

*To appear in the HANDBOOK OF EXPERIMENTAL FINANCE, Sascha Füllbrunn and Ernan Haruvy (eds), Edward Elgar Publishing*

## **Using results from Learning to Forecast laboratory experiments to predict the effect of futures markets on spot market stability**

**Johan de Jong<sup>a</sup>, Joep Sonnemans<sup>b</sup> and Jan Tuinstra<sup>b</sup>**

<sup>a</sup> *Erasmus School of Economics, Erasmus University Rotterdam*

<sup>b</sup> *Amsterdam School of Economics, University of Amsterdam*

In this chapter we first give a short overview of Learning to Forecast (LtF) experiments, thereby focusing on the differences between markets with positive and negative expectations feedback. Subsequently, we discuss how the results of these experiments can be used to predict behavior for more complicated market environments that exhibit both types of feedback. In particular, we will consider the case where a futures market is connected with a spot market.

### **1. Learning to Forecast experiments**

Market behavior is often driven by expectations about future market prices. An investor who expects a stock to increase (decrease) in value in the near future, will increase (decrease) her position in this stock in order to reap potential capital gains. A producer of an agricultural product who expects that the market price will be high (low) by the time he will be able to harvest and sell his crop, will sow more (fewer) seeds. To get a better understanding of the functioning of markets, we need to understand how market participants form expectations, and how these expectations affect price dynamics and market outcomes.

Analysis of individual market behavior alone, such as buying and selling decisions in financial markets or supply decisions in production markets, will typically be insufficient to determine the underlying price expectations. If an investor purchases a specific stock, we may conclude that this investor expects a high future price of this stock, but it does not reveal how high he expects this future price to be exactly. Even if we know this investor's exact willingness to pay for this stock, we will still not know his price expectation, unless we know the utility function of this investor exactly (and even then, we would have to rely on the assumption that his trading behavior is optimal given his beliefs).

This problem can be addressed by using laboratory experiments to measure expectations about future market prices directly. This can be done through so-called *Learning to Forecast* (LtF) laboratory experiments, where the only task of participants is to predict future market prices (or some other economic variable, such as inflation). Participants in these LtF experiments do not trade themselves; a computer algorithm calculates the optimal trading behavior given the forecasts of the participants active in the same market, aggregates these trading decisions and determines the resulting market price (typically the price that clears the market).<sup>1</sup> The earnings of the participants are determined by their forecast errors: if their predictions are more accurate, participants earn more. In that sense, the participants are more like professional analysts than private investors or producers.<sup>2</sup>

Participants typically have to predict prices for about 50 consecutive periods and can observe the past market prices and their own past predictions, but not the predictions of other participants. They also do not know the trading strategies of the investors or producers they advise, or the price generating mechanism. However, participants are given qualitative information about this, so that they have some idea about the effect their prediction will have on the market clearing price (note that exact knowledge of the

---

<sup>1</sup> Aggregate market behavior in laboratory experiments where participants do trade themselves is very similar to aggregate behavior in LtF experiments (see Bao et al., 2013, and Bao et al., 2017). This suggests that focusing on the expectation formation process still allows us to capture the most important features of market dynamics. This is further supported by laboratory experiments where participants have to trade and have to submit (long-term) price forecasts as well, see e.g. Haruvy et al. (2007) and Hanaki et al. (2018). These studies show that price forecasts are consistent with trading behavior, and that adding the forecast elicitation task does not change trading behavior (provided participants are only rewarded for trading, or only for their forecasts, see Hanaki et al., 2018).

<sup>2</sup> The task is typically explained to the participants as follows. In positive feedback markets: "You are a financial advisor to a pension fund that wants to optimally invest a large amount of money. The pension fund has two investment options: a risk-free investment and a risky investment. The risk-free investment is putting all money on a bank account paying a fixed interest rate. The alternative risky investment is an investment in the stock market. In each time period the pension fund has to decide which fraction of their money to put on the bank account and which fraction of the money to spend on buying stocks. In order to make an optimal investment decision the pension fund needs an accurate prediction of the price of stocks. As their financial advisor, you have to predict the stock market price during 52 subsequent time periods. Your earnings during the experiment depend upon your forecasting accuracy. The smaller your forecasting errors in each period, the higher your total earnings." (Hommes et al., 2005).

In negative feedback markets: "You are the adviser to a producer. The nature of the product that is being produced is not relevant in this experiment. At the start of each period you make a prediction of the price of the product in that period. The producer you are coupled with decides how much to produce, based upon your prediction of the price. Several producers are active in one market. Every producer is coupled with exactly one adviser (participant in this experiment) and every adviser with exactly one producer. The realized price is determined by the total production of all producers in a market and the total consumer demand (the realized price is such that total supply equals total demand)." (Sonnemans et al., 2004).

underlying price generating mechanism does not appear to have an effect on price forecasts and the resulting market dynamics, see Sonnemans and Tuinstra, 2010).<sup>3</sup>

The essential feature of LtF experiments is that *expectations feedback* is taken into account: participants' forecasts are based on previous market prices, and future values of these market prices in turn are partly determined, through trading decisions and market clearing, by participants' forecasts, as explained above.<sup>4</sup> Depending on the direction in which the realized price changes when the average price forecast increases, the expectation feedback can be characterized as positive or negative. In the next two sections we will discuss these two cases separately.

For both cases we are interested in the individual prediction strategies that participants use (and how these strategies depend upon the characteristics of the market environment), and in the resulting market dynamics, that is, whether prices converge to their equilibrium values, or whether they exhibit mispricing, excess volatility or even bubbles and crashes.

## **2. Positive expectations feedback: Financial markets**

Under *positive expectations feedback* high average expectations (i.e. higher than the equilibrium price) will lead to a realized price that is also high (i.e. higher than the equilibrium price), while low average expectations will result in low prices. A prime example is provided by financial asset markets, where an investor who expects the price of a particular asset to increase in the future will want to invest more in this asset, while investors who expect the price to decrease might be eager to divest. If, on average, investors expect the price to increase in the future, aggregate demand for the asset will go up, and the market clearing price will increase instantaneously. This self-confirming nature of financial markets leads to a positive correlation between average expectations and the price and is related to the concept of *strategic complements* from game theory (see Bulow et al., 1985): if the other investors, on average, forecast a high (low) price, the realized price

---

<sup>3</sup> This approach to investigate expectations was introduced in Marimon et al. (1993). See Hommes (2011) for an overview of the early literature.

