STUDY CENTER GERZENSEE

Limited Rationality and Strategic Interaction: A Probabilistic Multi-Agent Model

Yves Ortiz and Martin Schüle

Working Paper 11.08

Limited Rationality and Strategic Interaction: A Probabilistic Multi-Agent Model

Yves Ortiz* Martin Schüle†

November 15, 2011

Abstract

We develop a multi-agent framework based on probabilistic cellular automata theory to describe off-equilibrium dynamics in the context of the economic problem of price adjustment in different strategic situations as investigated experimentally by Fehr and Tyran (2001) and (2008). It is found that the main experimental findings, namely suboptimal aggregate behavior in terms of sluggish adjustment after a fully anticipated money shock, can be reproduced and largely explained by the interaction of sophisticated and naive agents. Furthermore, a range of conceptual issues as e.g. the source of endogenous beliefs on the other players rationality is addressed within our multi-agent framework. We find that, if costs/payoffs act as driver of rational behavior, then endogenous beliefs and consequential aggregate behavior are driven by the particular off-equilibrium time-dependent payoff/cost profile rather than by total off-equilibrium payoffs/costs that naive agents face in the respective strategic situation.

*University of Zurich and Study Center Gerzensee, yvesortiz@gmail.com
†Institute of Neuroinformatics, University of Zurich and ETH Zurich

1 Introduction

Nominal rigidities, i.e. the failure of firms to adjust to the frictionless equilibrium price after an exogenous shock, have been a central theme of economic research during the past decades. It is well known that under so-called strategic complementarity, i.e. when an agent's level of activity depends positively on the other agents' activity, multiple equilibria can arise, and sluggish adjustment of prices after exogenous shocks occur. Research activity on this latter issue has produced a wealth of competing theories to explain such sluggish adjustment of prices. Some theoretical approaches introduce frictions (e.g. costly adjustment, outdated information, costly information), others introduce different sorts of limited rationality (e.g adaptive expectations, rule of thumb agents). Haltinwanger and Waldman (1985) argue that strategic complementarity magnifies deviation from equilibrium disproportionally and leads to more pronounced sluggish adjustment to the new equilibrium allocation. This argument lead to a series of price setting experiments by Fehr and Tyran (2001), in which they report remarkable sluggish adjustment under strategic complementarity after a nominal shock, whereas adjustment takes place almost instantaneously under strategic substitutability. At first, it is not evident why a simple inversion of a nearly linear reaction function as used by Fehr and Tyran (2001), while keeping the post-shock equilibrium allocation unchanged, should have such a large impact on the adjustment process.

A few intuitive explanations of sluggish adjustment have been proposed in the literature, all of them sharing the assumption of the presence of limited rational agents. In our contribution, we suggest a multi-agent approach, based on a probabilistic cellular automata model. Such a, in the economic literature relatively unknown, probabilistic cellular automata model allows for an intuitive set-up of the model, while exhibiting at the same time very well the essential features of the experiment as we will see. In our model, agents play either a adaptive backward looking or a sophisticated forward looking price setting strategy. Sophisticated strategies take into account the presence of adaptive strategies. As time evolves, sophisticated strategies played in the game become predominant leading to convergence to the new equilibrium. Our model is able to reproduce exactly the aggregate adjustment dynamics displayed in the experiment by Fehr and Tyran (2001). In accordance with previous explanations reviewed in the next section, our model illustrates how agents playing sophisticated strategies optimally refrain from choosing the rational equilibrium in response to the presence of agents playing less sophisticated strategies given that the payoff structure exhibits strategic complementarity. According to our model, observed differences in price setting behavior right after the shock are only consistent with different believes depending on the payoff structure. Under strategic complementarity (substitutability), agents expect sophisticated strategies to be played with much lower (higher) probability.

Further, our model allows to effectively track the expected payoffs of either strategy during the entire adjustment process. This shows that although a adaptive strategy is heavily penalized during the first periods under strategic substitutability, it is less costly over the entire adjustment process under strategic substitutability. This fact gives support to the conjecture that, if costs/payoffs act as driver of rational behavior, then influence goes through the particular induced time-dependent payoff/cost structure rather than through total payoffs/costs prevalent in the respective strategic situation. In other words, high immediate costs incurred by adaptive strategies might exert a disciplining effect on agents prone to play a adaptive strategy, while moderate but persistent costs might encourage adaptive behavior.

In summary, we propose a model, based on a probabilistic cellular automata approach, that is set up very simply and intuitively, but nevertheless allows to reproduce exactly the experimental findings of Fehr and Tyran (2001). The model confirms the literature, e.g. Haltinwanger and Waldman (1989), insofar it stresses the role of strategic complementarity/substitutability and limited rationality in explaining nominal rigidities. It further suggests that closer attention has to be paid to time-dependent cost structures induced by the presence of adaptive strategies, if coordination failure is to be explained satisfactorily in stylized price setting experiments.

1.1 Related Literature

Many economic situations involve strategic complementarities in which the optimal action of a single individual is an increasing function of the actions of the others (see Cooper and Haltiwanger (1996), Cooper and Johri (1997) for macroeconomic evidence, see Milgrom and Roberts (1990) for examples related to industrial organization). It has been shown that, departing from quite general assumptions, strategic complementarities lead to multiple equilibria that can be pareto ranked (e.g. Vives (1990), Milgrom and Roberts (1990), Cooper and John (1988)). Apart from questions about the selection and stability of the resulting equilibria, theoretical work has established that strategic complementarities in conjunction with limited rational agents might magnify the response to real and nominal shocks and lead to slower adjustment to equilibrium (e.g. Haltinwanger and Waldman (1989), Haltinwanger and Waldman (1989)). The intuition underlying this theoretical result is the following: under strategic complementarity, rational agents are forced to account for the actions of limited rational agents and to mimic to a certain extent their suboptimal choices.

Fehr and Tyran (2001), tested these theoretical predictions in a price setting experiment. The results of the experiment indicate that the participants have difficulties to coordinate their actions to a new equilibrium after a fully anticipated nominal shock under strategic complementarity. Interestingly, adjustment occurred almost instantaneously under strategic substitutability and, thus, the findings of Fehr and Tyran (2001) confirmed some of the theoretical predictions. In a similar experiment Fehr and Tyran (2008), henceforth referred to as FT, demonstrate how agents effectively build different believes on the other's play depending on the strategic situation.[1] Subsequent experimental work has confirmed FT's findings while refining their results in different aspects (Davis (2009), Davis and Korenok (2010)).

There are hitherto only few explanations of these new experimental findings. The results of Haltinwanger and Waldman (1989) merely state that, given the presence of a fixed fraction of agents with adaptive expectations, the rational agents' optimization implies a choice value close to the value chosen by the adaptive agents. Indeed, Davis and Korenok (2010) find that around 20% of the players behave as price setters with adaptive expectations under strategic complementarity. Camerer and Ho (2004) calibrate their so-called Cognitive Hierarchy Model to predict the outcome of Fehr and Tyran (2001), which is accomplished by assuming that agents' forward reasoning involves more time steps under strategic substitutability than under strategic complementarity. Davis (2009) and FT execute simple simulations with agents following ad-hoc behavioral rules which cannot fully account for the adjustment dynamics observed in the experiments.

After all, it remains at first sight puzzling why people evidently form different believes on the other people's rationality depending on the strategic situation. FT presume that believes might be endogenous with respect to the payoff structure in that the mistake made by an adaptive strategy is salient under strategic substitutability. As a consequence, players playing an adaptive strategy will recognize their mistake more easily and, this fact being common knowledge, believes and actions will converge quickly to equilibrium.

In contrast to the existing explanations, our model emphasizes the role of off-equilibrium payoffs/costs by quantifying explicitly the costs from deviation (i.e. not playing best response) incurred by either strategy during off-equilibrium time. Our model is parsimonious and only a few simple behavioral assumptions are required: agents playing sophisticated, forward looking strategies take into account the presence of adaptive strategies and maximize their expected payoffs during off-equilibrium play. Adaptive strategies adopt the realized average price as their pricing strategy for the next period. Further, the fraction of either strategy played is not fixed but evolves endogenously such that the number of agents playing adaptive strategies tends to zero as the average price converges to the new equilibrium level. The interaction of agents playing these two strategies generates the adjustment process observed in FT. The resulting time dependent off-equilibrium payoff/cost distributions can then also shed light on the possible sources of endogenous belief forming.

