# Chapter 9 Analytic Explorations and Evolutionary Processes

Studying strategies for social interactions suggests a theoretical analysis, including a classical game-theoretical analysis as well as evolutionary simulations. Such a theoretical analysis may indicate the superiority of particular strategies about others. If a strategy outperforms others, it is psychologically plausible that individuals actually select it, promoted by either severe deliberation and/or learning processes. Especially if a strategy that outperforms others is simple to apply, it is reasonable to assume that this strategy is actually applied by individuals.

Therefore, the central goal of *study 3* reported in the chapter is to illustrate how strategies compete against one another and to show the advantages of particular strategies compared to others. Of special interest are those strategies which lead to efficient outcomes, and thereby, can explain trust and reciprocity. Unfortunately, this evaluation of different strategies will illustrate that the well-established Nash equilibrium concept is not sufficient to indicate the superiority of certain strategies about others. On the contrary, it will get clear that many strategy combinations are classified as equilibria, and a superiority of one single or at least a few strategies cannot be assessed. Therefore, another reasonable criterion for evaluating strategies is introduced, namely the evolutionary stability of strategies. An evolutionary stable strategy is a "strategy such that if all the members of a population adopt it, then no mutant strategy could invade the population under the influence of natural selection" (Maynard Smith, 1982, p. 10). Both concepts, namely the Nash equilibrium concept and the evolutionary stability concept, are applied for evaluating a selected set of strategies, and thereby, illustrate the advantages of particular features of the strategies.

Whereas in the first section of this chapter the analytic exploration is presented, in the second section the results of the simulations of evolutionary processes are reported. With these simulations, it is possible to evaluate a large set of strategies. In addition, the goal is to show how the potential complexity of strategies influences the outcome of an evolutionary process. The second goal is to illustrate how small errors in the application of strategies can influence an evolutionary outcome. Finally, it is analyzed how a potentially produced surplus is distributed between both players and if efficient outcomes are obtained.

*9.1 Folk Theorem for Indefinitely Repeated Games*

First, I will define the indefinitely repeated investment game, as is used for the following theoretical analysis and simulations: The game consists of two players, A and B. After each period, a new period follows with a continuation probability of $d$, so that the probability that the game will last for exactly $t$ periods is $d^{t-1}(1-d)$. The expected number of periods for one game is therefore $\sum_{t=0}^{\infty} t d^{t-1}(1-d)$, which is a geometrical series that converges to $1/(1-d)$, which gives an expected number of t=100 periods for $d$=0.99. In every period, both players earn an endowment of 10. Player A decides whether to invest any integer percentage of the endowment. If an investment is made, it is multiplied by 3 and sent to player B who decides whether to return any integer percentage of the trebled investment to player A. If one wishes to evaluate the performance of a particular strategy for the entire repeated game against another strategy, the problem arises that the payoff depends on the length of the game, which is determined randomly. There is, therefore, an established technique to evaluate the payoff for one strategy by using the expected average payoff a strategy obtains against another strategy (Binmore, 1992, pp. 360-366).

The game-theoretical equilibrium prediction for the investment game can be drawn from the folk theorem[8] for indefinitely repeated games (Fudenberg & Tirole, 1991). It restricts all outcomes to those that are "individually rational." All possible payoff combinations for the investment game are shown in Figure 9. It is obvious that, for player A, it is individually not rational to agree on any average payoff that is below the endowment of 10. For player B, it is individually not rational to agree on any average payoff that is below the endowment plus 1% of the trebled investment since it is more profitable for player B to retain the entire trebled investment in any period than to retain less than 1% in all periods (given a constant investment rate). These are the only two restrictions made by the folk theorem, yielding a large number of equilibria (see Figure 9).

[8] The theorem is called "folk" theorem, even though it is well-known nobody knows to whom it should be attributed (Binmore, 1992).

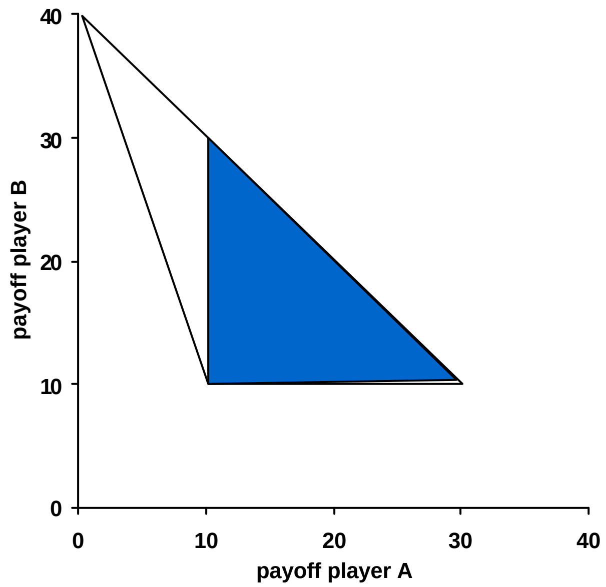


*Figure 9.* Payoff region and prediction of the folk theorem for indefinitely repeated games.
The large triangle indicates all possible payoff combinations (payoff region). The hatched triangle marks the payoff combinations for the investment game with a continuation probability of 0.99, predicted by the folk theorem for indefinitely repeated games. The diagonals between the two coordinates (30, 10) and (0, 40) represent efficient outcomes, that maximize mutual payoffs.

This prediction is unsatisfying for two reasons. First, the number of equilibria is very large, and which of the many equilibria people finally select is left open. Additionally, the prediction ignores with what kind of decision strategies the equilibria are reached. In order to explain how the different outcomes are finally reached, one has to study the repeated game strategies that define decisions for the game. These strategies can be rather complex as they can define a decision for any possible eventuality. A candidate group of strategies for describing repeated game strategies are finite state automata, which are precise strategies that can vary in their complexity and can result in being rather simple. By studying these automata and eliciting the type of strategies that perform well in competition against each other, one may gain some insight on the decision strategies people implement for the game and one obtains models that can be tested experimentally. I will first analyze strategies by the Nash equilibrium concept.

Restricting the set of possible strategies to finite state automata has been proposed by many authors (Abreu & Ariel, 1988; Aumann, 1981; Binmore & Samuelson, 1992). An automata consists of various states and each of the states determines an output. Given some input, the automata possibly moves to a different state. The state in which the automata commences is the initial state. For applying automata to the investment game, the automata's output is defined as the investment rates or return rates. The input is defined as the opponent's last decision. The number of states can be utilized as an indicator of the automata's complexity. Given the sequential game structure, automata representing player B determines their decision subsequent to player A, therefore, they could have several initial states depending on player A's first decision. Note, that if an automaton for player A invests nothing, then the automaton for player B cannot return anything. In this case, the automaton for player B may move to another state, but no decision can be made, and therefore, the automata for player A receives no input from the automata for player B. Hence, for any state with an investment of 0%, the automaton for player A receives no input from the other automaton for player B, so that there can only be one possible move to another state.

*9.2 Equilibria for the Investment Game*

What kind of Nash equilibrium exists if the strategies for the repeated investment game are restricted to finite automata? For investigating how the equilibria depend on the complexity of automata, first automata of the lowest complexity are considered and subsequently automata's complexity is increased.

If automata are restricted to one state, the automata can only differ in their output. If the output is restricted to integer percentages ranging from 0% to 100%, this gives 101 different automata for both players. For all player A's automata, with an investment greater than 0%, the best reply for player B is to return nothing. If the automata for player A make no investment, the best reply is any of player B's 101 automata as they all yield a payoff of 10. For all automata for player B, with return rates of at least 34%, the best reply for player A is the automaton with an investment of 100%. If player B returns less than 34%, the best reply automaton for player A is the one with no investment. Hence, 34 equilibria result, all consisting of player A's automaton that makes no investment and the 34 player B automata with a return lower than 34%. This shows that if the strategies for the investment game are restricted to automata with lowest complexity, that is one single state, it implies that the opponent's decision cannot influence the behavior of the automata which leads to the least efficient outcome of a payoff of 10 for both players. From this, it can be concluded that a minimum complexity of strategies does not enable efficient outcomes which explain trust and reciprocity.

How does this result change if automata's complexity is increased to two states? In general, if automata have more than one state, it has to be specified in which state the automaton is. Every automata commences in the initial state. From there, the input of the automaton, which is the opponent's decision of any integer percentage ranging from 0 to 100, determines to which state the automaton moves. Following the intention to propose automata as models for people's decision strategies, one has to be concerned about the psychological plausibility of automata's characteristics. It appears to be psychologically plausible that opponent's decisions are categorized in groups so that, for instance, no difference is made between a return rate of 51% or 52%. One could proceed even further and argue that people may apply an aspiration level to classify their opponent's decisions. The assumption of an aspiration level has often been made in decision-making research. Simon proposes an aspiration level for the "satisficing" rule, in which decision makers take the shortcut of setting an aspiration level when they search for options, hence, they search as long as they find an option that exceeds the aspiration level (Simon, 1956, 1990). Lopes (1996) suggested that individuals have an aspiration level when they make choices between risky options (lotteries) such that they refuse lotteries with possible outcomes which do not reach their aspiration levels. Tversky (1972) proposes that people use aspiration levels for each attribute when comparing multi-attribute options. Options that do not exceed the aspiration level for one attribute are eliminated from further consideration.

Concluding from this research, it is probable that people also apply an aspiration level for categorizing their opponents' behavior for the investment game. Individuals in the role of player A (player B) may apply an aspiration level with which they categorize the returns (investments) of their opponents. Following from this, each automaton implements an aspiration level with which the opponents' decisions are dichotomized. The aspiration levels can range from 0 to 100, so that an aspiration level of, for instance, 40% categorizes all returns lower than 40% as "Exploitation" and any higher or equal return as "Reciprocity." Likewise, an aspiration level for player B's strategies categorizes low investments as "Distrust" and high or equal investments as "Trust." The labels are for simplification: If, for instance, the return only slightly exceeds the aspiration level, the return is classified as "Reciprocity", although the individual may still regard the return as rather exploiting. However, the behavioral reactions to the small return may be similar to higher returns, justifying the same label. The dichotomized decisions are implemented as input for the automata and determine to which state the automata moves. Given that already in the first period, player B's decision can depend on player A's investment, player B's automata can have two possible initial states in which they commence subsequent to a low or high investment in the first period.

