WORKING paper

BANQUE DE FRANCE
EUROSYSTÈME

A Quantitative Easing Experiment

Adrian Penalver[1], Nobuyuki Hanaki[2], Eizo Akiyama[3], Yukihiko Funaki[4], Ryuichiro Ishikawa[5]

November 2017, WP #651

Abstract

This paper presents experimental evidence that quantitative easing can be effective in raising bond prices even if bonds and cash are perfect substitutes and the path of interest rates is fixed. Despite knowing the fundamental value of bonds, participants in the experiment believed that bond prices would exceed this value when they knew that a central bank would buy a large fraction of the market in a quantitative easing operation. By contrast, there was no average deviation of prices from fundamentals when trading only occurred between participants themselves.[6]

Keywords: Quantitative Easing, Experimental asset markets

JEL classification: C90, D84, G21

[1] Banque de France, adrian.penalver@banque-france.fr
[2] Université Nice Sophia Antipolis , nobuyuki.hanaki@gredeg.cnrs.fr
[3] University of Tsukuba, eizo@sk.tsukuba.ac.jp
[4] University of Waseda, funaki@waseda.jp
[5] University of Waseda, r.ishikawa@waseda.jp
[6] The views expressed herein are those of the authors and should under no circumstances be interpreted as reflecting those of the Banque de France. This research was funded by Project BEAM (Behavioral and Experimental Analyses in Macro-Finance). This paper has benefited from comments received at the SEF conference in Nice, the CEF conference in New York, the Stony Brook Experimental Macroeconomics workshop and a seminar at the Banque de France. We are particularly grateful for the discussion comments of Marianne Andries.

NON-TECHNICAL SUMMARY

Quantitative easing (QE) is an unconventional monetary policy instrument available to central banks when their policy interest rates hit the effective lower bound. QE in its most basic form involves the purchase of government bonds in exchange for central bank reserves with the intention to retain them for a significant length of time. In an era in which central banks pay interest on reserves, QE is an exchange of one interest-bearing liability of the state for another. Under the textbook expectations hypothesis of the yield curve, QE can have no effect on bond yields if there is no change in the expected path of the policy interest rate and therefore no effect on output, employment or wages. In these circumstances, QE would just be an irrelevant shortening of the average maturity of net public debt.

There is, however, strong evidence that QE programmes have moved bond prices and yields, although the scale and duration of such effects is still debated. The academic literature has focused on two departures from the textbook model to explain why QE reduced bond yields. One theory is that central bank money and government bonds are not perfect substitutes perhaps because markets are segmented due to investors' 'preferred habitat' or investors do not like holding the interest rate risk associated with long-term bonds. As a result, a fall in the volume of long-term bonds in private hands can raise the price and lower the yield relative to short-term rates. The alternative explanation is that QE is a means by which central banks can give credibility to forward guidance commitments on policy interest rates. Since central banks will make losses on their bonds if yields rise in proportion to the size of their bond portfolio and unwinding a large balance sheet will take some time, QE backs up the signal that the short-term interest rate will remain low for longer than a time-consistent policy rule would suggest. Lowering the expected path of short-term rates drags down long-term rates through the expectations hypothesis of the term structure.

The laboratory experiments in this paper are designed to see if QE still works even if bonds and reserves are essentially identical and there is no role for monetary policy. The subjects participating in the experiment are put in the situation of large commercial banks that hold a portfolio of bonds and central bank reserves. In a benchmark case, they trade bonds for reserves in a call market over 11 rounds without any involvement by the central bank. Two variants of QE are considered: one in which the central bank buys specified amounts on known dates and holds the purchases to the end of the experiment; and a second in which the central bank buys and then sells all of its accumulated portfolio on known dates.

Under fully rational expectations, neither of these QE operations should change the price. Although it might be thought that prices will rise because the central bank will buy in the future and therefore that prices should rise in advance, each trader has an incentive to try to price just under the existing marginal trader to sell all their bonds at any price above the fundamentals. If each trader follows the same strategy, any price above the fundamental price will be undercut until it disappears.

Figure 1 shows that on average the market price in the benchmark case was very close to the fundamental value. However, prices rose well above the fundamental value under QE with buy and hold, only converging on the fundamental price by period 11. By contrast, when the central bank bought and then sold, prices rose above fundamentals but then fall below as the central bank sold. Since the only difference between these experiments is the nature of the central bank's actions, we ascribe the different price paths observed to differences in beliefs about the impact of the central bank's operations.

These differences in beliefs are tracked throughout the experiment. In each period, subjects were asked to provide forecasts of the prices from that period onwards and were given a bonus on top of their trading performance for the accuracy of their predictions. This allows us to observe how the forecasts and outcomes interact and how they are influenced by the central bank. Subject's prior beliefs about the path of prices are significantly higher when they know there will be QE than in the benchmark case. Under the buy and hold scenario, prices track above median initial expectations, causing an upward revision in price forecasts which in turn sustains higher market prices. Price forecasts respond by less in the buy and sell scenarios, consistent with subjects anticipating that prices will be forced down when the central bank sells.

Overall, these experiments suggest that shifts in beliefs could be part of the explanation for the effect of QE on bond prices (although it doesn't exclude the other channels either). Accepting this possibility does not affect any aspects of the transmission mechanism from bond yields to economic activity and ultimately inflation. Where it does matter is whether QE has effects after net purchases have stopped. According to either of the prevailing theories, it is the stock of purchases that matters so a buy and hold strategy sustains a constant accommodative stance. The policy is sustainable fundamentally. In these experiments, prices only remain high after purchases have stopped because of the persistence in beliefs about prices, a phenomenon that in real markets could be revised or overturned by alternative shocks. The accommodative stance after purchases have stopped is thus sustained only by co-ordination of beliefs, which is a less stable equilibrium.

Figure 1: Average price paths for each treatment group across 11 periods

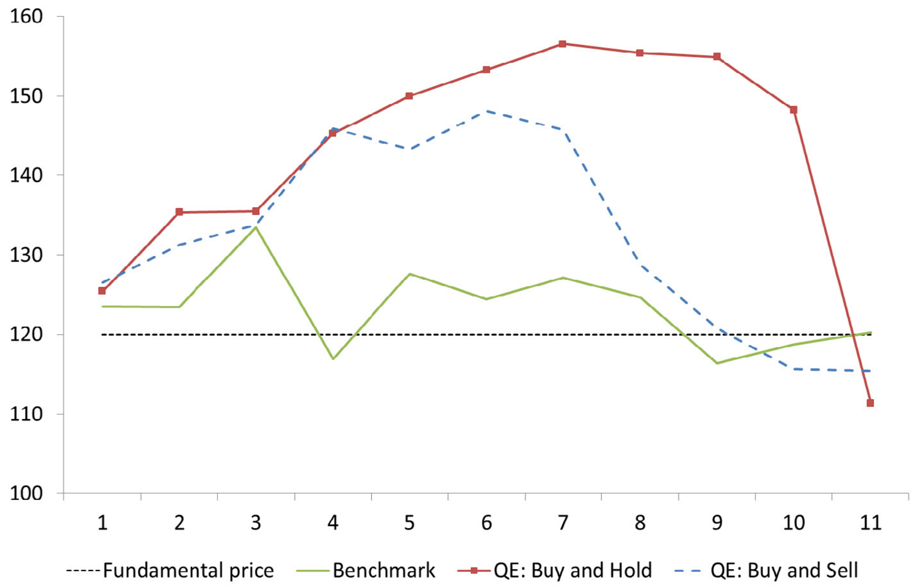

Effets de l'assouplissement quantitatif des banques centrales: apports d'une étude expérimentale

RÉSUMÉ

Cet article présente des éléments de preuve expérimentale attestant que le programme d'assouplissement quantitatif est efficient quand bien même les obligations et la monnaie sont de parfaits substituts et la trajectoire des taux d'intérêt reste inchangée. Bien qu'ayant connaissance de la valeur fondamentale des obligations, les participants à l'expérience sont convaincus que les prix des obligations seraient négociés au-dessus de cette valeur s'ils savaient qu'une banque centrale serait disposée à acquérir une importante part du marché durant les opérations d'assouplissement quantitatif. Il n'existe, en revanche, aucune variation moyenne des prix par rapport aux fondamentaux lorsqu'il est admis que les échanges se réaliseront uniquement entre les participants.

Mots-clés : assouplissement quantitatif, marchés financiers expérimentaux.

1. Introduction

Quantitative easing (QE) in its most basic form is the purchase of government bonds in exchange for central bank reserves with the intention to retain them for a significant length of time.[1] In an era in which interest is paid on reserves, this amounts to the exchange of one interest-bearing liability of the state for another. In textbook models with frictionless and complete markets and fully rational and infinitely living agents and no arbitrage, such a transaction can have no temporary or permanent effects on any macroeconomic variables (Eggertsson and Woodford, 2003). In particular, short-term and long-term interest rates will be unchanged and there will be no effect on output and inflation.[2]

There is, however, strong evidence that QE programmes have moved bond prices and yields, although the scale and duration of such effects is still debated (Joyce, Tong, and Woods, 2011, Krishnamurthy and Vissing-Jorgensen, 2011). The literature has focused on two departures from the textbook model to explain these effects. One theory is that central bank money and government bonds are not perfect substitutes (Tobin, 1958) perhaps because markets are segmented due to investors' 'preferred habitat' (Vayanos and Vila, 2009) or investors do not like holding the interest rate risk associated with long-term bonds. If long-term government bonds and central bank reserves are not perfect substitutes, a fall in the volume of long-term bonds in private hands can raise the price and drop the yield relative to short-term rates. The alternative explanation is that QE is a means by which central banks can give credibility to forward guidance commitments to deviate from established monetary policy behaviour, such as a Taylor rule (Eggertsson and Woodford, 2003). QE reinforces the signal that the short-term rate will remain low for longer than a time-consistent policy rule would suggest. Lowering the expected path of short-term rates drags down long-term rates through the expectations hypothesis of the term structure.

This paper considers a different explanation. As currently implemented, QE is a publicly announced commitment to buy relatively quickly a fixed value of bonds at any price. Indeed since the intermediate objective of the policy is to lower bond yields, the greater the rise in bond prices, the more "successful" the instrument. The central bank is thus an unusual participant in the bond market because it is not deterred from buying by a higher market price (at least up to some point). If there were only one seller of government bonds, he or she could offer to sell at a price at which the central bank was just indifferent between buying and reneging on its commitment at a reputational cost. In

[1] The scale of the purchases and the holding period distinguish quantitative easing from standard open market operations. QE programmes have also bought non-government bonds but this is outside the scope of this paper.

