

The effect of futures markets on the stability of commodity prices

Johan de Jong^{*a,b}, Jan Tuinstra^{a,b}, and Joep Sonnemans^{a,b}

^aAmsterdam School of Economics, University of Amsterdam, The Netherlands

^bTinbergen Institute, The Netherlands

April 8, 2021

Abstract

Do futures markets have a stabilizing or destabilizing effect on commodity prices? The empirical evidence is inconclusive. We try to resolve this question by means of a learning-to-forecast experiment in which a futures market and a spot market are coupled. The strength of the coupling depends positively on the number of speculators on the futures market and negatively on storage costs and speculator risk aversion. We find that the spot price volatility changes non-monotonically with the strength of the coupling, resulting in a stabilizing effect on spot prices for weakly coupled markets and a destabilizing effect when the coupling with the futures market is strong.

JEL classification: D84; G13; G41; Q02

Keywords: Expectations feedback; Futures markets; Price volatility; Experimental economics

Acknowledgements: We thank John Duffy, Myrna Hennequin, Cars Hommes, Harold Houba, Petko Kalev, Anita Kopányi-Peuker, Rosemarie Nagel, Theo Offerman, seminar participants at the University of Amsterdam and the participants at EF Heidelberg 2018, ESA Berlin 2018, TEM Berlin 2018, BEAM-ABEE Amsterdam 2018, QED Lisbon 2019, and the virtual EFW 2021 for helpful comments and discussion. Furthermore, we are grateful for financial support from the Research Priority Area Behavioral Economics of the University of Amsterdam. Jan Tuinstra also acknowledges financial support from Project BEAM (Behavioral and Experimental Analyses in Macro-finance) (NWO 464-15-143), which is partly financed by the Netherlands Organization for Scientific Research (NWO). The laboratory experiment was approved by the Ethics Committee Economics and Business of the University of Amsterdam under number 20180426120405 (April 26, 2018).

^{*}corresponding author: j.a.dejong@uva.nl, PO Box 15867, 1001 NJ Amsterdam, The Netherlands

Speculators may do no harm as bubbles on a steady stream of enterprise. But the position is serious when enterprise becomes the bubble on a whirlpool of speculation. When the capital development of a country becomes a by-product of the activities of a casino, the job is likely to be ill-done. (John Maynard Keynes, *The General Theory of Employment, Interest and Money*, 1936, Chapter 12, section VI)

1 Introduction

Much of the theoretical work on futures markets suggests that they have a stabilizing effect on commodity prices (Friedman, 1953). Yet, in the policy debate there continue to be calls for tighter regulations or even bans on all commodity future trading based on the argument that speculation in the futures market is a source of instability for commodity prices (Kennedy, 2012, April 11). The matter here may not simply be one of who is right. As was already acknowledged by Kaldor (1939) and many authors afterwards, there may be some merit to both claims in the sense that both a stabilizing and a destabilizing effect exist. The question is whether we can identify one effect that will always dominate.

The, at some times, fierce debate has inspired a large body of empirical work on the effect of futures markets on commodity price stability. Worth noting here are the studies using the introduction (or abolishment) of futures markets, such as the ones for onions (Working, 1960; Gray, 1963), pork bellies and beef (Powers, 1970), live cattle (Taylor and Leuthold, 1974), wheat (Netz, 1995), and potatoes (Morgan, 1999). In almost all cases the authors conclude that in the years following the introduction (abolishment) of a futures market, the volatility in the commodity spot prices was lower (higher) than in the years before. These results suggest that a stabilizing effect of futures markets does exist and can be dominant, at least in a certain period following a futures market's introduction. However, there is convincing empirical evidence of a destabilizing effect as well. Roll (1984) argues that a large part of the volatility in the prices of orange juice futures cannot be explained by any changes in real economic variables such as the weather or changes in demand. Because the spot and futures prices of a commodity are typically linked, we expect this excess volatility to spill over into the spot market. One indication that this is happening is provided by Pindyck and Rotemberg (1990) who find a stronger comovement in commodity spot prices than can be explained by macroeconomic variables.

More recently the question of the impact of future markets on commodity price stability again received a lot of attention. Masters and White (2008) put forward the hypothesis that the advent of commodity index investing was responsible for the surge in commodity prices between 2002 and 2008. Although most authors of empirical studies agree that the surge itself was due to other factors, they arrive at different conclusions regarding

the effect of the 'financialization' of commodity markets on price volatility. Du et al. (2011), Algieri (2012), and McPhail et al. (2012) claim that it increased the volatility of commodity prices, while Büyüksahin and Harris (2011), Bohl and Stephan (2013), and Brunetti et al. (2016) do not find evidence of this.

There has been some effort to include destabilizing effects of futures markets in theoretical models. Several authors look at situations of asymmetric information, in which spot market participants take decisions that depend on the futures price. In this case noise or biases injected by speculators on the futures market can adversely affect spot price stability (Stein, 1987; Sockin and Xiong, 2015; Goldstein and Yang, 2019). When there is no information asymmetry, futures markets can still affect spot prices through the storage or risk channels (Cheng and Xiong, 2014). If the first channel is important, shocks to inventory demand (Kawai, 1983) or growth shocks to the economy, as in Dvir and Rogoff (2009), could cause extra volatility in the spot market. An example of the risk channel is presented by Newbery (1987), who argues that futures markets can stimulate risk taking by producers, creating stronger fluctuations in supply and thus increased price volatility. However, these latter models do not address the main concern in the policy debate, which is that speculation on the futures markets itself can be a source of instability. To faithfully model destabilizing speculation in situations without asymmetric information has proven notoriously difficult with rational expectations.¹ At the same time, any other choice of expectations requires a careful justification, since there are many ways to be irrational (Sims, 1980).

In this paper we refrain from explicitly defining how expectations are formed. Instead, we use a learning-to-forecast experiment to elicit the expectations of the agents in a stylized model of coupled commodity spot and futures markets. The participants in the experiment are told that they act as 'advisors' to these agents and are asked to forecast spot prices based on information about past prices. These forecasts are then used in the model to generate a new set of prices on the spot and futures markets upon which the cycle repeats. Our aim in this paper is to investigate how the resulting price dynamics depends on the parameters in the model.

We gain some important insights with this approach. Our main result is that the volatility of the spot prices changes non-monotonically with the parameters in the model. As a consequence futures markets do not always have a stabilizing effect. Neither will they always be destabilizing. When the spot and futures markets are only weakly coupled, for example because it is very costly to store the commodity, there are few speculators, or these speculators are very risk-averse, the stabilizing effect of futures markets dominates

¹Hart and Kreps (1986) show that in a model without lagged production it is possible for speculators to destabilize spot prices under specific circumstances. This may be applicable to shorter timescales, on which producers cannot react to changes in demand. We only consider timescales on which production does play a role.

and spot price stability will increase with increasing coupling strength. However, as the coupling strength becomes stronger this trend reverses and eventually the net effect of the futures market becomes a destabilizing one. In our experimental treatment with the strongest coupling we even observe commodity price bubbles and crashes. None of these markets appear to stabilize towards the end of the experiment.

To understand this result we have to look at the individual decisions in the experiment. When the coupling strength is weak the price dynamics is dominated by negative expectations feedback: the higher the expectations of the price in the next period, the lower the price will be. This is due to the producers in the model, who will increase production when expecting higher prices. This increases supply and subsequently lowers prices. It turns out that under those circumstances the predictions of the participants stay very close to the fundamental price. This results in very stable price dynamics where the only deviations from the fundamental price come from small external demand shocks. Because futures markets help to smooth these external shocks, the price volatility decreases with increased strength of coupling between the spot and futures markets.

When the coupling strength is increased further, for example due to an increasing number of speculators on the futures market, the influence of the negative feedback from expectations diminishes because excess production by producers is more easily absorbed in the inventories. Instead, the expectations of the speculators on the futures market become more important. Unlike the producers' expectations, the expectations of the speculators have a positive impact on prices. In first instance it will affect the futures prices, which will be raised if the speculators adjust their expectations upwards. However, higher futures prices mean that buying the product in the next period becomes more expensive. This will induce inventory holders to buy more of the product on the spot market in the current period and increase the amounts they have in storage. The end result is that higher speculator expectations also raise prices on the spot market, a form of positive expectations feedback. When the positive feedback strongly dominates, the participants in the experiment do not coordinate their expectations on the fundamental price. Instead, their expectations can be best described as trend-following and it results in strongly fluctuating spot prices. This tendency to extrapolate past returns has also been found in investor surveys (Greenwood and Shleifer, 2014; Amromin and Sharpe, 2014) and gives rise to asset market bubbles when incorporated into (standard) models (Adam et al., 2017; Barberis et al., 2018).

Our results for the experimental markets with the weakest and strongest coupling are in line with previous results in the learning-to-forecast literature. Hommes et al. (2007) perform an experimental test of the cobweb theorem, in which the feedback from expectations to prices is purely negative. They find that many participants are able to learn the fundamental price and that their expectations can be classified as (close to)

rational. On the other side of the spectrum are a number of asset market experiments in which the expectations feedback is purely positive. Hommes et al. (2005, 2008) observe frequent bubbles and crashes in experimental markets under those circumstances. The positive feedback does need to be stronger though than about $2/3$ in order to observe permanently unstable prices (Sonnemans and Tuinstra, 2010; Bao and Hommes, 2019).² The work by Heemeijer et al. (2009) confirms that the striking difference in results of the different learning-to-forecast experiments is indeed due to the different direction of expectations feedback.

We are not the first to employ the laboratory experiment as a tool to study the effect of futures markets. Forsythe et al. (1984) and Friedman et al. (1984) use a design with assets that live for two or three periods and can be traded in each period through a continuous double auction. A futures market for the asset in the last period of its life significantly increases the speed at which the participants converge to coordination on the rational expectations equilibrium. Porter and Smith (1995) introduce a futures market in a call market experiment and conclude that the bubbles are reduced compared to the treatment without a futures market. However, in a survey paper Palan (2013) questions the validity of this result, also because in a replication with digital options instead of futures contracts the effect is not found (Palan, 2010). Finally, Noussair and Tucker (2006), in a call market experiment, open futures markets for each period before the spot market opens. This forces the participants to consider future prices of the commodity before trading and the authors find that this eliminates bubbles completely. There thus appears to be some evidence for a stabilizing effect in these experiments, but overall the results are still ambiguous. Note also that the stabilizing effect in these experiments is based on the information channel and not related to storage as in our experiment.

The rest of this paper is organized as follows. In the next section (Section 2) we introduce the model. This is followed by the experimental design in Section 3, the results of the experiment in Section 4, and the conclusion in Section 5.

2 Model of coupled spot-futures markets

The model used in this paper shares its essence with those of Muth (1961) and Sarris (1984), who also modeled the effect of speculation on the price of a storable commodity. It features four types of agents: producers, consumers, speculators, and inventory holders. The purpose behind this is not to model four different types of people, but to model four different types of roles. Combining some of these roles in a single decision maker, for example a producer who also keeps inventories or an inventory holder who also chooses

²This means that the deviation of the realized price from the fundamental price is two-thirds of the deviation of the average expectation from the fundamental price.

to speculate, would not alter the analysis below. The separation here is for analytical convenience only. Each period the agents in the model interact with each other on a spot and a futures market. They take decisions that maximize their expected (risk-adjusted) profits, given the information available and, in case of the producers and speculators, given their expectations about prices in the future.

2.1 The spot market

The main actors on the spot market are producers and consumers. There are K producers and they all need one period to produce the commodity. Therefore an individual producer needs to decide one period in advance how much to produce, before knowing the price she will get for her product. She chooses optimally, maximizing expected profit given her expectation of next period's price $p_{k,t}^{e,p}$ and some non-linear cost function, resulting in an S-shaped individual supply curve:

$$S_{k,t}(p_{k,t}^{e,p}) = c(1 + \tanh(\lambda(p_{k,t}^{e,p} - d))), \quad c, d, \lambda > 0. \quad (1)$$

The parameter λ , which is linked to the steepness of the supply curve, can be tuned to have either a locally stable fundamental price or a locally stable two-cycle under naive expectations (Hommes, 1994). There are two reasons to choose a supply curve of this form. First of all, it ensures that supply is bounded, which is a realistic feature for short-term production decisions (see Hommes (1994) for a more extensive justification). Second, when combined with a downward sloping demand curve, the S-shaped supply curve can give rise to a stable two-cycle in which prices alternate between high and low indefinitely under naive or adaptive expectations. It is worthwhile to allow for this possibility in the experiment.

The combined supply of the producers meets a consumer demand that linearly decreases with price p_t :

$$D_t(p_t) = a - bp_t + \epsilon_t, \quad a, b > 0, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2). \quad (2)$$

The ϵ_t represent small, independent demand shocks.³ Except for being a realistic feature of commodity markets, the shocks are a crucial component of the model. The shocks provide a basic level of volatility, without which there would be nothing to stabilize. We need them to discern a possibly stabilizing effect of futures markets.

³In principle it would have also been possible to independently add aggregate supply shocks. However, in this case having a single or multiple independent shocks in one period does not change the price dynamics.