With this notation, strategies for the investment game can be easily represented. Figure 10 portrays a selected set of strategies for player A. Figure 10 also shows two Grim strategies, whic h commence in the initial state with an investment of 100%. If the opponent returns a substantial amount (higher or equal to the aspiration level), Grim remains in the initial state and repeats the investment of 100% in the next period. However, if in any period player B returns less than the Grim's aspiration level, the strategy moves to state 2 in which no investment is made for all subsequent periods.

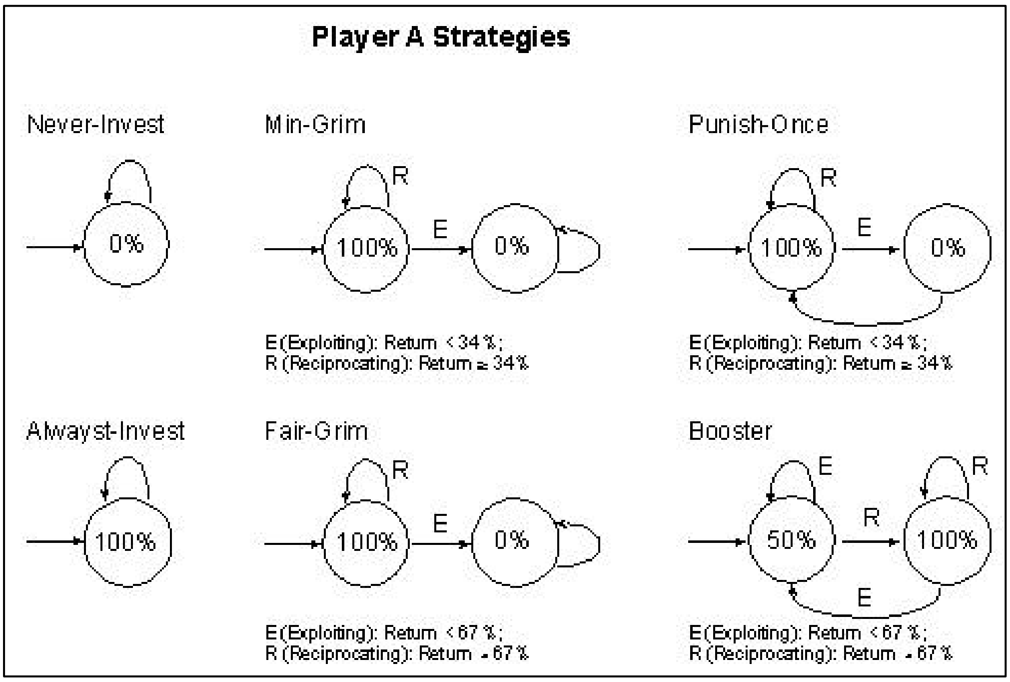


*Figure 10.* Selected set of strategies for player A.

In the case where an automata has a state in wh ich no investment is made (e.g. Min-Grim), only one unconditional move is possible for this state. The same holds for one-state automata.

The possible number of automata, even restricted to a maximum number of two states, is very large. Two-state automata can vary by 100 possible different aspiration levels, and by 101 different outputs of each state, yielding already 20,200 possible combinations. Furthermore, from each of the $k$ states an automata can move to $k$ states, either subsequent to a high or low return, yielding another $k^{(2k)}$ possibilities, which altogether produces 16,321,600 possible two-state automata. Given this large number of strategies the further analytical exploration is restricted to a very small set strategies.

For selecting a restricted set of strategies, first possible automata for player A are considered. Automata for player A have basically two crucial investment rates: Either the interaction with player B leads, in the long-term, to a payoff greater than the endowment so that an investment of 100% maximizes the payoff, or the interaction leads to a payoff lower than the endowment, in which case no investment maximizes the payoff. Hence, the automata can be restricted to these two investment rates. An aspiration level, for player A, of 34% guarantees a profitable relationship, whereas a level of 67% would lead to equal final payoffs for both players. With these restricted possibilities of investments and aspiration levels the total number of possible automata, for player A, is quite reduced. For instance, there are only two remaining one-state automata, which are used for the further analysis. The *Never-Invest* automaton never makes an investment and the *Always-Invest* automaton always makes an investment of 100%. In addition, four two-state automata, responsive to the behavior of the opponent, are added to the set. I select two Grim strategies; *Min-Grim*, a Grim strategy with an aspiration level of 34%, and *Fair-Grim*, a Grim strategy with an aspiration level of 67%. These Grim strategies commence with an investment of 100% and repeat this investment in all following periods unless they become exploited, hence, they move to state 2 in which they make no investment in all following periods. Since the Grim strategies react very sensitively to exploitation, as they do not return to any investment after they are exploited once, two other strategies, which are more tolerant to exploitation, are added. *Punish-Once* commences with an investment of 100% and repeats this investment in all following periods, unless a low return is made. In this case, it moves to state 2 with no investment. In contrast to the Grim strategy, in the next period Punish-Once returns to state 1. Punish-Once applies an aspiration level of 34%. *Booster* commences with an investment of 50% and increases the investment in the second state to 100%, if the other player makes a substantial return. If player B repeats high returns, Booster remains in the second state. If player B makes low returns Booster remains in or moves to the first state.

For player B, three representative return rates are selected: 0%, 34%, and 67%. If player B wishes to exploit player A, the most profitable exploitation is a return rate of 0%. If player B wishes to keep the relationship profitable for player A, a return rate of 34% should be selected and a return rate of 67% yields equal final payoffs for both players. The number of reasonable automata, for player B, can be quite restricted (see Figure 11). Three one-state automata for player B, that differ only with respect to their return rate, already provide a good set of strategies for player B. These three automata are: *No-Return, Min-Return*, and *Fair-Return* with a return rate of 0%, 34%, and 67% respectively. Additionally, the two-state automata *Min-Reactive* and *Fair-Reactive*, which are responsive to player A's decision, are selected. Both strategies move to state 1 with no return if an investment below the aspiration level is made (the aspiration level was set to 90% to investigate whether it promotes high investments). Subsequent to an investment of at least the aspiration level, the strategies move to state 2. In state 2 Min-Reactive makes a return of 34%, whereas Fair-Reactive makes a return of 67%.

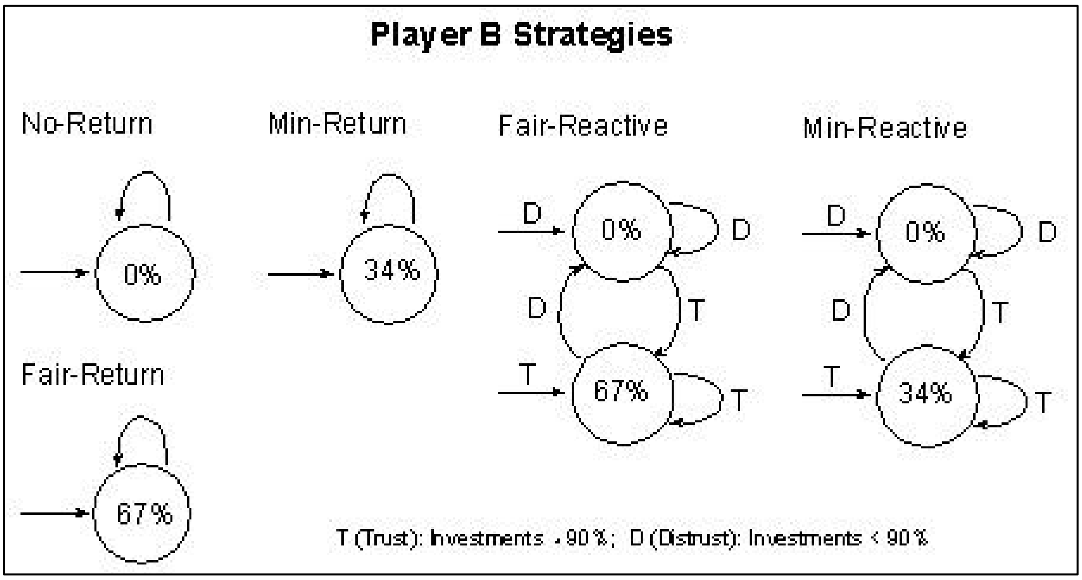


*Figure 11.* Selected set of strategies for player B.
In the case of one-state automata only one unconditional move is possible.

If one restricts the possible strategies for the investment game to the selected set of automata, the strategies and their expected average payoffs against each other can be represented in a 5x5 matrix (see Table 1).

The selected set of strategies includes several Nash equilibria as indicated in Table 1. The first equilibrium is a combination of Never-Invest and No-Return. Fair-Grim yields an equilibrium with Fair-Return and Fair-Reactive. Min-Grim and Punish-Once each form equilibria with Min-Return and Min-Reactive. Always-Invest forms no equilibrium. The equilibria demonstrate that, already with automata of low complexity, a large number of payoff combinations, predicted by the folk theorem, can be obtained and demonstrates with which types of strategies they can be realized.