[2] Older irrelevance propositions for open market operations were described in (Wallace, 1981) and (Sargent and Smith, 1987).

other words, once the central bank has publicly committed to buy, there is an exploitable opportunity for sellers collectively.

In a completely competitive market with fully rational agents, common knowledge, no segmentation and assuming the central bank is not buying the entire market, such an effect will not exist. Each participant has an incentive to offer marginally below the offer of the marginal seller to grab a share of this surplus until this is completely eroded away. Since the central bank will buy (and sell) at the fundamental price, there is no reason for the market price to deviate in periods when the central bank is not active.

But what if not everyone is fully rational and instead many are using level-k thinking Nagel (1995)? If $k = 0$ participants think the price will rise, then this will influence the views of $k = 1$ participants which together affect the expectations of $k = 2$ participants and so on. As long as it is believed that there are enough low k participants who believe prices will rise, even a fully rational participant might be induced to make an offer above the fundamental price. This structure is similar to the guessing game of Ledoux made famous by Nagel (1995) (see Buhren, Frank, and Nagel (2012) for historical details).[3]

The objective of this paper is to test whether QE can influence the market price for bonds in an environment in which cash and bonds have identical payoffs in the rational expectations equilibrium and there is no role for monetary policy. In other words, does QE work if both of the channels described in the literature are switched off by construction? And does the holding period of the central bank matter? These questions are investigated using a set of laboratory experiments.

Three treatments are set up in which participants in the experiment are given ”bonds” and ”central bank cash” with equivalent per unit payoffs per period in a rational-expectations equilibrium. The participants can trade bonds between each other in a market over eleven periods. The three treatments vary by the presence of a central bank and the announcements it makes. In the benchmark treatment, there is no central bank involvement and the participants simply trade amongst themselves. In experiments of similar quite simple games, there is a wealth of evidence that prices can diverge significantly from fundamentals (see, for example Bostian and Holt (2009) and Smith, Lohrenz, King, Montague, and Camerer (2014)), so these benchmark results can help us avoid attributing to QE the effects of a bubble that might be an inherent feature of the experimental trading environment.

In the remaining two treatments, the central bank announces at the beginning of the experiment that it will buy a significant fraction of the outstanding bonds in periods 4

[3] In one experiment in Nagel (1995) participants were asked to pick an integer between 0 and 100 and the winner was the closest to $\frac{1}{2}$ of the average response. Any person choosing the unique Nash equilibrium, 0, would have lost as the winning response was 13.

and 5. In one case it holds these bonds until the end of the experiment and in the other it commits to sell them in periods 8 and 9. We find statistically significant evidence that central bank buying raises prices but this is only sustained in the case that the central bank holds its portfolio to the end of the experiment. Prices subsequently fall below the benchmark when the central bank sells.

The paper is organized as follows. Section 2 describes the treatments in more detail and the Nash equilibrium which is identical across the three environments. Section 3 describes the conduct of the experiment and Section 4 presents the main results. Section 5 reports the crucial initial price expectations of the participants and their subsequent evolution both of which are shown to play a crucial role in explaining actual price dynamics. Section 6 discusses the relevance of the results for policy and Section 7 concludes.

2. Set-up

Each treatment had eight identical markets containing 6 participants each who traded between each other over 11 rounds. Each participant was given an initial allocation of 8 bonds and 800 units of cash denominated in ”marks”. Bonds paid a coupon of 6 marks of cash each period and matured at the end of the 11 periods at a face value of 120 marks of cash. Cash was held at the central bank as a deposit and there was a 5% deposit rate on balances at the end of each period. Cash returned its face value at the end of the 11 periods.[4] The participants in the game were thus representative of large commercial banks with settlement accounts at the central bank in countries where interest is paid on reserves. The deposit rate was fixed and known to be fixed throughout the 11 periods.

During each period, the participants had the option to buy or sell bonds for cash in a call market. Participants could post amounts and prices both to buy and to sell but could not borrow from each other nor short-sell either asset. A single market clearing price was calculated and all trades that could be executed at that price were done so. Trades were settled immediately and coupons and deposit interest were paid to the holder at the end of the period. All participants were told the market clearing price at the conclusion of each round of trading, although not the volume traded.

To calculate the fundamental price of the bond in period 11, consider the pay-offs to two strategies: 1) sell the bond for price P and have 1.05P marks at the end of the period; 2) keep the bond and have 120+6 marks at the end of the period. It is trivial to see that the price that makes the bond holder indifferent between the two strategies is P=120 marks. By backward induction, the fundamental price of the bond is 120 marks in each period. Since participants have no incentive to sell below or buy above the fundamental

[4] This set up without the central bank intervention is based on the one used by Bostian and Holt (2009), Smith, Lohrenz, King, Montague, and Camerer (2014).

price, the rational expectations equilibrium is a constant bond price of 120 marks and an indeterminate amount of trading at that price.

To simulate the QE purchase programme, participants were told that the central bank would buy 16 of the 48 bonds outstanding in the market during periods 4 and 5 in a discriminatory auction. The process was thus similar to the method used by the Bank of England and the Federal Reserve Bank of New York to purchase assets.[5] The central bank would attempt to buy 8 bonds at the beginning of each period. If it could not buy 8 bonds in period 4, it attempted to buy the remainder as well as its announced allotment in period 5. If it could not buy enough bonds in both periods at any price, it bought as many as it could and then stopped. (In practice, there was never any problem in buying the bonds according to the central bank's announcement.) When the central bank bought and then sold, it reversed the transactions in a symmetric way at the beginning of periods 8 and 9. In the discriminatory auction, the central bank paid the prices of the cheapest 8 bonds available. All participants were told the price and quantity of bonds bought (or sold) by the central bank but not the identities of the sellers (or buyers).

In the periods in which the central bank was active, there was then a subsequent call market in which participants traded amongst themselves so that a market clearing price of bonds could be established. Assuming common knowledge and rational expectations, the presence of the central bank should not cause the price to deviate from 120 marks in any period. Each participant would like to sell as many bonds to the central bank as possible at a price above the fundamental value but no one participant can buy enough bonds from the others at above the fundamental price to become a monopoly seller. Since the bond auctions will remain competitive, offers will be bid down to the fundamental price in the rational expectations equilibrium. It will not deviate from this price before or after the central bank's actions.

3. Experimental details

Computerized experiments were carried out between June and July 2016 at Waseda University in Japan.[6] 144 students were recruited from across the campus and each student participated in only one experimental session. Each experimental session consisted of 24 subjects divided into four groups of six who performed the experiment together. Thus 8 groups were exposed to the benchmark treatment, 8 groups were exposed to the QE buy and hold treatment and 8 groups were exposed to the QE buy and sell treatment.

Trading took place in call markets as in van Boening, Williams, and LaMaster (1993) Haruvy, Lahav, and Noussair (2007), Akiyama, Hanaki, and Ishikawa (2014), Akiyama,

[5] The Eurosystem purchases bonds directly in the market.

[6] Experiments were computerized using z-Tree (Fischbacher, 2007).

Hanaki, and Ishikawa (2015) and Bosch-Rosa, Meissner, and Bosch-Domènech (2015). In call markets, unlike in continuous double auctions, there will be a single market clearing price for the bond in each period. Subjects could submit buy as well as sell orders by separately specifying a pair of price and quantity for buy and sell orders. Therefore, if a subject decided to submit a buy order in a period, she had to specify the maximum price at which she was willing to buy a bond (b, for bid), and the maximum number of bonds she was willing to buy (d). Similarly, to submit a sell order in a period, a student had to specify the minimum price at which she was willing to sell a bond (a, for ask), and the maximum number of bonds she was willing to sell (s). We imposed three constraints on the orders subjects could submit:

(1) Prices for bids and asks were constrained to be integers between 1 and 2000 and the participants could only buy or sell integer units of bonds;
(2) Subjects could not short sell either asset
(3) a had to be greater than or equal to b

Subjects were asked to submit their orders within 60 seconds.[7] Once all the traders in the market had submitted their orders, the market-clearing price was calculated and all transactions were settled at that price among traders who submitted a maximum buying price no less than (or a minimum selling price no greater than) the market clearing price. If there were multiple market clearing prices, the lowest was chosen so there is no bias in favour of high prices. Participants were rewarded, in part, according to their total cash holdings after their bonds had matured at the end of period 11.

Subjects could not communicate with each other during the experiment to avoid any possibility of explicit collusion. Subjects were also not aware of the bids and offers which minimises the possibility that subjects could try to signal intentions to one another.

In each period, participants were also asked to make forecasts of the market clearing price in all future periods before submitting orders in the call markets.[8] That is, each participants forecasted prices for periods 1 to 11 at the beginning of period 1, forecasted prices for periods 2 to 11 at the beginning of period 2, and so on. This means that a participant submitted a total of 66 price forecasts during their experiment. At the end of each game, participants were given a bonus based on the number of forecasts that were within 10% of the market-clearing prices. Specifically, for each of the forecast that was within 10% range of the market-clearing prices, participants obtained 0.5% of their final cash holding (after the maturity of the bonds and the final interest payments on cash) as a bonus for their forecast performance. Thus, if all 66 forecasts a participant submitted

[7] However, this time limit was not binding in that when it is reached, subjects were encourage to submit their order now by a message on the screen.

[8] Haruvy, Lahav, and Noussair (2007) first implemented this method of eliciting long-run price forecasts.

were within 10% range of realized prices, the participant received an additional 33% of his/her final cash holding as a bonus payment.[9]

4. RESULTS

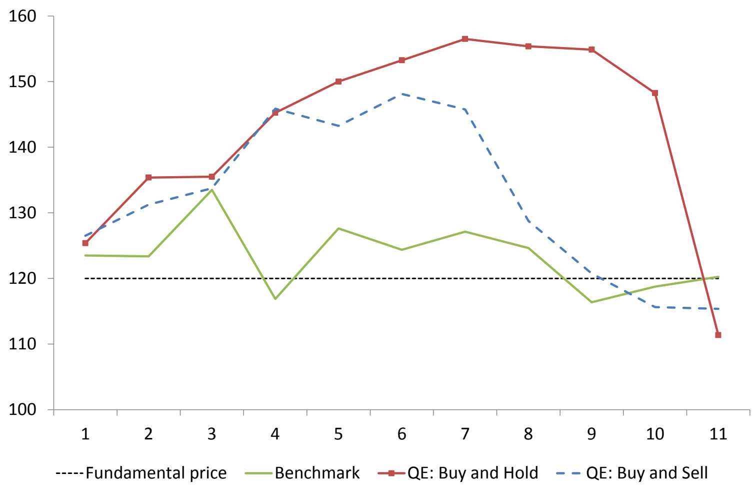

FIGURE 1. Price dynamics in the three treatments.

