

Imitation

– Theory and Experimental Evidence –*

Jose Apesteguia[†] Steffen Huck[‡]
Public University of Navarre *University College London*

Jörg Oechssler[§]
University of Bonn

February 17, 2004

Abstract

We introduce a generalized theoretical approach to study imitation models and subject the models to rigorous experimental testing. In our theoretical analysis we find that the different predictions of previous imitation models are due to different informational assumptions, not to different behavioral rules. It is more important whom one imitates rather than how. In a laboratory experiment we test the different theories by systematically varying information conditions. We find that the generalized imitation model predicts the differences between treatments well. The data also provide support for imitation on the individual level, both in terms of choice and in terms of perception. But imitation is not unconditional. Rather individuals' propensity to imitate more successful actions is increasing in payoff differences.

JEL codes: C72; C91; C92; D43; L13.

*We thank seminar audiences at Berkeley, UC San Diego, UC Santa Barbara, the Stockholm School of Economics, Stockholm University, Simon Fraser University, Texas A&M, UCL, and the ESA Meetings in Tucson for discussions. Gary Charness, Karl Schlag, and Nat Wilcox provided valuable comments. Financial support from the DFG, grant OE 198/1-1, the Leverhulme Trust, the Economic and Social Research Council via ELSE, and the TMR research network ENDEAR is gratefully acknowledged.

[†]Department of Economics, Public University of Navarre, Campus Arrosadia, 31006 Pamplona, Spain, email: josej.apesteguia@unavarra.es

[‡]Department of Economics and ELSE, University College London, Gower Street, London WC1E 6BT, United Kingdom, email: s.huck@ucl.ac.uk.

[§]Department of Economics, University of Bonn, Adenauerallee 24, 53113 Bonn, Germany, email: oechssler@uni-bonn.de

1 Introduction

Everyone who watches children growing up will attest that imitation is a main source of learning. And introspection shows that imitation plays a large role also for the learning behavior of adults, in particular when faced with unfamiliar environments. While social scientists and psychologists have long recognized the importance of imitation (see Ash, 1952, for an early example), imitation has only recently moved into the focus of economists.

Important theoretical advances have been made by Vega–Redondo (1997) and Schlag (1998 and 1999). Both approaches are based on the idea that individuals who face repeated choice problems will imitate others who obtained high payoffs in previous rounds. But despite this basic similarity, the two theories imply markedly different predictions when applied to specific games. For example, for games with a Cournot structure, Schlag’s model predicts Cournot–Nash equilibrium play,¹ while Vega–Redondo’s model predicts the Walrasian outcome.² The latter prediction is also obtained by Selten and Ostmann’s (2001) notion of an ‘imitation equilibrium’.

The current paper makes two main contributions. First, we introduce a generalized theoretical approach to imitation, which enables us to analyze why the models of Vega–Redondo (1997) and Schlag (1998, 1999) come to such different predictions. We show that the difference between the two models is due to different informational assumptions rather than different learning or adjustment rules. It is more important *whom* one imitates than *how* one imitates. In particular, if one imitates one’s own opponents, outcomes become very competitive. If, on the other hand, one imitates other players who face the same problem as oneself but play against different opponents, then Nash equilibrium play is obtained.

The second objective of our paper is to present rigorous experimental tests of the different imitation models. We chose to study imitation in a

¹Cournot–Nash is also predicted by imitation models in large population settings as studied by Björnerstedt and Weibull (1996).

²Variants of Vega–Redondo’s model in which players have more than one period memory produce different results as shown by Alos–Ferrer (2001) and Bergin and Bernhardt (2001).

normal form game with the payoff structure of a simple discrete Cournot game. This has the advantage that the theoretical predictions of the various imitation models are very distinct. Both traditional benchmark outcomes of oligopoly models (Cournot–Nash equilibrium and Bertrand equilibrium) are supported by at least one imitation model. Also, the games are easy to implement in an experiment, and we have a good understanding of how Cournot markets operate in laboratory environments under different circumstances.³

Despite being inherently “behavioral”, there have been few experiments of imitation models. In particular, Schlag’s imitation model has not been experimentally tested at all, while the models of Vega–Redondo and Selten and Ostmann have been subject to isolated experiments. Huck, Normann, and Oechssler (1999, 2000) and Offerman, Potters, and Sonnemans (2002) experimentally support Vega-Redondo’s model. Also, Abbink and Brandts (2002) provide data that are well-organized by a model closely related to Vega-Redondo’s. Finally, Selten and Apesteguia (2002) find some experimental support for Selten and Ostmann’s (2001) static model of imitation.

The experimental design of the current paper has the advantage that it allows to test all the above mentioned theories in one unified frame. Subjects interact in groups. In one treatment (called GROUP) they can only observe the performance of their opponents with whom they play in the same group. In a second treatment (ROLE), they can only observe the performance of subjects who are in the same role but play in a different group. Finally, there is a ‘FULL information’ treatment in which subjects can observe all of the above.

In our data, all qualitative predictions of the generalized imitation model are confirmed, both, at the aggregate and the individual level. Specifically, average outputs are ranked according to the theoretical predictions and significantly so. In particular, in line with Vega–Redondo, the treatment in which opponents can be observed is the most competitive. The treatment in which only subjects in other groups can be observed is roughly in line

³See e.g. Plott (1989), Holt (1995), and Huck, Normann, and Oechssler (2004) for surveys.

with the Cournot–Nash equilibrium prediction and is the least competitive. Intermediate outcomes result if subjects have access to both types of information.

On the individual level we find that, much in line with Schlag’s model, the likelihood of imitation increases in the difference between the highest payoff observed and the own payoff. Imitation is not unconditional. Moreover, we find that imitation is more pronounced when subjects observe others with whom they directly compete—rather than others who have the same role but play in different groups. The reason for this is still an open question.

All these results are obtained from studying choice data. Subjects do imitate and they do it in specific ways. Whether or not subjects are aware of this, is a different issue on which we shed some light by analyzing replies to a post–experimental questionnaire. Interestingly, many replies quite clearly reveal that subjects know what they are doing. Many also perceive themselves as imitating.

The remainder of the paper is organized as follows. Section 2 introduces the games and the experimental details. In Section 3 we review the imitation models, introduce a general framework, and derive theoretical results. In Section 4 the experimental results are reported and, finally, Section 5 concludes. Most proofs are collected in Appendix A. The instructions for the experiment are shown in Appendix B.

2 Experimental design and procedures

In our experiments subjects repeatedly play simple 3–player normal form games, with a payoff structure that is derived from a symmetric Cournot game. All players have five pure strategies with identical labels, a, b, c, d , and e . Subjects are, however, not told anything about the game’s payoff function apart from the fact that their payoff deterministically depends on their own choice and the choices of two others, and that the payoff function is the same throughout all of the experiment (see the translated instructions in Appendix B).

Interaction in the experiment takes place in populations of nine subjects.

Each subject has a *role* and belongs to a *group*. There are three roles, labelled X, Y , and Z , filled by three subjects each. Roles are allocated randomly at the beginning of the session and then kept fixed for the entire session. Sessions last for 60 periods. In each period, subjects are randomly matched into three *groups*, such that always one X -player is matched with one of the Y -players and one of the Z -players.⁴ While subjects know that they are randomly matched each period, they are not told with whom they are matched and there are no subject-specific labels. In each experimental session, two independent populations of nine subjects participate to increase anonymity. After each period, subjects learn their own payoff. Additional feedback information depends on the treatment. There are three treatments altogether.

Treatment ROLE In treatment ROLE a player is informed, after each period t , of the actions and payoffs in t of players who have the same role as himself but play in different groups.

Treatment GROUP In treatment GROUP a player is informed, after each period t , of the actions and payoffs in t of players in his own group.

Treatment FULL In treatment FULL a player can observe all the information given in treatments ROLE and GROUP.^{5 6}

The payoff function is based on a linear Cournot market with inverse demand, $p = 120 - X$, and zero costs. The strategies a, b, c, d , and e correspond to the output quantities 20, 23, 30, 36, and 40, respectively. That is, a corresponds to the symmetric joint profit maximizing output, c to the

⁴One might wonder why we introduce roles to study behavior in a symmetric game. The answer is twofold. First, this allows us to disentangle the effects of imitation rules and information. Second, we will be able to use the identical setup for studying asymmetric games in follow-up projects.