2.2 The futures market

The futures market features H speculators speculating on the spot price of the commodity in the next period by buying and selling futures. At the moment they trade, the current period's spot price is not yet known and therefore their decisions are based on two-period ahead forecasts by their advisors. In fact, the current period's spot price will depend on the speculators' decisions.

A speculator's only concern is her next period wealth $W_{h,t+1}$, which depends on prices and the quantity $z_{h,t}$ of the product that her position in futures represents:

$$E_{ht} [W_{h,t+1}] = W_{h,t} + \left(p_{h,t+1}^{e,s} - p_{t+1}^f \right) z_{h,t}. \quad (3)$$

$z_{h,t}$ may be positive or negative, depending on whether the speculator's position is long or short, respectively. In Eq. (3) $p_{h,t+1}^{e,s}$ is the speculator's prediction of next period's spot price and p_{t+1}^f is the price in the futures contract. Note that this is the current futures price, and known with certainty to the speculators at the moment of trade, but it is denoted with the subscript $t+1$ to emphasize that it is the price the speculator agrees to pay or be paid in the next period.

One of the challenges in trading commodity futures is that commodity prices may be excessively volatile and speculators must take this risk into account. We assume that they will do this by mean-variance maximization of their next period wealth $W_{h,t+1}$:

$$\begin{aligned} \max_{z_{h,t}} \left\{ E_{ht} [W_{h,t+1}] - \frac{\phi}{2} \text{Var}_{ht} [W_{h,t+1}] \right\} \\ = \max_{z_{h,t}} \left\{ W_{h,t} + \left(p_{h,t+1}^{e,s} - p_{t+1}^f \right) z_{h,t} - \frac{\phi}{2} z_{h,t}^2 \text{Var}_{ht} [p_{t+1}] \right\}. \end{aligned} \quad (4)$$

Here ϕ is the coefficient of absolute risk aversion.⁴ As in the asset market models by Brock and Hommes (1998) and Hommes et al. (2005, 2008), we make the additional assumption that the speculators' perception of the price volatility remains constant:

$$\text{Var}_{ht} [p_{t+1}] = \sigma^2. \quad (5)$$

One way to interpret this assumption is that the speculators have a good sense of the volatility they should expect because they know how the prices of the commodity varied historically. This leaves them with little reason to update as long as the parameters of the system do not change.

⁴It can be shown that when returns are normally distributed this mean-variance maximization is exactly the one performed by rational agents possessing a constant absolute risk aversion (CARA) utility function (e.g. Sargent (1987)).

The position that speculator h takes, follows from carrying out the maximization in Eq. (4):

$$z_{h,t} = \frac{1}{\phi\sigma^2} \left(p_{h,t+1}^{e,s} - p_{t+1}^f \right). \quad (6)$$

Total demand z_t by all speculators combined is then:

$$z_t = \frac{H}{\phi\sigma^2} \left(\bar{p}_{t+1}^{e,s} - p_{t+1}^f \right), \quad (7)$$

with $\bar{p}_{t+1}^{e,s}$ the speculators' average prediction of the next period's spot price. The result in Eq. (7) is intuitive. The speculators' aggregate demand for futures increases with the difference between their average expectation of the next period's price and the futures price. This difference is the risk premium responsible for the expected profit of the speculators. Also, the absolute value of z increases with the number of speculators active in the market and decreases with increased risk aversion (higher ϕ) and higher price volatility in the futures market (higher σ^2).

On an isolated futures market the speculators can only trade with each other ($z_t = 0$). As a result, the futures price is exactly equal to the speculators' average expectation of the next period's spot price and the risk premium vanishes. Without the prospect of profits for speculators, it is questionable whether a futures market can continue to operate. Some form of coupling with the corresponding spot market is therefore crucial for a futures market.

2.3 Market coupling by inventory holders

For many market participants keeping inventories is an essential part of their operations, for example because they sell out of inventory or in order to guarantee continuous operation of a production process. In the absence of speculation inventory holders will balance this convenience yield with interest and storage costs to find the optimal working inventory. However, the availability of futures contracts invites inventory holders to engage in arbitrage. When the futures price is high compared to the spot price, it is profitable to take a short position on the futures market and at the same time buy some of the commodity on the spot market to store it for one period. In the reverse situation, inventory holders can combine selling some of their inventory on the spot market with a long position on the futures market. Although keeping a lower than optimal inventory costs them in terms of convenience yield, it guarantees them a restock in the next period at a very favorable price. Note that for this to be pure arbitrage, the amount bought or sold on the spot market must exactly match the position taken in the futures market.

In practice the costs of keeping inventories are often (partially) determined by supply and demand on a market for storage. To model such a market in detail is beyond the

scope of this paper. Therefore, as in Sarris (1984), we introduce a single representative inventory holder, whose costs increase quadratically for any deviation from optimal working inventories. In case of a positive deviation this is due to rising interest and storage costs (including possible loss of quality of the commodity under storage), while for negative deviations the increase stems from a reduction of the convenience yield. We write the inventory holder's profit from arbitrage as:

$$\pi_t = \left(p_{t+1}^f - p_t \right) I_t - \frac{\gamma}{2} I_t^2, \quad (8)$$

where I_t is the aggregate deviation from optimal inventories and γ is a cost parameter which is higher when the costs associated with inventory deviation are higher. Maximization of this profit yields the total inventory deviation in period t :

$$I_t = \frac{p_{t+1}^f - p_t}{\gamma}. \quad (9)$$

The inventory deviation I_t (which may be negative as long as total inventories are positive) is the quantity of the commodity that the inventory holders demand on the spot market and, because they engage in pure arbitrage, supply on the futures market. Regardless of whether this supply is positive (short position) or negative (long position), the speculators will automatically take the other side of the market. Therefore

$$z_t = I_t. \quad (10)$$

Combining Eqs. (7), (9), and (10) leads to the following relation between spot prices, futures prices, and the speculators' expectations:

$$p_{t+1}^f = \frac{p_t + \frac{\gamma H}{\phi \sigma^2} \bar{p}_{t+1}^{e,s}}{1 + \frac{\gamma H}{\phi \sigma^2}}. \quad (11)$$

Using this result in Eq. (9) the inventory deviation becomes:

$$I_t = \frac{\bar{p}_{t+1}^{e,s} - p_t}{\frac{\phi \sigma^2}{H} + \gamma} = A \left(\bar{p}_{t+1}^{e,s} - p_t \right), \quad (12)$$

where

$$A = \frac{1}{\frac{\phi \sigma^2}{H} + \gamma}. \quad (13)$$

In the end the expression for the inventory deviation takes a simple form: it changes linearly with the difference between the speculators' expected price for the next period and the current period spot price. Muth (1961) and Sarris (1984) arrive at a similar

expression for the inventory deviation under speculative storage. Note that all model parameters related to the coupling between the spot and futures markets combine into one parameter: A . The size of A determines to what extent inventories respond to differences between expectations of prices in the next period and current spot prices and can therefore be interpreted as the strength of the coupling between the two markets. As expected, the coupling strength increases when storage costs go down (lower γ). However, increased speculative demand due to a larger number of speculators or decreased risk aversion can also strengthen the coupling.⁵ Financialization is therefore one of the processes that lead to an increase in A in this model.

The actions of the inventory holders affect the situation on the spot market. A change in inventories generates extra demand or supply, pushing spot prices up and down, respectively. The market clearing condition for the spot market becomes:

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

Using Eqs. (1), (2), and (9) in the market clearing condition and solving for the spot price yields the price equation:

$$p_t = \frac{a - cK - c \sum_{k=1}^K \tanh(\lambda(p_{k,t}^{e,p} - d)) + A\bar{p}_{t+1}^{e,s} - I_{t-1}(\bar{p}_t^{e,s}, p_{t-1})}{A + b} + \frac{\epsilon_t}{A + b}. \quad (14)$$

The structure of the equation above reveals the properties of the system. As one would expect, p_t decreases when the producers expect higher prices (negative feedback) and increases when the speculators expect higher prices (positive feedback). The positive feedback is moderated by the coupling strength A and can vary from being absent ($A = 0$) to completely dominating ($A \rightarrow \infty$). The system also contains a state variable in the form of the previous period inventory deviation $I_{t-1}(\bar{p}_t^{e,s}, p_{t-1})$. Naturally, if a large inventory has been build up, this has a negative effect on spot prices and vice versa for a large shortage. Finally, the coupling between the markets weakens the effect of the extrinsic demand shocks on the spot price. This is the stabilizing effect of futures markets.

⁵Note that A is also affected by the volatility of the spot prices via σ^2 . If the volatility changes with the strength of the coupling, then any attempt to change the coupling strength by changing one of the parameters ϕ , γ , or H will be either amplified or dampened by the accompanying change in σ^2 (the direction of the change will always stay the same). This is not a concern for our experiment, which is run for fixed values of A . However, in cases where one wants to investigate the dynamic response to a change in one of the parameters, it does need to be taken into account.

2.4 Limit cases

The limiting cases $A = 0$ and $A \rightarrow \infty$ of this model deserve some extra attention. For $A = 0$ the price equation reduces to the one for an isolated spot market:

$$\lim_{A \downarrow 0} p_t = \frac{a - cK - c \sum_{k=1}^K \tanh(\lambda(p_{k,t}^{e,p} - d))}{b} + \frac{\epsilon_t}{b}. \quad (15)$$

Without a coupling to the futures market, there is no role for inventories and no positive feedback from expectations. Also, other than the coupled system, the isolated spot market features a unique⁶ rational expectations equilibrium (REE), in which the producers' expectations are equal to the expected value of p_t . Unfortunately the REE price p^* cannot be written in closed form. However, despite of this, participants in the cobweb experiment by Hommes et al. (2007), which essentially uses Eq. (15) as the underlying equation, manage to find the REE price and coordinate on it.

When $A \rightarrow \infty$ the spot prices are fully determined by the expectations of the speculators:

$$\lim_{A \rightarrow \infty} p_t = \bar{p}_{t+1}^{e,s} - \bar{p}_t^{e,s} + p_{t-1}. \quad (16)$$

In this limit case there is no negative feedback from producers' expectations and no effect of extrinsic demand shocks on prices. The structure of the equation is such that the change in the spot price exactly equals the change in the expectation for the *next* period spot price by the speculators. An important consequence of this is that if speculators adopt trend-following expectations, price bubbles can continue to grow forever.

3 Experimental design

To elicit the expectations for our model we designed a laboratory experiment. In the instructions the participants are told that they will take up the role of advisor to either a producer or a speculator. Their task is to provide price forecasts, one period ahead if working for a producer or two periods ahead if working for a speculator. The more accurate their prediction is, the more points they earn. The participants learn about their exact role (i.e. working for a producer or a speculator) just before the start of the actual experiment.

Each instance of coupled experimental spot and futures markets consists of four advisors to producers and four advisors to speculators (the exact numbers are not known to the participants). Each advisor is coupled to exactly one producer or speculator and each producer or speculator is coupled to exactly one advisor. Producers and speculators

⁶This is due to the fact that a strictly increasing (supply) function and strictly decreasing (demand) function intersect exactly once.

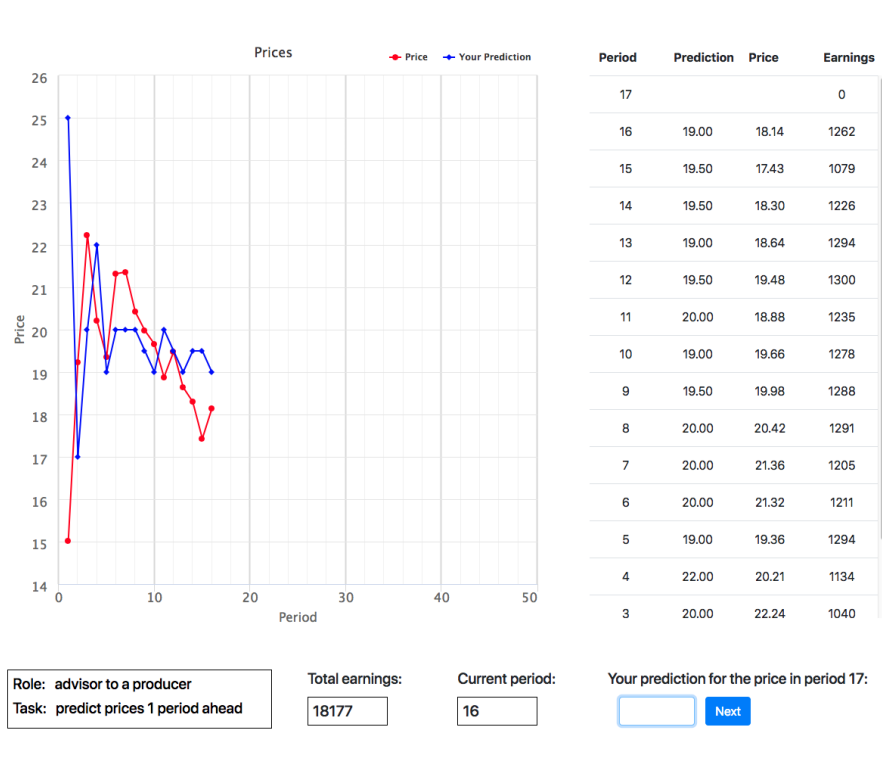


Figure 1: Screenshot of the interface. Past spot prices and the individual’s own past predictions are available both in graph and in table form. In the lower left corner the participants are reminded of their task (one-period-ahead or two-period-ahead prediction), their earnings up to that point, and the current period.

take the predictions of their advisors as their expectations without modification. Their decisions, and the decisions of the other agents in the model, are completely automated and optimal according to the model specification. When all participants in a session have submitted their forecasts, the computer calculates the spot price in the next period according to Eq. (14) and updates the participants’ screens with the new price. The experimental markets always last 50 periods.