The main and striking result of the experiment is reported in Figure 1 which shows the unweighted average price paths across the three different treatments. The average price in the benchmark treatment does not deviate far from the fundamental price across the 11 periods. By contrast, the average price in the QE buy and hold treatment rises far above the fundamental price before collapsing at the end. Even more different still, in the QE buy and sell treatment prices rise before the central bank makes its purchases and then fall as the central bank sells, with the average price at the end below the fundamental price. Since the only difference between the three treatments is the known behaviour of the central bank, we are confident in ascribing the differences in outcomes to perceptions about the consequences of the different treatments. The rest of this section substantiates this key finding, including formal statistical significance tests.

Behind the average price effect lies quite considerable variation across the markets in each treatment. Figure 2 illustrates the price paths over the 11 periods in the benchmark treatment for each of the 8 markets alongside the volume-weighted average price. The dynamics of individual markets confirm previous results in the literature that prices can deviate quite substantially from the fundamental price (see, for example, Bostian and

[9] This incentive scheme for the forecast performance is equivalent to the one employed by Akiyama, Hanaki, and Ishikawa (2014, 2015).

Holt (2009), Smith, Lohrenz, King, Montague, and Camerer (2014), Smith, Suchanek, and Williams (1988)). Palan (2013) provides a review of the literature.

However, in our experiment, unlike in the existing literature, the average price path exhibits no obvious pattern and is not statistically different from the fundamental price as one can see from the result of statistical tests reported in the table in Figure 2. There is initially some variance in prices but these converge very close to the fundamental price by period 11 in all 8 markets. This pattern is observed whether we use the volume-weighted price or the unweighted price.

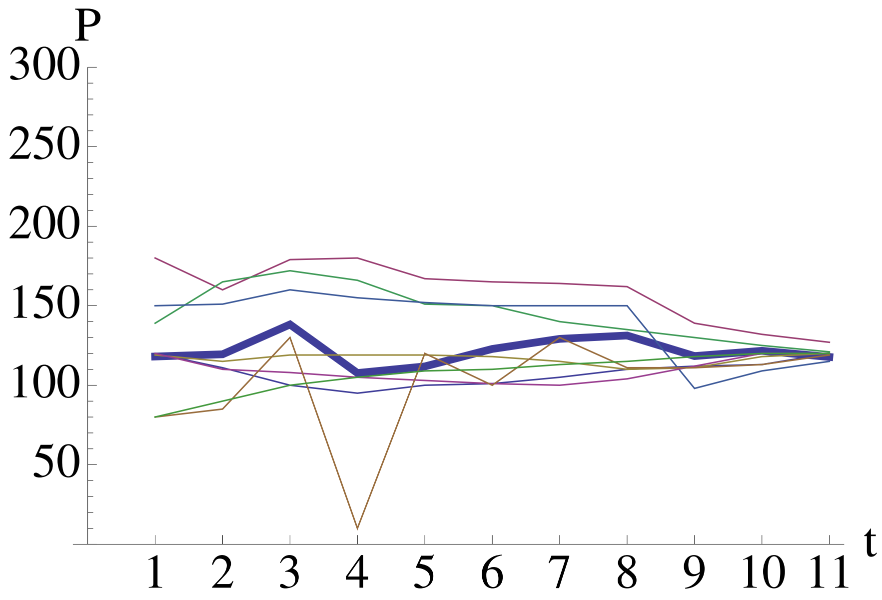

Note: Thin lines are the realization of each of the 8 markets. The thick line is the volume weighted average price across the 8 markets.

	1	2	3	4	5	6	7	8	9	10	11
Unweighted	124	123	134	117	128	124	127	125	116	119	120
P-value	0.78	0.77	0.28	0.87	0.43	0.65	0.41	0.56	0.45	0.65	0.84
Weighted	118	119	138	108	112	123	129	131	118	122	118
P-value	0.85	0.96	0.14	0.52	0.21	0.75	0.16	0.12	0.67	0.55	0.13

Note: P-values are the probability that the sample is drawn from a distribution with a mean equal to the fundamental value of 120 using a two-sided Student's t-test with 7 degrees of freedom.

FIGURE 2. Price dynamics in the benchmark treatment.

TABLE I. COMPARISON OF MEANS

	1	2	3	4	5	6	7	8	9	10	11
Unweighted											
QE Buy	125	135	136	145	150	153	157	155	155	148	111
Benchmark	124	123	134	117	128	124	127	125	116	119	120
P-value	0.92	0.50	0.89	0.20	0.07*	0.04**	0.06*	0.08*	0.05*	0.10	0.62
Weighted											
QE Buy	113	130	136	142	152	162	164	167	171	144	129
Benchmark	118	119	138	108	112	123	129	131	118	122	118
P-value	0.71	0.46	0.88	0.11	0.00***	0.01***	0.02**	0.05**	0.02**	0.19	0.14

Note: P values using a two-sided Student's t-test with 7 degrees of freedom. ***≤ 0.01, **≤ 0.05, *≤ 0.1

Figure 3 illustrates the disperse price paths for the buy and hold treatment. This shows a quite different profile to the benchmark treatment and if nothing else indicates that central bank actions can disturb the price determination process. In one market, the market price scarcely deviated from the fundamental price and in two others, prices were very low initially and then rose dramatically before collapsing in a path resembling a bubble. Prices remained at or above the fundamental price in all 8 markets until round 10.[10] Nevertheless, prices in almost all markets converged towards the fundamental price by period 11.[11] Average prices are statistically significantly above the fundamental price from periods 4 to 9 (reported in the table in Figure 3). This result remains true even if the two markets with an obvious bubble are excluded.

Table I presents simple comparison of means tests which confirm the visual impression that average prices are statistically significantly higher over periods 5 to 9 in the treatment in which the central bank undertakes QE than in the benchmark treatment. (Volume-weighted and simple averages are presented in Table I to remove any suspicion that these results come from unrepresentative prices.)

Figure 4 illustrates the price paths for the buy and then sell treatment. As reported in the table below the figure, average prices in this game were statistically significantly above fundamentals in periods 4 to 8 in the weighted-average case and periods 4, 5 and 8 in the unweighted case.[12] Again, this result holds if we exclude the market with an obvious bubble. Prices were then statistically significantly *below* the fundamental price in the weighted-average case in period 10. Despite these fluctuating dynamics, prices converged to the fundamental price in the final period.

[10] With the exception of 1 price in one period at 119.

[11] No trading takes place at the extremely low price in period 11 in one of the games.

[12] Periods 6 and 7 were quite close to the 10% significance threshold.

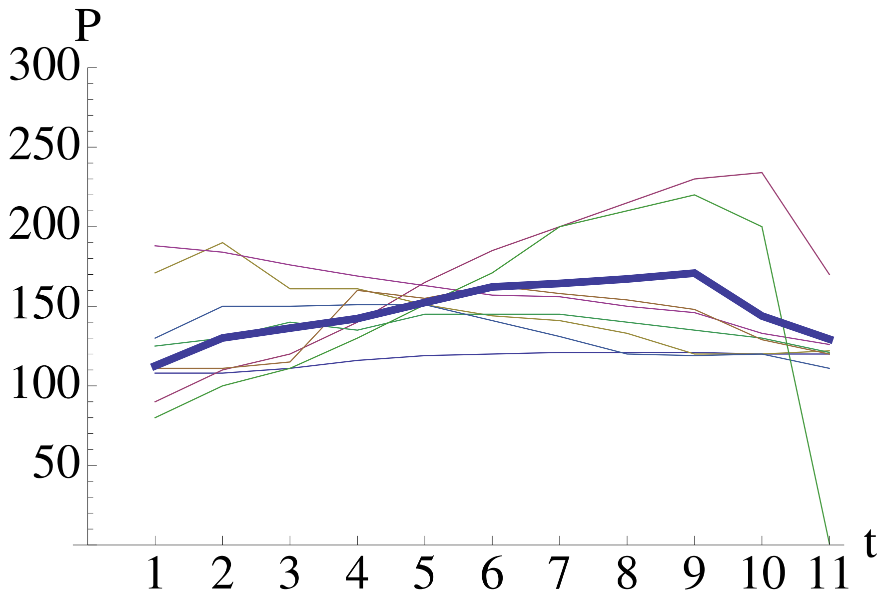

Note: Thin lines are realization of each of 8 markets. Thick line is the volume weighted average price across 8 markets.

	1	2	3	4	5	6	7	8	9	10	11
Unweighted	125	135	136	145	150	153	157	155	155	148	111
P-value	0.70	0.26	0.12	0.01***	0.00***	0.00***	0.01***	0.03**	0.06*	0.11	0.63
Weighted	113	130	136	142	152	162	164	167	171	144	129
P-value	0.48	0.31	0.03**	0.00***	0.00***	0.00***	0.00***	0.01**	0.02**	0.15	0.21

Note: P-values are the probability that the sample is drawn from a distribution with a mean equal to the fundamental value of 120 using a two-sided Student's t-test with 7 degrees of freedom. ***≤ 0.01, **≤ 0.05, *≤ 0.1

FIGURE 3. Price dynamics in the buy and hold treatment.

Table II reports a comparison of means test between the two QE treatments and it is evident that behaviour before, during and immediately after the common buying phase is indistinguishable. But prices fall in the case where the central bank subsequent sells and are statistically significantly below the buy and hold treatment in periods 8, 9 and 10.

It is clear, therefore, that in these experiments QE does make a significant difference for bond prices and that the two QE scenarios have starkly different effects.

The obvious next question to ask is where these differences come from. The next section will highlight the importance of price expectations and in particular the different initial expected paths between the two QE treatments and the benchmark treatment. The remainder of this section shows, by contrast, that there were no other differences in the behaviour of the markets.

