Individual evolutionary learning in a New Keynesian model

Olena Kostyshyna *

November 12, 2011

Abstract

In this paper I study the learnability of a rational expectations equilibrium solution under individual evolutionary learning in a New Keynesian model (Woodford (2003)). Woodford (2003) and Bullard and Mitra (2002) show that this model is determinate and expectationally stable (E-stable) when the Taylor principle is satisfied. Arifovic et al. (2008) show that the rational expectations equilibrium minimum state variable (REE MSV) solution is learnable in the case of social learning whether or not the Taylor principle is satisfied. In this paper, I apply individual evolutionary learning (IEL) in the same environment. Both social and individual learning allow for heterogeneity of expectations in contrast to recursive homogeneous learning in Bullard and Mitra (2002).

I find that agents using IEL are not able to learn the REE MSV solution whether or not the Taylor principle holds. The intuitive explanation is that the social aspect of learning is essential to agents' coordination on the REE MSV solution. Social interaction and ability to learn from the actual decisions of others allows agents to imitate the best ideas of the population. I perform a sensitivity analysis by varying the implementation of IEL.

1 Introduction

In this paper I study the learnability of a rational expectations equilibrium solution under individual evolutionary learning in a New Keynesian model (Woodford (2003)). Woodford (2003) and Bullard and Mitra (2002) show that this model is determinate and expectationally stable (E-stable) when the Taylor principle is satisfied. Arifovic et al. (2008) show that the rational expectations equilibrium minimum state variable (MSV) solution is learnable in case of social learning whether or not the Taylor principle is satisfied. In this paper, I apply individual evolutionary learning (IEL) in the same environment. Both social and individual learning allow for heterogeneity of expectations, whereas Bullard and Mitra (2002) use recursive homogeneous learning. The main objective of this paper is to study whether the REE MSV solution is learnable under individual evolutionary learning and to compare the outcome to the results with social learning.

The difference between social and individual learning is the difference between 'learning from others' and 'learning by doing'. Examples of research studying the implications of social and individual learning include Arifovic (1994), Chen and Yeh (1994), Dawid and Kopel (1998), Vriend (2000), Arifovic and Maschek (2006) and Arifovic and Karaivanov (2010). Arifovic (1994) studies a cobweb model with social and individual evolutionary learning, and finds that both kinds of learning lead to a Walrasian competitive equilibrium. Vriend (2000) shows that simulations with social learning find a Walrasian equilibrium, whereas simulations with individual learning go to a Cournot-Nash

*Portland State University

equilibrium. Vriend (2000) attributes this difference to the presence of the spite effect in the case of social learning. Arifovic and Maschek (2006) further investigate the details and characteristics underlying the difference in outcomes between social and individual learning in Vriend (2000) and point out that the combination of cost specification with zero cost and updating fitness only with the rules that were used mitigates the spite effect and makes it more difficult for individual learning to find the Walrasian equilibrium.

Arifovic and Karaivanov (2010) use social and individual learning in a principal-agent model, and find that the optimal contract is learnable under social learning, but not under individual learning. The key factor is the evaluation of the hypothetical payoffs for individual learning which is not suitable for the Stackelberg game in the principal-agent model, but is appropriate for playing Nash.

1.1 Main findings

Agents using IEL are not able to learn the REE MSV solution whether or not the Taylor principle holds. This result is different from the result obtained under social learning in Arifovic et al (2008): agents using social learning are able to coordinate on the REE MSV solution whether or not the Taylor principle holds. The intuitive explanation of this difference is that the social aspect of learning is essential to agents' coordination on the REE MSV solution. Social interaction and the ability to learn from the actual decisions of others allows agents to imitate the best ideas from the population.

In IEL, agents choose their decisions from their own pool of strategies based on hypothetical payoffs. In social learning, agents are exposed to the decisions of other agents and compare them based on actual payoffs. The difference in the computation of actual and hypothetical payoffs is important to explain the different results obtained under social and individual learning. When agents learn by doing (individual learning), they hypothesize about how each of their strategies would have performed, if it had been chosen, taking the decisions of the other agents as given. When agents learn from others (social learning), they learn about the other agents' actual decisions and actual payoffs, and so they are exposed to the information that contributed to the determination of actual outcomes. In social learning, agents can imitate the better performing solutions of others. So even if some agents are away from the equilibrium at any given point of time, they have the chance to meet another agent whose strategy is closer to the equilibrium, bringing a higher payoff, and they can adopt this better performing rule. This prevents the economy from wandering away from equilibrium.

2 Environment

I study the simple version of the New Keynesian model described in Woodford (2003). The model is derived from the optimizing behavior of households and firms. Households maximize utility by choosing consumption, labor and money holdings. Households consume a composite good, but produce only one good. Composite goods are produced by a continuum of monopolistically competitive firms that determine their price subject to a Calvo sticky prices friction. In the log-linear model, the output gap and deviation of inflation from target are determined by the following equations:

$$z_t = z^e_{t+1} - \sigma^{-1}[r_t - \pi^e_{t+1}] + \sigma^{-1} r^n_t \quad (1)$$

$$\pi_t = \kappa z_t + \beta \pi^e_{t+1} \quad (2)$$

where z_t is the output gap, π_t is the deviation of the inflation rate from a prespecified target, r_t is the deviation of the short-term nominal interest rate from the value that would hold in a steady state with the level of inflation at target and output at the level consistent with fully flexible prices. z^e_{t+1} and π^e_{t+1} denote the subjective expectations of agents formed in period t, which can differ from a rational expectation. The parameter $\beta \in (0,1)$ is the discount factor of the representative household, $\sigma > 0$ is the intertemporal elasticity of substitution of the household, and $\kappa > 0$ is the degree of price stickiness in the economy. The natural rate of interest, r^n_t, is a stochastic term which follows the AR(1) process

$$r^n_t = \rho r^n_{t-1} + \epsilon_t, \tag{3}$$

where ϵ_t is *i.i.d.* noise with variance σ^2_ϵ, and $0 \leq \rho < 1$ is a serial correlation parameter. The monetary authority sets the interest rate according to the following feedback rule:

$$r_t = \varphi_\pi \pi_t + \varphi_z z_t, \tag{4}$$

where φ_π and φ_z are policy parameters taken to be strictly positive. Equations (1), (2) and (4) determine the equilibrium values of z_t, π_t and r_t each period as functions of the exogenous shock r^n_t and current subjective expectations. The model above is given in log-linear form, and the steady state is $(z_t, \pi_t, r_t) = (0,0,0)$.

Equations (1), (2), and (4) can be written as:

$$y_t = \alpha + B y^e_{t+1} + \chi r^n_t \tag{5}$$

where $\alpha = 0$, $y_t = [z_t, \pi_t]'$,

$$B = \frac{1}{\sigma + \varphi_z + \kappa\varphi_\pi} \begin{bmatrix} \sigma & 1 - \beta\varphi_\pi \\ \kappa\sigma & \kappa + \beta(\sigma + \varphi_z) \end{bmatrix}, \tag{6}$$

and

$$\chi = \frac{1}{\sigma + \varphi_z + \kappa\varphi_\pi} \begin{bmatrix} 1 \\ \kappa \end{bmatrix}. \tag{7}$$

Bullard and Mitra (2002) show that the minimal state variable (MSV) solution is

$$y_t = \bar{a} + \bar{c} r^n_t \tag{8}$$

where $\bar{a} = 0$ and $\bar{c} = [I - \rho B]^{-1}\chi$.

Bullard and Mitra (2002) show that the necessary and sufficient condition for a rational expectations equilibrium to be determinate in the sense of Blanchard and Kahn (1980) is the same as the necessary and sufficient condition for the expectational stability of the rational expectations equilibrium:

$$\kappa(\varphi_\pi - 1) + (1 - \beta)\varphi_z > 0 \tag{9}$$

Condition (9) states the Taylor Principle: the nominal interest rate should respond more than one-for-one to a deviation of inflation from the target. Evans and Honkapohja (2001) show that expectational stability governs the stability of the real time system under recursive algorithms such as least squares.

3 Individual evolutionary learning

3.1 Overview

I study the behavior of this system where recursive least squares learning is replaced by individual evolutionary learning. Agents initial beliefs about the forecasting model are heterogeneous, and they update their beliefs using individual evolutionary learning. The question is whether agents are able to learn the equilibrium MSV solution.

Individual evolutionary learning is implemented using a genetic algorithm that is a numerical optimization technique first introduced by Holland (1975) and described in Goldberg (1987), Michalewicz (1996) and Back et. al. (2000). Among the many advantages of using a genetic algorithm for optimization is that it does not rely on the starting point and it is able to explore the entire search space through random choices. This aspect of a genetic algorithm allows it to be better at finding globally optimal solutions. The genetic algorithm works especially well in case of discontinuous, nondifferentiable, noisy, multimodal and other unconventional surfaces (Schwefel 2000). The genetic algorithm uses stochastic transition rules to guide the search toward regions of the search space with expected improved performance. Economic literature has many examples where genetic algorithms are used to model the learning behavior of heterogeneous economic agents. These include using genetic algorithms in cobweb model (Arifovic, 1994, Dawid and Kopel, 1998, Franke, 1998, Arifovic and Maschek, 2006, Vriend, 2000), overlapping generations monetary economies (Arifovic, 1995, Dawid, 1996, Bullard and Duffy, 1998, 2001), models of foreign exchange (Arifovic 1996, Lux and Schornstein 2002), financial markets (LeBaron, 1999, 2006, Lux and Marchesi, 2000), growth (Arifovic et al., 1997), matching (Haruvy et al. 2006), and Walrasian general equilibrium (Gintis, 2007).

3.2 Perceived law of motion

In the implementation of individual evolutionary learning, I assume that agents know the correct equilibrium form of the solution but do not know the equilibrium values of the coefficients and must learn about them. Agents are heterogeneous in their initial beliefs. Under individual evolutionary learning, each agent must learn on its own, and cannot learn from others as in social learning.

The economy is populated by N individuals indexed by j, $j = 1, ..N$. Each individual maintains I strategies indexed by i, $i = 1, ..I$. Each strategy $s^i = \{a1^i, a2^i, c1^i, c2^i\}$ is a set of coefficients in the perceived law of motion. Each agent j maintains a strategy pool S_t at each point of time t, $S_t = \{s_t^1, ..., s_t^I\}$. In each period t, each agent makes a decision about which set of coefficients to use. The decision of agent j in period t is a set of coefficients $(a1_{j,t}, a2_{j,t}, c1_{j,t}, c2_{j,t})$ used to formulate the perceived law of motion (PLM):

$$z_t = a1_{j,t} + c1_{j,t} r_t^n \quad (10)$$

$$\pi_t = a2_{j,t} + c2_{j,t} r_t^n \quad (11)$$

The perceived law of motion (10, 11) takes the correct form of the MSV equilibrium solution (8), however, the values of the coefficients in the PLM can differ from their REE MSV values. Each agent j learns about the vector of coefficients $(a1, a2, c1, c2)$ by updating its pool of rules in each period t. The natural rate of interest r_t^n is a stochastic term, and its presence complicates learning for the agents.

The initial pool of strategies, S_0, is randomly generated around equilibrium MSV values. The value of each coefficient is drawn from a normal distribution with the mean equal to the corresponding

MSV equilibrium value. The equilibrium values of the intercepts $a1$ and $a2$ are zero, and I use a smaller value for the standard deviation of the distribution to initialize and later update them. The standard deviation of coefficients $c1$ and $c2$ is set equal to the larger of the two MSV equilibrium values, and the standard deviation of $a1$ and $a2$ is equal to half of the standard deviation of the coefficients $c1$ and $c2$.[1]The decision in the first period $t = 1$ is random, the subsequent choice is explained in section 3.5.

3.3 Expectations and the actual law of motion

In period t agent j uses its decision $(a1_{j,t}, a2_{j,t}, c1_{j,t}, c2_{j,t})$ to form expectations of the output gap and deviation of inflation from target using (10), (11) and (3):

$$z^e_{j,t+1} = a1_{j,t} + c1_{j,t}\rho r^n_t, \tag{12}$$

$$\pi^e_{j,t+1} = a2_{j,t} + c2_{j,t}\rho r^n_t \tag{13}$$

The average expectations of the output gap and the deviation of inflation from target are computed as

$$z^e_{t+1} = \frac{1}{N}\sum_{j=1}^{N} z^e_{j,t+1}, \tag{14}$$

$$\pi^e_{t+1} = \frac{1}{N}\sum_{j=1}^{N} \pi^e_{j,t+1} \tag{15}$$

The actual values of the output gap and deviation of inflation from target are obtained using (5):

$$y_t = \alpha + B\begin{bmatrix} z^e_{t+1} \\ \pi^e_{t+1} \end{bmatrix} + \chi r^n_t. \tag{16}$$

3.4 Genetic operators

After the actual outcomes are determined based on (16), agents update their pool of rules using genetic operators. The genetic operators are intended to model agents' learning. This includes exploration of different strategies through experimentation and crossover, and an increase in frequency of the strategies with higher payoffs through tournament selection. Agent's choice of strategy is random where probability is based on the performance of the strategies.

I use three genetic operators: crossover, mutation, and tournament selection. I implement a real-valued genetic system.