To help the participants in making their forecasts, they are provided with a history of all spot prices up to the current period as well as their previous predictions. They get this information both in graph and in table form (see Fig. 1 for a screenshot of the interface). At the start of the experiment, when no past prices are available, the participants receive a hint that the spot price will most likely be between 8 and 50 in the first period (the full range of spot prices is from 0 to 1000). Also, in the instructions we include a qualitative description of the economy, emphasizing the role of producers and speculators and the negative and positive feedback from expectations (see Appendix A). Participants can only proceed to the experiment after answering a few questions about these concepts correctly. The exact equations of the model are not shared with the participants, nor are they given the commodity’s fundamental price.

During the experiment the participants can only see spot prices and the reason for this is twofold. First of all, we want to eliminate the possibility that participants confuse the spot and futures price information and by accident forecast futures prices instead of spot prices. Second, showing futures prices to the participants would open up an extra channel through which futures markets can influence spot prices, namely by affecting the forecasts by the producers' advisors. Including them immediately would make it more difficult to interpret the results. We consider it an interesting direction for future research.

Payment of the participants is based entirely on the accuracy of their predictions, using the function:

$$earnings_t = \max \left\{ 1300 - \frac{1300}{25} error_t^2, 0 \right\}, \quad (17)$$

where the error is the absolute difference between the realized and predicted prices in period t . At the end of the experiment, the total earnings in points are converted into euros, at a rate 1 euro for 2000 points.

The parameters of the model are chosen such that the fundamental price in the experiment is not a round number, in this case $p^* = 20.7$. This is achieved by taking $a = 41.4$, $b = 1$, $c = 5.175$, and $d = 20.7$. The slope parameter in the individual supply functions is $\lambda = 0.0725$, just large enough to make the price dynamics of an isolated spot market unstable under naive expectations.⁷ For the levels of optimal inventories we take high values, such that it is unlikely that the inventories ever reach zero during the experiment. Finally, we set $\sigma_\epsilon = 1.5$ as the standard deviation of the independent demand shocks. We have four treatments, each corresponding to a different strength of coupling between the spot and futures markets: almost isolated markets ($A = 0.01$), weakly coupled markets ($A = 0.5$), strongly coupled markets ($A = 10$), and very strongly coupled markets ($A = 30$).

The main question is whether futures markets always have a stabilizing effect on spot prices or whether they can also be a destabilizing force. To that end we will compare the standard deviations of the spot prices in the different treatments. For each pair of treatments, the null hypothesis is that the standard deviations of the spot prices are the same in each treatment. The alternative hypothesis is that they are not. In the case that we find significant differences, we are particularly interested to know whether the stability changes monotonically with coupling strength A or not. In Appendix B we show simulated results for a few cases in which all participants use the same prediction strategy. Naive and rational expectations lead to a monotonic decrease of price fluctuations with A , while in the case of trend-following expectations the fluctuations first decrease and

⁷The eigenvalue of the system with isolated spot and futures markets under naive expectations is equal to 1.50075 (see Appendix B)

Table 1: Standard deviations of spot prices

	All periods				Last 25 periods			
	$A = 0.01$	$A = 0.5$	$A = 10$	$A = 30$	$A = 0.01$	$A = 0.5$	$A = 10$	$A = 30$
Group 1	3.11	1.10	9.15	6.20	1.65	1.03	4.39	4.16
Group 2	2.01	1.08	10.57	3.15	1.40	0.78	9.85	2.07
Group 3	2.14	1.52	33.38	40.98	1.71	1.52	2.40	48.18
Group 4	2.45	1.30	17.93	223.68	1.77	1.10	5.45	216.94
Group 5	2.23	1.25	4.77	23.16	1.54	0.95	0.82	27.38
Group 6	2.35	1.51	1.00	2.95	1.69	1.37	0.56	3.17
Group 7	2.19	1.39	4.05	12.48	1.60	1.40	1.98	13.14
Group 8	3.02	1.53	1.64	75.03	1.77	1.18	0.63	90.26
Mean	2.44	1.34	10.31	48.45	1.64	1.17	3.26	50.66

then increase for larger values of A (a U-shape).

Our experiment was programmed in oTree (Chen et al., 2016) and run at the CREED laboratory of the University of Amsterdam in May and June 2018. 256 people participated, the majority being undergraduate students in Economics and Business Economics at the University of Amsterdam (58%), and most others were either enrolled in another social science program (18%), or studied law (11%). 44% of the participants was male and the average age was 21.8 years. Sessions lasted approximately 1.5 hours during which the participants earned on average €28.78.

4 Results

Fig. 2(a) shows semi-log plots of the spot prices in each period for the treatment in which the coupling between the spot and futures markets is negligible ($A = 0.01$). The results are in line with earlier learning-to-forecast experiments with isolated spot markets by Hommes et al. (2007) and Heemeijer et al. (2009). In each experimental spot market we find prices fluctuating around the fundamental value of 20.7. The first 5 to 10 periods are characterized by larger deviations as in this stage the participants are still learning the fundamental price (see the analysis by Heemeijer et al. (2009) and the enlarged plots including individual forecasts in Appendix C). Afterwards the price fluctuations stay at a similar level. This level lies, considering that the average standard deviation of prices in the last 25 periods is 1.64 (Table 1, 6th column), only slightly above the level of the external noise process ($\sigma_\epsilon = 1.5$).

Compared to the almost-isolated markets, the treatment with weak coupling ($A = 0.5$)




Figure 2: Semi-log plots of prices in different experimental spot markets. There are four treatments, differing only in the strength of the coupling between the spot and futures markets: $A = 0.01$ (a), $A = 0.5$ (b), $A = 10$ (c), $A = 30$ (d). The graphs in the lower two subfigures have larger scales than the graphs in the upper two subfigures.

changes the situation in two ways. On the one hand increased flexibility of inventories effectively reduces the external demand shocks by a factor of 1.5. On the other hand, through the market coupling the spot prices become subject to positive expectations feedback, which may make it harder for the participants to find the fundamental price and coordinate their expectations on it. Fig. 2(b) shows that at coupling strength $A = 0.5$ there is no sign of the latter. As in the almost-isolated treatment the spot prices fluctuate around the fundamental value at a slightly higher level than that caused by the external demand shocks (1.17 instead of 1.00, see Table 1, 7th column). However, because the external demand shocks are now effectively smaller, prices are on average considerably more stable in this treatment.

When the coupling strength is further increased to $A = 10$, spot prices follow a very different pattern (see Fig. 2(c)). Instead of seemingly random fluctuations, they move in multi-period cycles around the fundamental value. The amplitudes of these cycles differ in different experimental spot markets, but they are in general much larger than those of the fluctuations in the first two treatments (note the change in scale). Compared to the cycles that have been observed in learning-to-forecast experiments with asset markets, the price dynamics in this treatment stands out in two ways. First of all, the length of the cycles is not stable, but seems to decrease. Second, the amplitude of the cycles decreases over time. As a result the participants eventually learn the fundamental price (after about 30 to 40 periods) and start to coordinate their expectations on it. A learning time of 30 to 40 periods is quite remarkable for a learning-to-forecast experiment, in which participants usually learn to forecast the fundamental price within the first 10 periods or not at all.

Figure 2(d) shows the spot prices in experimental markets with the strongest coupling: $A = 30$. In this treatment the spot prices are almost completely determined by the speculators on the futures market. The observed dynamics varies. In some markets the spot prices move in multi-period cycles around the fundamental value until the end of the experiment, while some other markets are characterized by large bubbles and crashes. One feature that distinguishes these markets from the experimental markets in the treatment with $A = 10$ is that in none of the markets spot prices stabilize during the experiment. Even in the first market, where prices seemed to have stabilized after 40 periods, a new cycle or bubble is forming in the last few periods. In many aspects the observed price dynamics is similar to that in learning-to-forecast experiments with asset markets (Hommes et al., 2008; Heemeijer et al., 2009). Note that for the participants, who are being paid based on the accuracy of their predictions, this is very unfavorable. Considering that we observe a few attempts at market manipulation in this strongest coupling treatment, some participants are also aware of this. However, despite the incentive the participants are not able to coordinate on another prediction strategy.

To assess the effect of the futures market on the stability of spot prices, we calculated




Figure 3: Bar plots showing for each treatment the mean of the standard deviations of spot prices in all periods (a) or the last 25 periods (b). In each plot we also indicated between which pairs of treatments the difference in standard deviation is significant (Mann-Whitney U-test, two-sided).

the standard deviation of the prices in each experimental spot market. The results for all 32 spot markets, 8 per treatment, are presented in Table 1 and the mean standard deviations for each treatment are also plotted in Fig. 3. Note that the mean standard deviation changes non-monotonically with the coupling strength: between $A = 0.01$ and $A = 0.5$ it decreases but then it increases between $A = 0.5$ and $A = 10$ and between $A = 10$ and $A = 30$. We used the Mann-Whitney U-test (two-sided) to determine between which of these treatment pairs the difference with respect to the spot price standard deviation is significant. The results of these tests are also depicted in Fig. 3. Because both the decrease of the standard deviation between $A = 0.01$ and $A = 0.5$ and the increase between $A = 0.5$ and $A = 10$ are statistically significant (with $p < 0.001$ and $p = 0.010$, respectively), we conclude that there is a U-shaped dependence of the spot price volatility on the strength of the coupling between spot and futures markets. Moreover, the overall increase in spot price standard deviation between $A = 0.01$ and $A = 30$ is also significant ($p < 0.001$). This suggests that futures markets have a stabilizing effect on spot prices when they are weakly coupled to the spot markets and that they have a destabilizing effect when the coupling strength increases beyond a certain point.

It is not obvious that the outcome of the analysis above remains unchanged when learning effects are taken into account. We noted before that in treatments with low values of A participants learn the fundamental price in the first 5 to 10 periods, resulting in a lower spot price volatility afterwards. Judging from Fig. 2 this effect might be stronger for $A = 0.01$ than for $A = 0.5$, which makes it necessary to check whether the outcome changes when we exclude the beginning of the experiment. Another potential problem arises because of learning in the strong coupling treatment ($A = 10$). Towards the end of the experiment prices are quite stable in every $A = 10$ spot market. It is likely that for the last periods those markets are not significantly more volatile than the $A = 0.5$ markets. To still get a significant U-shape, the increase in mean standard deviation between $A = 10$ and $A = 30$ will need to be significant for those periods. In the end, we address both problems by performing the same analysis on only the last 25 periods. Fig. 3(b) shows for each treatment the mean standard deviation of the last 25 prices in each spot market. Also here we find that the dependence of mean standard deviation on the coupling strength is U-shaped. The decrease from $A = 0.01$ to $A = 0.5$ is still significant ($p < 0.001$). As expected this is not the case anymore for the increase between $A = 0.5$ and $A = 10$. However, now the increase between $A = 10$ and $A = 30$ is significant ($p = 0.021$). Therefore we arrive at the same conclusion: the dependence of the spot price volatility on the coupling strength is not monotonic, but U-shaped.

The origin of the left, downward-sloping part of the U-curve is clear. It is driven by increased flexibility of inventories, which mitigates the effect of external shocks in supply and demand on the spot price. How this trend is broken, cannot be explained well with the analysis done so far. Simulations with very simple forecasting rules (Appendix B) reveal that the use of a trend-following heuristic by participants can lead to a U-shape. However, this heuristic also causes extremely volatile spot prices in almost isolated spot markets, which contradicts our observations in treatment $A = 0.01$. One possibility, that would reconcile the use of simple heuristics with our observations, is that participants use different heuristics in different treatments.

Heemeijer et al. (2009) analyze the individual forecasting strategies of participants in LtF-experiments with both positive and negative expectations feedback markets. They find that in their experiments more than half of the participants use a prediction strategy equivalent to a heuristic of the form:

$$p_t^e = \alpha_1 p_{t-1} + \alpha_2 p_{t-1}^e + (1 - \alpha_1 - \alpha_2) p^* + \beta (p_{t-1} - p_{t-2}). \quad (18)$$

According to this equation participants' expectations are a weighted average of the last spot price p_{t-1} , the last prediction p_{t-1}^e , and the fundamental price p^* complemented by a trend term with coefficient β . Interestingly, these are also the four elements most often

Table 2: Significance and mean values of trend-following parameters in individual forecast analysis. The percentage expresses the fraction of participants for whose predictions the trend-following parameter is significant. The line below shows the mean trend-following parameter (averaged only over those cases in which the parameter was significant).