⁵Notice that in FULL a player can still not observe the choices and payoffs of players that are neither in his group nor in his role.

⁶Additionally, subjects can observe the average payoffs of all nine players in a population. The reason for this is that we plan to study in future work theories based on aspiration levels for which we would like to use FULL as a base treatment.

Table 1: **Payoff table**

		action combination of other players in group														
		aa	ab	ac	ad	ae	bb	bc	bd	be	cc	cd	ce	dd	de	ee
a		1200	1140	1000	880	800	1080	940	820	740	800	680	600	560	480	400
b		1311	1242	1081	943	851	1173	1012	874	782	851	713	621	575	483	391
c		1500	1410	1200	1020	900	1320	1110	930	810	900	720	600	540	420	300
d		1584	1476	1224	1008	864	1368	1116	900	756	864	648	504	432	288	144
e		1600	1480	1200	960	800	1360	1080	840	680	800	560	400	320	160	0

Note: The order in which the actions of the other group members is displayed does not matter.

Cournot output, and e to the symmetric Walrasian output. The payoff table (unknown to subjects) is displayed in Table 1. Subjects are told that the experimental payoffs are converted to Euros using an exchange rate of 3000:1.⁷

The computerized experiments⁸ were carried out in June 2002 in the Laboratory for Experimental Research in Economics in Bonn. Subjects were recruited via posters on campus. For each treatment we carried out three sessions — each with two independent populations of nine subjects, which gives us six independent observations per treatment. Accordingly, the total number of subjects was 162 ($= 9 \times 6 \times 3$). The experiments lasted on average 70 minutes, and average payments were 15.25 Euros.⁹

After the 60 rounds subjects were presented with a questionnaire in which they were asked for their major field of study and for the motivation of their decisions.

3 Imitation theories

In this section we will establish theoretical predictions for various imitation models in the context of our experimental design. Recall that the treatments vary with respect to the information subjects receive about actions and/or

⁷In the first session of treatment FULL we used an exchange rate of 4000:1.

⁸The program was written with z-tree of Fischbacher (1999).

⁹At the time of the experiment one Euro was worth about one US dollar.

payoffs in the previous round. Let player $(i, j)_t$ be the player who has role $i \in \{X, Y, Z\}$ in group $j \in \{1, 2, 3\}$ at time t , and let $s_i^j(t)$ be that player's strategy in t . We refer to the set of individuals whose actions and payoffs can be observed by individual $(i, j)_t$, as $(i, j)_t$'s *reference group*, $R(i, j)_t$. Individual $(i, j)_t$'s *set of observed actions* includes all actions played by someone in his reference group and is denoted by

$$O(i, j)_t := \{s_h^k(t) | (h, k)_t \in R(i, j)_t\}.$$

Notice that $(i, j)_t \in R(i, j)_t$ and $s_i^j(t) \in O(i, j)_t$ in all our experimental treatments.

Following Schlag (1999) we call a behavioral rule *imitating* if it prescribes for each individual to choose an observed action from the previous round. A *noisy* imitating rule is a rule that is imitating with probability $1 - \varepsilon$ and allows for mistakes with probability $\varepsilon > 0$. (In case of a mistake any other action is chosen with positive probability.) A behavioral rule with *inertia* allows an individual to change his action only with probability $\pi \in (0, 1)$ in each round. In the following we shall first characterize different imitation rules according to their properties without noise or inertia. Predictions for the Cournot game will then be derived by adding noise and inertia.

A popular and plausible rule is “imitate the best” (see e.g. Vega-Redondo, 1997) which simply prescribes to choose the strategy that in the previous period performed best among the observed actions. In our setting it is possible that an action yields different payoffs in different groups. This implies that it is *a priori* not clear how an agent should evaluate the actions he observes. An *evaluation rule* assigns a value to each action in a player's set of observed actions $O(i, j)_t$. When an action yields the same payoff everywhere in his reference group, there is no ambiguity and the action is evaluated with this observed payoff.¹⁰ When different payoffs occur for the same action, various rules might be applied. Below we will focus on two evaluation rules that appear particularly natural in a simple imitation setting with boundedly rational agents: the *max rule* where each strategy is evaluated according to the highest payoff it received, and the *average rule*

¹⁰This is always the case in treatment GROUP.

where each strategy is evaluated according to the average payoff observed in the reference group. Of course, other rules, such as a “pessimistic” min rule, might also have some good justification. Nevertheless, we shall follow the previous literature and focus on the max and the average rules.¹¹

Definition 1 *An imitating rule is called “imitate the best” if it satisfies the property that (without noise and inertia) an agent switches to a new action if and only if this action has been played by an agent in his reference group in the previous round, and was evaluated as at least as good as that of any other action played in his reference group. When several actions satisfy this, any of those is chosen with positive probability.*

- *“Imitate the best” combined with the average rule is called “imitate the best average” (IBA).*
- *“Imitate the best” combined with the max rule is called “imitate the best max” (IBM).*

Schlag (1998) shows in the context of a decision problem in which agents can observe one other participant that “imitate the best” and many other plausible rules do not satisfy certain optimality conditions. Instead, Schlag (1998) advocates the “Proportional Imitation Rule” which prescribes to imitate an action with a probability proportional to the (positive part of the) payoff difference between that action’s payoff from last period and the own payoff from last period. If the observed action yielded a lower payoff, it is never imitated.

The extension of this analysis to the case of agents observing two or more actions is not straightforward. Schlag (1999) considers the case of two observations and singles out two rules that are both “optimal” according to a number of plausible criteria, the “double imitation” rule (DI) and the “sequential proportional observation” rule (SPOR). In both cases, Schlag assumes that strategies are evaluated with the average rule. Specifying the two rules in more detail is beyond the scope of this study since our data do

¹¹For “imitate the best average”, see, e.g., Ellison and Fudenberg (1995) and Schlag (1999). For “imitate the best max”, see Selten and Ostmann (2001).

not allow to check more than some general properties of *classes* of rules to which DI and SPOR belong.

Schlag (1999, Remark 2) shows that with two observations both, DI and SPOR, satisfy the following properties:

- (i) They are imitating rules.
- (ii) The probability of imitating another action increases with that action's previous payoff and decreases with the payoff the (potential) imitator achieved himself.
- (iii) If all actions in $O(i, j)_t$ are distinct, the more successful actions are imitated with higher probability.

Furthermore, it can be shown that DI satisfies the following plausible properties.

- (iv) Never switch to an action with an average payoff lower than the average payoff of the own action.
- (v) Imitate the action with the highest average payoff in the sample with strictly positive probability (unless one already plays an action with the best average payoff).
- (vi) Never switch to an action with average payoff below the average payoff in the sample.

Property (iv) shows that DI belongs to the large class of imitating rules that use the average evaluation rule and can be described as “imitate only if better”. Combined with property (v) “imitate the best with positive probability” this is all we need for deriving the theoretical properties of DI and similar rules in the context of our experiment.

Definition 2 *An imitating rule is called a “weakly imitate the best average” rule (WIBA) if it satisfies (without noise and inertia) properties (iv) and (v).*

If we modify Properties (iv) and (v) to allow for the max rule, we obtain

- (iv') Never switch to an action with a maximal payoff lower than the maximal payoff of the own action.
- (v') Imitate the action with the highest maximal payoff in the sample with strictly positive probability (unless one already plays an action with the highest maximal payoff).

Definition 3 *An imitating rule is called a “weakly imitate the best max” rule (WIBM) if it satisfies (without noise and inertia) properties (iv') and (v').*

While IBA (“imitate the best average”) as well as DI (double imitation) belong to the class of WIBA (“weakly imitate the best average”) rules, IBM belongs to WIBM. The rule SPOR does not belong to either class of rules since it violates (iv) and (iv').

Before we proceed with deriving theoretical predictions, we need to introduce some further notation. The imitation dynamics induce a Markov chain on a finite state space Ω . A state $\omega \in \Omega$ is characterized by three strategy profiles, one for each group, i.e., by a collection $((s_1^1, s_2^1, s_3^1), (s_1^2, s_2^2, s_3^2), (s_1^3, s_2^3, s_3^3))$. Notice that there is no need to refer to specific individuals in the definition of a state, i.e., here s_i^j (without the time index) refers to the strategy used by whoever has role i and happens to be in group j .