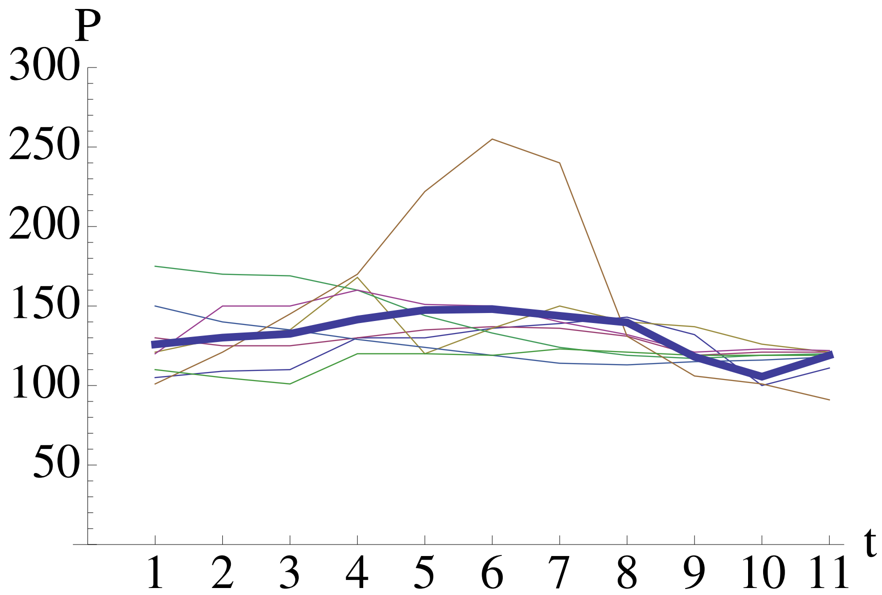

Note: Thin lines are realization of each of 8 markets. Thick line is the volume weighted average price across 8 markets.

	1	2	3	4	5	6	7	8	9	10	11
Unweighted	127	131	134	146	143	148	146	129	121	116	115
P-value	0.49	0.18	0.12	0.01***	0.09*	0.12	0.11	0.05**	0.83	0.25	0.25
Weighted	126	130	133	142	147	148	144	140	118	106	119
P-value	0.48	0.24	0.11	0.01**	0.04**	0.05**	0.08*	0.00***	0.61	0.00***	0.15

Note: P-values using a two-sided Student's t-test with 7 degrees of freedom. ***$\leq$ 0.01, **$\leq$ 0.05, *$\leq$ 0.1

FIGURE 4. Price dynamics in the buy and sell treatment.

TABLE II. COMPARISON OF MEANS

	1	2	3	4	5	6	7	8	9	10	11
Unweighted											
QE Buy	125	135	136	145	150	153	157	155	155	148	111
QE Buy and Sell	127	131	134	146	143	148	146	129	121	116	115
P-value	0.95	0.79	0.89	0.95	0.62	0.78	0.56	0.09*	0.07*	0.08*	0.83
Weighted											
QE Buy	113	130	136	142	152	162	164	167	171	144	129
QE Buy and Sell	126	130	133	142	147	148	144	140	118	106	119
P-value	0.33	1.00	0.69	0.93	0.68	0.33	0.24	0.09*	0.02**	0.04**	0.17

Note: P-values using a two-sided Student's t-test with 7 degrees of freedom. ***$\leq$ 0.01, **$\leq$ 0.05, *$\leq$ 0.1

One might wonder, for example, whether there were differences in market liquidity across the treatments or conjecture that the behaviour of a small number of participants mattered more in the QE treatments than in the benchmark case.

Figure 5 reports a histogram of the absolute value of transactions made by each participant across the three treatments. In two of the three treatments a small minority never bought or sold a bond in the market. And some were very active, buying or selling on average two bonds in each period. What is clear from the figure is that there is no obvious difference in the profile of market participation across the three treatments.

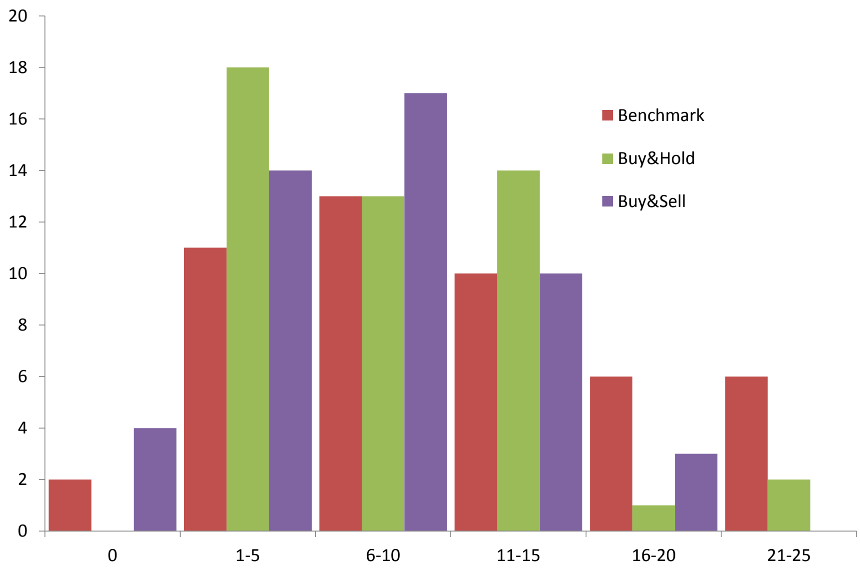

FIGURE 5. Market participation

To assess whether markets in certain treatments were more dominated by some participants than others, we first identified the marginal price setter in each market in each period. This generally corresponded with the individual whose supply or demand order was only partially filled at the market clearing price. On a few occasions two or more partipants posted the same marginal price and were recorded as the marginal price setter with the appropriate fraction. With this information we calculated the average proportion of times each individual participant was the marginal price setter. If there was no heterogeneity of bids and offers within each market, then each participant would be the marginal price setter $\frac{1}{6}$ of the time. If there was one participant who was always the only marginal price setter, then they would have a score of 1 and everyone else in that market, 0.We then calcuated a simple Gini co-efficient to measure the degree of concentration across each market.

A histogram of the Gini co-efficients across the three treatments is reported in Figure 6. There is somewhat more dispersion in the degree of market domination in the QE treatments than the benchmark treatment but little average difference in behaviour. Moreover, in fact, the more concentrated markets in the QE treatments have slightly lower average prices. Overall, therefore, we are confident in ruling out that the higher prices recorded in the QE treatments were due to the behaviour of a few individuals.

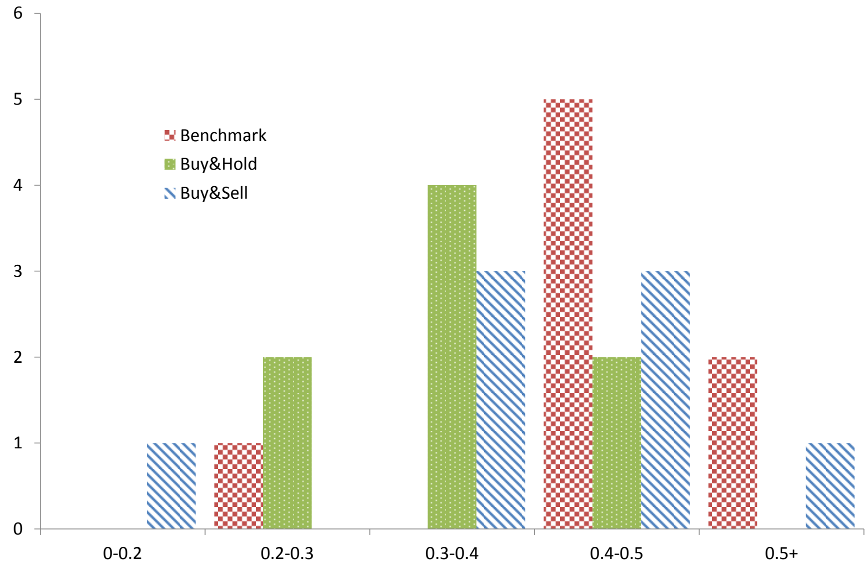

FIGURE 6. Market domination - Gini co-efficients

5. PRICE FORECASTS

To dig a bit deeper into the evolution of prices in the different treatments, it is useful to look at the forecasts made by the participants both at the start of the treatment and how they were updated as a result of experience.

Figure 7 plots the distribution of prior beliefs about market prices in the first period. There are a number of striking features about these distributions. First, the range is very wide. In all three treatments, less than a quarter of participants have initial forecasts within 20 marks either side of the fundamental price and over half predict less than 60 marks or more than 180 marks. Second, the distributions for the two QE treatments are remarkably similar and distinct from the distribution in the benchmark treatment. Third, the distribution in the benchmark treatment has a fat lower tail whilst the distributions of the two QE treatments have fat upper tails.

Figure 8 plots the average median initial price forecast paths for all 11 periods across the 8 markets in each treatment.[13] Illustrations of the initial distributions at different time horizons for each treatment are provided in the appendix. The Figure shows that

[13]This measure of presenting the forecasts was chosen because if the market price in each market was set by the median belief then this measure would deliver the average market price across the eight markets in each treatment. Other measures of the central tendency of the initial forecasts give similar results.

