

Imitation

– Theory and Experimental Evidence –*

Jose Apestegua[†]

Public University of Navarre

Steffen Huck[‡]

University College London & ELSE

Jörg Oechssler[§]

University of Bonn

April 14, 2004

Abstract

We introduce a generalized theoretical approach to study imitation and subject it 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 significant effects of seemingly innocent changes in information. Moreover, the generalized imitation model predicts the differences between treatments well. The data provide support for imitation on the individual level, both in terms of

*We thank seminar audiences at Texas A&M, UC Berkeley, UC San Diego, UC Santa Barbara, University College London, Basque Country University, Simon Fraser University, the Stockholm School of Economics, the joint theory seminar of WZB, Humboldt University, and Free University Berlin, and at the Universities of Bonn, Bristol, Edinburgh, Munich, and Stockholm as well as participants at the ESA Meetings in Tucson for discussions. Gary Charness, Dan Friedman, Heike Harngart, Karl Schlag, and Nat Wilcox provided valuable comments. We are grateful to Peter Urban for research assistance. 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 (FMRX-CT98-0238) is gratefully acknowledged.

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

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.

Keywords: Evolutionary game theory; Stochastic stability; Imitation; Cournot markets; Information; Experiments; Simulations.

1 Introduction

Everyone who watches children growing up will attest that imitation is a main source of learning. And introspection shows that imitation also plays a significant role for adult learning. In fact, imitation is prevalent in much of everyday decision making, in particular when the environment is complex or largely unknown. Openings in chess games are a good example of finding routes through traffic, or buying complex consumer items like cars, laptops, or digital cameras. But, 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 towards understanding imitation 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. 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, while Vega-Redondo's model predicts the Walrasian outcome. The latter prediction is also obtained by Selten and Ostmann's (2001) notion of an 'imitation equilibrium', while Cournot-Nash is also predicted by imitation models with large population as studied by Björnerstedt and Weibull (1996).

The current paper makes two main contributions. First, it introduces 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. Basically, the models differ along two different dimensions, the informational structure ("whom agents imitate") and the behavioral rule ("how agents imitate"). While agents in Vega-Redondo's model observe their immediate competitors, in Schlag's model they observe others who are just like them but play in different groups against different opponents. Additionally, agents in Vega-Redondo's model copy the most successful action of the previous period whenever they can. In contrast, Schlag's agents only imitate in a probabilistic fashion and the probability

with which they imitate is proportional to the observed difference in payoffs between own and most successful action. We show that the difference between the two models is due to the different informational assumptions rather than the different adjustment rules. In that sense, 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, 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 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.¹ The key design feature of our experiment is that we vary the feedback information subjects receive between rounds of play. In one treatment they observe their competitors' actions and profits, in another they observe the actions and profits of others who are like them but played against different people. And, finally, there is a treatment where agents have access to both types of information.

On some level, these variations appear to be very innocent and many (learning) models would not predict any difference between them. In that sense, the experimental part of our study examines whether (and if so how) slight variations in the informational structure of a repeated-game setting have an impact on behavior. We find that the variations indeed have significant effects. Moreover, the directions of these effects are well organized by our generalized imitation model. Specifically, average outputs are ranked according to the theoretical predictions and significantly so: The treatment

¹See e.g. Plott (1989), Holt (1995), and Huck, Normann, and Oechssler (2004) for surveys.

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

Analyzing individual adjustments, we find strong support for imitative behavior. Simple imitation can explain a surprisingly large fraction of subjects’ decisions. But subjects differ in their propensity to imitate. While some imitate not more often than a randomization device would, others are almost pure imitators. In general though, 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. In addition, we find that imitation is more pronounced when subjects observe their direct competitors—rather than others who have the same role but play in different groups.

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. Quite a number of subjects perceive themselves as imitating.

Despite being inherently “behavioral”, there have been few prior experiments on imitation. 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) find experimental support for 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 remainder of the paper is organized as follows. Section 2 introduces the games and the experimental details. In Section 3 we review the imita-

tion 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. Appendix B contains a treatment of Selten and Ostmann’s imitation equilibrium. The instructions for the experiment are shown in Appendix C and Appendix D contains additional regression results.

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 C).

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. One might wonder why we introduce roles to study behavior in a symmetric game. The answer is twofold. First, it is exactly this “trick” that 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.

