



Center for Empirical Macroeconomics

Working Paper No. 87

**How Robust is the Equal Split Norm? On the
Destabilizing Effect of Responsive Strategies**

by

Herbert Dawid and J. Dermietzel

University of Bielefeld
Department of Economics
Center for Empirical Macroeconomics
P.O. Box 100 131

33501 Bielefeld, Germany

<http://www.wiwi.uni-bielefeld.de/~cem>

How Robust is the Equal Split Norm?

On the De-stabilizing Effect of Responsive Strategies

Herbert Dawid* Jörn Dermietzel
Department of Business Administration and Economics
University of Bielefeld

October 2004

Abstract

In this paper the evolution of bargaining behavior is studied under the assumption that individuals might choose between obstinate and responsive strategies. Following Ellingson (1997) it is assumed that obstinate agents commit to a certain demand, whereas responsive agents adapt optimally to their opponents strategy. Agents change strategies due to imitation based on observations of the success of other individuals. An agent-based model, where the updating of the population profile is governed by tournament selection and mutation, is used to describe the evolution of behavior. In contrast to existing local and stochastic stability results, which predict robust convergence to an equal split norm in this and related frameworks, the simulations show persistent episodes of substantial deviation of behavior from the equal split. Furthermore, it is shown that parameters governing frequency and type of updating as well as selection pressure have significant impact on the qualitative features of the simulation results. This shows the importance of being able to attach economic interpretations to changes in these parameter values.

Keywords: Bargaining, Nash Demand Game, Equal-Split Norm, Evolutionary Algorithm, Agent-Based Simulation

1 Introduction

Bargaining is one of the central elements of interaction between individuals in many economic environments including markets, contractual relationships, organizations etc. It has been long realized that in a standard one-shot bargaining situation (the so-called Nash demand game) the non-cooperative concept of Nash equilibrium does not provide a unique prediction of the outcome of the bargaining process (Nash (1953)). As alternative approaches, which do provide unique outcomes in this standard bargaining problem, cooperative solutions (Nash (1950)) or the consideration of alternative offer games (Rubinstein (1982)) were suggested. More recently, it has been

*Corresponding Author: Department of Business Administration and Economics, University of Bielefeld, P.O. Box 100131, 33501 Bielefeld, email: hdawid@wiwi.uni-bielefeld.de

argued by Young (1993) that assumptions about common knowledge and coordination made in these approaches in many instances unrealistically strong and actual bargaining behavior is rather determined by myopic best reply based on past observations of bargaining behavior. He shows in a two population model that, if agents intend to act in such a way and there are small but persistent probabilities for strategy implementation errors, a unique bargaining convention evolves in the long run and this outcome corresponds to the (cooperative) Nash bargaining solution where bargaining power is determined by the size of the sample of past bargaining behavior an agent can observe. If individuals in both populations have identical characteristics (in particular identical sample sizes) the stochastically stable convention corresponds to an equal split between agents. Subsequent work using different types of best response dynamics (Binmore et al. (2003)) or a population also including more sophisticated agents (Saez-Marti and Weibull (1999)) suggests that the result about the stochastic stability of the Nash bargaining solution is quite robust with respect to small changes in the considered setup¹.

Whereas these evolutionary bargaining models assume myopic best response behavior of agents, the majority of studies in evolutionary game theory assumes that the frequency of a strategy in a population changes based on its past success. The traditional representation of the evolution of the population profile under such an assumption is the replicator dynamics (see e.g. Hofbauer and Sigmund (1998)) which was initially developed in biological settings but has subsequently been derived as a representation of the evolution of population profiles in the presence of reinforcement-type learning behavior (Borgers and Sarin (1997)) or imitative behavior of agents (Weibull (1995)). A prominent example of the use of the replicator dynamics to study the evolution of bargaining behavior is given in Ellingsen (1997), who argues that *... the approach is particularly justified when studying generic situations such as bargaining, and there is considerable evidence consistent with stories of social evolution*[p. 584]. Ellingsen is not only interested in the distribution of demands observed in the long run, but poses the question which *type* of strategy prevails in the population in the long run. The considered types are, on the one hand, 'obstinate' strategies where a pre-determined demand is made regardless of the opponent's strategy and, on the other hand, the 'responsive' strategy where the agent is assumed to correctly identify the opponent's strategy and to adapt optimally to the opponent's expected demand². The analysis of Ellingsen (1997) shows that with such a strategy set there is no evolutionary stable population profile. However there is a connected set of neutrally stable profiles consisting of all states where more than half of the agents use an obstinate strategy demanding exactly half of the surplus, and all the other agents use the responsive strategy. Standard results establish that each of profiles in this connected set are Lyapunov stable with respect to the replicator dynamics. Note that in a population consisting only of obstinate agents demanding half the surplus and responsive agents bargaining always results in an equal split. Ellingsen also shows that no profile outside the considered connected set can be neutrally stable and concludes: *The unique equilibrium outcome is an equal split*[p. 593].

¹The stochastic stability of the Nash bargaining solution disappears however in most settings where non-contractible surplus-increasing investments are incorporated into the analysis, see Ellingsen and Robles (2002), Tröger (2002) or Dawid and MacLeod (2001, 2004).

²Ellingsen refers to some kind of truthful communication which allows an agent to identify the strategy of her opponent. However, he does not give an explicit description of the underlying communication protocol. The interaction of obstinate and responsive strategies can be seen as a stylized reduced form representation of the different degrees of ex-ante commitment of negotiation parties to their demands. Such differences can indeed be observed in many bargaining situations, take for example collective bargaining between unions and management, where the strength of ex-ante made public statements differ quite a bit between different negotiations and between parties.

Based on all these positive findings concerning the emergence and stability of the equal split one could consider the standard two-person bargaining problem as a prime example of an economic environment where adaptive behavior of agents endogenously creates efficient norms of behavior inducing no disagreement and a fair split of the surplus³.

In this paper we will however show that the interplay of obstinate and responsive strategies – as considered in Ellingsen (1997) – might lead to repeated time windows of disagreement and substantial average deviation from the equal split rule if the strategy adaptation process is subject to small but persistent noise (as is assumed in the stochastic stability literature). In order to make our point we present results from agent-based simulation runs, where agents in a population repeatedly play a slight variation of the Nash demand game, the strategy set contains obstinate and responsive strategies and the population profile is updated using a selection and mutation operator. We will then go on to explore how different aspects of the strategy updating procedure and in particular the selection operator influence the average deviation of demands from the equal split. Indeed it will turn out that updating frequency, selection pressure or the fact whether updating is simultaneous or continuous has statistically significant impact on the simulation results. Since these parameters might be seen as purely technical aspects of the evolutionary algorithm employed in the simulation exercise, this might raise serious questions about the robustness of the simulation setup and the appropriateness of the tool of agent-based simulation to study the evolution of bargaining norms. Quite on the contrary, we argue in this paper, that our findings demonstrate that analytical local stability results based on aggregate dynamics are potentially misleading. The use of agent-based computer simulation allows to explicitly model the micro interaction structure underlying population level dynamics. Our results highlight that different interaction structures, which might all *ex-ante* be considered consistent with replicator type dynamics, can yield quite different results. This observation should actually sharpen the attention to be paid to micro foundations of the considered models. If it is possible to attach a clear economic interpretation to the variations in the used operators, sensitivity of results with respect to such variations can give rise to additional qualitative insights.

