Asset Price Bubbles and Crashes with Near-Zero-Intelligence Traders

Towards an Understanding of Laboratory Findings*

John Duffy
University of Pittsburgh

M. Utku Ünver
Koç University

First Draft: July 2003
Revised: February 2004

Abstract

We examine whether a simple agent–based model can generate asset price bubbles and crashes of the type observed in a series of laboratory asset market experiments beginning with the work of Smith, Suchanek and Williams (1988). We follow the methodology of Gode and Sunder (1993, 1997) and examine the outcomes that obtain when populations of zero–intelligence (ZI) budget constrained, artificial agents are placed in the various laboratory market environments that have given rise to price bubbles. We have to put more structure on the behavior of the ZI-agents in order to address features of the laboratory asset bubble environment. We show that our model of "near–zero–intelligence" traders, operating in the same double auction environments used in several different laboratory studies, generates asset price bubbles and crashes comparable to those observed in laboratory experiments and can also match other, more subtle features of the experimental data.

JEL Classification Nos. D83, D84, G12.

*We would like to thank an anonymous referee, Guillaume Frechette, David Laibson, Al Roth and participants in Harvard Experimental and Behavioral Economics Workshop for their comments, and Charles Noussair for providing his data set.

1 Introduction

Smith, Suchanek and Williams (1988) devised a laboratory double auction market that gives rise to asset price bubbles and crashes as evidenced by the behavior of inexperienced human subjects who are placed in this environment. The Smith et al. (1988) finding of price bubbles and crashes has been replicated by several other experimentalists and found to be robust to a number of modifications of the laboratory environment specifically aimed at eliminating bubbles.[1]

A difficulty with these laboratory asset markets is that they do not map easily into existing theories of price determination in markets with a single common–value good. Most of the laboratory markets that give rise to bubbles have finite horizons and are set up in such a way that rational, profit–maximizing agents would never choose to engage in any trade. By contrast, the theoretical bubble literature demonstrates how bubbles can arise in *infinite* horizon environments despite the fact that agents are (typically) homogeneous and have rational expectations.[2] These rational bubble theories are therefore of little use in understanding the laboratory asset bubble phenomenon. Surprisingly, the experimentalists themselves have little to say as to why bubbles and crashes regularly occur and appear to be puzzled by their own inability to eliminate asset bubbles in a wide range of laboratory environments. As Smith et al. (2000) notes, these "controlled laboratory markets price bubbles are something of an enigma."

Our aim in this paper is to take a further step toward understanding the laboratory asset price bubble and crash phenomenon, not by conducting additional experiments with paid human subjects, but by placing a population of artificial adaptive agents in the same laboratory environments that have given rise to price bubbles and determining how the artificial agents must behave so as to generate the asset price bubbles and related features found in the experimental data.[3] Theoretical analysis of individual behavior in the double auction market mechanism has turned out to be quite difficult due to the large multiplicity of equilibria that are possible in this environment (Friedman

[1]Smith et al.'s original (1988) bubble findings have been replicated by King et al. (1993), Smith et al. (2000) and Lei et al. (2001), among others, using similar experimental designs. In addition, these authors and others, (e.g. van Boening et al. (1993), Porter and Smith (1995), Fisher and Kelly (2000) and Noussair et al. (2001)), have also examined departures from the original Smith et al. (1988) protocol with an eye toward eliminating or attenuating asset price bubbles in experiments with inexperienced subjects.

[2]See, for example, Blanchard (1979), Tirole (1985), Diba and Grossman (1987), O'Connell and Zeldes (1988), Froot and Obstfeld (1991) and the references contained therein.

[3]This agent–based computational (ACE) approach represents a new "bottom–up" (as opposed to "top–down) methodology to understanding boundedly rational behavior in dynamic, stochastic environments with heterogeneous agents. See Tesfatsion's web site, `http://www.econ.iastate.edu/tesfatsi/ace.htm` for a thorough description of the ACE methodology, as well as bibliographies of and pointers to ACE research papers.

(1993)). Agent–based techniques provide an alternative means of gaining insight into the features of these environments that may be responsible for generating asset price bubbles and crashes in laboratory studies.[4]

At the same time, agent–based models are subject to a number of arbitrary modeling decisions. We address this difficulty in two ways. First, we attempt to use the *simplest* model of agent behavior. In particular, we follow Gode and Sunder's (1993, 1997) approach of using "budget-constrained zero–intelligence machine traders" as a means of focusing attention on the institutional features, e.g. the rules of the laboratory market environment. As we show later in the paper, we have to modify the basic zero–intelligence (ZI) approach in several respects in order to capture features of the experimental data we seek to understand. However, the modifications we make are, again, the simplest possible; indeed we explore the marginal contribution of the two modifications we have to make to the ZI methodology and show how both are critical to our findings.

Second, we impose further discipline on our modeling exercise by requiring that our artificial agent simulations match several key features of the experimental data as reported in the various laboratory studies that have given rise to bubbles. We then ask how the data from the simulations match other, more subtle features of the data. We also explore the performance of our calibrated baseline model in other experimental designs that have been proposed in an effort to eliminate bubbles without recalibrating our model to better fit data in these alternative environments. Our main finding is that asset price bubbles of the type observed in certain laboratory markets can be generated using a very simple agent–based model where trading is subject to the rules of the laboratory market and where individual bids and asks are subject to budget constraints.

Unlike Gode and Sunder (1993, 1997), we are not interested in the effect of various market procedures on *allocative efficiency*; instead our aim is to determine whether our calibrated agent-based model can deliver, both qualitatively and quantitatively, results that are similar to those found in a variety of different laboratory studies. Thus we examine the performance of our base-line, calibrated model in alternative market environments that experimentalists have proposed and examined in an effort to eliminate asset pricing bubbles. We find that our model continues to track experimental results well in these other environments even though it is not calibrated to match any of the features of these other environments. Finally, we redo our calibration exercise for a different version of the laboratory bubble environment proposed by Lei et al. (2001) where agents are restricted to be either buyers or sellers. For this environment, we eliminate the weak foresight aspect of our model, whereby the probability of being a buyer decreases over time. A calibration

[4]Researchers have only recently begun to use agent–based modeling techniques to understand and predict behavior in experimental studies with real human subjects. See, e.g., Duffy (2004) for a survey of this literature.

of this simpler model continues to perform well in tracking the features of the data observed in the Lei et al. experiment. We conclude that agent–based modeling approaches provide one means of assessing new experimental designs or market mechanisms designed to eliminate or reduce the frequency of asset price bubbles.

2 Laboratory Market Price Bubbles

The original market environment of Smith et al. (1988) involved 9 or 12 inexperienced traders who participated in $T = 15$ or 30 trading periods of a computerized market. Each subject began the experimental session with an endowment of x units of cash and y units of the single asset. In each trading period, subjects could submit both bid and ask prices for a unit of the asset (only one unit could be traded at a time) subject to budget/endowment constraints. Bid or ask prices that did not improve on pre–existing bid or ask prices were ranked relative to the current best bid and ask prices and placed in an order book queue. Agents were free to buy or sell a unit at a time at the current best bid or ask prices which were the only prices shown on each subjects' trading screens. When a unit was sold, the inventory and cash balances of the two traders were adjusted accordingly, and the transaction price was revealed to all traders. The next best bid and ask prices from the queue became the new best available bid and ask prices on all traders' screens. Trading was halted at the end of each 240 second trading period.

Following the completion of each trading period, subjects earned a dividend payment per unit of the asset that they owned at the end of the period. The dividend amount was a random variable consisting of a uniform draw from a distribution with support: $\{d_1, d_2, d_3, d_4\}$ where $0 \leq d_1 < d_2 < d_3 < d_4$, so the expected dividend payment was $\bar{d} = \frac{1}{4}\sum_{i=1}^{4} d_i$. After dividends were paid out, provided the last trading period T had not been reached, another 240 second trading period would commence.

At the beginning of each experimental session, i.e. before the start of trading period $t = 1$, player i's initial cash balance, x^i, and endowment of the single asset, y^i satisfy:

$$x^i + \overline{D}_1^T y^i = c$$

where c is a constant that is the same for all i. Given that all traders have the same expected value for their initial endowment at the start of the experiment, they should be indifferent between not trading in any period, or trading at the fundamental market price in period t, which earns them zero profits.[5] Since players are told the session will last T periods, the fundamental expected

[5]If there are small costs to such trades, then the no trade outcome is unique.

market price of the asset at the beginning of trading period t is:

$$\overline{D}_t^T = \bar{d}(T - t + 1) + \overline{D}_{T+1}^T,$$

and is decreasing as $t \to T$. Here, $\overline{D}_{T+1}^T$ denotes the expected default (buy–out) value of the asset after period T. If there are any trades, the traded prices should track the fundamental expected market price, $\overline{D}_t^T$ over time and should steadily decrease by an increment of $\bar{d}$ per trading period.

Following the start of the experiment, individual agents' cash endowments and inventories become endogenous, reflecting individual trading decisions. Endowments were not reinitialized at the start of each new trading period. Dividend earnings from the previous period become available for making cash purchases in the following period. All trades are allowed provided that the two parties to a trade have the necessary asset and cash endowments to complete the trade; these endowment amounts are updated in real time in the laboratory session using computerized software, and we follow the same practice in the artificial agent simulations. At the end of T trading periods, the standard practice was to pay out the period T dividend realization amounts per share and then pay out the default (buy-out) value of the asset.

The basic finding reported by Smith et al. (1988) is that with inexperienced subjects, there is a considerable volume of trade especially in the early periods of the market, and that the mean traded price exhibits a "hump-shaped" pattern. Initially the mean traded price lies below the fundamental price, but quickly soars above this fundamental price. Substantially higher than fundamental prices are sustained for some number of trading periods near the middle of the session despite the declining fundamental value of the asset. Such a sustained departure of prices from fundamentals is labeled a "bubble" by the experimenters.[6] Finally, sometime during the last few trading periods, a market crash becomes a likely event, with the mean traded price falling precipitously to values close to or even below the fundamental asset value. This hump-shaped path for mean traded prices is the most well–known signature of the laboratory asset market bubble. However, there are other, more subtle relationships between prices, bids and offers and trading volume that also characterize these laboratory asset market bubbles. We will address these relationships later in the paper.

Asset bubbles in laboratory markets tend to disappear when agents have garnered experience with the market environment as Smith et al. (1988) also demonstrate. Consequently, experimental research on the source of laboratory asset price bubbles has naturally focused on the behavior of *inexperienced* subjects.[7] We emphasize at the outset that the model we set forth below is *not*

[6]Somewhat surprisingly, the experimenters have not adopted a precise definition as to what constitutes a "bubble" or a "crash." Instead, there is a certain "call-them-as-they-see-them" approach to characterizing whether bubbles or crashes have occurred. An exception is Noussair et al. (2001) as discussed later in the paper.

[7]These inexperienced subjects might be likened to the noise–traders found in the finance literature on asset price

intended as a model of behavior by experienced market players; rather it is a model of play by *inexperienced* market participants only. Furthermore, our use of "near zero–intelligent" traders is not intended as a commentary on the rationality of the human subjects; rather it should be interpreted as providing a lower bound on possible behavior in market game environments with inexperienced players.[8]

In subsequent experimental research, Smith et al. (2000) and Lei et al. (2001) have come closest to eliminating bubbles with *inexperienced* subjects, though they are not completely successful in this effort. Smith et al. (2000) examine environments that differ in the frequency with which dividends are paid out, while Lei et al. (2001) consider environments where traders are prevented from making speculative trades and where another market for a non–asset, consumption good is also in operation. Smith et al. (2000) find that bubbles are most unlikely (but can still occur) when the asset only pays a single random dividend following the last trading period, T, while Lei et al. (2001) find that bubbles are most unlikely (but can still occur) when traders are prevented from speculating *and* another, non–asset market is in operation, so that subjects are less predisposed to be actively engaged in the asset market. Lei et al. (2001) motivate their two market treatment with the following statement:

> "Because the data are difficult to reconcile with the theory, it is natural to conjecture that aspects of the methodology of this type of asset market experiment are the sources of the errors in decision making (p. 846)."

After finding that their various manipulations of the market environment do not always prevent bubbles from occurring, they rule out one possible explanation:

> "We do not interpret our data as suggesting that the conscious pursuit of capital gains does not occur in experiments of this type (p. 857)."