Crossover is the first genetic operator. It is applied to the pairs of strategies (parents), and the outcome/offspring combines the characteristics of its parents. The advantage of crossover is that it allows parts of the solutions to combine to create new ideas, which in certain cases can lead to finding better strategies quicker than just using mutation. I use simple crossover. Each agent performs the following. In its set of I strategies, all strategies are randomly matched without replacement to form pairs. Each pair of randomly matched strategies is subject to crossover with probability p_{cross}. With probability 0.5, each coefficient is exchanged between two strategies. Some pairs of strategies will

[1] Initialization of the algorithm, setup of genetic operators and values of the model parameters are the same as in Arifovic et al. (2008) for easier comparison of the results of individual and social learning.

not be subjected to crossover, some pairs will not exchange any coefficients, some pairs can exchange some of their coefficients, or in the extreme case, all coefficients can be exchanged, i.e. parent rules just change their indices.

Mutation is implemented after crossover. Mutation models experimentation in learning. Agents can find good solutions by chance, or make mistakes. The mutation operator is crucial in genetic algorithms because it allows exploration of the entire search space to find globally optimal solutions. However, the effect of mutation can be destructive by introducing genetic variation, especially late in the simulation. To control this, I use a standard deviation of mutation that decreases over time as described below in (18).

Mutation is implemented as follows. With probability p_{mut}, each coefficient is changed as:

$$new = current + random * mutdeviation_t, \tag{17}$$

where $random$ is a random number drawn from a standard normal distribution, $current$ is the current value of the coefficient, and $mutdeviation_t$ is the standard deviation used for the mutation in period t. Standard deviation, $mutdeviation_t$, decreases over time as:

$$mutdeviation_t = deviation * (1 - decrease * t/T) \tag{18}$$

where $deviation$ is the standard deviation used to generate an initial set of rules, T is the total number of periods in the simulation, and $decrease$ is a constant. I set $decrease$ equal to 0.95, it is intended to allow for a non-zero mutation standard deviation even in the last period of the simulation.

Once the pool of strategies is refreshed with crossover and mutation, agents evaluate the performance of each strategy as described in section (3.5). The last genetic operator is replication implemented by tournament selection. Through replication, the agent selects the better performing strategies to remain in the pool of strategies. Tournament selection is implemented as follows. The agent randomly chooses 2 strategies from its pool of strategies, with replacement. The agent compares the performance of these strategies and selects the strategy with the higher fitness to put into its new pool. Each agent repeats this process I times. Tournament selection provides selection pressure - higher performing strategies are selected into the new pool, and lower performing strategies are discarded.

3.5 Forecasting performance

After crossover and mutation, the agent computes the performance of each strategy and uses it to evaluate strategies in tournament selection and to choose their decision in the next period. The agents use mean squared forecast error as a performance criterion. They compute the mean squared forecast error for the output gap and the deviation of inflation over all past periods. This environment is stochastic, and, therefore, it is important to use all available data to evaluate the performance, not only data from the last period. Otherwise, agents' performance would depend too much on the last period's shock, and this could prevent agents from learning the true model. [2]

In period t, the agent evaluates the hypothetical payoff of strategy i, $s_t^i = \{a1_t^i, a2_t^i, c1_t^i, c2_t^i\}$, as the mean squared forecast error that would have resulted if the agent had used this particular strategy in current and past periods. In this computation, the agent takes the average forecast of

[2] Branch and Evans (2004, p. 3) assume "... agents to make their decision [based] on unconditional mean payoffs rather than on the most recent period's realized payoff. This is more appropriate in our stochastic environment since otherwise agents would frequently be misled by single period anomalies."

others as given and recognizes that using a different decision in the past would have changed the average forecast and, thus, would have affected the actual values of the output gap, z, and the inflation deviation from target, π. The agent computes the fitness of each strategy as follows. First, agent j computes what its forecasts of z and π would have been in each previous period k if it had used strategy i, $s^i = \{a1_t^i, a2_t^i, c1_t^i, c2_t^i\}$, from the current, date t, pool of rules:

$$z_{i,k}^{e,h} = a1_t^i + c1_t^i \rho r_{k-1}^n \tag{19}$$

$$\pi_{i,k}^{e,h} = a2_t^i + c2_t^i \rho r_{k-1}^n \tag{20}$$

Second, using the hypothetical forecasts $z_{i,k}^{e,h}$ and $\pi_{i,k}^{e,h}$, agent j computes the hypothetical average forecasts $z_k^{e,h}$ and $\pi_k^{e,h}$ which would have resulted if agent j had used strategy i in period k taking the actual average forecasts in period k, z_k^e and π_k^e, as given:

$$z_k^{e,h} = (z_k^e * N - z_{j,k}^e + z_{j,k}^{i,h})/N, \tag{21}$$

$$\pi_k^{e,h} = (\pi_k^e * N - \pi_{i,k}^e + \pi_{j,k}^{i,h})/N \tag{22}$$

Using the hypothetical average forecasts, $z_k^{e,h}$ and $\pi_k^{e,h}$, the agent computes the modified outcomes of z and π that would have resulted if strategy i had been used in period k:

$$\begin{bmatrix} z_k^m \\ \pi_k^m \end{bmatrix} = \alpha + B \begin{bmatrix} z_k^{e,h} \\ \pi_k^{e,h} \end{bmatrix} + \chi r_t^n. \tag{23}$$

In this computation, I implicitly assume that each agent knows how the actual outcomes are determined in the economy. However, each agent does not know how the other agents make their forecasts. And this economy is self-referential - outcomes are determined based on the agents' expectations, and agents update their expectations based on the actual outcomes. This motivates agents' learning, with the goal of learning to make accurate forecasts.

Based on the above computations, the fitness of strategy i is computed as:

$$F_{i,t} = -\frac{1}{t}\sum_{k=1}^{t}(z_k^m - z_{i,k}^{e,h})^2 - \omega\frac{1}{t}\sum_{k=1}^{t}(\pi_k^m - \pi_{i,k}^{e,h})^2 \tag{24}$$

where constant ω gives equal weight to the forecast errors for the output gap and the deviation of inflation from target as in Arifovic et al. (2008). The values of mean forecast squared errors for the output gap and inflation deviation from target differ in orders of magnitude: the forecast squared error of the output gap is approximately 100 time larger than the forecast squared error of inflation deviation. Without weight adjustment, the fitness would not take into account the accuracy of the forecasts for the inflation deviation, resulting in the coefficients for inflation deviation moving away from their MSV values.

The agents make their decisions using hypothetical payoffs. The agent's choice is based on the performance such that the strategies with higher hypothetical payoffs have higher probabilities to be chosen. The probability for strategy i to be used in period $t+1$ is given by:

$$p_{t+1}(i) = \frac{F_{i,t}}{\sum_{i=1}^{I} F_{i,t}} \tag{25}$$

4 Computational experiments

I study the behavior of the model economy by performing computational experiments. Each simulation starts with 100 periods of initial history generated at MSV equilibrium values of coefficients and is 1000 periods long after the initial history. The parameter values are as in Woodford (2003): $\sigma = 0.157$, $\kappa = 0.024$, $\beta = 0.99$, and $\rho = 0.35$. The standard deviation of r^n is $\sigma_\epsilon = 3.72$. The simulations are performed for a range of values of coefficients in the Taylor policy rule. The values of the coefficient for the output gap are in the range $\varphi_z \in [0.2, 1.1]$. And values of the coefficient for inflation deviation are in the range $\varphi_\pi \in [0.5, 2]$. This parameter space includes combinations of the coefficients (φ_z, φ_π) that satisfy the condition for determinacy and E-stability (9) and combinations that do not satisfy it. The value of ω in the fitness criterion is 100.

The parameters of individual evolutionary learning are as follows.The number of agents is $N = 30$. Each agent has $I = 100$ strategies. The probability of mutation is $p_{mut} = 0.1$. A high probability of mutation can introduce too much noise and it has been shown in the literature that a low probability of mutation performs better. Therefore, I use a low probability of mutation set to 0.1. The probability of crossover is $p_{cross} = 0.5$. Related studies have shown that a high probability of crossover performs better.

In this section, I report the results of the computational experiments. I collect and report the following data. In each period and for each individual, I compute the deviation of each coefficient used in that period from its respective MSV equilibrium value, then I average these deviations across all individuals and plot the time series. If agents were able to learn the equilibrium coefficients, the deviations would be zero. Figure 1 presents the results of a typical simulation for a determinate and E-stable region with Taylor rule coefficients $\phi_\pi = 1.5$ and $\phi_z = 0.5$ and shows the average deviations from MSV solution $\pm$ one standard deviation. This figure shows that coefficients $a1$, $a2$, and $c2$ are quite close to their equilibrium values, but that the deviation of coefficient $c1$ from its equilibrium value is large (approximately -1.5). This means that agents are not able to learn the equilibrium values for all of the coefficients. Figure 2 illustrates the results from another simulation in the determinate and E-stable region with Taylor rule coefficients $\phi_\pi = 2.0$ and $\phi_z = 0.2$, in which individuals are not able to learn the MSV solution and the largest deviation from equilibrium is observed for coefficient $c1$ (about -2.5).

Figure 3 illustrates the results of a typical simulation for an indeterminate and E-unstable region with coefficients for the Taylor rule $\phi_\pi = 0.5$ and $\phi_z = 0.3$. This figure shows that the deviation of coefficient $a2$ from its equilibrium (zero) is approximately -1, and the deviation of coefficient $c1$ is slightly below -2. This means that agents cannot learn the equilibrium values. Figure 4 illustrates similar results from another simulation in the indeterminate and E-unstable region with the Taylor rule coefficients $\phi_\pi = 0.5$ and $\phi_z = 0.5$.

To further investigate the performance of this system, I perform 100 simulations for each parameter combination (ϕ_z, ϕ_π) and collect the following data. In each period, each agent uses a set of coefficients to form forecasts. I compute the deviations of these coefficients from their equilibrium values for each agent, and then the average deviations across all agents for each period. To study whether the coefficients go to their equilibrium values, I compute the average deviations of each coefficient during the last 100 periods, 10 periods and last period of the simulation. The collected data on deviations is then averaged over 100 simulations. If agents learn the equilibrium values, these deviations are zero.

Table 1 presents the 100-period average deviations of coefficients from their equilibrium values.[3]

[3]Results for the 10-period and 1-period averages are similar.

This table also shows the deviations of coefficients $c1$ and $c2$ as a percentage of their equilibrium values (percentage deviations are not computed for coefficients $a1$ and $a2$ because their equilibrium values are zero). Table 2 presents the 100-period average and standard deviations of absolute deviations of coefficients from equilibrium values. Table 1 shows that coefficient $c1$ is the furthest from the equilibrium value, while the rest of the coefficients are close to their equilibrium values. Coefficients $c1$ and $c2$ exhibit large percentage deviations of 80-100% from their equilibrium values. Similarly, Table 2 shows the largest deviations for coefficient $c1$, and the absolute deviations of all coefficients are larger than the absolute values of deviations in Table 1. These tables illustrate that agents are not able to learn the equilibrium values of all coefficients.

Data in the figures and tables illustrate that agents are not able to learn the equilibrium MSV solution, whether or not the Taylor principle holds. The inability of individual evolutionary learning to coordinate on the MSV solution differs from the results obtained in this model with social learning in Arifovic et al. (2008). In the case of social learning, agents are able to learn the REE MSV solution, whether or not the Taylor principle holds. The difference between social and individual evolutionary learning is the opportunity for social interaction with social learning. This aspect allows agents to imitate better solutions by observing the actions and payoffs of other agents. Thus even though some agents may be away from the equilibrium at any given point of time, they have the chance to meet another agent whose decision rule is closer to the equilibrium with higher payoff, which they can adopt. This prevents the economy from wandering away from the equilibrium.

It is important to consider the technical aspects underlying the difference between social and individual learning. In individual evolutionary learning, performance is evaluated based on the hypothetical payoff. As explained in section (3.5), the hypothetical payoff evaluates the performance of a strategy as if it had been used in the previous periods, taking the decisions of the other agents as given. This evaluation is similar to the computation of best response. The agent takes into account that using a different strategy would change the actual outcome, and the hypothetical payoff is based on the hypothetical outcome.

In individual learning, each agent learns on its own, without access to the experience of others. Each agent evaluates its ideas, taking the decisions of others as given, as if computing a best response for Nash. Moreover, each agent understands its impact on the actual outcomes, i.e. it understands that if it had used a different strategy, the actual outcome would have been different. Hence, the fitness of each strategy is computed taking into account its hypothetical impact on the 'actual' outcome, i.e. each strategy is evaluated using a *different* modified outcome.

In social learning, fitness is computed using the *same* actual outcome. In each period, agents' decisions determine the actual outcome in the economy. Through imitation, agents learn from others by observing their decisions and payoffs. These decisions are the decisions that determined the actual outcome, and their performance is evaluated based on the same actual outcome. What agents share through imitation is based on the common experience. Furthermore, imitation provides performance pressure that directs agents toward the equilibrium outcome.

In IEL, agents choose from strategies based on the performance evaluated under different hypothetical outcomes. A strategy can appear better-performing than others, but it can also lead agents away from learning the MSV equilibrium. Individual learning results in agents narrowly focusing on their own experience and impact. Thus, the agents can arrive to a suboptimal self-confirming equilibrium (Cho and Kasa (2010)).

It is an interesting question whether agents are aware of their impact on the actual outcome (Cho and Kasa (2010), p.3) and whether they can build counterfactual hypothesis - how the actual outcomes would have been different if the agent had used a different strategy (set of coefficients).

In this section, I used the assumption that agents are able to recognize their impact on the actual outcomes, and that they understand the feedback from their expectations to the actual outcomes. In the next section, I would like to study whether it makes a difference when agents are unaware of or ignore this feedback and impact.