There is already quite a bit of literature where agent-based models or evolutionary algorithms have been employed to study the development of negotiation strategies in repeated bargaining games (e.g. Oliver (1996), Dawid (1999), Carpenter (2002), Gerding et al. (2003), see also the literature review in Gerding et al. (2000)). Most closely related to our work is the analysis of Carpenter (2002), who compares the outcomes of agent-based simulations with the prediction of the replicator dynamics analysis in several bargaining models including the setup used in Ellingsen (1997). Several features of his simulation setup ensure a close resemblance of simulation results to the replicator dynamics. There is no selection at an individual level, rather the fraction of each strategy in the following period is directly determined to be proportional to the average payoff of this strategy in the current period. Additionally, a very large population (1000 agents) is considered and there are no mutations. Indeed it is shown that in such a setup in most cases the trajectory generated by the agent-based model is very close to that of the replicator dynamics with respect to convergence speed and long run behavior. Furthermore, it is shown that a certain minimal fraction of obstinate equal split players is necessary to yield convergence towards one of the neutrally stable states identified by Ellingsen (1997). As pointed out above, our approach is quite different from that chosen in Carpenter (2002). We are not so much interested in the

³We assume in what follows that agents do not differ with respect to preferences, strategy adaptation rules or available strategies.

question whether replicator dynamics can be well approximated by agent-based simulations, but rather consider different variations of a rather standard evolutionary simulation setup and explore how the probability for the emergence of an equal split norm is influenced by such variations.

The paper is organized as follows. In Section 2 the bargaining game is introduced and the different simulation setups we employ are discussed. Simulation results are presented in several subsections in Section 3 and we conclude with a brief discussion of our findings in Section 4.

2 The Model

2.1 The Bargaining Game

As pointed out above, we adopt a model very close to Ellingsen's (1997) formulation. Consider a population of n individuals who repeatedly bargain with changing opponents over a given joint surplus. More precisely, in each period t each agent is (uniformly) randomly matched with one other agent in the population and the two play a one-shot bargaining game similar to the Nash demand game (every period $n/2$ such games are played in the entire population). Both players simultaneously choose their demands from a finite set $D = \{0, \frac{1}{m}, \frac{2}{m}, \dots, 1\}$ with even $m \in \mathbb{N}$ and the resulting payoff for player i is given by

$$\pi_i(d_i, d_j) = \begin{cases} \frac{d_i}{d_i + d_j} & d_i + d_j \leq 1 \\ 0 & d_i + d_j > 1 \end{cases},$$

where d_i denotes the demand of player i .

Agents use pure strategies when playing the game, where the set of pure strategies is denoted by I . Following Ellingsen (1997) it is assumed that the opponent's can mutually identify the other's strategy after they have been matched. Accordingly, a strategy in this setup is a mapping from I to D . There are two types of strategies. First, for each possible demand $d \in D$ there is an *obstinate* strategy o_d committing to demand d regardless of the opponent's strategy, i.e. $o_d(i) = d, \forall i \in I$. Second, there is a *responsive* strategy r , which chooses the maximal feasible demand against an obstinate strategy and half the surplus against another responsive strategy, i.e. $r(o_d) = 1 - d \forall d \in D$ and $r(r) = 0.5$. Note that overall there are $m + 2$ pure strategies in I . A *population state* or *population profile* is given by a vector $s = (s_0, s_1, \dots, s_m, s_r)$, where the set of all population states is $S := \{s \in [0, 1]^{m+2} \mid \sum_{i \in I} s_i = 1, ns_i \in \mathbb{N} \forall i \in I\}$.

If we consider the Nash equilibria of this n -player game (assuming that all players simultaneously choose their strategies prior to being matched), it turns out that only two types of Nash equilibrium profiles exist. First, a state where all individuals choose o_m (i.e. every agent uses an obstinate strategy demanding $d = 1$) is a Nash equilibrium yielding a payoff of zero for all players. Second, any state s where all players either use the 'fair' obstinate strategy $o_{m/2}$ or the responsive strategy r is a Nash equilibrium. In any such equilibrium bargaining results in an equal split with payoff 0.5 for each player. Furthermore, Ellingsen (1997) shows that among these Nash profiles only the ones with $s_{m/2} > 0.5, s_r = 1 - s_{m/2}$ are neutrally stable. It should be noted that, if the strategy set would only include obstinate strategies, for each $d \in D$ there would be a Nash equilibrium profile s where $s_d + s_{1-d} = 1$.⁴ Hence, from a game-theoretic point of view the addition of the responsive strategy to I to a large extent removes the problem of multiplicity of equilibria which is typically

⁴To be precise, for each d there exists a population size n such that there exists a Nash equilibrium profile $s \in S$

encountered in one-shot bargaining. In this setup every stable equilibrium induces an equal split. This might suggest that a fast and smooth convergence of adaptive bargaining behavior towards the equal split norm should be expected. However, as we will see, in a dynamic setting the addition of r to the strategy set has on the contrary quite destabilizing effects.

2.2 Updating of Strategies

In our simulations a simple evolutionary algorithm is employed to model how the bargaining strategies of agents change over time. At a given point in time t an agent a is characterized by her bargaining strategy $i_{a,t} \in I$ and the *average* payoff $\hat{\pi}_{a,t}$ earned by this agent in all bargaining games she participated since the last updating of her bargaining strategy. The population is initialized randomly where each strategy in I has equal probability. During the simulation each agent from time to time reconsiders and updates the currently used bargaining strategy. We will examine two different ways to determine in which periods updating takes place:

- **Simultaneous Updating:** Every L periods all agents in the population simultaneously update their strategies. In this setup the average payoffs of all agents are at all times based on the identical number of bargaining games played.
- **Continuous Updating:** Each period every agent updates her strategy with a probability $1/L$, where draws leading to updating are independent across agents and time. In this setup the number of bargaining encounters on which the current average payoff value is based might vary significantly between agents.

We will sometimes refer to the assumption about the way updating periods are determined as an updating regime. Although simultaneous updating is the standard assumption in most applications of evolutionary algorithms and also in most studies in the field of agent-based computational economics, we believe that there are several good reasons to consider continuous updating as well. First, in many instances it is not clear why individuals who infrequently decide to change their current strategy should do so all at the same time. Second, we interpret the selection procedure in terms of an imitation process and here the heterogeneity of 'experience' of the different agents to be imitated is in our opinion an interesting, realistic but so far often neglected feature of the process. Third, many of the studies dealing with replicator dynamics, including Ellingsen (1997), use a continuous time version – thereby avoiding to have to deal with overshooting phenomena – and continuous updating is certainly the more natural equivalent for a continuous time dynamics.