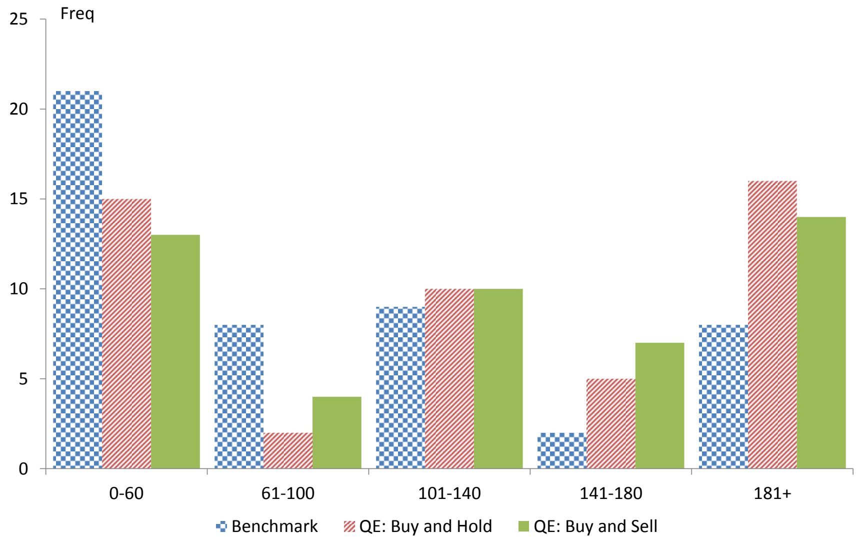

FIGURE 7. Initial forecast distribution - all treatments

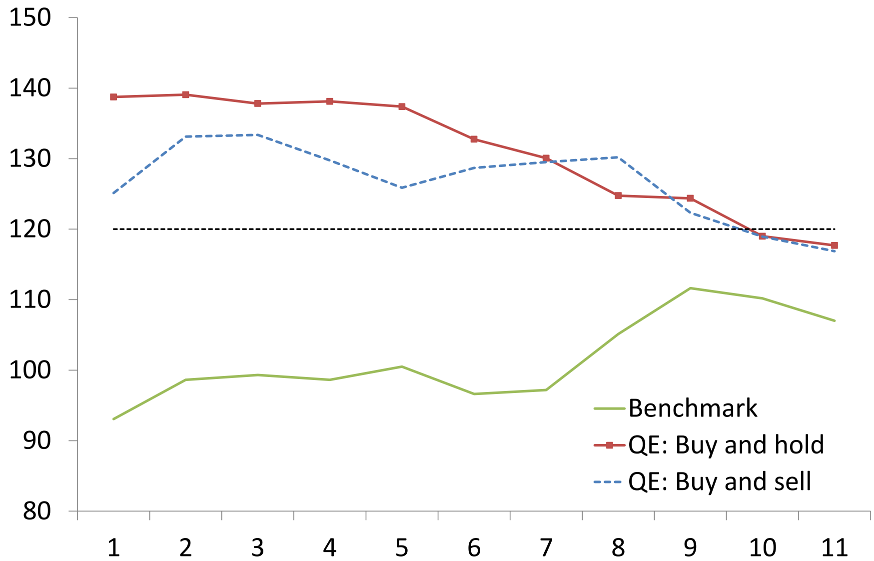

FIGURE 8. Median initial price forecast path

average price expectations are below those for the two QE treatments across the whole initial forecast horizon. Indeed, for a reason that is unclear, the average median price forecast is below the fundamental price throughout.

When the central bank will buy and then hold, there is clearly an expectation that prices will be above fundamentals for much of the experiment. The average median price forecast jumps immediately to around 140 marks and stays at this level until period 5 and then drifts down towards the fundamental value. The average median price forecast in the buy and sell treatment is above the fundamental price but slightly lower than the buy and hold treatment with the exception of period 8.

These initial expectations condition the opening trading round and all subsequent pricing behaviour. We present first the benchmark treatment which is useful background to understand the subsequent dynamics when there is central bank intervention. Figure 9 (best seen in colour) plots the market clearing price for each market in solid lines and the associated median forecasts for each market with crosses of the same colour. These are generally very closely related: high median forecasts are associated with high market clearing prices later in the same period and high prices elicit high forecasts in the subsequent period. These relationships are formalised in Tables III and IV.

Table III reports the co-efficients of panel regressions to explain individual contemporaneous price forecasts. The explanatory variables are the forecast made last period for the current period and last period's market clearing price (which occurs between when the two forecasts are made). Focusing on the first column it is clear that both explanatory variables are extremely important with slightly more weight applied to the lagged price.

Table IV reports the co-efficients of panel regressions explaining the market clearing price. Here the possible explanatory variables are the median contemporaneous price forecast and the lagged price. From regression A it is clear that one could make a very good prediction for the market clearing price if you knew the median price forecast for that period. Regression B shows that if you only had access to the lagged market price (which is the only public information) then this would still be a reliable guide. Regression C shows that all of the information from last period's price relevant for predicting the current period price is contained in the median forecast.

Overall, the foregoing analysis suggests that the outcomes in the early markets had a strong and lasting effect on the subsequent evolution of prices through the effect on beliefs. As a measure of this, the correlation between the realised market price in period 1 and the median forecast for period 10 made in period 10 is 0.75. The initial market prices condition the following period's price forecasts which influences market prices and so on. With relative stability in prices, individual forecasts strongly converge as can be seen in Figure 10. The key observation to note, however, is that without any central bank intervention, there is no average deviation in prices from fundamentals.

With this benchmark analysis in mind, we can now turn to the two cases with central bank intervention, starting with the buy and hold treatment. Figure 11 illustrates that, as in the benchmark case, expectations and prices are closely intertwined, if anything, the correlations are even closer.

TABLE III. PRICE FORECAST EQUATIONS

Dependent variable	**Forecast**
Method	Panel Least Squares
Periods	10
Cross sections	48

	Benchmark	QE: Buy and Hold	QE: Buy and Sell
Constant	13.7	-6.4	12.1***
s.e	(12.8)	(4.6)	(3.8)
Forecast(L)	0.42***	0.05***	0.10***
s.e.	(0.03)	(0.02)	(0.01)
Price(-1)	0.49***	1.01***	0.82***
s.e.	(0.10)	(0.03)	(0.03)
Adjusted R^2	0.38	0.71	0.70

Note: P-values: ***$\leq$ 0.01, **$\leq$ 0.05

TABLE IV. PRICE EQUATIONS - BENCHMARK TREATMENT

Dependent variable	**Price**
Method	Panel Least Squares
Periods	10
Cross sections	8

	A	B	C
Constant	19.4**	44.2***	19.4**
s.e	(9.5)	(9.8)	(9.5)
Median forecast	0.83***		0.83***
s.e.	(0.07)		(0.15)
Price(-1)		0.64***	-0.00
s.e.		(0.08)	(0.14)
Adjusted R^2	0.61	0.46	0.60
Akaike info crit.	8.45	8.78	8.48

Note: P-values: ***$\leq$ 0.01, **$\leq$ 0.05

The convergence of forecasts within each market is also much quicker, see Figure 10. This is consistent with the evidence in Table III that participants only relied on lagged

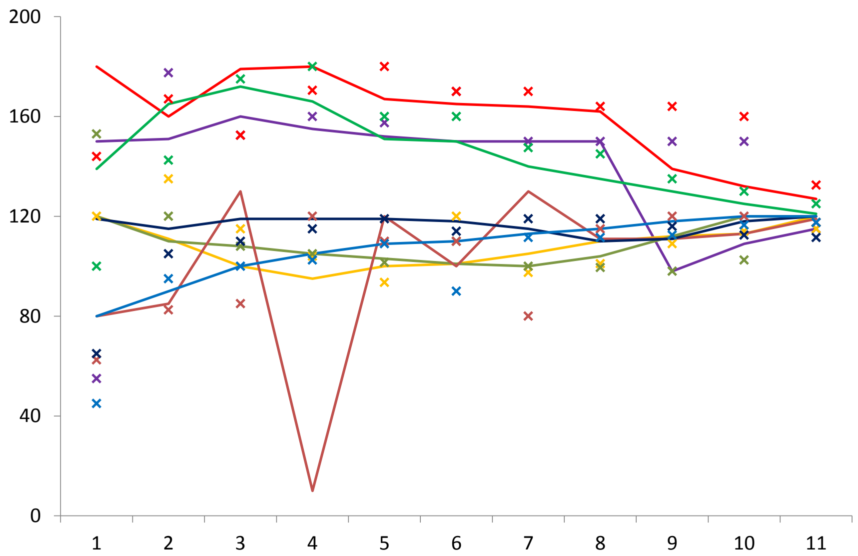

FIGURE 9. Dynamics of prices and the median forecasts in benchmark treatment.

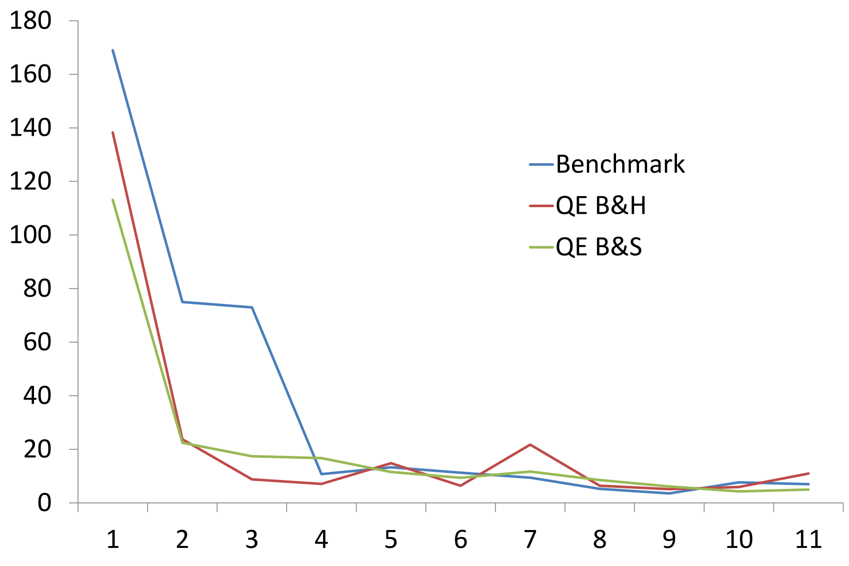

FIGURE 10. Average standard deviations of within market forecasts

prices in generating their forecast suggesting that relied more on market signals for guidance in an environment with greater strategic uncertainty.

Table V reports panel regressions to explain the evolution of market prices. Column A shows that as in the benchmark case, median forecasts are a fairly good predictor of

TABLE V. PRICE EQUATIONS - QE BUY AND HOLD TREATMENT

Dependent variable	**Price**
Method	Panel Least Squares
Periods	10
Cross sections	8

	A	B	C	D
Constant	36.5***	33.4***	33.9***	41.9***
s.e.	(11.7)	(11.8)	(13.5)	(12.5)
Median forecast	0.73***	0.74***		1.37***
s.e.	(0.08)	(0.08)		(0.35)
Price(-1)			0.75***	-0.70*
s.e.			(0.09)	(0.38)
Four		0.08	0.08	0.07
s.e.		(0.06)	(0.06)	(0.06)
Five		0.04	0.04	0.04
s.e.		(0.05)	(0.05)	(0.05)
Adjusted R^2	0.53	0.53	0.46	0.55
Akaike info crit	9.18	9.20	9.34	9.18

Note: P-values: ***$\leq$ 0.01, **$\leq$ 0.05, *$\leq$ 0.10

the market clearing price. Column B adds the average price paid by the central bank in the auctions in each market to see whether these influence subsequent market prices. There is no evidence of an effect in period 5 and in period 4 it is only significant at the 20% level. Columns C and D show that, as before, in the absence of median forecasts, the lagged price is an effective predictor but that the latter adds no independent information if the median forecast is available.