5 Robustness checks

5.1 Ignoring the feedback from expectations to the actual outcomes

In this section, I consider the case of performance evaluation where agents do not take into account that using different strategies changes actual outcomes. This means that when they compute the hypothetical payoff of a strategy, they do not calculate the hypothetical outcome that would have resulted if a given strategy had been chosen in the previous period. (I am going to call this implementation the simplified IEL.) Why might agents ignore the impact of their strategies? First, they could view their own impact on the actual outcome as negligible. Second, they may not know the actual law of motion (16), i.e. how the actual output gap, z, and the deviation of inflation from target, π, are determined. Third, agents cannot observe the expectations of other agents or do not know the values of average expected values z^e_{t+1} and π^e_{t+1}, and, therefore, cannot compute the hypothetical outcomes even if they know the actual law of motion. Finally, it could be all or a combination of these reasons.

In this case, fitness is computed as:

$$F_{i,t} = -\frac{1}{t}\sum_{k=1}^{t}(z_k - z^{e,h}_{i,k})^2 - \omega\frac{1}{t}\sum_{k=1}^{t}(\pi_k - \pi^{e,h}_{i,k})^2 \tag{26}$$

where z_k and π_k are the actual output gap and the deviation of inflation from the target in period k. $z^{e,h}_{i,k}$ and $\pi^{e,h}_{i,k}$ are the hypothetical forecasts that would have been made for period k if agent j had used the coefficients of current strategy i and are computed using (19) and (20). This performance (26) differs from the baseline performance measure described in the previous section because it uses the actual time series z_k and π_k, not their modified/hypothetical z^m_k and π^m_k values as in (24). This simplified implementation of individual learning is closer to the implementation of social learning in Arifovic et al. (2008).

Figures 5 and 6 present the results for typical simulations in the determinate and E-Stable region. We can see that the deviations from the MSV equilibrium values are much smaller than in the baseline case. Most of the coefficients are very close to their MSV equilibrium values, except coefficient $a1$ which deviates by about 0.5 in Figure 5 and by about 1.0 in Figure 6 (the MSV equilibrium value of $a1$ is 0). Coefficient $c1$ goes to its MSV equilibrium value in both figures, whereas it exhibited the highest deviations in the baseline treatment.

Figures 7 and 8 present results for typical simulations in the indeterminate and E-unstable region. Again, the deviations from the MSV equilibrium values are much smaller than in the baseline case, and most of the coefficients are very close to their MSV values. The largest deviations are observed for coefficient $a1$ - approximately 0.5 in both figures.

These figures illustrate that coefficients are much closer to the MSV solution than those in the baseline implementation of IEL (although not as close as in the case of social learning in Arifovic et al. (2008)). In the baseline IEL implementation, the agents know more about the economy than the agents in simplified IEL. They know about the feedback from expectations to actual outcomes, they know the actual law of motion and they can compute hypothetical outcomes for alternative strategies. They are more sophisticated, but they cannot learn the MSV equilibrium. Simplified IEL works better, and the key aspect is the use of the actual outcomes in the performance evaluation in the simplified IEL. All rules are evaluated relative to the same outcome - the actual outcome - which provides an 'anchor' for performance evaluation in contrast to using different modified/hypothetical outcomes in the baseline IEL. This prevents agents' wandering away from the equilibrium.

Simplified IEL performs better than baseline IEL, but it is not as close to the equilibrium as social learning. In addition to using the actual outcomes in performance evaluation, social learning also allows agents to learn from the actual decisions of others through imitation. This system is self-referential: agents' expectations determine the actual outcomes, and agents update their expectations based on the actual outcomes. The ultimate goal of learning is to learn about the data generating process (DGP). Agents' decisions determine outcomes, and so they are part of DGP. In social learning, agents observe the decisions of others, which brings them closer to learning about the DGP than hypothesizing in individual learning. The agents learn from the others' actual decisions which are evaluated based on the same actual outcome that is determined by the observed actual decisions. These characteristics of social learning help agents to coordinate on the equilibrium. In contrast, baseline IEL evaluates the performance of each strategy relative to its own modified/hypothetical outcome. It lacks the benchmark (actual outcome) in social learning. This can lead to agents wandering away from equilibrium.

5.2 Separate performance evaluation

I perform another robustness check with an alternative performance evaluation - agents evaluate the forecasting performance of the output and inflation equations separately. This performance evaluation avoids reliance on a single performance measure (24) that uses weight ω to equalize the importance of the forecast squared errors of output and inflation. Agent j computes the mean squared errors for the output gap and inflation separately as:

$$F^z_{i,t} = -\frac{1}{t}\sum_{k=1}^{t}(z^m_k - z^{e,h}_{i,k})^2 \tag{27}$$

$$F^\pi_{i,t} = -\frac{1}{t}\sum_{k=1}^{t}(\pi^m_k - \pi^{e,h}_{i,k})^2 \tag{28}$$

where $z^{e,h}_{i,k}$ and $\pi^{e,h}_{i,k}$ are computed as in (19) and (20), and z^m_k and π^m_k are computed as in the baseline IEL implementation (23).

This change in performance evaluation also affects the replication and selection of a decision. Replication is implemented via tournament selection and is modified in the following way. The agent randomly chooses 2 strategies from its pool of strategies, with replacement. The agent chooses coefficients for forecasting the output gap from the strategy with the higher payoff for output forecasting, $F^z_{i,t}$, and coefficients for forecasting the deviation of inflation from the target from the strategy with the higher payoff for inflation forecasting, $F^\pi_{i,t}$. The agent repeats this process I times, to obtain a pool of I strategies.

The selection of a new strategy is modified such that the agent chooses strategies to forecast the output gap and inflation separately. The probability of strategy i's coefficients for the output gap to be used in period $t+1$ is given by:

$$p^z_{t+1}(i) = \frac{F^z_{i,t}}{\sum_{i=1}^{I} F^z_{i,t}} \tag{29}$$

And the probability for strategy i's coefficients of the inflation deviation to be used in period $t+1$ is given by:

$$p^\pi_{t+1}(i) = \frac{F^\pi_{i,t}}{\sum_{i=1}^{I} F^\pi_{i,t}} \tag{30}$$

Figures 9, 10, 11 and 12 illustrate the results of typical simulations with separate fitness measures. These figures show the same kind of data as in figures 1, 2, 3 and 4 and present the results of the simulations with the same seeds, the only difference is the use of separate fitness. These figures show that agents are not able to learn the equilibrium coefficients. Figures 9 and 10 show large deviations of coefficient $c1$ from its MSV equilibrium value, approximately -2. Figures 11 and 12 show large deviations of the coefficients $a1$, $a2$ and $c1$ from their MSV equilibrium values. In Figure 11, the deviation of $a1$ is approximately 3, the deviation of $a2$ is above 2, and the deviation of $c1$ is approximately 2. In Figure 12, the deviation of $a1$ is approximately 1, the deviation of $a2$ is above 0.5, and the deviation of $c1$ is approximately 2. Agents are not able to learn the MSV equilibrium, and the coefficients are farther from the equilibrium than in the case of baseline IEL.

Separate performance evaluation allows each set of coefficients ($a1$ and $c1$ for output, $a2$ and $c2$ for inflation deviation) to wander off separately away from the equilibrium. Tournament selection works separately on the output coefficients and inflation coefficients based on the hypothetical payoffs. The hypothetical payoff is evaluated using a hypothetical outcome computed using both the output and inflation coefficients. It is possible that the inflation coefficients perform better than the output coefficients in a given strategy. The tournament selection 'breaks' the strategy of four coefficients by picking well-performing inflation coefficients and well-performing output coefficients to form a new strategy. This new strategy of four coefficients will have a different hypothetical outcome than the strategies from which each set originated from.

Similarly, the selection of a decision works such that the output and inflation coefficients are chosen separately. Each of them can perform well when evaluated as part of their original strategy, but when combined the hypothetical outcome changes and the performance is different. This complicates learning, and the above figures illustrate larger deviations from equilibrium than in the baseline implementation. This occurs because learning based on separate performance is even further from the data generating process: coefficients are chosen separately, whereas outcomes are determined jointly in the actual data generating process.[4]

6 Discussion and conclusion

This paper shows that in the model with individual evolutionary learning, agents are not able to learn the rational expectations equilibrium MSV solution. This result contrasts with the finding of Arifovic et al. (2008) where in the case of social learning, agents are able to learn MSV equilibrium, whether the Taylor principle holds or not.

It is interesting to determine the reason for this difference in the outcomes between individual and social learning. The basic structure of social and individual learning is very similar and relies on the following important elements. First, better performing ideas increase in the population over

[4] I include this robustness check for completeness of investigation.

time. Second, both algorithms explore the entire space of possible values through experimentation and crossover. The difference between social learning and individual learning is that social learning provides an opportunity for interaction with other agents through imitation. Performance pressure in social learning comes from imitation only. The implementation of individual learning includes tournament selection within the pool of rules for each individual agent, not among all agents. Performance pressure in individual learning comes from tournament selection and from the probabilistic choice of decisions based on fitness.

Generally, one of the strengths of evolutionary learning is in the exploration of the entire space of possible values which avoids the dependence of a solution on the initial value. In the dynamic stochastic environment presented here, what matters is not only the exposure to other ideas (exploring all space of values) but the exposure to the actual decisions of others and their actual performance, i.e. to learn from others based on the common experience. Individual hypothesizing lacks the exchange of actual experience. Therefore, what matters in this environment is to learn from "what *was* done" by others, not to conjecture "what *could* have been done" on your own.

Intuitively, the opportunity for social interaction and comparison with others allows agents to broaden their horizons and choose decisions which perform better in an environment where actual outcomes are determined jointly by the expectations of all agents. This environment is interesting because expectations determine actual outcomes, and the agents learn about the best model of the data generating process to make accurate forecasts. However, the actual data is generated through the diverse models of the heterogeneous agents which complicates learning. Therefore, social learning leads to improved learnability because the agents learn from the actual decisions of others that determine the actual outcomes. As the actual decisions of other agents determine the data generating process, and insofar as the agents want to learn about DGP, they are better off learning from the actual decisions of others that contribute to the DGP. Individual learning means that each agent learns on its own and takes the decisions of others as given instead of learning from them. Therefore, individual learning is not sufficient to learn about the equilibrium solution.

References

[1] Alkemade, F., Poutre, H.L., Amman, H.M. 2007. On social learning and robust evolutionary algorithm design in the Cournot oligopoly game.*Computational Intelligence*, 23(2): 162-175

[2] Arifovic, J. 1994. Genetic Algorithm Learning and the Cobweb Model. *Journal of Economic Dynamics and Control*, 18, 3-28.

[3] Arifovic, J. 1995. Genetic Algorithms and Inflationary Economies. *Journal of Monetary Economics*, 36, 219-243.

[4] Arifovic, J.1996. The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economics. *Journal of Political Economy*, 104 :510-541

[5] Arifovic, J. 2000. Evolutionary algorithms in macroeconomic models. *Macroeconomic Dynamics* 4(3): 373-414.

[6] Arifovic, J., Bullard, J., and J.Duffy. 1997. The Transition from Stagnation to Growth: An Adaptive Learning Approach. *Journal of Economic Growth*, 2:185-209

[7] Arifovic, J., and J. Ledyard. 2004. Computer Testbeds: The Dynamics of Groves-Ledyard Mechanisms. Manuscript, April.

[8] Arifovic, J., Maschek, M.K. 2006. Revisiting individual evolutionary learning in the cobweb model - an illustrationg of the virtual spite effect. *Computational Economics*, 28: 333-354

[9] Arifovic, J., Karaivanov, A. 2010. Learning by doing vs. learning from others in a principal-agent model. *Journal of Economic Dynamics and Control*, 43: 1967-1992

[10] Back,T., Fogel, D.B., and Z. Michalewicz. 2000. *Evolutionary Computation 1 Bacis Algorithms and Operators.* Institute of Physics Publishing, Bristol and Philadelphia

[11] Blanchard, O., and C. Kahn. 1980. The solution of linear difference models under rational expectations. *Econometrica* 48(5): 1305-11.

[12] Branch, W., and G. Evans. 2007. Model uncertainty and endogenous volatility. *Review of Economic Dynamics*, 19: 207-237.

[13] Bullard, J. and J. Duffy. 1998. Using Genetic Algorithms to Model the Evolution of Heterogenous Beliefs. *Computational Economics*, 1-20.

[14] Bullard, J., and J. Duffy. 2001. Learning and Excess Volatility, *Macroeconomic Dynamics* 5, 272-302.

[15] Bullard, J., and K. Mitra. 2002. Learning about monetary policy rules. *Journal of Monetary Economics* 49(6): 1105-1129.

[16] Chen, Shu-Heng. 2003. Agent-based computational macroeconomics: a survey. In T. Terano, H. Deguchi, and K.Takama, Meeting the Challenge of Social Problems via Agent-Based Simulation, post-proceedings of the Second International Workshop of Agent-Based Approaches in Economic and Social Complex Systems, New York: Springer, pp. 141-170

[17] Cho, In-Koo, and Ken Kasa. 2010. Learning and model validation: an example.

[18] Cho, In-Koo, and Kenneth Kasa. 2010. Learning and model validation. University of Illinois.