If an agent updates her current strategy a selection operator is used to choose the new strategy. We will consider two different types of selection operators. In order to stick to the most standard framework and also as close as possible to the analyses of Ellingsen (1997) and Carpenter (2002) we will initially present results obtained with proportional selection. In what follows we will then however focus on runs where tournament selection is used. The exact definitions of the two types of selection operators are as follows:

- **Proportional Selection:** One agent in the population is chosen randomly and her strategy is adopted by the updating agent. The probability for an agent a to be chosen is proportional to $\hat{\pi}_{a,t}$, where t is the period where the updating takes place.

with $s_d + s_{1-d} = 1$; see Dawid (1999) for a precise characterization of the Nash equilibria in the slightly more complex but closely related setup of a sealed-bid double auction game played between two populations.

- **Tournament Selection:** A subset K consisting of k agents is randomly selected, where each agent has equal probability to become a member of K . The members of K are ordered with respect to their payoffs and the strategy of the agent with the highest payoff is adopted. If several agents split the top, the strategy of each of them is adopted with equal probability.

The reason we prefer to deal with tournament selection in most of our simulation runs is that tournament selection has a very straight-forward interpretation as an imitation process where each agent can only observe the success of a certain subset of agents in the population. This approach also allows to interpret an increase in selection pressure directly as an increase of the number k of agents whose payoff can be observed in the updating period. Stronger selection pressure therefore corresponds to increased information flows and more transparency in the population.

After selection has been carried out, the mutation operator is applied with probability $\epsilon > 0$. In case of a mutation the selected strategy is replaced by a randomly drawn strategy $i \in I$, where the probability for each strategy to be drawn is $1/(m + 2)$.

This completes our description of the simulation model. The variations of the updating process we will consider in the next section concerns on the one hand the choice between simultaneous and continuous updating and on the other hand examines effects of changes in the updating frequency (L) and the selection pressure (k).

3 Simulation Results

In our simulations we try to capture transient and long run evolution of bargaining behavior under persistent noise. In this sense the approach taken here is similar to the philosophy underlying the stochastic stability analyses, namely that the existence of punctuated equilibria is a usual phenomenon for evolutionary processes of norm formation under persistent noise. There are long periods of apparent convergence of the process interrupted by sudden mutation driven changes to a different attractor. Processes of this kind are best described by their statistical properties, in the simplest case by considering time averages of the variables of interest. To be able to observe switches between punctuated equilibria the simulations must be run for an extensive number of periods⁵. Of course there is a trade-off with computing time. Based on experiments with different simulation durations we have chosen a duration of $T = 50000$ periods for all simulations reported in this paper. Furthermore, we use parameter values $n = 100, m = 10, \epsilon = 0.01$ in all the simulations. Varying these parameter in a sensible range around the current values will not qualitatively alter our findings.

3.1 The Destabilizing Effect of Responsive Strategies

In order to demonstrate the effect of the existence of responsive strategies on the evolution of bargaining behavior we start with considering a slightly changed setup where the set of pure strategies I consists only of obstinate strategies. In figure 1 the evolution of the fraction of agents using the fair obstinate strategy, which always demands $d = 0.5$, is depicted for two single runs, one under

⁵Considering such a long time horizon might raise concerns about the economic relevance of the findings. However, first, bargaining situations are very common in economic transactions meaning that one simulation period might well correspond to less than a day, and, second, dealing with norm formation, which is an inert process, consideration of long time windows sees quite adequate.

proportional selection and one under the strongest possible selection scenario, namely tournament selection with $k = 100$.⁶ In both cases simultaneous updating with low frequency ($L = 50$) is used.

Insert figure 1 here

Both frames of figure 1 indicate that very quickly the equal split norm emerges and (almost) all players demand $d = 0.5$. The presence of persistent mutations leads to small deviations of s_5 from $s_5 = 1$, where the deviations are slightly more pronounced in the case of proportional selection where selection pressure is weaker compared to imitate the best. The runs presented here are by no means exceptional. Quite on the contrary, although without the responsive strategy there exists a large number of Nash profiles inducing no equal split, in our simulations we observe convergence towards the equal split norm. This finding is very robust with respect to variation of the parameters L, k or the updating regime.

The observed behavior changes quite a bit if the responsive strategy is included in the strategy set. In figure 2 we show the evolution of the fraction of agents using o_5 , the fraction of agents using the responsive strategy r , the average absolute deviation of demands from the equal split, and the average population profit for a representative run where the simulation setup is exactly like the one used in figure 1 (a) except for the strategy set, which now includes all obstinate plus the responsive strategy.

Insert figure 2 here

Just considering the fraction of fair obstinate strategies the difference to the trajectory depicted in figure 1 (a) is indeed striking and one gets the impression that the equal split norm is rather unstable. Taking into account also the fraction of agents using the responsive strategy and the actual demands (figure 2 (c)) we realize however that many fluctuations of the frequency of fair obstinate players is due to a replacement of the fair obstinate strategy by the responsive strategy. In an environment of fair and responsive strategies such a replacement has no influence on the demand made by the agents. Put differently, we can observe some drift within the connected set of neutrally stable equilibria identified in the analysis of Ellingsen (1997). Average demands are close to the equal split for a very long time window starting after a short transient phase and lasting till about period 25000. At this point however there is a complete breakdown of the equal split norm and for almost 10000 periods average demands are substantially higher than $d = 0.5$ resulting in frequent disagreement and low profits. After these 10000 periods the equal split norm re-emerges and stays in place till the end of this run. Although the number of switches between existence and breakdown of the equal split norm and the duration of the different phases vary quite a bit between different individual runs, the qualitative pattern of 'regime switching' between the existence of an equal split norm and phases of inflated demands with frequent disagreement is very robust in this framework. To gain a better understanding of what triggers the breakdown of the equal split norm and its subsequent re-emergence, we present in figure 3 for the same run the evolution of the frequency of six different strategies in the crucial time window between $t = 15000$ and $t = 35000$.

⁶This type of selection is often also called 'Imitate the Best', since the best paying strategy is imitated by the entire population. Obviously, such strong selection leads to population distributions with zero or very small variance.

Insert figure 3 here

Almost till period 20000 we can observe population states consisting almost entirely of obstinate fair players. Shortly before $t = 20000$ however a substantial fraction of agents switches to the responsive strategy. This is due to some initial mutations towards responsive strategies followed by a short window of existence of obstinate strategies with $d = 0.2$ (inducing some selective pressure towards responsive strategies) and a certain amount of drift typical for proportional selection under weak fitness differentials. Nevertheless, the fair split stays the prevalent bargaining outcome till about $t = 23000$ when mutation creates a small pocket of obstinate players who demand the entire surplus. Due to the large fraction of responsive players in the population this strategy is very profitable and quickly spreads. The spread of this very immodest⁷ strategy quickly wipes out the fair obstinate and the responsive strategies from the population profile, but once the responsive players are gone, the immodest obstinate players have no potential partners to agree with. There is a complete breakdown of any bargaining norm with disagreement in virtually all bargaining encounters and an average payoff in the population close to zero. After this breakdown we see transient phases of populations consisting of 'matching' obstinate players where some demand $d > 0.5$ and others $1 - d$. Such population profiles are profitable environments for the responsive strategy, so we typically observe a re-emergence of responsive strategies which in turn opens the door for over demanding obstinate strategies and another breakdown. This pattern continues until a population of fair obstinate players emerges after a breakdown and we are back at the equal split norm which in general will then be upheld for an extensive period of time. The way the repeated breakdown of efficient behavior followed by a period of turmoil with low payoffs and finally a re-emergence of efficiency is generated in this model resembles observations which have been made concerning the evolution of finite memory strategies playing repeated Prisoner's Dilemma games (see Nowak and Sigmund (1992, 1995)). This discussion shows that the evolution and stability of the equal split norm in the symmetric one-shot bargaining problem is by no means a 'done deal' if agents have sufficient information such that they might react responsively to their opponent's strategy.