Our focus in this paper is also on whether features of the design of this type of asset market experiment are responsible for the seemingly anomalous price bubble findings. However, rather than supposing that subjects are rational or even conscious profit maximizers, we follow the approach of Gode and Sunder (1993, 1997) and assume just the opposite: that traders are unconscious,

volatility, see, e.g. DeLong et al. (1990). Populations with a mix of both inexperienced and experienced traders have yet to be considered experimentally.

[8]Indeed, in a very stimulating paper, Brewer et al. (2002) show that in double auction, buyer-seller markets where demand and supplies are continously refreshed, the behavior of inexperienced human subjects is quite different from that of zero-intelligence traders.

Players	Endowment (Cash; Quantity)	Number of Players
Class I	(\$2.25; 3)	3
Class II	(\$5.85; 2)	3
Class III	(\$9.45; 1)	3
Dividends	$d \in \{\$0, \$0.04, \$0.14, \$0.20\}^a$	$\overline{d} = 0.12$
Initial Value of a Share	$\overline{D}_1^{T\,b} = \$3.60$	
Buy-out Value of a Share	$D_{T+1}^{T} = \$1.80$	

[a] Each dividend outcome occurs with probability $\frac{1}{4}$.

[b]Each period's expected fundamental value is denoted by $\overline{D}_t^T$ for $t = 1, ..., T+1$. These values were calculated and displayed on the screen in each trading period in the human subject experiments.

Table 1: Smith et al. (1988) Experimental Design 2

irrational, near–zero intelligence agents who make random bids and offers subject to certain market constraints. This approach allows us to disentangle the potential sources of asset price bubbles in these laboratory environments by focusing attention on features of the experimental design rather than the (possibly strategic) behavior of subjects. It also allows provides a model for considering alternative experimental designs that might reduce or eliminate the incidence of price bubbles.

3 An Agent-Based Model

We begin with a description of artificial agent behavior in our baseline asset market environment. This environment corresponds to one of the experimental designs – Design #2 – examined in Smith et al. (1988). In this environment there are $N = 9$ agents who participate in $T = 15$ trading rounds. There is a single asset that pays a random dividend at the end of each period. The number of agents, their initial endowments, and the details of the dividend process are given in Table 1.

In each trading period, agents can either be buyers or sellers, and so we refer to them as "traders." Each trader can submit buy (bid) or sell (ask) orders within each trading period. Algorithmically, we choose a random sequence in which each of the 9 traders submits a single bid or ask. A trading period t consists of a total of S such sequences, and S is chosen later in our calibration so as to match the volume of trade. Let $s = 1, 2, ..., S$ index the random sequences within period t.

When it is trader i's turn in sequence s of trading period t, we first randomly determine whether he will be a buyer or seller. The probability that agent i is a buyer is π_t and the probability he is a seller is $1 - \pi_t$. The initial probability of being a buyer, π_0, is 0.5 and decreases over time.

Specifically the probability of being a buyer in trading period t is given by

$$\pi_t = \max\{0.5 - \varphi t, 0\}$$

where $\varphi \in \left[0, \frac{0.5}{T}\right)$ is a parameter.[9] This assumption endows the artificial agents with a certain *foresight* that there is a finite end to the market: since the fundamental value of the asset is declining over time and agents will have to sell any unit they have at the end of the session at the default value, their tendency to buy decreases over time. We refer to this as the "weak foresight" assumption. Its primary role is to generate a reduction in transaction volume over time, consistent with the experimental data. However, since it also means that asks will become more frequent than bids, it can also lead to a reduction in transaction prices as well.

If trader i is a seller in sequence *s and has a unit available for sale*, then trader i submits an ask price. If trader i is a buyer in sequence *s and has sufficient cash balances*, then trader i submits a bid price. We refer to the these trading restrictions as "loose" budget constraints, because, by contrast with Gode and Sunder's budget-constrained ZI model, bid or asks in our environment are not made contingent on the intrinsic, fundamental value of the asset. Gode and Sunder (1993) adopt a stricter, binding, *no loss constraint*, wherein an agent buys (sells) an item only if his private value (cost) is higher (lower) than his bid (ask). In our environment where individuals can be both buyers and sellers, this type of constraint would force agents to buy or sell at the intrinsic value, and consequently price bubbles would never be observed. Hence, our adoption of the loose budget constraint convention.

A second departure from Gode and Sunder's ZI approach, is our assumption that bid and ask prices are not completely random, but depend in part on the mean transaction price from the previous trading period, which we denote by $\overline{p}_{t-1}$. Specifically each trader's bid or ask price is a fixed convex combination of the mean traded price from the previous period and a random price. This assumption captures the behavioral notion that *anchoring* effects are important - here the relevant anchor for bids and asks is the previous period's mean traded price.[10] This latter departure from Gode and Sunder's ZI approach leads us to qualify our agents as "near–zero intelligence" traders, as our agents, by contrast with Gode and Sunder's agents, can recall the mean transaction price from the previous period.

[9] We choose the interval for φ as $\left[0, \frac{0.5}{T}\right)$ to ensure $\pi_t \in (0, 0.5]$.

[10] Anchoring effects are said to be operant if players' numeric estimates are related to some previously considered benchmark, often the initial numeric value. For instance, Genesove and Mayer (2001) report that the nominal price an owner paid for his Boston–area condominium is a critical reference point in the determination of that owner's subsequent asking price. Here, we take the benchmark to be the previous period's mean traded price.

The random component in trader i's ask or bid in sequence s of period t is denoted by $u^i_{t,s}$ and consists of a random draw from the interval $[\underline{\epsilon}_t, \overline{\epsilon}_t]$, where

$$\begin{aligned}\underline{\epsilon}_t &= 0,\\ \overline{\epsilon}_t &= \kappa \overline{D}^T_t,\end{aligned}$$

and $\kappa > 0$ is a parameter. Notice that while $\underline{\epsilon}$ is constant for all t, $\overline{\epsilon}_t$ will decrease over time since $\overline{D}^T_t$ decreases as $t \to T$.

If trader i is a seller in sequence s of period t his ask price is given by:

$$a^i_{t,s} = (1-\alpha)u^i_{t,s} + \alpha \overline{p}_{t-1}$$

where $\alpha \in (0,1)$ is a constant parameter that is the same for all traders that captures the weight given to the anchor, $\overline{p}_{t-1}$. We assume that $\overline{p}_0 = 0$, as traders in period $t = 1$ have no prior history upon which to base their pricing decisions. Seller i can submit an ask as long as he has a positive share holding at sequence s of period t, i.e. $y^i_{t,s} > 0$. Otherwise seller i does not submit an ask in sequence s of period t.

Similarly, if trader i is a buyer in sequence s of period t, his bid price is given by:

$$b^i_{t,s} = \min\left\{(1-\alpha)u^i_{t,s} + \alpha \overline{p}_{t-1}, x^i_{t,s}\right\}.$$

Trader i can submit a bid as long as he has a positive cash holdings at sequence s of period t, i.e. if $x^i_{t,s} > 0$; otherwise no bid is submitted.

An issue that immediately arises is the choice of the appropriate upper bound, κ to place on bid or ask prices. The intrinsic, fundamental value of each share in each trading period, $\overline{D}^T_t$, was displayed on computer screens in the human subject experiments and so can be presumed to have been public knowledge. Given our rule for bids and asks, prices should converge to $\frac{\kappa \overline{D}^T_t}{2}$, so one could argue that $\kappa = 2$ is an obvious choice. However, this choice would force agents to eventually buy and sell at the intrinsic value. Hence, the parameter κ was chosen to be greater than 2; the exact choice was determined on the basis of calibration to certain measures of the experimental data as discussed in section 3.1. While such an upper bound may seem arbitrary, we note that Gode and Sunder (1993, 1997) have to impose an analogous and similarly arbitrary upper bound on the ask range of sellers in the double auction environments they examine where agents are always either buyers or sellers. Our upper bound on the bid/ask range amounts to a straightforward generalization of Gode and Sunder's approach to the "trader" environment, where agents are free to be both buyers and sellers.

There are several things to note about our rules for bids and asks. Since traders can be both buyers and sellers, we have assumed they have a common belief about the range over which prices may lie, namely $\left[0, \kappa \overline{D}_t^T\right]$. The only source of agent heterogeneity is the random component to bids and asks which is necessary to generate trades; without it, given our anchoring assumption, buyers and sellers would all submit the same prices and everyone would be just indifferent between trading or not trading. Notice further that agents' pricing decisions are *non–strategic.* In particular, the speculative, or "greater fool" explanation for price bubbles – that agents buy at high prices because they believe that they can sell to another agent (greater fool) at even higher prices – is not operative here, as all players have a common view of the range of possible prices and they do not act strategically in any way. What *is* important is that previous period mean traded prices act as an anchor for current period price determination. If the initial price anchor, $\overline{p}_0 = 0$, and $\alpha > 0$ as we assume, then prices will necessarily increase over the first few trading periods. Indeed, if the probability of being a buyer or seller were fixed at 0.5 (i.e. if $\varphi = 0$), then we would find $\lim_{t \to T} \overline{p}_t \to \frac{\kappa \overline{D}_t^T}{2}$ for sufficiently large T, so prices will be greater than zero for all t. However, since the fundamental value, $\overline{D}_t^T$, decreases over time, mean traded prices can fall as well, due to the shrinking upper bound on the random component of bids and asks.

This explanation for why prices first rise and then fall holds regardless of the value of φ. As we show in section 4.2, we need $\varphi > 0$ primarily to reduce trading volume, consistent with the experimental results. With $\varphi = 0$, we would continue to get a hump–shaped path for mean traded prices but we would not get any decrease in transaction volume. Still, it would be incorrect to say that φ has no effect on traded prices. With $\varphi > 0$, there is a gradually increasing excess supply of units towards the end of the market which contributes to the reduction in mean transaction prices.

Figure 1 shows the mean transaction price from the experimental data of Smith et al. (1988) for Design #2 (labeled "Actual Price") along with several mean transaction price paths from simulations of our agent–based model. The mean price path from simulations of our optimally calibrated baseline model is labeled "Sim Price – Optimal Fit." The details of the optimization procedure we employed are discussed below. The path of mean transaction prices from our simulated model exhibits the same hump–shaped path as found in the experimental data.

In Figure 1, we contrast the path of prices from the "optimal fit" version of our model with that from two different variations on our model. The first, labeled "Sim Price for Phi = 0," is a simulation of our model where φ is set to zero and all other parameters are kept at their optimal values. Consistent with our earlier discussion, in the absence of any weak foresight (i.e. when $\varphi = 0$), the mean traded price path is indeed converging to $\frac{\kappa \overline{D}_t^T}{2}$, which is also plotted in Figure 1. When $\varphi > 0$, there are fewer buyers and more sellers as the asset market proceeds. This excess

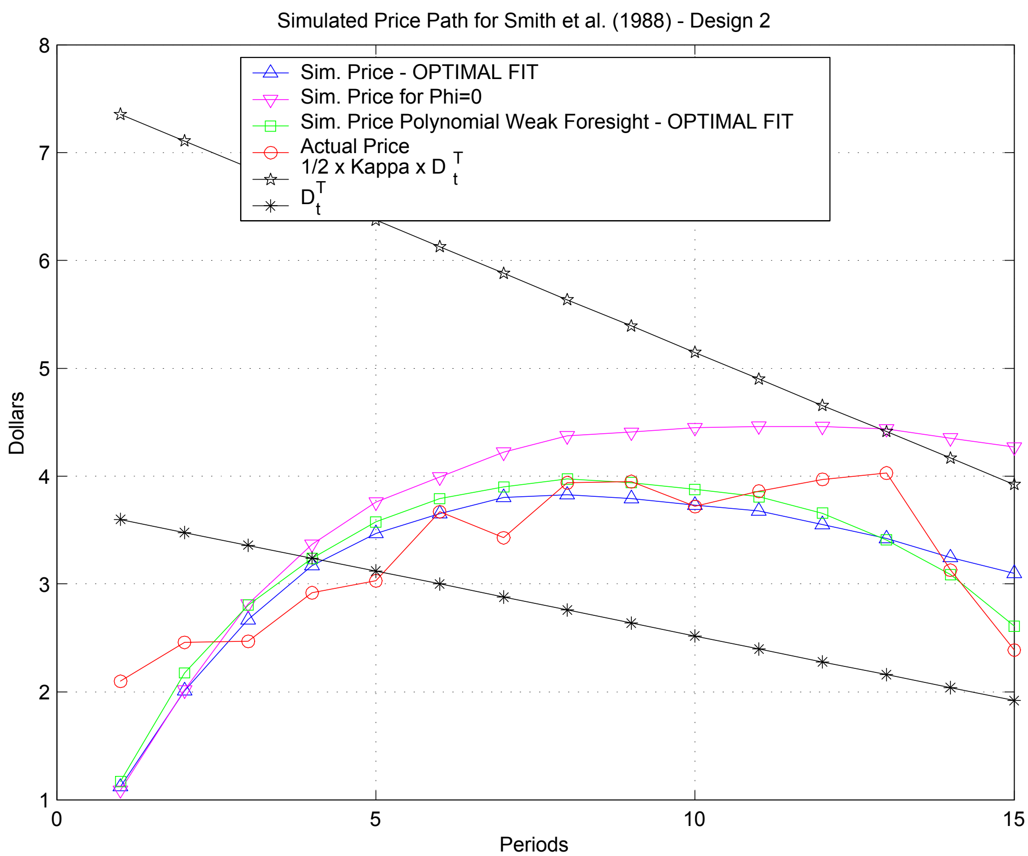

Figure 1: The mean transaction price path in the simulations and in the actual data.

supply causes a further decrease in traded prices.