We shall refer to *uniform states* as states where $s = s_i^j = s_h^k$ for all i, j, h, k and denote a uniform state by ω^s , $s \in \{a, b, c, d, e\}$. Two uniform states will be of particular interest. The state in which everybody plays the Cournot Nash strategy c , to which we will refer as the *Cournot state* ω^c ; and the state in which everybody plays the Walrasian strategy e , to which we shall refer as the *Walrasian state* ω^e .

To analyze the properties of the Markov processes induced by the various imitation rules discussed above, we shall now add (vanishing) noise and inertia. That is, whenever we refer in the following to some rule as, for example “imitate the best”, we shall imply that agents are subject to, both, inertia and (vanishing) noise. States that are in the support of the limit

invariant distribution of the process (for $\varepsilon \rightarrow 0$) are called *stochastically stable*. The (graph theoretic) methods for analyzing stochastic stability (pioneered in economics by Canning, 1992, Kandori, Mailath, and Rob, 1993, and Young, 1993) are, by now, standard (see e.g. Fudenberg and Levine, 1998, and Young, 1998, for text book treatments).

In the following we will state a number of propositions that show how the long-run predictions of the imitation rules we consider depend on the underlying informational structures. We begin by stating results for WIBA and WIBM. It will turn out that WIBA and WIBM rules lead to identical predictions if agents either observe other agents in their group *or* other agents in the same role. They differ if agents can observe both as in treatment FULL. Finally, we will analyze SPOR rules and show that they yield the same long-run predictions regardless of the treatment.

Our first proposition concerns WIBA and WIBM rules in treatment GROUP.

Proposition 1 *If agents follow either a WIBA (“weakly imitate the best average”) or a WIBM (“weakly imitate the best max”) rule and if the reference group is as in treatment GROUP, the Walrasian state ω^e is the unique stochastically stable state.*

Proof see Appendix A.

The intuition for this result is similar to the intuition in Vega-Redondo’s original treatment of the imitate the best rule. In any given group, the agent with the highest output obtains the highest profit as long as prices are positive. This induces a push toward more competitive outcomes.¹²

Let us now turn to treatment ROLE where $(h, k)_t \in R(i, j)_t$ if and only if $h = i$. We will see that the change of the informational structure has dramatic consequences. If agents can only observe others who are in the same role as they themselves but play in different groups, the unique

¹²Introducing constant positive marginal cost does not change the result. If price is below marginal cost, the agent with the lowest output is imitated which again pushes the process towards the Walrasian state.

stochastically stable outcome under a WIBA rule and a WIBM rule is the Cournot–Nash equilibrium outcome.

Proposition 2 *If agents follow a WIBA or a WIBM rule and if the reference group is as in treatment ROLE, the Cournot state ω^c is the unique stochastically stable state.*

Proof see Appendix A.

The intuition for Proposition 2 is that any deviation from the Cournot Nash equilibrium play lowers the deviator’s absolute payoff. Agents in the same role will observe this but will not imitate because they earn more using the equilibrium strategy. On the other hand, one can construct sequences of one-shot mutations that lead into the Cournot state from any other state.

Turning to treatment FULL one might expect that its richer informational structure (where agents have the combined information of treatments GROUP and ROLE) causes some tension between the Walrasian and the Cournot outcome. It turns out that this intuition is correct. In fact, with a WIBA rule there are two stochastically stable states in treatment FULL, the Cournot state (where everybody plays c), and the state where everybody plays d .

Proposition 3 *If agents follow a WIBA rule and if the reference group is as in treatment FULL, then both, the Cournot state ω^c and the state in which everyone takes action d , ω^d , are the stochastically stable states.*

Proof see Appendix A.

Comparing a WIBA rule with a WIBM rule, one might say that agents following WIBM are “more aggressive”. Hence, one might intuitively expect that WIBM leads to higher quantities than WIBA. As the next proposition shows this is true in the sense that, in addition to ω^c and ω^d , the Walrasian state, ω^e , is stochastically stable under WIBM.

Proposition 4 *If agents follow a WIBM rule and if the reference group is as in treatment FULL, then the Cournot state ω^c , the state in which everyone*

takes action d , ω^d , and the Walrasian state ω^e are the stochastically stable states.

Proof see Appendix A.

In contrast to the previous studied rules, the SPOR rule of Schlag (1999) also allows to imitate actions that do worse than the current action one is using. This has the consequence that, in the framework of stochastic stability, any uniform state can be a long run outcome of the process.

Proposition 5 *If agents follow a SPOR rule, all uniform states are stochastically stable regardless of their reference group.*

Proof Agents following SPOR imitate any strategy with positive probability except an action that yields 0, the absolutely worst payoff (see Schlag, 1999). Thus, we observe a) that only uniform states are absorbing and b) that it is possible to move from any uniform state to any other uniform state by just one mutation, which implies that all uniform states are stochastically stable. ■

3.1 Imitation Equilibrium

We shall now review the recently introduced notion of an *imitation equilibrium (IE)* (Selten and Ostmann, 2001), and derive its predictions for our treatments. Unlike the preceding models, imitation equilibrium is a static equilibrium notion. Following Selten and Ostmann (2001) we will say that player (i, j) has an *imitation opportunity* if there is an $s_h^k \neq s_i^j$, $s_h^k \in O(i, j)$, such that the payoffs of player (h, k) are the highest in $R(i, j)$ and there is no player in $R(i, j)$ playing s_i^j with payoffs as high as (h, k) .¹³ A *destination* is a state without imitation opportunities. An *imitation path* is a sequence of states where the transition from one element of the sequence to the next is defined by all players with imitation opportunities taking one of them. The imitation path continues as long as there are imitation opportunities.

¹³This requirement is the same as in IBM.

An *imitation equilibrium* is a destination that satisfies that all imitation paths generated by any deviation of any one player return to the original state. Two classes of imitation paths generated by a deviation (henceforth called *deviation paths*) that return to the original state are distinguished.

(i) *Deviation paths with deviator involvement*: the deviator himself takes an imitation opportunity at least once and the deviation path returns to the original state.

(ii) *Deviation paths without deviator involvement*: the destination reached by a deviation path where the deviator never had an imitation opportunity gives lower payoffs to the deviator than those at the original state, making that the deviator returns to the original strategy. This creates an imitation path that returns to the original state.

Proposition 6 *Imitation equilibrium (IE) is characterized by the following.*

(a) *In Treatment GROUP the Walrasian state ω^e is the unique IE.*

(b) *In Treatment ROLE the Cournot state ω^c is the unique IE.*

(c) *In Treatment FULL ω^c , ω^d , and ω^e are the only uniform IE.*

Proof see Appendix A.

The proposition reveals remarkable similarities between Selten and Ostmann’s imitation equilibrium and the dynamic class of WIBM rules. In fact, imitation equilibrium and the long-run predictions of WIBM coincide perfectly for the current game.

3.2 Some qualitative hypotheses and simulations

Table 2 summarizes the theoretical results and indicates for each behavioral rule considered above whether two easy-to-check properties are satisfied *with respect to the evaluation rule* used.

All imitation rules, with the exception of SPOR, have in common that they predict that agents should not switch to strategies that are evaluated as worse than the strategy they currently employ. Moreover, it can be shown that DI and IBA additionally satisfy the property that no strategy is imitated that (on average) yielded a payoff below the population average.

Table 2: Summary of predictions

Imitation Rule	never imitate worse than own*	never imitate worse than avg.*	long run prediction**
WIBA	✓	–	ω^e in GROUP ω^c in ROLE ω^c, ω^d in FULL
DI	✓	✓	as WIBA
IBA	✓	✓	as WIBA
WIBM	✓	–	ω^e in GROUP ω^c in ROLE $\omega^c, \omega^d, \omega^e$ in FULL
IBM	✓	–	as WIBM
SPOR	–	–	$\omega^a, \omega^b, \omega^c, \omega^d, \omega^e$
IE	✓	–	as WIBM

Note: A “✓” indicates that the theory in question satisfies the property *given* the rule to evaluate payoffs. “–” indicates that the theory does not in general satisfy this property. * this prediction is without noise. **In all cases except IE this is the set of stochastically stable outcomes.