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 treat-

ment. 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 and learn the average payoff in the entire population.²

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 Cournot output, and e to the symmetric Walrasian output, where price equals marginal cost (of zero). 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 3000:1.³

The computerized experiments⁴ were carried out 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.⁵ 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.

²Notice that in FULL a player cannot observe the choices and payoffs of players that are neither in his group nor in his role.

³In the first session of treatment FULL we used an exchange rate of 4000:1.

⁴The program was written with z-tree of Fischbacher (1999).

⁵At the time of the experiment one Euro was worth about one US dollar.

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.

3 Imitation models

3.1 Theory

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 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 and 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.⁶ 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* 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.⁷

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, each is chosen with positive probability.*

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

⁶This is always the case in treatment GROUP.

⁷For “imitate the best average” see, e.g., Ellison and Fudenberg (1995) and Schlag (1999). For “imitate the best max” see Selten and Ostmann (2001).

- “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 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 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).

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').

Both, WIBA and WIBM allow for a large variety of specific adjustment rules, including Vega-Redondo’s imitate the best rule as well as forms of probabilistic adjustment as considered by Schlag. In the following, we will state all results for these rather large *classes* of rules. Hence, it is in this sense

that we will conclude that the informational structure (whom to imitate) is more important than the specific rule (how to imitate).

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” (or, in short, IBM), 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 analogous 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.⁸ Insofar, Proposition 1 can be seen as generalization of Vega-Redondo’s original result to the case where agents might be randomly rematched. As long as the informational structure is such that agents observe only their competitors (in the last period) the Walrasian outcome results.

Let us now turn to treatment ROLE where $(h, k)_t \in R(i, j)_t$ 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 stochastically stable outcome under both, WIBA rule and WIBM rules, 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

⁸Introducing 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.

same role will observe this but will not imitate because they earn more using the equilibrium strategy. On the other hand, every non-equilibrium state can be left by a single mutation, namely by having an agent who is currently not playing his best reply switch to his best reply. This improves his payoff and will be observed by other agents in the same role who will follow suit. What remains to be shown is that one can construct sequences of one-shot mutations that lead into the Cournot state from any other state. To establish this claim we use the fact that the game at hand has a potential.

Comparing Propositions 1 and 2 establishes our earlier claim. While the specifics of an imitation rule do not matter as long as the rule falls in the rather large class of WIBA and WIBM rules, changing the informational structure has a profound effect on long-run behavior. Turning to treatment FULL one might expect that its richer informational structure (with agents having 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.

The proof of Propositions 3 shows that specifics of the payoff function matter for the exact prediction under WIBA. A generalization for a larger class of payoff functions would predict outcomes ranging from the Cournot to some more competitive outcomes (without exactly specifying the boundary). On the other hand, the proofs for Propositions, 1, 2, and 4 do not make use of anything that is specific to our chosen payoff function and it is easy to see that they could be generalized to a large class of Cournot games in exactly the same form as above.

Finally, 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. ¥

In the appendix we also analyze the predictions of Selten and Ostmann's (2001) *imitation equilibrium*. Interestingly, it turns out that, despite its static character, it makes the same predictions about behavior in the long run as the class of dynamic WIBM rules.

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.

Table 2: Summary of predictions

Imitation Rule	never imitate worse than own*	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	–	$\omega^a, \omega^b, \omega^c, \omega^d, \omega^e$

Note: A “X” 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. **This is the set of stochastically stable outcomes.

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 are currently using. With respect to total outputs in the Cournot market games, all imitation rules, except SPOR, suggest that quantities 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. Finally, the theoretical results suggest for treatment FULL quantities between GROUP and ROLE. Thus, we obtain the following qualitative hypothesis about the ordering of output levels:

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

Hypothesis \mathcal{QH} has, strictly speaking two parts. First, it suggests that there is a difference between the experimental treatments (what many other theories would not predict). Second, it suggests a particular order that would be expected if imitation is an important force for subjects’ adaptations.

The problem with long run predictions derived from stochastic stability analysis is that they are just that: long run predictions. Furthermore, in




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

general they crucially depend on the assumption of vanishing noise. Thus, the issue arises how imitation processes behave in the short run and in the presence of non-vanishing noise. In order to address this issue, we run simulations for the different treatments. In particular, we simulate 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. The noise level we use is substantial: with probability 0.8 in each round a player follows IBM and with probability 0.2 a player chooses randomly one of the five actions (each then with equal probability). For each treatment we simulate 100 such populations with starting actions chosen from a uniform distribution.

