UvA-DARE (Digital Academic Repository)

Carl's nonlinear cobweb

Hommes, C.

DOI

[10.1016/j.jedc.2017.12.007](https://doi.org/10.1016/j.jedc.2017.12.007)

Publication date

2018

Document Version

Final published version

Published in

Journal of Economic Dynamics and Control

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Hommes, C. (2018). Carl's nonlinear cobweb. *Journal of Economic Dynamics and Control*, 91, 7-20. <https://doi.org/10.1016/j.jedc.2017.12.007>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Carl's nonlinear cobweb

Cars Hommes

CeNDEF, University of Amsterdam and Tinbergen Institute, Netherlands

ARTICLE INFO

Article history:

Received 3 August 2017
Revised 18 October 2017
Accepted 20 December 2017
Available online 9 January 2018

JEL classifications:

D84
D83
E32
C92

Keywords:

Cobweb model
Chaos
Expectations
Heterogeneity
Laboratory experiments

ABSTRACT

This essay surveys some of my work on expectations, learning and bounded rationality within the classical cobweb model following early inspiring ideas from Carl Chiarella. In particular, I focus on the role of nonlinear dynamics, learning and heterogeneity within the cobweb framework and how price fluctuations in the cobweb theory fit with observations of individual and aggregate behaviour from laboratory experiments with human subjects.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

The ideas of Carl Chiarella have inspired generations of young scholars in economics and finance. In the early stages of my own career, during my PhD thesis work on chaos in economic models (Hommes, 1991), I have decisively been inspired by Carl's seminal contribution on chaos in the cobweb model with adaptive expectations (Chiarella, 1988). The current note summarizes some of my work on nonlinear cobweb models over the years following Carl's inspiration. This work includes both theory and laboratory experiments and is not restricted to the case of homogeneous expectations, but also discusses the important role of heterogeneity in forming expectations.

Let me start with a small anecdote. As a PhD student in Spring 1990 I attended a workshop on nonlinear economic dynamics in Vienna, organized by Gustav Feichtinger. At the publishers' book exhibition I came across a new book, *"The Elements of a Nonlinear Theory of Economic Dynamics"*, by Carl (Chiarella, 1990). I took the copy of the book and started reading in the first chapter with the intriguing title: *"The need for a nonlinear theory of economic dynamics"*. While reading, someone came standing next to me and said *"That's an interesting book you have there ..."*. I replied that I was a PhD student and the book was exactly about the topic of my thesis. The man looked at me and said with what became later known to me as his characteristic friendly smile: *"My name is Carl Chiarella, I am the author of that book ..."*. It was the start of many stimulating discussions I had with Carl about economic dynamics, nonlinearity, expectations, heterogeneity, etc., and a long and warm friendship over many years. This note reflects warm memories and many of his stimulating ideas and research that we have discussed over more than 25 years.

E-mail address: c.h.hommes@uva.nl

The article is organized as follows. Section 2 sets out the cobweb model, while Section 3 focusses on adaptive expectations with a seminal contribution from Carl. Section 4 discusses some cobweb models with heterogeneous expectations. Section 5 confronts the cobweb theory with data from laboratory experiments with human subjects. Section 6 discusses a repeated strategy experiment and the frequency of chaos in such a cobweb environment. Finally, Section 7 concludes.

2. The cobweb model

The cobweb model has developed into a textbook benchmark model since its introduction in the 1930s.¹ In particular the cobweb model served as a simple economic environment to investigate the role of expectations. For example, Ezekiel (1938) studied the cobweb model under *naive expectations*, Nerlove (1958) studied the stabilizing role of adaptive expectations and, Muth (1961) used the cobweb framework in his seminal paper introducing rational expectations. Most of these studies use a linear cobweb framework, that is, with linear demand and supply curves. Here, following Chiarella (1988), we will focus on *nonlinear* cobweb models.

The cobweb model is a partial equilibrium model of commodity price fluctuations of a perishable consumption good, such as corn or hogs, that takes some fixed time period to produce. Let p_t be the market price, p_t^e be the producers' expected price and q_t be the produced quantity. The model equations are:

$$D(p_t) = a - dp_t + \epsilon_t, \quad a > 0, \quad d \geq 0 \quad \text{demand} \quad (1)$$

$$S(p_t^e) = \operatorname{argmax}_{q_t} \{p_t^e q_t - c(q_t)\} = (c')^{-1}(p_t^e), \quad \text{supply} \quad (2)$$

$$D(p_t) = S(p_t^e) \quad \text{market clearing} \quad (3)$$

Here D is the consumer's demand curve (derived from utility maximization), S the producer's supply curve (derived from profit maximization with a strictly convex cost function $c(q)$) and (3) describes market clearing. To keep the model as simple as possible, we focus on a linearly decreasing demand curve, where $-d$ is the slope of the demand curve, a determines the demand level and ϵ_t are (small) exogenous independently and identically distributed (IID) demand shocks.

Using ((1)–(3)) and solving for the market clearing price yields

$$p_t = D^{-1}(S(p_t^e)) = \frac{a + \epsilon_t - S(p_t^e)}{d}. \quad (4)$$

The price dynamics in (4) thus depends upon the demand and supply curves, but it also depends crucially on the assumed *expectations hypothesis*. We first consider some benchmark cases, such as naive and rational expectations.

Naive expectations

The simplest case assumes that producers have *naive expectations*, that is, their prediction equals the last observed price: $p_t^e = p_{t-1}$. Under naive expectations, the price dynamics (4) reduces to a one-dimensional dynamical system

$$p_t = D^{-1}(S(p_{t-1})) = f(p_{t-1}). \quad (5)$$

When demand is decreasing and supply is increasing, the map $f = D^{-1}S$ is decreasing. The price dynamics (5) then has a unique steady state p^* , at the unique price where demand and supply intersect. According to the well known *cobweb theorem* (Ezekiel, 1938) the *stability* of the steady state depends upon the ratio of marginal supply and marginal demand at the steady state p^* . There are essentially two possibilities for the price dynamics:

1. if $-1 < S'(p^*)/D'(p^*) < 0$ the steady state p^* is (locally) stable, and prices (in some neighbourhood of the steady state) converge to the steady state;
2. if $S'(p^*)/D'(p^*) < -1$ the steady state p^* is unstable, prices diverge from the steady state and converge to a 2-cycle or unbounded price oscillations arise.

In the case of a nonlinear, bounded supply curve, in case 2 above when the steady state is unstable, prices will converge to a (noisy) *stable 2-cycle*, with regular up and down oscillations, as illustrated in Fig. 5 in Section 5. Hence, under naive expectations, in the cobweb model with nonlinear, but *monotonic* demand and supply curves the price dynamics is always simple: prices either converge to a stable steady state or prices diverge from the steady state converging to a stable 2-cycle (or to unbounded up and down price oscillations).