We now turn to the final treatment in which the central bank buys in periods four and five and then sells its portfolio in periods eight and nine. Figure 12 shows that, as both previous cases, price and forecasts are closely linked for each market. And column C of Figure III is very similar to B indicating that as in the benchmark case, partipants place much more weight on the new information contained in market prices than their previous forecast. Thus even though there were different students participating in the two treatments, their underlying behaviour was very similar but different from those in the benchmark case. Figures 12 illustrates the relationship between median forecasts and prices which is formalised in the panel regressions in Table VI. This, too, paints a very similar picture to the buy and hold treatment. The two things to note are that the price

TABLE VI. PRICE EQUATIONS - QE BUY AND SELL TREATMENT

Dependent variable	**Price**
Method	Panel Least Squares
Periods	10
Cross sections	8

	A	B	C	D
Constant	21.8**	24.0***	31.6***	25.1***
s.e.	(8.7)	(8.4)	(9.4)	(8.4)
Median forecast	0.83***	0.81		1.16***
s.e.	(0.06)	(0.06)		(0.24)
Price(-1)			0.76***	-0.36
s.e.			(0.07)	(0.24)
Four		0.05**	0.06**	0.04**
s.e.		(0.02)	(0.02)	(0.02)
Five		0.04	0.03	0.05
s.e.		(0.04)	(0.04)	(0.03)
Eight		-0.05	-0.08*	-0.04
s.e.		(0.04)	(0.05)	(0.04)
Nine		-0.04	-0.07	-0.03
s.e.		(0.05)	(0.05)	(0.05)
Adjusted R^2	0.68	0.71	0.63	0.71
Akaike info crit	8.29	8.26	8.49	8.24

Note: P-values: ***$\leq$ 0.01, **$\leq$ 0.05, *$\leq$ 0.10

paid by the central bank in period 4 is statistically significant at the 5% level but there are negative co-efficients on periods eight and nine which make no economic sense.

Joining the threads of this analysis together, participants in the two QE treatments had distributions of beliefs that put considerable weight on values above fundamentals. This influenced the initial market prices that were then used very heavily to update subsequent price forecasts and thereby the evolution of market prices. There was possibly some marginal information contained in the prices actually paid by the central bank in period 4 but in general the actual operations did not change the evolution of prices very much. The key differences between the QE treatments and the benchmark were the distribution of initial beliefs and the much stronger feedback between prices and subsequent beliefs.

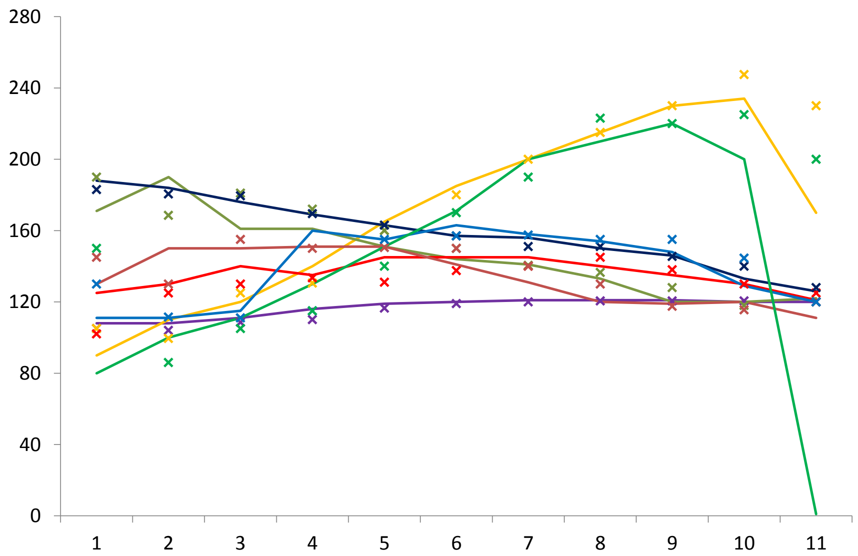

FIGURE 11. Dynamics of prices and the median forecasts in buy and hold treatment.

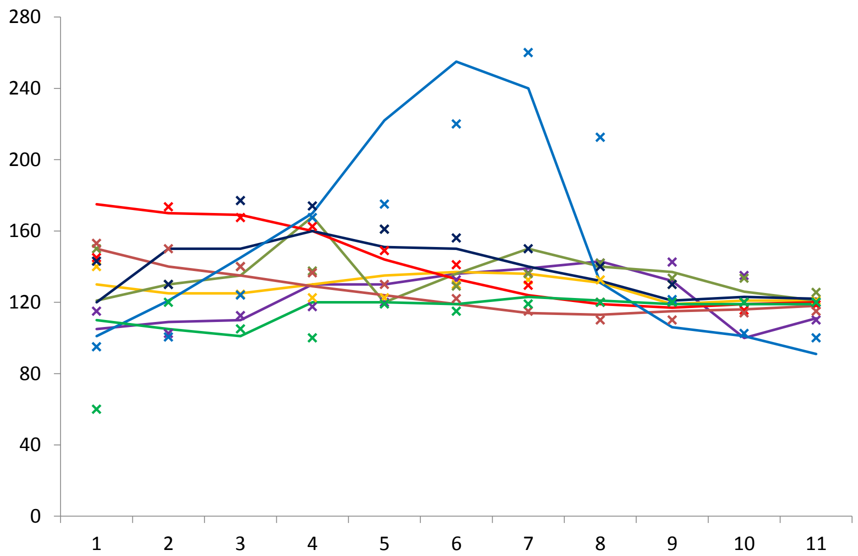

FIGURE 12. Dynamics of prices and the median forecasts in buy and sell treatment.

6. INSIGHTS

What insights do these results offer for the conduct of QE? If we take the results at face value, the first result is a positive one: QE can work even if bonds are completely substitutable with cash. All that is required is that enough traders believe that enough other traders believe that others will bid a higher price. Indeed it is even conceivable that each trader in isolation believes that QE should not work in principle but believes that enough other traders do believe it will work (or believes that others believe that it will work etc) to offer at a higher price. Accepting that shifts in beliefs could be part of the explanation for the effect of QE on bond prices does not affect any aspects of the transmission mechanism from bond yields to economic activity and ultimately inflation. This belief-driven channel also does not exclude the alternative explanations of imperfect substitutability between bonds and cash or the reinforcement of forward guidance.

If QE can work even when bonds and cash are very close substitutes, then this suggests that central banks can buy quite short-term maturity debt. This is consistent with the evolution of QE across several central banks. The Bank of England began its QE programme in March 2009 with a minimum residual maturity of 5 years that was subsequently reduced to 3 years. The European Central Bank began with a minimum residual maturity of 2 years and reduced this to 1 year in December 2016.

A more speculative conclusion one could draw is that central banks need to make QE dominate other potential factors in the minds of traders. For this to occur, QE programmes need to be large and long. Central banks also need to focus their communication on the assets purchased rather than the money created. This might seem trivial but it is notable that the original QE by the Bank of Japan, widely regarded at the time as ineffective, concentrated on the expansion of the money supply and not on the counterpart asset purchases. It might also be important to maintain a large stock of assets yet to be purchased.

It is also likely that the central bank must be prepared to "lose whatever it takes" for it to be fully effective. The belief that prices will rise above fundamentals is directly linked to how much the central bank is expected to overpay. The average student made a profit from the central bank of over 10% of their initial bond portfolio. And these central bank losses were higher in the case that the central bank bought and then sold. Thus it is crucial that central banks communicate that they are focused solely on the macroeconomic outcomes of the policy rather than the impact on the central bank balance sheet (see (Bean, 2009)).

QE can have a sustained effect on prices even after it has stopped buying despite the absence of any fundamental channel. In a situation of high strategic uncertainty, participants used the history of market prices as a strong guide for future prices. QE can have lasting effects because they change the frame through which traders view the market. This does, however, suggest that QE can work when it is the dominant market "narrative" but could lose traction in the event of significant other shocks.

A final implication is that QE can destabilise markets and create bubbles. 3 of the 16 markets in the QE treatments experienced booms and then busts which was not the case in any of the 8 benchmark markets. This possibly reflects the fact that the influence of lagged prices is stronger when there is more strategic uncertainty. Even though the central bank's actions were very simple and common knowledge, the lack of common knowledge of strategies between the participants made the choices of the participants much more difficult.

Of course, all this analysis comes with a large health warning. This is only a single experiment and it remains to be seen whether the results are replicated in future experiments. Bond traders are also likely to be more strategically sophisticated than undergraduate students exposed to bond trading for the first time. We cannot, for example, find any statistically significant factors that explain which participants made money. So there is some danger in extrapolating from this artificial environment to the real world.[14] That said, the world is considerably more complex than this environment and so these experiments can still offer lessons provided they approximate the thought processes and strategies of actual traders.

7. Conclusion

This paper has presented experimental evidence that quantitative easing can still work even if bonds and central bank cash are perfect substitutes. Despite knowing the fundamental value of bonds, participants in the experiment clearly believed that bond prices would trade above this value when they knew that a central bank would buy a large fraction of the outstanding market in a quantitative easing operation. By contrast, there was no average deviation of prices from fundamentals when it was known that trading only occurred between participants. The analysis showed that the central bank influences the interaction between actual prices and forecast prices in a way that boosts prices. The positive effect of quantiative easing purchases on prices could persist even after the central bank had stopped buying but quickly turned into reverse if the central bank subsequently sold off its portfolio.

[14] A similar problem applies in extrapolating from rational expectations models too.

References

Akiyama, E., N. Hanaki, and R. Ishikawa (2014): "How do experienced traders respond to inflows of inexperienced traders? An experimental analysis," *Journal of Economic Dynamics and Control*, 45, 1–18.

——— (2015): "It is not just confusion! Strategic uncertainty in an experimental asset market," *Economic Journal*, forthcoming.

Bean, C. (2009): "Quantitative Easing: An Interim Report," speech at the London Society of Chartered Accountants Annual Lunch, London 13 October 2009.

Bosch-Rosa, C., T. Meissner, and A. Bosch-Domènech (2015): "Cognitive Bubbles," Mimeo, Berlin University of Technology.

Bostian, A. A., and C. A. Holt (2009): "Price bubbles with discounting: A web-based classroom experiment," *Journal of Economic Education*, 40(1), 27–37.

Buhren, C., B. Frank, and R. Nagel (2012): "A Historical Note on the Beauty Contest," MAGKS Papers on Economics 201211.

Eggertsson, G. B., and M. Woodford (2003): "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 34(1), 139–235.

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

Haruvy, E., Y. Lahav, and C. N. Noussair (2007): "Traders' Expectations in Asset Markets: Experimental Evidence," *American Economics Review*, 97(5), 1901–1920.