Figure 1 shows average quantities in these simulations organized in blocks of ten periods each.⁹ 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, differences being significant at any conventional significance level. Therefore, importantly, the predictions we obtained for the long run and with vanishing noise appear rather robust also for the short run and in the presence of noise.

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} , we evaluate the data on the aggregate level. This will show whether and, if so, how the different informational structures affect subjects' behavior. While this will provide some indirect evidence for the relevance of imitation, a more thorough study of imitation must be based on data from individual adjustments. Thus, in Section 4.2 we analyze individual data by counting how often actual adjustments are in line with predicted adjustments. This is followed in Section 4.3 by a regression analysis that helps us to test whether the probability of imitating is indeed, as Schlag's models suggest, a function of the observed payoff differences. Finally, we conclude this section by analyzing the post-experimental questionnaire. This will provide additional insight into 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 first round, all 60 rounds of the experiment, and the last 30 rounds. Standard deviations of the six observations per treatment are shown in parentheses.

⁹Recall that actions a, b, c, d , and e correspond to output quantities 20, 23, 30, and 40, respectively.

Table 3: Summary statistics

	Treatment		
	ROLE	GROUP	FULL
avg. output, round 1	27.48 (3.56)	26.72 (3.10)	26.70 (3.20)
avg. output, rounds 1-60	30.32 (.452)	32.96 (.825)	31.71 (.829)
avg. output, rounds 31-60	30.71 (.514)	33.57 (.972)	32.31 (.881)
avg. var., rounds 1-60	54.43	49.78	51.31
avg. var., rounds 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 parentheses.

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 the seemingly innocent changes in information conditions have a systematic impact on behavior. While there is no significant difference between average quantities in the first round, the differences in quantities over all 60 and the last 30 rounds are highly significant.¹⁰ Moreover, the ordering is as predicted by the generalized imitation model. 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 previous section. However, the quantitative differences in average quantities are less

¹⁰Also, 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 ($r_s = 0.23$).




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

pronounced than predicted by imitation theory, though, which indicates that other factors play a role, too.

For now, we summarize our findings in the following two statements.

Result 1 The reference group has a significant impact on behavior.

Result 2 Outputs are ordered as predicted by hypothesis \mathcal{QH} .

Given the usual noise in experimental data from human subjects, Result 2 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 played last period’s best evaluated action in his reference group, (ii) ‘Better’: the subject switched to an action that was evaluated as better than his own action, but not as the best, (iii) ‘Same’: the subject did not change his action despite observing a better strategy in his reference group, (iv) ‘Worse’: the subject changed to an action that was evaluated as worse than his own action, and (v) ‘Different’: the subject changed to an action that was not observed in the reference group. Table 4 reports how many decisions fall into each of the categories (i) through (v) for each treatment and both evaluation rules. The differences between the max and the average rules are very small which 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. Hence, for ease of presentation we will focus on the max rule from now on.

There are a couple of observations which are immediate from inspecting Table 4:

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.

- There is very little switching to worse or better (but not best) actions. Most subjects either repeat their previous choice, imitate the most successful action, or experiment by switching to a new action.
- Imitation of the previously most successful action is most prevalent in treatment GROUP.

Recall that WIBA and WIBM predict that agents should not switch to actions evaluated as worse than the own action in the previous round. Table 4 shows that pooled over all treatments only 8.3% of choices violate this condition. To put this rate into perspective, we need a method that contrasts it with the corresponding rate 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 hypothesis relative to this simulated data. In order to give random behavior the best shot, we take the experimentally observed frequencies of actions as the theoretical distribution from which random behavior is generated. The simulations show that random behavior would violate the “never imitate worse than own” condition in 16% of cases, which is significantly higher than the actual rate at all standard significance levels according to a permutation test.

Result 3 *On average, the “never imitate worse than own” condition is vi-*

olated in only 8.3% of cases which significantly outperforms random predictions.