Table 1

*Payoff matrix for the indefinitely repeated investment game with a restricted set of strategies for both player roles*

| Player A strategies | Player B strategies | | | | |
|---|---|---|---|---|---|
| | No-Return | Min-Return | Fair-Return | Min-Reactive | Fair-Reactive |
| Never-Invest | 10.00*, 10.00* | 10.00, 10.00* | 10.00, 10.00* | 10.00, 10.00* | 10.00, 10.00* |
| Fair-Grim | 9.90, 10.30 | 10.01, 10.20 | 20.10*, 19.90* | 10.01, 10.20 | 20.10*, 19.90* |
| Min-Grim | 9.90, 10.30 | 10.200*,29.80* | 20.10*, 19.90 | 10.20*, 29.80* | 20.10*, 19.90 |
| Punish-Once | 5.00, 25.00 | 10.200*,29.80* | 20.10*, 19.90 | 10.20*, 29.80* | 20.10*, 19.90 |
| Booster | 5.00, 25.00 | 10.199, 29.70* | 20.05, 19.85 | 5.00, 25.00 | 5.00, 25.00 |
| Always-Invest | 0.00, 40.00* | 10.200*, 29.80 | 20.10*, 19.90 | 10.20*, 29.80 | 20.10*, 19.90 |

*Note.* The cells of the matrix show the expected average payoffs of the strategies for the indefinitely repeated investment game, with a continuation probability of 0.99. The first value represents player A's payoff and the second value player B's payoff. The asterisks indicate the best reply strategies for both players. The cells with two asterisks indicate the seven Nash equilibria. For a description of the strategies see Figure 10 and Figure 11.

The best reply strategy for player A, for any player B strategy with a return rate lower than 34%, is the Never-Invest strategy. The best reply for player A, for any player B strategy, with a return rate of at least 34%, is to select, for instance, any Grim strategy which has, at maximum, an aspiration level that equals the return rate (the Punish-Once strategy is only optimal for return rates lower than 50%). This analysis demonstrates that simple automata of player B, consisting of only one state, can lead player A to choose strategies that enable player B to obtain payoffs ranging from 10 to 40 (including player B's endowment).

If player A selects one of the Grim strategies, then the best reply strategies for player B depend on the aspiration level of the Grim strategies: If, for instance, Grim has an aspiration level of 67%, then the best reply for player B is to always return 67%. Of course, this could also be seen vice versa, given any constant return rate by player B of at least 34%, the best reply for player A is a Grim strategy with the corresponding aspiration level.

The Always-Invest strategy forms no equilibrium since player B's best reply is to select the No-Return strategy, which leads to player A's worst outcome. This demonstrates that, for player A, it is important to incorporate a "punishment mechanism" in the strategy, which leads player A to "abandon" the exchange relationship with player B if player B makes a low return. This implies that player A requires at least two states for an automaton to obtain payoffs above the endowment.

In sum, all equilibria, that lead to efficient outcomes, consist of strategies for player A with at least two states and of strategies for player B with one *or* two states. For player B, a second state is not necessary. On the contrary, the strategies for player A require a higher complexity to incorporate a punishment mechanism which prevents player B from exploiting player A, and thereby, allows player A to trust player B and to produce efficient outcomes.

*9.3 Evolutionary Stability*

The previous sections have demonstrated the large number of possible equilibria and strategies that form these equilibria for the indefinitely repeated investment game. The result is unsatisfying as it undetermined which of the many strategies forming equilibria people may apply. The evaluation of the strategies, by the Nash equilibrium concepts, considers whether a strategy is a best reply given another strategy and mutual best replies form equilibria. What other criteria are reasonable to evaluate strategies?

The concept of evolutionary stability has often been suggested for evaluating strategies (see Maynard Smith, 1984; Maynard Smith & Price, 1973; Samuelson, 1997; Weibull, 1995). The Nash equilibrium concept has the disadvantage in that for a given combination of strategies alternative best reply strategies may exist, which are strategies leading to the same payoff, so that a player may convert to these strategies. If a player converts to an alternative best reply strategy, the opponent may, as a consequence, also convert to yet another best reply strategy, so that finally, the original strategies are no longer applied and a new equilibria is reached. Therefore, one may argue that for evaluating strategies the ability to form an equilibrium, which is stable, is an additional important aspect of strategies. However, the evolutionary stability concept can only be applied to a symmetrical game in which the strategies for the different players are interchangeable. One solution to this problem is to transform the asymmetrical investment game into a symmetrical game (see Samuelson, 1997). It is helpful to assume a situation in which a population of agents select between different automata to play the game on their behalf. Agents play the indefinitely repeated investment game repeatedly and during this process different strategies evolve. Strategies that produce a low payoff in previous games are less likely to be selected again by the agents, whereas well performing strategies will be selected more frequently. Agents, thereby, are the driving force of the evolutionary process, which produces viable strategies.

In order to transform the asymmetrical investment game into a symmetric game G, it is assumed that player's strategies in the symmetrical game G are a combination of a strategy for player A and a strategy for player B of the indefinitely repeated investment game. Each player in the asymmetrical investment game has a finite set of strategies ($S_A$ for player A and $S_B$ for player B). The payoff of a strategy for a player is defined as the expected average payoff of the strategy given the other player's strategy. The asymmetrical game is then defined by the set of players, the set of possible strategies, and the set of possible payoffs.[9]

It is assumed, that prior to each game the player roles in the investment game are randomly assigned to the players in the symmetrical game G. Subsequent to the role assignment, the players in the symmetrical game apply the corresponding strategy from their combination of strategies to the investment game. A strategy $s$ for the symmetrical game G consists of a strategy $s_A$ for player A and a strategy $s_B$ for player B in the investment game. Let $S$ be the set of strategies, which are combinations of two strategies, for the symmetrical game G. The payoff $p_i$ for a strategy $s_i$ of player $i$, for the symmetrical game G, is defined as the expected average payoff for the player A strategy and for the player B strategy. Let $p_i(s_i, s_j)$ be the expected average payoff for player $i$ if player $i$ selects the strategy $s_i$ and player $j$ selects strategy $s_j$ of the game G. The symmetrical game G is then defined by the set of players, the set of strategies $S$, and the set of possible payoffs.

With the symmetrical game, it is possible to apply the evolutionary stability concept for evaluating strategies. Definition 1 provides the two conditions, which have to be fulfilled for a strategy to be an evolutionary stable strategy (ESS). The first condition of the definition corresponds to the Nash equilibrium criterion. It ascertains that a strategy $\sigma^*$ is an evolutionary stable strategy if it is a best reply against itself, so that no alternative strategy $s'$ leads to a higher payoff against $s^*$. Condition 1.2 is the crucial component of the evolutionary stability concept as it constitutes the stability criterion: If $s^*$ leads to the

[9] For a better illustration of the elaboration, the definitions are only provided for pure strategies (i.e. strategies which are selected by the players with a probability of 1.0). However, the definition also holds for mixed strategies. A mixed strategy is a probability distribution over the set of pure strategies, that defines the probability with which each of the pure strategies are played by the players.

same payoff against itself, as an alternative strategy $s'$ does against $s^*$, then $s^*$ must lead to a greater payoff against $s'$ than $s'$ leads against itself.

Definition 1. Evolutionary stable strategy (ESS)

A strategy $s^* \in \Sigma$ is an evolutionary stable strategy if

$$p(s^*, s^*) \geq p(s', s^*) \quad \forall s' \in \Sigma \quad \text{and if} \qquad (1.1)$$

$$p(s^*, s^*) = p(s', s^*) \Rightarrow p(s^*, s') > p(s', s') \quad \forall s' \in \Sigma, \; s' \neq s^*. \qquad (1.2)$$

In contrast to the Nash equilibrium concept, which is weak for indefinitely repeated games, the evolutionary stability concept is very strong. If the above presented sets of selected strategies are combined with each other, this yields 30 combined strategies. These strategies can then be applied for the symmetrical game and their payoffs can be determined. It results that none of these strategies are evolutionary stable, since the second stability criterion cannot be fulfilled. An alternative strategy always exists, which leads to an equal payoff against the evaluated strategy, and the evaluated strategy does not lead to a greater payoff against the alternative strategy, than the alternative strategy leads against itself.

For instance, a population of agents applying a strategy combination of Min-Grim and Min-Return is evolutionary not stable, since it can be invaded by a strategy combination of Always-Invest and Min-Return. The strategy combination Always-Invest and Min-Return leads to the same payoff against the combination Min-Grim and Min-Return than the combination Min-Grim and Min-Return does against itself, so that the stability criterion 1.2 remains to be proven. Against stability criterion 1.2, the strategy combination Min-Grim and Min-Return leads to the same payoff against the combination Always-Invest and Min-Return as the combination Always-Invest and Min-Return leads against itself, so it can be concluded that Min-Grim and Min-Return is evolutionary not stable. Analyses of the indefinitely repeated prisoner's dilemma have shown that no evolutionary stable strategy for the game exists (Boyd & Lorberbaum, 1987; Lorberbaum, 1994), indicating that the evolutionary stability concept is, in general, to strong for distinguishing strategies for an indefinitely repeated game.

Another criterion for evaluating strategies takes into account that agents may make small errors when selecting or applying a strategy. These errors appear realistic given that human decision-making can be characterized by small unsystematic errors. The concept of Limit evolutionary stable strategies (Limit ESS) incorporates the possibility of errors

(Leimar, 1997; Samuelson, 1991; Selten, 1983, 1988). The Limit ESS concept is stronger than the Nash equilibrium concept, but weaker than the ESS concept. An equilibrium strategy has often several best reply strategies, that lead to the same payoff against the equilibrium, which often leads to disqualifying an equilibrium strategy as an ESS. These ties in payoffs can often be broken if the plausible assumption of small errors is made. However, there are two interpretations on how errors could occur which shall be investigated. First "selection errors" can occur when agents select their strategy or "execution errors" when they execute a particular strategy.

The first type of "selection errors" implies that agents apply their strategies correctly, but, from time to time with low probability, erroneously select a strategy they did not intend to apply. Samuelson has elaborated this view (Samuelson, 1991, p. 122) and was able to indicate that Limit ESS form symmetric Nash equilibria, which are characterized by three conditions (see Definition 2). First, there is no alternative strategy that produces a higher payoff against the Limit ESS than the Limit ESS itself (this is the symmetric Nash equilibrium condition, see 2.1). Second, the Limit ESS is not weakly dominated[10] by alternative strategies (weak dominance condition, see 2.2). Third, the Limit ESS is not composed of role equivalent strategies for one of the player roles, that is, strategies that produce identical payoffs to the Limit ESS in the same player role (equivalence condition, see 2.3).