	Producers				Speculators			
	0.01	0.5	10	30	0.01	0.5	10	30
Significant trend parameter	28%	34%	69%	88%	28%	13%	59%	75%
Mean value	-0.07	0.21	0.54	0.63	0.02	0.16	0.74	0.80

mentioned by our participants when asked about their strategy in the questionnaire at the end of the experiment. Therefore, we estimated Eq. (18) for all forecasts of advisors to producers in our experiment and a similar one (with p_{t+1}^e and p_t^e instead of p_t^e and p_{t-1}^e , respectively) for forecasts of advisors to speculators. In case any of the p-values were above 0.05, the parameter with the highest p-value was removed and the equation was then estimated again. This was repeated until all remaining parameters were significant at the 5% level. The first 10 periods, which we consider as a learning phase, and predictions of 0 or 1000, which form the boundaries of the acceptable price range, are not included in the estimation.

Table 2 shows for each treatment the fraction of participants for whose predictions the trend-following parameter β is significant, as well as the mean value. For the almost-isolated markets ($A = 0.01$) we observe only a limited number of strategies with significant trend-following (28%). Moreover, because there are also some cases for which the parameter is negative, the mean value of the parameter is almost 0. It therefore seems that in this treatment trend-following behavior does not play a role in the price dynamics. For $A = 0.5$ the fraction of participants that demonstrably uses trend-following is also small. By contrast, in treatments with strong market coupling ($A = 10$ and $A = 30$) more than half the participants have trend-following parameters significantly different from 0 (more than 75% for $A = 30$) and the mean value of the parameter ranges from 0.54 for producers in $A = 10$ to 0.80 for speculators in $A = 30$. Clearly, in these markets trend-following does play a role in forecasting spot prices and increasingly so when the coupling strength increases. This offers an explanation for the high volatility of spot prices in strongly coupled markets.

5 Conclusion

We have studied the effect of futures markets on the stability of commodity spot prices using a stylized model of coupled spot and futures markets and a learning-to-forecast experiment. In the experiment half of the participants are asked to forecast prices for producers, who need to decide in advance how much to produce for the next period. The higher the price the producers expect to get for their product, the more they will produce and the lower the actual price will be (negative expectations feedback). The other half of the participants forecast prices for speculators on the futures market. In contrast to the producers, if speculators have higher expectations of spot prices in the next period, futures prices will rise, and, if the spot and futures markets are coupled, so will the spot prices (positive expectations feedback). A central role is played by the coupling between the spot and the futures market, which is based on storage. The stronger this coupling, the larger the influence of the futures market on spot prices. A very strong coupling arises when storage is cheap and speculators are numerous and relatively risk tolerant. The process of financialization is therefore characterized by an increasingly strong coupling.

In our experiment we find that the volatility of spot prices changes non-monotonically with coupling strength A . Experimental markets with weak coupling ($A = 0.5$) exhibit significantly smaller spot price fluctuations than almost isolated markets ($A = 0.01$). However, when the coupling strength is increased further to $A = 10$ or $A = 30$, prices deviate from their fundamental value considerably more. Instead of the seemingly random fluctuations observed in weakly coupled and isolated markets, spot prices in the strongly coupled markets typically move in persistent multi-period cycles or go through bubbles and crashes. This dynamics arises endogenously as a result of the use of trend-following in the forecasting strategies of the participants. We find more and stronger trend-following in the treatments with strongly coupled markets.

Our findings suggest that there is not one answer to the question whether futures markets work to stabilize or destabilize commodity spot prices. The answer depends on the circumstances, in particular on how strongly the spot and futures markets are coupled. In our setup, weakly coupled spot and futures markets show lower volatility of spot prices than an isolated spot market, while in the case of a strong coupling the net effect of the futures market is clearly a destabilizing one. Because the dependence of the volatility on the coupling strength is U-shaped, also the response of the volatility to an increase in the coupling strength is ambiguous. Therefore a process like financialization may reduce the volatility of spot prices for some commodities, but increase price volatility for commodities for which the futures and spot markets are more strongly coupled.

References

- ADAM, K., A. MARCET, AND J. BEUTEL (2017): “Stock price booms and expected capital gains,” *American Economic Review*, 107, 2352–2408.
- ALGIERI, B. (2012): “Price volatility, speculation and excessive speculation in commodity markets: Sheep or shepherd behaviour?” Discussion paper, ZEF Discussion Papers on Development Policy.
- AMROMIN, G. AND S. A. SHARPE (2014): “From the horse’s mouth: Economic conditions and investor expectations of risk and return,” *Management Science*, 60, 845–866.
- ANUFRIEV, M. AND C. HOMMES (2012a): “Evolution of market heuristics,” *The Knowledge Engineering Review*, 27, 255–271.
- (2012b): “Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments,” *American Economic Journal: Microeconomics*, 4, 35–64.
- BAO, T. AND C. HOMMES (2019): “When speculators meet constructors: Positive and negative feedback in experimental housing markets,” *Journal of Economic Dynamics and Control*, 107, 103730.
- BARBERIS, N., R. GREENWOOD, L. JIN, AND A. SHLEIFER (2018): “Extrapolation and bubbles,” *Journal of Financial Economics*, 129, 203–227.
- BOHL, M. T. AND P. M. STEPHAN (2013): “Does futures speculation destabilize spot prices? New evidence for commodity markets,” *Journal of Agricultural and Applied Economics*, 45, 595–616.
- BROCK, W. A. AND C. H. HOMMES (1998): “Heterogeneous beliefs and routes to chaos in a simple asset pricing model,” *Journal of Economic Dynamics and Control*, 22, 1235–1274.
- BRUNETTI, C., B. BÜYÜKŞAHİN, AND J. H. HARRIS (2016): “Speculators, prices, and market volatility,” *Journal of Financial and Quantitative Analysis*, 51, 1545–1574.
- BÜYÜKŞAHİN, B. AND J. H. HARRIS (2011): “Do speculators drive crude oil futures prices?” *The Energy Journal*, 167–202.
- CHEN, D. L., M. SCHONGER, AND C. WICKENS (2016): “oTree—An open-source platform for laboratory, online, and field experiments,” *Journal of Behavioral and Experimental Finance*, 9, 88 – 97.

- CHENG, I.-H. AND W. XIONG (2014): “Financialization of commodity markets,” *Annu. Rev. Financ. Econ.*, 6, 419–441.
- DU, X., C. L. YU, AND D. J. HAYES (2011): “Speculation and volatility spillover in the crude oil and agricultural commodity markets: A Bayesian analysis,” *Energy Economics*, 33, 497–503.
- DVIR, E. AND K. S. ROGOFF (2009): “Three epochs of oil,” National Bureau of Economic Research Working Paper.
- FORSYTHE, R., T. R. PALFREY, AND C. R. PLOTT (1984): “Futures Markets and Informational Efficiency: A Laboratory Examination,” *The Journal of Finance*, 39, 955–981.
- FRIEDMAN, D., G. W. HARRISON, AND J. W. SALMON (1984): “The informational efficiency of experimental asset markets,” *Journal of Political Economy*, 92, 349–408.
- FRIEDMAN, M. (1953): *Essays in positive economics*, University of Chicago Press.
- GOLDSTEIN, I. AND L. YANG (2019): “Commodity financialization and information transmission,” Working Paper.
- GRAY, R. W. (1963): “Onions revisited,” *Journal of Farm Economics*, 45, 273–276.
- GREENWOOD, R. AND A. SHLEIFER (2014): “Expectations of returns and expected returns,” *The Review of Financial Studies*, 27, 714–746.
- HART, O. D. AND D. M. KREPS (1986): “Price destabilizing speculation,” *Journal of Political Economy*, 94, 927–952.
- HEEMEIJER, P., C. HOMMES, J. SONNEMANS, AND J. TUINSTRA (2009): “Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation,” *Journal of Economic Dynamics and Control*, 33, 1052–1072.
- HOMMES, C., J. SONNEMANS, J. TUINSTRA, AND H. VAN DE VELDEN (2005): “Coordination of expectations in asset pricing experiments,” *The Review of Financial Studies*, 18, 955–980.
- HOMMES, C., J. SONNEMANS, J. TUINSTRA, AND H. VAN DE VELDEN (2007): “Learning in cobweb experiments,” *Macroeconomic Dynamics*, 11, 8–33.
- HOMMES, C., J. SONNEMANS, J. TUINSTRA, AND H. VAN DE VELDEN (2008): “Expectations and bubbles in asset pricing experiments,” *Journal of Economic Behavior & Organization*, 67, 116–133.

- HOMMES, C. H. (1994): “Dynamics of the cobweb model with adaptive expectations and nonlinear supply and demand,” *Journal of Economic Behavior & Organization*, 24, 315–335.
- (1998): “On the consistency of backward-looking expectations: The case of the cobweb,” *Journal of Economic Behavior & Organization*, 33, 333–362.
- KALDOR, N. (1939): “Speculation and economic stability,” *The Review of Economic Studies*, 7, 1–27.
- KAWAI, M. (1983): “Price volatility of storable commodities under rational expectations in spot and futures markets,” *International Economic Review*, 435–459.
- KENNEDY, JOSEPH P., II. (2012, April 11): “The High Cost of Gambling on Oil,” *The New York Times*, A23.
- KEYNES, J. M. (1936): *The General Theory of Employment, Interest and Money*, McMillan.
- MASTERS, M. W. AND A. K. WHITE (2008): “The accidental Hunt brothers: How institutional investors are driving up food and energy prices,” Accessed July 2017, Available at <https://www.bettermarkets.com/michael-w-masters>.
- MCPHAIL, L. L., X. DU, AND A. MUHAMMAD (2012): “Disentangling corn price volatility: The role of global demand, speculation, and energy,” *Journal of agricultural and applied economics*, 44, 401–410.
- MORGAN, C. W. (1999): “Futures markets and spot price volatility: a case study,” *Journal of Agricultural Economics*, 50, 247–257.
- MUTH, J. F. (1961): “Rational expectations and the theory of price movements,” *Econometrica*, 315–335.
- 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.
- NEWBERY, D. M. (1987): “When do futures destabilize spot prices?” *International Economic Review*, 291–297.
- NOUSSAIR, C. AND S. TUCKER (2006): “Futures markets and bubble formation in experimental asset markets,” *Pacific Economic Review*, 11, 167–184.
- PALAN, S. (2010): “Digital options and efficiency in experimental asset markets,” *Journal of Economic Behavior & Organization*, 75, 506 – 522.

- (2013): “A review of bubbles and crashes in experimental asset markets,” *Journal of Economic Surveys*, 27, 570–588.
- PINDYCK, R. S. AND J. J. ROTEMBERG (1990): “The Excess Co-Movement of Commodity Prices,” *The Economic Journal*, 100, 1173–1189.
- PORTER, D. P. AND V. L. SMITH (1995): “Futures contracting and dividend uncertainty in experimental asset markets,” *Journal of Business*, 509–541.
- POWERS, M. J. (1970): “Does Futures Trading Reduce Price Fluctuations in the Cash Markets?” *American Economic Review*, 60, 460–64.
- ROLL, R. (1984): “Orange juice and weather,” *The American Economic Review*, 74, 861–880.
- SARGENT, T. J. (1987): *Macroeconomic Theory*, Academic Press, 2nd ed.
- SARRIS, A. H. (1984): “Speculative storage, futures markets, and the stability of commodity prices,” *Economic Inquiry*, 22, 80–97.
- SIMS, C. A. (1980): “Macroeconomics and reality,” *Econometrica*, 1–48.
- SOCKIN, M. AND W. XIONG (2015): “Informational frictions and commodity markets,” *The Journal of Finance*, 70, 2063–2098.
- SONNEMANS, J. AND J. TUINSTRA (2010): “Positive expectations feedback experiments and number guessing games as models of financial markets,” *Journal of Economic Psychology*, 31, 964–984.
- STEIN, J. C. (1987): “Informational externalities and welfare-reducing speculation,” *Journal of Political Economy*, 95, 1123–1145.
- TAYLOR, G. S. AND R. M. LEUTHOLD (1974): “The influence of futures trading on cash cattle price variations,” *Food Research Institute Studies*, 13, 29–35.
- WORKING, H. (1960): *Price effects of futures trading*, Stanford University Food Research Institute.

Appendix A Instructions to participants

A.1 General instructions

General information

In this experiment your task is to predict prices. The better your predictions, the more

you will earn. Below you will find some general information about how the prices are formed. Carefully read this information. It will be followed by some questions to check your understanding of the experiment.

The market

We will focus on a single product for which there are several producers. Each period the producers sell everything they produced on a centralized market. The price they receive depends on the total production in that period (supply) and how much the customers want to buy (demand). The customers buy more if the price is lower. As a consequence, market prices will generally be lower in periods with a large production than in periods with a small production.

The producers

A producer sets production in order to maximize profit. His optimal decision depends on the price he will receive for his product. The higher this price, the more he should produce. A complicating factor for the producer is that it takes one period to produce the product. This means that he will need to decide how much to produce **one period in advance**, before he knows the price he will get. His solution is to get an estimate of the price in the next period and base his production decision on this estimate instead. All producers use this method. Note that this means that if on average producers expect higher prices in the next period, total production will be larger. Consequently, the realized prices will then be lower.