With respect to outputs, all imitation rules, except SPOR, suggest that outputs in treatment GROUP (where Walrasian levels are expected in the long run) should be rather high, whereas in treatment ROLE a lower output level, the Cournot outcome, is expected. Since treatment FULL provides both kind of information, one might expect an intermediate outcome between GROUP and ROLE. This suggests the following qualitative hypothesis about the ordering of output levels:

$$\mathcal{QH} : \text{ROLE} \preceq \text{FULL} \preceq \text{GROUP}.$$

In order to address justified concerns that stochastic stability analysis may fail to yield reasonable predictions for an experiment with only 60 rounds, we run simulations for the different treatments. In particular, we simulate a population of 9 players over 60 rounds when each player behaves according to the IBM rule (IBA yields almost identical results) given the reference group defined by the respective treatment. Since stochastic stability analysis is often criticized for the assumption of vanishing noise, we include




Figure 1: Block time series of average outputs per treatment (simulated data). Note: “Block 0” denotes average outputs in the first period.

in the simulation a substantial amount of noise. With probability 0.8 in each round a player follows IBM. With probability 0.2 a player chooses randomly one of the five actions with equal probability. For each treatment we simulated 100 such populations with starting actions chosen from a uniform distribution.

Figure 1 shows average quantities in those simulations.¹⁴ Already in the first 10 periods, the qualitative prediction expressed in \mathcal{QH} is apparent. After 20 rounds, behavior is fairly constant. Over all 60 periods, average outputs in the simulations were 30.3 for ROLE, 33.6 for FULL, and 37.7 for GROUP. In the following sections, we shall test hypothesis \mathcal{QH} against the aggregate results from the experiments.

¹⁴Recall that actions a, b, c, d , and e correspond to output quantities 20, 23, 30, and 40, respectively.

4 Experimental results

We now turn to the experimental analysis of the generalized imitation framework proposed above. We organize this section as follows. First, based on the qualitative hypotheses \mathcal{QH} derived above, we evaluate the data on the aggregate level. This will give a first idea of whether and how imitation shapes subjects' behavior. We then turn our attention to the study of individual behavior. In Section ?? all imitation rules are evaluated with respect to what they say about when *not* to imitate while in Sections ?? we analyze the issues of when and how to imitate. To evaluate rules that make explicit statements about the *probability* of imitating a certain action (like those of Schlag, 1998, 1999), we study in Section 4.3 a probit model. Finally, we conclude this section by analyzing the post-experimental questionnaires. This will provide additional insight whether subjects are intentional imitators or whether it just looks *as if* they are.

4.1 Aggregate behavior

We begin by considering some summary statistics on the aggregate level. Table 3 shows average outputs for all treatments, separately for the entire sixty rounds of the experiment and the last thirty. Standard deviations of the six observations per treatment are shown in parenthesis. The average of individual variances over all 60 and the last 30 periods are also given in Table 3.

Figure 2 shows average outputs per treatment in an average block time series, organized in blocks of ten periods.

Both, Table 3 and Figure 2, clearly show that output levels are ordered as predicted. In particular, outputs in ROLE are lower than in FULL, and in FULL lower than in GROUP. The p -values for (two-sided) permutation tests (see, e.g., Siegel and Castellan, 1988) on the basis of the average outputs per population are as follows:

$$\text{ROLE} \prec_{.008} \text{FULL} \prec_{.02} \text{GROUP}$$

This is exactly in line with the qualitative predictions derived in the pre-

Table 3: Summary statistics

	Treatment		
	ROLE	GROUP	FULL
avg. output ₁₋₆₀	30.32 (.452)	32.96 (.825)	31.71 (.829)
avg. output ₃₁₋₆₀	30.71 (.514)	33.57 (.972)	32.31 (.881)
avg. var. ₁₋₆₀	54.43	49.78	51.31
avg. var. ₃₁₋₆₀	53.02	43.00	49.07

Note: avg. outputs are calculated by using the output levels 20, 23, 30, 36, and 40. Standard deviations of avg. output of the 6 independent observations per treatment are given in parenthesis. $\text{avg. var}_{\tau-t}$ is the average over individual variances in output from round τ to t .




Figure 2: Block time series of average outputs per treatment (actual data).

vious section. We see that the type of feedback information available to subjects has a statistically significant impact on behavior — lending strong qualitative support to the generalized imitation model and stressing the importance of informational conditions. The differences in average quantities are less pronounced than predicted by theory though, which indicates that other factors besides imitation play a role too.

Also, Table 3 and Figure 2 indicate that there is an increasing trend in output. The Spearman rank order correlation coefficients between the time series of the average output levels and time are positive in all treatments, and significant at the 0.01 level (two-sided) in GROUP ($r_s = 0.49$) and FULL ($r_s = 0.62$), but not significant in ROLE ($r_s = 0.23$).

For now, we summarize our main finding in

Result 1 *Outputs are ordered as predicted by hypothesis \mathcal{QH} and significantly so.*

Given the usual noise in experimental data from human subjects, Result 1 seems quite remarkable. However, before drawing more definite conclusions about the viability of imitation it is necessary to analyze individual adjustments which we shall do in the following section.

4.2 Individual Behavior

A proper experimental test of imitation theories needs to consider individual data. Thus, in this section we evaluate the success of the imitation models by computing compliance rates of individual adjustment behavior with the predictions of the respective models. We begin by classifying individual behavior into the following categories: (i) ‘Best’: the subject switched to last period’s best action in his reference group, (ii) ‘Better’: subject switched to an action that was better than his own action, but not the best, (iii) ‘Same’: there was no change of action, (iv) ‘Worse’: change to an action that was worse than his own action, and (v) ‘Different’: change to an action that was not observed in the reference group. Table 4 reports the classification of behaviors according to categories (i) through (v) separated by treatment and the evaluation rule.

Table 4: Classification of Individual Behavior by Type of Change

	Best	Better	Same	Worse	Different
ROLE	34.9%	1.7%	13.4%	8.5%	41.5%
	35.9%	1.8%	12.3%	8.5%	
GROUP	41.2%	2.3%	18.1%	5.3%	33.1%
FULL	32.6%	7.0%	22.8%	11.1%	26.5%
	32.8%	7.1%	22.7%	10.9%	

Note: Reported are the percentages of subjects that switched to actions in the various categories. Upper values are calculated using the average rule, lower values by using the max rule.

Recall that the average and the max rule coincide in Treatment GROUP by definition. Obviously, observed rates under category ‘Different’ are the same irrespective of the imitation rule under consideration. For the other categories, the percentages may differ according to the evaluation rule. Note, however, that the observed rates for both evaluation rules are very close together. This is due to the fact that the two rules typically prescribe the same actions (because the strategy with the highest max is typically also the one with the highest average). Only in less than 2% of all cases do they diverge.

We begin by inspecting Table 4 with respect to predictions about when *not* to imitate. WIBA, WIBM, and IE predict that one should not imitate an action with payoffs evaluated as lower than those of the own action (applying the respective evaluation rule). To extract from Table 4 the compliance rates of WIBA, WIBM, and IE with respect to the prediction of when not to imitate, we sum the corresponding values in all categories except ‘Worse’. Table 5 reports the compliance rates sorted by treatment and by evaluation rule.

The compliance rates are rather high in all cases, particularly in treatment GROUP. This means that there are few strict violations of imitation rules belonging the set of WIBA/WIBM rules or to IE. However, in order to have a meaningful measure of success for evaluating the performance of the

Table 5: Compliance Rates

	ROLE	GROUP	FULL
when not to imitate (average rule)	91.5%	94.7%	88.9%
when not to imitate (max rule)	91.5%	94.7%	89.1%
imitating	58.5%	66.9%	73.5%
weighted imitating w_p	56.5%	65.0%	69.4%
IBA	48.3%	59.3%	55.4%
IBM	48.5%	59.3%	55.5%
IBA given imitation	82.5%	88.7%	75.4%
IBM given imitation	82.4%	88.7%	75.4%

imitation rules, we need a method that contrasts the observed compliance rates with those that would obtain *if there were no relation between behavior and imitation*. We use the following method. We randomly simulate the behavior of 100 populations of nine players for 60 periods, and calculate the success of the hypotheses relative to this simulated data. In order to give random behavior the best shot, we take the experimentally observed frequencies from the experiment as the theoretical distribution from which random behavior is generated. The average rates of compliance (and standard deviations) pooling the data in GROUP, ROLE, and FULL with the prediction on when not to imitate according to the average rule are 84% (.014), and according to the max rule are also 84% (.014). The permutation test shows that the compliance rates reported in Table 5 are higher than those randomly obtained at all standard significance levels.