Joyce, M., M. Tong, and R. Woods (2011): "The United Kingdoms quantitative easing policy: design, operation and impact," *Bank of England Quarterly Bulletin*, 51(3), 200–212.

Krishnamurthy, A., and A. Vissing-Jorgensen (2011): "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy," *Brookings Papers on Economic Activity*, 43(2 (Fall)), 215–287.

Nagel, R. (1995): "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85(5), 1313–26.

Palan, S. (2013): "A Review of bubbles and crashes in experimental asset markets," *Journal of Economic Surveys*, 27(3), 570–588.

SARGENT, T. J., AND B. D. SMITH (1987): "Irrelevance of Open Market Operations in Some Economies with Government Currency Being Dominated in Rate of Return," *American Economic Review*, 77(1), 78–92.

SMITH, A., T. LOHRENZ, J. KING, P. R. MONTAGUE, AND C. F. CAMERER (2014): "Irrational exuberance and neural crash warning signals during endogenous experimental market bubbles," *Proceeding of National Academy of Science, U.S.A.*, 111(29), 10503–10508.

SMITH, V. L., G. L. SUCHANEK, AND A. W. WILLIAMS (1988): "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56(5), 1119–1151.

TOBIN, J. (1958): "Liquidity Preference as Behavior Towards Risk," *Review of Economic Studies*, 25(2), 65–86.

VAN BOENING, M. V., A. W. WILLIAMS, AND S. LAMASTER (1993): "Price bubbles and crashes in experimental call markets," *Economics Letters*, 41(2), 179–185.

VAYANOS, D., AND J.-L. VILA (2009): "A Preferred-Habitat Model of the Term Structure of Interest Rates," NBER Working Papers 15487, National Bureau of Economic Research, Inc.

WALLACE, N. (1981): "A Modigliani-Miller Theorem for Open-Market Operations," *American Economic Review*, 71(3), 267–74.

APPENDIX

TABLE VII. COMPARISON OF MEANS - POOLED RESULTS

	1	2	3	4	5	6	7	8	9	10	11
Unweighted											
Benchmark	134	135	140	131	133	131	131	127	120	120	119
QE Buy	128	133	134	137	136	135	136	134	134	131	117
P-value	0.49	0.86	0.35	0.51	0.61	0.50	0.57	0.30	0.05*	0.07*	0.73

Note: P-values using a two-sided Student's t-test with 23 degrees of freedom. ***$\leq$ 0.01, **$\leq$ 0.05, *$\leq$ 0.1

TABLE VIII. COMPARISON OF MEANS - POOLED RESULTS

	1	2	3	4	5	6	7	8	9	10	11
Unweighted											
QE Buy	128	133	134	137	136	135	136	134	134	131	117
QE Buy and Sell	132	135	136	142	144	142	135	124	121	119	118
P-value	0.45	0.79	0.55	0.20	0.03**	0.10	0.96	0.08*	0.04**	0.04**	0.79

Note: P-values using a two-sided Student's t-test with 23 degrees of freedom. ***$\leq$ 0.01, **$\leq$ 0.05, *$\leq$ 0.1

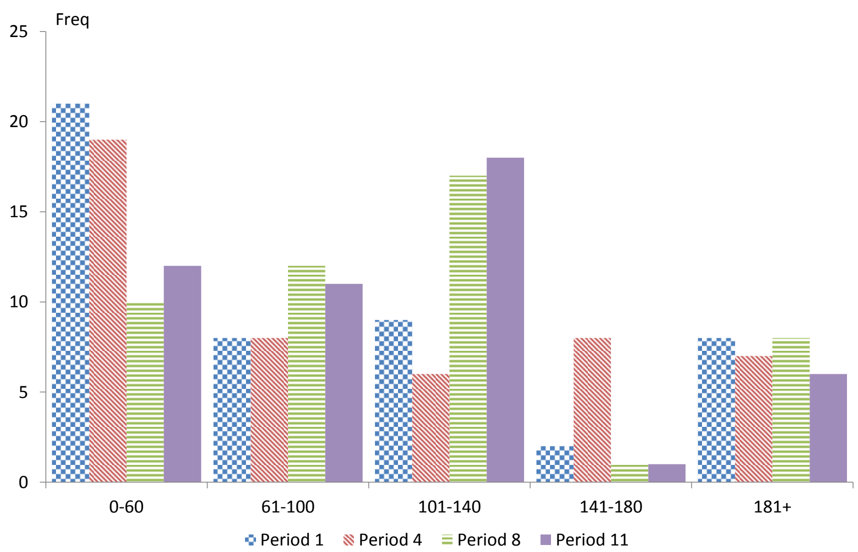

FIGURE 13. Initial forecast distribution - benchmark

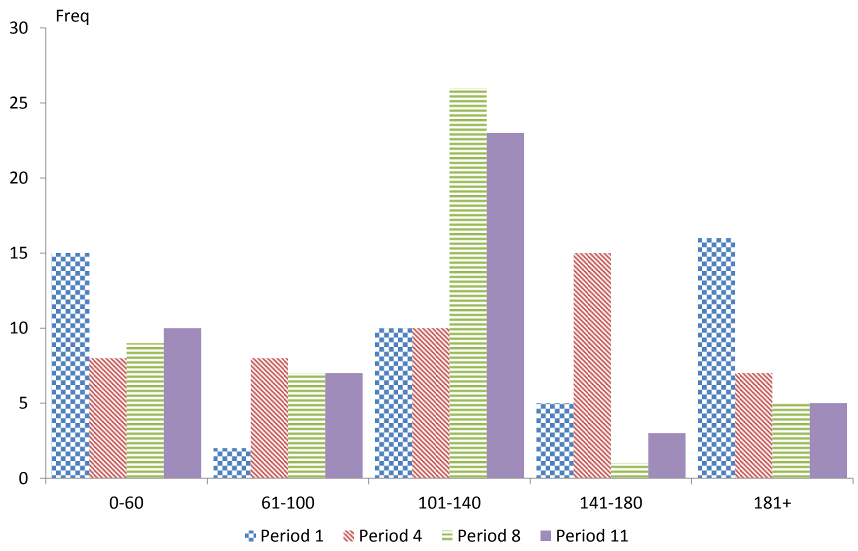

FIGURE 14. Initial forecast distribution - QE Buy and Hold Treatment

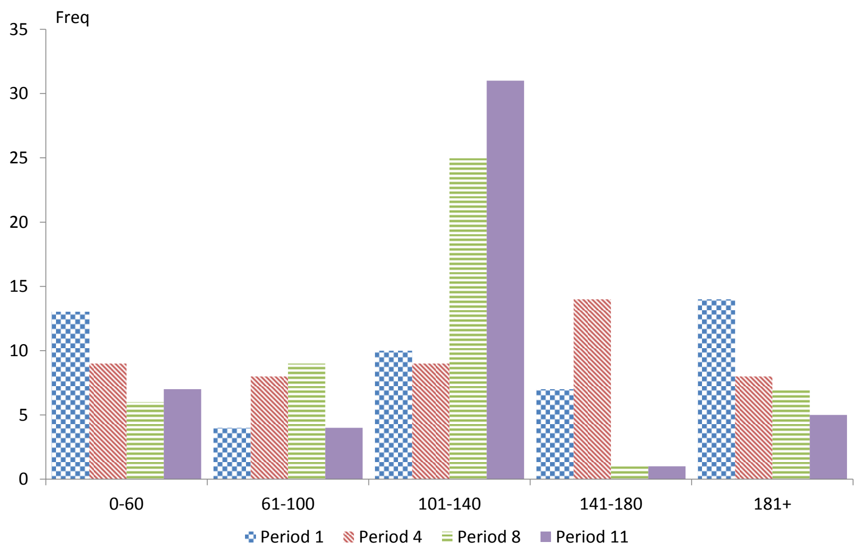

FIGURE 15. Initial forecast distribution - QE Buy and Sell Treatment

Text of the experiment. The instructions for the three treatments begin with information common to all three treatments. The common part of the instructions is for the baseline treatment. After a practice round, the instructions for each treatment are presented as follows. For the buy and hold treatment, an explanation of the buying operation is added after the common information. For the buy and sell treatment, an explanation of the selling operation is added after the instructions for the buy and hold treatment. The handout of the instruction below is distributed to the participants, and the instruction is explained by the movie with the sound that a computer reads out the sentences. Both the movie and sentences to be read out are identical to the instruction below.

7.0.1. *Instructions for Today's Experiment.* We first explain the instructions for todays game. There is a practice period for the game so that you may familiarize yourself with the computer interface before the real experiment. The experiment consists of three games, each of which has 11 periods. After completing the games, we will ask you to respond to a questionnaire and take some quizzes. Your earnings will be paid in cash at the conclusion of the experiment.

Your earnings will consist of a participation fee of 500 yen, and an amount that depends on the results of the games. The questionnaire and quizzes will not impact your earnings. The three games are independent of each other, so that the result of one game does not affect the other games. You will have a short bathroom break before the game begins.

7.0.2. *Today's experiment.* You will participate in a bond trading game in which you trade national bonds in an artificial market. Please listen to the instructions carefully. If you do not understand any part of an instruction, ask for clarification by raising your hand. Moreover, if you have any questions during the experiment, raise your hand and an instructor will come to you and answer your question privately. Throughout the experiment, please respect the following rules.

(1) Do not talk to the other participants during the experiment or the breaks.
    - This may affect the results of the experiment.
(2) Use your mouse or keyboard only when instructed to do so by the instructor; otherwise, it may cause a problem.
    - If any malfunction occurs, all participants will have to restart the game.

7.0.3. *Outline of bond trading game.* You will be divided into several groups. You will not know the identities of the members of each group. Each group will consist of six subjects. You will play the bond trading game with the other five people in the group to which you belong.

7.0.4. *Objectives of the game.* Your objective in this game is to make as much profit as you can. There are two ways of making a profit:

- First, you can realize a profit margin through buying and selling bonds, from dividends on your bond holdings, and from interest on your cash holdings.
- Second, you can make a profit by accurately predicting the future prices of the bonds.

We use Marks as the currency for the experiment. At the end of the experiment, your profit will be converted into Yen (1 Mark = 1 Yen) and paid out to you.

7.0.5. *Earning a profit margin.* You will be given eight bonds and 800 Marks at the beginning of the game. To earn a profit margin by trading, you need to buy bonds at a low price and sell them at a higher price. For example, suppose that you buy a bond for 100 Marks, and that the price of the bond then increases to 120 Marks. If you sell the bond, you earn 120 (selling price) - 100 (purchase price) = 20 Marks profit. In contrast, suppose that you buy a bond for 100 Marks, and that the price of the bond then decreases to 80 Marks. If you sell the bond, you will make 80 (selling price) - 100 (purchase price) = 20 Marks loss. We explain later how the prices are determined.