Definition 2. Limit evolutionary stable strategy (Limit ESS)

A strategy $s^* \in \Sigma$ is a Limit evolutionary stable strategy if

$$p(s^*, s^*) \geq p(s', s^*) \quad \forall s' \in \Sigma \qquad (2.1)$$

and if there exists no strategy $s' \in \Sigma$ such that

$$p(s', s) \geq p(s^*, s) \quad \forall s \in \Sigma \qquad (2.2)$$

were the inequality is strict for at least one σ.

In addition, let $s^*$ be composed of two strategies $s_A^*$ and $s_B^*$ for the player roles A and B, then there is no strategy $s_i'$ for $i, j \in \{A, B\}$ and $i \neq j$

such that $p_i(s_i', m_j) = p_i(s_i^*, m_j) \;\; \forall m_j \in S_j, \;\; s_i' \neq s_i^*$. (2.3)

[10] A strategy strictly dominates another strategy if the strictly dominating strategy leads to *higher* payoff against all strategies of the opponent than the other strategy does. A strategy weakly dominates another strategy if the weakly dominating strategy leads to *higher or equal* payoff against all strategies of the opponent than the other strategy does.

Definition 2 implies that the stability criterion 1.2 of Definition 1 could be violated, such that an alternative strategy leads to the same payoff against itself than the Limit ESS against the alternative strategy.

Which strategies, of the 30 strategy combinations, of the selected set of strategies are Limit ESS when selection errors occur? The symmetrical game has 49 Nash equilibria, but only seven of these are symmetrical and, thereby, qualify for a test by Limit ESS. All the seven symmetrical equilibria correspond to combinations of the seven Nash equilibria indicated in Table 1. For instance, the Nash equilibrium of the strategies Min-Grim and Min-Return for the asymmetrical game form a symmetrical equilibrium as a combined strategy in the symmetrical game. From the 7 symmetrical equilibria only two consist of Limit ESSs.

First, the strategy combination Never-Invest and No-Return is a Limit ESS. Although some strategies reach an equal payoff against Never-Invest and No-Return, no strategy combination reaches a higher or equal payoff against all other strategies compared to Never-Invest and No-Return. Second, the combination of Min-Grim and Min-Return can also be classified as a Limit ESS. No other strategy obtains a higher payoff against this Limit ESS and no alternative strategies weakly dominate the Limit ESS, that is, no alternative strategy reaches a higher or equal payoff against all other strategies, compared to the Limit ESS. Other strategies do not classify for a Limit ESS. For instance, the Fair-Grim and Fair-Return is weakly dominated by Min-Grim and Fair-Return, since Min-Grim reaches a higher payoff against Min-Return than Fair-Grim.

The other types of errors that can occur in a game are “execution errors.” Following this interpretation, agents do not make errors in selecting the intended strategy, but do make errors when they execute their selected strategy. These errors imply that with low probability agents sometimes deviate from the prescribed decisions of their selected strategy. For instance, an agent applying the Grim strategy may already make an error in the first period by making no investment, instead of the prescribed 100% investment. This interpretation of execution errors is consistent with Selten’s approach of a “trembling hand” (Selten, 1983, 1988). Errors (which Selten calls “trembles”) ensure that, even with very low probabilities, all possible combinations of decisions (all branches in the extensive representation of the game) are reached with low probability.[11] A game in which errors

[11] Selten’s (1983) approach has been formulated for finite games, whereas I use it for indefinitely repeated games, which causes no problems since the expected average payoff for one strategy against another is well defined as the continuation probability is smaller than 1.

occur with low probability is called a “perturbed game” (Selten, 1983). In a perturbed game every possible choice is made with minimum probability. A strategy $s*$ is a Limit ESS for a game in which no errors occur, if it is an evolutionary stable strategy (ESS) for the perturbed game (Selten, 1983, p. 304). For investigating whether the set of selected strategies consists of Limit ESS if execution errors occur, a Monte Carlo study was conducted. In this Monte Carlo study the selected set of strategies played the indefinitely repeated investment game repeatedly against each other (each strategy played two million games against all other strategies of the other player role). In each period of each game an error occurred with a probability of 0.01. If an error occurred, the strategies execute a random decision, such that a return or investment ranging from 0% to 100% was chosen with equal probability. Table 2 shows the average payoff the selected strategies obtained against each other.

Table 2

*Payoff matrix for the perturbed indefinitely repeated investment game - results of a Monte Carlo study*

| Player A strategies | Player B strategies | | | | |
|---|---|---|---|---|---|
| | No-Return | Min-Return | Fair-Return | Min-Reactive | Fair-Reactive |
| Never-Invest | 9.95*, 10.15* | 10.00, 10.10 | 10.05, 10.05 | 9.96, 10.14 | 9.97, 10.13 |
| Fair-Grim | 9.18, 12.47 | 10.02, 11.64 | 18.63, 18.55* | 9.98, 11.68 | 17.32, 17.41 |
| Min-Grim | 9.18, 12.49 | 10.228, 28.16* | 19.24, 19.15 | 10.15, 25.36 | 17.72, 17.80 |
| Punish-Once | 4.83, 25.66 | 10.246, 29.59* | 19.97, 19.87 | 10.20, 29.48 | 19.81, 19.87 |
| Booster | 5.05, 25.02 | 10.237, 28.86* | 19.59, 19.50 | 5.68, 25.47 | 6.83, 24.33 |
| Always-Invest | 0.20, 39.70* | 10.247*, 29.65 | 20.00*, 19.90 | 10.21*, 29.69 | 19.92*, 19.98 |

*Note.* The cells of the matrix show the average payoff the strategies obtained against each other. The strategies played 2 million indefinitely repeated investment games against each other with a continuation probability of 0.99. In each of these games, in each period, both strategies made an error with a probability of 0.01, so that on average a strategy made one error in one repeated game. For each repeated game the average payoff was determined and again for all 2 million games the average was taken.

Table 2 shows that only one Nash equilibrium could be obtained for the perturbed indefinitely repeated investment game. This equilibrium consists of the Never-Invest and the No-Return strategy. From Table 2, a symmetrical game and the strategies payoffs can be constructed by combining all strategies for player A with the strategies for player B. This symmetrical game has only one symmetrical Nash equilibria formed by the strategy combination Never-Invest and No-Return. No alternative strategy reaches a higher or equal payoff against Never-Invest and No-Return, therefore, the strategy is an evolutionary stable strategy for the perturbed game and, thereby, the only Limit ESS for the game without errors.

By allowing execution errors in the game, ties in payoffs obtained for the Min-Grim, Punish-Once, and Always-Invest strategies against several strategies for player B (see Table 1) could be broken. As can be seen in Table 2, Always-Invest outperforms Min-Grim and Punish-Once against four of the five strategies of player B. Interestingly, Min-Grim obtains a lower payoff compared to Punish-Once and Always-Invest against four of the five strategies of player B. This can be attributed to Grim's high sensitivity if a low return is made which leads to a breach of any investments for all following periods. Given that in a perturbed game low returns by player B occur by mistake, they induce the Grim strategy to produce an inefficient outcome in all periods after the error occurred and, thereby, reduces its payoff.

Concluding the present section, a selected set of strategies has been evaluated according to three concepts. The first, namely the Nash equilibrium concept, could not distinguish particular strategies, since the number of equilibria is too large. The second concept of evolutionary stability is too strong in evaluating strategies, since no strategy of the selected set could be classified as an ESS. The third concept of Limit ESS, is based on the psychological plausible assumption that small errors occur when strategies are selected or applied. This concept could indicate a few strategies of the selected set as outstanding. First, the Min-Grim and Min-Return strategy combination, which leads to an efficient outcome, and thereby, explaining trust and reciprocity, is a Limit ESS under the assumption of selection errors. In addition, the strategy combination Never-Invest and No-Return, which produces an inefficient outcome, could be classified as a Limit ESS for both kinds of errors.

However, these results should be interpreted carefully as the investigation was made only for a very small selected set of strategies. This set is quite representative by consisting of strategies which apply a certain principle up to the limit. Thereby, the selection covers a broad range of strategies, which are built up on similar principles. However, having shown that some strategies of the selected set are Limit ESS, this does not imply that these strategies are Limit ESS for a broader set of strategies. Therefore, in the following, the set of strategies for the game is enlarged. Since the number of automata as strategies for the game (and their combinations) becomes quite large,[12] when the number of states of the automata increases, the analytic exploration is not proceeded. Instead, the results of computer simulations of evolutionary processes are reported. These simulations can be classified according to two evolutionary approaches. Following the first heterogeneous approach, the evolutionary process commences with a population of a variety of strategies and records on how this population develops over the evolutionary process. The goal of the heterogeneous approach is to identify any stable outcomes of an evolutionary process. The second homogenous approach commences with a homogenous population of strategies and makes them subject to a mutation of strategies. The goal of this approach is to classify particular strategies as evolutionary stable.

### *9.4 Simulating an Evolutionary Process*

In the present section, the set of possible strategies is enlarged to all automata with a limited number of two states and subsequent four states. This extension leads to a large number of strategies, which can not be easily investigated by analytic deliberation. Therefore, the *outcome* of an evolutionary process, including potentially all strategies from the possible set, is studied by an evolutionary simulation. However, the main goal is identical: Identifying the type of strategies, which are more liable to establish a stable state in the evolutionary process and form the outcome of the process. The goal is not to describe the evolutionary *process* itself since the evolutionary or learning process, with which people might acquire their strategies, may be different.

Described, in the first section, is the method for simulating the evolutionary processes. The evolutionary process commencing with a population of agents equipped with heterogeneous strategies is reported on in the following two simulations. These simulations can be interpreted as implementing selection errors, since even in a stable state of the evolutionary process, in which agents apply one predominate strategy, there are always some agents who “select” alternative strategies (due to mutation and crossover). Additionally in the third section, errors are implemented in executing strategies. Finally, the evolutionary stability of a population of agents applying particular strategies is explored.