[1] Players form different believes when they know that they play against fully rational computerized agents. It turns out that believes and average price then converge quickly to equilibrium under complementarity as well.

1.2 The Fehr and Tyran Experiment

FT's experiment was conducted with several groups of n persons. Each person plays a price setting game choosing a price (i.e. acting as a firm) and getting a payoff in the next period. The game extends over $T = 30$ periods and the payoffs are realized according to a payoff table distributed to the players at the beginning of the game. Player i can easily infer that his own payoff, Π_i^{t+1}, strongly depends on his pricing decision, P_i^{t+1}, and the average price of all other players in the next period, P_{-i}^{t+1}. Payoff tables entail an unique Nash-Equilibrium in the pre-shock and post-shock period. Throughout the first $T/2$ periods, the game is played with the first payoff table and all players set the equilibrium price P_0^* after a short adjustment phase. In period $T/2$, a new payoff table is distributed to the players, identical to the first table up to the fact that all payoffs and prices have been divided by two. Since the players know that the game will be played according to the new table from period $T/2 + 1$ until the final period T, the subjects face an exogenous and fully anticipated nominal shock with the new equilibrium price $P_1^* = P_0^*/2$. Figure 1 visualizes the different adjustment paths. FT also run a test for statistical significance: Under complementarity, the hypothesis that the average price P_t is equal to the equilibrium price P_1^* can be rejected at the 1% level during the first 8 periods after the shock. Under substitutability however, the hypothesis that P_t is equal to the equilibrium price P_1^* can never be rejected. The key points of FT's experimental set-up are thus the following: First, all agents have full information about all relevant aspects of the game and they are free to set any price for which a payoff is defined without any further restrictions. Second, players are rewarded according to real payoffs, $\pi_i^t = P_i^t/\bar{P}^t$, that is, a nominal shock does not alter the equilibrium payoffs. Third, under the rational agent hypothesis, players should realize that a nominal shock leaves the real payoffs of the game unchanged and choose the new equilibrium price once the nominal shock becomes effective.

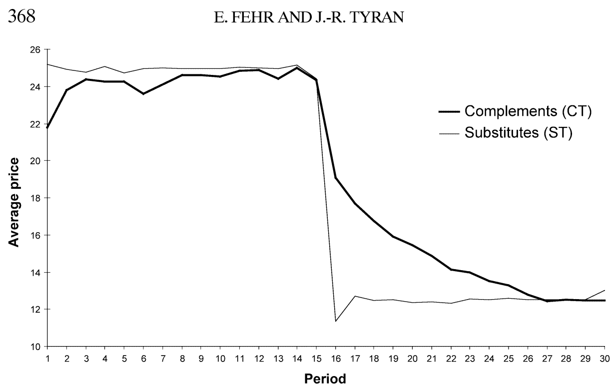

Figure 1: Nominal Average Prices. Source: Fehr and Tyran (2008)

2 Probabilistic Choice Approach

We employ a probabilistic choice approach to reproduce and analyze the main features of the adjustment dynamics reported in the experiment by FT. In the experiment, every agent's choice depends crucially on his believes about the other agents' price setting choices; the average of all agents' choices becomes common knowledge at the end of every period through the realized average price. This particular feature in FT's experiment led us to model the agents decision procedure by means of a simple probabilistic choice rule where the choice probability of an agent depends on the average choice of all other agents. More precisely, we model the choice procedure with a probabilistic cellular automata model where the choice probability of the next period depends on the agent's current choice and the current observed average price.

In the next section, we introduce probabilistic cellular automata and motivate the probabilistic choice rule we use. We also give an economic interpretation of the probabilistic choice rule, based on the related models of myopic learning.

2.1 Probabilistic Cellular Automata

The approach we use to model FT's experimental findings is based on a probabilistic cellular automata model. We first introduce cellular automata of which the probabilistic cellular automata are a stochastic generalization.

Cellular automata (CA) are discrete dynamical systems. The state of a CA is specified by a set of discrete values of the sites of a regular, uniform infinite lattice, i.e. by the states of the "cells" of a CA. A CA evolves in discrete time steps with the states of the cells determined by a rule acting on a neighbourhood of each cell. Cellular automata (CA) models have been widely studied and applied in physics, biology and computer science, see e.g Kari (2005) or Wolfram (2002). In the economic literature, there seem so far only few applications of CA, as e.g. a CA model for the dynamics of stock markets by Qiu et al. (2007).

We introduce here further only the restricted class of one-dimensional CA. The one-dimensional CA is further assumed to be finite, that is the lattice L of the CA is assumed to be of finite length N with periodic boundary conditions, i.e. $L = \mathbb{Z}/N\mathbb{Z}$. The configuration of such a finite, one-dimensional CA is given by the sequence $x = (x_i)_{i \in L}$ with $x_i \in S$ being elements of some finite set of states $S = \{0, 1, ...\}$. If the state set S is given by $S = \{0, 1\}$, we speak of a binary CA. The configuration space X is the set of all sequences x, i.e. $X = S^N$. The global map F is the map $F : X \to X$ where the local transition function is the map $f : S^{2r+1} \to S$, $r \geq 1$, with $F(x)_i = f(x_{i-r}, ..., x_i, ..., x_{i+r})$. The integer r specifies the range of the neighbourhood and is called the radius of the CA. The iterations of the map F acting on an initial configuration x completely determines the temporal evolution of x, hence such CA are called deterministic.

The model proposed below is based on a stochastic generalization of the finite, one-dimensional and deterministic CA introduced above, which is called a *probabilistic cellular automata* (PCA). The dynamics of PCA is still generated by the local transition function, but the next state of a cell is now only known with a given probability. Hence, the local transition function becomes a local probability transition function. In the binary case, which is used in our model, the local probability transition function is defined by $p : \{0, 1\}^{2r+1} \to [0, 1]$, where p is the probability that the cell i updates to state 1 at time $t + 1$ given the states of its neighbourhood (x_i^t) at time t, i.e. $p_i^t = p[x_i^{t+1}|(x_i^t)]$. The local probability transition function is subject to the normalization condition $\sum_{x_i^{t+1}=\{0,1\}} p[x_i^{t+1}|(x_i^t)] = 1$. The global map of a PCA is then given by $P(x^{t+1})_i = B_i^t(p[x_i^{t+1}|(x_i^t)]), \forall i \in L$, where B_i^t is a series of independent Bernoulli random variables, i.e. $B_i^t(p)$ is a random variable equal to 1 with probability p and 0 with probability $1 - p$. It is understood that x_i^t denotes the random variable given by observing the state of cell i at time t.
In general, the dynamics of a PCA is given by a master equation (see e.g. Van Kampen (1981)) which reads to

$$P(x^{t+1}) = \sum_{x^t} P(x^t) \prod_i p[x_i^{t+1}|(x_i^t)], \tag{2.1}$$

where $P(x^t)$ is the probability to observe the PCA in the state x at the time t and $p[x_i^{t+1}|(x_i^t)]$ is the local probability transition function. Evidently, any (finite) PCA is a *(finite) Markov chain.*

The PCA model proposed below is peculiar insofar the local probability transition function p_i^t depends only on the state of the cell i and on the so-called *density* of the configurations, i.e. $p_i^t = p[x_i^{t+1}|x_i^t, \rho_t]$, where the density is defined as $\rho_t = \frac{1}{N}\sum_i x_i^t$.