¹ See e.g. Ricci (1930) (discussing work of H.L. Moore), Schultz (1930), Tinbergen (1930) and Leontief (1934) for early (German) references, surveyed in Waugh (1964). The name cobweb model is due to Kaldor (1934) (pp. 134–135), referring to the price-quantity 'cobweb' diagrams. Pashigian (1987) gives an overview of "cobweb theorems". Rosser (2000) attributes the first reference to the cobweb model to Cheysson (1887).

Rational expectations

An often heard argument against simple forecasting rules such as naive expectations is that they are “irrational” in the sense that they lead to *systematic forecasting errors*. This argument seems particularly relevant in negative expectations feedback systems such as the cobweb model when prices converge to a (noisy) 2-cycle. When producers expect a high (low) price, they will supply a high (low) quantity and consequently, by the law of demand and supply, the realized market price will be low (high), inconsistent with their expectations. Along a ‘hog cycle’ of up and down price oscillations, expectations are *systematically wrong* and forecasting errors are strongly correlated. Smart agents would learn from their systematic mistakes and revise expectations accordingly, so the argument goes. Similar considerations lead (Muth, 1961) to introduce *rational expectations*, where producers’ subjective price expectations are model consistent and equal the objective conditional mathematical expectation of the market price, i.e. $p_t^e = E_t[p_t]$. Imposing rational expectations in the cobweb model implies that expectations must satisfy

$$p_t^e = E_t[p_t] = p^*, \quad (6)$$

where p^* is the (unique) price corresponding to the intersection point of demand and supply curves. Given producers’ rational price forecast $p_t^e = p^*$, the actual law of motion (4) becomes

$$p_t = p^* + \frac{\epsilon_t}{d}. \quad (7)$$

The cobweb model with monotonic demand and supply curves has a unique rational expectations equilibrium (REE), given by an IID process with mean p^* . Along a REE expectations are *self-fulfilling* and producers make no systematic mistakes, since forecasting errors are uncorrelated. However, in order to form rational expectations, perfect knowledge of the underlying market equilibrium equations as well as strong computing abilities to compute the equilibrium price are required. Agents must know the exact demand and supply curves and must be able to compute their intersection point p^* . These assumptions do not seem to be very realistic in a real complex market environment and at least some form of learning would be needed to coordinate expectations on the rational equilibrium.

Chaos

Since the 1980s it is well known that simple nonlinear (economic) models can generate irregular, chaotic fluctuations (Day, 1994; Grandmont, 1985; Hommes, 2013). However, as long as demand and supply curves are *monotonic* and bounded the cobweb dynamics is simple, with prices converging either to steady state or to a 2-cycle. In the 1980s and 1990s several authors have shown that with a non-monotonic supply curve, the cobweb model with naive expectations can generate complicated, chaotic price fluctuations (Artstein, 1983; Day and Hanson, 1991; Jensen and Urban, 1984; Lichtenberg and Ujihara, 1989). A non-monotonic supply curve could be justified by some kind of “income effect”, but does not seem very plausible since the nonlinearity needed to generate chaos is very high.² This brings us to a crucial insight by Carl (Chiarella, 1988): under adaptive expectations a cobweb with nonlinear, but *monotonic* demand and supply curves can generate cycles and chaos.

3. Adaptive expectations

Adaptive expectations, studied in the classical paper by Nerlove (1958), is another simple and frequently used expectations scheme and is given by

$$p_t^e = p_{t-1}^e + w(p_{t-1} - p_{t-1}^e), \quad 0 \leq w \leq 1. \quad (8)$$

The price forecast is “adapted” in the direction of the most recently observed actual price p_{t-1} , with an expectations weight factor w . Adaptive expectations in the form (8) is sometimes referred to as *error correction learning*. Note that for $w = 1$ adaptive expectations (8) reduces to naive expectations. An equivalent form for adaptive expectations is

$$p_t^e = (1 - w)p_{t-1}^e + wp_{t-1}, \quad 0 \leq w \leq 1, \quad (9)$$

that is, expected price is a weighted average of the most recently observed actual price and the most recent expected price. Using (9) repeatedly, the adaptive expectations predictor can be written in a third form, as a weighted sum, with geometrically declining weights, of all past prices:

$$p_t^e = wp_{t-1} + w(1 - w)p_{t-2} + w(1 - w)^2p_{t-3} + \dots = \sum_{i=1}^{\infty} w(1 - w)^{i-1} p_{t-i}. \quad (10)$$

Using (9) and (3), one easily derives that the dynamics of expected prices is described by

$$p_t^e = (1 - w)p_{t-1}^e + wD^{-1}S(p_{t-1}^e). \quad (11)$$

² Hommes and Rosser (2001) consider a fishery model with a strongly backward-bending, discounted supply curve for bioeconomic equilibrium sustained yield leading to chaotic price behaviour.

Fig. 1. Carl Chiarella's nonlinear S-shaped supply curve, convex to the left and concave to the right of its inflection point c .

The one-dimensional map generating the expected price behaviour is

$$f_w(x) = (1 - w)x + wD^{-1}S(x). \quad (12)$$

Fixed points or steady states of the map f_w correspond to intersection points of demand and supply. For decreasing demand and increasing supply the steady state is unique. A straightforward computation shows that the steady state $x = p^*$ is locally stable if

$$-\frac{2}{w} + 1 < S'(p^*)/D'(p^*) (< 0). \quad (13)$$

For $w = 1$ this stability condition reduces to that under naive expectations. For $0 < w < 1$, the stability condition (13) is less stringent and therefore Nerlove (1958), who used linear demand and supply curves, concluded that adaptive expectations has a stabilizing effect upon price dynamics.

Chiarella (1988) studied the cobweb model with adaptive expectations and nonlinear, but monotonic demand and supply curves. In particular, he considered an increasing, nonlinear S-shaped supply curve as illustrated in Fig. 1. The crucial observation made by Carl Chiarella is that for a nonlinear supply curve as in Fig. 1, the corresponding map f in (12) is non-monotonic and therefore may generate cycles or even chaotic dynamics (cf. p. 379, Fig. 2 in Chiarella (1988)).^{3,4}

Hommes (1991, 1994) studies the cobweb model with S-shaped supply curve of the form

$$S(p^e) = b + \tanh[\lambda(p^e - c)], \quad \lambda > 0, \quad b \geq 1, \quad c > 0, \quad (14)$$

where the parameter c is the inflection point, λ tunes the nonlinearity of the supply curve and $b > 1$ ensures that production is always non-negative.⁵ Under adaptive expectations chaotic price oscillations can arise even when both demand and supply are monotonic. There is a simple, intuitive mechanism of why chaotic price fluctuations can arise, even when demand and supply are monotonic. When demand is decreasing and supply is increasing, the composite map $D^{-1}S$ is decreasing. The graph of the map f_w in (12) is a weighted average of the increasing diagonal $y = x$ and the decreasing graph of the (nonlinear) map $D^{-1}S$. Hence, for $w = 0$ f_0 is increasing, while for $w = 1$ (i.e. for naive expectations) the graph of f_1 is decreasing. For nonlinear, monotonic demand and supply curves there typically exists an interval of intermediate w -values for which the map f_w is non-monotonic, as illustrated in Fig. 2. For such w -values chaotic price oscillations may arise.