[12] For instance, even automata consisting of only two states can vary according to their 100 possible aspiration levels and their 101 possible outputs, with 2 possible transitions subsequent to high and low reactions of the opponents for each state giving $(101x2x2)^2$x100 strategies for player A. For player B $(101x2x2)^2$x100x2 strategies are possible resulting in $5x10^{14}$ combinations altogether.

*Method: Representing Strategies and Applying a Genetic Algorithm*

Finite state automata are implemented as descriptions of strategies. Each state determines an output, which is restricted for simplification to multiples of 10.[13] Each automaton has an aspiration level for categorizing the opponent's decision. Technically, an automaton is represented by a one-dimensional array. Each state requires three elements determining the output and the two possible moves to other states. Two additional states, for player A, (three states for player B) are necessary for specifying the aspiration level and the initial state(s). Note, that the number of states specifies the maximum of states to which the automata is restricted, which does not imply that all states have to be implemented. Different automata, although they may have a different number of states, can be equivalent according to their behavioral output. The automaton from the set of equivalent automata with the minimum number of states is called the "minimal automaton." For a set of equivalent automata it has been proven that always only one single minimal automaton exists (see Hopcroft & Ullman, 1979, pp. 67-68). There is a requirement to transform automata to the minimal automaton for a better interpretation and, therefore, all automata reported in this paper are minimal automata.

A genetic algorithm is applied for simulating the evolutionary processes (Goldberg, 1989; Michalewicy, 1996; Mitchell, 1996). Genetic algorithms are commonly applied as techniques for solving optimization problems in which analytic solutions are too complex. Recently, they have also found applications in an evolutionary game theory (e.g. Axelrod, 1987; Hoffmann, 1999; Menczer & Belew, 1996). In general the simulated evolutionary process consists of a population of agents. In every single generation, the agents play the investment game against each other for determining agents' fitness. Agents with a higher fitness are more likely to be selected for the next generation, and their strategies eventually modified via a crossover or mutation process, which finally improves their fitness in the next generation.

In the following, the evolutionary process is described in more detail. A population of agents equipped with a pair of automata, one for player A and one for player B, is the subject of the evolutionary selection process. In the first generation of the evolutionary process, 100 pairs of automata are generated by the genetic algorithm. For the first generation, agent's automata are generated randomly, so that the output (investment or

[13] This simplification was necessary to reduce the computational time required to run the simulations. Additionally, it avoided a too larger variance between only slightly different strategies.

return rates) is drawn from a uniform distribution of multiples of 10 ranging from 0 to 100. The states, to which the automata move, are also drawn with equal probability from the set of possible states. The aspiration levels are drawn from a uniform distribution of integers ranging from 0 to 100. At every generation, each agent plays the game 50 times in both player roles. For each of these 100 games the opponent is another agent, who is drawn randomly from the population of agents. Due to the varying number of periods per game, the average payoff is determined for each game. The fitness of an agent is defined as the average payoff across the 50 games the agent plays as player A and the 50 games the agent plays as player B. Theoretically, an agent can reach a maximum payoff of 35 and a minimum of 5 as fitness (composed of the average of player A's payoff [0,30] and player B's payoff [10,40]).

After the fitness is determined, agents are selected for the next generation via a tournament selection procedure (for different selection procedures see Goldberg & Deb, 1991; Michalewicy, 1996). The tournament is conducted 100 times. For each tournament six agents are selected randomly. The agent (respectively their two automata), with the highest fitness, is selected from each tournament. If no agent outperforms the other five agents, one agent from the set is selected randomly with equal probability. By implementing this procedure, agents with a high fitness have a chance of being represented six times in the next generation. However, even agents with low fitness can be selected if they are grouped with agents of lower fitness. Subsequent to the selection procedure, agents' strategies are eventually, with a probability of 0.8, mixed via a two-point crossover procedure. If a crossover takes place, two array positions representing the automata are determined randomly, with equal probability. Subsequently, the elements of the array, from the lowest to the highest position, are interchanged between both automata. After the crossover procedure, automata are eventually modified via a mutation procedure. For each element of the array representing automata, a mutation occurs with a particular probability. The mutation probability depends on automata's number of states, so that with a probability of 0.33 a mutation occurs at least for one element of the array representing the automaton (e.g. for a two state automata this implies a mutation probability of 0.049). If a mutation occurs for the possible moves of the automata, new states, to which the automata could move, are drawn with equal probability from the set of possible states. If a mutation occurs for the automata's output, new values, as multiples of 10, are drawn with equal probability. A mutation of the aspiration level has, in general, a strong effect on the behavior of automata, since it could effect the automata's reactions to the opponents' behavior at every state. Therefore, in the case of a mutation for the aspiration level, new values are drawn from a normal distribution with the mean of the old aspiration level and a standard deviation of 5 (with 0 and 100 as distribution borders), thereby, the aspiration level changed only smoothly.

*Starting From a Heterogeneous Population*

For the first simulations, a heterogeneous population of agents' strategies is generated initially according to the above description. Agents' strategies are made subject to a selection and mutation process. Each evolutionary process is simulated for 10,000 generations. After each 100 generations, automata are transformed to minimal automata, and the means of the automata's parameters are recorded. For analyzing the outcome of the evolutionary processes, the last 10,000$^{th}$ generation is of particular interest. However, due to the permanent mutation and crossover procedure, the 10,000$^{th}$ generation is also only a snapshot of an ongoing evolutionary process. This implies that some new strategies are generated, which do not have high fitness and will not be transferred to the next generation. For sorting out these poor performing strategies, the evolutionary process continues for 200 generations only subject to the selection process without any strategy changes via crossover or mutation. The constitution of these homogenized populations is reported below. Fifty genetic algorithms are run to compensate for random features of particular runs. For transparency in the following, although the agents were subject to the evolutionary process, the strategies of the agents are described separately for player A and player B.

The simulation can be interpreted as incorporating the above described concept of selection errors. At any state of the evolutionary process, even in stable situations due to the mutation and crossover procedure, the population of agents always consisted of some agents who do not select the predominate strategy combination, but instead apply an alternative strategy for either one or both players. In accordance with this fact, the agents were always exposed to mutant strategies attempting to invade the population of agents applying the predominant strategy.

Two simulations are determined with different constraints on the complexity of the strategies. In the first simulation, the automata are restricted to a maximum of two states, whereas in the second simulation the number of states is restricted to four states. The interesting question is whether the evolutionary process will make any utilization of the potential complexity, that is, the higher number of states, which is offered to the strategies, so that other strategies in the condition with a restriction of two states then emerge.

*Simulation of automata with two states.*

The evolutionary processes can be characterized by some general patterns. In the beginning (first 100th generations), one-state automata evolved quickly: For player A the automata Never-Invest, and for player B No-Return, are the baseline automata which appear in the beginning. This formation is often retained for long periods until suddenly new automata appear. The new formation often consisted of the Min-Grim strategy for player A and the Min-Return strategy for player B. The new formation was often stable for a long period before it suddenly returned to the Never-Invest and No-Return formation. Figure 12 shows the development of the average payoffs for agents' strategies for two sample runs.