<sup>4</sup> Alternatively, expectation formation can be studied by letting participants predict the next value of an exogenously given time series (see Glaser et al., 2019, for a recent example). However, we are specifically interested in expectations in an environment with expectations feedback, where prices are endogenous and depend upon price forecasts, since this is relevant for many economic applications.

will be high (low) as well and the best response is then to predict in line with the other investors. This positive expectations feedback also characterizes financial market LtF experiments.<sup>5</sup>

Because *current* demand for an asset is based upon the expected return of investing in that asset and therefore also depends upon its *future* market price, in most of these financial market LtF experiments participants have to predict the market price two periods ahead. That is, because the market clearing price in period  $t$ , say  $p_t$ , is determined by demand for the asset in that period, which depends upon the price expected, by all participants in the market, for period  $t+1$ , the market clearing price  $p_t$  is not available yet when participants have to predict  $p_{t+1}$ . Therefore, participants have to base their predictions for period  $t+1$  on prices up until period  $t-1$ .

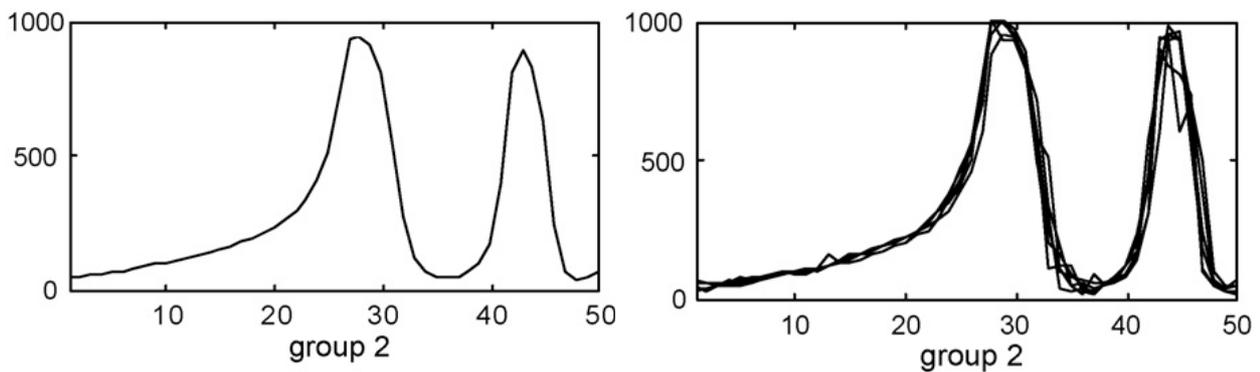
Now let us briefly consider, as an example, the positive feedback LtF experiment from Hommes et al. (2008). The price generating mechanism (which is based upon investors that are myopic mean-variance maximizers, fixed exogenous supply of the asset, and market clearing) in that experiment is given by

$$p_t = \frac{1}{1+r} [\bar{p}_{t+1}^e + \bar{y}] \quad (1)$$

Here  $r$  is the interest rate and  $\bar{y}$  is the mean dividend paid out by the asset. Moreover,  $\bar{p}_{t+1}^e = \frac{1}{6} \sum_{h=1}^6 p_{h,t+1}^e$  is the average price prediction for period  $t+1$ , where the average is taken over the six participants that are active in a market (see Hommes et al., 2008, for a derivation of this market clearing price and further details). For this experiment the interest rate and mean dividend are set at  $r=0.05$  and  $\bar{y} = 3$ , respectively, which results in a fundamental value of 60 (that is, if the average expectation equals 60, then the market clearing price will also be equal to 60). Participants to the experiment know the interest rate and the mean dividend but are not given Equation (1). They also do not observe the predictions of other participants. Their payment for each of the 50 periods for which they have to form a prediction is based upon their quadratic forecast error.

---

<sup>5</sup> Another game that is characterized by strategic complementarities and that is closely related to positive feedback LtF experiments is the number guessing game (or beauty contest game), see Nagel (1995). Hanaki et al. (2019) show that the strategic environment (whether actions are strategic complements or strategic substitutes) has a significant effect on the deviations from the Nash equilibrium in a laboratory experiment with the number guessing game when groups contain at least 5 participants (the minimal group size of the experiments we discuss in this chapter is 6). See Sonnemans and Tuinstra (2010) for a further discussion of the differences between the number guessing game (with strategic complements) and positive feedback LtF experiments.



**Figure 1:** Example of prices and predictions in a positive feedback market (see Figures 3 and 7 (market 2) in Hommes et al., 2008). The horizontal axis displays the periods. The prices (left panel) go through two large bubbles. The predictions of the six participants (right panel) exhibit a high degree of coordination. The fundamental price in this market is 60 and price predictions are capped at 1000.

In Hommes et al. (2008) five of the six markets show large bubbles in asset prices, and in none of these markets does the price converge to its fundamental value. Figure 1 shows the realized prices and the individual price predictions in market 2 from that experiment. Note that prices increase to a level that is sixteen times as high as the fundamental value of 60. The high level of coordination of price predictions that can be seen in the right-hand side of Figure 1 is remarkable, in particular because participants cannot see each other's predictions. This coordination is a consequence of the strategic complementarity of predictions in positive feedback LtF experiments. Moreover, participants tend to coordinate on strategies that extrapolate past trends in prices. In particular, simple trend-extrapolating prediction strategies of the form  $p_{h,t+1}^e = p_{t-1} + \theta(p_{t-1} - p_{t-2})$  are often encountered in positive feedback LtF experiments, where typically  $\theta$  is somewhere between 0.7 and 1.0.<sup>6</sup> Such strategies lead to self-confirming increases in asset prices and explain the bubbles and crashes that emerge in these positive feedback LtF experiments. Both features, large bubbles in asset prices and strong coordination of predictions, are quite robust. They can, for example, also be observed when fundamental robot traders are added to the market (see Hommes et al., 2005), or when the number of

<sup>6</sup> For example, for 42 of the 84 participants in the experiment presented in Hommes et al. (2005), individual predictions could be fitted by a rule of the type  $p_{h,t+1}^e = \alpha + \beta p_{t-1} + \theta(p_{t-1} - p_{t-2})$ , where the 25<sup>th</sup> percentile of the estimated  $\theta$  values is equal to 0.68 and the 75<sup>th</sup> percentile is equal to 1.05 (with the lowest (highest) estimate of  $\theta$  equal to 0.41 (1.60) and the average estimate equal to 0.88).

participants in each market is (much) larger than six (see Bao et al., 2020 and Hommes et al., 2021), and they do not disappear with experience (see Kopányi-Peuker and Weber, 2021). For a general and up-to-date overview of LtF experiments with positive expectations feedback we refer to the recent article by Bao et al. (2021).