Fig. 3 shows a bifurcation diagram with respect to the expectations weight factor w . For high values of w , sufficiently close to $w = 1$ (i.e. close to naive expectations) prices converge to a stable 2-cycle, whereas for small values of w , sufficiently close to $w = 0$, prices converge to the RE steady state. For intermediate w -values chaotic price oscillations arise. Notice also that, as w decreases from 1 to 0, the amplitude of the price oscillations decreases. In a nonlinear cobweb model, adaptive expectations therefore has a stabilizing effect upon price fluctuations in the sense that its amplitude decreases; however, at the same time the nature of these smaller amplitude price fluctuations may become more irregular and lead to close to the steady state chaotic price fluctuations.

When prices fluctuate chaotically, the corresponding forecasting errors will be highly unpredictable and the question arises whether boundedly rational agents would be able to detect any structure in these chaotic forecasting errors and improve upon their simple adaptive forecasts. Hommes (1998) shows that when price fluctuations are chaotic, they still

³ Interestingly, the non-monotonic maps involved have two critical points, a local maximum and a local minimum, leading to even more complex dynamics and bifurcation scenarios than for standard non-monotonic maps, such as the quadratic map, with one critical point.

⁴ In a similar spirit Finkenstädt and Kuhbier (1992) show chaotic behaviour under adaptive expectations in a cobweb model with a linear supply and a nonlinear, decreasing demand curve.

⁵ Note that an S-shaped supply curve is fully consistent with profit maximization with a convex cost function and can e.g. be derived from a fourth or higher order polynomial convex cost curve $c(q) = \frac{1}{d+1}(q-1)^{d+1} + q$, where d is an odd integer, e.g. $d = 3$. Optimal supply then becomes $q = S(p^e) = (p^e - 1)^{\frac{1}{d}} + 1$; see Hommes (2000).

Fig. 2. Non-monotonic maps for cobweb model with adaptive expectations and S-shaped supply curves as in (14). The corresponding map f_w in (12) has two critical points where a (local) maximum and (local) minimum occur.

Fig. 3. Bifurcation diagrams with respect to the expectations weight factor w without noise (a) and with noise (b), with the other parameters fixed at $a = 7.65$, $b = 1$, $c = 6$, $d = 1$ and $\lambda = 10$. For large values of w prices converge to a regular 2-cycle with large amplitude. As w decreases the amplitude of price fluctuations decreases and a bifurcation route to chaos occurs. When w becomes very small, chaotic fluctuations are stabilized and prices converge to REE.

exhibit (strong) negative first order autocorrelations. Hence, even though prices fluctuations are chaotic, there may still be room for improvements of forecasting performance by boundedly rational agents.

4. Heterogeneous expectations

So far we have focused on a representative agent cobweb model, where all producers are identical with homogeneous expectations. But agents are heterogeneous and laboratory experiments have shown that, even when facing the same information, individuals may disagree and take different consumption, production or investment decisions. In a complex market environment it seems more appropriate to model agents as *boundedly rational* and *heterogeneous*, using different types of forecasting rules.

A general idea, put forward in Brock and Hommes (1997), is that agents are heterogeneous and choose between different types of expectations rules varying in the degree of rationality, and ranging from very simple to highly sophisticated. More sophisticated rules require more effort, e.g. in terms of information gathering, and are therefore costly compared to simple rule of thumb forecasting rules, that are freely available to everyone. Agents switch between different expectations rules driven by *evolutionary selection* or *reinforcement learning* following the principle that agents tend to switch to more successful rules. Stated differently, rules that have performed better in the recent past attract more followers. Evolutionary selection and reinforcement learning are therefore the main forces to discipline a potentially large population of forecasting strategies.⁶

Brock and Hommes (1997) developed this idea originally in a cobweb model with costly rational versus freely available naive expectations. Here we focus on the cobweb model where agents choose between heterogeneous *linear* forecasting rules. To keep the model as simple as possible, we assume linear demand and supply curves. Market clearing in the cobweb

⁶ There is an extensive literature on models with heterogeneous expectations, see e.g. Hommes (2006) for an overview. Carl Chiarella has made important contributions to this literature, especially in asset pricing models with heterogeneous expectations, see e.g. the survey Chiarella et al. (2009). A recent book reflecting the breadth of Carl's work is Dieci et al. (2014).

model with linear demand and supply and heterogeneous expectations is given by

$$a - dp_t = \sum_{h=1}^H n_{ht} sp_{ht}^e, \quad (15)$$

where p_{ht}^e , $1 \leq h \leq H$, represents the price forecast of producer type h and n_{ht} is the (time varying) fraction of agents using strategy h at the beginning of period t . The only nonlinearity in (15) arises through the time-varying fractions n_{ht} measuring the impact of strategy h at time t .⁷

The evolutionary part of the model describes how beliefs are updated over time, that is, how the fractions n_{ht} of trader types evolve over time. Fractions are updated according to an *evolutionary fitness* or *performance measure*. Rules with better performance will attract more followers. The fitness measures of all forecasting strategies are publically available, but subject to noise. The probability or fraction of agents choosing strategy h is given by the *discrete choice model* with multinomial logit probabilities (or ‘Gibbs’ probabilities (Anderson et al., 1993; Manski and McFadden, 1981)):

$$n_{ht} = \frac{e^{\beta U_{h,t-1}}}{Z_{t-1}}, \quad Z_{t-1} = \sum_{h=1}^H e^{\beta U_{h,t-1}}, \quad (16)$$

where Z_{t-1} is a normalization factor, so that the fractions n_{ht} add up to 1. The crucial feature of (16) is that the higher the fitness $U_{h,t-1}$ of trading strategy h , the more traders will select strategy h . The parameter β in (16) is called the *intensity of choice*, measuring how sensitive the mass of traders is to selecting the optimal prediction strategy. The intensity of choice β is inversely related to the variance of the idiosyncratic noise. One extreme case, $\beta = 0$, corresponds to infinite variance noise, so that differences in fitness cannot be observed and all fractions (16) will be fixed over time and equal to $1/H$. The other extreme case, $\beta = +\infty$, corresponds to the case without noise, so that the deterministic part of the fitness can be observed perfectly and in each period, *all* traders choose the optimal forecast. An increase in the intensity of choice β represents an increase in the degree of rationality w.r.t. evolutionary selection of trading strategies.