The second variation on our baseline model is aimed at delivering a larger fall-off or "crash" in mean traded prices. For this variation of the model, the probability of being a buyer in period t is now:

$$\pi_t = \max\{0.5 - \varphi t^{\gamma}, 0\}$$

where $\gamma > 1$ is an additional behavioral parameter. The interpretation of this modification is that there is a polynomially increasing desire by agents to sell units of the asset as the known, finite horizon approaches. We discovered that many values of $\gamma > 1$ yield a higher percentage of crashes than in the baseline $\gamma = 1$ model, though these alternative models yield only a slightly better fit to the experimental data. In Figure 1, we present the mean traded price path from a simulation with $\gamma = 7$, where other model parameters optimally fit for this level of γ - the price path labeled "Sim Price - Polynomial Weak Foresight - Optimal Fit." We chose $\gamma = 7$ because it created the best fit (in terms of our sum of squared deviations (SS) objective function discussed in section 3.1) among γ values in the set $\{1, 3, 5, 7, 9\}$. However, the improvement in terms of fit to the experimental data of the $\gamma > 1$ version of our model was minimal. In the interest of keeping the number of behavioral parameters to a minimum, we have chosen to consider the simpler, $\gamma = 1$ baseline model in the remainder of the paper.

As in the laboratory studies and in actual markets, we use standard bid and ask improvement rules which require that buyers improve on (i.e. raise) the current best bid price and sellers improve on (i.e. lower) the current best ask price. If a bid price is submitted that is greater than or equal to the current best ask price, the convention adopted here is the same one used in the laboratory experiments: the unit is sold at the current best ask price. Similarly, if an ask price is submitted that is less than or equal to the current best bid price, the unit is sold at that current best bid price, again in line with the experimental practice. Once a unit is traded, we follow one of two conventions for updating the best bid and ask prices. In the first, "continuous order book" convention, the one used by Smith et al. (1988, 2000), the next best bid or ask price in the electronic order book becomes the current best available bid or ask price. In the second "cleared order book" convention, the one used by Lei et al. (2001) and Noussair et al. (2001), the order book is completely cleared following each trade, so the first new bids and asks submitted following a trade become the current best available. For the baseline simulations, we use the continuous order book convention since this is the one used by Smith et al. (1988). However, we find that our results are not sensitive to the type of order book convention.[11] Following the end of each *trading period*, the order book

[11] For both conventions, we use the following rule: If a player has an outstanding limit order to buy (or sell) and it is his turn again in the trading period to submit an order, we permit this player only to submit a bid (or ask).

is completely cleared, a convention that is adopted in all of the laboratory studies. Dividends are then paid out, and each agent's cash balances, x^i, are adjusted accordingly.

Of course, during trading period t, any trades that agents make happen immediately and result in an immediate (real–time) adjustment to their cash balances, $x^i_{t,s}$ and asset endowments, $y^i_{t,s}$. Such trades may also affect the bid ranges over which traders can submit bids, or whether they may submit asks (e.g. if they have no units left to sell). Specifically, an agent who has bought a unit has reduced cash holdings and is therefore prevented from submitting bids that would exceed current available cash holdings. In addition, an agent who has sold a unit, has one less unit to sell; if the unit most recently sold was that agent's last unit, then that agent cannot submit any further asks. These restrictions simply reflect the enforcement of budget constraints and are consistent with the rules of the laboratory studies.

We note that a trading period t ends after S random sequences have played out. We then calculate the **mean traded price** for the period, $\overline{p}_t$. The mean traded price in period t of session (or simulation run) k, $\overline{p}^k_t$, is constructed as follows. Let vol^k_t denote the **volume of transactions** measured as the number of shares traded in period t of session k. Define $\overline{p}^{b-a,k}_t$ as the **mean bid–ask spread price** in period t of session k and define p^k_{th} as the sale price of hth unit in period t of session k. The mean transaction price at the end of period t of session k is defined by:

$$\overline{p}^k_t = \begin{cases} \frac{1}{vol^k_t}\sum\limits_{h=1}^{vol^k_t} p^k_{th} & \text{if } vol^k_t > 0 \\ \overline{p}^{b-a,k}_t & \text{if } vol^k_t = 0 \end{cases}$$

The mean transaction price, $\overline{p}^k_t$, is the quantity we use to measure the market price of a share.

3.1 Model Calibration

We used a simulated method of moments estimation procedure to calibrate the parameters of our model φ, κ, S and α. Specifically, we adopted the following two step method of moments procedure to optimally determine these parameter values.

1. In Step 1, we performed a univariate optimization over κ in the interval $[0.5, 8]$ for given φ, S, and α, so as to minimize the weighted sum of squared deviations of the simulated mean transaction price path from the actual mean price path in the experimental data plus the weighted sum of the squared deviations of the simulated mean transaction volume path from

Therefore a player cannot have an ask and a bid in the order book simultaneously (and cannot buy from himself). In the continuous order book convention, we permit a player to have only one outstanding limit order at any moment. He can make a better bid or ask but then his older bid or ask is erased from the book.

the actual mean volume path in the experimental data (denoted with an E superscript). In particular, we sought to minimize the sum of squared deviations function defined by

$$SS(\kappa, \alpha, \varphi, S) = \sum_{t=1}^{T} \left(\frac{\overline{p}_t(\kappa, \alpha, \varphi, S) - \overline{p}_t^E}{\overline{D}_1^T} \right)^2 + \sum_{t=1}^{T} \left(\frac{\overline{vol}_t(\kappa, \alpha, \varphi, S) - \overline{vol}_t^E}{TSU} \right)^2,$$

where TSU denotes the **total stock of units** endowed to all agents in an experimental market, ($TSU = \sum_i y^i$). This function gives equal weight to fitting the mean transaction price and the mean trading volume that is reported in the experimental data. The mean transaction price $\overline{p}_t$ and the mean transaction volume $\overline{vol}_t$ in period t are defined by

$$\overline{p}_t = \frac{1}{K} \sum_{k=1}^{K} \overline{p}_t^k \qquad \text{and} \qquad \overline{vol}_t = \frac{1}{K} \sum_{k=1}^{K} \overline{vol}_t^k,$$

where K is the total number of simulated sessions.

The variables $\overline{p}_t^E$ and $\overline{vol}_t^E$ denote the corresponding mean transaction price and the mean volume in trading period t in the experimental data (over all sessions). This procedure is nested in a grid for $\varphi \in \{0, \frac{1}{180}, \frac{2}{180}, ..., \frac{5}{180}\}$, $S \in \{1, 2, ..., 10\}$, and $\alpha \in \{0, 0.05, 0.1, ..., 1\}$.

2. In Step 2, we use the sets of vectors $(\kappa, \alpha, \varphi, S)$ found in Step 1 as our starting points for a 3-dimensional optimization procedure. We search for optimal $(\kappa^*, \alpha^*, \varphi^*)$ values for each integer S selected. We choose these initial points according to how small the sum of squares function was for these vectors in Step 1. We use a simplex algorithm developed for MATLAB to calculate the local optima around these starting points. Among all the locally optimal values found, we pick the vector that implements the global minimum.

In the current problem, we are able to pin down the optimal values as:

$$\kappa^* = 4.0846, \alpha^* = 0.8480, \varphi^* = 0.01674, \text{ and } S^* = 5.$$

The basic purpose of Step 1 is to explore the surface of the probabilistic sum of squares function. Although we use an unconstrained minimization algorithm in Step 2, we do not encounter any locally optimal points outside the range of the parameters.

3.2 Statistics

In this section, we define some statistics that we will use in exploring the simulation results. An important signature of an asset price bubble is persistently high prices – prices in excess of what would be predicted by market fundamentals. The **price amplitude** is a commonly used measure of the existence of bubbles. It is defined as:

$$PA_k = \max_{t\in\{1,..,T\}}\{\frac{\overline{p}_t^k - \overline{D}_t^T}{\overline{D}_t^T}\} - \min_{t\in\{1,..,T\}}\{\frac{\overline{p}_t^k - \overline{D}_t^T}{\overline{D}_t^T}\}$$

for session k. An alternative measure of a bubble is the **absolute intrinsic value deviation** which is defined as

$$AIVD_k = \sum_{t=1}^{T}\sum_{h=1}^{vol_t^k}\frac{\left|p_{th}^k - \overline{D}_t^T\right|}{TSU}$$

for session k. Several authors also use an alternative measure to the absolute intrinsic value deviation. This measure is called the **intrinsic value deviation** which is defined for session k as:

$$IVD_k=\sum_{t=1}^{T}\sum_{h=1}^{vol_t^k}\frac{p_{th}^k - \overline{D}_t^T}{TSU}.$$

High transaction volume is another feature of bubbly asset markets. Following the literature, we adopt a statistic known as the **turnover rate** which is defined as the percentage of the total stock of units that is sold in the entire market as a measure of transaction volume.

We also report statistics on transaction and price dynamics using our simulated data. The transaction dynamics are captured by the mean volume of trade in each trading period. The price dynamics are reflected in the **normalized mean price deviation**. The normalized mean price deviation in period t for session k is defined as

$$NPD_t^k = \frac{\overline{p}_t^k - \overline{D}_t^T}{\overline{D}_1^T}.$$

We plot the average normalized mean price deviation and the average volume paths versus the trading periods for our simulations.

While we report all of these statistics for our simulations, not all of these statistics are reported in the various experimental studies.[12]

4 Simulation Findings

4.1 Baseline Model

As mentioned above, our baseline model is that of Smith et al. (1988), Design #2, as described in Table 1. We have simulated our artificial agent model with 9 traders in this environment for

[12] Different authors have used different bubble, price and volume statistics to present their results. Since we want to compare our simulation results with existing experimental results, we calculate all statistics that have been reported in the laboratory asset bubble literature for our simulated data.

a total of $K = 100$ independent market sessions, each consisting of $T = 15$ periods, using the optimal parameter vector we obtained from our simulated method of moments procedure. The mean transaction price path from this simulation exercise (averaged over all 100 sessions) and for the actual experimental data were presented earlier, in Figure 1. It should be no surprise that the simulated mean price path tracks the actual mean price path rather well, as minimization of the squared deviation between the simulated and the actual price path was one component of the objective function for our simulated method of moments procedure.

