

Case-Based Learning in the Cobweb Model

Andreas Pape* Wei Xiao †

February, 2014

Abstract

In this paper, we propose a new approach to model bounded rationality in macroeconomics. This method is based on Case-based Decision Theory (Gilboa and Schmeidler, 1995), which been shown to match human choice behavior in experiments (Pape and Kurtz, 2013; Guilfoos and Pape, 2014). The idea is that case-based learning agents behave like ‘average people’ who do not possess structural knowledge of the economy and respond primarily to variables that are directly relevant to their own wellbeing. These agents accumulate a memory bank that stores past cases that they have encountered, and, when they are confronted with a new choice, they judge how similar the current circumstances to cases in memory, and use that judgement to forecast payoffs of alternative actions. We apply this learning approach to the Cobweb model. We find that market prices converge to the rational expectations prices if agents search for a better outcome persistently. If agents are not sufficiently persistent, then multiple equilibria abound, and the rational expectations equilibrium becomes a special case. On the other hand, if they are too persistent, they can search forever and never achieve convergence.

Keywords: Bounded Rationality; Rational Expectations; Adaptive Learning; Case-Based Decision Theory.

JEL Classification Numbers: C62, D84, E31, E37.

*Department of Economics, State University of New York at Binghamton, Binghamton, NY 13902. Email: apape@binghamton.edu

†Department of Economics, State University of New York at Binghamton, Binghamton, NY 13902. Email: wxiao@binghamton.edu

The main reasoning technique that people use is drawing analogies between past cases and the one at hand.

Isaac Gilboa and David Schmeidler
Case-based Decision Theory, 1995

From causes which appear similar we expect similar effects. This is the sum of all of our experimental conclusions.

David Hume
*Philosophical Essays Concerning
Human Understanding*, 1748

1 Introduction

Expectations play an important role in macroeconomics. Modern macroeconomic models build on a dynamic general equilibrium framework, in which future aggregate outcomes are affected by private agents' expectations and expectations depend on aggregate outcomes. Much of recent development in macroeconomics aims at resolving this "self-referential" nature of expectations. The leading paradigm is rational expectations, initially advanced by Muth (1961), and made popular by Lucas (1972) and Sargent (1973). The central idea is that profit-driven behavior eliminates non-rational expectations, and in equilibrium, the subjective distribution of outcomes (expectations) should coincide with the objective distribution.

While rational expectations remains the leading paradigm in modeling expectations, in recent years there has been a burgeoning literature that challenges its premises and practicality. The common criticism is that the rational assumption endows market participants cognitive abilities that are beyond those of their real-life counterparts. In order to form rational expectations, agents are often assumed to be able to observe the true distributions of all shocks, all parameters of the economy, and the true solutions of the structural models. Not even the most sophisticated economists claim to possess such knowledge in real life. The alternative is the "bounded rationality" approach, a methodology put forth by Herbert Simon in 1957. There are various ways bounded rationality has been modeled in macroeconomics. The adaptive learning approach, for example, characterize agents as

econometricians who use observed data to form forecasts of future economic aggregates (Sargent (1993), Evans and Honkapohja (2001)). The genetic algorithm approach simulates learning by applying the genetic mechanisms of reproduction, crossover, and mutation to human agent learning (Arifovic, 1994). Another approach to learning is to use neural networks. In neural networks, the relationship between inputs and outputs are represented with a network structure, and learning takes place when signals are passed through different layers of the network (Cho (1995) and Cho and Sargent (1996)).

In this paper, we propose a new approach to model bounded rationality in macroeconomics. This method is based on Case-based Decision Theory (Gilboa and Schmeidler, 1995). Case-based Decision Theory was developed to be a more realistic decision theory than expected utility by dispensing with the requirement that agents know the entire state space of the problem. CBDT has received some recent empirical support as a model which can explain individual choice behavior in psychological and economic experiments and so can be thought of as providing a model of how an ‘average person’ might respond to a generic choice situation.¹ Keeping with this principle of modeling an ‘average person,’ we assume the case-based learners in our model do not possess knowledge of the structural model nor the distribution of fundamental shocks. Moreover, and in contrast to agents in some of the aforementioned learning models, these agents also do not attempt to estimate a structural model. Instead, they use a memory bank of past cases that they have encountered. When the agent encounters a new case and a decision needs to be made, she looks into her memory bank for cases that are *similar* to the current case. By comparing the payoffs associated with past cases, agents select a choice that delivers the best payoff in similar cases. Agents then accumulate cases in memory as they make more choices, and it is this accumulation of memory which drives learning. This memory bank could be, for example, a farmer’s memories of the prices she charged for wheat in each year and the profits associated with each pricing decision.

We apply this case-based learning approach to the Cobweb model. We choose this model for two reasons. First, we need a very simple model in order to offer

¹For recent work, see Pape and Kurtz (2013) and Guilfoos and Pape (2014). See Section 2 for more details.

a clear and succinct description of our methodology. Two, the Cobweb model has the basic self-referential nature that most macroeconomic models have, and is well-suited for the study of expectation formations. It was the model that Muth (1961) used to promote rational expectations, and was also what Evans and Honkapohja (2001) and Arifovic (1994) used to study econometric learning and genetic algorithms.

A key question we ask is with case-based learning, whether or not the economy can converge to an equilibrium, and if so, what type of equilibrium it would converge to. The result can be summarized as follows: The key parameter is the patience of the agents, which is here captured by a parameter called the *annealing rate* δ . The annealing rate is analogous to the utility discount rate (see Section 2 for more details). For an intermediate level of δ , around .99, we find that the economy converges to an equilibrium which corresponds to the rational expectations equilibrium. When agents are insufficiently patient (δ is too low) we find that agents converge much quicker to an equilibrium, but the price level which results may be far from the REE price. In this case, there are multiple equilibria in the economy, and they are path or history dependent. In other words, with the same initial conditions, the economy can converge to a different equilibrium simply because agents take an alternative path in their learning and decisions processes. On the other hand, we find that when agents are too patient, that is, when δ is too high ($\delta = 1$), agents continue to search for an optimal quantity and there is no convergence in prices at all. This result is reminiscent of the stability result in other learning literatures. For example, with econometric learning, a general finding is that the rational expectations equilibrium is learnable as long as the demand curve has a slope less than one (Evans and Honkapohja, 2001). The path-dependent nature of equilibria is similar to the defining characteristic of a self-confirming equilibrium (Sargent, 2008).

Our result suggests that in the Cobweb economy, the case for a unique rational expectations equilibrium weakens when agents are heuristic learners. Multiple equilibria abound, and the unique equilibrium is a special case. We believe this approach is useful for the study of other macroeconomic topics.

The rest of the paper is structured as follows. Section 2 describes the Case-based decision theory and related literature. Section 3 sets up the Cobweb model

and explains how case-based learning works in this environment (Section 3.3). Section 4 presents our main result. Section 5 concludes.

2 Learning under Case-based Decision Theory

Decision-makers are *case-based learners* if they make choices dynamically consistent with Case-based Decision Theory (Gilboa and Schmeidler, 1995). Case-based Decision Theory postulates that when an agent is confronted with a new problem, she asks herself: How similar is today's circumstance to other circumstances I have experienced? What actions were taken in those cases? What were results? She then forecasts payoffs of actions using her memory, and chooses the action with the highest forecasted payoff.

'Decision theories' are representation theorems: an agent's choice behavior is observed and, if the choice behavior follows certain axioms, a mathematical representation of utility, beliefs, et cetera can be constructed. The dominant decision theory in economics is Expected Utility Theory (von Neumann and Morgenstern, 1944; Savage, 1954). The primitives of expected utility are: a set of actions or 'acts' available to the agent, a set of outcomes, and a set of states of the world. The outcome of a particular action is contingent on the state of the world, and agents may be uncertain about the true state, and instead have a belief distribution over possible states. An example: an agent may choose to take or not take an umbrella to work (actions); the agent might get wet or not (outcomes); and the weather might be rainy or sunny (states of the world).