### **3. Negative expectations feedback: Production markets**

In production markets, in particular when the production process is time consuming (such as is the case, for example, for many agricultural products) firms have to decide their production level long before they can actually supply their product to the market. Their decisions are therefore based upon the price they *expect* to get when their product is ready to be sold. Such markets are typically characterized by *negative expectations feedback*: If firms on average expect a price that is higher (lower) than the equilibrium price, their supply will be high (low), and the market price – that is, the price for which aggregate supply equals (the exogenously given) consumer demand – will then be lower (higher) than the equilibrium price. In such markets predictions are *strategic substitutes*: If other participants on average predict a high price, the best response for a participant is to predict a low price, and vice versa.

In negative expectations feedback LtF experiments participants typically predict only one period ahead, where this period is interpreted as the time between the production decision and selling the product.<sup>7</sup> A classic example is the hog cycle (Tinbergen, 1930) where farmers decide how many pigs to raise, and one period corresponds to the time needed to fatten a pig. If farmers would use a naïve prediction strategy, that is, if they expect next period's price to be equal to the current price, prices will oscillate around their equilibrium values, with these oscillations eventually converging to the equilibrium price, diverging away from it, or remaining constant, depending on the slope of the supply curve relative to the slope of the demand curve.

When predictions are strategic substitutes, it pays off for participants to disagree with the majority of the other participants, which inhibits coordination of predictions. Indeed, where predictions in positive feedback LtF experiments tend to be strongly correlated, forecasts in the first couple of periods in negative feedback LtF experiments are

---

<sup>7</sup> Exceptions are models where investments decisions have to be made that last for more years, like planting a multi-year crop or building a power station (see for example Arango and Moxnes, 2012).

typically scattered around the equilibrium price, with some predictions above and some below that equilibrium price. There will therefore be a tendency for the average forecast to be close to the equilibrium price and prices often converge to that equilibrium price quickly. In fact, persistent fluctuations or bubbles and crashes rarely occur in negative feedback LtF experiments (see, for example, Hommes et al., 2007, Heemeijer et al., 2009, and Bao et al., 2013).

Let us now consider an example of such a negative expectations feedback market in a bit more detail. Suppose there are  $K$  firms active in the market for a certain commodity, and each firm needs one period to produce that commodity. The decision on how much to supply to the market in period  $t$  will depend upon the firm's expected price for that period, denoted by  $p_{k,t}^e$  for firm  $k$ . These expectations are formed at the end of period  $t-1$ , after the firm has observed the market clearing price for that period,  $p_{t-1}$ . For a given cost function (assumed to be the same for each firm in the market) the firm optimally sets its quantity such that marginal costs are equal to the expected price, resulting in individual supply  $S(p_{k,t}^e)$  for firm  $k$ .

Aggregate supply meets a consumer demand that is assumed to linearly decrease with the price  $p_t$ :

$$D(p_t) = a - bp_t + \mu_t, \quad (2)$$

where  $a, b > 0$  are strictly positive demand parameters, and  $\mu_t$  is a normally and independently distributed random series of small periodic demand shocks, with mean zero.

The market clearing price in period  $t$  can now be found as

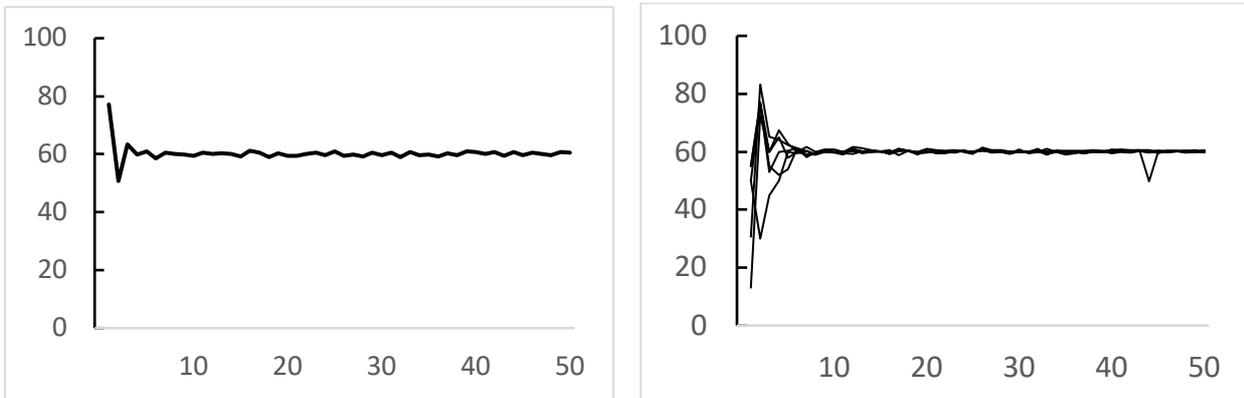
$$p_t = \frac{a - \sum_{k=1}^K S(p_{k,t}^e)}{b} + \varepsilon_t$$

where  $\varepsilon_t = \mu_t/b$  is a noise term.

In Heemeijer et al. (2009) the supply function is assumed to be linear as well, which gives rise to the following pricing equation

$$p_t = \frac{20}{21}(123 - \bar{p}_t^e) + \varepsilon_t,$$

with  $\varepsilon_t \sim N(0, 1/4)$ , and where  $\bar{p}_t^e$  is the average price forecast for period  $t$ , with the average taken over the six participants that are active in the same market.<sup>8</sup> Note that the fundamental value is equal to 60: if all participants predict a price of 60, the market clearing price will be 60 in expectation.



**Figure 2:** Example of prices (left panel) and individual predictions of the six participants (right panel) in a negative feedback LtF experiment (market 1 from Heemeijer et al., 2009). The horizontal axis displays the periods.

Figure 2 shows a typical experimental result (market 1 from Heemeijer et al., 2009). In the first couple of periods predictions go through an initial phase of high volatility, after which they neatly converge to the equilibrium price, only to be disturbed occasionally by the impact of a mistake by one of the participants. The same behavior is observed in the other five negative feedback markets in Heemeijer et al. (2009). Moreover, if the supply curve is steeper than the demand curve, and the hog cycle model would predict large fluctuations, prices will still be relatively close to their equilibrium value (see Hommes et al., 2007). Individual prediction strategies can often be classified as *naive* (the prediction is equal to the last observed price, that is,  $p_{k,t}^e = p_{t-1}$ , for firm  $k$ ), or *adaptive* (the prediction is adapted in the direction of the last observed price, that is  $p_{k,t}^e = w p_{t-1} + (1 - w) p_{k,t-1}^e$ , with  $w \in (0, 1)$  the updating parameter), see Heemeijer et al. (2009). Note, however, that due to the high level of convergence in negative feedback markets, individual prediction strategies are more difficult to uncover than in positive feedback markets. As soon as the market has converged to the rational expectations equilibrium, many different prediction strategies

<sup>8</sup> One would get this price generating mechanism by setting, for example,  $a = 123$ ,  $b = 21/20$  and a cost function for production equal to  $c(q) = 3q^2$  (leading to an individual supply curve of  $S(p_{k,t}^e) = \frac{1}{6} p_{k,t}^e$ ).