3.2 The Effect of Selection Pressure

A crucial step on the way to the breakdown of the equal split norm for proportional selection is the emergence of a heterogenous population consisting of obstinate fair and responsive players during a period of time where equal split is still the prevalent outcome of bargaining. If selection pressure is sufficiently strong, the emergence of such a mixed population is highly unlikely and therefore one might expect that increased selection reduces the danger of a breakdown of the equal split norm.

In figure 4 we show the equivalent set of trajectories as in figure 2 under tournament selection with $k = 100$ ('Imitate the Best'), $L = 50$ and simultaneous updating.

Insert figure 4 here

⁷Ellingsen (1997) calls strategies demanding less than $d = 0.5$ *modest*, strategies with $d > 0.5$ *immodest*.

Compared to the run with proportional selection we observe a larger number of complete bargaining breakdowns of the equal split norm, but the duration of the breakdowns is significantly smaller. The reason for both effects is that in this setup even small payoff differentials induce an immediate switch to an almost uniform population and hence few mutations can have a very pronounced effect on the population distribution. In a uniform state where all agents follow the fair obstinate strategy two simultaneous mutations, one to a responsive strategy and one to some obstinate strategy with $d < 0.5$ can be sufficient to generate a payoff differential in favor of the responsive strategy⁸ inducing a shift to a uniform state consisting only of responsive strategies. As we know, such a state is then very vulnerable to the invasion of immodest obstinate strategies and a breakdown of the equal split norm. On the other hand, during time intervals where the equal split norm has broken down the population goes through a fast transition of uniform states consisting only of obstinate agents making identical demands. Any such state where the demand differs from 0.5 is highly unstable and therefore after a relatively short number of periods the population 'samples' the uniform state with only fair obstinate strategies and the equal split norm re-emerges.

A casual comparison of the two runs considered in figures 2 and 4 shows substantial qualitative differences but does not give a clear indication whether increased selection pressure has in general a positive or negative effect on the average deviation (in the sense of time averages) from the efficient equal split norm. In order to examine this question more systematically we compare simulations under tournament selection with tournament size $k = 5, k = 30$ and $k = 100$. In terms of an economic interpretation these three cases correspond to scenarios where individuals can observe the success of a randomly selected group of agents – e.g. due to random meetings in the street – and imitates the best in this group. The larger the size of the group individuals can observe the stronger the selective pressure on strategies and the faster becomes behavior uniform. In this sense, the parameter k might be interpreted as a measure of how 'closely' connected the considered population of people is. The qualitative findings we are going to present does not crucially depend on the fact that we stick to Tournament selection in our presentation of the statistical results. We did also carry out extensive simulations under proportional selection and comparison of these results with Imitate-the-Best runs is in accordance with what we discuss here.

Our previous discussion of the mechanisms responsible for the breakdown and re-emergence of the equal split rule suggests that the effect an increase of selection pressure has might sensitively depend on the updating frequency and also the updating regime used in the simulation. In order to account for that, we consider four different scenarios when comparing the results for different k -values:

- I:** Simultaneous updating, $L = 50$
- II:** Simultaneous updating, $L = 5$
- III:** Continuous updating, $L = 50$
- IV:** Continuous updating, $L = 5$

In each of these four scenarios we have carried out 20 runs for $k = 5, 30, 100$. To analyze the deviation of the population profile from the equal split norm in a concise way, we focus on the expected absolute deviation of actual demands from $d = 0.5$ under the population profile at each

⁸In order for this differential to materialize it is necessary that the two mutants actually meet till the next updating period, so the value of parameter L is of some importance here.

period, and average over the entire run. Put formally, for each run we calculate

$$\begin{aligned}\Delta d &= \frac{1}{T} \sum_{t=1}^T \left[\sum_{i=0}^m s_{i,t} \left| \frac{i}{m} - 0.5 \right| + s_r \sum_{i=0}^m s_i \left| \left(1 - \frac{i}{m} \right) - 0.5 \right| \right] \\ &= \frac{1}{T} \sum_{t=1}^T (1 + s_r) \left[\sum_{i=0}^{m/2-1} \left(0.5 - \frac{i}{m} \right) + \sum_{i=m/2+1}^m \left(\frac{i}{m} - 0.5 \right) \right]\end{aligned}$$

Note that the second term in the first line is due to the fact that the expected demand of a responsive agent and therefore also the average deviation from the equal split depends on the distribution of strategies in the population. We denote by Δd_{av} the average value of Δd across the 20 runs we have carried out. Two comments are in order here. First, we use the time average of the entire run because the initial transient period is negligible for a simulation duration of $T = 50000$ and therefore the time average should be a good proxy of the expected demand under the limit distribution of the process. Second, ex-ante it is not obvious that the size of the average deviation of demands from the equal split norm also gives a good indication of the amount of time the equal split norm is upheld. We have also calculated direct measures of the time the population adheres to the equal split norm, but since it turned out that large average deviations are strongly correlated with long time windows of deviations from the norm we present all results in terms of average deviations.

In table 1 we present the values of Δd_{av} for $k = 5, 30, 100$ in all four scenarios described above. The values of Δd_{av} range from $\Delta d_{av} = 0.0334$ – an average deviation of 6.7% from the equal split to $\Delta d_{av} = 0.2192$ – an average deviation of more than 40%. In each scenario we have also tested whether differences in the Δd values between the cases with $k = 5$ and $k = 30$ as well as with $k = 30$ and $k = 100$ are statistically significant. A Wilcoxon signed rank test was used for this analysis. The values in bracket underneath the inequality signs between the different means in the table indicate the p-value⁹ by which the hypothesis of equal means can be rejected.

	$L = 50$			$L = 5$						
	$k = 5$	$k = 30$	$k = 100$	$k = 5$	$k = 30$	$k = 100$				
sim.	0.2192	>	0.0826	>	0.0334	0.0886	>	0.0770	>	0.0496
upd.		(1)		(1)			(0.9570)		(1)	
cont.	0.1561	<	0.1689	>	0.0830	0.1312	>	0.0725	>	0.0406
upd.		(0.6410)		(0.9928)			(1)		(1)	
	0.1561	<	0.0830							
		(0.9999)								

Table 1: Comparison of average absolute deviation of demands from the equal split (Δd_{av}) for increasing selection pressure.