Result 2 *On average, behavior of subjects is in line with the predictions of when not to imitate in more than 90% of all cases. This result significantly outperforms random predictions.*

We now turn to predictions about *when to imitate*. Before checking particular imitation rules, we first ask how often subjects' behaviors can be classified as imitating in general. Recall that all rules described in Section 3 are imitating behavioral rules, i.e. they prescribe to imitate an action that has been observed in the previous period (abstracting from noise). Summing the corresponding values in all categories except the one for 'Different'

in Table 4 gives the percentages of decisions that are imitating. Table 5 reports the corresponding values by treatment. Clearly, there are many non-imitating choices (between 26.5 and 41.5%). Those may be viewed as instances of experimentation or as behaviors based on own past payoffs. Thus, it is certainly not appropriate to consider imitation as the exclusive explanation for observed behavior. Furthermore, since this measure may sometimes give a misleading impression, we compute imitating compliance rates weighted with respect to the number of different actions a player observes. The weighted compliance rate w_p for population p is defined as

$$w_p = \frac{\sum_{i,j,t} (5 - \#O(i,j)_t) I(i,j)_t}{\sum_{i,j,t} (5 - \#O(i,j)_t)}, \quad (1)$$

where $\#O(i,j)_t$ is the number of different actions in player (i,j) 's set of observed actions at period t , and $I(i,j)_t$ is a dummy variable that equals 1 if player (i,j) was imitating at period t and 0 otherwise. The weight on an imitating choice is decreasing in the number of actions a player observed. (If all five actions are observed, any choice is imitating; hence there is zero weight.) The weighted compliance rates for treatments ROLE, GROUP, and FULL are reported in Table 5.

Now we can analyze the compliance rates with two particular imitation rules belonging to the sets of WIBA/WIBM rules; IBA and IBM. According to Definition 1, to get the compliance rates with the predictions of IBA and IBM we should add the values obtained in 'Best' and 'Same'. Table 5 reports the corresponding compliance rates sorted by treatment. Given that there are many non-imitating choices, it is not surprising that the compliance rates for IBM or IBA are not terribly high, although we will show below that they are significantly higher than under random play. In the last two lines In Table 5 we compute the compliance rates for IBM and IBA *given that subjects do imitate*. Clearly, these are remarkably high indicating that *when* players imitate, IBA and IBM capture behavior to a great extent.

In order to compare the data to random behavior using the method described above, we again simulated 100 populations of 9 players. The average (absolute) rate of compliance thus obtained for IBA and IBM pooled for

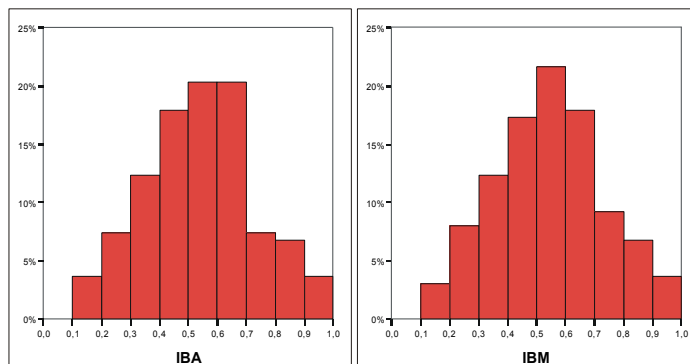


Figure 3: Distribution of individual players on the basis of the compliance rates with IBA and IBM for ROLE, GROUP, and FULL, pooled together.

GROUP, ROLE, and FULL under random choice is 34.6% with a standard deviation of 3.3%. Permutation tests on the basis of the average rates of compliance for the populations show that both imitation rules outperform random predictions at any conventional significance level. This further confirms that imitation is present in our data, and that, in particular, IBM and IBA play a significant role in explaining it. Figure 3 shows the distribution of individual players on the basis of the (absolute) rates of compliance with the predictions of IBA and IBM for treatments GROUP, ROLE, and FULL pooled together. About 10% of the players show a percentage of absolute compliance with IBM and IBA higher than 80%. This suggests that there is a sizeable number of almost pure imitators. It is also worth noting that more than 35% of the participants show a compliance rate with IBM and IBA higher than 60%.

Result 3 *IBM and IBA do about equally well, and both outperform random predictions significantly. Moreover, 10% of subjects are almost pure imitators whose choices are in line with IBM/IBA in more than 80% of all decisions.*

Our analysis of individual adjustments confirms the insight obtained by looking at aggregate data. If individuals can imitate actions, most of them

do so, and some do so almost all the time.

Finally, Table 5 shows that the (absolute) rates of compliance with IBM and IBA are higher in GROUP than in ROLE (in fact, Table 5 shows that GROUP gets higher levels of compliance in all the dimensions we have studied here). The permutation test yields significance at the .05 level (two-sided).¹⁵ This is an interesting finding that will gain further support below. Intuitively, one might expect that imitation of others who are in the same role as oneself is more appealing than imitation of a competitor who, after all, might have a different payoff function. Recall that, at least initially, our subjects do not know that they are playing a symmetric game. Also, subjects are randomly rematched every period and can not expect to face the same opponents as last period.

Result 4 *Imitation is significantly more pronounced when subjects can observe their immediate competitors (as in treatment GROUP) than when they can observe others who have the same role in different groups (as in treatment ROLE).*

4.3 Estimating imitation rules

The predictions of Schlag’s imitation rules “Proportional Imitation”, DI and SPOR explicitly refer to the *probability* of imitating an action. To do justice to these predictions, we present in this section probit estimates for subjects’ choice functions. In particular, we analyze how subjects’ decisions to change their action depends on their own payoff and the best payoff they observe. Furthermore, we also analyze how the likelihood of following IBM depends on a subject’s own payoff and the best payoff the subject observes.¹⁶

Table 6 shows the result of estimating the following probit model with random effects for all treatments,

$$\Pr(s_i^t \neq s_i^{t+1}) = \Phi(\alpha + \beta\pi_i^t + \gamma\pi_{i\max}^t + v_i + \varepsilon_i^t), \quad (2)$$

¹⁵All other pairwise comparisons are not statistically significant.

¹⁶Due to the high correlation of the best max and the best average, results for IBA are very similar and, therefore omitted.

Table 6: Estimating the likelihood that subjects change their action

	ROLE	GROUP	FULL
Constant α	1.103*** (.1239)	.2197** (.0960)	.1360 (.1074)
Own payoff β	-.0013*** (.0001)	-.0019*** (.0002)	-.0011*** (.0001)
Max payoff γ	.0003*** (.0001)	.0014*** (.0001)	.0005*** (.0001)
# of obs.	3186	3186	3186

Note: *** denotes significance at the 1% level, ** denotes significance at the 5% level.

where s_i^t denotes subject i 's strategy in period t , π_i^t the subject's payoff, $\pi_{i \max}^t$ the maximal payoff the subject observed in his reference group. Φ is the standard normal distribution, v_i the subject-specific random effect, and ε_i^t the residual.

In agreement with imitation rules DI and SPOR, the regressions show that coefficients for own payoffs are significantly negative while those for maximal observed payoffs are significantly positive. This holds for all treatments but it is most pronounced for GROUP.

After analyzing *when* subjects switch to a different action, we shall now analyze *where* they switch to *if* they switch. Table 7 reports subjects' likelihood of following IBM dependent on their own payoff and the maximal observed payoff. These regressions are only run for cases in which a subject actually switched to another action (since the theories allow for inertia, not switching is always in line with the prediction). The estimated equation is

$$\Pr(s_i^{t+1} = s_{i \max}^t) = \Phi(\alpha + \theta\pi_{i \max}^t + \lambda(\pi_{i \max}^t - \pi_i^t) + v_i + \varepsilon_i^t), \quad (3)$$

where $s_{i \max}^t$ is the action that had the highest maximal payoff (IBM) in period t in subject i 's reference group. All other variables are as defined before. Note, however, that in contrast to regression (2) we now include $\pi_{i \max}^t$ directly, and also in form of a variable which represents the difference between the maximal observed payoff and the own payoff. We chose this form to be able to test whether only the difference matters, as predicted e.g.

Table 7: Estimating the likelihood that subjects follow IBM

	ROLE	GROUP	FULL
Constant α	-1.1354*** (.1593)	-1.0767*** (.0982)	-1.0008*** (.1442)
Max payoff θ	-.00005 (.0001)	-.0003* (.0001)	-.0002 (.0001)
Difference to max λ	.0009*** (.0001)	.0023*** (.0002)	.0003*** (.0001)
# of obs.	2079	1644	1920

Note: Only cases with $s_i^{t+1} \neq s_i^t$ included. *** denotes significance at the 1% level, * significance at the 10% level.

by Schlag's Proportional Imitation rule, or whether own payoff and max payoff enter independently.¹⁷ If θ is not significantly different from zero, then only the payoff difference matters.