(e.g. naive, adaptive or trend-extrapolating expectations) all correctly predict the equilibrium price. In the strategy experiment discussed in Sonnemans et al. (2004) prediction strategies are measured directly, that is, participants formulate algorithms that predict prices in simulations where they play against the algorithms of the other participants. Most strategies are quite complicated, but typically (weighted) average previous prices or adaptive strategies are also used there.<sup>9</sup>

#### **4. Connecting production and financial markets**

The dichotomy between positive and negative expectations feedback presented above is not so clear cut for actual markets. Some markets may have elements of both positive and negative expectations feedback, in particular when different types of market participants are active on this market. It might also be the case that a positive feedback financial market is connected to a negative feedback production market. For example, consider a spot market for a particular commodity, and a financial market where speculators trade in futures on the commodity that is sold in the spot market. An important question then is what the effect is of connecting the positive feedback futures market and the negative feedback spot market on the dynamics of spot market prices. Existing empirical evidence on the impact of futures markets on the stability of spot markets is mixed. On one hand, there is some evidence that the introduction of a futures market decreases volatility in the spot market (Working, 1960; Gray, 1963; Powers, 1970; Netz, 1995), but it should be noted that these new futures markets are relatively small at the start. On the other hand, Roll (1984) finds that that the volatility in the prices of orange juice futures is much larger than can be explained by changes in the weather and other factors that influence supply and demand.

We investigate the effect a futures market has on the associated spot market by developing a simple model along the lines of Muth (1961) and Sarris (1984) and combine this with our knowledge of the prediction strategies that are typically used in positive and negative feedback LtF experiments. Our focus here will be on the intuition behind the model. The interested reader can find a formal treatment in de Jong et al. (2019).

---

<sup>9</sup> Note however that these algorithms may differ from the implicit strategies used in laboratory experiments because it may be difficult for some participants to formulate the heuristic they would use.

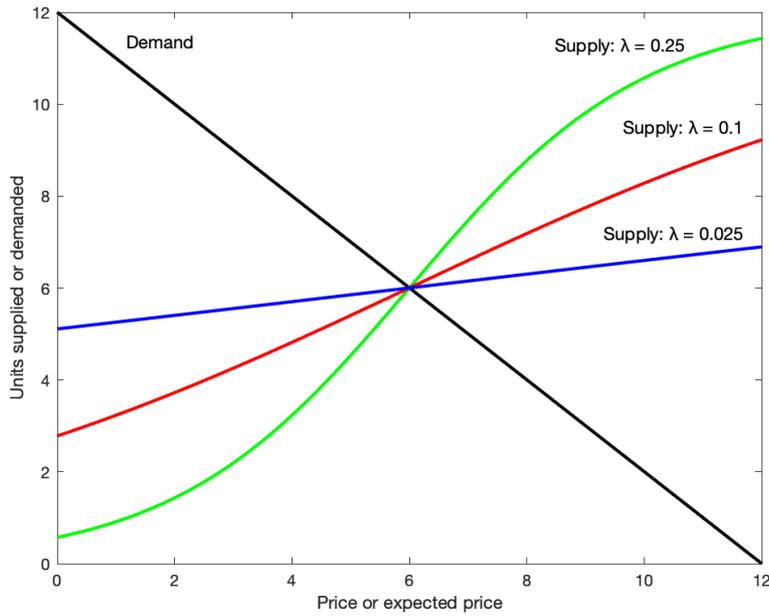
As in the models by Muth and Sarris (and many others) storage plays a central role in connecting the spot and futures markets. We therefore introduce *inventory holders*, next to the *consumers* and *producers* active on the spot market, and the *speculators* active on the futures market. Inventory holders act as arbitrageurs and trade on both markets. Of the four types of agents that we distinguish, only producers and speculators take decisions that depend upon their price expectations, in line with the negative and positive feedback markets described in Sections 2 and 3.

The *consumers'* decisions only depend on the current price in the spot market (and their price expectations therefore do not play a role). Aggregate demand by consumers of the commodity is represented by the linear demand function (2) that we discussed in Section 3, and that we repeat here for convenience:

$$D(p_t) = a - bp_t + \mu_t.$$

Recall that  $a, b > 0$  are strictly positive demand parameters, and that  $\mu_t$  is a normally and independently distributed random series of small demand shocks. We assume that consumers cannot store the commodity themselves in order to consume it in a later period.

The role of the *producers* is also similar to the one they played in the negative feedback market discussed in Section 3. In particular, there are  $K$  producers, each having to decide one period in advance how much of the commodity they are going to supply to the spot market in period  $t$ . The optimal production decision will depend positively on the price the producer will get for its product in the next period. However, this price is not yet known at the time the production decision needs to be taken and therefore each producer chooses the quantity that, given his expectation of the spot market price in period  $t$ ,  $p_{k,t}^{e,p}$ , maximizes his expected profit. We depart from Heemeijer et al. (2009) by assuming that the marginal costs of production are a nonlinear function of the output level, and that this results in an upward sloping S-shaped individual supply curve,  $S_\lambda(p_{k,t}^{e,p})$ , where  $\lambda > 0$  is a parameter that captures the curvature of the supply function. Such a production function implies that individual production cannot increase without bound, which seems to be relevant for many industries. Moreover, it ensures that prices remain bounded when the equilibrium price is unstable.



**Figure 3:** The aggregate supply function, for different values of  $\lambda$ , together with the linear demand function.

Figure 3 depicts the demand function (when the demand shock,  $\mu_t$ , is equal to zero) and the aggregate supply function for different values of the parameter  $\lambda$ .<sup>10</sup> Higher values of this parameter correspond to steeper supply curves, meaning that the price elasticity of supply is higher, and producers adjust their output more aggressively in response to a change in their expectations. When the value of  $\lambda$  passes a certain threshold, the price dynamics under naive expectations (and in the absence of a futures market) changes from oscillations converging to the equilibrium price, to convergence of prices to a stable two-cycle. Note that if both supply and demand functions would have been linear, a stable two-cycle generically does not exist under naive expectations. Instead, high values of  $\lambda$  would

<sup>10</sup> For this figure, and the simulations later in this section, we use the specification  $S_\lambda(p_{k,t}^{e,p}) = c \left( 1 + \tanh \left( \lambda (p_{k,t}^{e,p} - d) \right) \right)$  for the individual supply function, which is also used in the negative feedback LtF experiment discussed in Hommes et al. (2007). The results from that experiment are very much in line with those discussed in Section 3, where a linear supply function was employed. For Figure 3 (and for the simulations later in this section) we use  $a = 12$ ,  $b = c = 1$ ,  $K = d = 6$  and  $\lambda = 0.025$ ,  $\lambda = 0.1$  and  $\lambda = 0.25$ , respectively. Note that, as long as  $a - bd = Kc$ , the rational expectations equilibrium is given by  $p^* = d = 6$ , which will be stable under naive expectations for  $\lambda < b/(Kc) = 1/6$ . At  $\lambda = 1/6$  a stable two-cycle is created, with prices perpetually moving between two values  $p^1$  and  $p^2$ , where  $p^1 < p^* < p^2$  and firms predicting  $p^1$  when the price will be  $p^2$  and the other way around.

give rise to diverging oscillations that grow without bound (until hitting a non-negativity constraint on price or quantity).