Case-based Decision Theory was developed with an appeal to *a priori* realism in its structure relative to expected utility theory: that is, Gilboa and Schmeidler say '[a] theory that will provide a more faithful description of how people think would have a better chance of predicting what they will do.' To that end, Gilboa and Schmeidler point out that 'in many decision problems under uncertainty, states of the world are neither naturally given, nor can they be simply formulated.' In the umbrella example, the states of the world are naturally and simply defined. However, consider a more complex decision; Gilboa and Schmeidler suggest considering President Clinton deciding on a military intervention in Bosnia-Herzegovina. The acts may be clear—nothing, economic sanctions, or vari-

ous military interventions—but the the possible states are very difficult to describe or list: factors that determine the numbers and types of casualties that would result from various strategies, relative strength of different warring factions, factors that determine the public opinion response, the internal politics of Russia that may effect their response. Since the states of the world are difficult to describe, and it is an even harder task to assign probabilities to these states, it is hard to imagine that expected utility theory “describe[s] the way people ‘really’ think about [such] problems.”² To find a more realistic formulation they resort to Hume (1748), who states, “From causes which appear *similar* we expect similar effects.” That is, agents make choices by drawing analogies from past cases to present cases. Gilboa and Schmeidler propose Case-based Decision Theory (CBDT) as a decision theory which corresponds to this view. This is actualized in the structure of CBDT by agents maintaining a ‘memory’ of past cases and, when confronted with a new choice, judging the similarity between past cases and the current case to form expectations of the payoff of alternative actions. In this formulation, agents do not need to have access to the correct state space, and in fact do not think of the problem as having a state space at all.³

The increased *a priori* realism of Case-Based Decision Theory has resulted in some early empirical success in explaining human choice behavior. Most relevant to our investigation here, Pape and Kurtz (2013) introduce a computational ‘software agent’ which implements Case-based Decision Theory.⁴ They use this case-based software agent, called CBSA, to find that imperfect memory, accumulative (not average) utility, a similarity function consistent with research from psychology, and a 80 – 85% target success rate renders CBSA a good fit for human data in

²A more recent example on the topic of military intervention would be Donald Rumsfeld’s distinction between ‘known knowns,’ ‘known unknowns,’ and ‘unknown unknowns.’ ‘Unknown unknowns,’ that is, “the things we do not know we don’t know,” could be thought of as states of the world that an expected utility maximizer is not aware of, and therefore assigns a zero probability to. Since no amount of Bayesian updating would move that probability away from zero, this would suggest that, to describe decisions which may involve ‘unknown unknowns,’ one might want a decision theory that did not require agents to have the correct list of possible states of the world.

³Other decision theories have dispensed with the state space without resorting to the Hume “similarity” formulation, such as Karni (2011) and Pape (2013).

⁴A *software agent* is “an encapsulated piece of software that includes data together with behavioral methods that act on these data (Tesfatsion, 2006).” Software agents are used in Agent-based Computational Economics (ACE), of which this paper is a part.

the classification learning experiment from the psychology literature.⁵ The same implementation, in Guilfoos and Pape (2014), was shown to match human data in a repeated prisoner’s dilemma experiment. This software agent is applied to the macro learning problem in this paper, so we are able to use the parameterization found by Pape and Kurtz and Guilfoos and Pape, and from other sources in Psychology.⁶

Let us consider CBDT formally. There are three primitives. The first primitive is a set \mathcal{P} , with typical element p , which is the set of ‘problems’ or circumstances that the decision-maker might face. A problem can be thought of as a vector of exogenous variables that the agent is able to observe before her choice. In the umbrella example, the problem vector might include the prediction of the local weather report that the agent listens to before leaving for work. The second primitive is a set \mathcal{A} , with typical element a , which is the set of actions available to the decision-maker. Third, there is a set \mathcal{R} , with typical element r , which is the set of possible results or outcomes. The action and result sets correspond to the set of actions and outcomes in EUT. A triplet (p, a, r) is called a *case*, and it can be thought of as a complete learning experience. Decision-makers accumulate a set of such cases, called a *memory*, which the decision-maker uses to forecast the outcome of choices.

A decision-maker makes choices of actions a in response to new circumstances p , given her memory \mathcal{M} . When these choices satisfy certain axioms, then one can write down a mathematical representation of that choice, called Case-based Utility $CBU : \mathcal{A} \rightarrow \mathbb{R}$. That mathematical representation has three components. The

⁵In particular, the ‘SHJ’ series of classification learning experiments, starting with Shepard et al. (1961).

⁶There is other empirical evidence in support of Case-based Decision Theory matching human behavior. Gayer et al. (2007) investigate whether case-based reasoning appears to explain human decision-making using housing sales and rental data. They hypothesize and find that that sales data are better explained by rules-based measures because sales are an investment for eventual resale and rules are easier to communicate, while rental data are better explained by case-based measures because rentals are a pure consumption good where communication of measures are irrelevant. Ossadnik et al. (2012) run a repeated choice experiment involving unknown proportions of colored and numbered balls in urns. They find that CBDT explains these data well compared to alternatives such as minimax (Luce and Raiffa, 1957) and reinforcement learning (Roth and Erev, 1995). Golosnoy and Okhrin (2008) use CBDT to construct investment portfolios from real returns data and compare the success of these portfolios to investment portfolios constructed from EUT-based methods, and find some evidence that using CBDT aids portfolio success.

first component is utility $u : \mathcal{R} \rightarrow \mathbb{R}$, which has the usual definition. Second, the agent has an aspiration level $H \in \mathbb{R}$, which represents a target level of utility of the agent (see below). The third component is a similarity function $s : \mathcal{P} \times \mathcal{P} \rightarrow [0, 1]$. The similarity function describes, quite simply, how similar two circumstances are in the mind of the decision-maker, and it captures the spirit of similarity described by Hume. The mathematical representation is:

$$CBU(a) = \sum_{(q,a,r) \in \mathcal{M}(a)} s(p,q) [u(r) - H] \quad (1)$$

Where $\mathcal{M}(a)$ is defined as the subset of the agents' memory \mathcal{M} in which action a was taken. This CBU measure represents the agent's preference in the sense that, for a fixed memory \mathcal{M} and problem p , a is chosen over a' only if $CBU(a) \geq CBU(a')$.

In macroeconomic models of markets, behavior convergence is an important property: When does the price converge to the rational expectations equilibrium level, or some other level, or not at all? Price convergence requires convergence in firm behavior; and in CBDT, convergence is governed by the aspiration value. The aspiration value H , as described above, represents the agent's target level of utility. Functionally, the agent uses the aspiration value as a default value for forecasting utility of new alternatives. As can be seen in the formulation of $CBU(a)$ in Equation 1, there is a switch in sign when agents find a utility level $u(r)$ which exceeds their aspiration level H . Suppose that agents chose an action a which resulted in result r , where $u(r) < H$. Since the term $[u(r) - H]$ is negative, this would *discourage* future attempts at action a , all else equal. On the other hand, if action a' lead to result r' , yielding utility $u(r') > H$, then future attempts at action a' would be *encouraged*. So the level of H determines a switch (on the margin) between *experimentation* with new actions and *exploitation* of current best actions. This is similar to the reservation value, and stopping when that value is reached, that arises from the 'optimal searching' literature starting with (Weitzman, 1979). (The search literature is of course premised on expected utility instead of case-based utility.)

Given its role in determining convergence, a key part of this model is the level and dynamics of the aspiration level. We adopt a simple updating rule for the as-

piration level, by adapting a computational optimization method called ‘simulated annealing’ (Kirkpatrick et al., 1983). The name and inspiration for simulated annealing comes from metallurgy: annealing is process involving heating and cooling metals to decrease defects. Proper annealing in metallurgy requires a correct cooling speed. ‘Slow cooling’ is implemented in the Simulated Annealing algorithm as a slow decrease in the probability of accepting worse solutions: initially, when the ‘temperature’ is hot, the algorithm searches for new solutions and rarely stays at any one solution. Later, as the ‘temperature’ cools, the algorithm settles on the best available solution. Analogously, the method we introduce in this paper is to start agents with an ambitious (i.e. high) aspiration level, and let it fall on a steady schedule. Let H_t be the aspirational value at time t , with exogenous initial level H_0 and *annealing rate* $\delta \in (0, 1]$.⁷ Then we define H_t as:

$$H_t = \delta H_{t-1}$$

i.e. $H_t = \delta^t H_0$

In both macroeconomic and microeconomic dynamic models, it is common to use a utility discount rate $\beta \in (0, 1)$ to discount future utility or profits into present values. Generically, if \vec{x} is a vector of consumption values over time,

$$U(\vec{x}) = \sum_{t=0}^T \beta^t u(x_t) \tag{2}$$

It is common to consider this β as a parameter of utility. It can typically be thought of as ‘patience;’ a β close to one corresponds to an agent who values the future highly, while a β close to zero can be thought of as an agent who puts little value on the future. The annealing rate δ can be thought of as analogous to β . Like β , δ represents ‘patience.’ A δ close to one means that agents keep striving toward a high level of utility, despite never or rarely finding it. This implies that the agent is patient in seeking its goal. On the other hand, a δ close to zero means that agents will quickly compromise their aspirations and settle on a (relatively) low payoff set of actions. Taking the analogy further, formally one can think of β

⁷ $\delta = 1$ implies no annealing.

as the discount rate that an agent applies to her *expectations* of the future; here δ is the discount rate the agent applies to her *aspirations* of the future.