2.2 A Probabilistic Choice Rule

Key to our model is the mapping of FT's experimental set-up to a setting which can be handled within the mathematical framework of PCA. In FT's experiment, there are two states in the economy, say $\{s_0, s_1\}$, represented by the two different payoff tables. We identify the N agents in the economy with the N cells of a PCA. The agents in the economy can choose between two prices, denoted by $\{x_{i0}, x_{i1}\}$. Choosing the prices is identified with the states of the cells (agents) of the PCA and consists in choosing the old or the new equilibrium prices, P_0^* or P_1^* respectively. According to FT, the prices are given by the values $\{P_0^*, P_1^*\} = \{25, 12.5\}$, which are encoded in our model as $\{0, 1\}$. We fix the time when the shock occurs ($s_0 \to s_1$) at $t = 0$. Agents observe the realized average price at t=0,1,2,... and decide on their price setting actions shortly afterwards, namely at $t = t + \epsilon$ with $0 < \epsilon \lll 1$. For time periods $t \geq 0$, the individual choice probabilities p_i^t are assumed to be governed by the following rule:

$$p_i^t = \begin{cases} \rho_t + (1 - \rho_t)\hat{\beta} & \text{if} \quad x_i^t = 1 \\ \rho_t + (1 - \rho_t)\hat{\alpha} & \text{if} \quad x_i^t = 0, \end{cases} \tag{2.2}$$

where ρ_t denotes the realized density in period t and β and α take values in the interval $[0, 1]$. The probabilistic rule (2.2) specifies a PCA with the local probability transition function p_i^t, since it can be written as $p_i^t = P[x_i^{t+1} = 1 | x_i^t, \rho^t]$. The initial transition probabilities p_i^0 are exogenously given. As shown later, the parameters α and β are determined such that they reflect an "adaptive" or "sophisticated" strategy. For the adaptive strategy, the parameter α is determined such that agents with $x_i^t = 0$ take the observed price in the economy at time t as their expected choice for the next period $t + 1$. For the sophisticated strategy, the parameter β is determined such that agents with $x_i^t = 1$ maximize their expected payoff for the entire adjustment period, given the payoff structure of the game. From now on, we refer to an agent as a sophisticated agent (SA) if he chooses the sophisticated strategy, i.e. if $x_i^t = 1$, and as an adaptive agent (AA) if he chooses the adaptive strategy, i.e. if $x_i^t = 0$. In the following section, we give an economic interpretation of our probabilistic choice rule.

2.3 Economic Interpretation

Our model shares several features with a model of myopic adjustment proposed by Droste et al. (2003) and Kosfeld et al. (2002). In their model, players increase the probability of playing strategies which are a best reply to the actions played by the other players. This assumption is related to the idea of reinforcement learning, where pure strategies that were more successful in past periods are more likely to be played in future periods. If a strategy was not the best response, then Kosfeld et al. (2002) assume that agents play the strategy with the same probability as in the previous period. This particular probabilistic adjustment rule determines a discrete time Markov process for the set of mixed strategies.[2]

In our model, agents playing the sophisticated strategy (i.e. SA) have a positive probability to change their state and thus to play the myopic strategy in the next period; the respective argument applies for agents playing the myopic strategy (i.e. AA). Such an adjustment rule can be interpreted as a simple model of individual learning. In our particular case, agents playing the sophisticated strategy forget their strategy with the probability $1 - [\rho_t + (1 - \rho_t)\hat{\beta}]$ and recall it with the probability $\rho_t + (1 - \rho_t)\hat{\alpha}$. A more in-depth discussion and an additional interpretation based on social learning can be found in Kosfeld et al. (2002).

As shown in the appendix, the fraction of agents playing the sophisticated strategy will only grow (i.e. $E[\rho_{t+1}] > \rho_t$) as long as $\beta\rho_t > \alpha(1 - \rho_t)$ and thus, our rule displays reinforcement for suitable values of ρ_t, β and α. Ceteris paribus, the latter reinforcement condition is more easily satisfied for higher values of ρ_t. That is, the more agents are playing the sophisticated strategy, the more the rule will display reinforcement.

In the next section, we explain in detail the simulation procedure and how the parameters α and β, corresponding to an adaptive and a sophisticated strategy, can be determined.

[2] In a recent paper, Rivas (2008) shows that the probabilistic adjustment rule proposed by Droste et al. (2003) can be viewed as a particular case of a more general class of Markovian learning models (Stochastic Better Response).

3 Simulation Procedure

We describe now the driving process behind the key variable ρ_t and the resulting price formation process. The random variable ρ_t is simulated iteratively departing from an initial value $\tilde{\rho}_0$:

i) Every x_i^t is drawn from a uniform distribution with parameter p_i^{t-1}.

ii) The obtained configuration $x^t = (x_1^t, x_2^t, \ldots, x_N^t)$ allows to compute $\rho_t = \frac{1}{N}\sum_i^N x_i^t$.

iii) ρ_t is used to specify p_i^t.

vi) Iteration is restarted at $t+1$.

Evidently, the parameters α and β in (2.2) are the key determinants of the random variable ρ_t. The parameter values will partly result as solution of an optimization problem described in the next section. In a first step, optimization calculus requires the construction of the price series P_t implied by the price setting decisions of our agents. We define the prevailing price in the economy as

$$P_t = \rho_t P_1^* + (1-\rho_t)P_0^* \tag{3.1}$$

and thus P_t is the weighted average of those agents who choose the new equilibrium ($P_i^t = P_1^*$) and those who choose the old equilibrium ($P_i^t = P_0^*$). The average price P_t is by construction the last stochastic element of the series $\{P\}_t \equiv (P_1, P_2, \ldots, P_t)$ which can be examined through its statistical measures such as its average $\{\bar{P}\}_t$ and quantiles $\{Q\}_t$.[3] A closer look at the relation between the price variable (3.1) and the probabilistic rule (2.2) reveals the underlying intuition. As demonstrated in the next section, the expected payoffs in the period $t+1$ depend on the expected distance between the agent's price and the average price, $|P_{t+1}$-$P_i^{t+1}|$. This distance depends again on the choice of β and α in the choice probability rule p_i^t. If e.g. the group of AA happens to be non-empty, then the expected average price $E_t[P_{t+1}]$ will differ from the new equilibrium price. Under certain circumstances, the presence of agents playing the adaptive strategy will induce the group of SA to refrain from setting their price too far away from the average price during the adjustment phase, causing sluggish adjustment of the entire system towards the new equilibrium price.

In the next section, we put this intuition on more formal grounds. In particular, we describe the ensuing optimization problem of SA facing actions of AA and obtain their optimal response in terms of an optimal choice of β.

3.1 Determining the Model Parameters

The standard approach to derive optimal behavior in a dynamic price setting situation would involve game theoretic concepts. An agent would optimize his pricing decision for the next period subject to the probable actions of all other agents. The complexity of such a task is evident. Therefore, we simplified in our model several aspects of the problem. First, we consider only two types of agents, adaptive agents (AA) and sophisticated agents (SA), both tied to a linear probabilistic choice rule with α and β as the parameters of choice. Second, SA optimize their choice only subject to the actions of the AA, i.e. the choice has the character of an optimal collective choice. In what follows, we formulate and solve SA's optimization problem. We break this task up into three sub-tasks:

1. Construction of a payoff function similar to the payoff function employed in FT.
2. Formulation of SA's optimization problem $\Rightarrow$ deriving $\hat{\beta}$ depending on α and $\tilde{\rho}_0$.
3. Choice of α and $\tilde{\rho}_0 \Rightarrow$ determining $\hat{\beta}$.

[3] A detailed description how $\{\bar{P}\}_t$ and $\{Q\}_t$ are constructed can be found in appendix A.

3.1.1 The Payoff Function

Optimization typically requires the definition of the underlying payoff structure. In our case, we stick as far as possible to the payoff function which FT employ to generate their payoff tables. In fact, we use an accurate linear approximation of FT's payoff function for the sake of simplicity. This simplification leaves our results unaltered since we reproduce the main properties of FT's function and both functions differ from each other only in marginal aspects. The following derivation of the payoff function becomes clearer if read together with the accompanying figures in appendix C.1. As in FT, our payoff function is build along a reaction function rf which gives the optimal individual price response $\hat{P}_t$ for every realization of the average price P_t, that is

$$\begin{aligned}\hat{P}_t &= rf(P_t) \\ &= a + bP_t.\end{aligned}$$

We then define the functional form of the payoffs along the reaction function

$$\Pi^{rf}_{t+1} = \begin{cases} A + \gamma_1 (1+b^2)^{\frac{1}{2}} P_{t+1} & \text{for} \quad P_{t+1} \leq \frac{a}{1-b} \\ A + \gamma_2 (1+b^2)^{\frac{1}{2}} P_{t+1} & \text{for} \quad P_{t+1} > \frac{a}{1-b}. \end{cases} \tag{3.2}$$

The costs of deviating from the optimal response is assumed to be linear in the distance:

$$c_i^{t+1} = \Psi |rf(P_{t+1}) - P_i^{t+1}|. \tag{3.3}$$

Finally, the payoff of an agent i for a given P_{t+1} can be expressed as

$$\Pi_i^{t+1}(P_i^{t+1}, P_{t+1}) = \begin{cases} \Pi^{rf}_{t+1} - c_i^{t+1} & \text{for} \quad \Pi^{rf}_{t+1} \geq c_i^{t+1} \\ 0 \quad \text{else.} \end{cases} \tag{3.4}$$

As shown in the appendix C.2, our payoff function (3.4) provides a good approximation of FT's payoff function once suitable values for the parameters $\{A, a, b, \gamma_1, \gamma_2, \Psi\}$ are chosen.