On the futures market *speculators* predict the spot market price two periods ahead. Like the investors in the positive feedback LtF experiment discussed in Section 2, they are mean-variance maximizers of their next-period wealth. Therefore, they are willing to take a larger position in the futures market in period  $t$  when the futures market price  $p_{t+1}^f$  (that is, the price they receive in period  $t$  for the contract to deliver one unit of the commodity in period  $t+1$ ) is further away from their spot market price prediction for period  $t+1$ . Their position in the futures markets will also depend on their risk attitude (represented by  $\varphi$ , the coefficient of absolute risk aversion) and on their beliefs about the variance of spot market prices,  $\sigma^2$  (which we assume to be constant). Aggregate demand of the  $H$  speculators can be shown to be equal to:<sup>11</sup>

$$z_t = \frac{H}{\varphi\sigma^2} (\bar{p}_{t+1}^{e,s} - p_{t+1}^f),$$

where  $\bar{p}_{t+1}^{e,s} = \frac{1}{H} \sum_{h=1}^H p_{h,t+1}^{e,s}$  is the average expectation of the speculators about the spot market price in period  $t+1$  (recall that speculators have not yet observed  $p_t$  when forming their prediction for the price in period  $t+1$ ). If the futures market is isolated from the spot market we have  $z_t = 0$ , as the speculators can only trade with each other in that case. When storage is possible, speculators can also trade with inventory holders.

The storage part of the model is based upon Sarris (1984). There exists an optimal level of inventories for the *inventory holders*, and deviations from that optimum are costly (higher interest and storage costs for a positive deviation, and a reduction in convenience yield for a negative deviation). We assume these deviation costs are quadratic. Let  $I_t$  be the deviation of the (representative) inventory holder from his optimal level of inventories. If the current price in the spot market differs from the price on the futures market, the inventory holder can make a risk-free profit that equals

$$\pi_t = (p_{t+1}^f - p_t)I_t - \frac{1}{2}\gamma I_t^2,$$

where  $\gamma > 0$  is a cost parameter. The inventory holder profit from arbitrage is maximized when the inventory deviation  $I_t$  equals

---

<sup>11</sup> The derivation of this demand function is similar to that of the demand by the investors in the model discussed in Section 2 (the derivation of which can be found in Hommes et al., 2008).

$$I_t = \frac{p_{t+1}^f - p_t}{\gamma}.$$

Note that  $I_t$  can be either positive or negative. When the futures price is higher (lower) than the spot market price, the inventory is increased (decreased) and futures are sold (bought). The speculators will take the other side of the market. This means that the influence of the speculators on the spot market price runs through the storage of the commodity: if the speculators predict a high future spot market price they demand more futures, which are supplied by the inventory holders. If the inventory holders supply more futures, they demand more of the commodity in the current period and change their inventory accordingly.

Imposing market clearing on the futures market allows one to solve for the futures market price  $p_{t+1}^f$ , which reveals that the inventory deviation  $I_t$  will be a linear function of the difference between the average price prediction of the speculators and the current spot market price:

$$I_t = A(\bar{p}_{t+1}^{e,s} - p_t), \quad \text{with } A = \left[ \frac{1}{H} \varphi \sigma^2 + \gamma \right]^{-1}.$$

The composite *coupling parameter*  $A$  measures the strength of the connection between the futures market and the spot market. It depends negatively on the costs and possibilities of storage through the cost parameter  $\gamma$ . If deviations in the optimal inventory level are very costly, the futures market will have a limited effect on the spot market, but if it is easy and cheap to deviate from that level, the connection between the two markets will be much stronger. The coupling parameter  $A$  also increases with the size of the futures market (that is, the number of speculators  $H$ ), and decreases with the level of risk aversion of the speculators,  $\varphi$ , and the expected volatility of spot market prices,  $\sigma^2$ . For goods that cannot be stored,  $\gamma$  will be infinite, meaning that  $A = 0$  and the two markets are effectively disconnected. By contrast,  $A \rightarrow \infty$  means that spot market prices are completely determined in the futures market.

The spot market price is determined by equilibrium between aggregate supply of the producers of the commodity and aggregate demand. Here aggregate demand consists of final consumer demand (represented by  $D(p_t)$ ), plus the (positive or negative) change in the inventories of the inventory holders:

$$\sum_{k=1}^K S_{\lambda}(p_{k,t}^{e,p}) = D(p_t) + I_t(\bar{p}_{t+1}^{e,s}, p_t) - I_{t-1}(\bar{p}_t^{e,s}, p_{t-1}).$$

The price that clears the spot market can now be derived as:

$$p_t = \frac{a - \sum_{k=1}^K S_{\lambda}(p_{k,t}^{e,p}) + A\bar{p}_{t+1}^{e,s} - I_{t-1}(\bar{p}_t^{e,s}, p_{t-1})}{A + b} + \frac{\mu_t}{A + b}. \quad (3)$$