Gilboa and Schmeidler (1996) advocate for an aspiration level with two properties: the aspiration level should be *realistic*, and *ambitious*. By “realistic,” they mean that the aspiration value should be set to be an average of its previous value and the best average performance so far encountered. “Ambitious” means one of two things. What they call *static* ambitiousness calls for an initial aspiration level that is ‘sufficiently’ high. So-called *dynamic* ambitiousness calls for an aspiration level if the aspiration level is (stochastically) set to exceed the maximal average performance by some constant infinitely often. It could be said that the simulated annealing process we follow here is ‘statically ambitious’ and somewhat ‘realistic,’ in that it begins ‘too’ high and falls toward the average, at least over long stretches of the agent’s life.

The functional form of similarity is from the psychology literature; in particular Shepard (1987), in the journal *Science*, who finds that “[e]mpirical results and theoretical derivations point toward two pervasive regularities of generalization.” He finds that similarity “approximates an exponential decay function of distance in psychological space.” In this case, Shepard’s result has remarkably specific implications about the functional form of similarity: similarity ought to be measured by the inverse exponential of vector distance. Applying this result, we use the similarity function $s(p, q) = \frac{1}{e^{d(p,q)}}$ where $p, q \in \mathcal{P}$ and $d(p, q)$ is Euclidean distance.⁸ This form was used in CBSA in Pape and Kurtz (2013) and was tested against some alternatives, where it was found support.

Pape and Kurtz (2013) also find that accumulative, instead of average, similarity provides a better fit for human data. Accumulative similarity is the functional form provided here; average similarity declares that the similarity between two circumstances be normalized by the total similarity so far accumulated in memory. Pape and Kurtz find that, in a simulation setting, average similarity puts undue weight on the possibility that only one action is best regardless of circumstance, as evidenced by a non-trivial fraction of agents, when endowed with perfect memory, exclusively choosing one action even as the circumstance changed. This behavior seems uninteresting in our case, where we want agents to not abandon

⁸Billot et al. (2008) provide an axiomatic foundation to this functional form.

the possibility that the choice of quantity matters to ones payoff.

Along with the primitives and dynamics described above, CBSA defines the decision environment: i.e. those parts of the choice problem that are external to the agent. These need not be defined for a decision theory, so are not a formal part of CBDT. They need only be formally defined when one seeks to generate simulated choice behavior to compare to empirical data, as we do here.

In CBSA the decision environment is represented by function (algorithm) called the *problem-result map* or *PRM*. The *PRM* is the transition function of the environment. It takes as input the current problem $p \in \mathcal{P}$ the agent is facing, the action $a \in \mathcal{A}$ that the agent has chosen, and some vector $\theta \in \Theta$ of environmental characteristics. The *PRM* returns the outcome of these three inputs: namely, it returns a result $r \in \mathcal{R}$; the next problem $p' \in \mathcal{P}$ that the agent faces; and a potentially modified vector of environmental characteristics $\theta' \in \Theta$. I.e.:

$$PRM : \mathcal{P} \times \mathcal{A} \times \Theta \rightarrow \mathcal{R} \times \mathcal{P} \times \Theta$$

For example, consider the firm choosing quantity in a setting with a state variable which affects costs of production. The problem is the current level of the state variable, and the similarity function could describe how similar different levels of the state variable are. Now consider a series of such problems. The *PRM* can be thought of as a device which delivers state variable levels to the agent. It provides a state variable level (which embeds the result of each action) and maintains a θ which describes, for example, the distribution of future state variables. (The *PRM* which corresponds to the macro learning problem is described in greater detail in Section 3.3.)

In the remained of this section, we briefly present the formal statement of the algorithms (functions) which govern CBSA, including how exactly the time-dependent aspiration value H_t enters the agent's choice.

Figure 1 describes the choice algorithm which implements the core of CBSA. It is an algorithmic description of the choice process defined by CBDT, with two modifications. The modifications allow for imperfect memory. In Pape and Kurtz (2013), it was found that a match between CBSA and human data was only achieved by allowing for imperfect memory: otherwise CBSA solves the classifica-

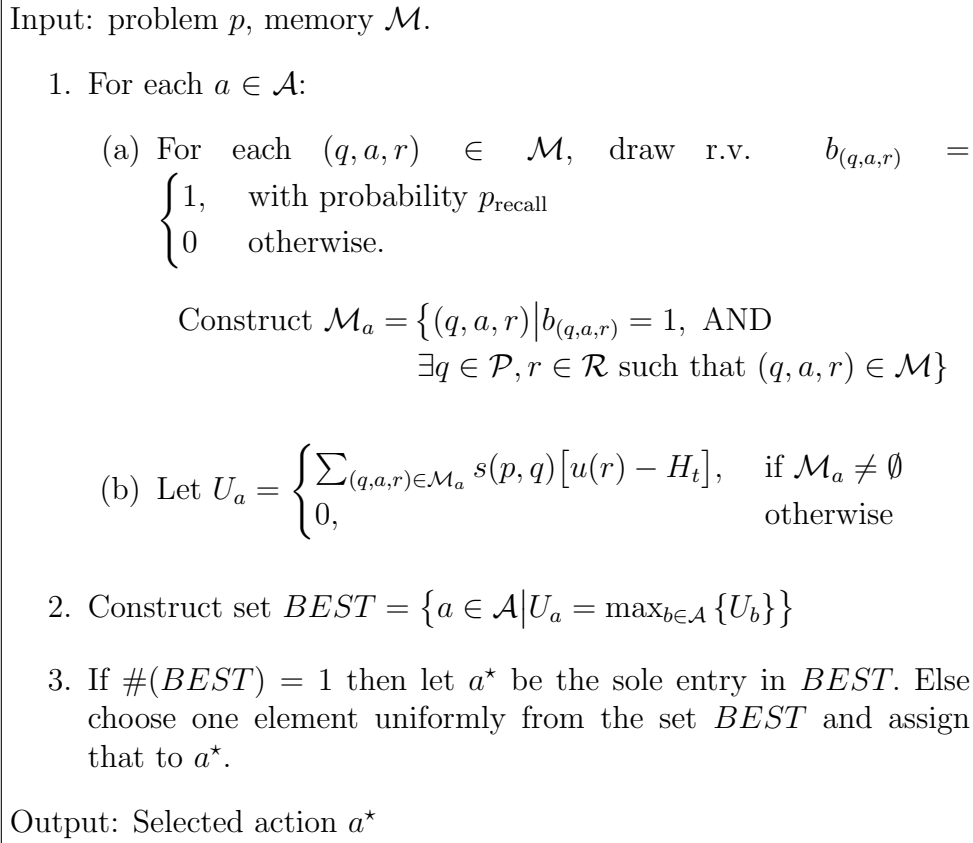


Figure 1: The Choice Algorithm

tion learning problem much faster than humans. There are two kinds of imperfect memory. First, there is *imperfect recall*, governed by a probability $p_{\text{recall}} \in [0, 1]$. Imperfect recall corresponds to an inability to access all memory at any given time, and it is therefore associated with limited cognitive capacity. Second, there is *imperfect storage*, governed by a probability $p_{\text{store}} \in [0, 1]$. Imperfect storage corresponds to a failure to add some experiences to memory after they are experienced, and it is therefore associated with limited memory storage capacity.

In Figure 1, the agent faces a problem $p \in \mathcal{P}$ and has a memory $\mathcal{M} \subseteq \mathcal{C}$. In Step 1a, for each action a , she collects those cases in which she performed this act. Since her recall is imperfect, relevant cases are selected into the set \mathcal{M}_a with probability p_{recall} , where relevant cases which are not recalled are simply ignored.⁹ In Step 1b, she uses this subset of her memory \mathcal{M}_a to construct a utility forecast of that act, called here U_a . The agent then chooses the action which corresponds to the maximum U . There is an additional step, left unspecified in the original CBDT: In the case of a tie, the agent randomizes uniformly over the acts which achieve this maximum.