3.1.2 Optimization Problem

In this section we state SA's optimization problem. We need the following definition.

Definition 1 *Let $E_t[P^{t+1}_{SA}] \equiv E[P_{t+1}|x_i^t = 1, \beta, \alpha, \rho_t]$ and $E_t[P^{t+1}_{AA}] \equiv E[P_{t+1}|x_i^t = 0, \beta, \alpha, \rho_t]$ be the expected price set by an agent, given its type and $\{\beta, \alpha, \rho_t\}$ at time t. Analogously, $E_t[P_{t+1}] \equiv E[P_{t+1}|\beta, \alpha, \rho_t]$ denotes the expected value of the average price given $\{\beta, \alpha, \rho_t\}$.*

Definition 1, together with the linearity of our payoff function (3.4), allows to express the expected payoff for each type of agent as

$$E_t\left[\Pi_i^{t+1}\right] = \Pi_i^{t+1}(E_t[P_i^{t+1}], E_t[P_{t+1}]), \quad i = SA, AA. \tag{3.5}$$

In our particular setting, we are able to compute $E_t[P^{t+1}_{SA}]$, $E_t[P^{t+1}_{AA}]$ and $E_t[P_{t+1}]$. This result is summarized in the following proposition.

Proposition 1 *The expected prices $E_k[P^{t+1}_{SA}]$, $E_k[P^{t+1}_{AA}]$ and $E_k[P_{t+1}]$ can be computed iteratively and depend solely on the parameters α, β and an initial value ρ_k, k=0,1,2,....*

The reader is referred to appendix B for a derivation. Proposition 1 enables us to compute the cumulated expected ("nominal") payoffs of an SA and AA along the entire adjustment path characterized by β, α and $\tilde{\rho}_0$, that is

$$\begin{aligned} E_0\left[\Pi_i(\beta, \alpha, \tilde{\rho}_0)\right] &\equiv E\left[\sum_{t=0}^{T-1} E_t\left[\Pi_i^{t+1}\right] |\beta, \alpha, \tilde{\rho}_0\right] & (3.6) \\ &= \sum_{t=0}^{T-1} E\left[\Pi_i^{t+1}|\beta, \alpha, \tilde{\rho}_0\right] \quad i = AA, SA, & (3.7) \end{aligned}$$

where the second equality follows from Proposition 1, the law of iterated expectations and the fact that our simulation takes $\tilde{\rho}_0$ as initial information. The price level adjusted ("real") payoff is then given by

$$E_0\left[\pi_i(\beta,\alpha,\tilde{\rho}_0)\right] \equiv \sum_{t=0}^{T-1} \frac{E[\Pi_i^{t+1}|\beta,\alpha,\tilde{\rho}_0]}{E[P_{t+1}|\beta,\alpha,\tilde{\rho}_0]} \quad i = AA, SA. \tag{3.8}$$

SA solve the following optimization problem:

$$\max_{\beta} E_0\left[\pi_{SA}(\beta;\alpha,\tilde{\rho}_0)\right]. \tag{3.9}$$

Note that SA maximizes the "real" payoffs, leaving no space for money illusion. We now proceed to pin down the free parameters α and $\tilde{\rho}_0$.

3.1.3 Choice of α and $\tilde{\rho}_0$

To determine the parameter α we resort to a behavioral assumption. AA are assumed to set their price such that $E_t[P_{AA}^{t+1}] = P_t$. Roughly speaking, AA choose the prevailing price P_t as their own price for the period $t+1$. This assumption leads to the following result:

Proposition 2 *$E_t[P_{AA}^{t+1}] = P_t$ is satisfied if and only if $\alpha = 0$.*

As demonstrated in appendix 1, this result follows straightforward from definition 1. In defense of the underlying assumption, we would like to point out the following: In our case, we consider two extreme strategies, reflected in the adaptive and the sophisticated probabilistic choice rule. Undoubtedly, one could come up with numerous choice rules, especially for a naive agent type. We believe however that the adaptive choice, as introduced above, stands out due to its extreme simplicity. The execution of the adaptive choice rule demands virtually no computational skills from a naive agent, since he does nothing but to adopt the average choice as his own. We therefore argue that our behavioral assumption is a simple but appropriate representation of a naive agent's behavior. Finding a starting value $\tilde{\rho}_0$ however, goes beyond tackling a mere technical difficulty; it calls for a discussion on the nature of heterogeneity itself. At time $t = 0$, when the payoff shock becomes effective, individual price setting decisions are allocated at rational equilibrium level. In spite of this displayed uniformity, it seems implausible to infer that agents are sophisticated to the same degree. Once agents are confronted with a new choice situation, e.g. due to a change of state, heterogeneity will become prevalent and will affect future aggregate behavior. Sophisticated agents will recognize how to choose optimally in the new situation, while less sophisticated will not. In terms of our model, we would therefore like to infer the distribution of AA and SA that emerge shortly after the payoff shock, that is, we would like to determine $\tilde{\rho}_0 \equiv \rho_{0+\epsilon}$ which allows us to compute $E_0\left[\pi_i(\beta;\alpha,\tilde{\rho}_0)\right]$.

As derived in appendix B, the dynamics of ρ_t is governed by the following relation

$$E[\rho_{t+1}|\beta,\alpha,\rho_t] = \rho_t + \rho_t(1-\rho_t)\beta + (1-\rho_t)^2\alpha. \tag{3.10}$$

We naturally identify the expected value at $t = 1$ by setting $E[\rho_1|\beta,\alpha,\tilde{\rho}_0] = \rho_1^{data}$. Along with $\alpha = 0$, our model requires $\tilde{\rho}_0$ and $\hat{\beta}$ to jointly satisfy the following equations

$$\rho_1^{data} = \tilde{\rho}_0^2 + (\tilde{\rho}_0 - \tilde{\rho}_0^2)(\hat{\beta}+1) \tag{3.11}$$

$$\hat{\beta} = \arg\max_{\beta} E_0[\pi_{SA}(\beta;0,\tilde{\rho}_0)], \tag{3.12}$$

where the first equation (3.11), as already stated, is a special case of (3.10) and the second equation (3.12) stands for the necessity of β to satisfy the optimality condition (3.9). The final parameters which solve the optimization problem are easily found by a fixed point iteration departing from starting values ρ_1^{data} and $\beta_1^{stval} = \arg\max_{\beta} E_1[\pi_{SA}(\beta;0,\rho_1^{data}])$. The starting values as well as the final values are summarized in Table 1. Figure 2 visualizes the relation between $\beta \in [0,1]$ and

$(\alpha = 0)$	ct	st
β_1^{stval}	0.450	1.000
ρ_1^{data}	0.472	0.950
$\tilde{\rho}_0$	0.383	0.776
$\hat{\beta}$	0.378	1.000

Table 1: Initial Values and Optimal Values

the expected 'real' payoffs of SA in both strategic situations. An explanation for the curves that result can be formulated as follows: A small value of β (e.g. $\beta = 0.1$) will lead to slow convergence whenever some AA participate and to costs due to coordination failure. On the other hand, SA are free to enforce faster convergence by setting e.g. $\beta = 1$. Under strategic complementarity (ct), this strategy entails large initial costs originated by large temporary deviations from the optimal response in early periods. Clearly, under ct, the payoff maximizing β must be an intermediate value in $[0, 1]$ that accounts for this trade-off between inefficient off-equilibrium time and deviation costs. As illustrated in appendix C.2, optimality of $\hat{\beta}$ is reflected by the fact that SA's expected price choice will move almost perfectly along the optimal reaction function during the entire adjustment phase. In other words, SA is able to anticipate the entire expected adjustment path of the average price depending on $\beta, \alpha, \tilde{\rho}_0$ and to choose β such that expected costs are minimized during the adjustment phase.

Under strategic substitutability (st), things are simple: SA's optimal response is to set $\beta = 1$ no matter how many AA are present in the game.[4] We now turn to the results of our simulations.