Table 7 shows that in fact only the payoff difference matters. In all three treatments the coefficient of the difference variable has the expected sign and is significant at the 1% level. In contrast, the coefficient of the maximal payoff observed is only (weakly) significant in treatment GROUP and not significantly different from zero in the other treatments. This is strong support for all rules that satisfy Property (ii) above, in particular for Schlag's Proportional Imitation rule.

Result 5 *In line with Schlag's imitation models, probit estimations show that the probability with which a subject changes his action decreases in his own payoff and increases in the maximal observed payoff. Further, imitation of the best action becomes increasingly likely when the difference between own and maximal observed payoff increases.*

4.4 Questionnaire results

While the choice data we collected clearly show that many of our subjects behave *as if* they imitate, one cannot be sure whether subjects are aware

¹⁷We performed the same test for regression (2) and found that own payoff and max payoff enter independently.

Table 8: Multiple choice questions

Number of subjects influenced by...	Treatment		
	ROLE	GROUP	FULL
own past payoff(s)	37	34	32
payoffs of others in group	–	39	30
payoffs of others in role	33	–	19

Note: There were 54 subjects per treatment. All subjects chose at least one category, but multiple answers were possible.

of what they are doing and imitate intentionally. Thus, at the end of the experiment we asked subjects to fill in a computerized questionnaire. Apart from asking for their major field of studies,¹⁸ we asked subjects to explain in a few words how they made their decisions and to answer a multiple choice question regarding the variables they based their decisions on. In particular, we asked: “Please sketch in a few words how you arrived at your decisions!”. The multiple choice question asked: “Which of the following was of particular importance to your decision (multiple answers possible)? a) the results of your past decisions; b) the average payoff of all participants (only in FULL); c) the results of the participants who were randomly matched with you (only in FULL and GROUP); d) the results of the other participants playing the same role (only in FULL and ROLE).

Table 8 summarizes subjects’ responses to the multiple choice question. In all treatments own past payoffs were of importance to a majority of subjects and in all but treatment GROUP own payoffs were the most frequently named factor. More than 50% of subjects took also payoffs of other players into consideration. Interestingly, we again find that subjects are more interested in imitation when they can observe payoffs of their immediate competitors (compare Result 4 above).

Some of the free-format answers sketching the decision criteria employed are also quite instructive. To summarize them we have classified the answers into 8 main categories which are shown in Table 9 together with selected typical answers. Some subjects argued exactly as assumed by the various

¹⁸There are no significant effects with respect to the field of studies.

Table 9: Classification of questionnaire answers

classification	typical answer
role	“Answer with highest payoff of other players in previous round”
group	“When I had the highest payoff, kept the action for the next round. Otherwise switched to the action that brought the highest payoff. Sometimes had the impression that convergent actions of all players yielded lower payoffs.”
random	“by chance since all attempts of a strategy failed!”
differentiate	“tried to act anti-cyclically, i.e. not to do what the other <i>Z</i> -players have done” (in treatment <i>ROLE</i>)
pattern	“tried to find out whether an action yielded high payoffs in a particular order — but pattern remained unknown” “...proceeded according to the scheme: <i>ADBECADBEC</i> ...”
clever	“My impression of the rule was that low letters correspond to low numbers. The sum of payoffs seemed to be correlated with the sum of the letters but those with higher letters got more. I attempted to reach <i>AAA</i> but my co-players liked to play <i>E</i> ...”
own	“found out empirically where I got most points on average”

Note: These answers are typical because they are very descriptive of the categories not because they are typical for all answers in this category.

imitation theories (classifications “group” and “role”). But other subjects simply chose at random, tried to differentiate themselves from the behavior of others, or followed obscure patterns. There were also subjects who were clever enough to find out the payoff structure of the game (but were often in despair about their opponents’ play). Finally, some subjects reported to follow only their own past payoffs.

Table 10 lists by treatment the frequency of answers that fall into these 8 categories. Imitation of others in the same group is again a frequently cited motivation in both, *GROUP* and *FULL*, whereas role-imitation is less prevalent. Random behavior and own-payoff driven behavior is frequent in all treatments. But there are also types that like to differentiate themselves, types that believe in pattern or pattern recognition, and there are some clever types that guessed the payoff structure correctly.

The key finding in this section is

Table 10: Frequency of questionnaire answers

classification	Treatment		
	ROLE	GROUP	FULL
role	6	–	3
group	–	10	12
random	17	9	15
differentiate	5	2	5
pattern	2	–	6
clever	–	8	2
own	9	13	11

Note: A few answers were classified into two categories.

Result 7 *Subjects not only behave as if they imitate but many imitate intentionally. Other behaviors like random choices, pattern driven behavior, or behavior determined by own past payoffs can also be observed.*

5 Conclusion

In contrast to traditional theories of rational behavior, imitation is a behavioral rule with very “soft” assumptions on the rationality of agents. Imitation is typically modelled by assuming that subjects react to the set of actions and payoffs observed in the last period, by choosing an action that was evaluated as successful.

Recent theoretical results have increased the interest in economics about imitation. Of particular importance are the results due to Vega–Redondo (1997) and Schlag (1998). In Cournot games, the former predicts the Walrasian outcome while the latter predicts the Cournot-Nash equilibrium. In principle, these differences could be due to the different adjustment rules the models employ and/or the different informational conditions they assume. We study both rules in a generalized theoretical framework and show that the different predictions depend *exclusively* on the different informational assumptions.

We derive testable predictions for various classes of imitation rules (to which Vega-Redondo’s and Schlag’s belong) for discrete Cournot games and

test them in an experiment. More specifically, we study populations of nine players who are assigned to one of three different roles. In each period, players are randomly matched into three groups to play a simple symmetric Cournot game, such that always one player from each role is assigned to a group.

If agents only receive information about others with whom they interact, all rules that imitate successful actions imply the Walrasian outcome as the unique stochastically stable state. If agents only receive information about others who have the same role as they themselves but interact in other groups, Cournot-Nash play is the unique stochastically stable state. If agents have both types of information, the set of stochastically stable states depends on the specific form of the imitation rule. But, in general, stochastically stable states range from Cournot to Walrasian outcomes in such settings.

The experimental results provide clean evidence on both, the aggregate and individual level, for the relevance of imitation rules. Average outputs per treatment are exactly, and significantly, ordered as suggested by the generalized imitation model. Alternative learning models that are exclusively based on own past payoffs cannot account for this difference. Similar support for imitation models is found on the individual level by analyzing compliance rates for individual adjustments. Additionally, estimations of individual adjustment rules are in line with the basic principles of all imitation rules. In particular, we find support for Schlag's models that suggest that the likelihood of imitating another more successful action increases in the difference between own and other's payoff.

Finally, we observe that imitation of actions seems to be more prevalent when subjects observe others with whom they interact as opposed to others who have the same role but play in different groups. There is no theoretical model that would account for such a difference. Moreover, one might think that imitation of others who are identical to oneself is more meaningful than imitation of others with whom we play but who might be different. (After all, subjects in our experiment did not know that they were playing a symmetric game.) But this is not supported by the data. One conjecture that

might explain the difference we observe is that imitation of more successful actions might be particularly appealing when one directly competes with those who are more successful. In other words, there might be a link between imitation and the relevance of relative income. In environments where imitation prevents agents to do worse than their immediate competitors, there is an obvious “evolutionary” benefit from imitating. Thus, evolution might have primed us towards imitative behavior if we compete with others for the same resources. This would explain our data but more theoretical work is needed to study the evolutionary advantages and disadvantages of imitative behavior.

A Proofs

Proof of Proposition 1. First notice that if agents observe only strategies played in the own group, IBM coincides with IBA. By standard arguments (see e.g. Samuelson, 1994) only sets of states that are absorbing under the unperturbed ($\varepsilon = 0$) process can be stochastically stable. A straightforward generalization of Proposition 1 in Vega–Redondo (1997) shows that only uniform states can be absorbing (in all other states there is at least one agent who observes a strategy that fared better than his own), which is why we can restrict attention in the following to uniform states.¹⁹ We will show that ω^e can be reached with one mutation from any other uniform state $\omega^s \neq \omega^e$. The proof is then completed by showing that it requires at least two mutations to leave the Walrasian state.