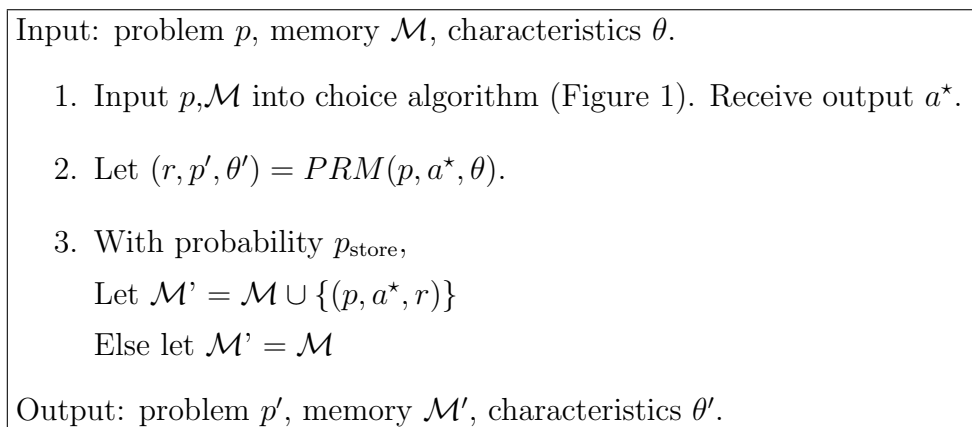


Figure 2: A Single Choice Problem.

Figure 2 describes a single choice problem faced by the agent. It imbeds a reference to the choice algorithm described in Figure 1. Figure 2 embeds the agent in an environment and explicitly references that environment, in the call to PRM . In Step One, the agent selects an act, a^* . In Step Two, the action is

⁹When $p_{\text{recall}} = 1$, the agent has perfect recall. It then corresponds to CBDT as it appears in Gilboa and Schmeidler (1995).

performed, in the sense that the environment of the agent reacts to the agent's choice: the *PRM* takes the current problem p , the action a^* selected by the agent, and the characteristics unobserved by the agent θ , and constructs a result r , a next problem p' , and a next set of characteristics θ' . In Step Three, the agent's memory is augmented by the new case which was just encountered, so long as the agent does not have a 'write-to-memory error:' i.e., with probability p_{store} , the case that was just experienced is added to the set \mathcal{M} . With probability $(1 - p_{\text{store}})$, that case is discarded.

Since the choice problem depicted in Figure 2 maps a problem, characteristic, memory vector to another vector in the same space, it can be applied iteratively. A series of such iterations, along with initial conditions and ending conditions, can then be used to produce a single time series of agent behavior, called a 'run.' Here, the initial conditions specify that agents have an empty memory, although it is simple to modify the algorithm such that the agent starts with some non-empty memory. The ending conditions can take on a variety of forms, and can be exogenous or endogenous. Often times the ending condition is simply a predetermined number of periods, as is the case in this paper.

3 The Cobweb Model

3.1 The environment

This is a competitive economy that produces a single good. While demand responds to the price instantaneously, supply cannot. Production lags force producers to forecast future prices, and make decisions based on the forecasts. There are n firms. The cost function c_i of firm i at time t is

$$c_{it} = aq_{it} + \frac{1}{2}bq_{it}^2 - cw_{t-1}q_{it}, \quad (3)$$

where $a, b, c > 0$, q_{it} is the quantity it produces at time t , and w_{t-1} is a state variable that affects the profit at time t . If we think of the good as an agricultural product such as wheat, then w can be thought of as a value that represents weather conditions. We assume $Ew = 0$ and $Eww' = \Omega$.

Firms maximize their expected profit:

$$\pi_{it}^e = p_{it}^e q_{it} - a q_{it} - \frac{1}{2} b q_{it}^2 + c w_{t-1} q_{it}, \quad (4)$$

At time $t - 1$, firm i chooses quantities q_{it} to maximize its profit. From the first order condition of this maximization problem, we can derive the individual supply curve

$$q_{it} = -\frac{a}{b} + \frac{1}{b} p_{it}^e + \frac{c}{b} w_{t-1}. \quad (5)$$

We assume that the demand schedule is:

$$p_t = \theta - \beta \sum_{i=1}^n q_{it}, \quad (6)$$

where $\theta, \beta > 0$.

Equating market demand and supply:

$$\sum_{i=1}^n q_{it} = n \left(-\frac{a}{b} + \frac{c}{b} w_{t-1} \right) + \frac{1}{b} \sum_{i=1}^n p_{it}^e = \frac{\theta - p_t}{\beta}, \quad (7)$$

we obtain the behavioral equation that describes the evolution of prices:

$$p_t = \mu + \alpha^* \sum_{i=1}^n p_{it}^e + \delta w_{t-1}, \quad (8)$$

where $\mu = \frac{an\beta}{b} + \theta$, $\alpha^* = -\frac{\beta}{b}$, and $\delta = -\frac{cn\beta}{b}$.

3.2 Rational expectations solution

Rational expectations require that agents' subjective expectations match the objective distributions of the actual economic system. The amounts to

$$p_{it}^e = E_{t-1} p_t, \quad (9)$$

where E_{t-1} represents the conditional expectations based on information available at time $t - 1$. Note that since agents are assumed to possess the same information,

their expectations are homogeneous.

Hence, (8) becomes

$$p_t = \mu + \alpha E_{t-1} p_t + \delta w_{t-1}, \quad (10)$$

where $\alpha = -\frac{\beta n}{b}$.

Taking conditional expectations on both sides of (10), we get

$$E_{t-1} p_t = \mu + \alpha E_{t-1} p_t + \delta w_{t-1}. \quad (11)$$

We can then solve for $E_{t-1} p_t$ as

$$E_{t-1} p_t = \frac{\mu}{1 - \alpha} + \frac{\delta}{1 - \alpha} w_{t-1}. \quad (12)$$

It is also clear from (10) and (11) that $p_t = E_{t-1} p_t$. Therefore the rational expectations solution for p_t is

$$p_t = \frac{\mu}{1 - \alpha} + \frac{\delta}{1 - \alpha} w_{t-1}. \quad (13)$$

A closely related solution is that of adaptive learning agents. Suppose agents behave like econometricians who use observed data to estimate the law of motion of the economy. Let's assume that learning agents understand that w_{t-1} has an impact on prices. Their perceived law of motion of the economy is

$$p_t = a + b w_{t-1}. \quad (14)$$

They run regressions of p_t against w_{t-1} to estimate the parameters a and b . Will agents learn the rational expectations solution? Evans and Honkapohja (2001) show that as long as the demand curve satisfies a slope condition, prices will eventually converge to the rational expectations equilibrium prices. Specifically, learnability requires that

$$\alpha < 1. \quad (15)$$

This result is often used to justify the rational expectations solution.

3.3 Case-based Learning

In order to use CBSA to produce choice behavior for the firm in this economy, we must define each component of CBDT/CBSA to correspond to the firm's choice problem. As described in Section 2, the three primitives are the set of problems or circumstances faced by the agent, the set of results or outcomes, and the set of actions which the agent selects in response to the problem.

The set of actions is easiest to define. Each period, a given firm i must choose a quantity q_{it} . Therefore, we let the set of actions Q be the finite set of allowable quantities. The definition of Q —its minimum and maximum allowable values, the coarseness of its coverage—must be specified for any given simulation run and can be thought of as a parameter of the setting. In this paper, we choose a Q which includes the REE quantity, and ranges from a quantity of zero to twice the REE quantity. We discretize this interval into 61 steps.

The set of results is the next easiest to define. Given the level of the state variable last period, w_{t-1} , and the current price p_t , the *result* of action q_{it} is simply the profit level. Let the set of results Π be the set of possible profit levels.

Finally, we define the set of problems or circumstances. The *circumstance* that surrounds the choice can be thought of as a vector of all information the firm is able to observe at time t about the current state of the world before his choice is made. Here, we primarily investigate the scenario where the circumstance is defined to be last period's weather: (w_{t-1}) . However, we also investigate alternative scenarios, such as weather including a few lags, i.e. $(w_{t-1}, w_{t-2}, w_{t-3})$. Let W be defined as the set of possible circumstances, with typical element \vec{w} .

As defined above, a *case* is a triplet consisting of a circumstance, the action taken in response to that circumstance, and the result of that action. So let \mathcal{C} be the set of all possible cases:

$$\mathcal{C} = W \times Q \times \Pi$$

with typical element $(\vec{w}_{t-1}, q_{it}, \pi_{it})$.

The *memory* of a firm is the set of all cases that they have experienced; at time 1, that memory is empty, and in each subsequent period, exactly one case is added to memory. For example, suppose firm i is called upon to choose quantity

q_t , and suppose the circumstance is defined to be the previous two periods of the state variable. Then her memory when making her time t decision consists of:

$$\begin{aligned} \mathcal{M}_{it} = \{ & ((w_0, w_{-1}), q_{i1}, \pi_{i1}), \\ & ((w_1, w_0), q_{i2}, \pi_{i2}), \\ & \dots \\ & ((w_{t-2}, w_{t-3}), q_{it-1}, \pi_{it-1}) \} \end{aligned}$$

Note that, because of the overlapping lag structure, each observation of the state variable appears twice, just as it would in a time-series regression with overlapping lags. From this point of view, the cases in memory can be thought of as observations in a time-series data set, where the circumstance and action are assumed to cause the result.