Another interesting way of slicing through the data shown in Table 4 is to compute how often subjects are in line with the predictions of a simple imitation rule like IBM. We classify behavior as compliant with IBM if either the best action was imitated or there was no change in action (inertia). Thus, by summing the values obtained for ‘Best’ and ‘Same’ in Table 4 we find a compliance rate of 58.3% pooled over all treatments. Given that there are many non-imitating choices, it is not surprising that this rate is not terribly high, although it is significantly higher than under random play, which would yield a compliance rate of 34.6% (using the method described above).¹¹ This further confirms that imitation is present in our data, and that, in particular, IBM and IBA play a significant role in explaining it. One can also compute a compliance rate for IBM given that subjects play an action they have previously observed.¹² In ROLE this yields a compliance rate of 82.9%, in GROUP 88.6%, and in FULL 75.5%. These rates are very high and indicate that *when* players imitate, they mostly imitate the best.

So far, we have only examined averages across subjects. But, as one would expect, there is substantial heterogeneity in subjects’ propensity to imitate. Figure 3 shows the distribution of individual players on the basis of the (unconditional) compliance rates for IBM (for all treatments pooled together). About 10% of the players show a percentage of unconditional compliance with IBM above 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 comply with IBM in more than 60% of all decisions.

Let us summarize this by stating a further result.

Result 4 IBM and IBA do about equally well, and both outperform random predictions significantly. Moreover, 10% of subjects are almost pure

¹¹Permutation tests on the basis of the average rates of compliance for the populations show that IBM outperforms random predictions at any conventional significance level.

¹²By dividing the sum of “Best” and “Same” through (100 minus “Different”).




Figure 3: Distribution of individual players on the basis of the compliance rates with IBM, all treatments pooled.

imitators whose choices are in line with IBM/IBA in more than 80% of all decisions.

Finally, let us briefly discuss the second observation we made after inspecting Table 4. There is more compliance with IBM (or IBA) in treatment GROUP than in ROLE. A permutation test yields significance at the 5% level (two-sided).¹³ 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 cannot expect to face the same opponents as last period.

Result 5 *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).*

¹³All other pairwise comparisons are not statistically significant.

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 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 observed payoff.¹⁴

Table 5 shows regression results for the first question—what makes subject change their strategy. The first column for each treatment shows estimations for a simple linear probability model with random effects:

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

where s_i^t denotes subject i ’s strategy in period t , π_i^t the subject’s own payoff, $\pi_{i\max}^t$ the maximal payoff the subject observed in his reference group, while v_i is the subject-specific random effect, and ε_i^t is the residual. Note that we include π_i^t directly and also in form of the payoff difference between max payoff and own payoff. This allows to test whether only the difference matters, as predicted e.g. by Schlag’s Proportional Imitation rule, or whether own payoff and maximal payoff enter independently. If β is not significantly different from zero, then only the payoff difference matters.

As a robustness check Table 5 also shows estimation results for a model that includes an additional term borrowed from the reinforcement learning literature. Reinforcement learning could be seen as the main rival to imitation in our experiment where subjects know very little about their environment. But including a term capturing an element of reinforcement learning is here not so much a step toward a more complete model of what our subjects really do but rather a check whether imitation remains a significant force when one allows for other ways of learning. As in the basic model of Erev and Roth (1998) the propensity of a strategy is simply the sum of

¹⁴Due to the high correlation of the best max and the best average, results for IBA are very similar and, therefore omitted.

Table 5: Estimating the likelihood that subjects change their action

	ROLE		GROUP		FULL	
constant	886*** (42.6)	997*** (40.4)	579*** (26.9)	730*** (26.4)	611*** (44.1)	756*** (37.3)
own payoff	-.316*** (.033)	-.289*** (.033)	-.197*** (.024)	-.164*** (.024)	-.121*** (.029)	-.077*** (.029)
payoff diff.	.098*** (.035)	.100*** (.034)	.476*** (.043)	.454*** (.043)	.211*** (.032)	-.208*** (.031)
relative propensity	—	-387*** (37.5)	—	-418*** (37.6)	—	-467*** (36.5)
R^2	.075	.131	.077	.146	.042	.174
# of obs.	3186	3186	3186	3186	3186	3186

Note: All coefficients and standard errors multiplied by 10^3 . Standard errors in parentheses. ***denotes significance at the 1% level, ** denotes significance at the 5% level.

all past payoffs a player obtained with that strategy. The *relative propensity* is the propensity of a strategy divided by the sum of the propensities of all strategies. The regressions in Table 5 include the relative propensity of the currently used strategy. Thus, the expected sign of the coefficient is negative. For further robustness checks Appendix D shows that the results for all regressions are essentially the same for linear fixed-effects models and random-effects probit models.