Consider any uniform state $\omega^s \neq \omega^e$ and suppose that some player $(i, j)_t$ switches to the Walrasian strategy e . As a consequence $(i, j)_t$ will have the highest payoff in group j which will be observed by the other group members. By property (v) all players who were in group j at time t will play e in $t + 1$ with positive probability. Moreover, due to the random matching it is possible that the three players who were in group j at time t will be in three distinct groups in $t + 1$. In that case, each of them will

¹⁹Notice that the random rematching of agents into groups is crucial here. If group compositions were fixed, different groups could, of course, use different strategies.

achieve the highest payoff in their respective group which will be observed by their group members who then can also switch to the Walrasian strategy e , such that ω^e is reached. (If there are more than three groups, it will simply take a few periods more to reach ω^e .)

It remains to be shown that ω^e cannot be left with a single mutation. This is straightforward. In fact, it follows from exactly the same argument as in Vega–Redondo’s result. If a player switches to some strategy $s \neq e$, he will have the lowest payoff in his group and will therefore not be imitated. Moreover, he observes his group members who still play e and earn more than himself. Thus, he will switch back eventually. ■

Proof of Proposition 2. Although with reference groups as in treatment ROLE, IBM does not always coincide with IBA, we can use identical arguments for both rules to prove the claim. This is due to the fact, that we can establish the claim by restricting attention to one-shot mutations that do not induce different payoffs for any particular strategy an agent observes.

By a similar argument as above, only states, in which all role players in a given role receive the same payoff, can be candidates for stochastic stability. We will show that the Cournot state ω^c can be reached with a sequence of one-shot mutations from any other absorbing state. The proof will be completed by showing that it requires at least two mutations to leave ω^c . It is easy to see that every non-equilibrium state can be left with one mutation. One of the players who is currently not best replying, say (i, j) , must simply switch to his best reply. This will increase (i, j) ’s payoff which will also be observed by all other players in role i . Hence, in the next period all players in role i may have switched to their best replies against their opponents. Thus, for the first claim it remains to be shown that there exists for any state $\omega \neq \omega^c$ a sequence of (unilateral) best replies that leads into ω^c . This is easy to see by inspecting the payoff matrix, but follows more generally from the observation that the game has a potential (see Monderer and Shapley, 1996).

Now, consider ω^c and see what happens when a single player (i, j) switches to some other strategy. As he moves away from his best reply, he will earn less than the other agents in the same role i . As he can observe

these other agents, he will not be imitated and will eventually switch back. Thus, it is impossible to leave ω^c with one mutation which completes the proof. ■

Proof of Proposition 3. Note again that only uniform states can be candidates for stochastic stability. We will show that it takes one mutation to reach the set $\{\omega^c, \omega^d\}$ from any absorbing state not in this set while it takes two mutations to leave this set. Consider first a possible transition from ω^e to ω^c . With 1 mutation a transition to the state $\omega = (cee)(eee)(eee)$ is possible. The two e -players in group 1 observe two e -players (including themselves) that earn 400 and two others that earn 0, which is on average 200. But they also observe one c -player who gets 300. Thus, with positive probability in the next round all players in group 1 play c and one round later everyone plays c . We denote this possible transition in short as:

$$\omega^e \xrightarrow{1} (cee)(eee)(eee) \rightarrow (ccc)(eee)(eee) \rightarrow \omega^c,$$

where the number above the arrow denotes the required number of mutations.

It is easy to see that the following transitions from $x = a, b$ to $y = c, d$ require one mutation only,

$$\omega^x \xrightarrow{1} (yxx)(xxx)(xxx) \rightarrow (yxx)(yxx)(yxx) \rightarrow \omega^y$$

as well as the transition from ω^e to ω^d ,

$$\omega^e \xrightarrow{1} (dee)(eee)(eee) \rightarrow (ddd)(eee)(eee) \rightarrow \omega^d.$$

Any transition from a state ω^y , $y = c, d$ to some states ω^x , $x \neq y$, is impossible with one mutation as the process must return to ω^y

$$\omega^y \xrightarrow{1} (xyy)(yyy)(yyy) \rightarrow \omega^y.$$

Transitions from $\{\omega^c, \omega^d\}$ to ω^e require 2 mutations:

$$\begin{aligned} \omega^c &\xrightarrow{2} (ccc)(ccc)(aec) \rightarrow (cec)(cec)(aec) \rightarrow \omega^e \\ \omega^d &\xrightarrow{2} (ddd)(ddd)(ead) \rightarrow (edd)(edd)(eae) \rightarrow \omega^e. \end{aligned}$$

Transitions inside the set $\{\omega^c, \omega^d\}$ also require 2 mutations in both directions,

$$\begin{aligned}\omega^d &\xrightarrow{2} (ccd)(ddd)(ddd) \rightarrow (ccc)(ddd)(ddd) \rightarrow \omega^c \\ \omega^c &\xrightarrow{2} (ccc)(ccc)(adc) \rightarrow (cdc)(cdc)(adc) \rightarrow \omega^d.\end{aligned}$$

Thus, $\{\omega^c, \omega^d\}$ is the set of stochastically stable states. ■

Proof of Proposition 4. Again notice first that in treatment FULL a state is absorbing if and only if it is uniform. (Otherwise there are still some actions that will eventually be imitated.) We will first show that we can construct sequences of one-shot mutations that lead from any of the two “collusive” uniform states (where everybody plays a or everybody plays b) into one of the others (which we claim to be stochastically stable). Then we will show that it requires three simultaneous mutations to leave the more competitive states (where everybody plays c , everybody plays d , or everybody plays e).

The first step is easy. Consider one of the two collusive states and suppose that one agent, say (i, j) switches at time t to either c, d , or e . Clearly, this agent will have the highest overall payoff and can be imitated by everybody in $R(i, j)$. Now suppose that in $t + 1$ agent (i, j) will only be imitated by agents who are also in role i but not by those in his group (due to inertia). Then each group in $t + 1$ will have one player with a competitive strategy and two with collusive strategies (regardless of the matching). The highest payoffs are, of course, obtained by those who now play the more competitive strategy and everybody can observe at least one of these agents. Hence, in $t + 2$ everybody will play the competitive strategy.

Next we show that it is not possible to leave one of the competitive states with a single mutation. Take, for example, the Walrasian state, ω^e , and suppose that one agent (i, j) switches at some time t to some strategy other than e . This will have two consequences: (i, j) will earn less than the other agents in group j but more than the other agents in role i . Now suppose that the other agents in role i imitate (i, j) in $t + 1$, but that (i, j) himself, does not immediately switch back to e (due to inertia). Then in $t + 1$

all players in role i will play the same strategy other than e while everybody else will still play e . Clearly, the latter earn more than the former such that now everybody can revert to playing e .

The same argument applies to states where everybody plays d or everybody plays c . Moreover, a similar argument applies for the case of two simultaneous mutations. (Again inertia can be used to compose identical strategy profiles in all groups after the mutations and the first round of imitation.) The proof is completed by the observation that any uniform state can be reached from any other uniform state by exactly three simultaneous mutations. For movements from less to more competitive states we can make such a transition if all players who have the same role i simultaneously switch to higher quantities. For reverse movements from more to less competitive states we can construct the transition if all players in the same group j simultaneously switch to lower quantities.²⁰ This completes the proof. ■

Proof of Proposition 6.

(a) Only uniform states can be imitation equilibria, otherwise there would be an imitation opportunity. To see that ω^e is an imitation equilibrium note that if (i, j) deviates from ω^e will experience lower payoffs than any other player; nobody follows and (i, j) returns to e . To see that any other uniform state is not an imitation equilibrium consider the deviation of (i, j) to the immediate higher production level. This creates an imitation opportunity to players in group j . By random matching this deviation may spread out the whole population, in which case a destination is reached. At the destination the payoffs of (i, j) are lower than at the original distribution. Player (i, j) returns to the original action. Now players in group j have higher payoffs than (i, j) , do not imitate him, and (i, j) has an imitation opportunity to go back to the deviation strategy.