The statistical results are very clear cut. In all but one case (increase from $k = 5$ to $k = 30$ for $L = 50$ and continuous updating) an increase in selection pressure reduces with statistical significance the average deviation from the equal split. This is quite in contrast to insights developed in optimization applications of evolutionary algorithms, where due to problems of premature convergence too strong selection pressure is considered counterproductive, and also in contrast to some

⁹If the exact p-value is larger than $1 - 10^{-4}$ we give a p-value of '1' in the tables.

work on the emergence of efficient norms in social systems (see e.g. March (1991)) where the merits of slow learning is stressed.

The size of the positive effect of strong selection seems to vary quite a bit between the four scenarios and is strongest for simultaneous updating with low frequency ($L = 50$). In this scenario comparing the process under small group imitation ($k = 5$) with imitate the best in the entire population reduces the average deviation from 40% to less than 7%!

As already discussed above, intuitively, the positive effect of strong selection seems to be due to two effects. First, fast convergence to uniform populations induces fast emergence of norms and prevents the slow invasion of fair obstinate populations by responsive strategies. Second, this is not a problem where extensive exploration of the strategy space is needed, but rather a problem of coordination. The mutation operator ensures some sampling of the strategy space and accordingly, the negative effect strong selection has with respect to an exhaustive coverage of the search space is not crucial here. This might change if significantly larger values of m were considered, but it seems to us that the exploration of the strategy space is indeed not the crucial issue when thinking about the evolution of norms in bargaining problems.

To close the discussion of the effect of selection pressure on the simulation outcomes let us briefly consider the impact on average payoffs. To keep the discussion short we focus here only on the comparison of profits for runs with $k = 5$ and $k = 100$ in the scenario with $L = 50$ and simultaneous updating. For $k = 5$ the time average of the mean population payoff averaged over the 20 runs reads $\bar{\pi}_{k=5} = 0.297$, whereas the corresponding value for $k = 100$ is $\bar{\pi}_{k=100} = 0.46$. The difference in average payoffs between the two cases is statistically significant with a p-value of virtually 1. Consistent with our observations about average deviation the average payoff under strong selection is significantly larger than under weak selection. Given that in this setup any bargaining game resulting in agreement yields an average payoff of 0.5 for the two partners, this observations follows directly from the fact that the time windows of disagreement following norm breakdowns are shorter under strong selection.

3.3 The Effect of Updating Frequency and Updating Regime

The results presented in table 1 indicate that the characteristics of the population evolution does not only depend significantly on the selection pressure but also on the frequency of updating of strategies and on the fact whether updating is simultaneous or continuous.

In this subsection we further explore this matter starting with the impact of the updating frequency. Intuitively, frequent updating implies that the choice of strategy of an agent relies on average payoffs based typically only on a small number of bargaining encounters. Hence there should be quite a strong stochastic component in the average payoffs of the individual agents. In particular if selection pressure is strong and updating is simultaneous, this might imply that the whole population quickly adopts the strategy of an agent who got lucky with respect to her partner in only a few bargaining games. On the other hand, frequent updating allows to quickly get away from very inefficient population profiles where almost all agents use immodest obstinate strategies and hence disagreement prevails.

This intuition is indeed confirmed by our simulation results. In table 2 we compare average deviations of demands from equal split for runs with $L = 5$ and $L = 50$. This comparison is carried out in four different scenarios:

I: Simultaneous updating, $k = 5$

II: Simultaneous updating, $k = 100$

III: Continuous updating, $k = 5$

IV: Continuous updating, $k = 100$

Again, the figures in the table give values averaged over 20 runs and the values in brackets the p-value of the corresponding Wilcoxon test.

	$k = 5$		$k = 100$		
	$L = 5$	$L = 50$	$L = 5$	$L = 50$	
sim. upd.	0.0886	< (1)	0.2192	> (0.9984)	0.0334
cont. upd.	0.1312	< (0.9087)	0.1561	< (0.9997)	0.0830

Table 2: Comparison of average absolute deviation of demands from the equal split (Δd_{av}) under $L = 5$ and $L = 50$.

Several interesting observations can be made. First, we observe that the effect of an increase of L on the average deviation from equal split does not only quantitatively differ between the four scenarios but also qualitatively. As anticipated from the discussion above, a large value of L decreases the average deviation of the demand from the equal split if there is strong selection and simultaneous updating. However, if either selection pressure is low or updating is continuous the deviation of the average demand from the equal split *increases* statistically significant as the average time window between updating periods is increased from 5 to 50. Actually, it is quite understandable that in these two settings the danger of frequent updating, namely that the whole population switches strategy due to a few lucky draws, is not so relevant. If updating is continuous and L is small agents decide based on small samples of bargaining payoffs for each strategy, but the crucial point is that not all of them decide based on the identical small samples. So even if $k = 100$ a lucky run for a strategy will only influence a certain subset of agents who happen to update during this lucky run. Also if selection pressure is small and updating is continuous an increase in L seems to result in an increase of average demand, but here the variance between runs is quite substantial and based on our 20 runs we have a significance level of only about 90% for this observation.

Let us now turn to the effect of the type of updating regime on average demands. Here we compare runs with simultaneous updating and continuous updating. Again, four scenarios are considered:

I: $k = 5, L = 50$

II: $k = 5, L = 5$

III: $k = 100, L = 50$

IV: $k = 100, L = 5$

The average deviation of demands from the equal split and the corresponding p-values are given in table 3.

We see that simultaneous updating might lead to smaller average deviations from the equal split, but that this effect is not present in all scenarios. If selection pressure is low and updating

	$L = 50$		$L = 5$			
	sim.	cont.	sim.	cont.		
$k = 5$	0.2192	> (0.9995)	0.1561	0.0886	< (0.9982)	0.1312
$k = 100$	0.0334	< (0.9999)	0.0830	0.0496	> (0.9928)	0.0406

Table 3: Comparison of average absolute deviation of demands from the equal split (Δd_{av}) under $L = 50$ and $L = 5$.

has low frequency ($L = 50$) simultaneous updating yields very substantial deviations from the equal split and continuous updating leads to better results in this respect, where the effect has high statistical significance. The same holds true if updating is frequent ($L = 5$) and selection pressure is strong, although in this case deviations are rather small under both updating regimes.

An interesting aspect of the four scenarios considered in table 3 is that they can be nicely ranked with respect to the amount of information an agents processes. If we consider the average number of observations an agent takes into account for his updating steps in a given time window, say 100 periods, we get 10 observations in scenario I, 100 in scenario II, 200 in scenario III and 2000 in scenario IV. Putting this together with our observations above, we can conclude that continuous updating is helpful for the emergence of the equal split norm if the amount of processed information is very low, for an intermediate range of the number of used observations simultaneous updating has positive effects, and if the amount of processed information is very large the type of updating regime seems to make only minor difference. Comparing the results across the four scenarios we observe that with continuous updating the average deviation from the equal split norm decreases as the average amount of information processed by each agent goes up. For simultaneous updating the same observation can be made if we only consider scenarios I, II and III or only I, II and IV. However, comparing III and IV, an increase in processed information leads to larger deviations from the equal split norm and accordingly to efficiency losses. Among all considered scenarios a setup with simultaneous updating with low frequency and strong selection pressure yields the best results with respect to the emergence of the efficient equal split norm. This is particularly notable because the number of observations processed by each agent under this scenario is by a factor 10 smaller than under scenario IV.