The speculators

A speculator specializes in predicting the market price **two periods in advance** as accurately as possible. If his prediction is correct, the speculator can make money. To see how this is possible, consider the following example. Suppose that at the end of period 3, after the market has closed, the speculator expects that the price of sugar in period 5 will be 4 cents per kg. When the market opens in the next period (period 4), he notices that he can buy sugar for 1 cent per kg. He buys a big amount and stores it for one period at a cost of 1 cent per kg. If in period 5 the sugar is indeed worth 4 cents per kg, the speculator made the following profit: $4 \text{ (selling price)} - 1 \text{ (buying price)} - 1 \text{ (storage costs)} = 2 \text{ cents per kg}$.

Also in situations that prices are falling a speculator can make money. This is illustrated in the example below. At the end of period 15 the speculator expects that the price of sugar in period 17 will be 7 cents per kg. In period 16 he notices that people have to pay at least 9 cents for a kg of sugar. He then approaches one of the buyers and proposes to deliver 30 kg of sugar to him in the **next period** for 8 cents per kg. For the

buyer, receiving his order in the next period is not as good as getting it immediately. He may get low on inventories in that period, which can for example frustrate the production process (if the buyer is a manufacturer) or disappoint customers (if the buyer is a shop owner). However, he may still agree because he gets the sugar for a lower price. If he agrees to the deal and the speculator was right with his prediction, the speculator buys 30 kg of sugar for 7 cents per kg on the market in period 17 and delivers it to the buyer for 8 cents per kg. This gives the speculator a profit of 1 cent per kg.

It is important to realize that the activities of speculators have real consequences for market prices. If on average the speculators expect prices to rise in period 5 (as in the first example), prices will already increase in period 4 because of higher demand for the product. And when they expect prices to fall in period 17, prices in period 16 will already be lower than they otherwise would be because people postpone buying the product. To summarize, on average higher expectations of speculators for the price in two periods raises prices in the next period and on average lower expectations of speculators lowers prices in the next period.

Role of the participants

The participants in this experiment will work as **financial advisors** for the producers and the speculators. These employers will base their expectations for 100% on the advice that they receive. Each producer and each speculator is advised by exactly one advisor and also each advisor works for one employer only (either a producer or a speculator). Those who advise a producer are asked to predict the price **one period ahead** and those who advise a speculator are asked to predict the price **two periods ahead**.

A.2 Extra instructions specific to advisors to producers

Your role

You will work as a financial advisor to a **producer**. This means your task is to predict the price of the product **one period ahead**. You will do this for 50 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.

Your task

The experiment starts with period 0. Below you can see a screen shot in which all the important elements are marked. The starting situation is such that in the first period the price will likely be between 8 and 50. You will start by giving a prediction for the price in period 1. After all participants have given their first prediction, the price for the

first period will be revealed and, based on the forecasting error of your prediction, your earnings in the first period will be given.

Earnings

The earnings shown on the computer screen will be in points. If your prediction is p_t^e and the price turns out to be p_t in period t your earnings are determined by the following equation:

$$earnings_t = \max \left\{ 1300 - \frac{1300}{25} (p_t^e - p_t)^2, 0 \right\}.$$

The maximum possible points you can earn for each period (if you make no prediction error) is 1300, and the larger your prediction error is, the fewer points you earn. You will earn 0 points if your prediction error is larger than 5. The earnings table on your desk shows the number of points you earn for different prediction errors. At the end of the experiment, your total earnings in points will be converted into euros, at an exchange rate of 1 euro for 2000 points. In addition you will receive 7 euros for participating in the experiment.

A.3 Extra instructions specific to advisors to speculators

Your role

You will work as a financial advisor to a **speculator**. This means your task is to predict the price of the product **two periods ahead**. You will do this for 50 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.

Your task

The experiment starts with period 0. Below you can see a screen shot in which all the important elements are marked. The starting situation is such that in the first period the price will likely be between 8 and 50. You will start by giving a prediction for the price in period 2. After all participants have given their first prediction, the market price for the first period will be revealed. At that point you are asked to give your prediction for the price of the product in the third period. After all participants have given their predictions again, the spot price in the second period will be revealed and, based on your forecasting error of the first prediction, your earnings for period 2 will be given.

Earnings

The earnings shown on the computer screen will be in points. If your prediction is p_t^e and the price turns out to be p_t in period t your earnings are determined by the following

equation:

$$earnings_t = \max \left\{ 1300 - \frac{1300}{25} (p_t^e - p_t)^2, 0 \right\}.$$

The maximum possible points you can earn for each period (if you make no prediction error) is 1300, and the larger your prediction error is, the fewer points you earn. You will earn 0 points if your prediction error is larger than 5. The earnings table on your desk shows the number of points you earn for different prediction errors. At the end of the experiment, your total earnings in points will be converted into euros, at an exchange rate of 1 euro for 2000 points. In addition you will receive 7 euros for participating in the experiment.

Appendix B Price dynamics with naive and trend-following forecasting heuristics

B.1 Introduction

The price dynamics resulting from the model with futures market critically depends on the agents' forecasting strategies. On the one hand the proper use of storage has the potential to lower price volatility in the presence of uncertainty in supply or demand. This is for example the case if inventories adjust upwards in times of low demand/high production and downwards in times of high demand/low production. On the other hand, the presence of a futures market and storage arbitrage also creates a positive feedback mechanism: the expectation of higher spot prices in the next period by speculators will increase spot prices in the current period. This positive expectations feedback can potentially create extreme price fluctuations (Heemeijer et al., 2009). Forecasting strategies may trigger predominantly the first or the second effect, resulting in, respectively, lower and higher volatility compared to a model without futures market. In this appendix the consequences of two boundedly rational strategies are outlined.

B.2 Naive expectations

One example in which the first, volatility lowering, effect will dominate is the case when both the producers and the speculators use a naive forecasting strategy. Note that this means that the producers use the current period's price as their forecast, while the speculators will use the previous period's price. As in a pure cobweb model, a high (low) price will lead naive producers to produce more (less) next period than consumers will wish to purchase at the fundamental price. This behavior drives prices to swing from one side of the fundamental value to the other, either with a decreasing amplitude or an increasing

amplitude. If the speculative influence is small (small A), for example because of high storage costs or strongly risk averse speculators, these producer driven price swings dominate and the two-period-ahead predictions of the speculators will be quite accurate. As a result in a period with a high commodity spot price the futures price will be low, and vice versa. inventory holders will respond by increasing inventories when prices are low and selling from inventory when prices are high. This increases the stability of the system and lowers price volatility.

If speculative influence is large (large A), low storage costs and low risk premia will drive the current spot price towards the futures price. Again the predictions of naive speculators will be approximately correct, but this time it is 'by construction', because prices will change only slowly. The drivers of the price change in this case will be the producers. Their naive predictions will also be approximately correct, causing high production and considerable inventory build-up when prices are high. The small cost increase associated with the larger inventories brings the price down a bit. When eventually the price reaches the fundamental price p^* a considerable amount of the product will be in storage. The price will keep going down until all excess inventory is sold and then the process reverses. One would thus expect this regime to be characterized by slow oscillations around the fundamental price. Still, the overall price volatility decreases compared to the pure cobweb model.

Formally, the naive forecasting strategy for this model is given by:

$$p_{k,t}^{e,p} = p_{t-1}, \quad p_{k,t}^{e,s} = p_{t-2}.$$

Substituting these in Eq. (14) immediately gives the equation for the price dynamics:

$$p_t = \frac{a - Kc - Kc \tanh(\lambda(p_{t-1} - d)) + A(2p_{t-1} - p_{t-2})}{A + b} + \frac{\epsilon_t}{A + b}. \quad (19)$$

Thus in each period the price is established based on the last two prices plus a stochastic term. The non-stochastic part bears some resemblance to the case of linear backward-looking expectations with two lags treated in Hommes (1998) and we will employ the same technique used there to determine the stability conditions of the steady state at the fundamental price p^* .

Let $p_{t-2} = x_{t-1}$ and $p_{t-1} = y_{t-1}$, such that the non-stochastic part of Eq. (19) can be written:

$$\begin{aligned} x_t &= y_{t-1}, \\ y_t &= \frac{a - Kc - Kc \tanh(\lambda(y_{t-1} - d)) + A(2y_{t-1} - x_{t-1})}{A + b}. \end{aligned}$$

Then $F_{\lambda,A}(x, y)$ maps x and y to their values in the next period:

$$F_{\lambda,A}(x, y) = \left(y, \frac{a - Kc - Kc \tanh(\lambda(y - d)) + A(2y - x)}{A + b} \right). \quad (20)$$

The stability conditions depend on the eigenvalues of the Jacobian of the map $F_{\lambda,A}(x, y)$ at the steady state price:

$$JF_{\lambda,A}(p^*, p^*) = \begin{pmatrix} 0 & 1 \\ -\frac{A}{A+b} & -\frac{Kc\lambda}{(A+b) \cosh^2(\lambda(p^* - d))} + \frac{2A}{A+b} \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ -\frac{A}{A+b} & \frac{-B+2A}{A+b} \end{pmatrix}, \quad (21)$$

where B is defined as the slope of the supply curve at the fundamental price:

$$B = -\frac{Kc\lambda}{\cosh^2(\lambda(p^* - d))}. \quad (22)$$

Only if the absolute values of both eigenvalues are smaller than one, the steady state is stable.

Solving the characteristic equation

$$\xi^2 + \frac{B - 2A}{A + b}\xi + \frac{A}{A + b} = 0 \quad (23)$$

yields the two eigenvalues ξ_1 and ξ_2 :

$$\xi_1(A) = -\frac{B - 2A}{2(A + b)} - \sqrt{\left(\frac{B - 2A}{2(A + b)}\right)^2 - \frac{A}{A + b}}, \quad (24)$$

$$\xi_2(A) = -\frac{B - 2A}{2(A + b)} + \sqrt{\left(\frac{B - 2A}{2(A + b)}\right)^2 - \frac{A}{A + b}}. \quad (25)$$

For $A = 0$ the second eigenvalue is zero while the first eigenvalue reduces to $-\frac{B}{b}$. As A increases the eigenvalues move closer together until the point A_1^* where the term under the root becomes negative and both ξ_1 and ξ_2 become complex:

$$A_1^* = \frac{B^2}{4(B + b)}. \quad (26)$$

At A_1^* both eigenvalues are equal and take a value:

$$\xi_1(A_1^*) = \xi_2(A_1^*) = \frac{B^2 + 2Bb}{2(B^2 + 2Bb + 2b^2)} < 1. \quad (27)$$

Because both eigenvalues are smaller than one, the steady state at the fundamental price will be stable at A_1^* . Therefore if the parameters in an isolated spot market are such

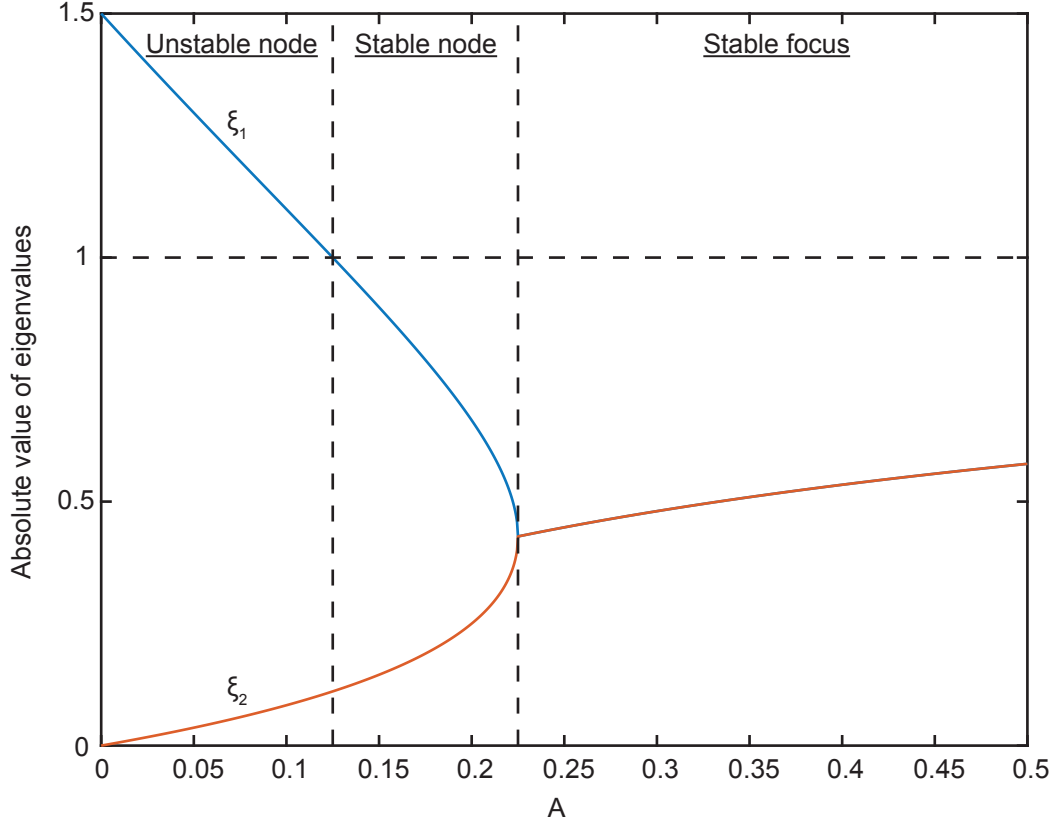


Figure 4: The absolute values of the eigenvalues ξ_1 and ξ_2 of the map corresponding to the model with naive agents and $\lambda = 0.25$. When all eigenvalues have absolute values smaller than one, the steady state at the fundamental price is stable. The regions for A for which this holds are marked 'Stable node' and 'Stable focus'. The latter arises if the two eigenvalues are complex.

that the steady state is unstable, it will become stable when the speculative influence A increases.⁸

Figure 4 shows a plot of the absolute values of the eigenvalues as a function of A for the case of 4 producers and 4 speculators ($K = H = 4$) and supply/demand parameter values $a = 12$, $b = 1$, $c = 1.5$, $d = 6$, and $\lambda = 0.25$. These values are different than used in the experiment, where we wanted to avoid a round number for the fundamental price. However, the conclusions of our analysis is equally valid for those values. In the rest of this appendix, unless explicitly mentioned, we will use the same settings in graphs and simulations. At these settings the fundamental price $p^* = 6$ and the slope of the supply curve $B = 1.5$. Three regions are indicated: an unstable node, a stable node, and a stable focus region. Figure 5 presents a simulated time series for one value of A within each region. In these simulations the standard deviation of the noise term is $\sigma_\epsilon = 0.1$.