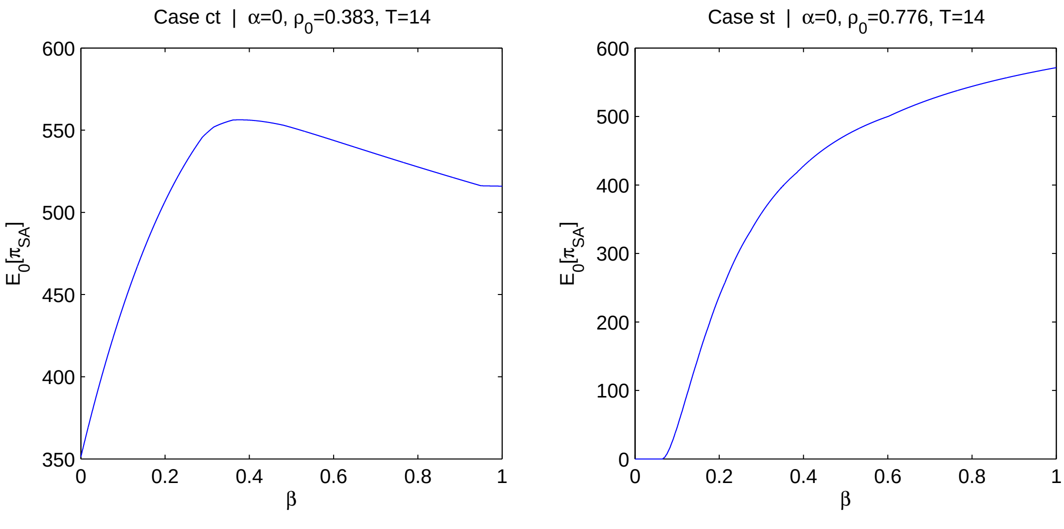

Figure 2: Sum of Expected Real Payoffs for $\beta \in [0, 1]$

3.2 Simulation Results

The simulation of $\{P\}_t$, $\{\bar{P}\}_t$ and $\{Q\}_t$ is based on the price setting decisions of $N = 100$ agents and robust with respect to the number of agents N. Even though the obtained series are more volatile for smaller number of agents the general pattern is always preserved. Figure 3 contrasts FT's empirical findings on the left panel with our simulation results in the right panel. Both, the empirical price series and the simulated realization of the price series, $\{P\}_t$ (solid line), are

[4] Actually, it would be even optimal to overshoot slightly by setting a price below P_1^*, an action not supported in our model.

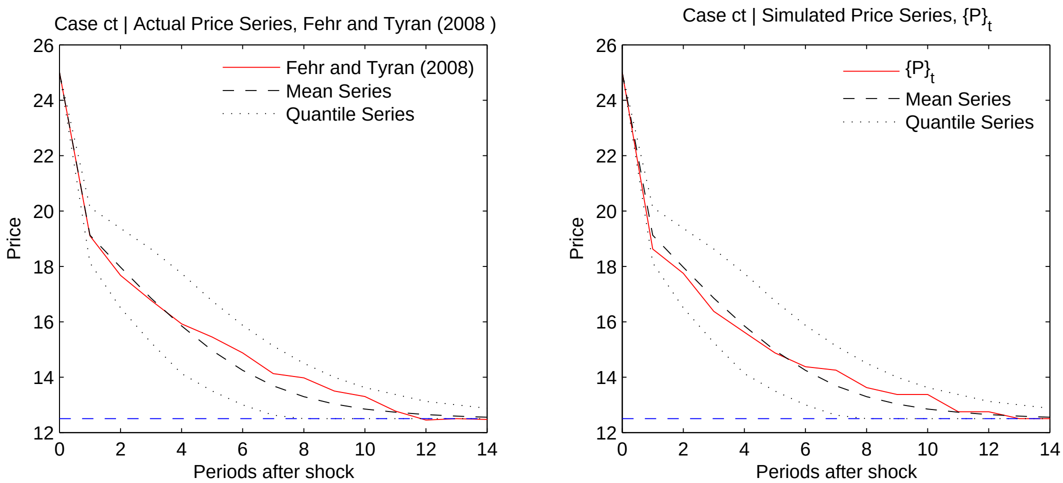

Figure 3: Simulation on $\{P\}_t$, $\{\bar{P}\}_t$, $\{Q\}_t$

shown embedded within the average price series $\{\bar{P}\}_t$ (dashed line) and the quantile series $\{Q(q)\}_t$ (dotted lines). The time window covered in the figures starts at the time when the shock occurs, i.e $t = 0$ and extends for $T = 14$ periods after the shock. Our specification manages to reproduce the behavior of the average price series found by FT in both strategic situations. Under strategic complementarity (ct), it takes 14 periods for the mean price series to adjust to the new equilibrium price $P_1^* = 12.5$ whereas the single realization of $\{P\}_t$ adjusts after 13 periods. Interestingly, the 0.05-quantile intercepts the equilibrium price level line at period $t = 8$, in agreement with rejection of the null hypothesis, $(H_0 : P_t = 0)$, during the first 8 periods based on Fehr and Tyran's data.

The strategic substitutability (st) case appears as a trivial outcome. On one hand, SA choose the new equilibrium price with certainty already in the first period. On the other hand, the group of AA immediately after the shock, i.e. $1 - \tilde{\rho}_0$, is small. Thus, there is nothing else to expect than immediate adjustment of the average price to the new equilibrium prize as illustrated in appendix C.3.

Our model manages to reproduce several features of the empirical price series during the adjustment phase as found by FT. First, we obtain the price series as the realization of a stochastic process with variance and mean close to those in the data. Second, convergence of a price series realization is reached with 90% probability within periods $t \in [8, 15]$, which, in view of FT's results, seems to be a good approximation. Third, our model generates an exponentially falling average price series that matches the price series found by FT.

4 Further Implications

Recalling that ρ_t gives the fraction of agents playing the sophisticated strategy leads to the natural interpretation of ρ_t as the *degree of rationality* of the (stylized) economy at time t. With this concept in mind, we can proceed to examine and interpret our findings with regard to $\tilde{\rho}_0$. As exposed earlier in this paper, agents' allocations are identical at $t = 0$, displaying apparent uniformity. If we put ourselves in SA's position at time $t = 0$, we barely recognize any valuable information from which to infer how many agents will play adaptive. Consequently, $\tilde{\rho}_0$ remains essentially an unobservable variable in our model[5]. On the other hand, FT's experimental results clearly report agents to form different expectations on the other agent's actions depending on the strategic situation. In an additional control experiment, FT observe a much faster price adjustment in ct when agents are confronted with completely rational computerized agents (of course,

[5]Remember that we estimated $\tilde{\rho}_0$ ex-post using the first price observation.

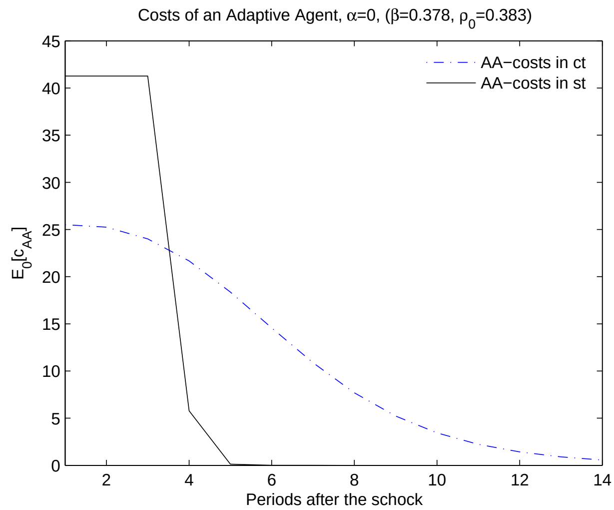

Figure 4: Costs of AA in ct and st

players were truthfully informed about the nature of their opponents). Noteworthy, strong (but not perfect) adjustment takes place already in the first period after the shock.

The cited evidence suggests that believes on the other agents' ability indeed affects aggregate outcome. In line with the concept sketched above, $\tilde{\rho}_0$ can interpreted as the degree of rationality within the economy that emerges just an instant after the shock, or alternatively, as the (correct) believes of agents playing the sophistacated strategy on the fraction of agents that will play the sophisticated strategy right after the shock. From a purely technical discussion, we obtained $\tilde{\rho}_0 = 0.38$ in ct and $\tilde{\rho}_0 = 0.77$ in strategic substitutability (st), meaning that the rationality induced by st is larger than by strategic complementarity (ct). Although this findings seem to confirm our views as well as those expressed in FT, our model does not give any deeper explanation for this result. In a first step, one could think of two simple cases, to explain irrational or suboptimal behavior in strategic choice situations namely, people are just not able recognize the optimal choice or, alternatively, people are perfectly aware of optimal choice but they (correctly or mistakenly) believe that other people are not. We see two main factors that may affect individual rationality as as well as beliefs on the other people's rationality. First, we think that cognitive accessibility to optimal choice plays a role. The harder a problem, i.e. the less accessible, the more agents will struggle to find the optimal choice. In an exemplary difficult situation, not a single person will be able to infer the optimal choice. As the problem becomes less difficult, some people will recognize the optimal choice, but correctly believe that other people will not. In the case where optimal choice is obvious, agents may mistakenly believe that some other agents may opt to deviate from optimal behavior and thus, irrational aggregate behavior will turn into a matter of coordination failure. Second, monetary incentives may be a valid driver of both individual rationality and beliefs on other agents' rationality. Agents will typically, but not necessary, put more effort in finding optimal choice or action when payoffs or costs at stake are higher. Moreover, monetary incentives could help to correct mistaken beliefs that lead to the coordination failure problem described above by punishing deviations from optimal behavior.