Evolutionary fitness here is (a weighted average of) *realized profits*. For the cobweb model with a linear supply curve $q_t = S(p_t^e) = sp_t^e$ (or equivalently a quadratic cost function $c(q) = q^2/(2s)$) realized profit by forecasting strategy h in period t is given by

$$\pi_{ht} = p_t q_t - (q_t)^2/(2s) - C_h = \frac{s}{2} p_{ht}^e (2p_t - p_{ht}^e) - C_h, \quad (17)$$

where C_h represents the average costs per time period for obtaining predictor h . For a simple habitual rule of thumb predictor, such as naive or adaptive expectations, these costs C_h will be zero, whereas for more sophisticated predictors such as fundamentalists beliefs, based on fundamental analysis, or rational expectations costs C_h may be positive.

The strategy switching model (16) assumes *synchronous* updating of strategies, that is, in each period *all* agents update their strategies. Agents, however, do not change strategies every period, but rather update their strategies occasionally and only gradually. Diks and Weide (2005) and Hommes et al. (2005) introduced asynchronous updating in the multi-nomial logit model, as

$$n_{jt} = (1 - \delta) e^{\beta U_{j,t-1}} / Z_{t-1} + \delta n_{j,t-1}, \quad (18)$$

where Z_{t-1} is a normalization factor as in (16). For $\delta = 0$, we are back in the case of synchronous updating, where for a high value of the intensity of choice almost all agents switch to the best strategy. For $\delta > 0$ strategy updating is more gradual with some inertia to stick to the previous strategy.

Competing linear forecasting rules

The cobweb model with costly rational expectations versus free naive expectations should be viewed as a stylized example, with rational expectations representing a costly, sophisticated stabilizing forecasting rule and naive expectations a freely available simple, but destabilizing, forecasting rule. Rational agents have perfect knowledge about market equilibrium equations and, in particular, about the beliefs of other, non-rational agents. While this case is theoretically appealing, it seems highly unrealistic in real markets that some agents have perfect information about the beliefs of other agents. Here we consider other, more realistic examples, where agents only use information extracted from observable quantities, such as prices. As a starting point of the discussion, we consider the case of two simple linear predictors, following Hommes (2013). Two special cases will be discussed, the case of *fundamentalists* versus naive expectations and the case of *contrarians* versus naive expectations.

Consider the two linear AR(1) prediction rules

$$p_{jt}^e = \alpha_j + \beta_j (p_{t-1} - \alpha_j), \quad j = 1, 2, \quad (19)$$

⁷ Goeree and Hommes (2000) show that, for nonlinear (but monotonic) demand and supply curves in a heterogeneous expectations cobweb model with rational versus naive expectations, the dynamical behaviour is very similar.

with parameter α_j representing the perceived long run average and β_j the perceived first-order autocorrelation. The market clearing price in the cobweb model with linear demand and supply and two trader types with linear predictors is given by

$$a - dp_t = n_{1t}s(\alpha_1 + \beta_1(p_{t-1} - \alpha_1)) + n_{2t}s(\alpha_2 + \beta_2(p_{t-1} - \alpha_2)), \quad (20)$$

where n_{1t} and n_{2t} denote the time varying fractions of agents using the linear forecasting rules.

Fundamentalists versus naive

The linear predictors (19) specialize to the case of fundamentalists versus naive expectations when $\alpha_1 = p^* = a/(d+s)$ (the steady state price), $\beta_1 = 0$, $\alpha_2 = 0$ and $\beta_2 = 1$. Fundamentalists always predict the steady state price p^* , where demand and supply intersect. Hence, if all agents would be fundamentalists, the realized market price would immediately jump to the rational expectations price p^* . In a heterogeneous world, fundamentalists are not perfectly rational, because they do not take into account that there are non-rational (in this example naive) agents in the market. Fundamentalists act “as if” all other agents are rational. A fundamentalists strategy however requires structural knowledge of the economy and information about “economic fundamentals”, and therefore information gathering costs for fundamentalists are positive. In the cobweb model the fundamental forecast requires knowledge of demand and supply curves in order to compute the fundamental steady state price p^* . While a fundamentalist strategy may not be perfectly rational, it may be more realistic in a heterogeneous world, because it does not require (perfect) knowledge of the behaviour, in particular the beliefs, of other agents in the market.

This version of the model generates a rational route to randomness when the intensity of choice becomes high (Hommes, 2013, pp. 150–151). The chaotic dynamics have two important features: (i) The sample average quickly settles down to a value close to the fundamental steady state, implying that in the long run both forecasting rules are unbiased; (ii) the first-order sample autocorrelation coefficient is strongly negative, typically of the order -0.85 .

Contrarians versus naive expectations

In the case of fundamentalists versus naive expectations, the price time series exhibit strong first order negative autocorrelations, even when the dynamics is chaotic. An agent who behaves as a time series econometrician could easily detect this strong negative autocorrelation and adapt her forecasts. Even without the use of any statistical software, an agent might detect negative autocorrelation, simply by observing that positive (negative) price deviations from the average are always followed by negative (positive) price deviations. What would happen if agents recognize such regularities from observed market prices and adapt their behaviour?

Consider a group of *contrarians*, who recognize that there is negative first order autocorrelation in realized prices and adapt their forecast by predicting that next period's deviation from the fundamental price will be on the opposite side of the steady state. If we continue to assume that agents are boundedly rational and only use simple linear forecasting rule, it is natural to replace the fundamental forecast by a *linear contrarian rule*

$$p_{1t}^e = p^* + \beta_1(p_{t-1} - p^*), \quad -1 < \beta_1 < 0. \quad (21)$$

This 2-type example with contrarians versus naive expectations also generates rational routes to randomness. Chaotic price fluctuations still exhibit negative first order autocorrelation, but less strong than before, with first order sample autocorrelation around $\rho_t \rightarrow -0.57$ (for $\beta_1 = -0.85$) (Hommes, 2013, pp. 152–153). Apparently, by increasing the level of “rationality” of the boundedly rational agents, that is, assuming that their linear forecasting rule takes (negative) observed autocorrelations into account, the exploitable/forecastable structure in the chaotic price series becomes weaker.