⁸This will happen at some point before A reaches the value in Eq. (26). This point is given by $A_2^* = \frac{1}{4}(B - b)$.

In the first region, where both eigenvalues are real, the steady state is unstable. As a consequence in each next period the price lies further away from the fundamental price p^* . The first 15 periods of the time series in Fig. 5(a) illustrate this. After this point divergence stops and the time series shows convergence to a stable two-cycle. Also in the second region both eigenvalues are real. However, here the steady state is stable. Figure 5(b) shows a possible time series of prices in case $A = 0.2$. The scale is 5% of the one in Fig. 5(a). If it weren't for the demand shocks each period, one would see the prices rapidly converging, each iteration being closer to p^* . In the third region the steady state is also stable, but the eigenvalues are complex. With two complex eigenvalues, prices may not converge to the fundamental price directly, but show a damped oscillation towards it (hence a stable focus instead of a stable node). This is exactly the expected behavior of prices when storage is cheap. Figure 5(c) shows an example of a time series for $A = 10$.

The analysis above shows that when all agents are naive a larger influence of the futures market can only increase stability, not decrease it. However, does this also mean that price volatility decreases with larger A ? This is a question that cannot readily be answered from the stability analysis. Therefore we will compare the variances of prices in simulated time series to address this. A simple measure is the variance of the prices in a series compared to the variance of the demand shocks (σ_ϵ^2):

$$\rho = \frac{\text{Var}(p_t)}{\sigma_\epsilon^2}. \quad (28)$$

Each simulated time series had sufficient length (10000 periods) such that possibly different dynamics in the first few periods could be neglected.

The results, ρ as a function of A , are plotted in Fig. 6 for several values of λ . Consistent with the results from the stability analysis, when speculative influence is small the relative variance decreases rapidly with increasing A . Fig. 6 shows that this decrease continues also after the steady state has become stable. Even after the two eigenvalues have become complex, and their absolute values start increasing again (around 0.22 for $\lambda = 0.25$) the relative variance keeps decreasing with A . However, whether this trend continues to infinity does not become clear from the simulations (also not for simulations that extend to larger values of A). In general, the simulations support the intuition that with naive agents price volatility decreases with increasing influence of futures market trading.

B.3 Trend-following expectations

Compared to the naive expectations, stable prices near the fundamental value are much less likely when agents use a trend-following forecasting strategy. A trend-following strategy by producers tends to create even stronger price swings than the naive predictions.

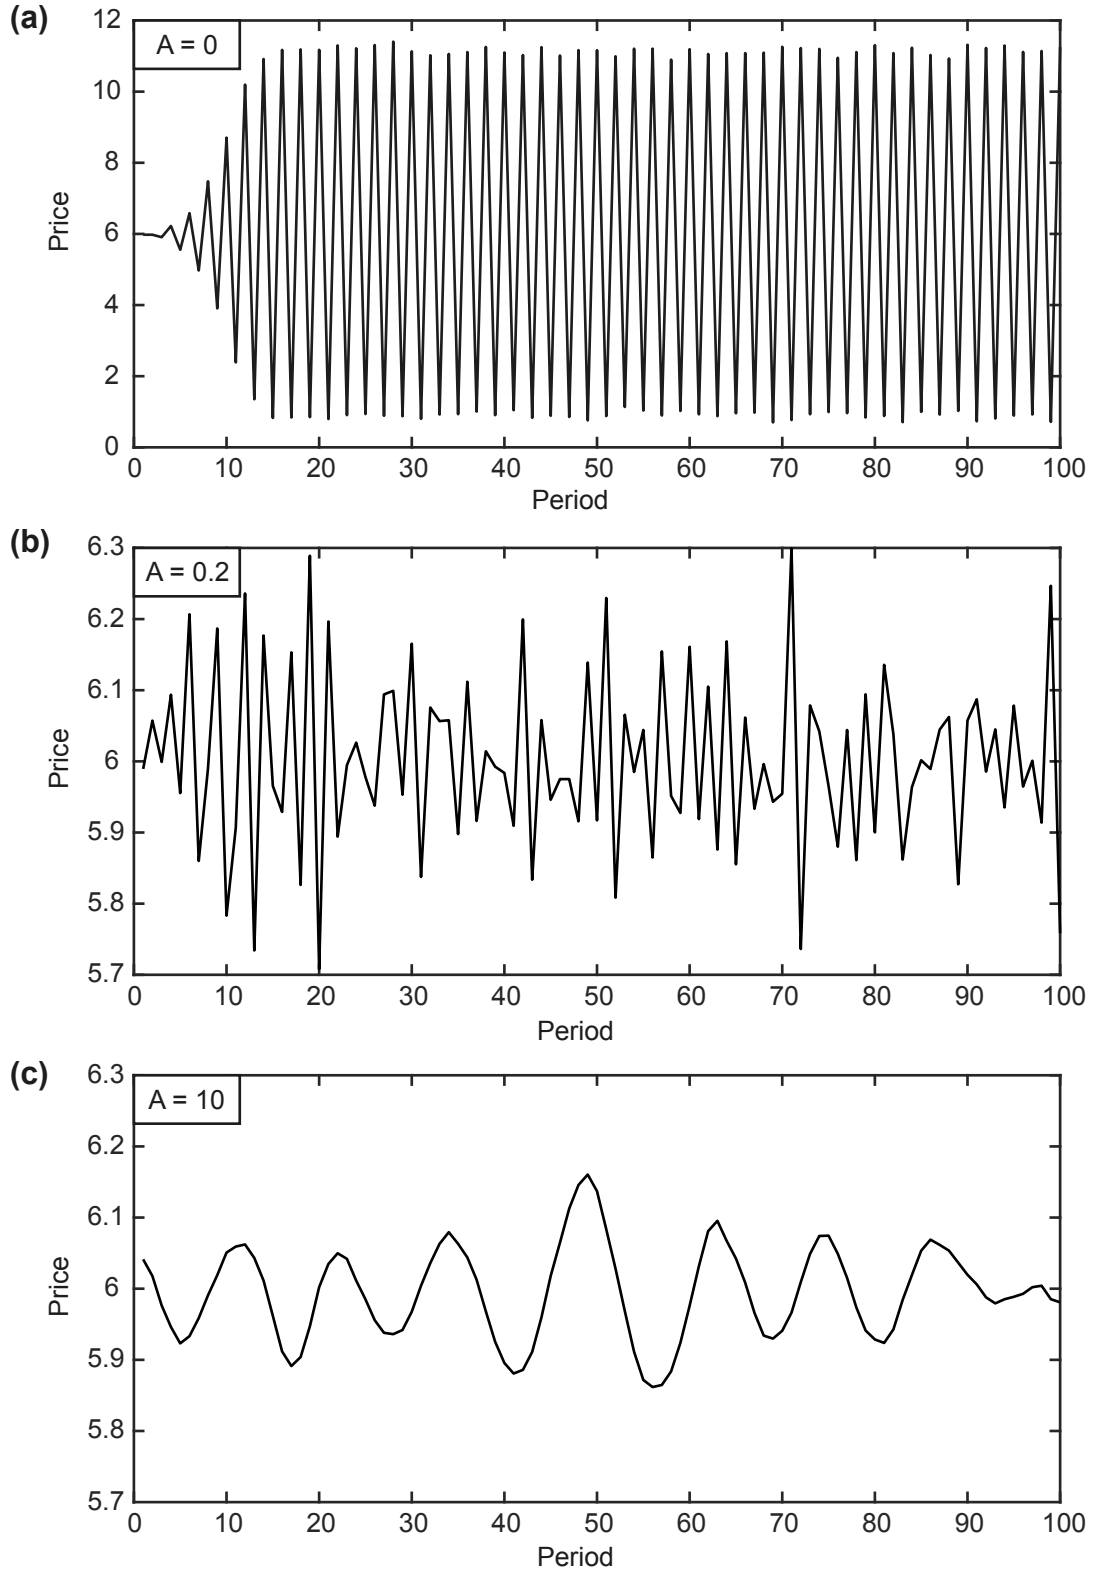


Figure 5: Time series of prices for naive expectations. For increasing values of the futures market influence A the price dynamics changes from a stable two-cycle with alternately very high and very low prices (a), to a stable steady state with small price variations around the fundamental price (b) and (c). The vertical scale in (b) and (c) is 5% of the scale in (a).




Figure 6: Variance of prices relative to the variance of the external noise process in the model with naive expectations. For each value of λ shown, the variance decreases with futures market influence A .

Now, when a positive demand shock drives prices upwards, producers will expect an even higher price next period, leading to an even larger production. Hence the drop in price will also be larger then, giving rise to an even lower prediction for next period's price by the producers. Again if speculative influence is small, the speculators may get the trend approximately right, lowering price volatility and increasing the stability of the system. However, when speculative influence grows this changes quickly. As spot prices start to follow the futures prices closer, a trend set by the speculators will not reverse easily, causing large inventory deviations and prices far from the fundamental value.

A trend-following forecasting strategy for the producers can be written as follows:

$$p_{k,t}^{e,p} = p_{t-1} + \alpha (p_{t-1} - p_{t-2}) = (1 + \alpha) p_{t-1} - \alpha p_{t-2},$$

where α can be any positive number. In literature on heuristic switching models trend-following expectations with values of α of 0.4 and 1.3 have also been called weak trend rule (WTR) and strong trend rule (STR) (Anufriev and Hommes, 2012a,b). In the model discussed here the speculators make two-period ahead forecasts. There are various ways in which a trend-following strategy for two-period-ahead forecasts can be defined. Here we choose the speculators to expect that the most recent price change they observed will repeat itself twice:

$$p_{k,t}^{e,s} = p_{t-2} + 2\alpha (p_{t-2} - p_{t-3}) = (1 + 2\alpha) p_{t-2} - 2\alpha p_{t-3}.$$

With the above defined forecasting strategies the price dynamics of Eq. (14) becomes:

$$p_t = \frac{a - Kc - Kc \tanh(\lambda((1 + \alpha)p_{t-1} - \alpha p_{t-2} - d))}{A + b} + \frac{A(2(1 + \alpha)p_{t-1} - (1 + 4\alpha)p_{t-2} + 2\alpha p_{t-3})}{A + b} + \frac{\epsilon_t}{A + b}. \quad (29)$$

This equation allows for a similar analysis of steady state stability as with naive expectations. There is one important difference: for trend-following expectations the new price depends on the past three prices instead of the past two. This results in a 3-D map for prices from one period to the next:

$$G_{\lambda,A}(x, y, z) = \left(y, z, \frac{a - Kc - Kc \tanh(\lambda((1 + \alpha)z - \alpha y - d))}{A + b} + \frac{A(2(1 + \alpha)z - (1 + 4\alpha)y + 2\alpha x)}{A + b} \right). \quad (30)$$

Here x , y , and z represent the last three prices in time series, with z being the most recent.