*Figure 12a.*

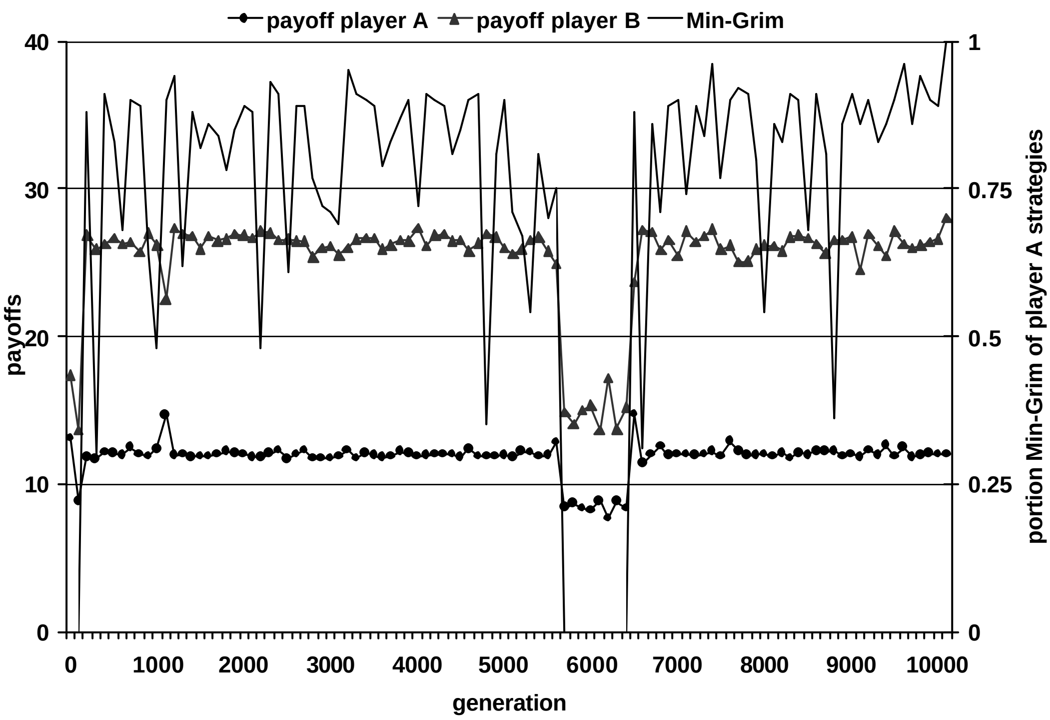

*Figure 12b.*

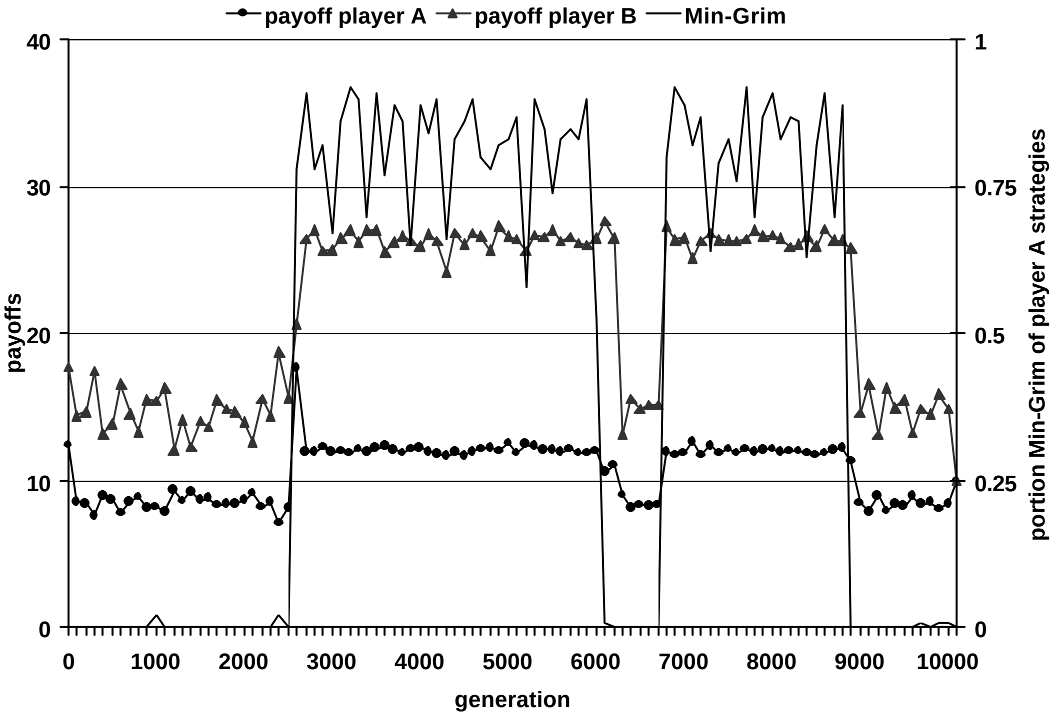


*Figure 12.* Evolutionary process of two sample runs.
The figures show the development of the average payoff for both strategies of the agents. Additionally, the portion of the Min-Grim strategy of player A strategies is shown mapped on the right ordinate.

Generations dominated by the Never-Invest strategies for player A and No-Return strategies for player B led to a payoff of approximately 10 for both players. “Attempts” by player A’s strategies to gain higher payoffs, by making any investment, failed since nothing would be returned. These attempts can explain why player B’s payoffs are, on average, a little above 10 compared to player A’s payoffs, which are a little below 10. Figure 12 also shows the portion of the Min-Grim strategy of the strategies for player A. It becomes apparent that an increase of payoffs for both players corresponds to an increase of the portion of the Min-Grim strategy.

For investigating a potential outcome of the evolutionary process, the last homogenized generation is studied. In sum, from the 50 runs, inefficient outcomes with payoffs of approximately 10 for both players were observed in 12 runs, whereas in the other 38 runs the payoffs were efficient. Surprisingly, the efficient outcomes always lead to a payoff of approximately 12 for player A and of approximately 28 for player B. Due to the restricted outputs of multiples of 10, a payoff of 12 is the minimum highest payoff player A could obtain above player A’s endowment of 10. This is a crucial result as it shows that an efficient outcome can be obtained frequently, however, the outcomes lead to an unequal distribution of the payoff for the two players, hence, an equal or "fair" distribution of payoff is not observed.

For player A, two predominant strategies were observed, that is, the majority of the population of strategies in the last generation consisted of these strategies. For the 12 inefficient outcomes, player A's strategies consisted of 11 times the one-state strategy Never-Invest. In the 38 efficient runs, the Min-Grim strategy was observed 32 times. For the remaining six efficient runs, the last generation consisted of different strategies. For instance, the Punish-Once could also be observed (for all strategies see Figure 10). The aspiration level outcome for player A strategies were between 30% to 40% in 33 runs from the 38 efficient runs. This level ensures that if player A makes an investment of 100% a minimum return above the aspiration level will provide player A with a higher payoff than player A's endowment. The average aspiration level for the 38 efficient runs was 34% with a standard deviation of 4%.

For player B, two predominant strategies were observed. In the 12 inefficient runs, the populations of the last generation consisted of 11 times the No-Return strategy. For the 38 efficient runs, the population consisted of Min-Return strategies in 20 runs. All strategies can also be found in the set of selected strategies described in Figure 11. (In contrast to the description of Figure 11, the Min-Return strategy in the simulations implemented a return rate of 40%, since all return and investment rates were restricted to multiples of 10.) In the remaining 18 runs, from the efficient runs, various strategies were observed, which could not be clustered easily. However, in 12 of these runs the strategies started with a return of 40%, if a high investment was made, and repeated this return rate unless lower investments were made. The aspiration level for classifying the investments as "Trust" or "Distrust" varied substantially across the different strategies. The average aspiration level was 25% with a standard deviation of 25%, indicating that the evolutionary process does not move to a specific level and dilute the role of the aspiration level for player B.

The agents, in 12 inefficient runs, in 11 cases, combined the Never-Invest strategy for player A and the No-Return strategy for player B. In 17 of the 38 efficient runs, in which agents used the Min-Grim strategy for player A, they applied the Min-Return strategy for player B. These were the two predominant combinations of strategies, and as stated above, they form mutual best reply strategies, that is, they form Nash equilibria.

*Increasing the potential complexity of strategies.*

In the second simulation, the number of states for each automaton was increased to four states. This change had a strong effect on the evolutionary process. In only 2 of the 50 runs, compared to 12 runs in the first simulation, the evolutionary process led to an inefficient result. In all the other 48 runs, the agents concluded with an average payoff of approximately 20, producing efficient outcomes, composed of an average payoff of approximately 12 for player A and approximately 28 for player B.

In contrast to the first simulation, these efficient outcomes were predominately not only produced by Min-Grim strategies, which was the case in 15 runs. In 17 runs, variants of the Grim strategy were observed. All of the variants also commenced with an investment of 100%, and repeated the investment if the return was above the aspiration level. However, in contrast to the pure Grim strategy the variants do not punish low returns as strongly as the pure Grim strategy. For instance, one variant moves to the second state subsequent to a low return, in which only half of the endowment is invested in all the following periods, and following a repeated low return, the strategy moves to the third state, in which no investment is made in all the following periods. Although these variants differ from the pure Grim strategy, the best reply strategy for player B is still the Min-Return strategy. Therefore, these variants do not develop any substantial differences to the pure Grim strategy. In contrast to these variants of the Grim strategies, new strategies were also observed (see Figure 13). In 11 of the 50 runs a strategy, called *Cautious*, was observed. This strategy is partly created by the Grim strategy. In comparison to the Grim strategy, it has an additional preliminary initial state with no investment. From this state, it always moves to the second state, which would represent the initial state of the Grim strategy. In five other runs, a strategy called *Forgiving-Once* (or variants) was observed.

This strategy commences with an investment of 100%, and maintains this investment for the duration where player B's returns are above or equal to the aspiration level. If player B's returns are less than the aspiration level, the strategy moves to the second state where again an investment of 100% is made and is repeated unless player B returns less than the aspiration level. After a second low return, the strategy moves to a third state in which no return is made for all the following periods. A best reply strategy for Forgiving-Once, returns nothing in the first period, and then always makes a return equal to the aspiration level of Forgiving-Once. In 43 runs, the aspiration level of the strategies for player A concluded to be quite narrow, between 30% and 40% (*M*=34%, *SD*=4%), hence, only returns of at least 40% were categorized as reciprocity.

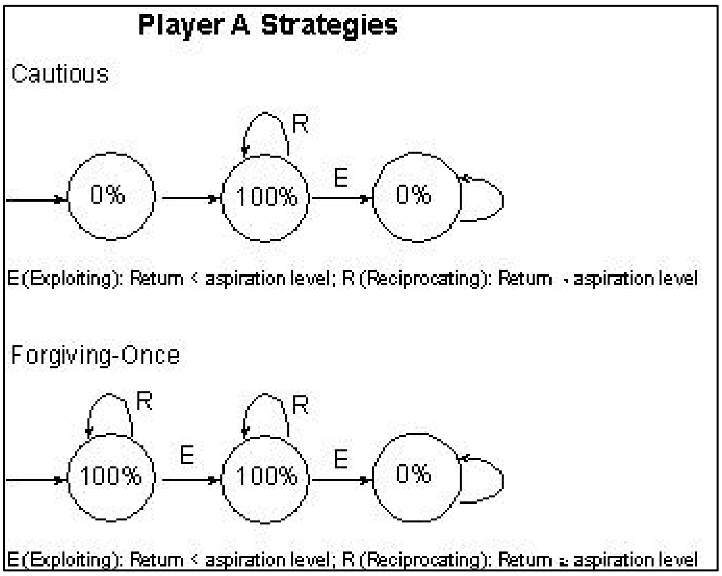


*Figure 13.* Evolved strategies under the condition of a maximum of 4 states.
Since the aspiration levels varied between the different runs no specific level was chosen.

Compared to the first simulation, the No-Return strategy was only observed in two runs from the 50 runs. The Min-Return strategy, which would have been a possible best reply strategy for player A strategies in 40 runs, was observed in only eight runs from the 50 runs. In the remaining 40 runs from 50 runs, no single strategy was observed frequently, making additional classifications difficult. However, most strategies share a common feature: In 30 runs from 40 runs, the strategies for player B always return 40% in the initial state, if player A makes an investment above the aspiration level of the player B strategy. When player A repeats this investment, a return of 40% is always made. In five of the 40 runs, the strategies for player B commence with a low return of 0% and increase this return to 40% in the following period. In sum, although diverse strategies evolve, most of them only differ in their reaction to low investments. The aspiration level for classifying investments varied substantially (*M*=41%, *SD*=29%).