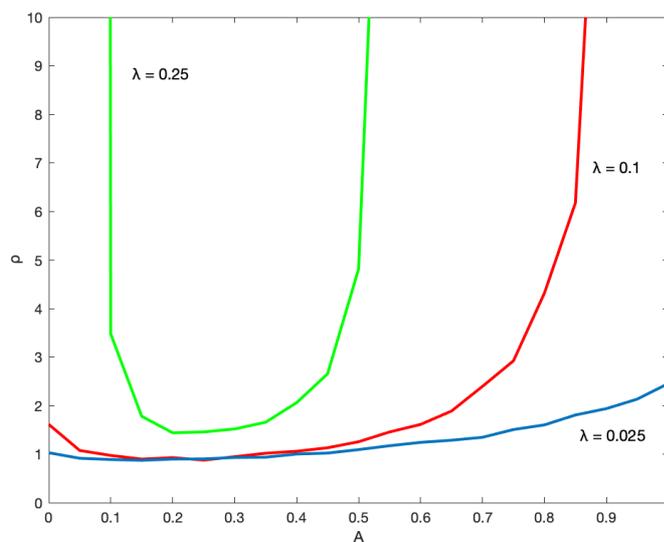
In Equation (3) the demand parameters  $a$ ,  $b$  and the demand shocks  $\mu_t$  come from the consumer demand function (2) and  $S_{\lambda}(p_{k,t}^{e,p})$  is the individual supply of producer  $k$  in period  $t$ , given that he predicts price  $p_{k,t}^{e,p}$  for that period. The final part in the numerator consists of the direct effect of the futures market on the spot market price,  $A\bar{p}_{t+1}^{e,s}$ , which depends upon the average expected price of the speculators for the spot market price in period  $t+1$ ,  $\bar{p}_{t+1}^{e,s}$ , and the indirect effect, which comes via last period's inventory  $I_{t-1}(\bar{p}_t^{e,s}, p_{t-1})$ . Note that, the more the inventory holders stored of the commodity in the previous period, the less they need to buy from the producers on the spot market to satisfy their current demand for the commodity. Also observe that the pricing equation has elements of both negative expectations feedback (through the producers' expectations for period  $t$ ) and positive expectations feedback (through the speculators' expectations for period  $t+1$ ).

We are interested in the effect that the level of connectedness between the futures market and the spot market, represented by the coupling parameter  $A$  and measuring the strength of positive feedback versus that of negative feedback in the model, has on stability of spot market prices. One effect of the coupling parameter  $A$  can be seen in the last part of the pricing equation: A larger value of  $A$  mitigates the effect of the demand shocks  $\mu_t$  on the spot market price. However, when the price expectations of the speculators in the futures market are more volatile, this will also increase the volatility in the spot market through  $A\bar{p}_{t+1}^{e,s}$ .

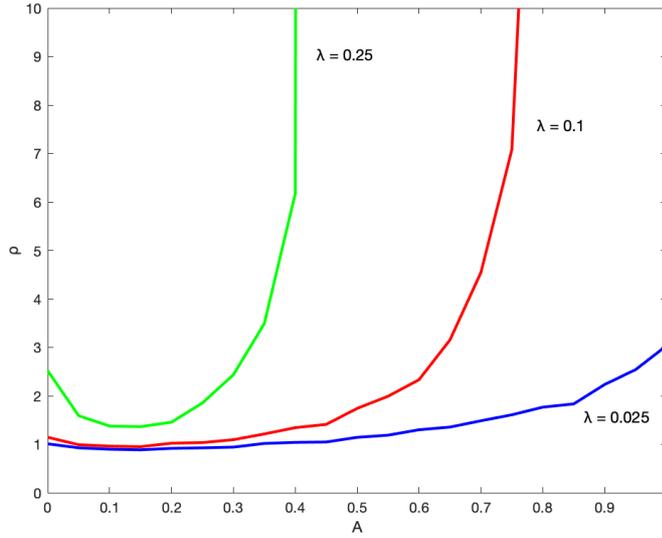
In general, the dynamics in spot market prices is governed by the strength of the connection between the markets,  $A$ , the steepness of the supply function, represented by the parameter  $\lambda$ , and the prediction strategies of the producers and the speculators. Under rational expectations of all traders, only the exogenous demand shocks would cause some variation in the prices. In that case an increase in the connection strength would decrease price volatility since it dampens the effect of the demand shocks. However, from the LtF experiments with positive and negative feedback markets that we discussed in Sections 2

and 3, we know that prediction strategies of human subjects are typically at odds with rational expectations. In order to obtain a good understanding of the type of dynamics that we may encounter in our model we run simulations in which producers use naive or adaptive expectations (as human subjects tend to do in negative feedback LfF experiments) and where speculators are endowed with trend-extrapolating expectations (as almost invariably used by human subjects in positive feedback experiments)

More specifically, for our simulations we use markets with  $K=6$  producers (with S-shaped supply functions, see footnote 10) and  $H=6$  speculators, and set the demand and supply parameters as  $a = 12$ ,  $b = 1$ ,  $c = 1$ , and  $d = 6$ . The spot market equilibrium price (in absence of demand shocks) then equals  $p^* = 6$ , and for the demand shocks we assume  $\mu_t \sim N(0,0.01)$ . For the first set of simulations we assume naive expectations for the producers (that is,  $p_{k,t}^{e,p} = p_{t-1}$  for all  $k$ ) and the following specification of trend-extrapolating expectations for the speculators:  $\bar{p}_{t+1}^{e,s} = p_{t-1} + 0.8(p_{t-1} - p_{t-2})$ . We run each simulation for 10,000 periods, and between simulations we vary the values of the coupling parameter  $A$  and the parameter that regulates the steepness of the supply function,  $\lambda$ . For each simulation we then determine the relative volatility of the spot market prices, as measured by  $\rho$ , which is the variance of spot market prices, divided by the variance of the demand shocks.



**Figure 4:** Simulations with producers with naive expectations and speculators with trend-extrapolating expectations, for different values of  $A$  and  $\lambda$ .



**Figure 5:** Simulations with producers with adaptive expectations and speculators with trend-extrapolating expectations for different values of  $A$  and  $\lambda$ .

Figure 4 shows the results of the simulations, with the coupling parameter  $A$  on the horizontal axis and the volatility measure  $\rho$  on the vertical axis. It follows that, for  $A$  relatively small, an increase in the coupling strength  $A$  leads to a decrease in volatility in spot market prices. However, when the coupling strength is increased further price volatility eventually increases again – suggesting a nonmonotonic effect of the coupling strength on market stability. For Figure 5 we change the expectations of the producers from naive to adaptive. In particular, we assume that each producer  $k$  forms expectations according to  $p_{k,t}^{e,p} = 0.65 p_{t-1} + 0.35 p_{k,t-1}^{e,p}$ .<sup>12</sup> Again, we observe a U-shape in the relationship between the coupling strength  $A$  and price volatility  $\rho$ : The spot market first becomes more stable if the coupling strength between the spot market and the futures market increases, but if that coupling strength increases even more, the spot market is destabilized. Note that the nonmonotonic shape is consistent with the mixed empirical evidence that we discussed above. That is, starting with a small futures market (represented by a small value of  $H$  and consequently a small value of the coupling parameter  $A$ ) an increase in the futures market initially reduces volatility in spot market prices, but if the futures market becomes too large

<sup>12</sup> For this specification we follow Anufriev and Hommes (2012), who use an adaptive expectations rule with a weight equal to  $w = 0.65$  as one of the heuristics in their heuristic switching model. This particular adaptive expectations rule is one of the rules used by participants in the LtF experiment described in Hommes et al. (2005), and the heuristic switching model developed in Anufriev and Hommes (2012) and using that rule is able to explain all of the data from that experiment quite well.

volatility increases again. In our model, informed by the outcomes of positive and negative feedback LtF experiments, this increase in volatility is created by the trend-extrapolating expectations used by the speculators that are active on the futures market.