Given this structure, and following the definitions in Section 2, we can define the case-based utility for the firm in this setting. Suppose that the agent has a memory \mathcal{M} and a faces a problem/circumstance \vec{w}_t , and must choose $q_t \in Q$. The agent constructs a measure of the desirability of each $q \in Q$ called case-based utility. The case-based utility of some $q \in Q$ is:

$$\begin{aligned} CBU(q) &= \sum_{(\vec{v}, q, \pi) \in \mathcal{M}(q)} s(\vec{w}_t, \vec{v}) [u(\pi) - H_t] \\ \text{where } \mathcal{M}(q) &= \{(\vec{v}, q, \pi) | (\vec{v}, q, \pi) \in \mathcal{M}\} \\ \text{and } H_t &= \delta^t H_0 \end{aligned} \tag{16}$$

The similarity function, as described in Section 2, is assumed to be inverse Euclidean distance: $s(\vec{w}, \vec{v}) = e^{-d(\vec{w}, \vec{v})}$. The utility function is assumed, for simplicity, to be risk-neutral profit: $u(\pi) = \pi$.

4 Results

4.1 Set-up

The macroeconomic parameters of this model are adapted from Arifovic (1994). Arifovic (1994) evaluates the convergence properties of the genetic algorithm, in a simple cobweb model without a state variable. We choose a similar setting as the second parameter set described in that paper. On the cost side, we assume $a = 2$, $b = 1$, and $c = 2$, so that the cost function of all firms, given a value of the state variable w_{t-1} , is:

$$c(q) = 2q + \frac{1}{2}q^2 - 2w_{t-1}q \quad (17)$$

The values of a and b are from that paper; c is chosen to be equal in magnitude to a .

On the demand side, we assume $\theta = 10$, and we let $\beta = .03$ in the base case. These values follow Arifovic (1994).

$$p_t = 10 - .03 \sum_{i=1}^n q_{it}, \quad (18)$$

Primary object of inquiry is the effect of the annealing rate δ on convergence, so we test values of δ ranging from close to zero, to 1. The base case is $\delta = .99$.

The secondary of inquiry is the implication of the slope of demand. As mentioned above (Section 3.2), the slope of demand has implications for the learnability of the rational expectations solution. In the simple cobweb model, $\beta > 0$ is learnable, and $\beta < 0$, or upward-sloping demand, is not learnable. However, under econometric learning, $\alpha = -\frac{n\beta}{b} < 1$ is learnable, while $\alpha > 1$ is not. As a consequence, we test three levels of β : $\beta = 0.03$, which is the base case and corresponds to normal, downward-sloping demand; $\beta = -0.03$, which corresponds to upward-sloping demand, but not demand that is so strongly upward sloping as to undermine econometric learnability; and $\beta = -1.2$, which is upward-sloping demand that disrupts econometric learning.

We evaluate two distributions of the state variable w_t . The first is a ‘simple’ formulation. In the simple formulation, we assume that w can take on three possible values; -1 , 0 , and 1 . The purpose of this formulation is investigate a form

of the state variable that readily provides a way to evaluate whether some levels of the state variable are easier to learn than others. It also provides a way to measure convergence, even if that convergence is away from the REE price: we can look at the prices that emerge for each discrete level of the state variable and evaluate price variance for each level: a lowering variance by state variable level would indicate some kind of convergence. The second is the ‘normal’ formulation, where it assumed that:

$$w_t = .9w_{t-1} + \epsilon_t \tag{19}$$

where ϵ_t is assumed Normal with a mean of zero and a variance of one. The purpose of this formulation is to investigate a form of the state variable that is more realistic and more typical of other macro models.

4.2 Presentation and Interpretation of Results

4.2.1 Benchmark result

Our primary interest is to understand the behavior of equilibrium prices under case-based learning. We focus on two questions: does agent learning eventually cause the market price to converge? If it does, does it converge to the rational expectations equilibrium price or some other values?

We start by examining the simulation results from our simplest setup, in which there are no exogenous shocks to the system and therefore no intrinsic uncertainty. The only source of price fluctuations is agents’ learning behavior. The simulation is run for 2000 periods, and is repeated 50 times. We find that with our benchmark calibration, the answers to the two questions are both positive. In all 50 simulations, the agents gradually learn to choose one optimal quantity of production, which results in one unique market price. Moreover, this price invariably coincides with the rational expectations solution.¹⁰ In Figure 3, we plot the market prices from a typical simulation. In the early stages of learning, prices vary quite significantly. But after about 150 periods, agents learn to choose an optimal quantity,

¹⁰In some of our simulations, prices converge to the REE with a small error. Their values are away from the REE values by a very small margin.

and the price converges to the REE price.

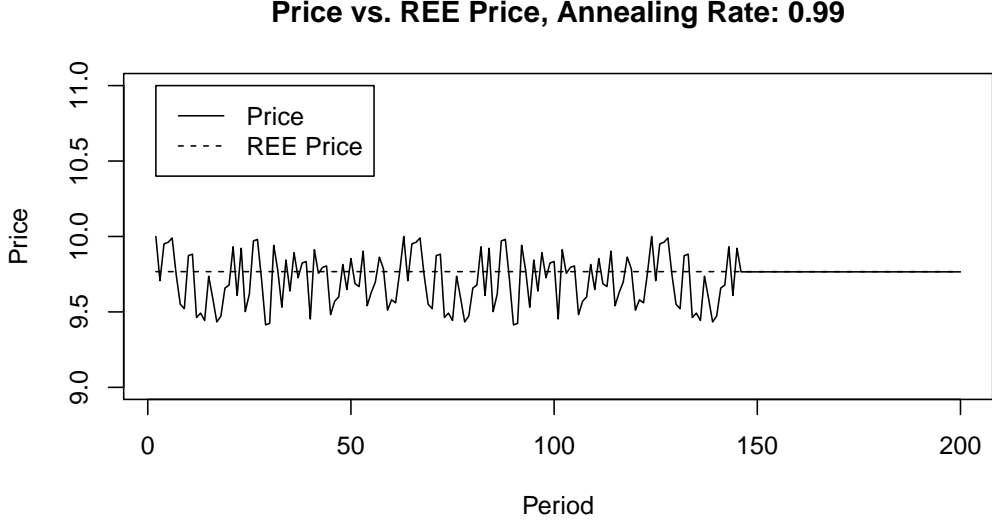


Figure 3: Convergence to the REE. *Annealing rate* $\delta = 0.99$.

Recall that in the case-based decision theory, an important element of agent behavior is the aspiration level. A high aspiration level requires that agents be satisfied with their choices only when the payoffs are relatively high, while a low aspiration level relaxes such constraints. In our model setup, we let agents start with a high aspiration level (a profit level that is high and unattainable), but allow the aspiration level to gradually decrease. The parameter that controls the rate of aspiration reduction is the annealing rate. Our benchmark calibration for the annealing rate is 0.99, which is analogous to a discount rate of 0.99 in a utility-maximizing framework. Thus, our agents will experiment with different choices in the early stages of the simulation, because they are not easily satisfied with their profits. Over time, they will learn from experience what levels of profits are feasible. As the aspiration level gradually decreases, agents will eventually settle with a choice that is both feasible and meets their aspiration. This is the underlining mechanism that leads to convergence. The fact that the equilibrium price coincides with the REE price is quite remarkable, because agents' learning behavior is completely heuristic and experience-based – there is no mechanical

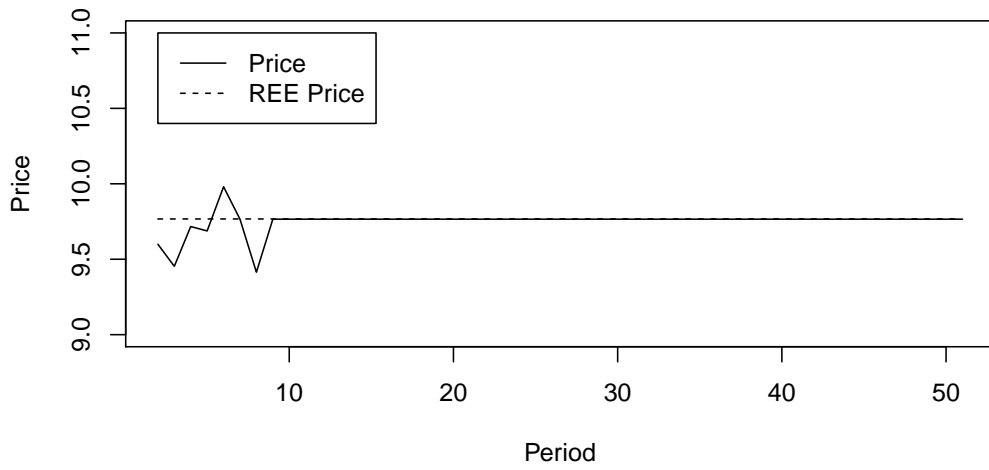
resemblance between case-based learning and the rational expectations framework. It is the structure of the economy that leads up to the convergence to the REE.