To summarize the results of the second simulations, two aspects should be emphasized. First, the number of efficient outcomes increased. This result may be due to the circumstance that strategies, which incorporate a punishment mechanism, are more liable to evolve if more states are available for building up such a mechanism. Second, due to the higher number of possible states more variants of strategies were observed. However, the decision pattern, produced by the pair of player A and player B strategies, were quite similar. In 35 of the 48 efficient runs, the strategies for player A commenced with an investment of 100%, and strategies for player B always returned 40% of the trebled investment. Given that different kinds of strategies produce this behavior, they are interchangeable. For a given population of strategies, strategies, which produce the same behavior, are able to invade the population. For instance, a three-state automaton for player A, which starts with an investment of 100% and only reduces the investment successively to 50% and 0% in two additional states if a return below 34% is made, produces the same behavior as the Min-Grim strategy and leads to the same payoffs if paired with the Min-Return strategy.

*Errors in Execution Strategies*

The above-mentioned simulations can be interpreted as incorporating selection errors. Another possibility on how alternative behavior could emerge, and thereby, influence the evolutionary process, are execution errors, also described above. This implies that strategies are not applied deterministically, but execution errors occur with low probability. For instance, an agent executing a strategy may invest only 10% by mistake instead of the prescribed 100% investment. When agents make errors in executing a strategy, this could strongly affect the opponent's subsequent decisions. For instance, an agent applying the Grim strategy reacts to low returns by the opponent with unconditional low investments in all following periods. This strong reaction may be unjustified if the opponent's low return only occurred accidentally. The analytical exploration described above has shown, for the selected set of strategies, that only one combination of strategies, Never-Invest and No-Return, are a Limit ESS if execution errors occur. Does this result hold true if the set of strategies is extended to all automata with a maximum of two (four) states? To investigate the effect of execution errors, the above simulations are repeated with the only difference being that, with a probability of 0.01, the prescribed decision of a strategy was not performed, but a decision was randomly determined instead. Given the expected number of 100 periods per game, an error occurred, on average, in only one period.

*Two-state automata.*

For simplification, I only report the results for the last generations. The result was that none of the population of agents in the last generation, from the 50 runs, obtained efficient outcomes. The population of strategies for player A always consisted of Never-Invest automata. The population of strategies for player B consisted in 47 runs of No-Return automata. This result is consistent with our theoretical analysis and shows that only the Never-Invest and No-Return strategies appeared to be Limit ESS when the strategies were restricted to two-state automata. If execution errors occur, the frequently obtained strategy combination of Min-Grim and Min-Return of the previous simulations are not observed.

*Four-state automata.*

The above simulations were repeated, in which the automata were restricted to a maximum of four states, with the only difference being execution errors occurring with a probability of 0.01. The introduction of errors had, again, a strong effect on the evolutionary outcome. In the last generation of the 38 runs from 50 runs, inefficient outcomes were obtained, hence, both players reached a payoff of approximately 10. In 37 of the inefficient 38 runs, the Never-Invest strategy was the predominate strategy for player A in the last generation, and in 29 runs from the 38 runs the No-Return strategy evolved for player B. This result is, again, a strong reduction of efficiency compared to the 48 efficient outcomes of the previous simulations without execution errors.

However, efficient outcomes were obtained in 12 runs. In these runs, the Min-Grim strategy, observed frequently in the previous simulations without errors, was not once obtained in the last generation. On the contrary, a variety of strategies evolved, which could hardly be clustered. Most of the strategies for player A commence with an investment of 100%, and repeat this investment if high returns are made by player B. If a low return is made, most strategies decrease their investments. Subsequently, high returns allow them to return to the initial state with high investments. Only two of the strategies comprise of a "terminal state," which is a state where the automaton will remain for all following periods (e.g. the second state of the Grim strategy is a terminal state). Also, for player B various strategies were observed. The frequently obtained Min-Return strategy of the previous simulations without errors was predominately observed in the last generation for four runs from the 12 efficient runs. In the other runs, different strategies were observed. All returned at least 40% of the trebled investment if repeated high investments were made. In the case of low investments, the returns were reduced—frequently to 0%.

The simulations show that when execution errors occur, the Never-Invest and No-Return is observed more frequently as an outcome of the evolutionary process. With the increased maximum number of states, compared to the previous simulation, efficient outcomes were also obtained. However, in comparison to the simulation without execution errors, the observed strategies for player A are more "forgiving" as they only seldom incorporate a terminal state such as the Grim strategy does. The increase in the maximum number of states allow the strategies to incorporate forgivingness mechanisms, which appeared to be important for obtaining efficient outcomes if execution errors occur.

*Starting From a Homogenous Population*

The previous simulations have shown that, when commencing from a heterogeneous population, the evolutionary process will frequently lead to an efficient outcome if no execution errors occur, which can predominantly be attributed to the Min-Grim and the Min-Return strategies. To test the stability of the Min-Grim and Min-Return strategy combination, the "homogenous approach" is followed.

*Min-Grim and Min-Return.*

An evolutionary simulation is conducted to analyze whether a homogenous population of agents, who apply only the Min-Grim and Min-Return strategy combination, could maintain their strategies during an evolutionary process, in which no execution errors occur. Given that the Min-Grim and Min-Return strategy combination produce efficient outcomes, the only possibility to invade a population of agents who use this combination are strategy combinations, which also produce efficient outcomes.

A simulation was conducted identical to the one reported above without execution errors and a maximum number of two states of the automata, with the difference that all agents in the first generation applied the Min-Grim and Min-Return strategy combination. Additionally, the evolutionary process was only observed for 400 periods, where this number of generations appeared to be sufficient for a potential invasion by mutant strategies. Again, the last generation was homogenized by 200 additional generations, without any strategy modifications.

It turned out that for all runs efficient outcomes were sustained. For player A, the Min-Grim strategy was sustained in 44 runs from the 50 runs, and for player B, the Min-Return strategy could be sustained in 41 runs. In 9 runs, the strategies for player B were most frequently replaced by two-state automata, which reduced the return rate if a low investment was made by player A. In four runs, Punish-Once invaded and replaced the Min-Return strategy, for player B, as the predominant strategy. The evolutionary simulation shows that in most of the 50 runs the Min-Grim and Min-Return strategy combination was the predominate combination of strategies implemented by the agent subsequent to a mutation and crossover process, and that the combination often resisted invasion by other strategies.

*Fair-Grim and Fair-Return.*

One striking result from the above-mentioned simulations without execution errors commencing with a heterogeneous population is that, if an efficient outcome was obtained in an evolutionary process, in most cases it led to a higher payoff for player B than for player A. In contrast to this finding, as the experiments reported in chapter 8 have shown, people often strive for equal payoffs. To explore why this outcome of "fair payoffs" was not obtained more frequently in the above reported evolutionary simulations, an additional simulation was conducted. This second simulation again focuses on the Grim strategy for player A. The simulation was conducted almost identically to the one in the previous section. However, this time Grim's aspiration level was changed to 67%. The strategies for player B consisted of one-state automaton with a return rate of 70%. These strategies lead to a payoff of approximately 20 for both players. The interesting question is whether the population of agents using the "fair" strategies will survive under selection pressure.

Figure 14 shows the development of the median of central main strategies' parameters. The proportion of Grim strategies falls to 62% after 10 generations, which can be attributed to the mutation procedure modifying 33% of all strategies in each generation. Subsequent to this reduction, the proportion of Grim strategies remained constant. The median investment rate of player A's strategies also remained constant at approximately 100%. On the contrary, the aspiration level of player A's strategies reduced to approximately 34% after 100 generations and remained at this level. During the entire evolutionary process, the strategies for player B mainly consisted of one-state automata (with a median state of 1.05 in the last generation), although the 70% return rate from player B's strategies, in the first generation, reduced to approximately 40% after 70 generations.

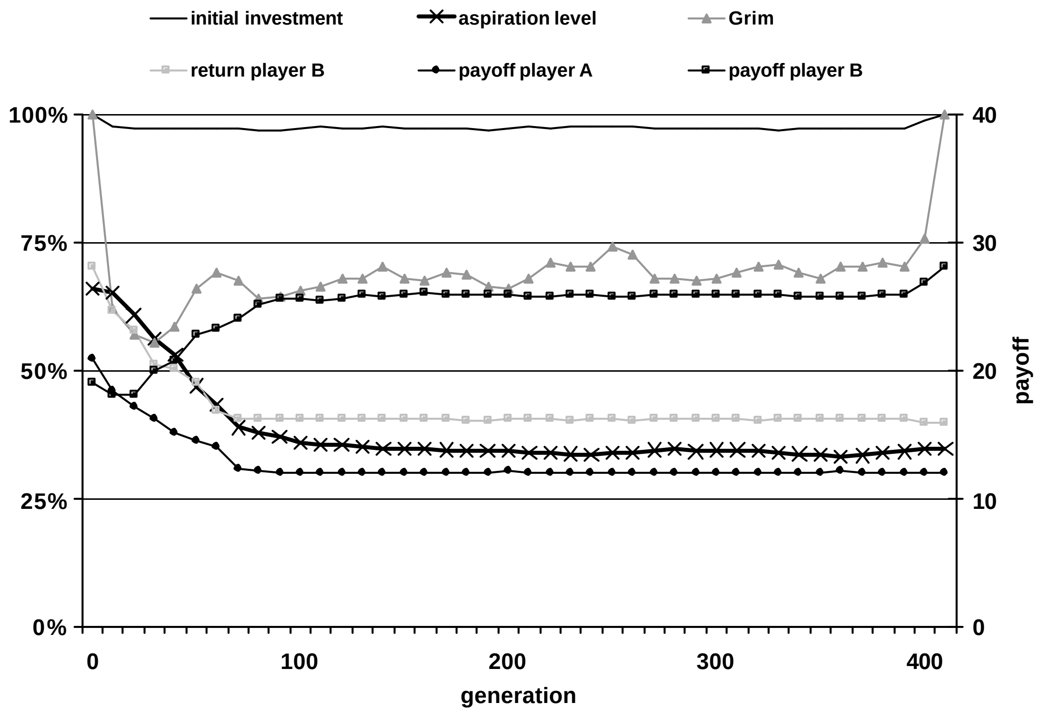

*Figure 14.* Evolutionary process commencing from a homogenous population of agents applying the strategy combination Fair-Grim and Fair-Return.