In Figure 2 we present a plot of the normalized mean price deviation, NPD, and transaction volume over time from our simulation and we also show the corresponding series from Smith et al.'s (1988) experimental data. The normalized deviation for the simulated market starts out 69%

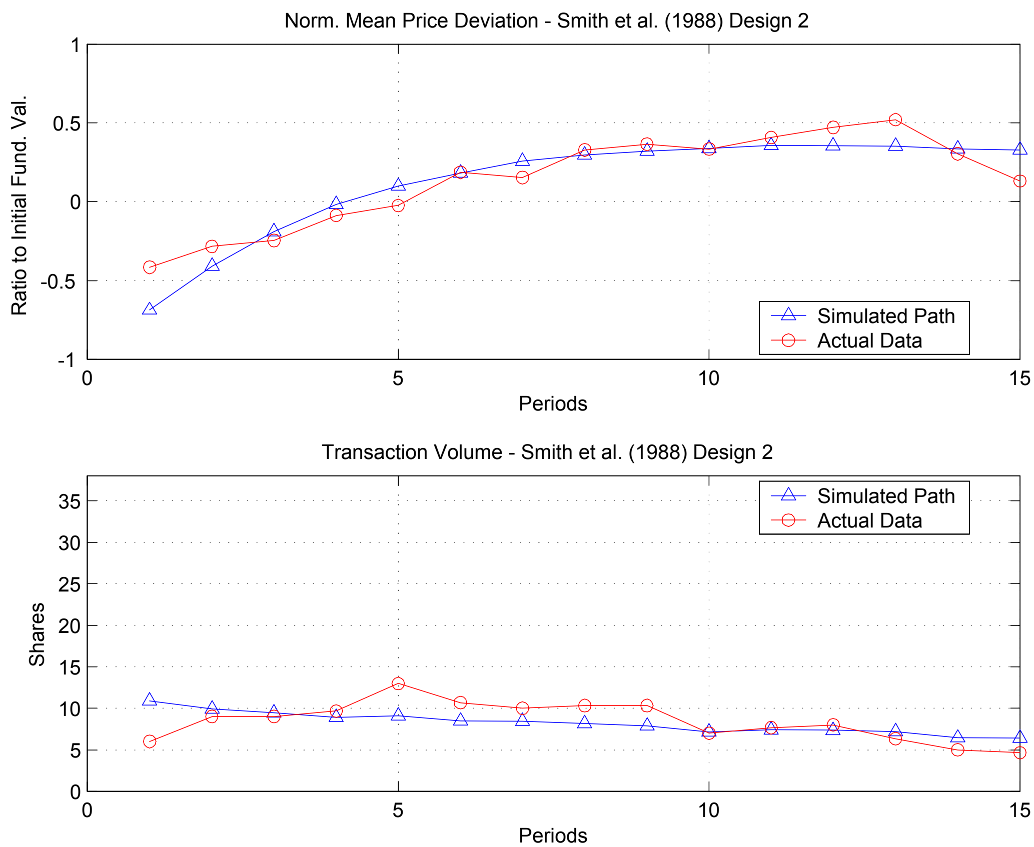

Figure 2: Simulation and Experiment Results.

below the intrinsic value in period 1 and increases up to 36% of the intrinsic value before fall off a little in the last few trading periods. Transaction volume starts out averaging 10.90 units in period

1 and monotonically decreases to an average of 6.41 units by the final period 15. These paths compare favorably with the experimental data, although again, this should not be too surprising as our calibration was chosen so as to minimize deviations from actual price and volume paths.

We next turn to a comparison of some statistics calculated using our simulated data with comparable statistics calculated using Smith et al.'s (1988) experimental data, that we did not attempt to explicitly match in our calibration exercise. Table 2 reports these statistics for both the simulation data and Smith et al.'s (1988) experimental data (if available).

Statistics	**Simulations**	**Experiments**
turnover %	685%	703%
$\overline{PA}$	1.35	1.38
$\overline{AIVD}$	7.94	5.68
$\overline{IVD}$	2.13	N/A
$\overline{p}_T - \overline{p}_{T-1}$	< 0	< 0

Table 2: Statistics in the simulations and the experiments.

In the simulated data, the turnover in shares and price amplitude (PA) statistics are a close match to the corresponding statistics in the experimental data. The absolute intrinsic value deviation (AIVD) calculated using the experimental data is less closely matched by the simulated data statistic. We note, however, that the experimental AIVD statistic reported in Table 2 is for all designs, not just Design 2, of Smith et al. (1988).[13]

Smith et al. report rising traded prices in all 3 sessions with inexperienced subjects reported for Design 2. They further report that mean traded prices fall in two of the three sessions towards the end of the market. We also observe a similar hump-shaped pattern in mean traded prices in all of our simulated markets.

Other authors have reported the experimental finding that many transactions are recorded at prices above the maximum fundamental value of a share or below the minimum fundamental value of a share. They have pointed to this finding as a sign of irrational behavior on the part of agents. Indeed our simulation results also capture this feature of the experimental data. As Table 3 reveals, 34.42% of the total turnover is realized at prices higher than the maximum value of the asset (calculated using the highest possible dividend realization in every period) and 10.91% of the total turnover is realized at prices lower than the minimum value of the asset (using the lowest possible dividend realization in every period).

[13]We found the absolute intrinsic value deviation for the Smith et al. (1988) data reported in the Noussair et al. (2001) study. We calculated the mean price amplitude for the three Smith et al. (1988) Design 2 sessions by ourselves using the data reported in their paper.

Turnover Composition	Simulations
under minimum fundamental value	10.91%
between min. and max. fund. val.	54.67%
above maximum fundamental value	34.42%

These data are not reported by Smith et al. (1988) for the experiments.

Table 3: Turnover in the Simulation Data

A real test for a simulation model such as ours is whether it captures more detailed features of the experimental data. Repeatedly in laboratory bubble experiments, authors have found that there is a significantly positive relationship between changes in the mean traded price and the difference in the number of bids and asks recorded in the previous period. We next look at this relationship using our simulated data.

Denote the number of bids in session k, period t, by B_t^k and the number of asks (or offers) in session k, period t, by O_t^k. Consider the following regression model:

$$\begin{aligned} \overline{p}_t^k - \overline{p}_{t-1}^k &= a + b(B_{t-1}^k - O_{t-1}^k) + \varepsilon_{kt} \quad (1) \\ & \varepsilon_{kt} \overset{iid}{\backsim} N(0, \sigma^2) \\ k &= 1, ..., N \text{ sessions} \\ t &= 2, ..., T \text{ periods} \end{aligned}$$

In the fully rational setting with risk neutral players, the estimate of the coefficient a should be equal to the negative of expected dividend payment, which is -0.12 in design #2 of Smith et al. (1988), and the estimate of the coefficient b should be equal to zero.

We estimate equation 1 using our entire simulation data set (100 simulations each consisting of 14 periods for t =2,...,15). Coefficient estimates are given in Table 4. Using the simulated data, the regression model cannot be rejected at the 1% level ($F = 131.41$, 1400 observations).

Regression	Sessions	$\widehat{a}$	$t-$stat	$p-$value	$\widehat{b}$	$t-$stat	$p-$value (one-sided)
Simulations	Cumulative	0.19	16.21	<0.01	0.033	7.92	<0.01
Experiments	#10 / growing price	0.2	5.4	< 0.05	$-$ 0.01	$-$ 2.1	< 0.05
	#16 / bubble-crash	0.058	0.76	> 0.05	0.038	2.2	< 0.05
	#18 / bubble-crash	-1.6	-0.17	> 0.05	0.029	1.8	< 0.05

Table 4: Coefficient Estimates of the Simulation and Experiment Data

We observe that the estimate of $\widehat{b}$ is significantly positive. Furthermore, the artificial agents do not discount the price of the asset in a rational manner, i.e. the estimated coefficient $\widehat{a}$ is also significantly positive, in contrast to the rational prediction that $a = -0.12$. Smith et al. (1988) run similar regressions separately for each session. These regression results are reproduced in Table 4 for comparison purposes. As this table reveals, consistent with our findings, Smith et al. find a significantly positive estimate for $\widehat{b}$ in 2 out of 3 sessions and a significantly positive estimate for $\widehat{a}$ in 1 out of 3 sessions. Moreover, our estimates of $\widehat{b}$ and $\widehat{a}$ both lie within the range of estimates reported by Smith et al. (1988). We conclude that experimental subject and simulated agent behavior is not dissimilar. In particular, when bids exceed (fall below) offers, subsequent period traded prices change in a predictable direction.

4.2 Comparative Statics

We performed some additional simulations using the Smith et al. (1988) Design 2, but with extreme values of α or φ in place of the optimal choices for these parameter values. The purpose of this exercise is to better comprehend the role played by these two key "behavioral" parameters in the determination of agent behavior. In particular we consider how our model fares under the alternative parameter vectors $(\kappa^*, \alpha = 0.95, \varphi^*, S^*)$, $(\kappa^*, \alpha = 0, \varphi^*, S^*)$, and $(\kappa^*, \alpha^*, \varphi = 0, S^*)$. The results of these simulations are compared with the paths obtained using the optimal parameter vector $(\kappa^*, \alpha^*, \varphi^*, S^*) = (4.0846, 0.8480, 0.01674, 5)$ for prices and volume in Figure 3. The left panel of this figure plots the normalized mean price deviation path from the simulations while the right panel plots the mean transaction volume path from the simulations. The optimal paths are shown in the first row, the laboratory data are shown in the last (fifth) row, and the other three rows present results from the various nonoptimal choices for α or φ.

Consider first the two extreme values for α. Setting $\alpha = 0$ (row 2 of Figure 3) eliminates the anchoring effect, so there is no reference point for the simulated agents' bids and asks. The simulated agent bids and asks are random numbers in $\left[0, \kappa^* \overline{D}_t^T\right]$. Since $\varphi^* > 0$, the mean price does not remain constant at $\frac{\kappa^* \overline{D}_t^T}{2}$ but falls below this value over time. Transaction volume declines slightly over time as well for the same reason. At the other extreme, when $\alpha = 0.95$ (row 3) there is a heavy anchor at the previous period's mean transaction price. Since the initial price, $\overline{p}_0 = 0$, mean traded prices rise only very slowly above 0. The mean traded price eventually rises above the fundamental value, but this rise does not coincide with the more rapid price rise that occurs earlier in an experimental session. Furthermore, since the rise in prices takes longer, a fall-off in prices is not observed within the same time-frame (15 periods) of the experimental markets. Finally, consider the case where $\varphi = 0$, (row 4) so there is no foresight of the approaching finite horizon. Consequently there are always, on average, equal numbers of buyers and sellers in this environment. Prices increase higher than in the optimal case, up to $\frac{\kappa^* \overline{D}_t^T}{2}$. Transaction volume

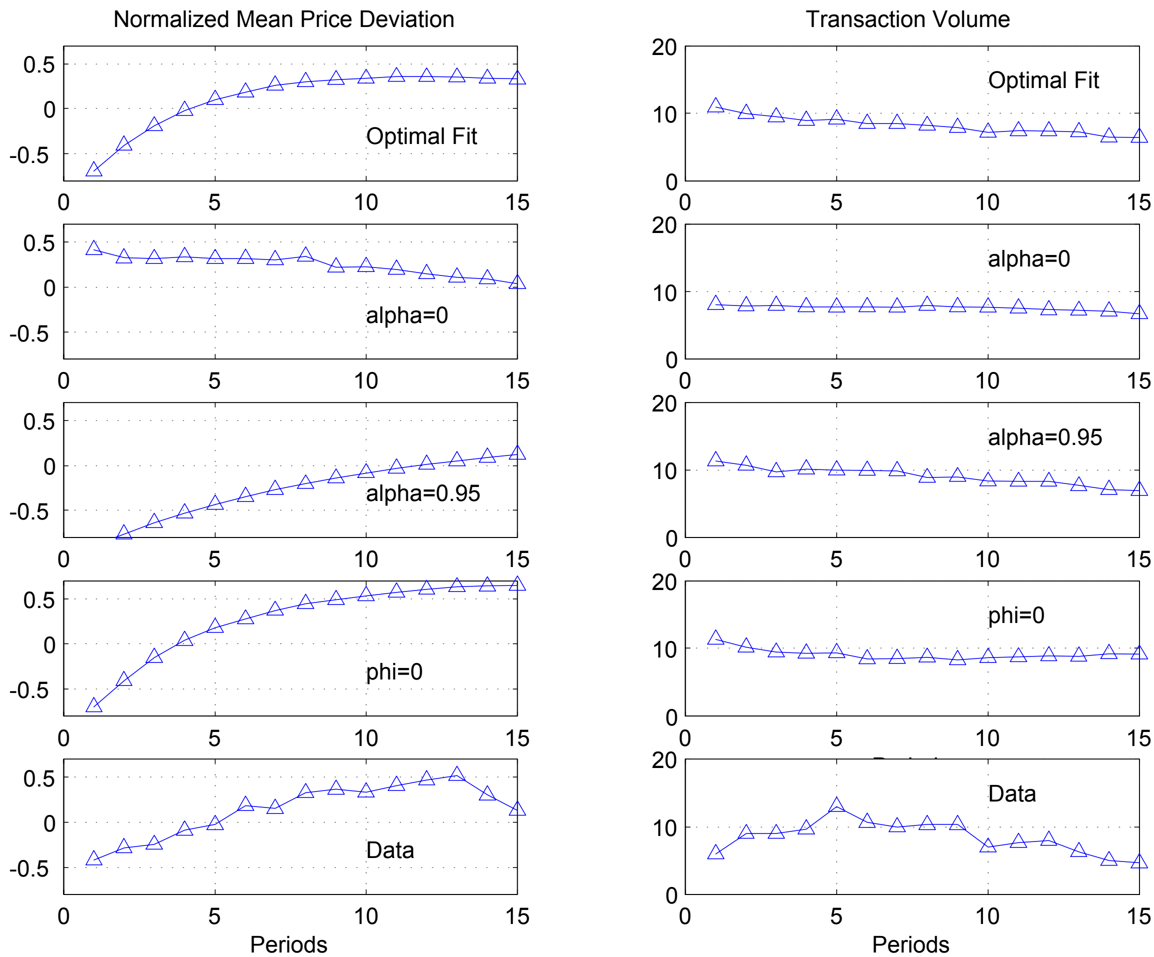

Figure 3: Comparative Statics on α and φ in the Simulations.

exhibits no downward trend.