Also for a 3-D map the stability of the steady state at the fundamental price depends on the eigenvalues of the Jacobian at p^* :

$$\begin{aligned}
JG_{\lambda,A}(p^*, p^*, p^*) &= \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ \frac{2\alpha A}{A+b} & \frac{\alpha K c \lambda}{(A+b) \cosh^2(\lambda(p^*-d))} - \frac{(1+4\alpha)A}{A+b} & -\frac{(1+\alpha)K c \lambda}{(A+b) \cosh^2(\lambda(p^*-d))} + \frac{2A(1+\alpha)}{A+b} \end{pmatrix} \\
&= \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ \frac{2\alpha A}{A+b} & \frac{\alpha B - (1+4\alpha)A}{A+b} & -\frac{(1+\alpha)(B-2A)}{A+b} \end{pmatrix}. \tag{31}
\end{aligned}$$

Only if all (up to three) eigenvalues have absolute values smaller than one the steady state is stable. The characteristic equation is:

$$\zeta^3 + \frac{(1+\alpha)(B-2A)}{A+b}\zeta^2 + \frac{(1+4\alpha)A - \alpha B}{A+b}\zeta - \frac{2\alpha A}{A+b} = 0. \tag{32}$$

Although analytical solutions for the eigenvalues exist, they are too long to reproduce here. Instead they are plotted in Fig. 7 as a function of A for three different values of α : 0 (naive expectations), 0.4 (WTR), and 1.3 (STR). The stability for the case $\alpha = 0$ was discussed before: if not yet stable for $A = 0$, the steady state will become stable for some value of A and it will remain stable also if A is increased further. This changes when the trend-following component in the forecasts becomes stronger. When agents use the strong trend rule (Fig. 7(c)) the steady state will be unstable for any A . The case in between, for $\alpha = 0.4$, shows a richer picture. As with naive expectations increasing A can make the steady state stable. However, further increases may undo this again. Unlike with naive expectations, the focus does not remain stable for large A .

Figure 8 shows simulated time series for several values of A in case agents use the weak trend rule ($\alpha = 0.4$). The dynamics for $A = 0$ shows a stable two-cycle, similar to the one for naive expectations (compare Fig. 5(a)), but with even larger price swings. At $A = 0.25$ the dynamics takes place around the stable steady state. However, there is still a pronounced alternation in prices visible at most times, which indicates strong negative first-order autocorrelation. For $A = 0.9$ the fast, alternating dynamics is replaced by slower oscillations that are close to a four-cycle. This difference is due to the change of a stable node (real eigenvalues) to a stable focus (2 complex and 1 real eigenvalue). For the largest value of A the steady state has become unstable again (Fig. 8(d)). Prices still show the slower oscillations also present in the stable focus, but now the amplitude of these slower oscillations diverges.

Finally, Fig. 9 shows, both for the weak and the strong trend rule, plots of the relative variance ρ in simulated time series as function of A . As expected, at the value of A where




Figure 7: The absolute values of the eigenvalues of the map corresponding to the model with several types of trend-following agents and $\lambda = 0.25$. When all eigenvalues have absolute values smaller than one, the steady state at the fundamental price is stable. (a) shows repeats the case of naive agents ($\alpha = 0$). The cases for WTR ($\alpha = 0.4$) and STR ($\alpha = 1.3$) are shown in (b) and (c) respectively.

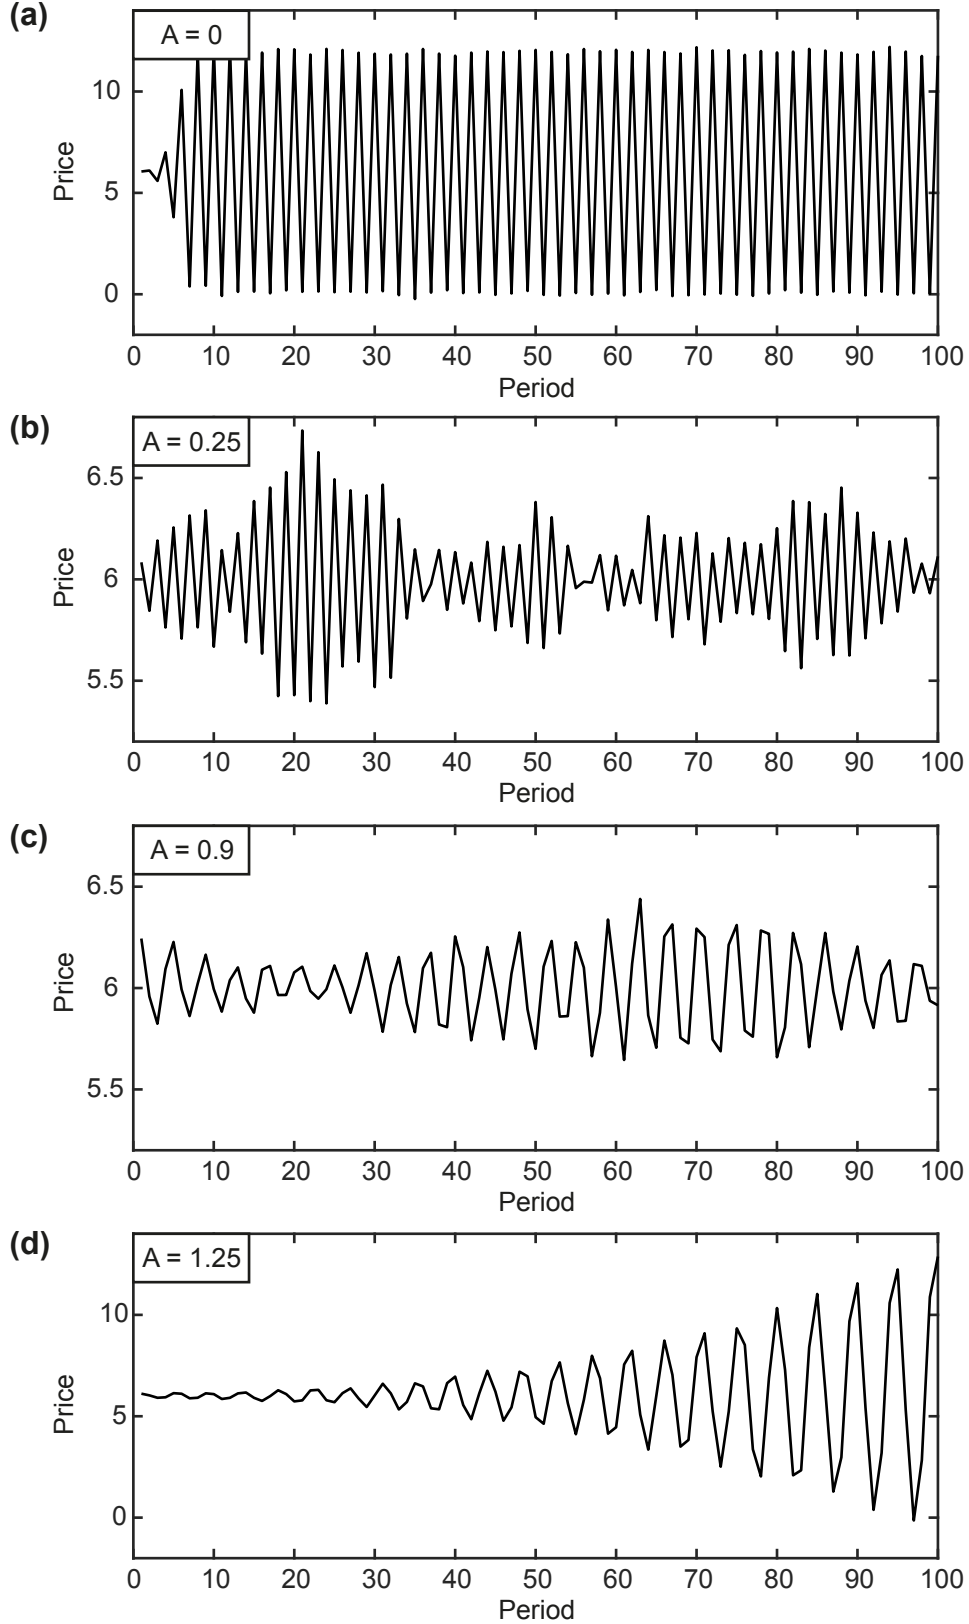


Figure 8: Time series of prices for agents following a weak trend rule ($\alpha = 0.4$). For increasing values of the futures market influence A the price dynamics changes from a stable two-cycle with alternately very high and very low prices (a), to a stable steady state with small price variations around the fundamental price (b) and (c) and finally to an unstable steady state again (d). The vertical scale in (b) and (c) is 10% of the scale in (a) and (d).

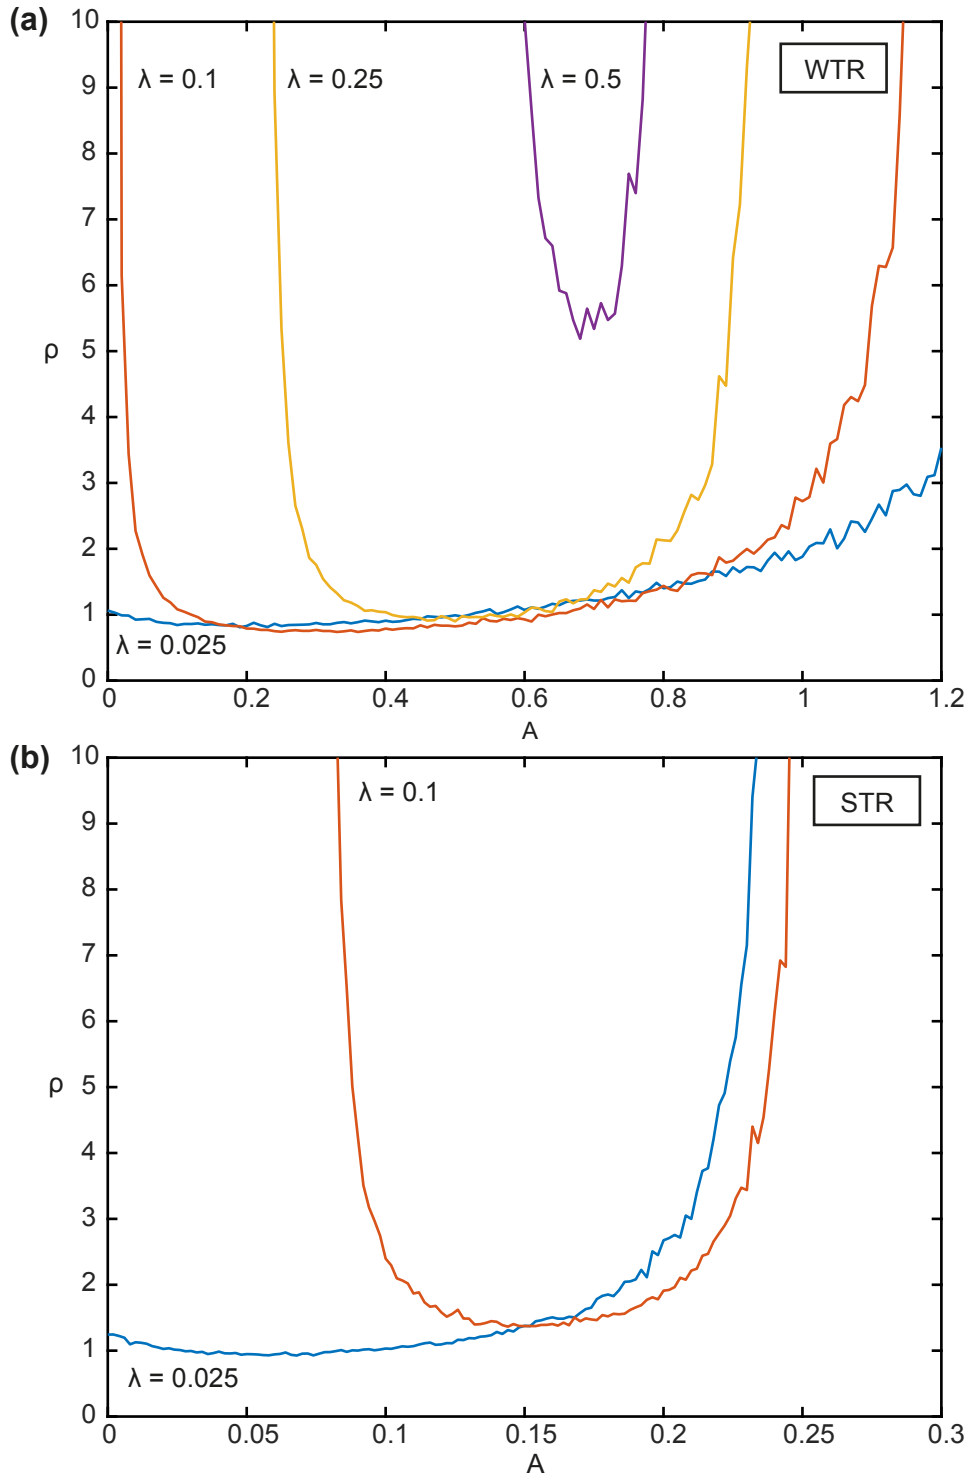


Figure 9: Variance of prices relative to the variance of the external noise process in the model with trend-following expectations. Both curves for the weak trend rule (a) and the strong trend rule (b) are shown. All curves are U-shaped, meaning that the relative variance first decreases with increasing A and then increases again when A gets too large.

the steady state becomes stable (unstable) one can observe a large decrease (increase) in the relative variance. However, all plotted curves (for several values of λ) are U-shaped. This means that in general for trend-following forecasting strategies, the volatility first decreases with increasing A and then increases again. The existence of U-shaped volatility curves has important implications. It shows that under some circumstances it is possible that increased influence of future market trading leads to larger price volatility on the commodity spot markets, despite the experience that the introduction of futures market generally reduces price volatility.