A recent LtF experiment that also combines positive and negative expectations feedback is Bao and Hommes (2019). However, the model underlying their design is quite different from our model. The price generating mechanism that they use is essentially equivalent with our Equation (1) from above, but with a feedback strength (which is equal to  $1/(1+r) = 0.95$  in the experiment that we discussed in Section 2) that depends on the weight of the suppliers, relative to that of the speculators. Bao and Hommes (2019) run three treatments characterized by different values of this feedback strength, in particular they consider 0.95 (as in Hommes et al., 2008), 0.86 and 0.71, respectively. For the first treatment they find large fluctuations in realized market prices, which is consistent with the results from Hommes et al. (2008). For the second treatment there are still substantial, but smaller, fluctuations in market prices, and for the third treatment market prices typically converge to their fundamental value quickly. Feedback strength therefore plays an important role in the price dynamics.<sup>13</sup> Also note that in the model used in Bao and Hommes (2019) the feedback strength is equal to the slope of the price generating mechanism under naïve expectations.

In contrast to Bao and Hommes (2019), the predictions of suppliers and speculators in our model are not directly aggregated into one aggregate prediction, but play different roles. This implies that, even in the simple case of naïve expectations for both speculators and suppliers, our model gives rise to a two-dimensional dynamical system, where the realized market price in period  $t$ ,  $p_t$ , will depend upon both  $p_{t-1}$  and  $p_{t-2}$ . For small values of  $A$  this system's eigenvalues are real and negative (as in the cobweb model)<sup>14</sup>. As  $A$  increases, the most negative of the two eigenvalues first moves closer to 0, eventually resulting in stable

---

<sup>13</sup>The relationship between feedback strength and price dynamics was also established in Sonnemans and Tuinstra (2010). Two recent LtF experiments, Bao and Zong (2019) and Hennequin and Hommes (2019), use this relationship to study whether monetary policy rules may help in deflating bubbles in asset prices. These studies start from Equation (1) in Section 2, but let the interest rate  $r$  vary over time. In particular, they assume that the monetary policy rule is such that, when the deviation of the market price from its fundamental value becomes too large, the interest rate increases (or decreases, if the deviation is negative). Both studies find that monetary policy rules of this type are quite effective in stabilizing market prices.

<sup>14</sup> For example, when  $A = 0$  the eigenvalues are 0 and -0.15, -0.6, and -1.5 for  $\lambda = 0.025$ ,  $\lambda = 0.1$  and  $\lambda = 0.25$ , respectively.

dynamics under naive expectations. When  $A$  is increased further the eigenvalues become complex with absolute values that increase with  $A$ . So, unlike the situation in Bao and Hommes (2019), these eigenvalues are never real and positive in our model.

## 5. Conclusion

In this chapter we first reviewed the results of Learning to Forecast experiments with either positive expectations feedback or negative expectations feedback. Subsequently, we discussed how the results from these experiments may inform us about the behavior of more complex markets. In particular, we considered a spot market with an associated futures market, leading to a market environment that has elements of both negative feedback (from the expectations of the firms that produce the physical commodity traded on the spot market) and positive feedback (from the expectations of the speculators that are active in the futures market). An important factor is the connection strength (as represented by the coupling parameter  $A$ ), which depends on the size of the futures market (relative to the spot market) and the cost of storage, and which measures the effect that the futures market has on the spot market. We run simulations, with prediction strategies typically encountered in existing positive and negative feedback LtF experiments as an input, and find that, depending on the price sensitivity of the producers' aggregate supply curve, and the strength of the connection between the two markets, connecting the futures market to the spot market will be stabilizing (if the connection is relatively weak) or destabilizing (when the connection is strong, and the price sensitivity of supply is relatively high).

A caveat is in order here. We cannot be sure that participants in these connected markets will use precisely the same type of prediction strategies we see in unconnected markets. For example, it could be that when the spot market is less stable than a stand-alone negative feedback market, producers will use different prediction rules, for example the type of expectations that are typically seen in positive feedback markets. The best way to find out whether our predictions will turn out to be correct, is by running a laboratory LtF experiment with connected markets. Such an experiment is inevitably more complex than one with a single market as it requires participants for at least two different roles: one as advisor to producers and one as advisor to speculators. Both roles require participants to

forecast spot prices, but advice to producers requires one-period-ahead forecasts while two-periods-ahead forecasts are needed for speculators. A first set of results of an experiment like this can be found in de Jong et al. (2019).

Obviously, instead of studying the effect of futures markets on price behavior in an LtF experiment, this effect can also be studied in a laboratory experiment where participants can trade in the relevant commodities themselves (which is sometimes referred to as Learning to Optimize experiments). Indeed, Porter and Smith (1995) and Noussair and Tucker (2006), for example, introduce futures markets in the well-known asset market experiment by Smith et al. (1988). Whereas bubbles and crashes in the asset price are typical in that original experiment (and have been replicated many times), introducing futures markets tends to significantly reduce or even eliminate these bubbles.<sup>15</sup> There are some important differences with respect to the experiment that we suggest (apart from the possibility to trade in the commodity). First, the underlying commodity in these Learning to Optimize experiments is a financial asset as well, meaning that – as opposed to our setting – both the spot and the futures market are characterized by positive expectations feedback and are conducive to speculation. Moreover, because in our design the commodity needs to be stored we can vary the strength of the connection between the two markets. Finally, in Porter and Smith (1995) and Noussair and Tucker (2006) all participants trade on both markets, whereas in our design participants are active either on the spot market or on the futures market. The stabilizing effect in these earlier experiments seems to be due to the fact that participating in the futures markets helps the participants to form better predictions about the spot market price, whereas storage plays a more important role in our experimental design.

There is a tradeoff in designing financial experiments. On one hand, experiments are necessarily limited in complexity and duration. The decision situation should be simple enough for participants to understand after a short instruction and preferably they will also have opportunities for learning during the experiment. On the other hand, the dynamics of interacting markets can be quite complex. The challenge of the designer is to simplify the task of the decision maker while keeping the most interesting aspects of the situation. In LtF

---

<sup>15</sup> In the experiment of Porter and Smith (1995) there is only one futures market (for one period, midway through the asset's maturity) which is sufficient to reduce bubbles, whereas in the experiment of Noussair and Tucker (2006) there is a futures market for each period, which has a much stronger effect on stabilizing prices.

markets participants only make predictions and don't have to worry about trading and the predictions are incentivized in a clear way. This leaves room for studying more complex situations, like described in this chapter. The designer of experiments has to make many decisions: do participants predict prices in only one market or in more markets simultaneously, which roles are played by participants and which by computerized players, what information is available to the participants and in what format, and what is the exact incentive structure? We conclude with the observation that there is a lot to learn from LtF experiments, and that much work is waiting for eager financial experimentalists!