SAC learning versus naive

In the previous example with contrarians versus naive expectations, there is still an inconsistency between what agents believe the first-order autocorrelation to be (i.e. -0.85) and what is observed in price realizations (i.e. first-order autocorrelation -0.57). A careful time series econometrician might try to exploit such forecastable structure. It seems natural to go one step further and introduce a type of agent with *adaptive learning* within the class of linear AR(1) rules, that is, introduce agents with adaptive learning in order to optimize the parameters of the linear forecasting rule. Type 1 agents then use a linear forecasting rule with time varying parameters:

$$p_{1t}^e = \alpha_{t-1} + \beta_{t-1}(p_{t-1} - \alpha_{t-1}), \quad (22)$$

where α_t and β_t are determined through sample autocorrelation (SAC-)learning (Hommes and Sorger, 1998; Hommes and Zhu, 2014):

$$\alpha_t = \frac{1}{t+1} \sum_{i=0}^t p_i, \quad t \geq 1 \quad (23)$$

Fig. 4. SAC-learning versus naive expectations. Top left: strange attractor in (x_t, n_{1t}) -space, with $x_t = p_t - p^*$ the deviation from the fundamental steady state price p^* . Top right: rational route to randomness, i.e. bifurcation route to chaos w.r.t. intensity of choice; Bottom panels: Agents learn to be contrarians, as the sample average of prices converges to its fundamental value (bottom left) and the first order autocorrelation coefficient converges, $\beta_t \rightarrow -0.62$ (bottom right). Parameters $a = 0$, $d = 0.5$, $s = 1.35$, $\delta = 0.5$, $C_1 = 1$, $\alpha_2 = 0$, $\beta_2 = 1$ and $C_2 = 0$.

$$\beta_t = \frac{\sum_{i=0}^{t-1} (p_i - \alpha_t)(p_{i+1} - \alpha_t)}{\sum_{i=0}^{t-1} (p_i - \alpha_t)^2}, \quad t \geq 1. \quad (24)$$

SAC-learning means that type 1 agents are learning the parameters of their linear forecasting rule by the long run sample average α_t and the first-order autocorrelation β_t of prices. This more sophisticated agent type 1 tries to learn an optimal AR(1) rule through adaptive learning, within a heterogeneous agent environment.⁸

Fig. 4 illustrates the dynamics in the 2-type model with costly SAC-learning versus naive expectations. In the long run, agents *learn to be contrarians*, as the sample average price converges to the fundamental steady state (or equivalently, the sample average of the deviation x_t converges to 0) and the first order sample autocorrelation $\beta_t \rightarrow -0.62$, consistent with the observed autocorrelation in realized prices. SAC-learners therefore learn the correct fundamental steady state price level and the correct first order sample autocorrelation of realized prices.

In summary, introducing more sophistication in the learning of linear rules in an unstable cobweb environment leads to chaotic price fluctuations around the steady state, where SAC-learners learn the optimal linear rule in a complex nonlinear environment. Due to adaptive learning and evolutionary selection, much of the strong negative first-order autocorrelations in prices is washed out, but in these simple 2-type examples some negative first-order autocorrelation characteristic for the cobweb dynamics remains even when prices fluctuate chaotically.

⁸ This form of learning fits with a large literature on so-called misspecification equilibria under learning, e.g. Branch (2006), Branch and Evans (2006), Hommes and Zhu (2014).

Fig. 5. Cobweb price dynamics in the strongly unstable treatment ($\lambda = 2$) in two benchmark simulations. Left Panel: convergence to a (noisy) 2-cycle under naive expectations. Right Panel: convergence to (noisy) RE equilibrium price (blue graph) under learning by sample average (red graph). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

5. Laboratory experiments

Accurate data on expectation formation are hard to obtain. Laboratory experiments in a controlled environment can provide useful insights into how individuals form expectations and learn in a market environment. A substantial literature on learning-to-forecast experiments (LtFEs) to study expectations in lab experiments has developed, see e.g. the surveys in Hommes (2011) and Assenza et al. (2014). A LtFE is an empirical test how individuals form expectations in a controlled lab environment, where all other modelling assumptions are computerized, for example consistent with producers' profit maximization and consumers utility maximization, conditional on the individual forecasts formed by subjects in the lab.

Hommes et al. (2007) used LtFEs to study how subjects form expectations in the classical cobweb framework.⁹ The participants were asked to predict next period's price of a commodity under limited information of market characteristics. Participants were only informed about the qualitative features of the cobweb-type market. They were advisors to producers, whose only task is to accurately forecast the price of the good for 50 subsequent periods. Pay-offs were defined as a quadratic function of squared forecasting errors, truncated at 0. Participants were informed that the price would be determined by market clearing and that it would be within the range $[0, 10]$. Furthermore, they knew that there was (negative) feedback from individual price forecasts to realized market price in the sense that if their forecast would increase, the supply would increase and consequently the market clearing price would decrease. Subjects however did not know how large these feedback effects would be, as they had no knowledge of underlying market equilibrium equations. Subjects thus had qualitative information about the market, but no quantitative details.

The realized price p_t in these experiments was determined by the (unknown) market equilibrium between demand and supply:

$$D(p_t) = \sum_{i=1}^K S(p_{i,t}^e), \quad (25)$$

with $p_{i,t}^e$ the price forecast of participant i at time t . In each market there were 6 subjects, i.e. $K = 6$. Demand was linear, $D(p_t) = a - dp_t + \eta_t$, with η_t a small stochastic IID shock. Supply $S_\lambda(p_{i,t}^e)$ was determined by the nonlinear, S-shaped supply curve

$$S_\lambda(p_{i,t}^e) = \tanh(\lambda(p_{i,t}^e - 6)) + 1, \quad (26)$$

as in Chiarella (1988) (see Section 3). Subjects in the experiment thus do not participate themselves in production decisions, but supply is computerized as if each individual producer maximizes expected profit, given his/her individual price forecast. The parameter λ tunes the nonlinearity of the supply curve and the stability of the underlying cobweb model. The resulting equilibrium price is obtained as:

$$p_t = D^{-1} \left(\sum_{i=1}^K S_\lambda(p_{i,t}^e) \right) = \frac{a - \sum_{i=1}^K S_\lambda(p_{i,t}^e)}{d} + \epsilon_t, \quad (27)$$

where $\epsilon_t = \eta_t/d$. The aggregate realized price p_t depends on individual price expectations as well as the realization of the (small) stochastic shocks. While the parameters of the demand function and the realizations of the noise component remained unchanged across all treatments at $a = 13.8$, $d = 1.5$ and $\epsilon_t = \frac{\eta_t}{d} \sim N(0, 0.5)$, the slope parameter of the supply function was varied. Here we consider two treatments. A stable treatment had $\lambda = 0.22$, for which under naive expectations the price converges quickly to the rational expectations equilibrium. In another strongly unstable treatment, with $\lambda = 2$, under naive expectations the RE price is unstable and prices converge to a (noisy) 2-cycle, as illustrated in Fig. 5 (left panel).

⁹ See Carlson (1967), Holt and Villamil (1986), Wellford (1989), Arifovic (1994) and Hommes et al. (2000) for earlier cobweb experiments.

Fig. 6. Cobweb Learning to Forecast Experiments for stable treatment ($\lambda = 0.22$; left panels) and unstable treatment ($\lambda = 2$; right panels). Top panels show prices and individual forecasts, while bottom panels show prices and average price forecast. The stable treatment (left panels) shows coordination on the RE price, while the unstable treatment (right panels) shows persistent heterogeneity and excess volatility.

Along the 2-cycle producers are “irrational” in the sense that they make systematic forecasting errors, predicting a high (low) price when realized market price will be low (high).