4 Conclusions

In this paper we have used an agent-based simulation model to study the evolution of bargaining behavior in a repeated standard bargaining problem. The existing evolutionary and adaptive learning literature has argued using quite different specifications of strategy space and updating mechanism that in this simple game the emergence of an equal split norm is a very robust phenomenon. Our simulations show that the consideration of two types of strategies – obstinate and responsive – as introduced and analyzed with respect to local stability in Ellingsen (1997) does often lead to persistent and substantial deviations of demands from the equal split rule. The equal split norm does emerge and typically persists for long periods of time but there are repeated episodes of substantial duration where, triggered by a complete breakdown of the norm, average demands substantially

exceed the fair split and accordingly there is a lot of disagreement.

We have then analyzed in some detail the effect different characteristics of the used evolutionary algorithm have on the average deviation of evolving behavior from the equal split norm. Our results show that in many cases this effect is very significant from a statistical point of view. Accordingly, we believe that it is important to have a sensible economic interpretation of the different parameter settings. In our setup the parameters governing the updating schedule and selection operator can be directly interpreted as the general propensity of agents to change their strategy and the number of observations about others' payoffs an agent can use when deciding about the new strategy. Accordingly, we believe that the observed sensitivity of results with respect to these parameter is rather a source for additional insights about the fine-structure of the emergence and breakdown of efficient bargaining norms than a 'flaw' of the algorithm.

In this paper the parameters k and L were taken as given characteristics of the agents. A natural extension of our analysis here would be to make these parameters endogenous choices of the agents¹⁰. Such an approach would also allow to introduce explicit costs of observing payoffs of other individuals and to study which kind of mix between information collection and updating strategy prevails (update frequently with only a few observations or update with low frequency but a lot of information at each step).

Finally, we have shown that it also makes a qualitative difference with respect to the simulation results whether updating is simultaneous or continuous. From a pure modelling point of view this shows that one should be cautious when using a standard evolutionary algorithm with simultaneous updating to represent processes where there is no good reason to assume that agents should update in a coordinated way. From a more economic point of view, one might argue that the economic environment has strong influence on whether updating is simultaneous or continuous (for example by the way the information flow about other payoffs is organized). Hence our results about the efficiency effects induced by a switch from continuous to simultaneous updating or vice versa can be seen as a contribution to mechanism design from a dynamic evolutionary perspective.

References

- Binmore, K., Samuelson, L. and P. Young (2003), 'Equilibrium Selection in Bargaining Models', *Games and Economic Behavior*, 45, 296-328.
- Borgers, T. and R. Sarin (1997), 'Learning through Reinforcement and Replicator Dynamics', *Journal of Economic Theory*, 77, 1-14.
- Carpenter, J.P. (2002), 'Evolutionary Models of Bargaining: Comparing Agent-based Computational and Analytical Approaches to Understanding Convention Evolution', *Computational Economics*, 19, 25-49.
- Dawid, H. (1999), 'On the Convergence of Genetic Learning in a Double Auction Market', *Journal of Economic Dynamics and Control*, 23, 1545 - 1567.
- Dawid, H. and MacLeod, W.B. (2001), 'Holdup and the Evolution of Bargaining Conventions', *European Journal of Economics and Social Systems*, 15, 3, 153-169.

¹⁰See e.g. Srinivas and Patnaik (1994) for an analysis of the endogenous determination of evolutionary algorithm parameters in the framework of optimization problems.

- Dawid, H. and MacLeod, W.B. (2004), 'Hold-up and the Evolution of Investment and Bargaining Norms', Working Paper, University of Bielefeld.
- Ellingsen, T. (1997), 'The Evolution of Bargaining Behavior', *Quarterly Journal of Economics*, 112, 581-602.
- Ellingsen, T. and J. Robles (2002), 'Does Evolution Solve the Hold-Up Problem?', *Games and Economic Behavior*, 39, 28-53.
- Gerding, E.H., van Bragt, D.D.B. and J.A. La Poutre (2000), 'Scientific Approaches and Techniques for Negotiation: A Game Theoretic and Artificial Intelligence Perspective', CWI Report SEN-R0005.
- Gerding, E.H., van Bragt, D.D.B. and J.A. La Poutre (2003), 'Multi-Issue Negotiation Processes by Evolutionary Simulation, Validation and Social Extension', *Computational Economics*, 22, 39-63.
- Hofbauer, J. and K. Sigmund (1998), *Evolutionary Games and Population Dynamics*, Cambridge University Press, Cambridge, UK.
- March, J.G. (1991), 'Exploration and Exploitation in Organizational Learning', *Organization Science*, 2, 71-87.
- Nash, J.F. (1950), 'The Bargaining Problem', *Econometrica*, 18, 155-162.
- Nash, J.F. (1953), 'Two-Person Cooperative Games', *Econometrica*, 21, 128-140.
- Nowak, M.A. and K. Sigmund (1992), 'Tit for tat in heterogenous populations', *Nature*, 355, 250-253.
- Nowak, M.A. and K. Sigmund (1995), 'Invasion Dynamics of the Finitely Repeated Prisoner's Dilemma', *Games and Economic Behavior*, 11, 364-390.
- Oliver, J.R. (1996), 'A machine learning approach to automated negotiation and prospects for electronic commerce' *Journal of Management Information Systems*, 13, 83-112.
- Rubinstein, A. (1982), 'Perfect Equilibrium in a Bargaining Game', *Econometrica*, 50, 97-109.
- Saez-Marti, M. and J.W. Weibull (1999), 'Clever Agents in Young's Evolutionary Bargaining Model', *Journal of Economic Theory*, 86, 268-279.
- Srinivas, M. and L.M. Patnaik (1994), 'Adaptive Probabilities of Crossover and Mutation in Genetic Algorithms', *IEEE Transactions on Systems, Man and Cybernetics*, 24, 656-667.
- Tröger, T. (2002), 'Why sunk Costs Matter for Bargaining Outcomes: An Evolutionary Approach', *Journal of Evolutionary Theory*, 102, 375-402.
- Weibull, J. (1995), *Evolutionary Game Theory*, MIT Press, Cambridge.
- Young, P. (1993), 'An Evolutionary Model of Bargaining', *Journal of Economic Theory*, 59, 145-168.





Figure 1: The evolution of the fraction s_5 of agents using the obstinate strategy demanding $d = 0.5$ for (a) proportional selection and (b) tournament selection with $k = 100$.