While examining the influence of cognitive accessibility (and other psychological factors) on individual rationality and beliefs clearly lies beyond the scope of this work[6], our model allows us to calculate the costs an adaptive agent has to bear in both strategic situations. Motivated by the previous discussion, we assume that $\tilde{\rho}_0$ mainly depends on individual rationality and beliefs on other people's rationality, both driven by psychological factors and the costs from being AA. To compare the costs of being AA on equal grounds, we assume that both strategic situations

[6]FT for example, argue that the error of behaving adaptive is much more salient under substitutability, which may induce people to think more carefully about their price setting decisions.

are perceived equally, i.e. we calculate the costs of being AA departing from the same $\tilde{\rho}_0$ in both strategic situations. Figure 4 illustrates the period by period costs of being AA in both strategic situations. Total costs in ct are 25% higher than in st, which clearly contradicts the hypothesis that total costs could come into question as a driver of rational behaviour. On the other hand, a closer look at the cost structure leads to interesting considerations. Under st, AA incurs a very large part of the costs (95%) during the first three periods. It seems plausible to assume that readily identifiable costs would influence the behavior of less sophisticated agens. Or putting it differently, it seems implausible to assume that payoffs/costs far off the period $t = 0$ could have a major influence on the behaviour of less sophisticated agents as the estimation of these payoffs/cost would require more sophisticated reasoning. Summarizing, if payoffs/costs play any role as a driver of individual rationality and beliefs, then, as we find, the influence comes through the payoff/cost structure rather than the overall payoff/costs that adaptive strategies face.

5 Conclusions

Fehr and Tyran (2008) demonstrated experimentally that in a certain strategic situation, known as strategic complementarity, suboptimal aggregate behavior can occur in terms of sluggish price adjustment after a fully anticipated money shock. In this paper, we propose a probabilistic cellular automata model that is capable to reproduce this interesting observation. In addition, and by means of our model, we can provide a further analysis of the economic problem apparent in the experiment of Fehr and Tyran.

As laid out in detail in the paper, our model says that suboptimal aggregate behavior, which expresses itself through sluggish adjustment to rational equilibrium, arises from the interaction of agents playing adaptive strategies with agents playing sophisticated strategies. We find that, under strategic complementarity, the presence of agents playing adaptive strategies induces agents with sophisticated strategies to (optimally) deviate from rational equilibrium in order to minimize their costs during the inefficient transition phase. Our analysis suggests that the off-equilibrium payoff structure of agents playing adaptive strategies is key for understanding the suboptimal aggregate behavior in a repeated game experiment as the one set up by Fehr and Tyran (2008). Under substitutability, agents playing the adaptive strategy incur 95% of total costs within the first three periods of the transition phase. Under complementarity, total costs are higher but evenly distributed across periods. Ceteris paribus, high immediate costs punishing suboptimal adaptive strategies seem to exert a greater disciplining effect on aggregate behavior than high total costs.

The almost perfect match between the predictions of our model and the experimental findings of Fehr and Tyran (2008) suggests that probabilistic cellular automata models provide a valuable tool to model the off-equilibrium dynamics manifest in certain economic problems. Generally, our results indicate that the off-equilibrium payoffs of adaptive strategies in repeated games experiments that display non-negligible off-equilibrium effects should be evaluated more systematically. The resulting time-dependent off-equlibrium payoff profiles may then help to identify and categorize the payoff structures which can cause coordination failures in real experiments and real economies.

References

Camerer, C. F. and Ho, T.-H. (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119(3):pp. 861–898.

Cooper, R. and Haltiwanger, J. (1996). Evidence on macroeconomic complementarities. *The Review of Economics and Statistics*, 78(1):pp. 78–93.

Cooper, R. and John, A. (1988). Coordinating coordination failures in keynesian models. *The Quarterly Journal of Economics*, 103(3):441–463.

Cooper, R. and Johri, A. (1997). Dynamic complementarities: A quantitative analysis. *Journal of Monetary Economics*, 40(1):pp. 97–119.

Davis, D. (2009). Do strategic substitutes make better markets? a comparison of bertrand and cournot markets. *Working Paper, Virginia Commonwealth University Richmond.*

Davis, D. and Korenok, O. (2010). Nominal price shocks in monopolistically competitive markets: An experimental analysis. *Working Paper, Virginia Commonwealth University Richmond.*

Droste, E., Kosfeld, M., and Voorneveld, M. (2003). Best-reply matching in games. *Mathematical Social Sciences*, 46:pp. 291–309.

Fehr, E. and Tyran, J.-R. (2001). Does money illusion matter? *The American Economic Review*, 91(5).

Fehr, E. and Tyran, J.-R. (2008). Limited rationality and strategic interaction: The impact of the strategic interaction on nominal inertia. *Econometrica*, 76(2):pp. 353 – 394.

Haltinwanger, J. and Waldman, M. (1985). Rational expectations and the limits of rationality: An analysis of heterogeneity. *The American Economic Review*, 75(3):pp. 326–340.

Haltinwanger, J. and Waldman, M. (1989). Limited rationality and strategic complements: The implications for macroeconomics. *The American Economic Review*, 104(3):pp. 463–483.

Kari, J. (2005). Theory of cellular automata: A survey. *Theoretical Computer Science*, (334):pp. 3–33.

Kosfeld, M., Droste, E., and Voorneveld, M. (2002). A myopic adjustment process leading to best-reply matching. *Games and Economic Behavior*, 40:pp. 270–298.

Milgrom, P. and Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):pp. 1255–1277.

Qiu, G., Kandhai, D., and Slot, P. (2007). Understanding the complex dynamics of stock markets through cellular automata. *Physical Review.*

Rivas, J. (2008). Learning within a markovian enviroment. *EUI Woking Papers*, (13).

Van Kampen, N. (1981). *Stochastic Processes in Physics and Chemistry.* North Holland Publishing Company.

Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, 19:305–321.

Wolfram, S. (2002). *A New Kind of Science.* B&T.

A Description of Average Price and Quantile Series

We generate $N = 1000$ price series $\{P\}_T^j = \{P_1^j, P_2^j, \ldots, P_t^j\}$, with $j = 1, 2, \ldots, N$ and $t = 1, 2, \ldots, T$ and estimate the elements $\bar{P}_t$ of the average price series $\{\bar{P}\}_T = \{\bar{P}_1, \bar{P}_2, \ldots, \bar{P}_T\}$ as

$$\bar{P}_t = \sum_{j=1}^{N} P_t^j.$$

The elements $Q(q)_t$ of the q-quantile series $\{Q(q)\}_T = \{Q_1, Q_2, \ldots, Q_T\}$ are defined as

$$Q(q)_t = \sup_l \{\tilde{P}\}_t^l \quad \text{s.t.} \quad \frac{l}{n} \leq q, \quad q = \frac{k}{N}, \quad k \in \{1, 2, \ldots, N\}, \forall t$$

where $\{\tilde{P}\}_t^l$ denotes the l-th element of the vector $\{\tilde{P}\}_t$ which contains all $\left[\{P\}_t^1, \{P\}_t^2, \ldots, \{P\}_t^N\right]$ in ascending order.