Under rational expectations, all individuals would predict the unique price p^* , at which demand and supply intersect. Given that all individuals have rational expectations and forecast p^* , realized prices will be given by

$$p_t = p^* + \epsilon_t, \quad (28)$$

that is, small random fluctuations around the RE steady state. Given the limited market information one cannot expect that all individuals have rational expectations at the outset, but one can hope that in such a simple, stationary environment individuals would *learn to have rational expectations*. For example, if price expectations are given by the *sample average* of past prices, convergence to the RE-price is enforced, as illustrated in Fig. 5 (right panel). The LtFE has been designed to test whether individuals are able to learn from their systematic mistakes under naive expectations and coordinate on a learning algorithm, such as learning by average, enforcing convergence to the RE steady state.

Fig. 6 shows time series of the individual forecasts (top panels) as well as the realized prices together with the average price forecast (bottom panels) for two typical experimental groups, one stable treatment (left panels) and one strongly unstable treatment (right panels). An immediate observation is that in the stable treatment, after a short learning phase of about 10 periods, the price volatility is low and individual forecasts as well as average forecasts are very close to the RE benchmark, with price fluctuations almost entirely driven by the small random shocks in the experiments. This is an example where in an unknown market environment with limited information individuals are able to learn and coordinate on a rational expectations equilibrium. Aggregate price behaviour and individual forecasts are very different however in the strongly unstable treatment. Realized prices exhibit large fluctuations, while individual forecasts are very volatile, even towards the end of the experiment. Hence, our LtFEs show that *only* in the stable cobweb case, the interaction of individual forecasting rules enforces convergence to the RE-benchmark. In the unstable treatment, heterogeneity in forecasting is persistent and has an aggregate effect upon prices characterized by excess volatility.

The behaviour in Fig. 6 is typical for all cobweb experiments. Hommes et al. (2007) summarize the *stylized facts* of the cobweb LtFE experiments as follows: (1) the sample mean of realized prices is close to the RE benchmark p^* in all treatments; (2) the sample variance of realized prices depends on the treatment: it is close to the RE benchmark in the

stable treatment, but significantly higher in the unstable treatments; (3) realized market prices are irregular and do not exhibit significant linear autocorrelations.

These stylized facts across different treatments appear hard to explain by standard learning mechanisms in the theoretical literature. For example, naive expectations are inconsistent with the experiments, because in the unstable treatment it predicts too much regularity (convergence to a 2-cycle) in aggregate price behaviour. Average price expectations, which is just the simplest form of *adaptive learning*, are also inconsistent with the experiments, because for both treatments it predicts convergence to the RE-benchmark (see Fig. 5, right panel). Adaptive expectations with chaotic dynamic are consistent with the first two stylized facts, but not with the third as the chaotic price fluctuations still exhibit strongly negative first-order autocorrelations (Hommes, 1998).

Heterogeneity in forecasting rules is needed to explain the stylized facts of the cobweb experiments across different treatments. Apparently, agents are learning and do *not* coordinate on a simple stable 2-cycle with systematic mistakes. In the experiments, the interaction of agents' individual forecasting and learning rules washes out all linear predictability in aggregate price behaviour. In the stable treatment, this interaction leads to coordination on the “correct” RE benchmark steady state, but in the unstable treatment heterogeneity persists, prices are excessively volatile, while all linear autocorrelation in prices has been washed out.

A GA learning model

Hommes and Lux (2013) present a model of heterogenous individual learning via genetic algorithms (GAs) explaining all three stylized facts in the cobweb LtFEs. Genetic algorithms require a functional specification of the forecasting rule, whose fitness-maximizing parameter values are searched for via the evolutionary algorithm. They use a simple first order autoregressive rule:

$$p_{i,t}^e = \alpha_i + \beta_i(p_{t-1} - \alpha_i). \quad (29)$$

Such a first order autoregressive (AR(1)) rule seems a natural forecasting scheme as agents could implement it using the sample average as their estimate of α_i and the first order sample autocorrelation as the estimate of β_i .¹⁰

The interaction of six individual GA-learning rules simultaneously reproduces all stylized facts in aggregate price behaviour observed in the experiments across the different treatments. Fig. 7 shows typical price time series under GA-learning as well as time series of the two parameters in the AR(1) forecasting rule, averaged over six GA-agents, for the stable treatment (left panel) and the strongly unstable treatment (right panel). In the stable treatment the parameters converge to a neighbourhood of the RE benchmark, consistent with the observed coordination of individual forecasts in the experiments, while in the strongly unstable treatment parameters continue to fluctuate, with α_t fluctuating around the RE steady state and β_t around 0, and prices keep moving away from the RE-benchmark, consistent with the persistent heterogeneity of expectations and excess volatility in the strongly unstable treatment (cf. Fig. 6).

6. A strategy experiment and chaos

In order to study learning in the long run, Sonnemans et al. (2004) performed a strategy experiment in the spirit of Selten et al. (1997)¹¹ in the cobweb framework. This four-round strategy experiment lasted seven weeks. In the cobweb strategy experiment, subjects are asked to formulate a complete strategy, that is, a description of all their forecasts in all possible states of the world (e.g. history of prices). In each period all strategies that participate in the market forecast the next price. The realized market equilibrium price is then determined by a fixed, but unknown, (linear) demand curve and (nonlinear, S-shaped) supply, depending upon individual expected market prices, aggregated over all producers. The realized market price thus depends on all individual strategies. Subjects gain experience in forecasting next period's price in an introductory experiment before submitting their first strategy. These strategies are then programmed and simulated. After each round, subjects receive feedback about the relative performance of their strategy, and the outcomes of five randomly selected simulations in which their strategy is included. Subjects had one week to revise their strategy for the next round. In each of the four rounds of the strategy experiment (as well as in the introductory experiment), financial incentives, based upon prediction performance, were used to motivate the subjects.

Subjects use many different strategies ranging from simple adaptive expectations to more complex forms of adaptive strategies and if-then conditional forecasting strategies. Fig. 8 shows how the distance of the realized market price to the RE price varies over time and across rounds. Within each round there is some learning, as the distance to REE decreases in the first 7 periods, and thereafter stays more or less constant. Furthermore, clearly there is learning over the rounds, with prices much closer to the RE price in rounds 3 and 4. Yet, even in these final rounds, prices do not converge to the RE price, but the squared distance to the RE price remains relatively large (about 50).

What can be said about the price dynamics in the long run? Fig. 9 shows the histograms of long run price dynamics ranging from convergence to a steady state, 2-cycle, 3-cycle, 4–10-cycle, non-convergence to chaos (with positive Lyapunov

¹⁰ In similar cobweb LtFE experiments Heemeijer et al. (2009) estimated individual forecasting rules, and many individuals actually used forecasting rules of the simple AR(1)-form (29).

¹¹ See also the classic work of Axelrod (1984).