[19] Dawid, H. 1996. Learning of cycles and sunspot equilibria by Genetic Algorithms. *Journal of Evolutionary Economics*, 6(4): 361-373

[20] Dawid, H. 1999. *Adaptive Learning by Genetic Algorithms: Analytical Results and Applications to Economic Models*, ed-n 2, Berlin: Springer

[21] Dawid, H. 2006. Agent-based Models of Innovation and Technological Change. in *Handbook of Computational Economics*, ed. by L.Tesfatsion amd K. L. Judd, Elsevier, 1235-1272

[22] Dawid, H., and M.Kopel. 1998. The Appropriate Design of a Genetic Algorithm in Economic Applications Exemplified by a Model of the Cobweb Type. *Journal of Evolutionary Economics*, 8:297-315

[23] Evans, G., and S. Honkapohja. 2001. *Expectations and Learning in Macroeconomics.* Princeton University Press.

[24] Evans, G.W., Honkapohja, S. 2009. Learning and macroeconomics. *Annual Review of Economics*, 1: 421-449

[25] Franke, R. 1998. Coevolution and Stable Adjustment in the Cobweb Model. *Journal of Evolutionary Economics*,8: 383-406

[26] Gintis, H. 2007. The Dynamics of General Equilibrium. *The Economic Journal*, 117 (October): 1280-1309

[27] Goldberg, D.E. 1989. *Genetic Algorithms in Search, Optimizations, and Machine Learning.* Reading, Mass.: Addison-Wesley Pub.Co.

[28] Haruvy, E., Roth, A.E., and U. M. *Ü*nver The Dynamics of Law Clerk Matching: An Experimental and Computational Investigation of Proposals for Reform of the Market. *Journal of Economic Dynamics and Control*, 30: 457-486

[29] Holland, J.H. 1975. *Adaptation in Natural and Artificial Systems.* University of Michigan Press, Ann Arbor

[30] LeBaron, B., Arthur, W. and Palmer, R. 1999. Time Series Properties of an Artificial Stock Market. *Journal of Economic Dynamics and Control*, 23 : 1487-1516

[31] LeBaron, B. 2000. Agent-based Computational Finance: Suggested Readings and Early Research. *Journal of Economic Dynamics and Control*, 24: 679-702

[32] LeBaron, B. 2006. Agent-based Computational Finance, in *Handbook of Computational Economics*, ed. by L. Tesfatsion and K. Judd, Elsevier, 1187-1232.

[33] Lux, T., and M. Marchesi. 2000. Volatility clustering in financial markets: a micro-simulation of interacting agents. *International Journal of Theoretical and Applied Finance*, 3:675-702

[34] Lux, T. and S. Schornstein. 2002. Genetic learning as an explanation of stylized facts of foreign exchange markets. *Journal of Mathematical Economics*, 41(1-2): 169-196

[35] Michalewicz, Z. 1996. *Genetic Algorithms + Data Structures = Evolution Programs.* 3rd ed. Springer-Verlag

[36] Schwefel, H.P. 2000. Advantages (and disadvantages) of evolutionary computation over other approaches. In *Evolutionary Computation 1 Bacis Algorithms and Operators.* ed. by Back,T., Fogel, D.B., and Z. Michalewicz. Institute of Physics Publishing, Bristol and Philadelphia

[37] L. Tesfatsion and K. L. Judd, eds. 2006. *Handbook of Computational Economics: Volume 2, Agent-Based Computational Economics*, Handbooks in Economics Series, North-Holland, Elsevier, Amsterdam, the Netherlands.

[38] Vriend, N.J. 2000. An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control*, 24: 1-19

[39] Woodford, M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton University Press.

Parameter	φ_π	0.5		1.0		1.5		2.0	
φ_z		mean	std deviation	mean	std deviation	mean	std deviation	mean	std deviation
1.1	a1	0.016	0.334	0.021	0.233	0.020	0.192	0.019	0.173
	a2	0.007	0.503	0.013	0.327	0.013	0.239	0.014	0.189
	c1	-0.828	0.005	-0.817	0.005	-0.804	0.005	-0.793	0.005
	c2	-0.021	0.001	-0.021	0.001	-0.021	0.001	-0.020	0.001
	c1,%	-100.038	0.643	-100.097	0.602	-100.025	0.637	-100.078	0.633
	c2,%	-99.658	2.242	-99.581	2.241	-100.324	2.289	-99.911	2.525
1.0	a1	0.024	0.385	0.024	0.255	0.024	0.210	0.021	0.183
	a2	0.017	0.580	0.011	0.352	0.015	0.253	0.016	0.200
	c1	-0.903	0.005	-0.889	0.005	-0.874	0.006	-0.860	0.006
	c2	-0.023	0.001	-0.023	0.001	-0.022	0.000	-0.022	0.001
	c1,%	-100.030	0.552	-100.050	0.581	-100.024	0.651	-100.030	0.690
	c2,%	-99.445	2.317	-100.117	2.228	-99.932	1.890	-100.283	2.302
0.9	a1	0.022	0.445	0.028	0.280	0.027	0.226	0.021	0.196
	a2	0.003	0.672	0.015	0.397	0.013	0.271	0.018	0.210
	c1	-0.992	0.006	-0.976	0.006	-0.958	0.005	-0.942	0.006
	c2	-0.025	0.001	-0.025	0.001	-0.025	0.001	-0.025	0.001
	c1,%	-99.971	0.616	-100.115	0.609	-100.043	0.580	-100.086	0.671
	c2,%	-99.171	2.712	-99.524	2.087	-99.922	2.252	-100.222	1.943
0.8	a1	0.025	0.541	0.029	0.315	0.029	0.247	0.023	0.215
	a2	0.013	0.820	0.017	0.441	0.017	0.297	0.017	0.225
	c1	-1.101	0.007	-1.080	0.006	-1.060	0.007	-1.039	0.006
	c2	-0.028	0.001	-0.028	0.001	-0.028	0.001	-0.027	0.001
	c1,%	-99.901	0.635	-100.040	0.551	-100.073	0.625	-100.034	0.616
	c2,%	-99.301	2.461	-99.893	2.154	-99.641	2.289	-99.718	2.057
0.7	a1	0.012	0.696	0.034	0.357	0.036	0.276	0.028	0.241
	a2	-0.008	1.049	0.025	0.500	0.020	0.322	0.019	0.238
	c1	-1.238	0.008	-1.211	0.008	-1.185	0.007	-1.160	0.007
	c2	-0.032	0.001	-0.032	0.001	-0.031	0.001	-0.031	0.001
	c1,%	-99.938	0.612	-100.009	0.662	-100.050	0.617	-100.088	0.643
	c2,%	-99.603	2.304	-99.886	2.258	-99.831	2.140	-100.027	2.199

Table 1: 100-period averages and standard deviations of deviations of coefficients from equilibrium values for a variety of policy rules.

Table 1 continued.

Parameter	φ_π	0.5		1.0		1.5		2.0	
φ_z		mean	std deviation	mean	std deviation	mean	std deviation	mean	std deviation
0.6	a1	0.007	0.924	0.034	0.417	0.037	0.306	0.028	0.268
	a2	-0.023	1.388	0.016	0.583	0.020	0.348	0.020	0.252
	c1	-1.411	0.008	-1.378	0.008	-1.345	0.009	-1.311	0.009
	c2	-0.037	0.001	-0.037	0.001	-0.036	0.001	-0.035	0.001
	c1,%	-99.854	0.596	-100.053	0.604	-100.090	0.688	-99.948	0.713
	c2,%	-98.923	2.349	-99.927	1.843	-100.165	1.913	-100.062	2.001
0.5	a1	-0.041	1.433	0.044	0.480	0.041	0.352	0.030	0.299
	a2	-0.101	2.134	0.022	0.688	0.026	0.380	0.021	0.273
	c1	-1.643	0.013	-1.597	0.010	-1.554	0.009	-1.509	0.009
	c2	-0.043	0.001	-0.043	0.001	-0.042	0.001	-0.041	0.001
	c1,%	-99.845	0.765	-99.976	0.645	-100.105	0.605	-100.019	0.584
	c2,%	-98.769	2.820	-99.978	1.969	-100.187	2.015	-100.239	2.107
0.4	a1	-0.143	2.625	0.051	0.585	0.043	0.397	0.039	0.335
	a2	-0.235	3.768	0.034	0.837	0.035	0.435	0.024	0.301
	c1	-1.963	0.021	-1.904	0.012	-1.837	0.012	-1.778	0.009
	c2	-0.053	0.002	-0.052	0.001	-0.051	0.001	-0.050	0.001
	c1,%	-99.623	1.046	-100.109	0.608	-99.986	0.630	-100.046	0.515
	c2,%	-98.148	3.451	-99.822	2.229	-100.367	2.103	-100.061	2.407
0.3	a1	-0.491	8.206	0.063	0.744	0.063	0.483	0.029	0.399
	a2	-0.714	11.121	0.016	1.070	0.036	0.492	0.030	0.328
	c1	-2.387	0.166	-2.349	0.015	-2.251	0.012	-2.162	0.014
	c2	-0.065	0.008	-0.066	0.001	-0.064	0.001	-0.062	0.001
	c1,%	-97.275	6.756	-100.032	0.617	-100.026	0.518	-100.013	0.623
	c2,%	-95.268	11.276	-99.793	2.001	-99.837	1.736	-100.106	1.801
0.2	a1	-1.649	32.208	0.074	1.042	0.065	0.605	0.054	0.484
	a2	-2.297	41.834	0.006	1.530	0.041	0.581	0.027	0.371
	c1	-2.713	1.188	-3.070	0.017	-2.907	0.019	-2.757	0.016
	c2	-0.076	0.041	-0.090	0.002	-0.086	0.001	-0.082	0.002
	c1,%	-83.432	36.532	-100.053	0.563	-100.077	0.666	-99.998	0.591
	c2,%	-80.053	43.456	-100.164	1.749	-100.339	1.459	-99.822	1.813

Parameter	φ_π	0.5		1.0		1.5		2.0	
φ_z		mean	std deviation	mean	std deviation	mean	std deviation	mean	std deviation
1.1	a1	0.252	0.219	0.176	0.152	0.147	0.123	0.135	0.109
	a2	0.384	0.323	0.243	0.217	0.176	0.162	0.138	0.129
	c1	0.828	0.005	0.817	0.005	0.804	0.005	0.793	0.005
	c2	0.021	0.001	0.021	0.001	0.021	0.001	0.020	0.001
1.0	a1	0.291	0.251	0.195	0.166	0.164	0.133	0.141	0.117
	a2	0.442	0.373	0.263	0.232	0.186	0.171	0.148	0.134
	c1	0.903	0.005	0.889	0.005	0.874	0.006	0.860	0.006
	c2	0.023	0.001	0.023	0.001	0.022	0.000	0.022	0.001
0.9	a1	0.344	0.281	0.217	0.178	0.177	0.143	0.150	0.127
	a2	0.509	0.435	0.295	0.264	0.200	0.181	0.157	0.140
	c1	0.992	0.006	0.976	0.006	0.958	0.005	0.942	0.006
	c2	0.025	0.001	0.025	0.001	0.025	0.001	0.025	0.001
0.8	a1	0.416	0.344	0.242	0.203	0.193	0.157	0.167	0.137
	a2	0.626	0.527	0.328	0.293	0.218	0.202	0.167	0.150
	c1	1.101	0.007	1.080	0.006	1.060	0.007	1.039	0.006
	c2	0.028	0.001	0.028	0.001	0.028	0.001	0.027	0.001
0.7	a1	0.529	0.448	0.275	0.229	0.217	0.174	0.189	0.152
	a2	0.810	0.663	0.371	0.334	0.236	0.219	0.176	0.160
	c1	1.238	0.008	1.211	0.008	1.185	0.007	1.160	0.007
	c2	0.032	0.001	0.032	0.001	0.031	0.001	0.031	0.001

Table 2: 100-period averages and standard deviations of absolute deviations of coefficients from equilibrium values for a variety of policy rules.

Table 2 continued.