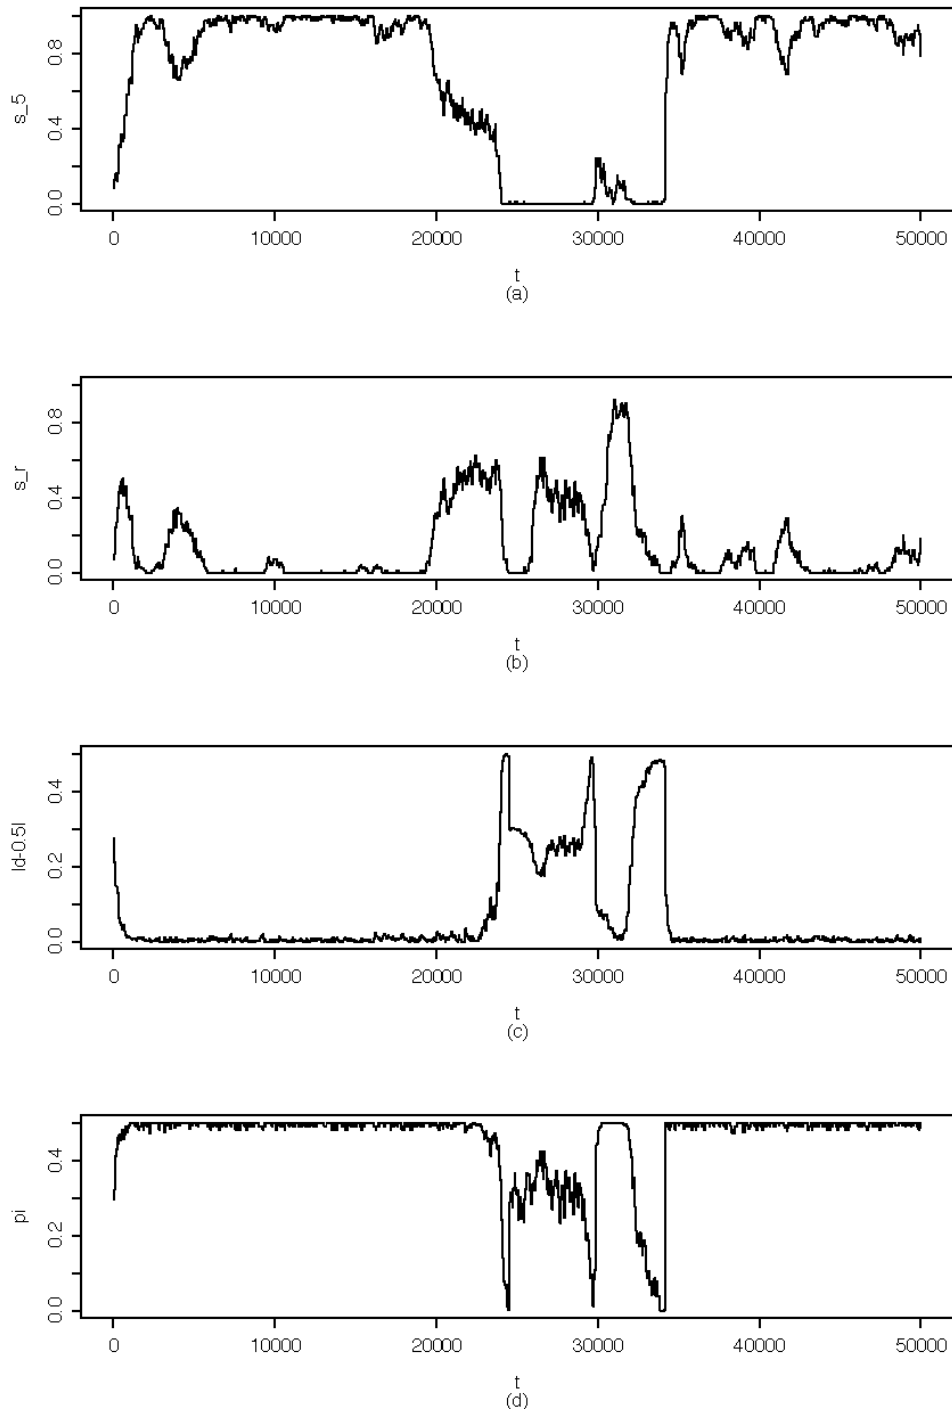


Figure 2: The evolution of (a) the fraction s_5 of agents using the obstinate strategy demanding $d = 0.5$, (b) the fraction s_r of agents using the responsive strategy, (c) the average demands and (d) the average payoff of all agents for a run with proportional selection.