Summarizing these comparative static exercises, it seems that careful choices for our two main "behavioral" parameters, α and φ, are important for our results. In particular, setting either parameter equal to zero worsens the performance of our model in terms of replicating the important features of the experimental data. In the following sections, we use our calibrated model to predict behavior in other asset market experiments that have been designed in an effort to prevent bubbles from occurring.

4.3 Asset Markets without Dividend Payments after Each Period

A recent paper by Smith et al. (2000) comes closest to eliminating laboratory bubbles in environments where agents can be both buyers or sellers. Their A1 design, involves a $T = 15$ period market with no dividend payments. The only money paid to subjects for asset holdings is the default value of the asset at the end of the market, following the end of period 15. Their hypothesis is that dividend payments at the end of each trading period, as in Smith et al. (1988), focuses traders' attention too myopically on the near term; by concentrating the dividend payoff into a single end-of-market payment the hope was that agents would be more far-sighted (and homogeneous in their expectations) and, as a consequence, bubbles would become less likely. There is some support for this hypothesis in their experimental data as we discuss below. Still, they report some market sessions where price bubbles continue to arise.

The design specifications of Smith et al.'s (2000) A1-1 to A1-6 sessions are given in Table 5.

Players	Endowment (Cash;Quantity)	Number of Players
Class I	$(\$3.5; 4)$	3
Class II	$(\$9.9; 2)$	3
Class III	$(\$13.1; 1)$	2
Class IV	$(\$16.3; 0)$	2
Dividends	$d = \$0$	
Intrinsic Value of a Share	$\overline{D}_1^T = \$2.4$	
Buy-out Value of a Share	$D_{T+1}^T \in \{\$1.8, \$2.4, \$3\}^a$	$\overline{D}_{T+1}^T = \$2.4$

[a] Buy-out value \$1.8 will occur with $p = \frac{1}{4}$, \$2.4 will occur with $p = \frac{1}{2}$, \$3.0 will occur with $p = \frac{1}{4}$.

Table 5: Smith et al. (2000) Experimental Design A1 Sessions 1-6

In applying our near-zero-intelligence agent model to this environment, we do not re-calibrate the model parameters to best fit the traded price and volume paths in the experimental data. Instead we use the parameter values for our model that were optimal for the Smith et al. (1988) experiment. Our aim is to use our calibrated baseline model to *predict* behavior in the Smith et al. (2000) experiment and then compare it with the actual data. This provides a more rigorous test of our artificial agent model than if we were to re-calibrate it to match features of the data reported by Smith et al. (2000). Using the optimal parameters for the Smith et al. (1988) design, but the experimental design given in Table 5 for Smith et al. (2000), we conducted a simulation exercise similar to the one previously discussed: $K = 100$ independent market sessions, each consisting of T=15 periods with 10 traders of the various classes given in Table 5.

In Table 6, we display some statistics from our simulation of the Smith et al. (2000) environment and compare these with the corresponding statistics from the experimental data. While our fit is

Statistics	**Simulations**	**Experiments**
turnover %	741%	559%
$\overline{PA}$	1.17	0.78
$\overline{AIVD}$	6.29	N/A
$\overline{IVD}$	2.42	3.96
$\overline{p}_T - \overline{p}_{T-1}$	< 0	< 0

Table 6: Simulation and Experiment Statistics.

not exact, we do observe comparable values for the turnover percentage, the price amplitude and the intrinsic value deviations in both the simulated and the experimental data. Note in particular that mean price amplitude as reported in Table 6 falls relative to the same measure reported for our baseline simulation calibrated to match features of the data reported in Smith et al. (1988): compare the mean price amplitude reported Table 2 with that in Table 6. A similar drop in price amplitude is found in the Smith et al. (2000) experimental data relative to the Smith et al. (1988) experimental data, (again, compare Tables 2 and 6) which supports the claim that price bubbles are less likely in the Smith et al. (2000) environment.

Turnover Composition	**Simulations**
under minimum fundamental value	13.16%
between min. and max. fund. val.	52.34%
above maximum fundamental value	34.50%

These statistics are not reported by Smith et al. (2000) for the experimental data.

Table 7: Turnover in the Simulation Data.

In Table 7 we decompose the turnover in units. We see that 34.50% of all turnover in units is realized at prices higher than the maximum fundamental value while 13.16% of all turnover is realized at prices lower than the minimum fundamental value of the asset in the simulation. This finding simply reflects the irrationality of our simulated agents.

The paths of transaction prices (normalized deviation of prices from intrinsic value) and volume in the simulations are given in Figure 4. In this model, the anchoring effect causes the mean transaction price deviation to start low and to get higher as trading proceeds. The wide bidding window causes the mean transaction price to rise over the fundamental value. As the bidding window stays constant, the fall in traded prices at the end of the asset market is caused by the positive value of the weak foresight parameter φ^*.

We re-estimate regression equation (1) using the data generated under this design. With rational, risk neutral bidders, we should observe $a = 0$, corresponding to the dividend payment per

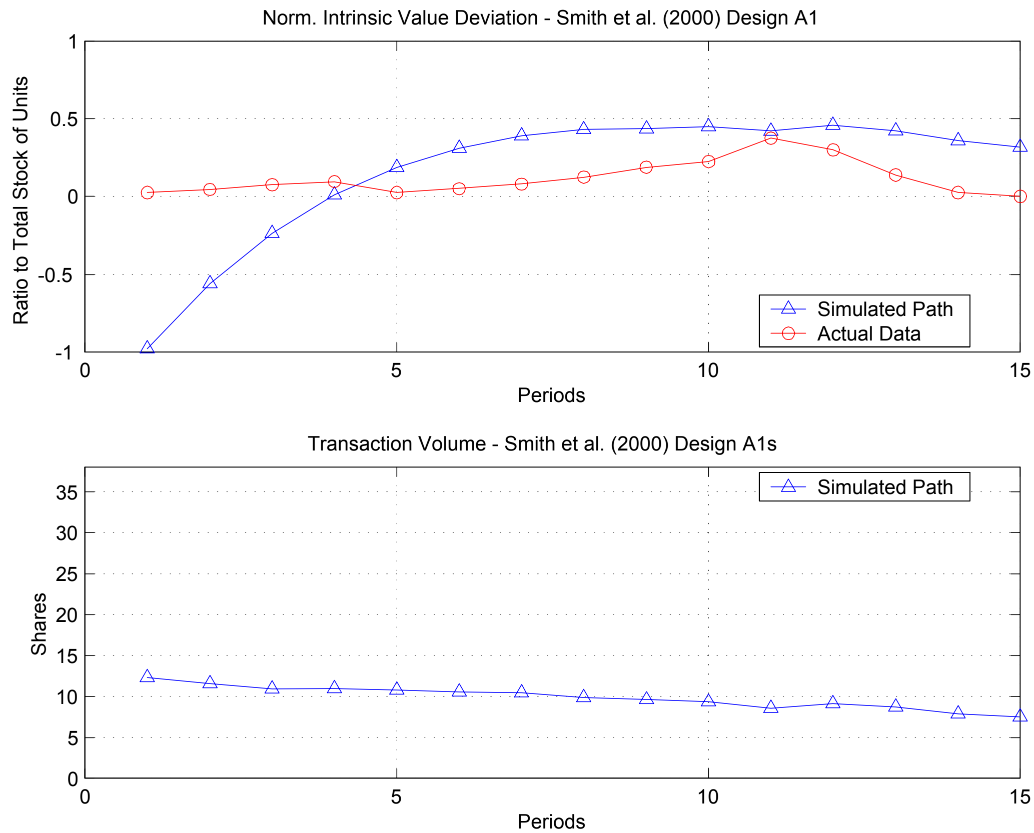

Figure 4: Normalized Intrinsic Value Deviation and Volume Paths in the Smith et al. (2000) Design and Simulations. Normalized Intrinsic Value Deviation at period t for session k is defined as $\sum_{h=1}^{vol_t^k} \frac{p_{th}^k - \overline{D}_t^T}{TSU}$. We are using this measure instead of the Mean Price Deviation, as this was the statistic reported by Smith et al. (2000). We note that Smith et al. (2001) do not report transaction volume for any of their experimental sessions.

period in this environment as well as $b = 0$. The coefficient estimates are given in Table 8. Using the simulation data, we cannot reject the model at the 1% level ($F = 125.41$, 1400 observations). We observe a significantly positive coefficient estimate for b using the simulation data. Smith et

Regression	Sessions	$\widehat{a}$	$t-$stat	$p-$value	$\widehat{b}$	$t-$stat	$p-$value
Simulations	Cumulative	0.21	26.41	< 0.01	0.032	11.20	< 0.01
Experiments[a] [b]	#1	N/A	N/A	N/A	-0.005	N/A	N/A
	#2	N/A	N/A	N/A	0.016	N/A	N/A
	#3	N/A	N/A	N/A	-0.001	N/A	N/A
	#4	N/A	N/A	N/A	0.078	N/A	N/A
	#5	N/A	N/A	N/A	0.008	N/A	N/A
	#6	N/A	N/A	N/A	0.003	N/A	N/A

[a] Smith et al. (2000) do not report estimates of $\hat{a}$ and do not report t-ratios or p-values for their estimates of $\hat{b}$.
[b] Sessions 7 and 8 of Smith et al. (2000) are not reported in this table as their dividend and endowment structures were quite different. However, the price dynamics of Sessions 7 and 8 are incorporated in the Design A1 price plot presented in Figure 4, because this figure was the only one available showing price dynamics.

Table 8: Coefficient Estimates of the Simulation and Experiment Data.

al. (2000) report coefficient estimates for b separately for each session, and these are reproduced in Table 8 for comparison purposes. Most of their coefficient estimates for b are positive (4 out of 6); they do not report $p-$values for these coefficients. As Table 8 reveals, our estimate for b lies within the range of estimates reported by Smith et al. (2001). We conclude that our model yields the same relationship between price changes and bid–offer volume found in the experimental data.

4.4 Asset Markets with a Constant Fundamental Value

Noussair et al. (2001) report on a laboratory asset market experiment similar to the Smith et al. (1988) design, but where the fundamental value of the asset remains constant over all T trading periods. Their main finding is that in half (4 out of 8) of their experimental sessions price bubbles arise, but in the other half, bubbles are not observed. Therefore, having a constant fundamental value as opposed to a decreasing fundamental value does not eliminate bubbles, though it might reduce their frequency. They have random dividend payments in each period, but their dividend payments have a constant expected value of zero. The design specifications (for the Purdue sessions) are given in Table 9:[14] In this environment, there are usually 10 players each with 10 units of the asset and a \$12.5 cash endowment. The dividends are drawn from the set {− \$0.03, −\$0.02, \$0.005,

[14] In the other sessions run at the University of Grenoble, the same design was used with a slight difference in the money endowment and the payoff conversion rate into cash. We report findings from all of their results, not just the Purdue sessions.