A natural inference we can draw from the above analysis is that there cannot be convergence to the REE at all levels of the annealing rate. When the annealing rate is 1, for example, agents will not be content with their search unless their profit reaches the predefined unattainable target. Consequently, their experiments with different choices will never end, and there will never be convergence. Our simulations confirm this point. What about annealing rates between 0.99 and 1? Our experiments show that as long as the annealing rate is not exactly 1, there will be convergence to the REE price. This is logical because as long as the aspiration level is decreasing over time, it will eventually reach a feasible level, and agents will select an optimal choice from experience and stick to it. What is different is the speed of convergence. It takes longer for the price to converge the higher the annealing rate.

What is not so obvious is what happens when the annealing rate is between 0 and 0.99. We experiment with several lower levels of the annealing rate in our simulations: 0.9, 0.8, 0.7, 0.2, and 0.01. We find that as the annealing rate goes down, different equilibrium outcomes start to occur: the market price still converges to a single value, but it is not necessarily equal to the REE value. Moreover, this value varies in each run of the experiment. The equilibrium price is path-dependent, in the sense that its value is affected by what choices agents experiment with in the early stages of learning. We plot the prices from a typical experiment in Figure 4. The annealing rate is 0.8. As we can see from the lower panel, the price converges to a certain value after just a few trials, and that it is far away from the REE price level. The top panel depicts a situation in which the agents' choices happen to be consistent with the REE price. In this case, convergence to REE happens "by luck." It occurs when agents happen to experiment with the REE price in the first few trials of their learning process.

Thus, when the annealing rate is low, multiple equilibria abound and the REE becomes a special case. This result can be understood as follows. With a low annealing rate, agents' aspiration level drops quickly, and it gets increasingly easy for them to be satisfied with given profit level. If the REE price has been considered before they are satisfied with a given profit level, they will choose to stick to the

Price vs. REE Price, Annealing Rate: 0.8



Price vs. REE Price, Annealing Rate: 0.8

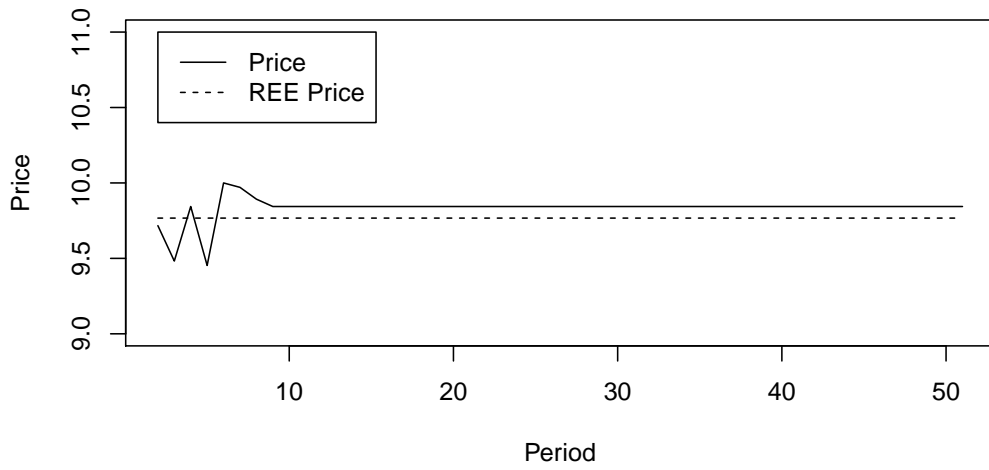


Figure 4: Lower annealing rate $\delta = 0.8$. *Converges quickly, sometimes to REE*

REE price. But if they have never tried the REE price when their aspiration level is sufficiently reduced, they will select the price that delivers the highest profit, and stick to it. This is also why the equilibrium price is path-dependent: the price that agents eventually choose depends on what prices they have experimented with in those periods before their aspiration is satisfied.

Our agents are subject to two different aspects of “bounded rationality.” The first is that they do not have any structural knowledge of the economy, and learn about their optimal choices heuristically by comparing cases in their memories. This is universal to all our learning agents. The second aspect is unique to the low-aspiration agents. There are more price and profit information to be explored, but they choose to be satisfied with their given level of profit, and stop searching for a better outcome. This second aspect is responsible for the existence of multiple equilibrium prices in the Cobweb model. In the macroeconomics literature, such behavior has been associated with “restricted perception” or cognitive limitations. Hommes and Zhu (2012), for example, argue that even if agents observe all the useful information relevant to the economic system, they simply may not possess the knowledge that associates their observations with what they are trying to forecast. An example they give is asset price forecasting. Even if consumers observe the fundamental shocks to the economy, they may not understand that stock prices are a function of them. Therefore, they forecast stock prices using only simple behavioral rules. Adam (2007)’s experimental work is consistent with this argument: in his experiments, the forecasting rule that best describes the subjects’ forecasting behavior is a simple AR(1), despite the fact that the subjects were given much richer information that could potentially improve the forecasts. Another possible reason for agents to forgo useful information is that information is costly to obtain and process. This point is made forcefully by Mankiw and Reis (2002), and has led to a series of research into the importance of sticky information.

4.2.2 Robustness

We add uncertainty to the economy. First we consider the case where the state variable w takes three values: $-1, 0$, and 1 . All simulations are run under the benchmark calibration. When there are multiple states of the world, the REE

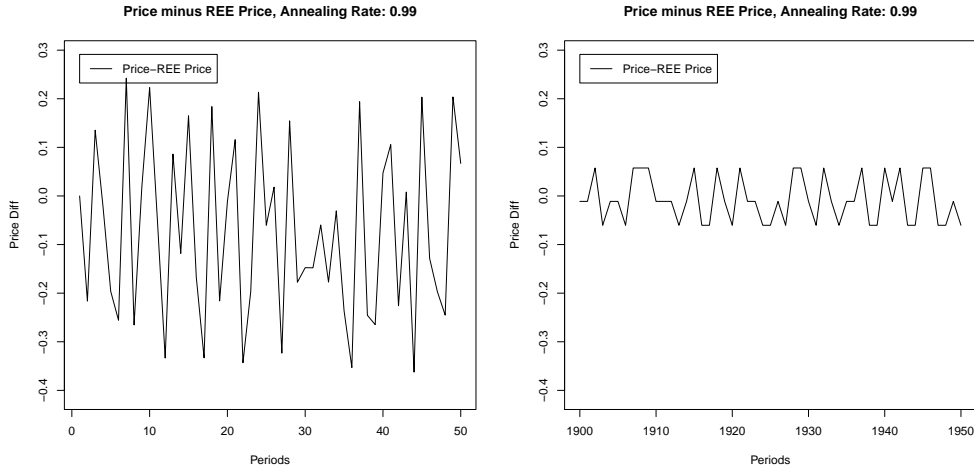


Figure 5: Approximate Convergence of Price to REE price: annealing rate 0.99

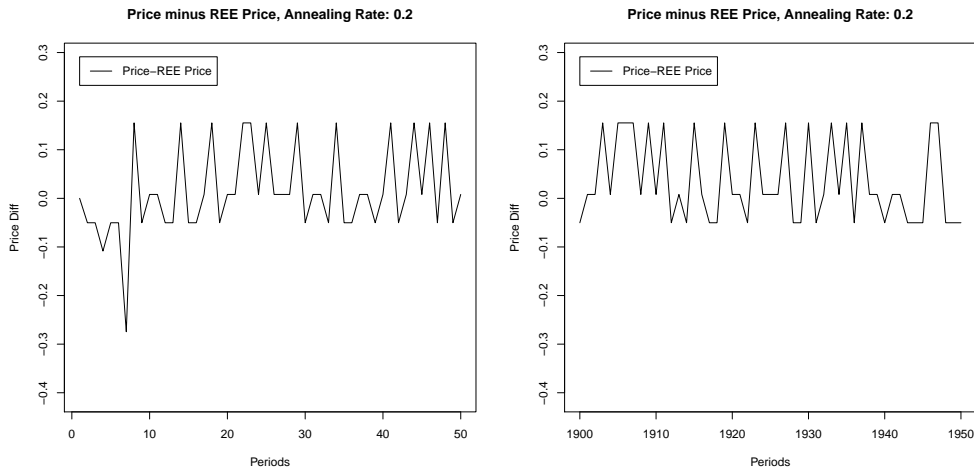


Figure 6: No Convergence of Price to REE price: annealing rate 0.2

price is no longer a single value. It is a function of the state variables. There is one REE price for each state.¹¹ We need a strategy to understand whether or not learning has resulted in the convergence to an equilibrium. A direct approach is to check how many different prices are charged at the beginning of the simulation, and how many are charged after learning takes place for a while. Specifically, we check the number of prices in the first vs. last 100 periods for different runs of the simulation. We find that in all 50 simulations, agents experimented with all of the 61 different prices in the first 100 periods, but in the last 100 periods, agents settle down with only 3 prices in 35 of our 50 simulations, and 2 prices in the remaining 15 simulations. Clearly, they have learned to choose only 2 or 3 optimal quantities via learning.