Parameter	φ_π	0.5		1.0		1.5		2.0	
φ_z		mean	std deviation	mean	std deviation	mean	std deviation	mean	std deviation
0.6	a1	0.712	0.584	0.321	0.268	0.241	0.191	0.209	0.170
	a2	1.081	0.864	0.440	0.380	0.258	0.234	0.188	0.169
	c1	1.411	0.008	1.378	0.008	1.345	0.009	1.311	0.009
	c2	0.037	0.001	0.037	0.001	0.036	0.001	0.035	0.001
0.5	a1	1.088	0.927	0.369	0.308	0.273	0.224	0.229	0.193
	a2	1.616	1.389	0.518	0.450	0.280	0.257	0.204	0.181
	c1	1.643	0.013	1.597	0.010	1.554	0.009	1.509	0.009
	c2	0.043	0.001	0.043	0.001	0.042	0.001	0.041	0.001
0.4	a1	1.976	1.722	0.448	0.376	0.307	0.253	0.265	0.207
	a2	2.809	2.507	0.628	0.551	0.320	0.294	0.229	0.196
	c1	1.963	0.021	1.904	0.012	1.837	0.012	1.778	0.009
	c2	0.053	0.002	0.052	0.001	0.051	0.001	0.050	0.001
0.3	a1	5.495	6.089	0.568	0.481	0.376	0.307	0.306	0.256
	a2	7.388	8.309	0.810	0.695	0.365	0.329	0.252	0.210
	c1	2.387	0.166	2.349	0.015	2.251	0.012	2.162	0.014
	c2	0.065	0.008	0.066	0.001	0.064	0.001	0.062	0.001
0.2	a1	23.811	21.620	0.800	0.668	0.475	0.378	0.382	0.299
	a2	30.453	28.612	1.155	0.996	0.436	0.385	0.288	0.233
	c1	2.825	0.884	3.070	0.017	2.907	0.019	2.757	0.016
	c2	0.081	0.028	0.090	0.002	0.086	0.001	0.082	0.002

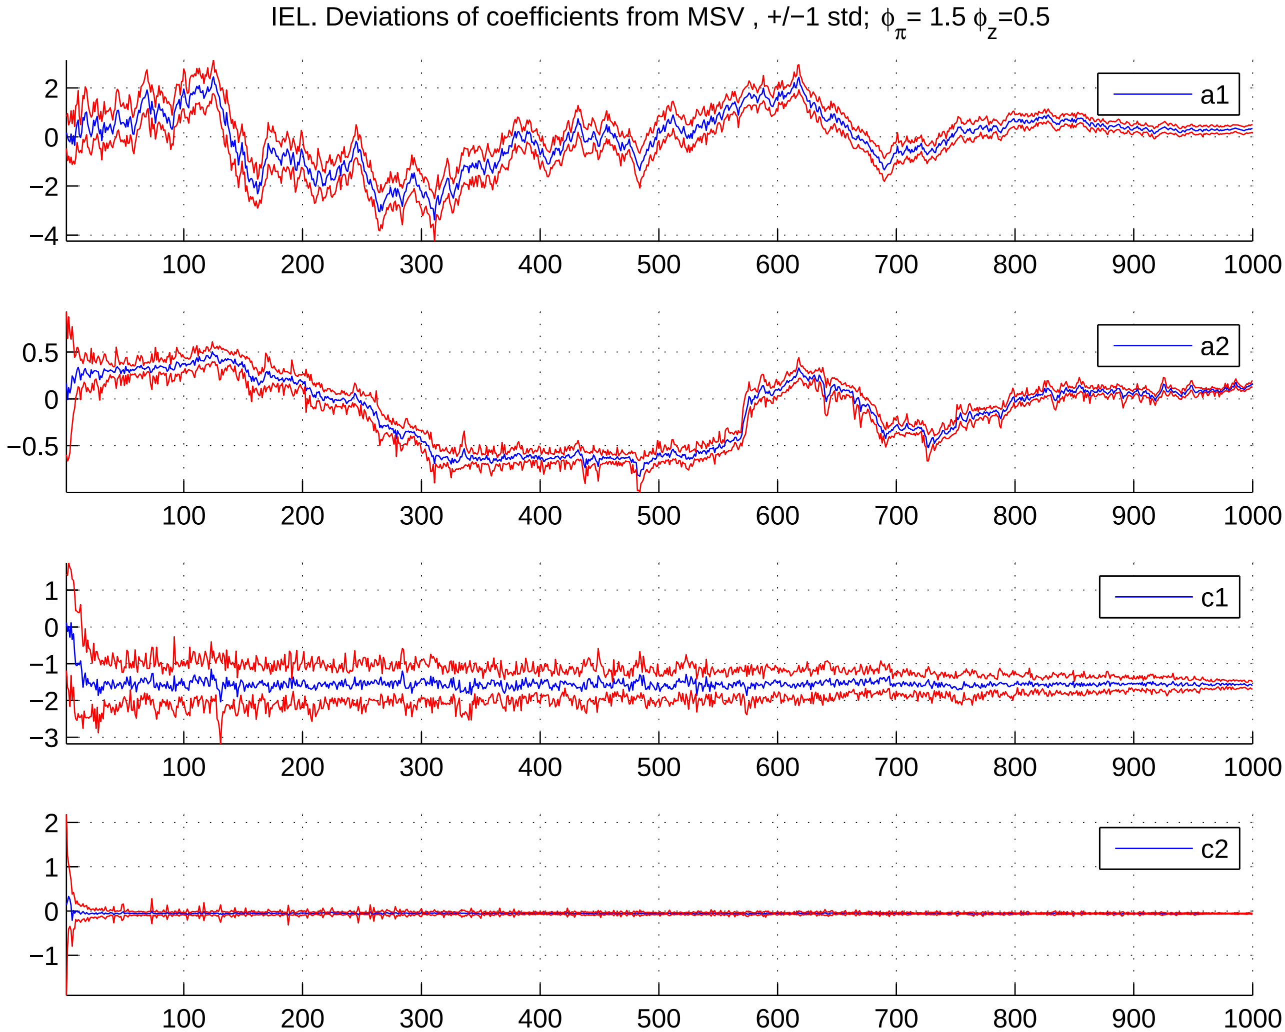

Figure 1: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in determinate and E-stable region with coefficients $\phi_\pi = 1.5$ and $\phi_z = 0.5$.

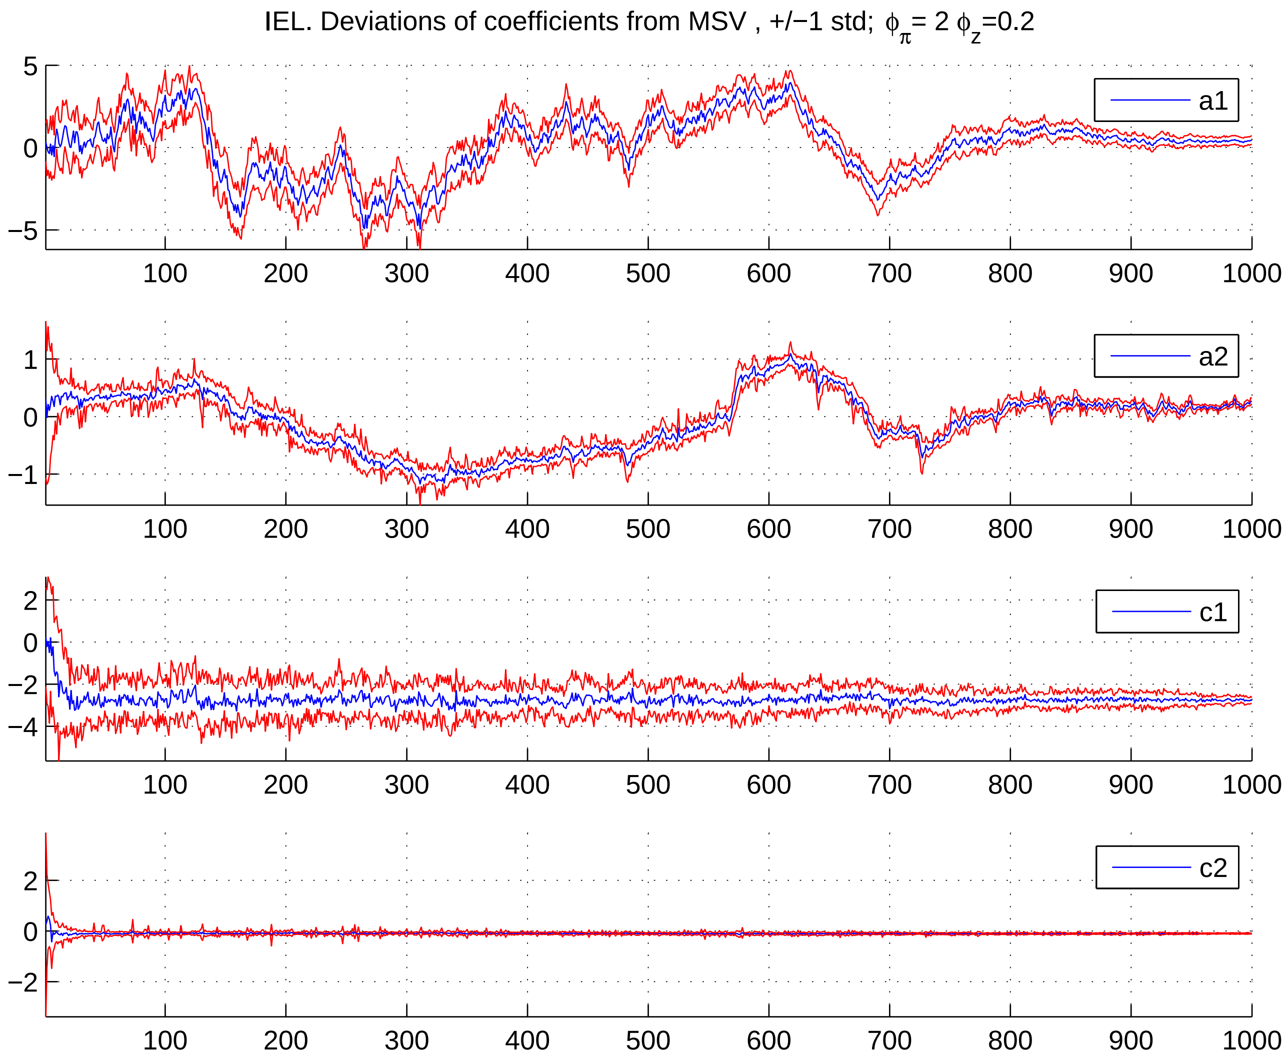

Figure 2: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in determinate and E-stable region with coefficients $\phi_\pi = 2.0$ and $\phi_z = 0.2$.

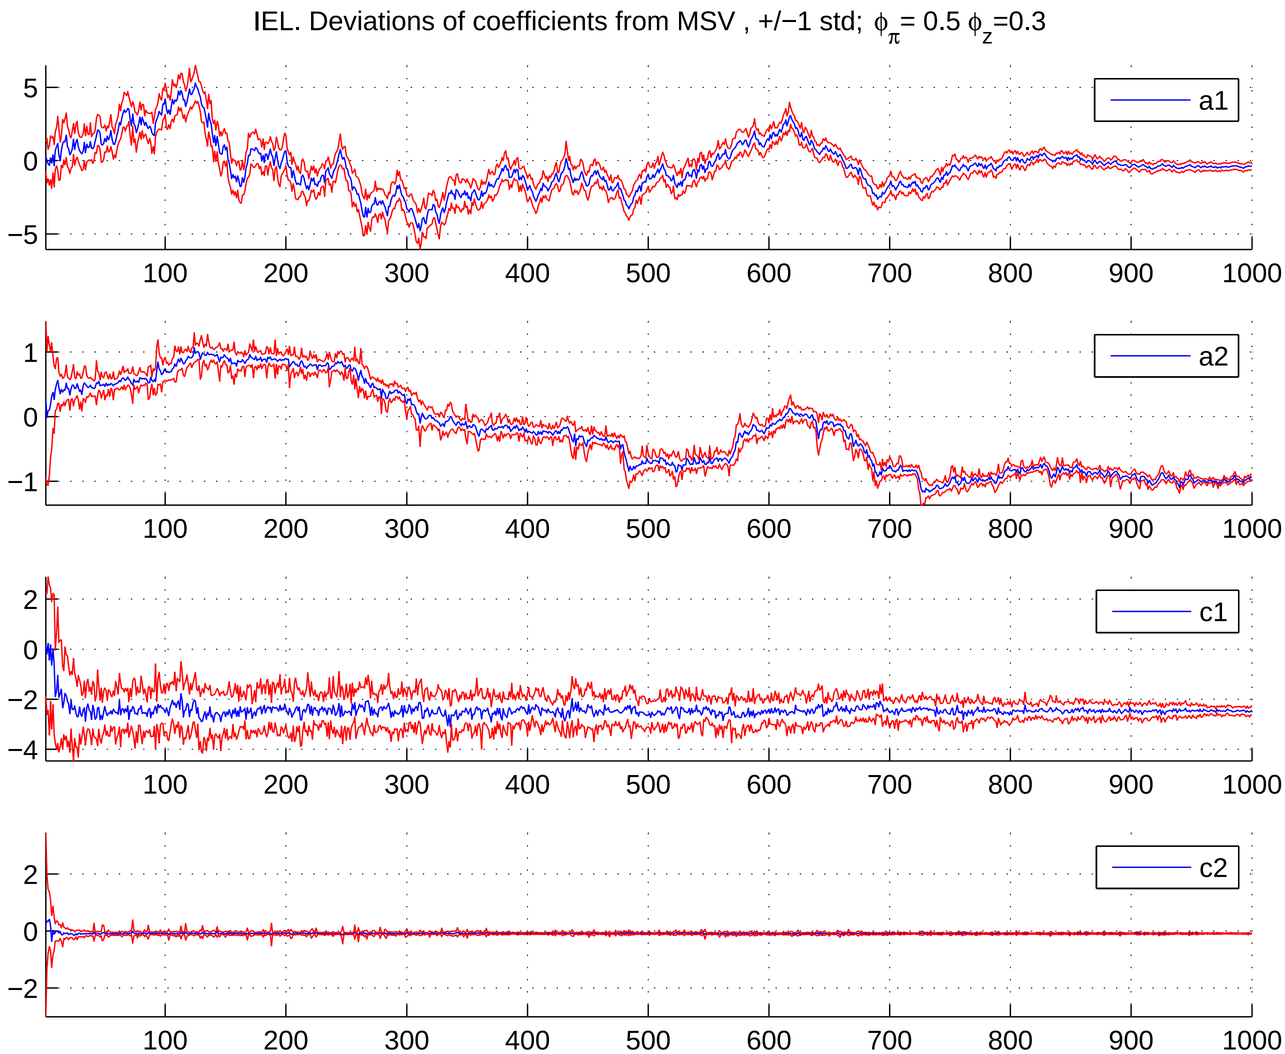

Figure 3: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in indeterminate and E-unstable region with coefficients $\phi_\pi = 0.5$ and $\phi_z = 0.3$.