Players	Endowment (Cash;Quantity)	Number of Players
Class I	($12.5; 10)	10
Dividends	$d \in \{-\$0.03, -\$0.02, \$0.005, \$0.045\}^a$	$\overline{d} = \$0$
Intrinsic Value of a Share	$\overline{D}_1^T = \$0.45$	
Buy-out Value of a Share	$D_{T+1}^T = \$0.45$	

[a] Each dividend outcome will occur with probability $\frac{1}{4}$.

Table 9: Noussair et al. (2001) Purdue Design.

\$0.045} with equal probability so that the expected payment is zero (and therefore constant). The buy-out value of a unit of the asset is \$0.45.

Applying our model to this environment, we did not re-calibrate the model parameters. Instead, we again used the optimal parameter vector we found for the Smith et al. (1988) experiment. We conducted $K = 100$ independent market sessions each involving $T = 15$ periods and 10 traders. Our aim was to assess the performance of our calibrated baseline model to a different experimental environment by comparing statistics from the simulated data with those reported by Noussair et al. (2001) for the experimental data. Table 10 reports these statistics for both the simulations and for Noussair et al.'s (2001) experimental data.

Statistics	**Simulations**	**Experiments**
turnover %	142%	419%
$\overline{PA}$	1.07	0.52
$\overline{AIVD}$	0.19	0.63
$\overline{IVD}$	0.015	N/A
$\overline{p}_T - \overline{p}_{T-1}$	< 0	< 0

Table 10: Simulation and Experiment Statistics.

We see that the rate of turnover of shares in the simulated data is much lower than in the experimental data. The amplitude of bubbles in the simulation are around twice the amplitude of bubbles in the experiment. As Table 11 reveals, we observe that a small percentage of the

Turnover Composition	**Simulations**
under minimum fundamental value	0.02%
between min. and max. fund. val.	81.98%
above maximum fundamental value	18.00%

These data are not reported by Noussair et al. (2001) for the experiments.

Table 11: Turnover in the Simulation Data.

transactions in the simulated data, about 18.00%, are realized at prices higher than the maximum value of the asset. Similarly, almost none of the transactions are realized at prices lower than the minimum fundamental value.

Noussair et al. (2001) are the only authors studying asset price bubbles in the laboratory who provide an operational definition of a bubble. They say that a bubble occurs if one of the following two conditions is met:

> "(a) The median transaction price in five consecutive periods is at least 50 units of experimental currency (about 13.9%) greater than the fundamental value. (b) The average price is at least two standard deviations (of transaction prices) greater than the fundamental value for five consecutive periods."(Noussair et al. (2001), p. 94).

Using this definition, Noussair et al. find, as noted above, that price bubbles obtain in 4 out of 8 sessions. Of these 4 bubbly sessions, Noussair et al. report that just 2 experienced a price crash. However, in all bubbly sessions, prices decrease towards the end of each experimental asset market. Adopting the same criterion, we find a much higher percentage of bubbles in our simulation exercise: 95 of our 100 simulated markets met either criterion a or b.

Figure 5 shows the normalized mean price deviation and transaction volume paths for the Noussair et al. (2001) environment. This figure plots the mean experimental price and volume data only for the Noussair et al. (2001) sessions where bubbles obtained. Though the fit of the simulated data to the experimental data is not so impressive, the fact remains that even with our generic calibration, we observe bubbles in the simulation data for the Noussair et al. (2001) design.

In this laboratory environment, as in the others, the anchoring effect causes transaction prices to start low and to rise as trading proceeds. The prices rise over the expected fundamental value due to there being a wide bid range. The bid range stays constant in the Noussair et al. (2001) environment since the fundamental value is constant. The prices only start to fall because of the weak foresight of the simulated agents who submit less and less buy orders and more and more sell orders as the market proceeds.

We have again estimated regression equation (1) for this design. In the rational case with risk neutral bidders we should find that $a = 0$ corresponding to the dividend payment per period and that $b = 0$. The coefficient estimates are given in Table 12 for the simulation and the experimental data respectively. We observe significantly positive estimates for both a and b using the simulation data. Again, using the simulation data, we cannot reject the model at the 1% level ($F = 140.32$, 1400 observations). Noussair et al. (2001) find that $\widehat{a}$ and $\widehat{b}$ are significantly different from zero in only three sessions, and their values are positive in those three sessions.

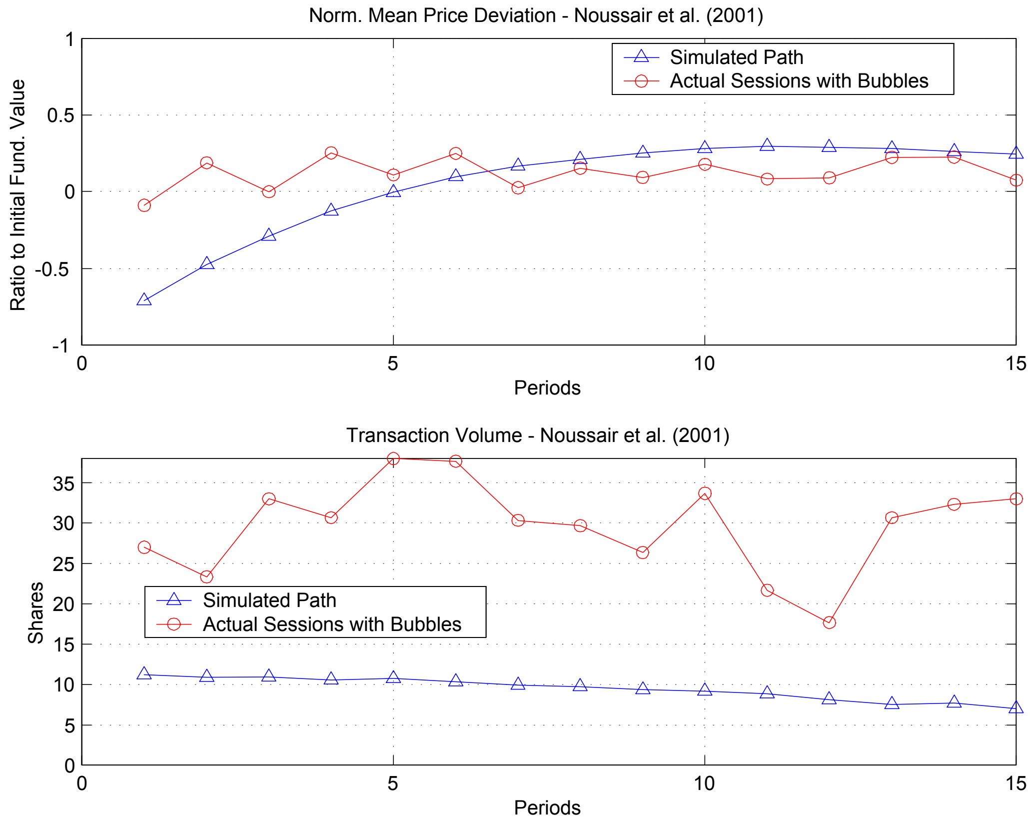

Figure 5: Normalized Mean Price Deviation and Volume Paths in the Noussair et al. (2001) Design (with Bubbles) and Simulations

Regression	Sessions	$\widehat{a}$	$t-$stat	$p-$value	$\widehat{b}$	$t-$stat	$p-$value
Simulations	Cumulative	0.046	26.40	< 0.01	0.0030	11.85	< 0.01
Experiments	G#1/ no bubble	0.0018	0.44	> 0.05	0.00040	0.35	> 0.05
	G#2/ no bubble	0.015	2.26	< 0.05	0.0013	2.68	< 0.05
	G#3/ bubble-crash	0.013	1.29	> 0.05	0.0018	2.05	< 0.05
	G#4/ no bubble	0.022	1.43	> 0.05	0.0011	0.77	> 0.05
	G#5/ bubble-crash	0.0039	0.059	> 0.05	-0.00013	-0.17	> 0.05
	P#1/ bubble	0.0078	0.48	> 0.05	-0.00023	-0.17	> 0.05
	P#2/ no bubble	0.0014	2.36	< 0.05	-0.00025	-1.13	> 0.05
	P#3/ bubble	0.0080	2.35	< 0.05	0.00075	2.61	< 0.05

Table 12: Coefficient Estimates of the Simulation and Experiment Data.

4.5 Discussion

The bubbles observed in the simulations of the Noussair et al. (2001) and Smith et al. (2000) designs have smaller price amplitude, smaller absolute intrinsic value deviation than in the simulations based on the Smith et al. (1988) design – compare Tables 6 and 10 with Table 2. The experimental bubbles were indeed smaller in the other two designs as compared with the experimental bubbles reported in Smith et al. (1988). Hence, our simulations capture this feature of the experiments quite well.

On the other hand, our simulations are far from perfect at tracking the price and volume paths in the Smith et al. (2000) and Noussair et al. (2001) designs using the parameters calibrated for the Smith et al. (1988) design. Setting aside the obvious explanation – that we use a calibration that is not optimized for these different experimental designs – an alternative explanation might be that the market environment of Smith et al. (2000) and, especially, that of Noussair et al. (2001) are considerably simpler for subjects to comprehend than the original Smith et al. (1988) environment. In these simpler environments, subjects may come to understand the fundamental value of the asset more clearly, and may take a more rational approach to submitting bids or asks. If this were the case, our model of near–zero–intelligence traders might be less appropriate as a model of human subject behavior, even among relatively inexperienced subjects.

Still, while the simulation model, calibrated to fit the Smith et al. (1988) data does not predict the data from other experimental designs so well, our main finding, that price bubbles and crashes persist in these other environments is consistent with the experimental findings. Using our model, one would be led to the same conclusion that Noussair et al. make: "the bubble phenomenon is a characteristic of a more general class of environments, not only the commonly studied declining fundamental value case," (Noussair et al. p. 97). This consistency gives us some confidence in the predictive power of our model which we explore further in the following sections.

4.6 Asset Markets with an Indefinite Horizon

One difficulty with the laboratory asset price bubble designs discussed above is their use of a known, finite horizon for the asset market. If the asset is known to be worth zero after a certain date or to have some fixed cash-out value, it seems likely that prices will fall to these values as the finite horizon looms, as subject eventually apply backward induction. Of course in real asset markets, the horizon over which assets may generate dividends and capital gains is indefinite, and so it would seem to be of some interest to examine whether asset price bubbles and crashes can arise in such indefinite horizon environments. Camerer and Weigelt (1993) present results from a double auction market where the asset that subjects traded was stochastically lived. They were interested in whether prices converged to the fundamental competitive equilibrium price, equal to the present discounted value of dividends. They found that with either inexperienced subjects, convergence to equilibrium was "slow and unreliable," primarily because traders are more uncertain of what prices will be in future periods, a problem that does not arise in repetitions of the static double auction environment. Camerer and Weigelt (1993) report that in some sessions prices slowly converged to the fundamental value while in other sessions there was no tendency towards convergence at all, with prices remaining either persistently above or persistently below the fundamental value for sustained periods of time.

We implement an indefinite horizon asset market in the same way that Camerer and Weigelt do: we specify a constant probability that the market (or asset) will continue (or live) from one trading period to the next. However we depart from Camerer and Weigelt in how we model the stochastic dividend process. In Camerer and Weigelt's paper, the dividend an agent receives from holding a unit of the asset at the end of each period depends on the agent's type for that period (high, medium or low dividend recipient). Heterogeneity of types is what promotes trade in their model. Since we already have agent heterogeneity in the form of the random components to bids and asks, we suppose instead that the dividend process is the same for all traders, which is more in line with the prior literature on bubbles in laboratory asset markets.