Fig. 7. Simulated prices (top panels) and learning parameters under GA-learning of the AR(1) forecasting rule (29). In the stable treatment (left panels) prices converge to RE, while in the unstable treatment (right panels) prices exhibit excess volatility with sample average α_t close to the RE steady state and first-order autocorrelations β_t close to 0, with occasional and persistent departures.

Fig. 8. The mean quadratic distance to RE steady state in each period across the four rounds of the strategy experiment (620 simulations per round).

exponent ($LE > 0$). Over the four rounds convergence to the steady state is relatively rare and occurs only for about 10% of the simulations. Complex chaotic dynamics occurs frequently, and increases to about 60% in round 4. What causes these chaotic price fluctuations? Fig. 9 (right panel) shows that individual strategies do not explain chaotic behaviour. Only a small fraction (about 10%) of individual strategy simulations are chaotic, while individual strategy simulations most often converge to a steady state (30%), a 2-cycle (20%) or a high period cycle (25%). Hence, complicated chaotic behaviour around the steady state is caused by the *heterogeneity* and the *interaction* of different simple or more complex forecasting strategies.

7. Concluding remarks

In this note I have surveyed some of my work on nonlinear cobweb dynamics, inspired by ideas from Carl Chiarella. Nonlinear cobweb models exhibit rich dynamical behaviour under plausible expectations formation processes by boundedly

Fig. 9. Frequency distributions (based on 620 simulation runs) of long run behaviour (steady state, cycles or chaos) in strategy experiment with heterogeneous rules (left panel) and homogeneous individual strategies (right panel). The classification of steady states and cycles is based upon periods 951–1000. Two different kinds of non-convergence simulations are distinguished based upon the estimated largest Lyapunov exponent (1000 periods were used to estimate the Lyapunov exponent).

rational agents. In the case of monotonic demand and supply curves, under naive expectations the dynamics is simple with prices either converging to a stable steady state or to a stable 2-cycle. Along the 2-cycle, however, naive expectations are systematically wrong with high (low) realized market prices when forecasts were low (high). Such an equilibrium is unlikely to survive as smarter agents could exploit such regular price oscillations.

A crucial insight, originating from Chiarella (1988), is that under adaptive expectations, when demand and supply are nonlinear but monotonic, irregular chaotic price fluctuations around the steady state may arise. In such a nonlinear market environment with chaotic price fluctuations it is more difficult to outperform adaptive expectations. Chaotic price fluctuations around the steady state with adaptive expectations may thus be a more reasonable outcome of agents' learning processes. Within the cobweb hog cycle framework these chaotic up-and-down price fluctuations however typically still exhibit strong negative first-order autocorrelations. Such regularities can still easily be observed from market prices and give rise to arbitrage opportunities. Smart agents may enter the market taking these negative autocorrelations into account by contrarian strategies or adaptive learning through sample autocorrelations. The interaction of different types of agents, using simple adaptive expectations as well as more sophisticated learning strategies and switching between them leads to more complex, chaotic price fluctuations. While for 2-type agents systems the chaotic prices still exhibit some regularities, such as (weak) negative first-order autocorrelations, when the number of agent types in the market increases to say five or six, prices become even more irregular and autocorrelations disappear. Interacting agents systems with a small number of heterogeneous agent types may thus explain the complexity of price fluctuations observed in laboratory experiments with human subjects in the cobweb framework.

Heterogeneous agents cobweb models have recently also been used for policy analysis. For example Schmitt and Westerhoff (2015, 2017) have studied how to manage rational routes to randomness by profit taxes, that is, which tax schemes are most effective in stabilizing a (cobweb) economy.

Thirty years after Carl's work, nonlinear cobweb models remain an active area of research and it is likely that Carl's early stimulating ideas, either on nonlinear cobwebs or more broadly on heterogeneous agent models, will continue to inspire the next generations of young economists.

Acknowledgement

This paper has been presented in a special session for Carl Chiarella at the Computation in Economics and Finance (CEF2017) conference, June 28–30, 2017, New York. I would like to thank the special session organizers Herbert Dawid, Tony He and Willi Semmler for the opportunity to present this work and the participants for stimulating discussions.

References

- Anderson, S., de Palma, A., Thisse, J., 1993. *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge.
- Arifovic, J., 1994. Genetic algorithm learning and the cobweb model. *J. Econ. Dyn. Control* 18 (1), 3–28.
- Artstein, Z., 1983. Irregular cobweb dynamics. *Econ. Lett.* 11, 15–17.
- Assenza, T., Bao, T., Hommes, C., Massaro, D., 2014. Experiments on expectations in macroeconomics and finance. In: Duffy, J. (Ed.), *Experiments in Macroeconomics*. In: *Research in Experimental Economics*, 17. Emerald Publishing, pp. 11–70.
- Axelrod, R., 1984. *The Evolution of Cooperation*. Basic Books, New York.
- Branch, W.A., 2006. Restricted perceptions equilibria and learning in macroeconomics. In: Colander, D. (Ed.), *Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model*. Cambridge University Press, New York, pp. 135–160.
- Branch, W.A., Evans, G.W., 2006. Intrinsic heterogeneity in expectation formation. *J. Econ. Theory* 127, 264–295.
- Brock, W.A., Hommes, C.H., 1997. Models of complexity in economics and finance. In: Hey, C.E.A. (Ed.), *System Dynamics in Economic and Financial Models*. Wiley, New York, pp. 3–41.
- Carlson, J., 1967. The stability of an experimental market with a supply response lag. *South. Econ. J.* 33, 305–321.
- Cheysson, E., 1887. La statistique géométrique: Méthode pour la solution des problèmes commerciaux et industrielles. *Legenie Civil*. Paris.
- Chiarella, C., 1988. The cobweb model: its instability and the onset of chaos. *Econ. Modell.* 5, 377–384.

