay for it.

When you sell a share your money increases by the price of the sale. You can only sell a share if you owned one to begin with.

In the middle column of the screen, labelled 'Transaction prices', you will see the prices at which shares have traded in the current period.

Practice period

You now have ten minutes to practice buying and selling shares. Your actions in this practice period will not influence your earnings or your position later in the experiment. The only goal is to master the use of the interface.

Please make sure that you successfully submit offers to buy and offers to sell. Also make sure that you successfully accept other people's offers to buy and sell shares.

If you have any questions, please raise your hand and an experimenter will assist you.

[DISTRIBUTED AFTER THE PRACTICE PERIOD]

Specific Instructions for this Experiment

[ALL TREATMENTS]

In each market there are ten participants. *Although there may be more than ten participants in the lab today, you will always be in the same market of ten participants, consisting of yourself and the same set of nine others.*

The market will consist of fifteen trading periods. In each period there will be four minutes during which you can trade shares in exchange for ECU.

At the beginning of the first trading period, your screen will display your initial holdings of money and/or shares. These will not necessarily be the same for all participants in the market.

Any trade that you make will change your holdings of money and shares. These holdings will carry over from one trading period to the next.

[CN ONLY]

You will either be a buyer or a seller in this experiment. Whether you are a buyer or seller will be shown on your trading screen. Each market will consist of five buyers and five sellers.

Specific Instructions for Buyers

If you are a buyer, you will only have the opportunity to buy shares. It is not possible for you to sell shares at any time, so you cannot resell what you buy.

Please note that if you are a buyer then the two buttons labelled 'Submit offer to sell' and 'Sell' will not be available on your trading screen.

Specific Instructions for Sellers

If you are a seller, you will only have the opportunity to sell shares. It is not possible for you to buy shares at any time, so you cannot buy back what you sell.

Please note that if you are a seller then the two buttons labelled 'Submit offer to buy' and 'Buy' will not be available on your trading screen.

[ALL TREATMENTS]

Dividends

Recall that the market consists of fifteen trading periods. Shares are assets with a life of fifteen periods. Each share will pay a dividend to its current owner at the end of each period.

[US ONLY]

The dividend is randomly determined by the computer, and will be the same for all shares. In particular, each share that you own at the end of a period will pay:

- a dividend of 0 ECU with probability 1/4;
- a dividend of 8 ECU with probability 1/4;
- a dividend of 28 ECU with probability 1/4; and
- a dividend of 60 ECU with probability 1/4.

Since each outcome is equally likely, the average dividend is (0+8+28+60) / 4 = 24 ECU in every period.

[CN ONLY]

The dividend will be the same for all shares. In particular, each share that you own at the end of a period will pay a dividend of 24 ECU.

[ALL TREATMENTS]

Dividends will be added to your money balance automatically at the end of each period. After the dividend is paid at the end of the fifteenth trading period, all shares will be worthless and there will be no further earnings possible from them.

[US ONLY]

Average Holding Value Table

You can use your AVERAGE HOLDING VALUE TABLE to help you make decisions.

The first column indicates the Ending Period of the market. The second column indicates the Current Period for which the average holding value is being calculated. The third column gives the Number of Holding Periods from the Current Period to the Ending Period.

The fourth column gives the Average Dividend per Period for each share that you hold. The fifth column gives the Average Holding Value per Share that you hold from the Current Period until the end of the market.

That is, for each share that you hold for the remainder of the market, you will earn on average the amount listed in column five. The value in column five is calculated by multiplying the values in columns three and four.

AVERAGE HOLDING VALUE TABLE

Ending Period	Current Period	Number of Holding Periods	×	Average Dividend Per Period	=	Average Holding Value Per Share
15	1	15		24		360
15	2	14		24		336
15	3	13		24		312
15	4	12		24		288
15	5	11		24		264
15	6	10		24		240
15	7	9		24		216
15	8	8		24		192
15	9	7		24		168
15	10	6		24		144
15	11	5		24		120
15	12	4		24		96
15	13	3		24		72
15	14	2		24		48
15	15	1		24		24

[CN ONLY]

Holding Value Table

You can use your HOLDING VALUE TABLE to help you make decisions.

The first column indicates the Ending Period of the market. The second column indicates the Current Period for which the holding value is being calculated. The third column gives the Number of Holding Periods from the Current Period to the Ending Period.

The fourth column gives the Dividend per Period for each share that you hold. The fifth column gives the Holding Value per Share that you hold from the Current Period until the end of the market.

That is, for each share that you hold for the remainder of the market, you will earn the amount listed in column five. The value in column five is calculated by multiplying the values in columns three and four.

HOLDING VALUE TABLE

Ending Period	Current Period	Number of Holding Periods	×	Dividend Per Period	=	Holding Value Per Share
15	1	15		24		360
15	2	14		24		336
15	3	13		24		312
15	4	12		24		288
15	5	11		24		264
15	6	10		24		240
15	7	9		24		216
15	8	8		24		192
15	9	7		24		168
15	10	6		24		144
15	11	5		24		120
15	12	4		24		96
15	13	3		24		72
15	14	2		24		48
15	15	1		24		24

[ALL TREATMENTS]

Your Earnings

At the end of the market, your earnings will equal the amount of money you have at the end of period fifteen, after the last dividend has been paid.

This amount of money will be equal to:

Any money you had at the beginning of period one
\+ *Any money you received from sales of shares*
− *Any money you spent on purchases of shares*
\+ *Any dividends you received*

At the conclusion of the experiment this amount will be converted into Danish kroner at the rate specified on page one of these instructions, and paid to you in cash.