(b) If (i, j) deviates from ω^c , he will get lower payoffs than players in role i . Nobody follows the deviation, and (i, j) returns to c . This shows that ω^c

²⁰Hence, a generalization of our statement for arbitrary numbers of groups and arbitrary group sizes is not possible. The set of stochastically stable states will, in general, depend on whether there are more roles or more groups.

is an imitation equilibrium. It is easy to show that any state other than ω^c where members of the same role play the same action, but where differences between roles are not excluded, is not an imitation equilibrium. Note then that there is a (i, j) that is not best-replying, then a deviation of (i, j) to his best-reply gives to him higher payoffs, creating an imitation opportunity to players in role i . At this destination (i, j) has higher payoffs than at the original state, and hence does not return to the original action. It remains to be shown that a state where at least one role whose members play different actions is not an imitation equilibrium. If in such a case, in any random matching any player has an imitation opportunity, then the assertion holds. Assume the opposite, then since there are not two different best-replies that give the same payoffs, at least one player is not best-replying, and hence the above argument shows that such a state is not an imitation equilibrium.

(c) To show in FULL that non-uniform states are not imitation equilibria is tedious, and hence we concentrate on uniform states. We first show that ω^c is an imitation equilibrium. At ω^c let (i, j) deviate to $s_i^j \neq c$. Then players in role i will have higher payoffs than (i, j) and players in group j will observe that those players in their respective role have higher payoffs than (i, j) . Hence, nobody follows. Then, (i, j) observes that c gives higher payoffs to players in role i and hence returns to c .

Now we show that ω^e is an imitation equilibrium. At ω^e let (i, j) deviate to $s_i^j \neq e$. In $t + 1$ players in role i will follow since will have lower payoffs than (i, j) and will observe that their respective group players also have lower payoffs than (i, j) , but players in group j will not follow since will have higher payoffs than (i, j) . In $t + 2$ all players in role i including (i, j) will imitate their respective group players and hence ω^e is reached.

We now show that ω^d is an imitation equilibrium. If at ω^d (i, j) deviates to $s_i^j \in \{a, b, c\}$, then a deviation path that returns to ω^d , analogous to the one analyzed for the case of ω^e , is generated. If at ω^d , (i, j) deviates to $s_i^j = e$, then a deviation path that returns to ω^d , analogous to the one analyzed for the case of ω^c , is generated.

To show that ω^a and ω^b are not imitation equilibria it is enough to show that there exists a sequence of random matchings that makes that the

imitation paths do not return to the original state. Let $x = a, b$ and $y = b$ if $x = a$ and $y = c$ if $x = b$. Then, one can check that the following path can be generated: $\omega^x \rightarrow (yxx)(xxx)(xxx) \rightarrow (yyx)(yxy)(yxx) \rightarrow \omega^y \rightarrow (xyy)(yyy)(yyy) \rightarrow \omega^y$. ■

B Instructions

Welcome to our experiment! Please read these instructions carefully. Do not talk with the person sitting next to you and remain quiet during the entire experiment. If you have any questions please ask us. We will come to you.

During this experiment, which takes 60 rounds, you will be able to earn points in every round. The number of points you are able to earn depends on your actions and the actions of the other participants. The rules are very easy. At the end of the experiment the points will be converted to Euros at a rate of 3000:1.

Always 9 of the present participants will be evenly divided into three roles. There are the roles X, Y, Z , taken in always by 3 participants. The computer randomly allocates the roles at the beginning of the experiment. You will keep your role for the course of the entire experiment.

In every round every X -participant will be randomly matched by the computer with one Y - and one Z -participant. After this, you will have to choose one of five different actions, actions A, B, C, D , and E . We are not going to tell you, how your payoff is calculated, but in every round your payoff depends uniquely on your own decision and the decision of the two participants you are matched with. The rule underlying the calculation of the payoff is the same in all 60 rounds.

After every round you get to know how many points you earned with your action and your cumulative points.

In addition, you will receive the following information:

You get to know which actions the other two participants who have the same role as you (and who were matched with different participants) have chosen, and how many points each of them earned.

You get to know which actions the other two participants you were matched with have chosen, and how many points each of them earned.

Furthermore you get to know how many points all 9 participants (in all the 3 roles) on average earned in this round.

Those are all the rules. Should you have any questions, please ask now. Otherwise have fun in the next 60 rounds.

References

- [1] Abbink, K., and Brandts, J. (2002), “24”, University of Nottingham and IAE, mimeo.
- [2] Alos-Ferrer, C. (2001), “Cournot vs. Walras in Dynamic Oligopolies with Memory”, Working Paper 0110, Department of Economics, University of Vienna.
- [3] Arrow, K.J., and Hurwicz, L. (1960), “Stability of the Gradient Process in n -Person Games”, *Journal of the Society of Industrial and Applied Mathematics*, 8, 280-294.
- [4] Asch, S. (1952), *Social Psychology*, Englewood Cliffs: Prentice Hall.
- [5] Bergin, J. and Bernhardt, D. (2001), “Imitative Learning”, Queen’s University, Canada, mimeo.
- [6] Björnerstedt, J. and Weibull, J.W. (1996), “Nash Equilibrium and Evolution by Imitation”, in K. Arrow et al. (eds.), *The Rational Foundations of Economic Behaviour*, London: Macmillan, 155-171.
- [7] Canning, D. (1992), “Average Behavior in Learning Models”, *Journal of Economic Theory*, 57, 442-472.
- [8] Ellison, G., and Fudenberg, D. (1995), “Word of Mouth Communication and Social Learning”, *Quarterly Journal of Economics*, 110, 93-126.
- [9] Fischbacher, U. (1999), “z-Tree. Zürich Toolbox for Readymade Economic Experiments”, University of Zürich, Working Paper no. 21.

- [10] Fudenberg, D., and Levine, D. (1998), *The Theory of Learning in Games*, Cambridge: MIT Press.
- [11] Holt, C.A. (1995), "Industrial Organization: A Survey of Laboratory Research", in: John Kagel and Alvin Roth (eds.): *The Handbook of Experimental Economics*, Princeton, Princeton University Press.
- [12] Huck, S., Normann, H.T., and Oechssler, J. (1999), "Learning in Cournot Oligopoly: An Experiment", *Economic Journal*, 109, C80-C95.
- [13] Huck, S., Normann, H.T., and Oechssler, J. (1999), "Does Information about Competitors' Actions Increase or Decrease Competition in Experimental Oligopoly Markets?", *International Journal of Industrial Organization*, 18, 39-57.
- [14] Huck, S., Normann, H.T., and Oechssler, J. (2004), "Two are Few and Four are Many: Number Effects in Experimental Oligopoly", *Journal of Economic Behavior and Organization*, 53, 435-446.
- [15] Kandori, M., Mailath, G. and Rob, R. (1993), "Learning, Mutation, and Long Run Equilibria in Games", *Econometrica*, 61, 29-56.
- [16] Monderer, D. and Shapley, L. (1996), "Potential Games", *Games and Economic Behavior*, 14, 124-143.
- [17] Offerman, T., Potters, J., and Sonnemans, J. (2002), "Imitation and Belief Learning in an Oligopoly Experiment", *Review of Economic Studies*, 69, 973-997.
- [18] Plott, C.R. (1989), "An Updated Review of Industrial Organization: Applications of Experimental Economics", in: R. Schmalensee and R.D. Willig (eds.): *Handbook of Industrial Organization*, vol. II, Amsterdam, North Holland.
- [19] Samuelson, L. (1994), "Stochastic Stability with Alternative Best Replies", *Journal of Economic Theory* 64, 35-65.

- [20] Schlag, K. (1998), “Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits”, *Journal of Economic Theory* 78, 130-56.
- [21] Schlag, K. (1999), “Which One Should I Imitate?”, *Journal of Mathematical Economics*, 31, 493-522.
- [22] Selten, R., and Apesteguia, J. (2002), “Experimentally Observed Imitation and Cooperation in Price Competition on the Circle”, University of Bonn, mimeo.
- [23] Selten, R. and Ostmann, A. (2001), “Imitation Equilibrium”, *Homo Oeconomicus*, 43, 111-149.
- [24] Siegel, S. and N. Castellan, J. Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences*, Singapore: McGraw-Hill.
- [25] Vega-Redondo, F. (1997), “The Evolution of Walrasian Behavior”, *Econometrica*, 65, 375-384.
- [26] Young, H.P. (1993), “The Evolution of Conventions”, *Econometrica*, 61, 57-84.
- [27] Young, H.P. (1998), *Individual and Social Structure*, Princeton: Princeton University Press.