- Chiarella, C., 1990. The Elements of a Nonlinear Theory of Economic Dynamics. Lecture Notes in Economics and Mathematical Systems, 343. Springer Verlag, Heidelberg.
- Chiarella, C., Dieci, R., He, X.-Z., 2009. Heterogeneity, market mechanisms, and asset price dynamics. In: Hens, T., Schenk-Hoppé, K.R. (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*, 231. North-Holland, Amsterdam, pp. 277–344.
- Day, R.H., 1994. *Complex Economic Dynamics Volume I: An Introduction to Dynamical Systems and Market Mechanisms*. MIT Press, Cambridge, MA.
- Day, R.H., Hanson, K.A., 1991. Cobweb chaos. In: Kaul, T.K., Sengupta, J.K. (Eds.), *Economic Models, Estimation and Social Systems, Essays in Honor of Karl A. Fox*. North-Holland, Amsterdam.
- Dieci, R., He, X.Z., Hommes, C.H., 2014. *Nonlinear Economic Dynamics and Financial Modelling: Essays in Honour of Carl Chiarella*. Springer Publishing Company, Incorporated.
- Diks, C.G.H., Weide, R.v. d., 2005. Herding, a-synchronous updating and heterogeneity in memory in a CBS. *J. Econ. Dyn. Control* 29, 741–763.
- Ezekiel, M., 1938. The cobweb theorem. *Q. J. Econ.* 52, 255–280.
- Finkenstädt, B., Kubbier, P., 1992. Chaotic dynamics in agricultural markets. *Ann. Oper. Res.* 37, 73–96.
- Goeree, J.K., Hommes, C.H., 2000. Heterogeneous beliefs and the non-linear cobweb model. *J. Econ. Dyn. Control* 24, 761–798.
- Grandmont, J.-M., 1985. On endogenous competitive business cycles. *Econometrica* 53 (5), 995–1045.
- Heemeijer, P., Hommes, C., Sonnemans, J., Tuinstra, J., 2009. Price stability and volatility in markets with positive and negative expectations feedback: an experimental investigation. *J. Econ. Dyn. Control* 33 (5), 1052–1072. doi:10.1016/j.jedc.2008.09.009.
- Holt, C.A., Villamil, A.P., 1986. A laboratory experiment with a single person cobweb. *Atlantic Econ. J.* 14, 51–54.
- Hommes, C., 2006. Heterogeneous agent models in economics and finance. In: Judd, K., Tesfatsion, L. (Eds.), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics. Handbooks in Economics Series*. Elsevier/North-Holland, Amsterdam, pp. 1109–1186. Chapter 23.
- Hommes, C., 2011. The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics and Control* 35, 1–24.
- Hommes, C., Huang, H., Wang, D., 2005. A robust rational route to randomness in a simple asset pricing model. *J. Econ. Dyn. Control* 29, 1043–1072.
- Hommes, C., Lux, T., 2013. Individual expectations and aggregate behavior in learning to forecast experiments. *Macroecon. Dyn.* 17 (2), 373–401.
- Hommes, C., Sonnemans, J., Tuinstra, J., Velden, H.v. d., 2007. Learning in cobweb experiments. *Macroecon. Dyn.* 11 (S1), 8–33.
- Hommes, C.H., 1991. Chaotic dynamics in economic models: Some simple case studies. University of Groningen Ph.D. thesis. Groningen Theses in Economics, Management & Organization. Wolters-Noordhoff, Groningen.
- Hommes, C.H., 1994. Dynamics of the cobweb model with adaptive expectations and nonlinear supply and demand. *J. Econ. Behav. Organ.* 24, 315–335.
- Hommes, C.H., 1998. On the consistency of backward-looking expectations: the case of the cobweb. *J. Econ. Behav. Organ.* 33, 333–362.
- Hommes, C.H., 2000. Cobweb dynamics under bounded rationality. In: Dockner, E.J., Hartl, R.F., Luptacik, M., Sorger, G. (Eds.), *Optimization, Dynamics and Economic Analysis: Essays in Honor of Gustav Feichtinger*. Springer-Verlag, Berlin, pp. 134–150.
- Hommes, C.H., 2013. *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge University Press.
- Hommes, C.H., Rosser, J.B.J., 2001. Consistent expectations equilibria and complex dynamics in renewable resource markets. *Macroecon. Dyn.* 5, 180–203.
- Hommes, C.H., Sonnemans, J., Velden, H.v. d., 2000. Expectation formation in an experimental cobweb economy. In: DelliGatti, D., Gallegati, M., Kirman, A. (Eds.), *Interaction and Market Structure: Essays on Heterogeneity in Economics*. In: *Lecture Notes in Economics and Mathematical Systems*, Vol. 484. Springer-Verlag, Berlin, pp. 253–266.
- Hommes, C.H., Sorger, G., 1998. Consistent expectations equilibria. *Macroecon. Dyn.* 2, 287–321.
- Hommes, C.H., Zhu, M., 2014. Behavioral learning equilibria. *J. Econ. Theory* 150, 778–814.
- Jensen, R.V., Urban, R., 1984. Chaotic price behaviour in a nonlinear cobweb model. *Econ. Lett.* 15, 235–240.
- Kaldor, N., 1934. A classificatory note on the determinateness of equilibrium. *Rev. Econ. Stud.* 1, 122–136.
- Leontief, W.W., 1934. Verzögerte Angebotsanpassung und Partielles Gleichgewicht. *Z. Nationalökonomie* 5 (5), 670–676.
- Lichtenberg, A.J., Ujihara, A., 1989. Application of nonlinear mapping theory to commodity price fluctuations. *J. Econ. Dyn. Control* 13, 225–246.
- Manski, C., McFadden, D., 1981. Alternative estimators and sample designs for discrete choice analysis. In: Manski, C., McFadden, D. (Eds.), *Structural Analysis of Discrete Data*. MIT Press.
- Muth, J.F., 1961. Rational expectations and the theory of price movements. *Econometrica* 29 (3), 315–335.
- Nerlove, M., 1958. Adaptive expectations and cobweb phenomena. *Q. J. Econ.* 72, 227–240.
- Pashigian, B.P., 1987. Cobweb theorem. In: Eatwell, J., Milgate, M., Newman, P. (Eds.), *The New Palgrave: A Dictionary of Economics*, Volume 1. MacMillan, Basingstoke.
- Ricci, U., 1930. Die "synthetisch ökonomie" von Henry Ludwell Moore. *Z. Nationalökonomie* 1 (5), 649–668.
- Rosser, J.B.J., 2000. *From Catastrophe to Chaos: A General Theory of Economic Discontinuities*, second ed. Kluwer Academic Publishers, Boston.
- Schmitt, N., Westerhoff, F.H., 2015. Managing rational routes to randomness. *J. Econ. Behav. Organ.* 116, 157–173.
- Schmitt, N., Westerhoff, F.H., 2017. Evolutionary competition and profit taxes: market stability versus tax burden. *Macroecon. Dyn.* in press. doi:10.1017/S1365100516000985.
- Schultz, H., 1930. Der Sinn der statistischen Nachfragekurven. Veröffentlichungen der Frankfurter Gesellschaft für Konjunkturforschung (10). Schroeder, Bonn.
- Selten, R., Mitzkewitz, M., Uhlich, G.R., 1997. Duopoly strategies programmed by experienced players. *Econometrica* 65, 517–555.
- Sonnemans, J., Hommes, C.H., Tuinstra, J., Velden, H.v. d., 2004. The instability of a heterogeneous cobweb economy: a strategy experiment in expectation formation. *J. Econ. Behav. Organ.* 54, 453–481.
- Tinbergen, J., 1930. Bestimmung und Deutung von Angebotskurven: Ein Beispiel. *Z. Nationalökonomie* 1 (5), 669–679.
- Waugh, F.V., 1964. Cobweb models. *J. Farm Econ.* 46, 732–750.
- Wellford, C.P., 1989. A laboratory analysis of price dynamics and expectations in the cobweb model. Discussion Paper. Department of Economics, University of Arizona. 89-15.