The regressions consistently show that the coefficients for own payoffs are significantly negative while those for the observed payoff difference between own and best strategy are significantly positive, which is in line with the theoretical prediction. This holds for all treatments and the coefficients have the same order of magnitude. However, confirming what we have seen in other parts of the data analysis, the coefficients are largest in treatment GROUP. Moreover, the estimated coefficients turn out to be very robust to the inclusion of the propensity term, which is significant and has the expected sign in all treatments. Thus, reinforcement learning seems to be a factor and it helps to improve the explanation of the observed variance. Nevertheless, the inclusion of the propensity term does not diminish the significance of the variables related to imitation.

Table 6: Estimating the likelihood that subjects follow IBM

	ROLE		GROUP		FULL	
constant	127*** (41.2)	146*** (41.9)	145*** (22.5)	113*** (25.4)	164*** (43.6)	166*** (45.3)
own payoff	-.001 (.038)	.004 (.038)	-.043 (.030)	-.058* (.030)	.056 (.038)	.056 (.039)
payoff diff.	.248*** (.038)	.246*** (.038)	.551*** (.045)	.586*** (.047)	.156*** (.040)	.156*** (.040)
relative propensity	—	-90.1* (4.77)	—	131*** (49.1)	—	-13.4 (61.6)
R^2	.038	.038	.080	.087	.009	.009
# of obs.	2079	2079	1644	1644	1920	1920

Note: All coefficients and standard errors multiplied by 10^3 . Standard errors in parentheses. Only cases with $s_i^{t+1} \neq s_i^t$ included. *** denotes significance at the 1% level, * significance at the 10% level.

After analyzing *when* subjects switch to a different action, we shall now analyze what makes them switch to the action with the best payoffs *if* they switch at all. Table 6 reports subjects' likelihood of following IBM (contingent on switching to another action)¹⁵ as a function of their own payoff and the observed payoff difference. As before the estimation results shown here are for linear probability models with random effects. Appendix D contains fixed effects and probit models. The first column for each treatment shows results for:

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

where $s_{i \max}^t$ is the action that had the highest maximal payoff (IBM) in period t in subject i 's reference group and all other variables are as defined before. The second column shows, as before, estimation results for a model that includes a propensity variable, this time the propensity of the action with the highest observed payoff (and thus, the expected sign of the coefficient is positive).

Table 6 shows that, as Schlag's models suggest, for IBM only the payoff difference matters. In all three treatments the coefficient of the difference

¹⁵Since the theories allow for inertia, not switching is always in line with the prediction.

variable has the expected sign and is significant at the 1% level. In contrast, the coefficient of own payoff 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. Moreover, the results are, as before, robust to the inclusion of the propensity term (although this time the propensity term does not improve the explanation of the observed variance, has an unexpected sign in treatment ROLE, and fails to be significant in treatment FULL).

We briefly summarize in

Result 6 *In line with Schlag’s imitation models, 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, the probability of imitating the best action is driven mainly by the difference between maximal observed and own payoff.*

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 of what they are doing and imitate intentionally. But we have additional evidence in form of a post-experimental questionnaire. Apart from asking for their major field of studies,¹⁶ 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.” In addition, there was a multiple choice question about which variables had impact on their decisions. Table 7 summarizes subjects’ responses to this 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

¹⁶There are no significant effects with respect to the field of studies.

Table 7: 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.

find that subjects are more interested in imitation when they can observe payoffs of their immediate competitors (compare Result 5 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 seven main categories which are shown in Table 8 together with selected typical answers. Some subjects argued exactly as assumed by the various 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 9 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 subsection is

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

Table 8: 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!”
contrarian	“tried to act anti-cyclically, i.e. not to do what the other Z-players have done” (in treatment ROLE)
pattern	“tried to find out whether an action yielded high payoffs in a particular order — but pattern remained unknown” “...proceeded according to the scheme: ADBECADBEC...”
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 AAA but my co-players liked to play E...”
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.