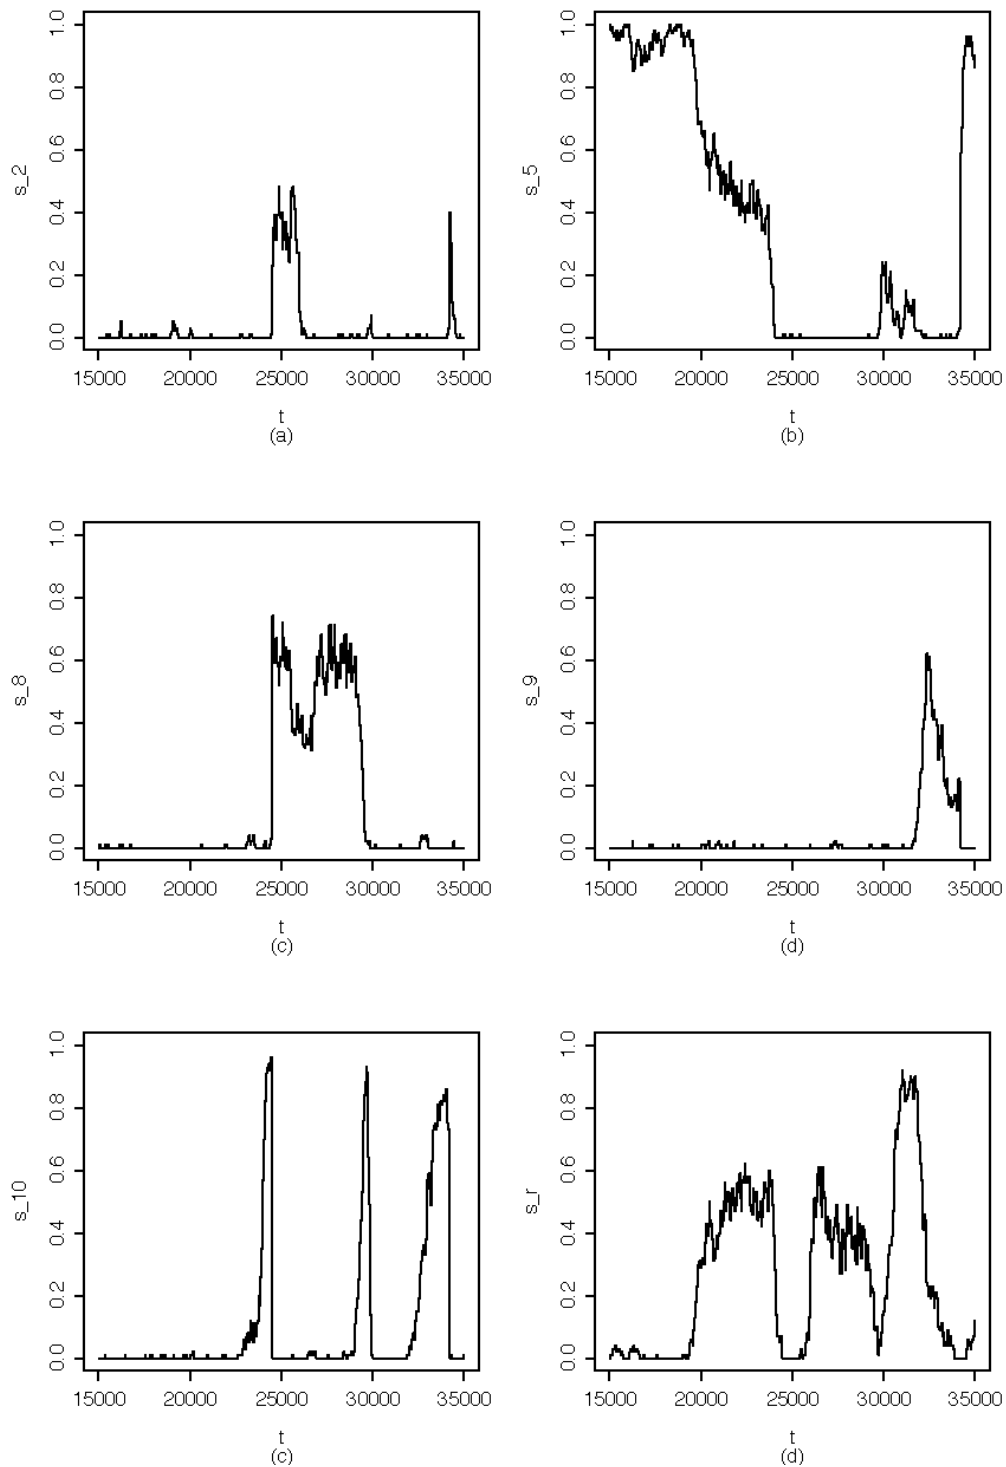


Figure 3: The evolution of the frequencies (a) s_2 , (b) s_5 , (c) s_8 , (d) s_9 , (e) s_{10} , (f) s_r for the run shown in figure 2 and $t \in [15000, 35000]$.

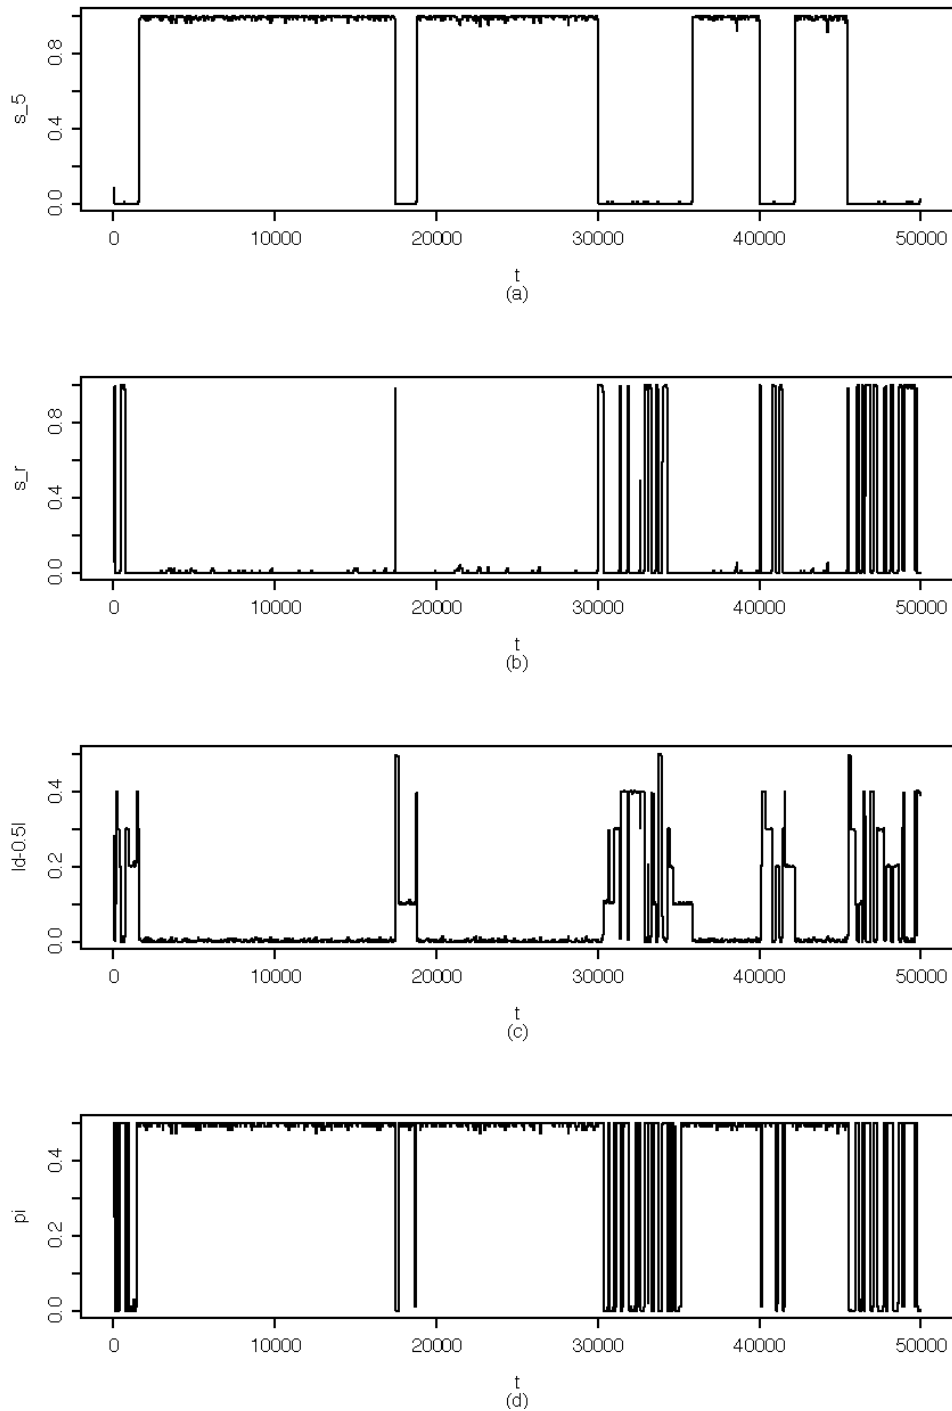


Figure 4: The evolution of (a) the fraction s_5 of agents using the obstinate strategy demanding $d = 0.5$, (b) the fraction s_r of agents using the responsive strategy (c) the average demands and (d) average payoffs of all agents for a run with tournament selection and $k = 100$.