### **Acknowledgements:**

We thank the participants of the Third Virtual Experimental Finance Workshop in January 2021, with special thanks to our discussant Yilong Xu, for their comments on our project.

### **References**

- Anufriev, M. & Hommes, C. (2012). Evolutionary selection and individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics*, 4(4), 35-64.
- Arango, S., & Moxnes, E. (2012). Commodity cycles, a function of market complexity? Extending the cobweb experiment. *Journal of Economic Behavior & Organization*, 84(1), 321-334.
- Bao, T., Duffy, J. & C. Hommes (2013). Learning, forecasting and optimizing: An experimental study. *European Economic Review*, 61, 186-204.
- Bao, T., Hennequin, M., Hommes, C. & Massaro, D. (2020). Coordination on bubbles in large-group asset pricing experiments. *Journal of Economic Dynamics & Control*, 110, Article Number: 103702.
- Bao, T., & Hommes, C. (2019). When speculators meet suppliers: Positive versus negative feedback in experimental housing markets. *Journal of Economic Dynamics & Control*, 107, Article Number: 103730.
- Bao, T., Hommes, C. & Makarewicz, T. (2017). Bubble formation and (in)efficient markets in learning-to-forecast and optimize experiments. *Economic Journal*, 127(605), F581-F609.
- Bao, T., Hommes, C., & Pei, J. (2021). Expectation Formation in Finance and Macroeconomics: A Review of New Experimental Evidence. Working Paper.
- Bao, T. & Zong, J. (2019). The impact of interest rate policy on individual expectations and asset bubbles in experimental markets. *Journal of Economic Dynamics & Control*, 107, 103735.
- Bulow, J.I., Geanakoplos, J.D. & Klemperer, P.D. (1985). Multimarket oligopoly – strategic substitutes and complements. *Journal of Political Economy* 93(3), 488-511.

- Glaser, M., Iliewa, Z., and Weber, M. (2019). Thinking about prices versus thinking about returns in financial markets. *Journal of Finance*, 74(6), 2997–3039.
- Gray, R. W. (1963). Onions revisited, *Journal of Farm Economics*, 45, 273-276.
- Hanaki, N., Akiyama, E. & Ishikawa, R. (2018). Effects of different ways of incentivizing price forecasts on market dynamics and individual decisions in asset market experiments. *Journal of Economic Dynamics & Control*, 88, 51-69.
- Hanaki, N., Koriyama, Y., Sutan, A. & Willinger, M. (2019). The strategic environment effect in beauty contest games. *Games and Economic Behavior*, 113, 587-610.
- Haruvy, E., Lahav, Y. & Noussair, C.N. (2007). Traders' expectations in asset markets: Experimental evidence. *American Economic Review*, 97(5), 1901-1920.
- Heemeijer, P., Hommes, C., Sonnemans, J., & Tuinstra, J. (2009). Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic Dynamics & Control*, 33(5), 1052-1072.
- Hennequin, M. & Hommes, C. (2019). Managing bubbles in experimental markets with monetary policy. *Mimeo*.
- Hommes, C. (2011). The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics & Control*, 35(1):1–24.
- Hommes, C. (1994). Dynamics of the cobweb model with adaptive expectations and nonlinear supply and demand. *Journal of Economic Behavior & Organization*, 24(3), 315-335.
- Hommes, C., Kopányi-Peuker, A. & Sonnemans, J. (2021). Bubbles, crashes and information contagion in large-group asset market experiments. *Experimental Economics*, 24(2), 414-433.
- Hommes, C., J. Sonnemans, J. Tuinstra, & van de Velden, H. (2007). Learning in cobweb experiments. *Macroeconomic Dynamics*, 11, 8–33.
- Hommes, C., Sonnemans, J., Tuinstra, J., & van de Velden, H. (2005). Coordination of expectations in asset pricing experiments. *The Review of Financial Studies*, 18(3), 955-980.
- Hommes, C., Sonnemans, J., Tuinstra, J., & van de Velden, H. (2008). Expectations and bubbles in asset pricing experiments. *Journal of Economic Behavior & Organization*, 67(1), 116-133.
- Jong, J. de, Sonnemans, J. & Tuinstra, J. (2019). The effect of futures markets on the stability of commodity prices. *Tinbergen Institute Discussion Paper 2019-028*.
- Kopányi-Peuker, A. & Weber, M. (2021). Experience does not eliminate bubbles: Experimental evidence. Forthcoming in *The Review of Financial Studies*.
- Marimon, R., Spear, S. E., & Sunder, S. (1993). Expectationally driven market volatility: An experimental study. *Journal of Economic Theory*, 61(1), 74–103.
- Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica*, 29(3), 315-335.
- Nagel, R. (1995). Unravelling in guessing games: An experimental study. *American Economic Review*, 85, 1313-1326.
- Netz, J. S. (1995). The effect of futures markets and corners on storage and spot price Variability. *American Journal of Agricultural Economics*, 77, 182-193.
- Noussair, C. & Tucker, S. (2006). Futures markets and bubble formation in experimental asset markets *Pacific Economic Review*, 11(2), 167-184.
- Powers, M. J. (1970). Does futures trading reduce price fluctuations in the cash markets? *American Economic Review*, 60, 460-464.

- Porter, D.P. & Smith, V.L. (1995). Futures markets and bubble formation in experimental asset markets. *Journal of Business*, 68(4), 509-541.
- Powers, M. J. (1970). Does futures trading reduce price fluctuations in the cash markets? *American Economic Review*, 60, 460-464.
- Sarris, A. H. (1984). Speculative storage, futures markets, and the stability of commodity prices. *Economic Inquiry*, 22(1), 80-97.
- Smith, V., Suchanek, G. & Williams, A. (1988). Bubbles and crashes and endogenous expectations in experimental spot asset markets. *Econometrica* 56, 1119-1151.
- Sonnemans, J., Hommes, C., Tuinstra, J., & van de Velden, H. (2004). The instability of a heterogeneous cobweb economy: a strategy experiment on expectation formation. *Journal of Economic Behavior & Organization*, 54(4), 453-481.
- Sonnemans, J. and Tuinstra, J. (2010). Positive expectations feedback experiments and number guessing games as models of financial markets. *Journal of Economic Psychology*, 31(6), 964–984.
- Tinbergen, J., (1930). Bestimmung und Deutung von Angebotskurven Ein Beispiel. *Zeitschrift für Nationalökonomie* 1, 669–679.
- Working, H. (1960). Price effects of futures trading, Stanford University Food Research Institute.