Appendix C Graphs of individual forecasts

Figures 10 to 13 provide plots of the prices in the spot markets along with the individual forecasts of these prices.

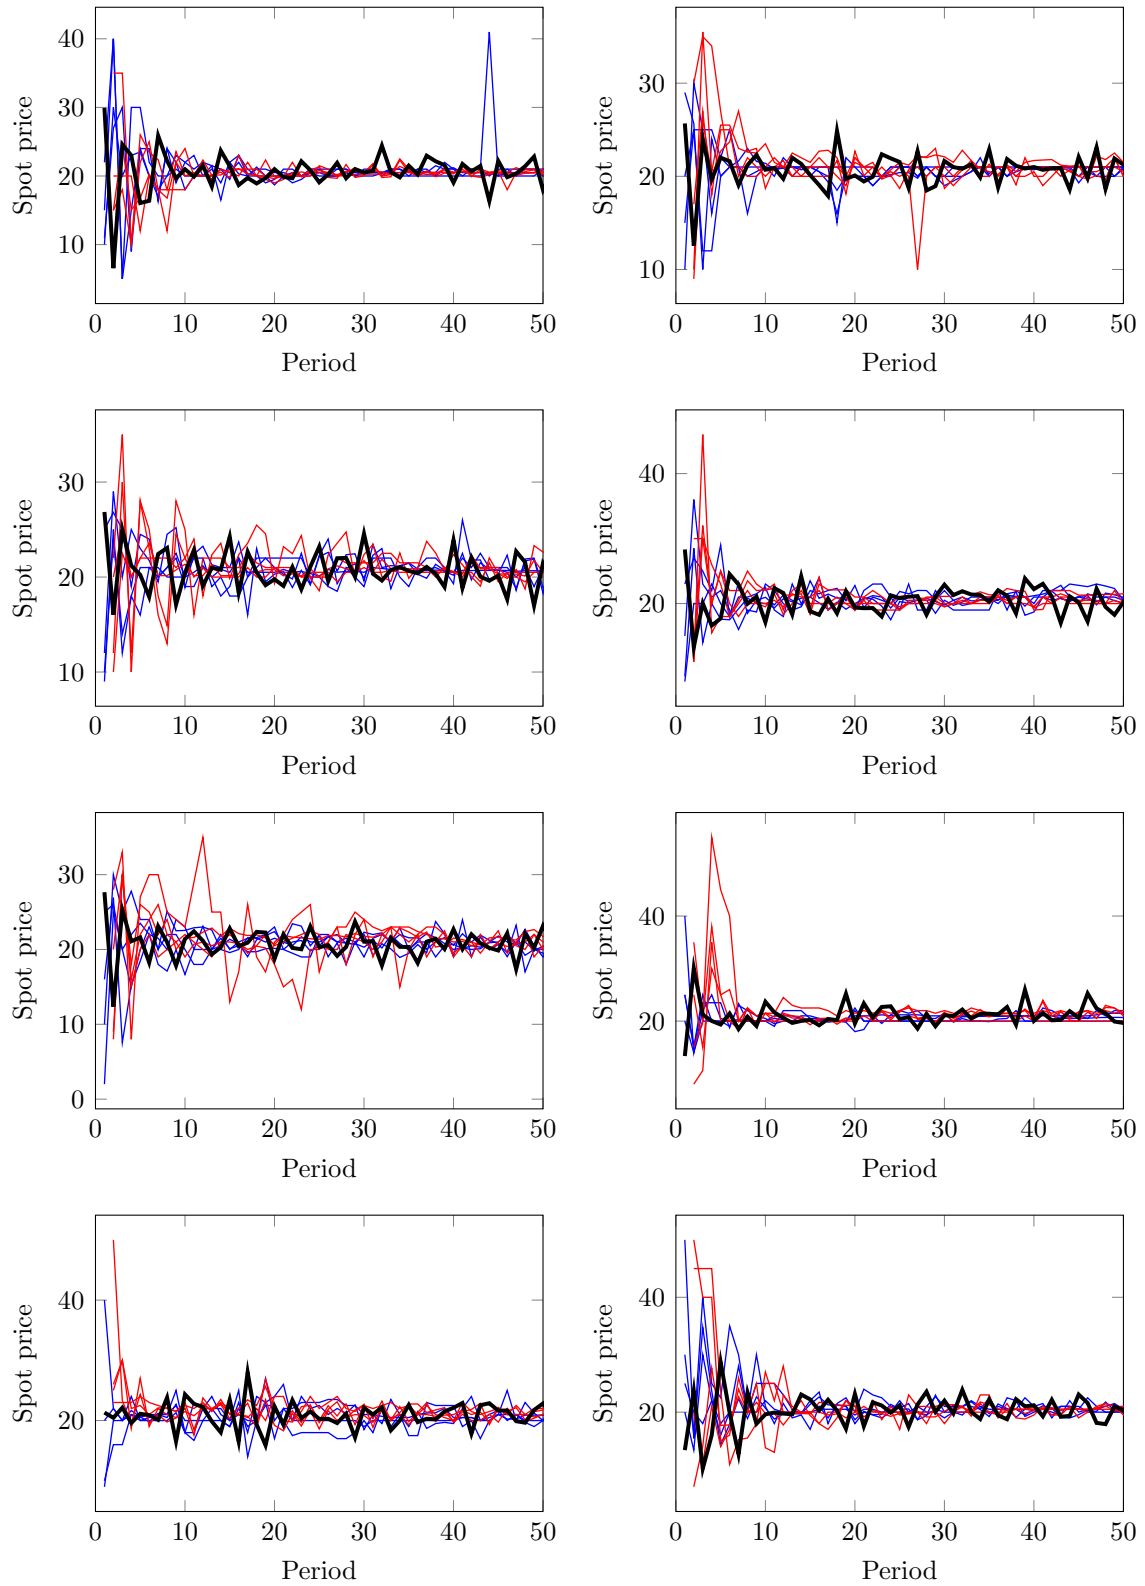


Figure 10: Spot prices (thick line) and individual forecasts of these prices (thin lines) in the almost isolated markets treatment ($A = 0.01$). Forecasts by advisors to producers (in blue) were produced one period before, while forecasts by advisors to speculators (in red) were produced two periods before. Unlike the graphs in Fig. 2, the scales in these graphs are linear and can be different for different graphs.




Figure 11: Spot prices (thick line) and individual forecasts of these prices (thin lines) in the weakly coupled markets treatment ($A = 0.5$). Forecasts by advisors to producers (in blue) were produced one period before, while forecasts by advisors to speculators (in red) were produced two periods before. Unlike the graphs in Fig. 2, the scales in these graphs are linear and can be different for different graphs.

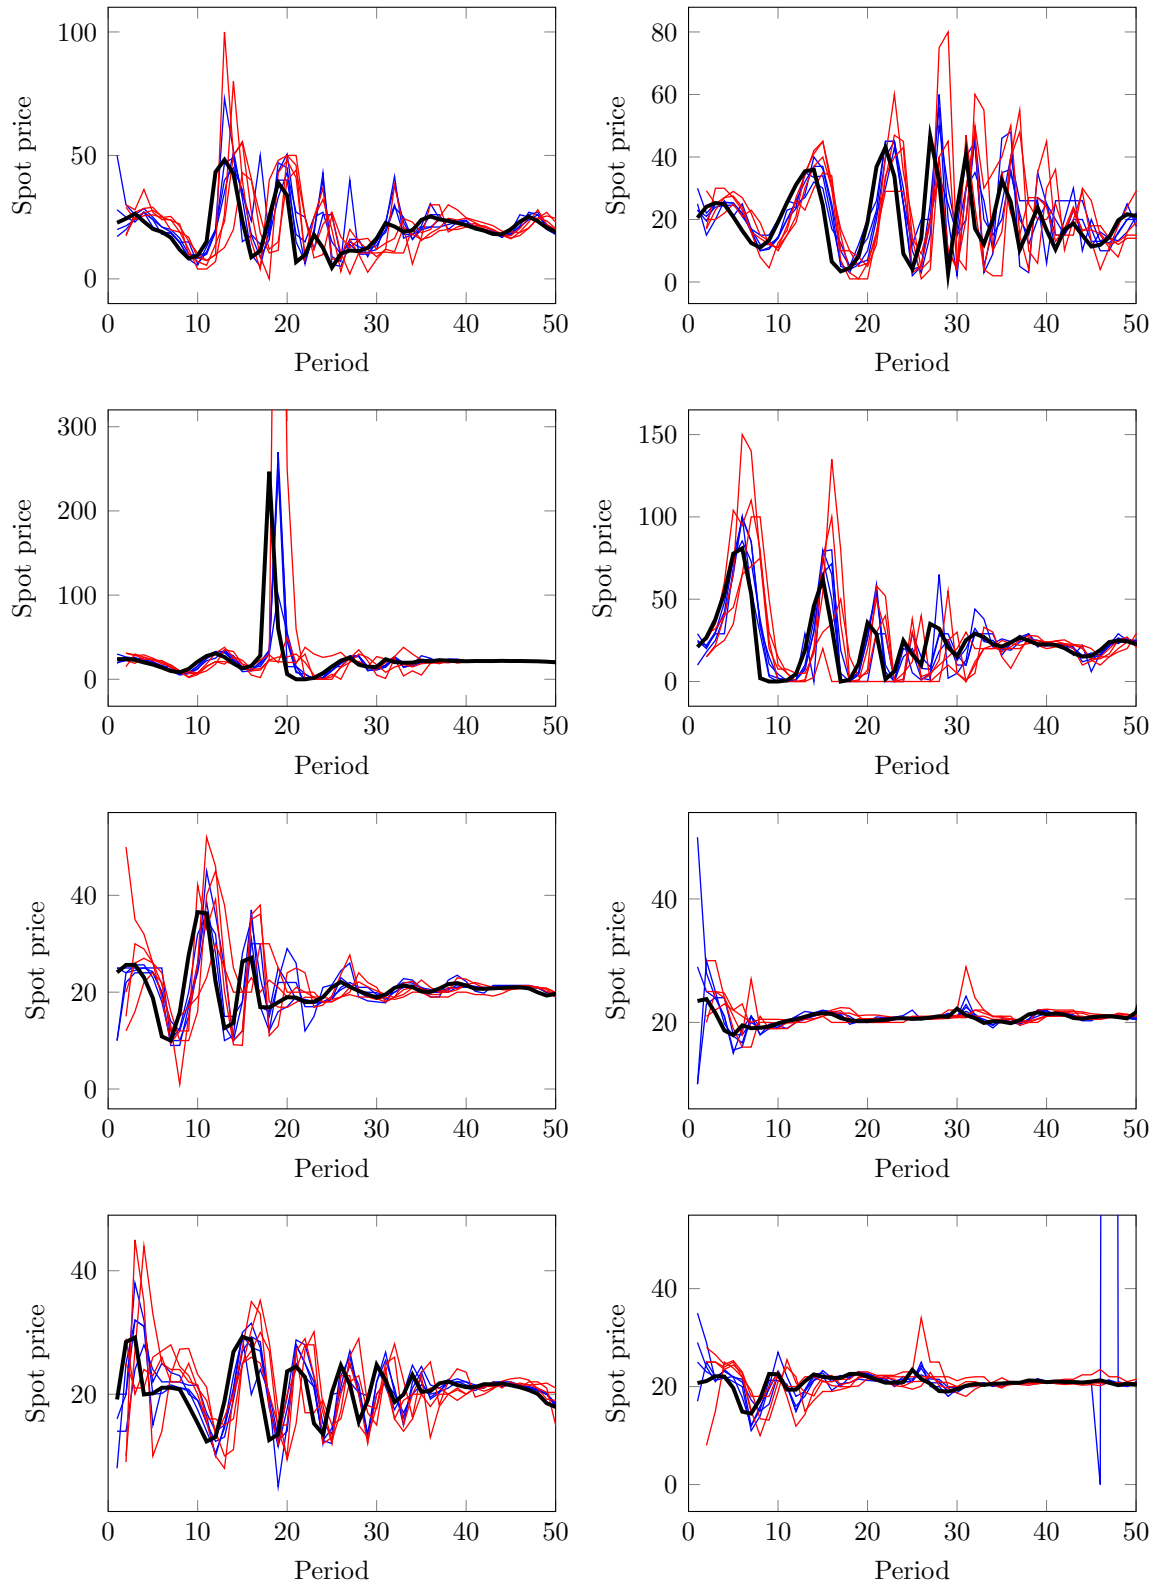


Figure 12: Spot prices (thick line) and individual forecasts of these prices (thin lines) in the strongly coupled markets treatment ($A = 10$). Forecasts by advisors to producers (in blue) were produced one period before, while forecasts by advisors to speculators (in red) were produced two periods before. Unlike the graphs in Fig. 2, the scales in these graphs are linear and can be different for different graphs.




Figure 13: Spot prices (thick line) and individual forecasts of these prices (thin lines) in the very strongly coupled markets treatment ($A = 30$). Forecasts by advisors to producers (in blue) were produced one period before, while forecasts by advisors to speculators (in red) were produced two periods before. Unlike the graphs in Fig. 2, the scales in these graphs are linear and can be different for different graphs.