Table 9: Frequency of questionnaire answers

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

Note: A few answers were classified into two categories.

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 economists’ interest in imitation. Of particular importance are results due to Vega–Redondo (1997) and Schlag (1998). Remarkably, the models make quite different predictions in many games, most notably in Cournot games, where 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 mainly depend on the different informational assumptions.

Comparatively slight changes in feedback information are, thus, predicted to affect behavior. Behavior is predicted to be more competitive if agents observe their immediate rivals than if they observe others who play in different groups against different opponents. From the vantage point of many other (learning) theories these differences appear surprising. Yet, in an experiment we provide clear evidence for the relevance of the information structure.

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 that changing feedback

in this manner significantly alters behavior. Learning models that do not take into account the observation of others' payoffs cannot explain this effect. Moreover, the differences between treatments are ordered as the generalized imitation model suggests. Direct support for the role of imitation is found by analyzing individual adjustments. We find that imitation can explain a substantial number of adjustments and that some subjects are almost pure imitators. Moreover, estimating subjects's choice functions we find support for Schlag's result that suggests that the likelihood of imitating a 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 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.

References

- [1] Abbink, K., and Brandts, J. (2002), "24", University of Nottingham and IAE, mimeo.
- [2] Alos-Ferrer, C. (2004), "Cournot vs. Walras in Dynamic Oligopolies with Memory", *International Journal of Industrial Organization*, 22, 193-217.

- [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] Erey, I. and Roth, A. (1998). “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria”, *American Economic Review*, 88, 848-881.
- [10] Fischbacher, U. (1999), “z-Tree. Zürich Toolbox for Readymade Economic Experiments”, University of Zürich, Working Paper no. 21.
- [11] Fudenberg, D., and Levine, D. (1998), *The Theory of Learning in Games*, Cambridge: MIT Press.
- [12] 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.
- [13] Huck, S., Normann, H.T., and Oechssler, J. (1999), “Learning in Cournot Oligopoly: An Experiment”, *Economic Journal*, 109, C80-C95.

- [14] Huck, S., Normann, H.T., and Oechssler, J. (2000), “Does Information about Competitors’ Actions Increase or Decrease Competition in Experimental Oligopoly Markets?”, *International Journal of Industrial Organization*, 18, 39-57.
- [15] 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.
- [16] Kandori, M., Mailath, G. and Rob, R. (1993), “Learning, Mutation, and Long Run Equilibria in Games”, *Econometrica*, 61, 29-56.
- [17] Monderer, D. and Shapley, L. (1996), “Potential Games”, *Games and Economic Behavior*, 14, 124-143.
- [18] Offerman, T., Potters, J., and Sonnemans, J. (2002), “Imitation and Belief Learning in an Oligopoly Experiment”, *Review of Economic Studies*, 69, 973-997.
- [19] 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.
- [20] Samuelson, L. (1994), “Stochastic Stability with Alternative Best Replies”, *Journal of Economic Theory* 64, 35-65.
- [21] Schlag, K. (1998), “Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits”, *Journal of Economic Theory* 78, 130-56.
- [22] Schlag, K. (1999), “Which One Should I Imitate?”, *Journal of Mathematical Economics*, 31, 493-522.
- [23] Selten, R., and Apesteguia, J. (2002), “Experimentally Observed Imitation and Cooperation in Price Competition on the Circle”, University of Bonn, mimeo.

- [24] Selten, R. and Ostmann, A. (2001), “Imitation Equilibrium”, *Homo Oeconomicus*, 43, 111-149.
- [25] Siegel, S. and N. Castellan, J. Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences*, Singapore: McGraw-Hill.
- [26] Vega-Redondo, F. (1997), “The Evolution of Walrasian Behavior”, *Econometrica*, 65, 375-384.
- [27] Young, H.P. (1993), “The Evolution of Conventions”, *Econometrica*, 61, 57-84.
- [28] Young, H.P. (1998), *Individual and Social Structure*, Princeton: Princeton University Press.

Appendix

A Proofs

Proof of Proposition 1. First notice that if agents observe only strategies played in the own group, the max and average evaluation rules coincide. 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.¹⁷ 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 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

¹⁷Notice that the random rematching of agents into groups is crucial here. If group compositions were fixed, different groups could, of course, use different strategies.

than himself. Thus, he will switch back eventually. \nexists

Proof of Proposition 2. Although with reference groups as in treatment ROLE, the max and average evaluation rules do not coincide, 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. \nexists

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. \forall

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.¹⁸ This completes the proof. \forall

B Imitation Equilibrium

We shall briefly 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) .¹⁹ 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.