[... Explanation of the computer interface ...]

In practice, the price you actually pay for a bond may not be the same as the maximum price you are willing to pay. This is because the market price depends on all the orders placed by the market participants. If the market price is greater than the maximum you are willing to pay, then your order will not be processed. This will be further clarified at a later stage.

In practice, the price at which you sell a bond may not be the same as the minimum price at which you are willing to sell. This is because the market price depends on all the orders placed by the market participants. If the market price is lower than your minimum price, then your order will not be processed. This will be further clarified at a later stage.

The most important points for buying and selling bonds are summarized below.

- You can simultaneously place buy and sell orders, or you can place only a buy order or only a sell order. It is also possible to not submit any order.
- If you do not want to submit a buy order, please enter 0 as the quantity to buy. If you do not want to submit a sell order, please enter 0 as the quantity to sell.
- The screen displays an error message if any of the following conditions are violated.

(1) The maximum quantity to sell must be less than or equal to the number of units you hold.
(2) The maximum purchase price multiplied by the quantity to buy must be less than or equal to the cash you have available.
(3) If you simultaneously place buy and sell orders, the maximum purchase price must be less than or equal to the minimum selling price.

The price is set according to the order book within your market. There is a single price for all bonds bought and sold in each period. The price is set to equate the number of buy orders and the number of sell orders.

We explain how the market prices are set by using the following two examples.

Example 1: how the market price is determined Consider the following buy/sell orders placed by four traders.

- Trader 1: One sell order, which can be exmarkted at 10 Marks or higher
- Trader 2: Two sell orders, which can be exmarkted at 40 Marks or higher
- Trader 3: One buy order, which can be exmarkted at 60 Marks or lower
- Trader 4: One buy order, which can be exmarkted at 20 Marks or lower

A seller is willing to sell at the price requested or higher. A buyer is willing to buy at the price specified or lower. As shown above, there is only one bond supplied at 10 Marks. If the price rises to 40 Marks, the number of bonds supplied increases to three. However, only one bond is demanded at 60 Marks. If the price falls to 20 Marks, the quantity demanded increases to two. Therefore, the quantity demanded is equal to the quantity supplied at prices between 21 Marks and 39 Marks. The market price is set at the minimum price in this interval; that is, 21 Marks.

Next we consider the second example.

Example 2: how the market price is determined Consider the following buy/sell orders placed by five traders.

- Trader 1: One sell order, which can be exmarkted at 10 Marks or higher
- Trader 2: One sell order, which can be exmarkted at 30 Marks or higher
- Trader 3: One sell order, which can be exmarkted at 30 Marks or higher
- Trader 4: One buy order, which can be exmarkted at 60 Marks or lower
- Trader 5: One buy order, which can be exmarkted at 30 Marks or lower

As shown above, only one bond is supplied at 10 Marks as in the previous example. If the price rises to 30 Marks, the number of bonds that are supplied increases to three. However, there is only one bond demanded at 60 Marks. If the price falls to 30 Marks, the quantity demanded increases to two. As a result, two transactions can be completed at 30 Marks. In this case, the market price is set at 30 Marks. The orders that are fulfilled are determined as follows.

Priority is given to Trader 1, because he/she requested a price less than the market price. In addition to the order of Trader 1, the order of either Trader 2 or Trader 3 is fulfilled. The choice between Trader 2 and Trader 3 is determined randomly by a computer.

At the end of each period, the following screen is displayed with the information described below.

- the market price
- the number of bonds you purchased in the current period or the number of bonds you sold in the current period.
- the interest payments for the current period.
- the dividend payments for the current period.
- your cash holdings after the transactions, the interest payments and the dividend payments have been processed for the current period.
- the number of bonds you currently hold.
- the number of market prices that you have predicted correctly.

7.0.6. *Earning interest from cash holdings.* In each game, there are 11 periods in which you can submit buy/sell orders and trade with other traders in your market. You will be paid interest of 5% on the amount of cash you hold at the end of each period.

The interest earned is rounded up to a whole number. For instance, suppose that at the end of the 5th period your cash holdings are 90 Marks and you have two bonds. You are paid interest of 5% on the cash holdings of 90 Marks. The interest of 4.5 Marks is rounded to 5 Marks, so your cash holdings are 95 Marks after adding the interest.

Likewise, dividends from bonds are added to your cash holdings when you hold bonds. This is explained below.

Earning dividends from bonds You will be paid a dividend of 6 Marks per bond for the bonds you hold at the end of each period. The dividend income at the end of each period is calculated as: 6 Marks (number of bonds you hold).

In the example above, 12 Marks (= 6 Marks * 2 bonds) is added to your cash holdings of 95 Marks after adding the interests. Thus you start the 6th period with 107 Marks and two bonds. If you hold bonds at the end of the game (after the 11th period), the bonds you hold are bought for 120 Marks each after any dividend payments.

Earning a profit by predicting future prices correctly Before each period begins, you will be asked to predict the market prices in the remaining periods.

The time limit for predicting the market prices is (number of periods remaining) x (20 seconds). Before the beginning of the first period, the time limit is 220 seconds. After that, the time limit decreases by 20 seconds each period. The time limit before the 11th period is 20 seconds.

Prediction of future prices You will be asked to predict the prices for all the remaining periods before each period begins. That is:

- before the beginning of period 1, there are 11 periods remaining so you must predict 11 prices;
- before the beginning of period 2, there are 10 periods remaining so you must predict 10 prices;
- ...
- before the beginning of period 11, only one period remains so you must predict one price.

Thus, you will make a total of 66 predictions of market prices.

7.0.7. *Earning a profit by predicting future prices.* The computer keeps a record of the number of accurate predictions (that is, when the market price realized is between 90% and 110% of your predicted price for the corresponding period).

At the end of each game, you will be paid a bonus based on the number of accurate predictions according to the following formula: (your final cash balance) x 0.5% x (the number of accurate predictions). The maximum bonus percentage is 0.5% x 66 = 33%. Please be aware that your final cash balance depends on earnings made from profit margins, interest and dividends , so the size of your bonus decreases as your earnings from profit margins, interest and dividends decrease.

7.0.8. *Summary of ways to make a profit.* There are two ways of making a profit: (1) earning a profit margin, earning returns from dividends and earning interest on cash holdings, and (2) predicting market prices of bonds.

7.0.9. *Practice.* We start with a practice round so that you can familiarize yourself with the software. In particular, you will learn how to enter the required information. The first screen displayed is for predicting future prices. Press the OK button after you have entered all your price forecasts. The computer will display the order entry screen once everyone has pressed OK.

The practice round ends when everyone has entered their orders and pressed the OK button. The results of the practice round will not be displayed. Rewards do NOT take the practice round into consideration.

Before starting the game, we will announce the following: Lets start the game.

- There are six people in the market.
- All the people in the market are in this room.
- You will be given eight bonds and 800 Marks at the beginning of the game.

[The instructions for the baseline treatment finish here.]

7.0.10. *Instructions for buying operation.* During this game, the computer will buy bonds you hold. The buying operation will be conducted before the beginning of the 4th and 5th periods. During these two operations, you will be asked to submit a sell order. The computer will buy bonds in ascending order of the prices submitted by the market participants.

The target quantity for the first buying operation is eight bonds, and a total of 16 bonds are to be bought in the two operations. If fewer than eight bonds are purchased in the first buying operation, then the shortfall will be added to the second buying operation at the beginning of the 5th period. For instance, suppose that the computer purchases only six bonds in the first buying operation, which is two units short of the target. Then the target quantity for purchasing in the second buying operation is 10 bonds.

Even if fewer than 16 bonds are purchased during the two buying operations, there is no additional buying operation.

For the buying operations, the following screen is displayed.

If you want to sell the bonds you hold, you need to enter a selling price and the maximum number of bonds to sell.

How the computer purchases bonds during the buying operation The sales prices submitted by the market participants are ordered from lowest to highest. The computer will purchase bonds in ascending order of the specified price until the target quantity is reached. If there are orders with identical selling prices, and meeting all of them will exceed the target quantity, then the computer will randomly choose which orders to meet (some orders may be partially met).

After the buying operations, neither interest nor dividends will be offered for cash and bond holdings.

[The instructions for the buying operation finish here.]

7.0.11. *Instructions for selling operation.* In this game, the computer will also sell the bonds it purchases during the two buying operations. The selling operations will be conducted before the beginning of the 8th and 9th periods. You will be asked to submit a buy order prior to these selling operations. The computer will sell the bonds in descending order of the prices submitted by the market participants.

The target quantity for the first selling operation is eight bonds, and a total of 16 bonds are to be sold in the two operations. If fewer than eight bonds are sold in the first selling operation, then the unsold bonds are added to the second selling operation before the beginning of the 9th period. For instance, suppose that the computer sells only five bonds in the first selling operation, so that there are three unsold bonds. Then there are 11 bonds available for the second selling operation.

Even if fewer than 16 bonds are sold during the two selling operations, there is no additional selling operation.

...

7.0.12. *How the computer sells the bonds.* The buying prices submitted by the market participants are ordered from highest to lowest. The computer will sell the bonds in descending order of the specified prices until the target quantity is reached. If there are orders with identical buying prices, and meeting all of them will exceed the target quantity, then the computer randomly chooses which orders to meet (some orders may be partially met).

After the selling operations, neither interest nor dividends will be offered for cash and bond holdings.

[The instructions for the buying & selling operations finish here.]

7.0.13. *Additional Information for readers:* After the round with the first buying operation, the following information is provided.

- How many bonds the computer has bought during the buying operation.
- How many bonds you sold to the computer, and the payment received (Marks) from the trade.
- The number of bonds and the amount of cash (Marks) you hold after the buying operation.

The center of the screen shows the prices, in ascending order, at which the bonds were bought by the computer. After the second buying operation, the list of prices appears on the right of the screen while the center disappears.

After the round with the selling buying operation, the following information is provided.

- How many bonds the computer has sold during the selling operation.
- How many bonds you bought from the computer, and the payment made (Marks) from the trade.
- The number of bonds and the amount of cash (Marks) you hold after the selling operation.

The screen displayed after the selling operations is similar.

The center of the screen shows the prices, in descending order, at which the bonds were sold by the computer. After the second selling operation, the list of prices appears on the right of the screen while the center disappears.