B Proofs

Proof 1 *From the definition of the agent's choice rules we know that*

$$\begin{aligned}
p_{SA}^t &= P[x_i^{t+1} = 1|\beta, \alpha, x_i^t = 1] = \rho_t + (1-\rho_t)\beta \\
1 - p_{SA}^t &= P[x_i^{t+1} = 0|\beta, \alpha, x_i^t = 1] \\
p_{AA}^t &= P[x_i^{t+1} = 1|\beta, \alpha, x_i^t = 0] = \rho_t + (1-\rho_t)\alpha \\
1 - p_{AA}^t &= P[x_i^{t+1} = 0|\beta, \alpha, x_i^t = 0].
\end{aligned}$$

The expected value of ρ_t given $\{\beta, \alpha, x^{t-1}\}$ is calculated as

$$\begin{aligned}
E[\rho_t|\alpha, \beta, x^{t-1}] &= \frac{1}{N}\sum_{i=1}^{N} E[x_i^t|\beta, \alpha, x^{t-1}] \\
&= \frac{1}{N}\sum_{i=1}^{N} (x_i^t = 0)(1 - p_{SA}^{t-1})x_i^{t-1} + (x_i^t = 1)p_{SA}^{t-1}x_i^{t-1} \\
&+ (x_i^t = 0)(1 - p_{AA}^{t-1})(1 - x_i^{t-1}) + (x_i^t = 1)p_{AA}^{t-1}(1 - x_i^t) \\
&= \frac{1}{N}\sum_{i=1}^{N} p_{SA}^{t-1}x_i^{t-1} + p_{AA}^{t-1}(1 - x_i^{t-1}) \\
&= p_{SA}^{t-1}\rho_{t-1} + p_{AA}^{t-1}(1 - \rho_{t-1}) \\
&= (\rho_{t-1} + (1-\rho_{t-1})\beta)\rho_{t-1} + (\rho_{t-1} + (1-\rho_{t-1})\alpha)(1 - \rho_{t-1}) \\
&= \rho_{t-1} + \rho_{t-1}(1-\rho_{t-1})\beta + (1-\rho_{t-1})^2\alpha.
\end{aligned}$$

Let us denote $E_{t-1}[\rho_t] = E[\rho_t|\alpha, \beta, \rho_{t-1}]$, then the result is written as

$$\begin{aligned}
E_{t-1}[\rho_t] &= \rho_{t-1} + \rho_{t-1}(1-\rho_{t-1})\beta + (1-\rho_{t-1})^2\alpha && \text{(B.1)} \\
&= f(\beta, \alpha, \rho_{t-1}). && \text{(B.2)}
\end{aligned}$$

Let us now calculate $E_{t-1}[\rho_{t+1}]$ and note that the transitions probabilities p_{SA}^t and p_{AA}^t are only known in expectations, i.e. $E_{t-1}[p_{SA}^t]$ and $E_{t-1}[p_{AA}^t]$. Because the process is Markovian, the Chapman - Kolmogorov equation holds Van Kampen (1981), that is

$$E_t[\rho_{t+1}] = f(\beta, \alpha, E_{t-1}[\rho_t]) \qquad \text{(B.3)}$$

holds.

Let us assume ρ_k *to be known. By denoting* $f^l = \underbrace{f \circ f \circ \cdots \circ f}_{l}$ *and keeping* (B.3) *in mind we can then build up successively expected values* ρ_t *for an arbitrary time* t

$$\begin{aligned} E_k[\rho_{k+1}] &= f(\beta, \alpha, \rho_k) \\ E_k[\rho_{k+2}] &= f(\beta, \alpha, E_k[\rho_{k+1}]) \\ &= f(\beta, \alpha, f(\beta, \alpha, \rho_k)) \\ &\vdots \\ E_k[\rho_t] &= f^{t-k}(\beta, \alpha, \rho_k) \end{aligned}$$

Our proposition for $E_k[P_{SA}^{t+1}]$, $E_k[P_{AA}^{t+1}]$, $E_k[P_{t+1}]$ *follows straight forward from the fact that these price variables are linear combinations of* ρ_t.

Proof 2 *From the definition of the expected price, we know* $E_t[P_{AA}^{t+1}] = p_{AA}^t P_1^* + (1 - p_{AA}^t) P_0^*$ *and the LHS must be equal to* P_t*, i.e.*

$$\begin{aligned} P_0^* + (P_1^* - P_0^*) p_{AA}^t &= P_0^* + (P_1^* - P_0^*) \rho_t \\ \Leftrightarrow p_{AA}^t &= \rho_t \\ \Leftrightarrow \rho_t + (1 - \rho_t)\alpha &= \rho_t, \end{aligned}$$

which is solved by $\alpha = 0$ *as long* $\rho_t < 1$.

C Additional Graphs

C.1 Schematic Illustration of Payoff-Function

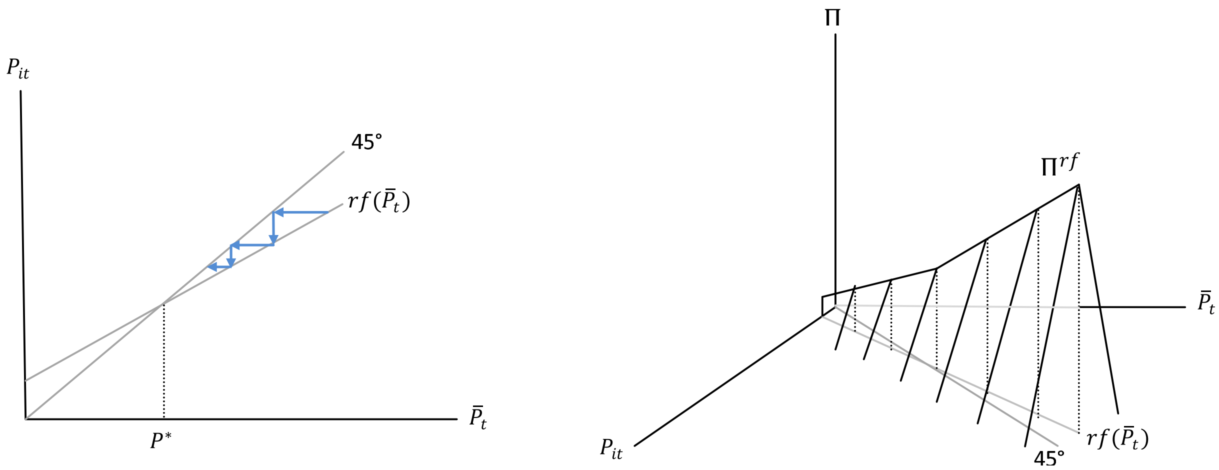

Figure 5: Schematic Illustration of Payoff-Function

C.2 Own vs. FT's Payoff Function

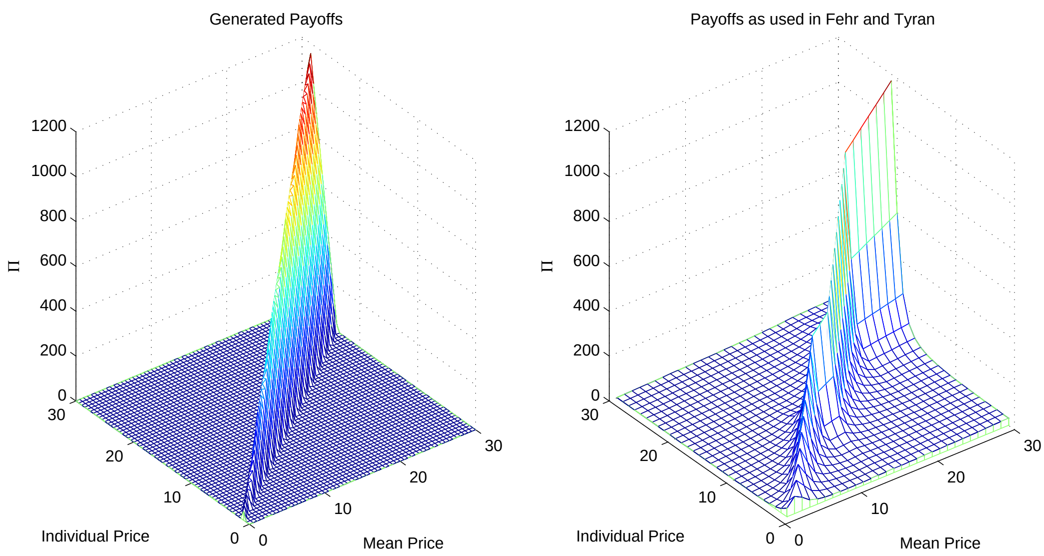

Figure 6: Own payoff function vs. payoff function in FT

$(\alpha = 0)$	ct	st
A	39	39
a	2.50	22.5
b	10/12.5	-(10/12.5)
$\gamma_1 = \gamma_2$	30	30
Ψ	200	200