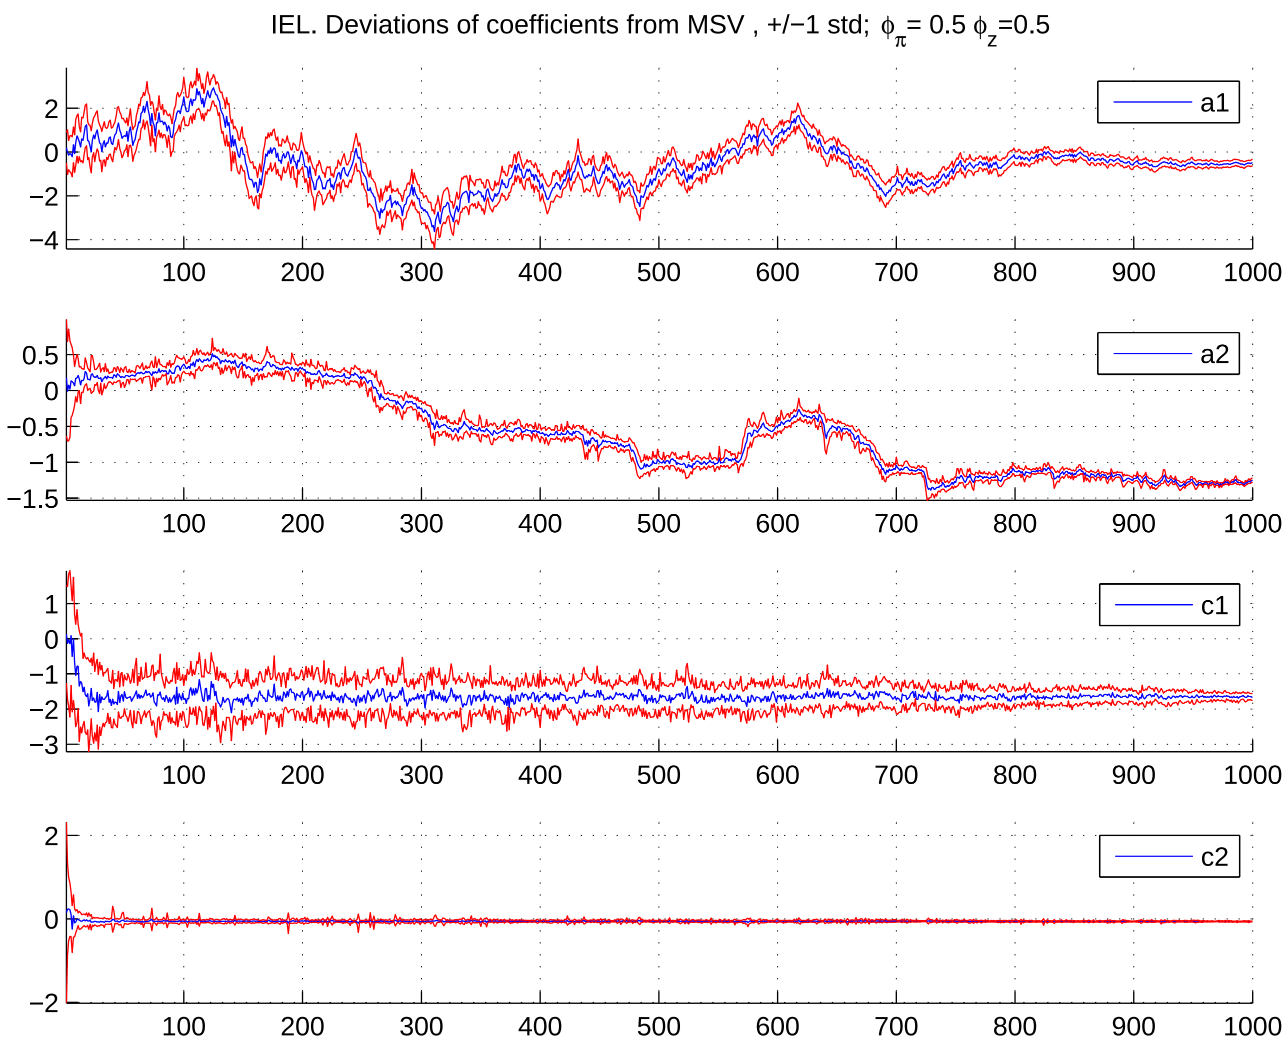

Figure 4: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in indeterminate and E-unstable region with coefficients $\phi_\pi = 0.5$ and $\phi_z = 0.5$.

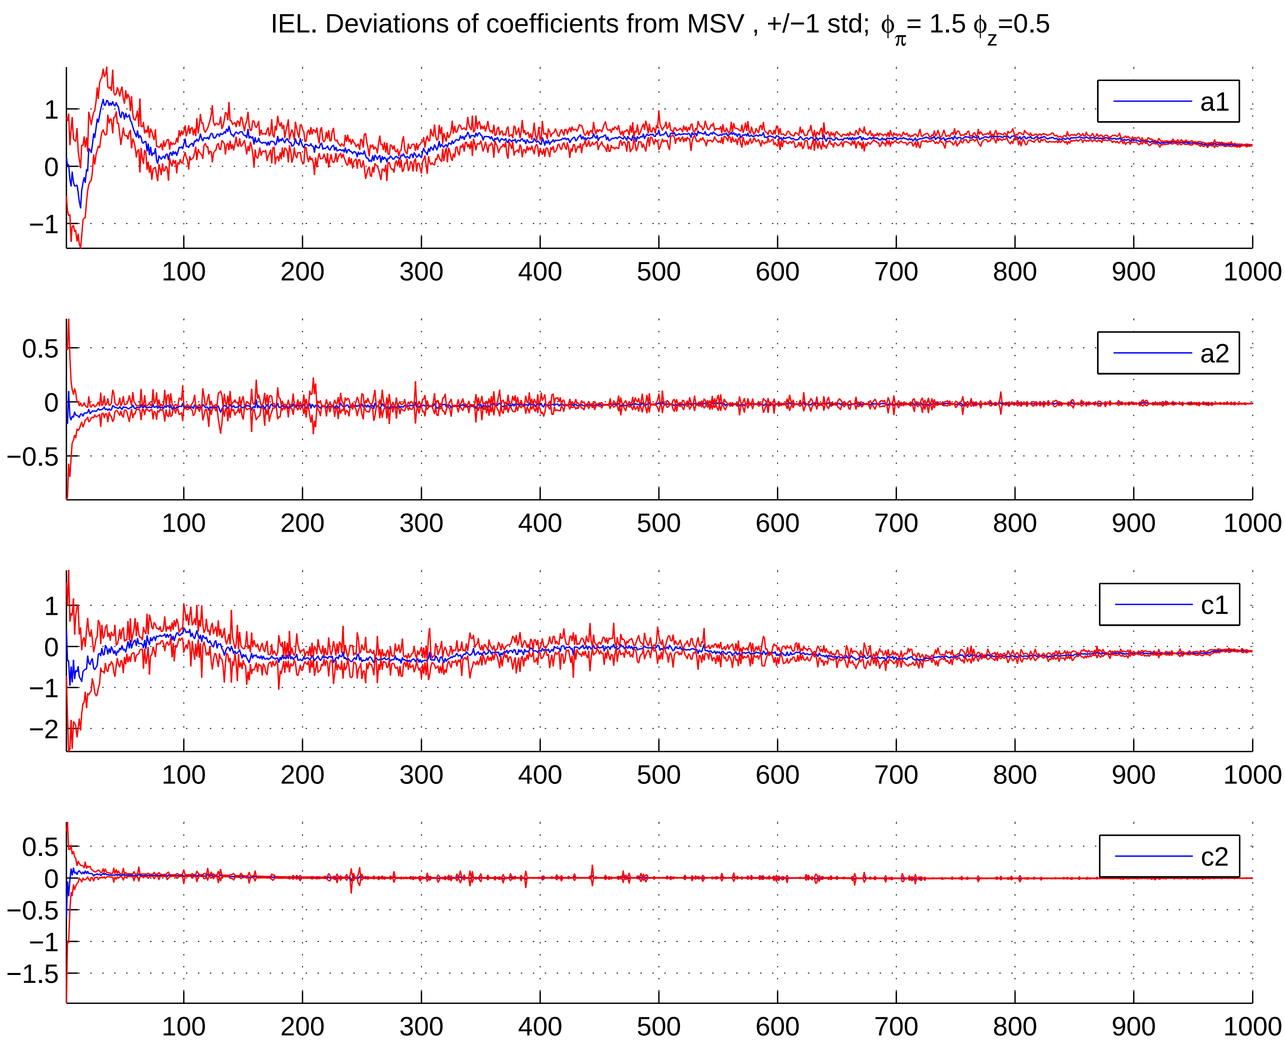

Figure 5: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in determinate and E-stable region with coefficients $\phi_\pi = 1.5$ and $\phi_z = 0.5$ and with performance computed using actual outcomes (26).

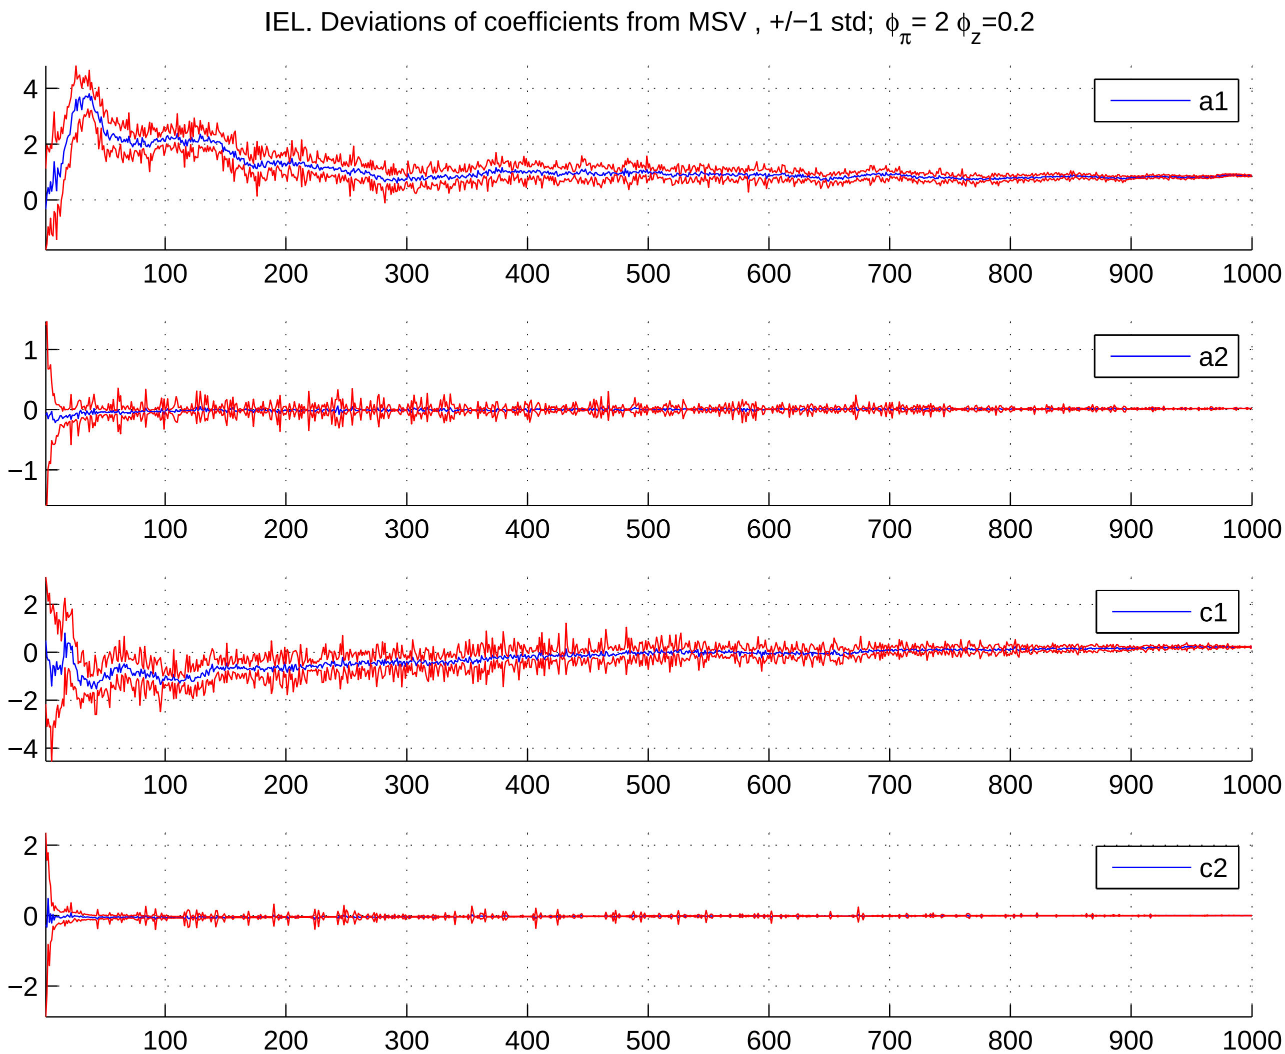

Figure 6: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in determinate and E-stable region with coefficients $\phi_\pi = 2.0$ and $\phi_z = 0.2$ and with performance computed using actual outcomes (26).