An alternative way to look at this issue is to examine price volatilities state by state. Suppose over time, agents learn that a certain quantity of production is optimal in a given state, then they will choose the same quantity whenever the economy is in that state. This decision will be reflected in market prices: after convergence occurs, we should observe zero or very low volatility for prices given each state.

Our results are presented in Table 1. We compare the standard deviation of prices in the first 100 periods of the simulation and the last 100 periods, for each state. All numbers in the table are averages from the 50 runs of the simulation. The table shows that when $w = -1$, the volatility of prices is 0.1659 in the first 100 periods, and is 0 in the last 100 periods. The same pattern holds for $w = 0$ and $w = 1$. In these two cases the volatility does not decrease to exactly 0 in the last 100 periods, but are extremely small in values (0.04 and 0.018). Essentially, the economy has converged to a dynamic equilibrium in which the equilibrium prices are functions of the state variables. How close is the equilibrium to the REE? In the last column of Table 1, we compute the standard deviations of the gap between actual prices and the REE prices, and compare them across the two sub-periods. The standard deviation decreases from 0.17 in the first 100 periods to about 0.05 in the last 100 periods. This indicates that equilibrium prices are getting increasingly closer to the REE prices.

¹¹In models with multiple equilibria, there may be more than one equilibrium price in each state. The Cobweb model does not have multiple REEs.

Annealing Rate δ	Period Range	Std. Deviation			
		from REE	$w = -1$	$w = 0$	$w = 1$
1	1-100	0.1812	0.1666	0.1675	0.1655
	1901-2000	0.1726	0.1661	0.1636	0.1601
0.999	1-100	0.181	0.1631	0.1699	0.1646
	1901-2000	0.054	0.0585	0.0015	0
0.996	1-100	0.1806	0.1652	0.1664	0.1661
	1901-2000	0.0444	0	0.0023	0.0032
0.993	1-100	0.1798	0.1636	0.1697	0.1587
	1901-2000	0.0482	0	0.0411	0.0159
0.99	1-100	0.1728	0.1659	0.1668	0.1388
	1901-2000	0.0521	0	0.0427	0.0188
0.9	1-100	0.0971	0.1042	0.0882	0.1077
	1901-2000	0.0685	0.034	0.0508	0.1024
0.7	1-100	0.1047	0.0775	0.0889	0.1534
	1901-2000	0.0954	0.0484	0.0774	0.1525
0.2	1-100	0.1032	0.0819	0.1114	0.1035
	1901-2000	0.0997	0.0618	0.1073	0.1053
0.01	1-100	0.113	0.0913	0.1317	0.1284
	1901-2000	0.1075	0.0674	0.127	0.1305

1 firm, $\beta = .3$.

Table 1: Effect of Annealing Rate on Convergence

Figure 5 presents a time series plot of the deviation of equilibrium prices from the REE prices in a typical simulation with an annealing discount rate of .99. In the first 50 periods (left plot), the differences between the equilibrium prices and the REE prices are much larger than those in the last 50 periods (right plot). Contrast this with Figure 6, in which the annealing rate is .2; there is little reduction in the variance of difference between the price and the REE price over time.

These results are similar to those from our benchmark simulations: agents are able to learn to optimize their choices in each state of the economy; moreover, their choices resulted in the asymptotic convergence of the equilibrium prices to the REE prices.

We next do the same experiment with lower annealing rates. We find that there is clear evidence of price convergence for all levels of the annealing rate. For example, when the annealing rate is 0.7, there are typically 4-9 different market prices in the first 100 periods of the simulation. In the last 100 periods, this number is reduced to 1-3, with 2 and 3 prices being the most frequently occurred outcome. Are these prices close to the REE prices? In Table 1, we can see that for annealing rates 0.7, 0.2, and 0.01, the standard deviations of the price-REE gap are virtually the same for the first and last 100 periods of the simulation. It indicates that the market prices generally do not converge to the REE prices. This is also reflected in the state-by-state volatility of prices. The table shows that for each value of the state, there is virtually no reduction in the standard deviation of the market prices. We know that the number of market prices have been reduced via learning. But because there are multiple equilibria, and these equilibrium prices are quite far away from each other, the standard deviations for some of our 50 simulations are large. When averaged over all 50 simulations, the standard deviation remains high.

Finally, we consider the case where the state variable follows an AR(1) process:

$$w_t = .9w_{t-1} + \epsilon_t. \quad (20)$$

This is an assumption often used in macroeconomics for this and similar type of models. It creates richer dynamics for the state variable, and is likely to generate more sophisticated price movements. We run the same simulations with this new

setup, and re-examine all our results. We find that our conclusions still hold. Convergence to the REE takes place with high annealing rates, and multiple equilibria occur with lower annealing rates. In Figure 7, we plot the market prices vs. the REE prices for the low annealing rate case (upper panel) and the high annealing rate case (lower panel). In the lower panel, market prices trace the REE prices fairly well, while in the upper panel, market prices have little correlation with REE prices.

4.2.3 A no convergence result

An interesting result from the bounded rationality literature is that convergence to the REE price is affected by the values of some structural parameters. For example, Evans and Honkapohja (2001) show that with econometric learning, there is no convergence to REE prices if $\alpha < 1$. Our following experiment is inspired by their result.

We let the slope of the demand curve change from a regular negative value to an increasingly larger positive value, and examine the behavior of prices. Since $\alpha = \frac{\beta}{b}$, and $b = 1$, α turns positive when the slope of the demand curve is higher than 1. Our results are presented in Table 2.

With all three levels of the slope of demand, prices do eventually converge. This is evident from the sharp reduction in the state-by-state volatilities. These prices, however, are very far from the REE prices. This is clear from the high standard deviations of the price-REE gaps in the last 100 periods of the experiment.


Therefore, case-based learning makes the same prediction as econometric learning in terms of the stability of REEs in the Cobweb model: the REE price is not learnable if the slope of the demand curve is higher than 1.

4.2.4 The number of firms

The above analysis is conducted under the assumption that there is a representative firm. Next, we relax this assumption and allow multiple firms (agents) to learn and make decisions independently. We experiment with $n = 1, 2, 20$, and 50. All other parameters are calibrated the same way as described above.

Table 3 shows the result. With more firms in the market, our main conclusion

Price vs. REE Price, Annealing Rate: 0.01



Price vs. REE Price, Annealing Rate: 0.999

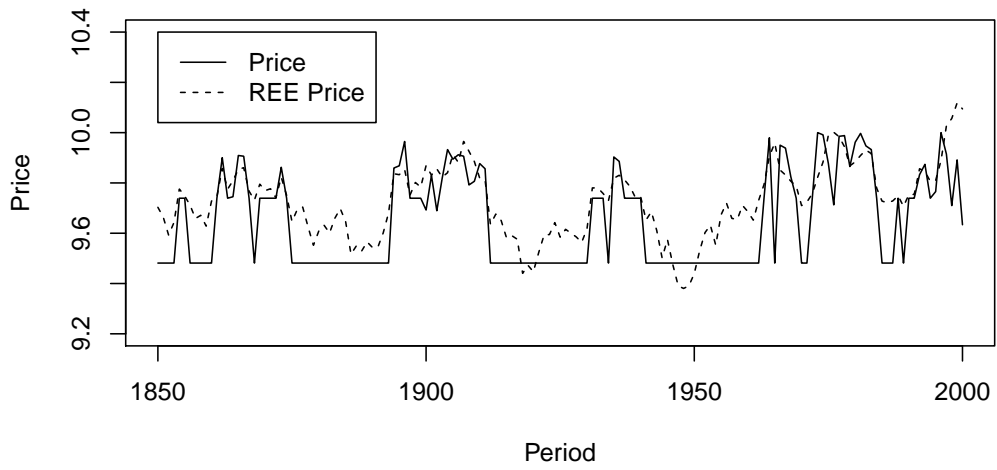


Figure 7: Convergence of Price to REE price: Normally-distributed state variable

Coeff. on Demand β	Period Range	Std. Deviation			
		from REE	$w = -1$	$w = 0$	$w = 1$
0.3	1–100	0.1737	0.1679	0.1645	0.1435
	1901–2000	0.0522	0	0.0435	0.0236
-0.3	1–100	0.1818	0.1738	0.1739	0.1509
	1901–2000	0.0666	0.0048	0.0523	0.0249
-1.2	1–100	114.9014	42.0486	41.8736	40.8387
	1901–2000	138.1442	1.2745	1.7441	2.1582

1 firm, $\delta = .99$.