Table 2: Parameter Values of Payoff Function

C.3 Simulation Results for st

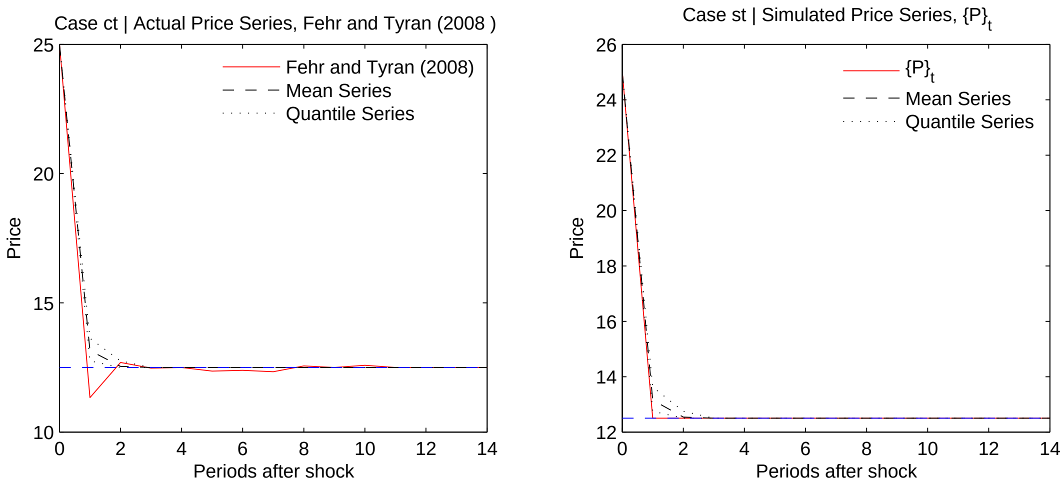

Figure 7: Simulation on $\{P\}_t$, $\{\bar{P}\}_t$, $\{Q\}_t$

C.4 Various Adjustment Paths and their Corresponding Payoffs

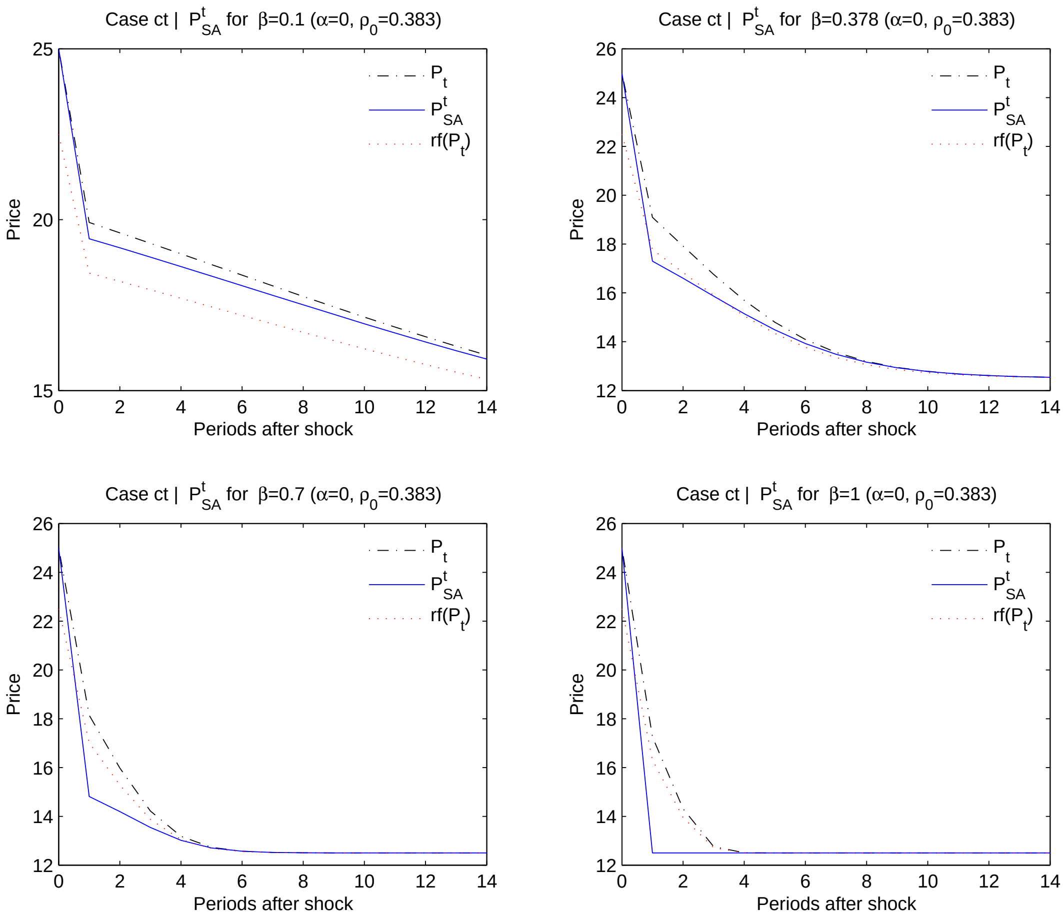

Figure 8: Adjustment Paths for Various Values of β

Note: To shorten notation we set $P_t \equiv E_0[P_t]$, $P^t_{SA} \equiv E_0[P^t_{SA}]$ and $rf(P_t) \equiv rf(E_0[P_t])$.

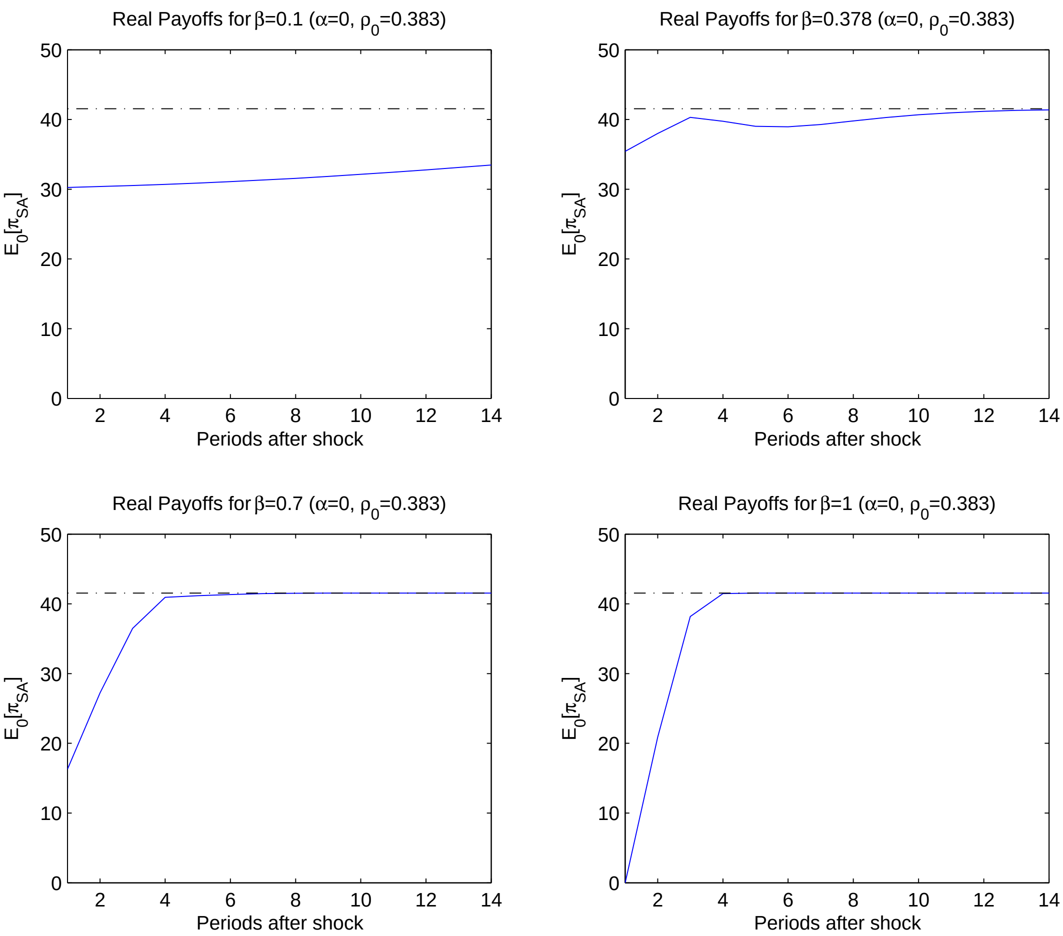

Figure 9: Corresponding Period by Period Payoffs