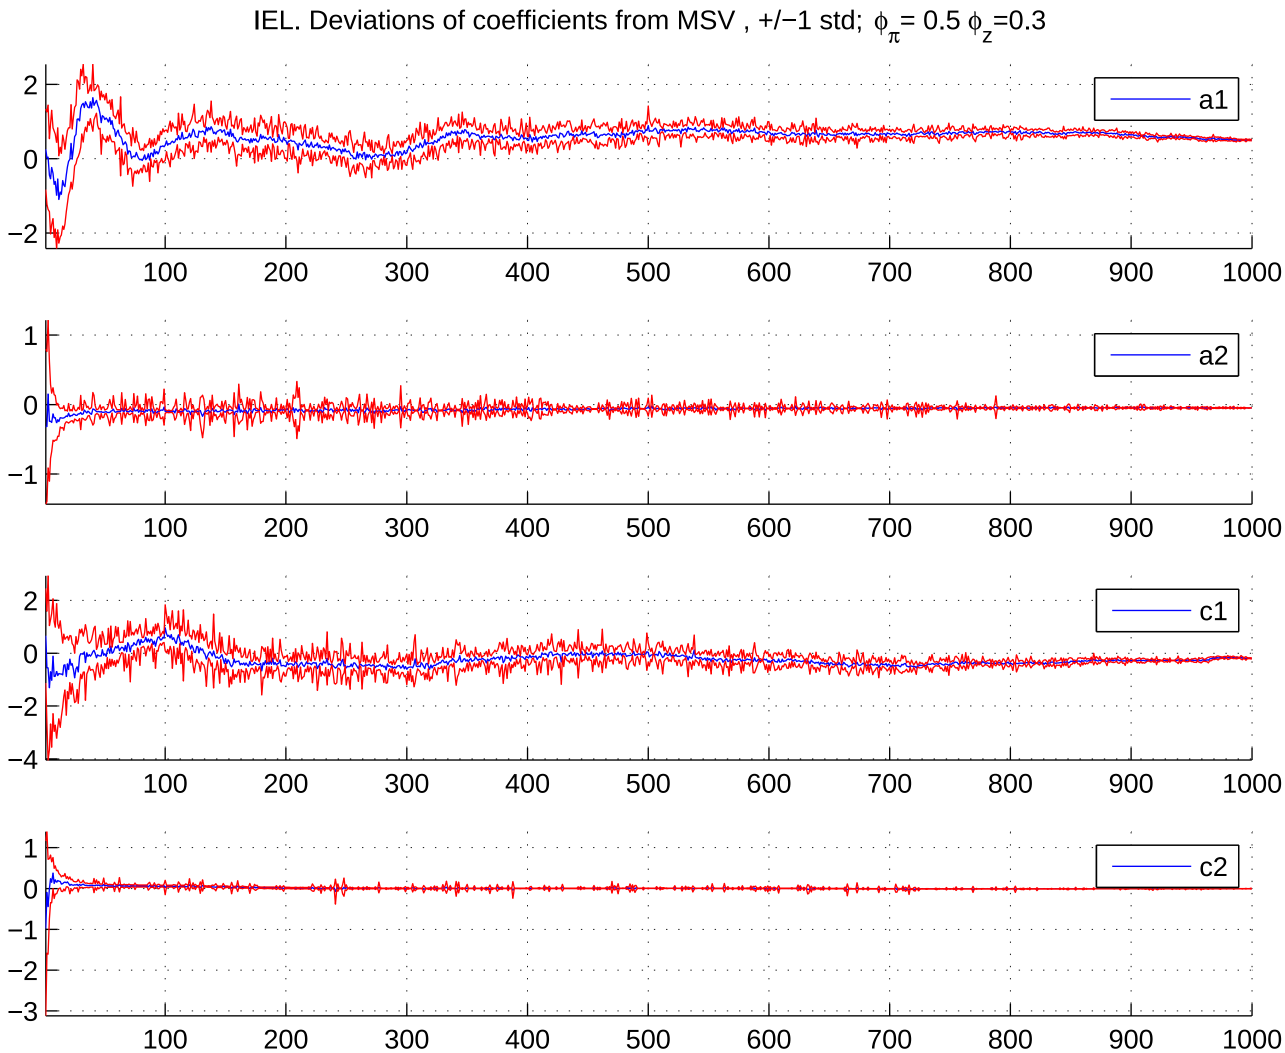

Figure 7: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in indeterminate and E-unstable region with coefficients $\phi_\pi = 0.5$ and $\phi_z = 0.3$ and with performance computed using actual outcomes (26).

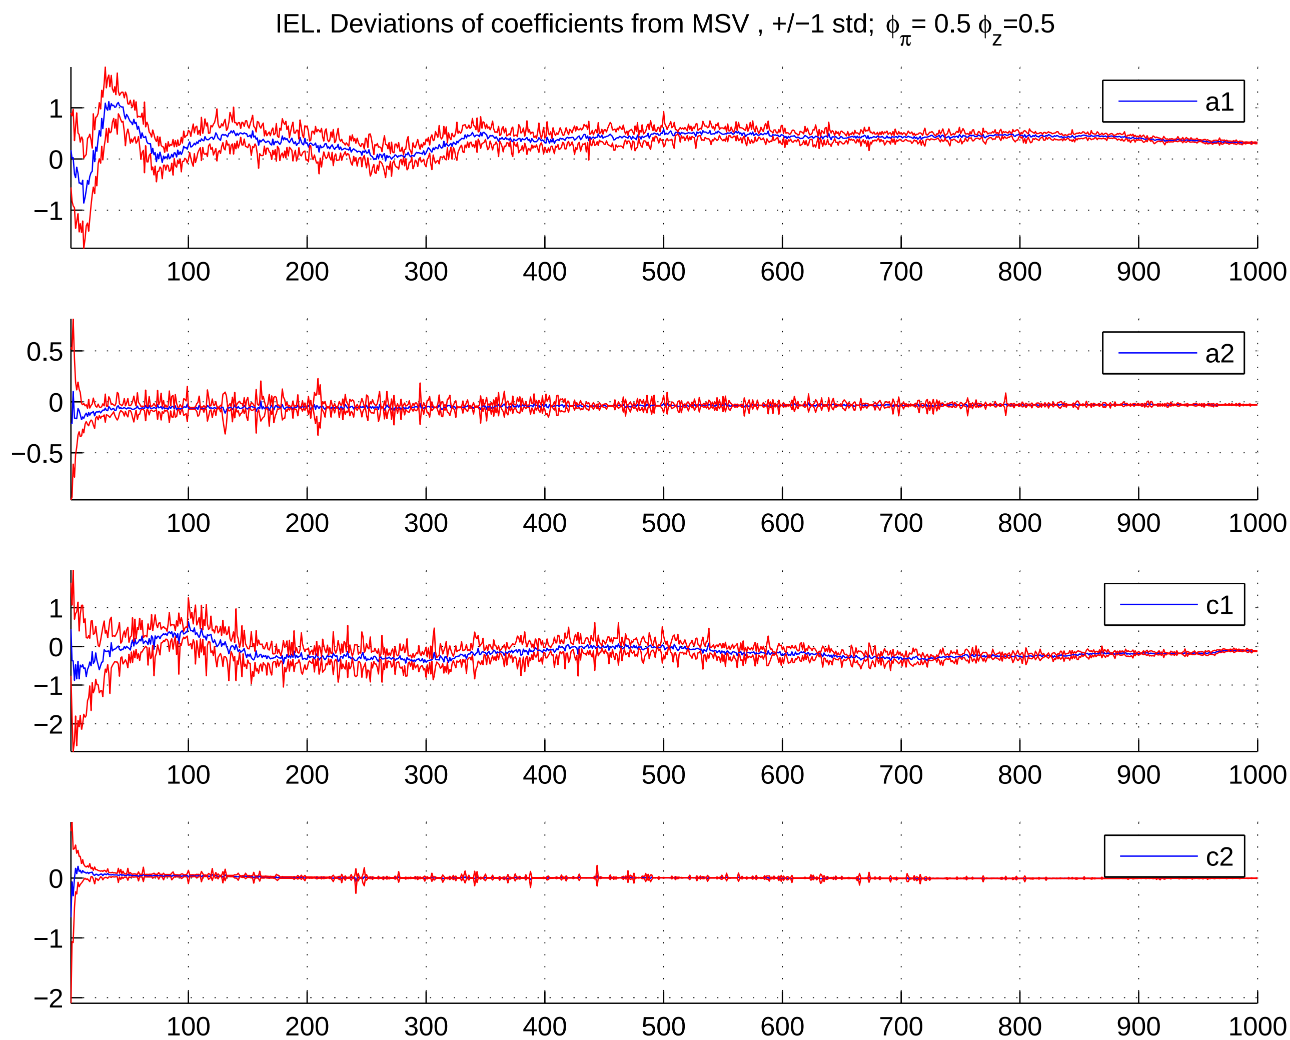

Figure 8: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in indeterminate and E-unstable region with coefficients $\phi_\pi = 0.5$ and $\phi_z = 0.5$ and with performance computed using actual outcomes (26).

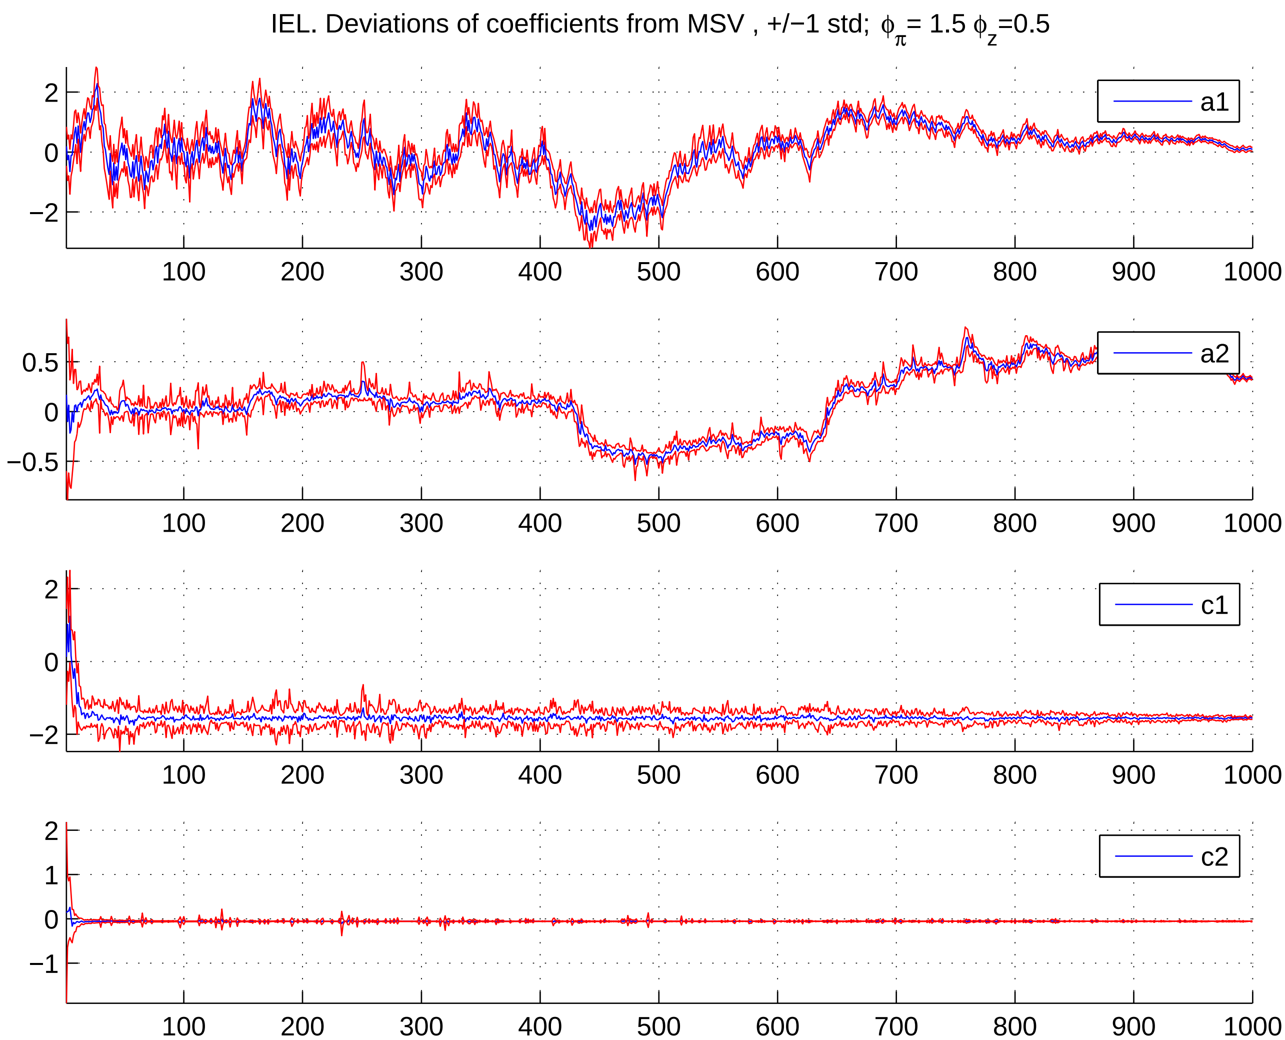

Figure 9: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in determinate and E-stable region with coefficients $\phi_\pi = 1.5$ and $\phi_z = 0.5$ and with separate fitness evaluation (28).