Specifically, we build on the design of Noussair et al. (2001), since the fundamental value of the asset in the indefinite horizon environment will also be constant. Table 13 gives our experimental design for the indefinite horizon environment. We suppose there are 10 traders in each session and each one of them is endowed with 10 units of the asset and \$12.50 in cash. A market consists of an indefinite number of trading periods. When a market ends, each remaining unit of the asset has a buy–out value of 0. There are also dividend payments after each trading period having expected value $\overline{d} = \$0.03$. The distribution of possible dividend values reported in Table 13 was obtained by simply adding \$.03 to the dividend values in the Noussair et al. (2001) design (c.f. Table 9). At the end of each trading period t the market continues with another trading period with probability $p_c = 0.9$ (Camerer and Weigelt chose $p_c = .85$), which can also be interpreted as a discount factor. Thus, the number of trading periods T remaining in a market session at the start of period t becomes a random variable with $E_t[T] = \frac{1}{1-p_c} = 10$. It follows that the fundamental value of a share at the start of any trading period is $\frac{\overline{d}}{1-p_c} = \frac{0.03}{0.1} = 0.3$.

We do not use the weak foresight parameter in this experimental design, i.e. we set $\varphi = 0$. The weak foresight parameter exogenously decreases the probability of being a buyer after each period. In the current design, after each completed trading period, there remain, on average, 10

Players	Endowment (Cash;Quantity)	Number of Players
Class I	($12.5; 10)	10
Dividends	$d \in \{\$0, \$0.01, \$0.35, \$0.075\}^a$	$\overline{d} = \$0.03$
Intrinsic Value of a Share	$\overline{D}_1^{E(T)} = \$0.30$	
Buy-out Value of a Share	$D_{E(T)+1}^{E(T)} = \$0$	

[a] Each dividend outcome will occur with probability $\frac{1}{4}$.

Table 13: Indefinite Horizon Design with Continuation Probability $p_c = 0.95$.

more trading periods to be played. Hence, there is no need for the weak foresight parameter. Aside from this change, the parameters are the ones we found to be optimal for the Smith et al. (1988) design, ($\kappa^* = 4.0846, \alpha^* = 0.8480,$ and $S^* = 5$).

Table 14 gives the simulation statistics over 100 sessions. The mean number of trading periods for each session (market) in the 100 session sample was 11.12, with a standard deviation of 9.84 periods. The sample maximum number of trading periods in any single market was 48 and the minimum number was 1 trading period. We observe both high turnover and price amplitude in these simulation results which are signatures of bubbly asset markets. As seen in Figure 6, the

Statistics	**Simulations**
turnover %	122%
$\overline{PA}$	1.34
$\overline{AIVD}$	0.21
$\overline{IVD}$	0.12
$\overline{p}_T - \overline{p}_{T-1}$	≈ 0

Table 14: Simulation Statistics with Infinite Horizon.

price level increases steadily to almost one and a half times the fundamental value within the first 8 periods of a session. Here, the parameter κ is still set at 4.0846, but $\varphi = 0$, so in Figure 6 the normalized price deviation should converge to $(\kappa/2) - 1 \approx 1.04$. Since the number of observations with more than 25 trading periods was few, we show data for the first 25 periods only. However, a careful inspection of the data after period 25 shows that indeed the normalized price deviation fluctuates around 1.04 as predicted by the model. The interesting prediction of this simulation is that although we have a bubble, the bubble never bursts. The transaction price does not decrease from period $T - 1$ to T in general. This is due to the fact that, after each trading period that has been reached, one can expect 10 more periods in the session. Hence, the fundamental value of the asset is constant. This leads to a constant bid interval. Since there is no weak foresight in the indefinite horizon model, prices do not change, i.e. the environment is stationary.

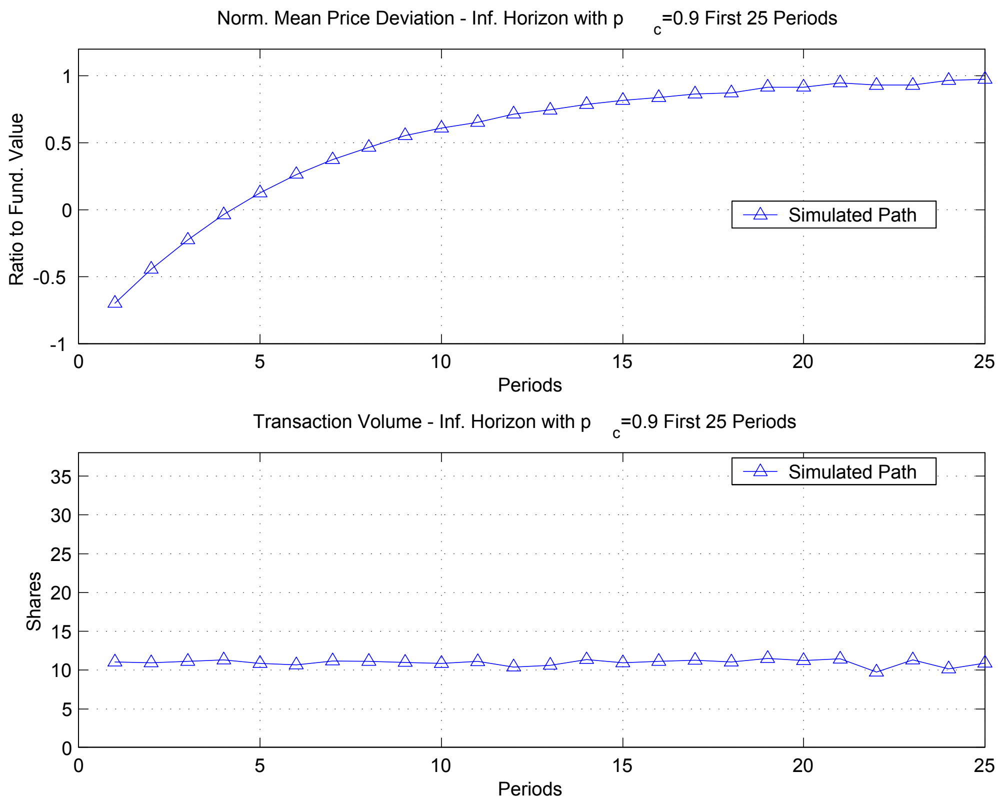

Figure 6: Normalized Mean Price Deviation and Volume Paths in the Indefinite Horizon Design Simulations.

Turnover Composition	Simulations
under minimum fundamental value	0.00%
between min. and max. fund. val.	100.00%
above maximum fundamental value	0.00%

Table 15: Turnover in the Simulation Data.

As seen in Table 15, all of the activity is observed at prices between the maximum and minimum fundamental value. At any period, the maximum value of the asset is determined as the maximum income an individual holding the asset would earn if he held the asset for the expected duration of the session: $\max D_t^{E(T)} = \frac{\max\{\$0,\$0.01,\$0.35,\$0.075\}}{1-p_c}$ =\$0.75. Similarly, the minimum value of the asset is determined as $\min D_t^{E(T)} = \frac{\min\{\$0,\$0.01,\$0.35,\$0.075\}}{1-p_c}$ =\$0.

Next, we estimate the regression equation (1) relating the change in traded prices to the difference between buy and sell offers in the previous period, equation. We find that $\widehat{a}$ is both positive and significant, but $\widehat{b}$ is negative and not statistically significant. (See Table 16). This time using the simulated data, we can reject the model at 10% level ($F = 1.54$, 1012 observations from 100 sessions). It should be noted that this result is obtained in the absence of the weak foresight parameter in the model. With risk neutral rational bidders, we should observe $a = 0$ corresponding

Regression	Sessions	$\widehat{a}$	$t-$stat	$p-$value	$\widehat{b}$	$t-$stat	$p-$value
Simulations	Cumulative	0.030	29.96	< 0.01	-0.00026	-1.24	> 0.1

Table 16: Coefficient Estimates of the Simulation Data.

to the change in the dividend payment per period and $b = 0$.

These findings are not inconsistent with those reported by Camerer and Weigelt (1993), even though there are some differences between their approach and our own as noted above. In particular, they report sessions where prices remained well above the fundamental equilibrium price with no tendency toward convergence (c.f. our findings with Figures 2, 3, 13 and 14 of the Camerer-Weigelt paper). While we have not tested the indefinite horizon market design we examine here with human subjects, we think it would be of interest to do so. If our prediction is correct, we should see a sustained departure of asset prices above their fundamental value and no crashes, at least among inexperienced subjects.

4.7 Asset Markets Without Opportunities for Speculation

In an important paper, Lei et al. (2001) show that speculative motives are not needed to generate bubbles in experimental asset markets. They consider a design where players are ex-ante classified as either buyers or sellers. A buyer can only buy shares and a seller can only sell shares. Players remain in the same role for the duration of a session. By preventing buyers from acting also as sellers, speculative explanations for bubbles (e.g. the greater fool hypothesis) are effectively shut down. Despite shutting down the possibility of speculation, Lei et al. (2001) report that asset price bubbles arise in their experimental sessions. This result refutes the conjecture that speculative behavior is the source of the laboratory asset price bubble phenomena. In this section we apply our model to the Lei et al. (2001) environment. As noted earlier, our asset price bubble explanation does not rely on any speculative behavior. Instead it relies on the anchoring and weak foresight assumptions.

In the Lei et al. environment, sellers are endowed with shares and buyers are endowed with money. Since buyers and sellers are fixed throughout a session, we do not use a decreasing probability for being a buyer (weak foresight assumption), or for that matter, any probabilistic device for determining who is a buyer or a seller. In this sense, the Lei et al. model is quite different from the laboratory environments we have previously considered. Aside from this one change, however, our model is the same as before. In particular, loose budget constraints remain in place and buyers and sellers continue to use the mean traded price of the previous period as an anchor for their current bids and asks, which continue to also have a random component. There continues to be a common upper bound to the bid ask range equal to κ times the fundamental value of the asset. Therefore, the loose budget constraint parameter κ, the anchoring weight parameter α, and the number of trading sequences in a period S are the only model parameters that need to be calibrated for this design.

The number of periods in a market is set at $T = 12$, as in Lei et al. (2001). The initial endowments, dividend payments and the number of players are given in Table 17.[15] As this environment differs considerably from the ones we have previously examined, we chose to re-calibrate our simulation model using a similar optimization algorithm to the one provided for the designs with speculation. We used the same two–step simulated method of moments procedure described in section 3.1, modified for the fact that we no longer search over values of the weak foresight parameter, φ. The optimal parameters were as follows: the number of trading sequences $S^* = 3$,

Players	Endowment (Cash;Quantity)	Number of Players
Buyers	$(\$24; 0)$	4
Sellers	$(\$0; 20)$	4
Dividends	$d \in \{\$0.0667, \$0.133\}^a$	$\overline{d} = 0.10$
Intrinsic Value of a Share	$\overline{D}_1^T = \$1.20$	
Buy-out Value of a Share	$D_{T+1}^T = \$0$	

[a] Each dividend outcome will occur with probability $\frac{1}{2}$.

Table 17: Lei et al. (2001) Experimental No-Speculation Design.

[15] Lei et al. (2001) use the cleared order book convention in their experiments. In the simulations we report here we also use this convention. We note that the particular order book convention used does not affect the simulation results extensively.

the loose budget constraint parameter $\kappa^* = 3.5123$, and the anchoring parameter $\alpha^* = 0.6$.

We ran $K = 100$ independent market sessions of our model with the calibration as described in the last section. Table 18 reports the statistics both for the simulations and for the actual experimental data.

Statistics	**Simulations**	**Experiments**
turnover %	92%	82%
$\overline{PA}$	3.63	3.48
$\overline{AIVD}$	0.49	N/A
$\overline{IVD}$	0.40	N/A
$\overline{p}_T - \overline{p}_{T-1}$	< 0	< 0

Table 18: Simulation and Experiment Statistics.

We observe that asset turnover is lower relative to other designs, as it also is in the experimental data. We further note that values for the price amplitude, the absolute intrinsic value deviation and the intrinsic value deviation in the simulated data are a very close match to the experimental data.

In Figure 7 we plot the normalized mean price deviation and transaction volume from the simulated data and for the actual Lei et al. (2001) experimental data. The normalized mean-price deviation is about -30% of initial fundamental value in period 1 and it increases up to less than 60% in period 5 before falling back to just 28% in last period. This provides strong evidence of a bubble-crash pattern in this design. Indeed, Lei et al. (2001) claim that there is a bubble–crash pattern in two of their three No-Speculation treatments.