Table 2: Effect of Demand Slope on Convergence

Number of Firms	Period Range	Std. Deviation			
		from REE	$w = -1$	$w = 0$	$w = 1$
1	1–100	0.1738	0.1681	0.1651	0.1433
	1901–2000	0.0522	0	0.0466	0.0192
2	1–100	0.2673	0.236	0.2414	0.2079
	1901–2000	0.1033	0.0093	0.0515	0.0503
20	1–100	1.5581	0.8375	0.8141	1.342
	1901–2000	0.5713	0.0398	0.1139	0.1552
50	1–100	3.1713	1.8934	1.6228	2.3911
	1901–2000	0.737	0.0331	0.1508	0.3508

$\beta = .3$, $\delta = .99$.

Table 3: Effect of Number of Firms on Convergence

continues to hold. Prices converge asymptotically to REE prices over time. The only difference is that the volatility of the price-REE gaps seems to increase as the number of firms increases. This indicates that more decision-makers add to the variety of prices charged and prolonged the convergence process.

5 Conclusion

In this paper, we propose modeling bounded rationality in macroeconomics using Case-based Decision Theory (Gilboa and Schmeidler, 1995). Case-based Decision Theory was developed to be a more realistic decision theory than expected utility by dispensing with the requirement that agents know with the entire state

space of the problem. CBDT has received some recent empirical support as a model which can explain individual choice behavior in psychological and economic experiments and so can be thought of as providing a model of how an ‘average person’ might respond to a generic choice situation. Some of the empirical support comes from work using the Case-Based Software Agent (Pape and Kurtz, 2013; Guilfoos and Pape, 2014), which we use in this paper to simulate learning by firms. When a Case-Based Software Agent encounters a new case and a decision needs to be made, she looks into her memory bank for cases that are *similar* to the current case. By comparing the payoffs associated with past cases, she selects a choice that delivers the best payoff in similar cases. She then accumulates cases in memory as she makes more choices, and it is this accumulation of memory which drives learning.

When we apply this case-based learning approach to the Cobweb model, we find that agent patience is the key parameter which determines the convergence properties of the economy. ‘Patience’ is captured by a parameter called the *annealing rate* δ . The annealing rate is analogous to the utility discount rate β . For an intermediate level of δ , around .99, we find that the economy converges to an equilibrium which corresponds to the rational expectations equilibrium. When agents are insufficiently patient (δ is too low) we find that agents converge much quicker to an equilibrium, but the price level which results may be far from the REE price. In this case, there are multiple equilibria in the economy, and they are path or history dependent. On the other hand, we find that when agents are too patient, that is, when δ is too high ($\delta = 1$), agents continue to search for an optimal quantity and there is no convergence in prices at all. We also find that the slope of the demand curve also plays a role in determining the convergence properties: like econometric learning (Evans and Honkapohja, 2001), we find that the rational expectations equilibrium is learnable under Case-based Decision Theory as long as the demand curve has a slope less than one.

Our result suggests that in the Cobweb economy, the case for a unique rational expectations equilibrium weakens when agents are heuristic learners. Multiple equilibria abound, and the unique equilibrium is a special case. Due to its usefulness in modeling macroeconomic models and its microeconomic empirical foundations, Case-based decision theory is a learning method which should be in-

corporated into future work in macroeconomics. It has other qualities that we seek to explore in future work: for example, social learning can be implemented by case-based agents sharing parts of their memories. This simulation setting also allows for ‘scaling up’ the economy to a larger, more complicated economy with heterogenous agents. Consumer or worker agents could also be case-based in a more complete model of the macroeconomy.

References

- Klaus Adam. Experimental evidence on the persistence of output and inflation. *The Economic Journal*, 117:603–636, 2007.
- J. Arifovic. Genetic algorithm learning and the cobweb model. *Journal of Economic Dynamics and Control*, 18:3–28, 1994.
- Antoine Billot, Itzhak Gilboa, and David Schmeidler. Axiomatization of an exponential similarity function. *Mathematical Social Sciences*, 55(2):107–115, 2008.
- I.-K. Cho. Perceptions play the repeated prisoner’s dilemma. *Journal of Economic Theory*, 67:266–284, 1995.
- I.-K. Cho and Thomas J. Sargent. *Handbook of Computational Economics*, chapter Neural Networks for Encoding and Adapting in Dynamic Economies. Elsevier, 1996.
- George W. Evans and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University press, 2001.
- Gabrielle Gayer, Itzhak Gilboa, and Offer Lieberman. Rule-based and case-based reasoning in housing prices. *The B.E. Journal of Theoretical Economics (Advances)*, 7(1), 2007.
- I. Gilboa and D. Schmeidler. Case-based Optimization. *Games and Economic Behavior*, 15:1–26, 1996.
- Itzhak Gilboa and David Schmeidler. Case-Based Decision Theory. *The Quarterly Journal of Economics*, 110(3):605–39, August 1995.

- Vasyl Golosnoy and Yarema Okhrin. General uncertainty in portfolio selection: A case-based decision approach. *Journal of Economic Behavior & Organization*, 67(3):718–734, 2008.
- Todd Guilfoos and Andreas D Pape. Predicting cooperation in the prisoner’s dilemma game using case-based decision theory. 2014.
- C. Hommes and M. Zhu. Behavioral learning equilibria. CeNDEF Working Paper, 2012.
- David Hume. *Philosophical essays concerning human understanding*. A. Millar, 1748.
- Edi Karni. A theory of bayesian decision making with action-dependent subjective probabilities. *Economic Theory*, 48(1):125–146, 2011.
- S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):pp. 671–680, 1983. ISSN 00368075.
- Robert E., Jr. Lucas. Expectations and the neutrality of money. *Journal of Economic Theory*, 4:103–124, 1972.
- R. D. Luce and H. Raiffa. *Games and Decisions*. Wiley, New York, 1957.
- N. Gregory Mankiw and Ricardo Reis. Sticky information versus sticky prices: A proposal to replace the new Keynesian Phillips curve. *Quarterly Journal of Economics*, 117:1295–1328, 2002.
- John F. Muth. Rational expectations and the theory of price movements. *Econometrica*, 29:315–335, 1961.
- Wolfgang Ossadnik, Dirk Wilmsmann, and Benedikt Niemann. Experimental evidence on case-based decision theory. *Theory and Decision*, pages 1–22, 2012. ISSN 0040-5833. doi: 10.1007/s11238-012-9333-4. URL <http://dx.doi.org/10.1007/s11238-012-9333-4>.
- Andreas D. Pape and Kenneth J. Kurtz. Evaluating case-based decision theory: Predicting empirical patterns of human classification learning. *Games and Economic Behavior*, 82:53–65, 2013.

- Andreas Duus Pape. Action-independent subjective expected utility without states of the world. *Theoretical Economics Letters*, 3:17, 2013.
- A. E. Roth and I. Erev. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, 8(1):164–212, 1995.
- Thomas J. Sargent. Rational expectations, the real rate of interest and the natural rate of unemployment. *Brookings Papers on Economic Activity*, 2:429–472, 1973.
- Thomas J. Sargent. *Bounded Rationality in Macroeconomics*. Oxford University press, 1993.
- Thomas J. Sargent. Evolution and intelligent design. *American Economic Review*, 98:5–37, 2008.
- L. J. Savage. *The Foundations of Statistics*. Wiley, 1954.
- R.N. Shepard. Toward a Universal Law of Generalization for Psychological Science. *Science*, 237(4820):1317, 1987.
- R.N. Shepard, C.I. Hovland, and H.M. Jenkins. Learning and Memorization of Classifications. *Psychological Monographs*, 75:1–41, 1961.
- Herbert A. Simon. *Models of Man: Social and Rational; Mathematical Essays on Rational Human Behavior in Society Setting*. John Wiley, 1957.
- Leigh Tesfatsion. *Handbook of Computational Economics*, volume 2, chapter 16. Elsevier B.V., 2006.
- J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.
- Martin L Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.