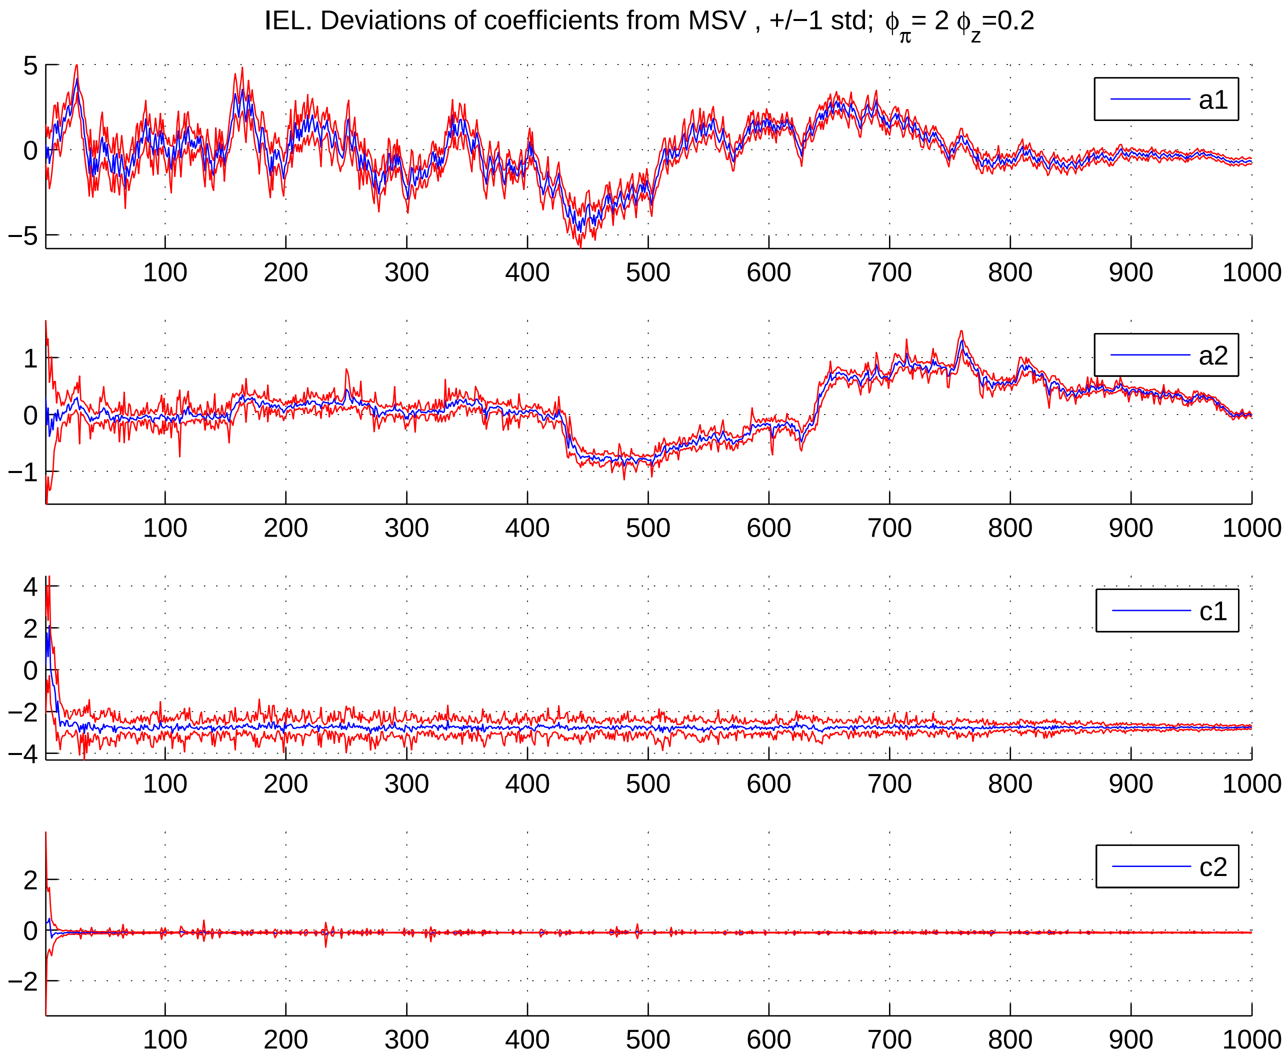

Figure 10: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in determinate and E-stable region with coefficients $\phi_\pi = 2.0$ and $\phi_z = 0.2$ and with separate fitness evaluation (28).

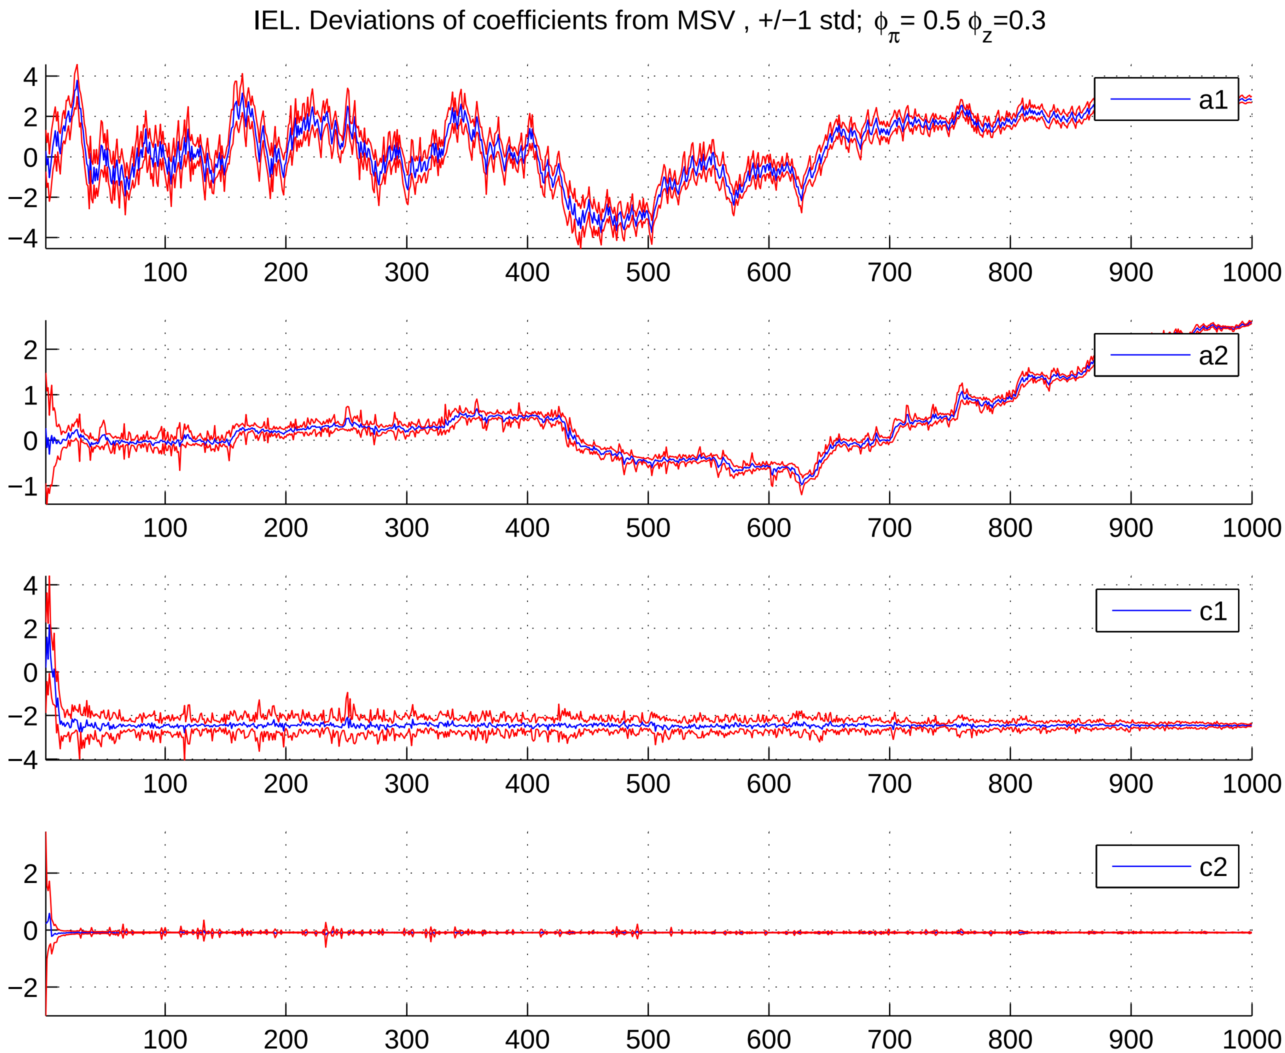

Figure 11: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in indeterminate and E-unstable region with coefficients $\phi_\pi = 0.5$ and $\phi_z = 0.3$ and with separate fitness evaluation (28).

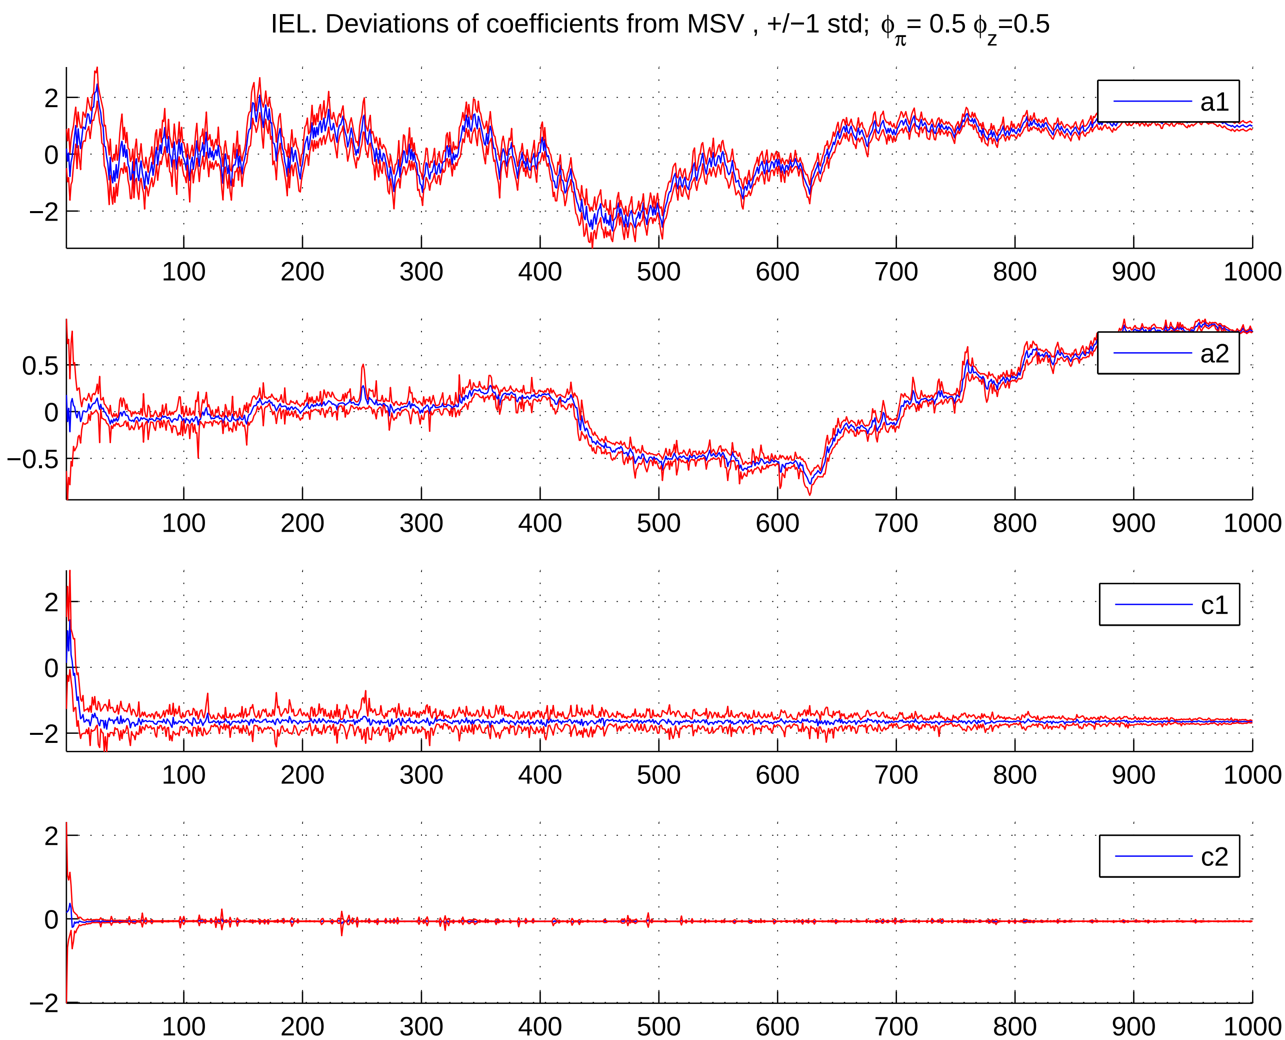

Figure 12: This figure presents the average deviations of actual coefficients in individual decisions from their respective MSV equilibrium values for the simulation in indeterminate and E-unstable region with coefficients $\phi_\pi = 0.5$ and $\phi_z = 0.5$ and with separate fitness evaluation (28).