The figure shows the development of the median values of main variables. The payoffs of the players are plotted on the right y-axis whereas all other variables are plotted on the left y-axis. Initial investment stands for the investment rate of player A in the first period. Aspiration level indicates the level of player A strategies for categorizing the returns of player B as "Reciprocity" or "Exploitation." Grim indicates the proportion of Grim strategies regardless of their aspiration levels. Return player B indicates player B's return rate in the first period, which is a constant return rate in the case of one-state automata. The payoff player A and payoff player B indicate the strategies' payoffs for both player roles.

The simulations demonstrate that strategies, which produce an almost equal payoff for both players, will not sustain for a long period during an evolutionary process. Why is this the case? Given that the strategies for player A consist of Fair-Grim strategies, with an aspiration level of 67%, the best reply strategy for player B is to return at least 67% (i.e. 70% due to the restriction to multiples of 10). However, due to the crossover and mutation process, there will always be some proportion of strategies for player B that will return less. Given the existence of these strategies, an alternative Grim strategy with a lower aspiration level can invade a population of Fair-Grim strategies: On the one hand, these Grim strategies earn the 70% returns from the Fair-Return strategies, on the other hand, they also obtain lower returns from other strategies, unless they are lower than the aspiration level of the alternative Grim strategy. Therefore, the alternative Grim strategy finally outperforms the Grim strategies with high aspiration levels. Once Grim strategies with low aspiration levels have replaced Fair-Grim strategies for player A, the strategies for player B will also reduce their return rates. This co-evolutionary cycle continues until there is no additional possibility for the Grim strategies to reduce their aspiration levels any further, which is the case when the return rates are a little more than player A's endowments. This process can be seen in Figure 14. In the initiation of the process, the median return rate of strategies for player B is below the median aspiration level of player A, and subsequently, the return rate is reduced. The aspiration level is also reduced in the initiation of the evolutionary process but it is always above the return rate until it can not be reduced any further, since player A would then obtain a payoff below the endowment. In this situation, the return rate of player B is adapted to player A's aspiration level, hence, it exceeds the aspiration level.

*9.5 Discussion*

The goal of the present chapter was to investigate strategies for the indefinitely repeated investment game and to indicate those strategies that outperform the others. As a criterion for evaluating the strengths of strategies, the Nash equilibrium concept and two evolutionary stability concepts were applied.

What are "well performing" strategies for the indefinitely repeated investment game? To illustrate this, an analytic exploration was followed for a restricted set of strategies, which demonstrates the strengths and weaknesses of particular strategies. By applying the Nash equilibrium to the restricted set of strategies, it became apparent that many strategies can form equilibria. In principle, all payoff combinations, predicted by the folk theorem, could be obtained with simple finite automata. If, for instance, player A applies a Grim strategy with an aspiration level of an arbitrary return rate of 88%, then it is optimal for player B to apply a strategy with a return rate of 88%, and given this strategy, the Grim strategy is then also optimal for player A. Therefore, the Nash equilibrium concept could not classify any payoff combinations as particularly appealing. However, it became apparent that for reaching an equilibrium a strategy for player A needs to incorporate a punishment mechanism for low returns. If this punishment mechanism is not present, it is never optimal for player B to make any return. In contrast, a punishment mechanism for low investments is not necessary for player B, since if player B makes a substantial return it is in the self-interest of player A to make a high investment.

Additionally, the strategies were evaluated by considering their evolutionary stability. The concept of evolutionary stable strategies resulted in being too strong, since no strategy could be identified as being an ESS. If the psychological plausible assumption is made, that small errors occur when agents apply or select strategies, the concept of Limit evolutionary stability can be applied. A strategy is a Limit ESS for a game if it is an ESS for the game in which these errors occur. For the selected set of strategies, two combinations of strategies could be distinguished as a Limit ESS if errors in selecting strategies occur: the combination of the Never-Invest and No-Return strategies and the combination of the Min-Grim and Min-Return strategies.

Evolutionary simulations were conducted to evaluate a larger set of strategies. The results of the simulations are consistent with the analytic explorations. Beginning with a heterogeneous population of agents using different strategies, the evolutionary process frequently led to two outcomes. Either the Never-Invest and No-Return combination or the

Min-Grim and Min-Return combination was applied by the agents in the last generation of the evolutionary process.

How does an increase in potential complexity change the results? It was most striking that when the maximum number of automata's states was increased Min-Grim, or variants, and Min-Return was still the modal outcome. Strategies applying all four available states were observed rather seldom. The reason for this may be that two states are sufficient for obtaining efficient outcomes, but any additional states will not improve strategies. For instance, the two strategies Cautious and Forgiving-Once do not outperform the simpler Grim strategy if player B makes a constant return of 40%. On the contrary, given a constant return of 40% Cautious is not even a best reply strategy for player A, since player A could make a higher payoff in the first period. The Forgiving-Once strategy has the disadvantage in that it could be exploited in one period. In sum, a strategy, which incorporates more states, does not necessarily improve in enabling a higher payoff, and on the contrary, the larger number of states often provides exploitation possibilities.

What is the reason for the frequent emergence of the Grim strategy? One important reason is that Grim is almost non-exploitable, since it can only be exploited once. However, this advantage is, at the same time, a disadvantage, since it disables Grim from obtaining efficient outcomes again once the strategy has moved to an inefficient decision (i.e. moved to it’s terminal state). What would more cooperative decision strategies, than the Grim strategy for the investment game, look like? Possible candidates are the Punish-Once and the Forgiving-Once strategies obtained in a few simulations. Agents who apply Punish-Once or Forgiving-Once reach the same payoff as Min-Grim does against Min-Return. However, neither strategy is a Limit ESS. The Min-Grim strategy weakly dominates both strategies, since, if the No-Return strategy was played by a few agents, the Min-Grim strategy would earn higher payoffs than the alternative strategies. Therefore, both the Punish-Once and the Forgiving-Once strategy do not appear to be a stable outcome of an evolutionary process.

The general problem for a strategy of player A is that it is only possible for player A to resume efficient outcomes by making a risky outlay that has no negative consequences for player B. In this respect, the asymmetric investment game varies from the prisoner’s dilemma. In the prisoner’s dilemma a player, who causes an uncooperative situation, can initiate cooperation by a cooperative decision in a subsequent period, which will give them a lower payoff than the other player, who does not cooperate. Thereby, a gain, by means of an exploitive decision, is compensated by a loss in another period by resuming cooperation. In contrast, in the investment game, the gain of an exploitive decision by player B is not compensated by a loss for player B when returning to a reciprocal interaction. It is even worse for player A: A loss, resulting from player B's exploitation, will not be compensated when the players return to a reciprocal interaction. Furthermore, in order to return to the reciprocal interaction, player A needs to take the risk of a repeated exploitation with an additional loss into account. This appears to be the main reason why strategies, kinder than Grim, do not evolve more frequently.

However, if agents make errors in executing their strategy, the outcome of an evolutionary process changes. If the automata's number of states were restricted to two states, efficient outcomes would not be obtained under the condition of execution errors. If the number of states was increased to four states, efficient outcomes would, again, be frequently obtained. However, the Min-Grim strategy was not observed once. In contrast, strategies appeared to be more tolerant concerning low returns, since terminal states with no investments for all following periods were seldom observed. Although the strategies reduce their investments if player B makes a low return, they often return to high investments in following periods. The Forgiving-Once strategy, for instance, appears reasonable as it tolerates one single error by the opponent (i.e. no return) during the whole game. Therefore, in the case of execution errors, higher complexity measured in the number of states allows the incorporation of a forgivingness mechanism that tolerates these errors, and thereby, improves the strategies.

Another striking result of the simulations is that in almost all cases, in which an efficient outcome evolved, the players obtained unequal payoffs. The player A only obtained a payoff, which is a little above player A's endowment, whereas the player B received almost all of the produced surplus. This result was expected from the theoretical analysis, since the Fair-Grim and Fair-Return strategy combination is weakly dominated by the Min-Grim and Min-Return combination. The last simulation, in which all agents in the first generation applied the Fair-Grim and Fair-Return strategy combinations, could illustrate the development of unequal outcomes. Fair-Grim is unable to prevent other Grim strategies, with lower aspiration levels, from invading the population. In the beginning, upcoming strategies for player B that return less than Fair-Return produce a lower payoff for their agents. However, these strategies clear the way for Grim strategies with low aspiration levels, and subsequently, for strategies with lower returns for player B. The simulation shows that, from generation to generation, the average aspiration level decreases, and, after a short period of 80 generations, it already reaches the Min-Grim level. In other words, there are Grim strategies that "underbid" the Fair-Grim strategy with a lower aspiration level, and thereby, undermine the "social norm" of high returns. This process resembles a public good problem (Ledyard, 1995) in which the public good is represented by all possible returns of player B. If agents in the role of player A restrict themselves to accepting only high returns, the public good of high returns will be maintained in the long-term, and therefore, this restriction is socially preferable. However, it is individually rational for agents, in the role of player A, to also accept low returns in every period to increase the individual payoff. If a large number of agents make this individual rational decision, it also becomes individually rational for agents, in the role of player B, to lower their return rates. Not surprisingly, the other simulation, which commenced with a homogenous population of agents applying the Min-Grim and Min-Return strategy combination, employ a "social norm" of low returns from the beginning, resulting in having a greater evolutionary stability.