In the simulations, transaction volume starts high at 6.5 units in the first period, and stays fairly steady for the first 9 periods, before falling off and ending up below 5 in period 12. We observe in Table 19 that 79.20% of the total turnover is traded at prices higher than the maximum value of the asset. On the other hand, only 5.38% of the total turnover are traded at prices lower than the minimum value of asset. In the actual experimental data fewer units are traded at prices higher than the maximum value and more units are traded at prices lower than the minimum value - see Table 19. Still, our simulation findings are broadly consistent with the laboratory findings: some agents are trading irrationally.

Turnover Composition	**Simulations**	**Experiments**
under minimum fundamental value	5.38%	16.67%
between min. and max. fund. val.	15.42%	44.44%
above maximum fundamental value	79.20%	38.89%

Table 19: Turnover in the Simulation and Experiment Data.

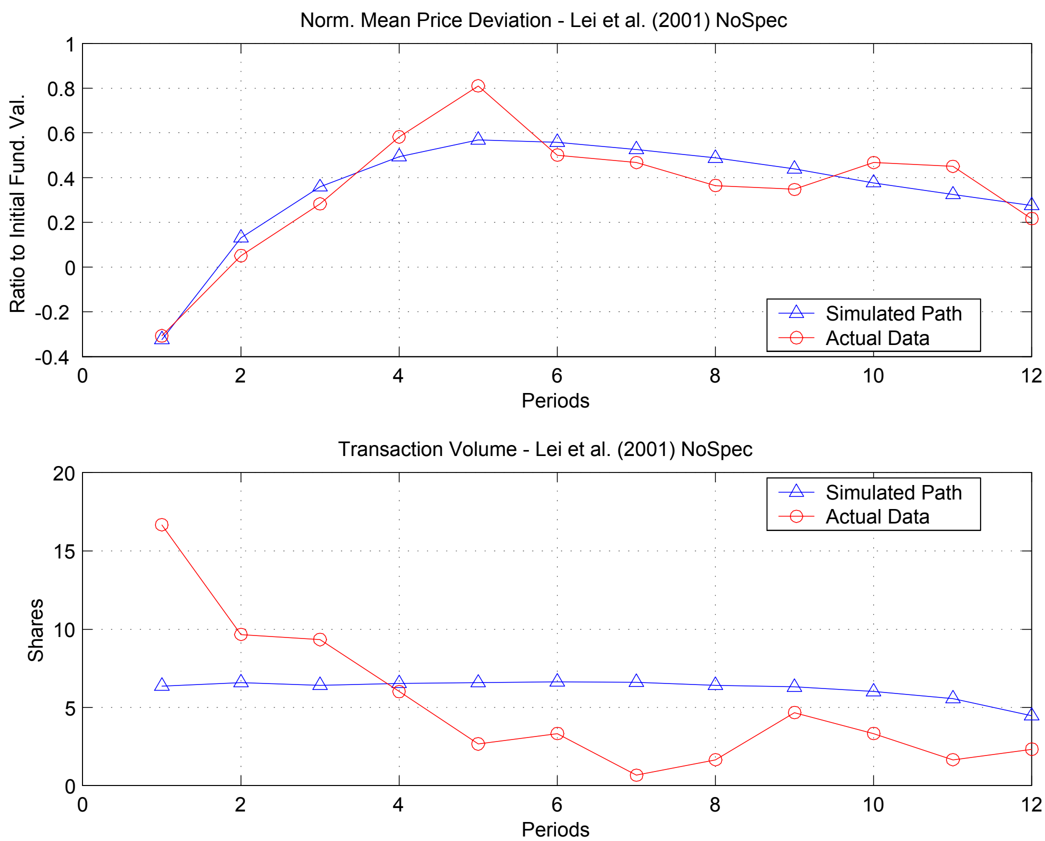

Figure 7: Average Mean Price and Volume Paths in the Simulation Data

In our agent-based model, the anchoring effect causes traded prices to start low and to become higher as trading proceeds. However, at some point, buyers start running out of money, since they no longer have the option of re–selling units and replenishing their cash balances with capital gains as in the trader market environment. Consequently, transaction prices fall, as sellers compete to transact with fewer buyers. Transaction volume falls a little as well, especially towards the end of the market and the market experiences a "crash". Notice that this explanation for the bubble-crash phenomenon does not rely on any foresight of the finite horizon.

We again estimate regression equation (1) using simulated data for this design. With rational, risk neutral bidders we should observe $a = -0.1$ corresponding to the change in the expected dividend payment per period and $b = 0$. The coefficient estimates are given in Table 20 for the simulation and experimental data, respectively. We find that $\widehat{a}$ is significantly negative and $\widehat{b}$ is positive, though not significantly different from zero. Lei et al. report that for the experimental data $\widehat{a}$ is significantly negative and $\widehat{b}$ is significantly positive.[16] We can reject the model in the simulation data at 10% level with F =0.83 (with 1100 observations obtained from 100 sessions and 11 periods for t =2,...,12). These results suggest once again that our model does a fair job of characterizing the major features of laboratory asset bubble phenomenon, and is also capable of capturing some of the more subtle features of the data.

Regression	Sessions	$\widehat{a}$	$t-$stat	$p-$value	$\widehat{b}$	$t-$stat	$p-$value
Simulations	Cumulative	$-$ 0.036	-5.65	< 0.01	0.0008	0.22	> 0.05
Experiments	Cumulative	-0.2147	-9.6293	< 0.05	0.002	11.80	<0.05

Table 20: Coefficient Estimates of the Simulation and Experiment Data

5 Conclusion

We have developed a simple, agent-based behavioral model with the aim of understanding the phenomenon of asset price bubbles and crashes among inexperienced subjects as reported in a number of laboratory studies beginning with Smith et al. (1988). We turned to agent–based modeling because standard theoretical models impose assumptions that are either unsatisfied or are too restrictive for the double auction environment of the laboratory studies. Indeed, rational asset pricing models would predict that agents placed in the laboratory environments would not engage in any trade whatsoever.

[16] We note that Lei et al. (2001) use the median transaction price in their estimates and statistics.

Our simple agent-based behavioral model builds on the idea of using budget constrained, zero–intelligence traders as pioneered by Gode and Sunder (1993, 1997). The ZI approach is not intended as a commentary on the rationality of human subjects. Rather, this approach effectively lays bare the importance of institutions, e.g. trading rules, procedures, and other features of the market environment relative to human decision-making skills in the determination of observed market outcomes. We have had to modify the Gode and Sunder ZI methodology in several ways to address the laboratory price bubble phenomenon but the resulting model continues to focus attention on the role played by the features of the laboratory asset market, in particular, whether there is a finite or infinite horizon, whether the fundamental value of the asset is decreasing over time and the frequency/timing of dividend payments. We have also addressed the "greater fool" hypothesis that asset price bubbles arise from speculative behavior on the part of players who buy at high prices in hopes of selling at even higher prices. Our anchoring explanation for price increases in the initial periods of a market session does not rely on any kind of speculative motive; it is purely backward looking. Accordingly, we are able to obtain price bubbles in buyer–seller markets, such as the one considered by Lei et al. (2001) where speculation motives are explicitly precluded. Our behavioral model, utilizing both anchoring and weak foresight, is the first, and to our knowledge, the only model that has been offered to explain the robust finding of laboratory asset price bubbles that has fascinated so many experimental economists. Our model not only generates price bubbles and crashes that are qualitatively similar to those found in the experimental data, but in many instances we are close to obtaining the right magnitudes in price changes, volume and other statistics as well. Our model can also replicate other, more subtle findings from the experimental studies such as the regression estimated relationship between changes in traded prices and differences in the volume of bids relative to offers. A testable implication of our model is that both anchoring effects and finite horizons matter; eliminating the ability of traders to condition on past transaction prices (e.g. by severely restricting the information they receive) and replacing the finite horizon with an indefinite horizon might work to eliminate bubbles and crashes with inexperienced subjects. Of course, we emphasize that our model cannot explain why bubbles and crashes cease to obtain as players gather experience; to explain that we would have to add further structure to the model, e.g. that agents are concerned with the payoff consequences of their actions. As we mentioned at the beginning of the paper, our model seeks only to address the behavior of *inexperienced* players, which has also been the primary focus of the experimental literature on bubbles since Smith et al.'s (1988) paper. We leave the modeling of how experience affects trading behavior in laboratory asset markets to future research.

References

[1] Blanchard, O. (1979), "Speculative Bubbles, Crashes and Rational Expectations," *Economics Letters* 3, 387–389.

[2] P.J. Brewer, M. Huang, B. Nelson and C. Plott (2002), "On the Behavioral Foundations of the Law of Supply and Demand: Human Convergence and Robot Randomness," *Experimental Economics* 5, 179–208.

[3] Camerer, C.F. and K. Weigelt (1993) "Convergence in Experimental Double Auctions for Stochastically Lived Assets" in D. Friedman and J. Rust, eds., *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Proceedings Volume XIV, (New York: Addison-Wesley), 355-396.

[4] N.T. Chan, B. LeBaron, A.W. Lo and T. Poggio (1999), "Agent Based Financial Markets: A Comparison with Experimental Markets," working paper.

[5] DeLong, J.B. A. Shleifer, L.H. Summers and R.J. Waldman (1990), "Noise Trader Risk in Financial Markets," *Journal of Political Economy* 98, 703–738.

[6] Diba, B.T. and H.I. Grossman (1987), " On the Inception of Rational Bubbles," *Quarterly Journal of Economics* 102, 697–700.

[7] Duffy (2004), "Agent–Based Models and Human Subject Experiments," to appear in K.L. Judd and L. Tesfatsion, (Eds.), *Handbook of Computational Economics* vol. 2, Amsterdam: Elsevier Science.

[8] Fisher, E.O'N. and F.S. Kelly (2000) "Experimental Foreign Exchange Markets," *Pacific Economic Review* 5, 365–387.

[9] Friedman, D. (1993), "The Double Auction Market Institution: A Survey," in D. Friedman and J. Rust, eds., *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Proceedings Volume XIV, (New York: Addison-Wesley), 3–25.

[10] Froot, K.A. and M. Obstfeld (1991), "Intrinsic Bubbles: The Case of Stock Prices," *American Economic Review* 81, 1189–1214.

[11] Genesove, D. and C. Mayer (2001), "Loss Aversion and Seller Behavior: Evidence from the Housing Market," *Quarterly Journal of Economics* 1233–1260.

[12] Gode, D.K. and S. Sunder (1993), "Allocative Efficiency of Markets with Zero–Intelligence Traders: Markets as a Partial Substitute for Individual Rationality," *Journal of Political Economy* 101, 119–137.

[13] Gode, D.K. and S. Sunder (1997), "What Makes Markets Allocationally Efficient?" *Quarterly Journal of Economics* 112, 603–630.

[14] King, R.R., V.L. Smith, A.W. Williams and M. Van Boening (1993), "The Robustness of Bubbles and Crashes in Experimental Stock Markets," in R.H. Day and P. Chen, (eds.) *Nonlinear Dynamics and Evolutionary Economics* (New York: Oxford University Press), 183–200.

[15] Lei, V., C.N. Noussair and C.R. Plott (2001), "Nonspeculative Bubbles in Experimental Asset Markets: Lack of Common Knowledge of Rationality vs. Actual Irrationality," *Econometrica* 69, 831–859.

[16] Noussair, C., S. Robin and B. Ruffieux (2001), "Price Bubbles in Laboratory Asset Markets with Constant Fundamental Values," *Experimental Economics* 4, 87–105.

[17] Porter, D.P. and V.L. Smith (1995) "Futures Contracting and Dividend Uncertainty in Experimental Asset Markets," *Journal of Business* 68, 509–541.

[18] Smith, V.L., G. Suchanek, and A. Williams (1988), "Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica* 56, 1119-1151.

[19] Smith, V.L., M. van Boening, and C.P. Wellford (2000), "Dividend Timing and Behavior in Laboratory Asset Markets," *Economic Theory* 16, 567–583.

[20] Tirole, J. (1985), "Asset Bubbles and Overlapping Generations," *Econometrica* 53, 1499–1528.

[21] van Boening, M., A.W. Williams and S. LaMaster (1993), "Price Bubbles and Crashes in Experimental Call Markets," *Economics Letters* 41, 179–185.