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.

¹⁸Hence, 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.

¹⁹This requirement is the same as in IBM.

(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.

The following 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.

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 (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 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: $\omega^x \rightarrow (yxx)(xxx)(xxx) \rightarrow (yyx)(yxy)(yxx) \rightarrow \omega^y \rightarrow (xyy)(yyy)(yyy) \rightarrow \omega^y$. \nexists

C 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:

[In ROLE and FULL] 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.

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

[In FULL] 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.

D Regressions

In this appendix we show all regression results for models (1) and (2). Tables 10, 11, and 12 show the results for what makes subjects switch to another strategy (model 1). Table 10 contains the estimations for treatment GROUP, Table 11 for ROLE, and Table 12 for FULL. The first two columns in each table show the results from the linear random effects model, also shown in the main body of the paper. The third and fourth columns show results obtained from a linear model with subject-specific fixed effects and the fifth and sixth column show estimates from a random effects probit model (marginal effects at population means).

Tables 13, 14, and 15 show the results for what makes subjects follow IBM (model 2). Table 13 contains the estimations for treatment GROUP, Table 14 for ROLE, and Table 15 for FULL. Again, the first two columns in each table show the results from the linear random effects model, also shown in the main body of the paper. The third and fourth columns show results obtained from a linear model with subject-specific fixed effects and the fifth and sixth column show estimates from a random effects probit model (marginal effects at population means).

Table 10: Estimating the likelihood that subjects change their actions in treatment **ROLE**.

ROLE	linear, random effects		linear, fixed effects		probit, random effects marginal effects only	
constant	886*** (42.6)	997*** (40.4)	879*** (35.6)	979*** (36.8)	—	—
own payoff	-.316*** (.033)	-.289*** (.033)	-.311*** (.033)	-.284*** (.033)	-.369*** (.04)	-.325*** (.04)
payoff diff.	.098*** (.035)	.100*** (.034)	.105*** (.035)	.108*** (.034)	.129*** (.04)	.133*** (.04)
relative propensity	—	-387*** (37.5)	—	-353*** (38.0)	—	-413*** (45.9)
R^2	.075	.131	.075	.129	—	—
# of obs.	3186	3186	3186	3186	3186	3186

Note: All coefficients and standard errors multiplied by 10^3 . Standard errors in parentheses. ***denotes significance at the 1% level.

Table 11: Estimating the likelihood that subjects change their actions in treatment **GROUP**.

GROUP	linear, random effects		linear, fixed effects		probit, random effects marginal effects only	
constant	579*** (26.9)	730*** (26.4)	581*** (17.8)	709*** (22.5)	—	—
own payoff	-.197*** (.024)	-.164*** (.024)	-.195*** (.024)	-.165*** (.024)	-.225*** (.03)	-.185*** (.03)
payoff diff.	.476*** (.043)	.454*** (.043)	.448*** (.044)	.429*** (.043)	.552*** (.05)	.538*** (.05)
relative propensity	—	-418*** (37.6)	—	-355*** (39.3)	—	-457*** (46.1)
R^2	.077	.146	.077	.145	—	—
# of obs.	3186	3186	3186	3186	3186	3186

Note: All coefficients and standard errors multiplied by 10^3 . Standard errors in parentheses. ***denotes significance at the 1% level.

Table 12: Estimating the likelihood that subjects change their actions in treatment FULL

FULL	linear, random effects		linear, fixed effects		probit, random effects marginal effects only	
constant	611*** (44.1)	756*** (37.3)	613*** (31.2)	736*** (32.9)	—	—
own payoff	-.121*** (.029)	-.077*** (.029)	-.123*** (.029)	-.089*** (.029)	-.148*** (.04)	-.104*** (.04)
payoff diff.	.211*** (.032)	.208*** (.031)	.208*** (.032)	.204*** (.031)	.275*** (.04)	.277*** (.04)
relative propensity	—	-467*** (36.5)	—	-389*** (38.1)	—	-516*** (50.1)
R^2	.042	.174	.042	.166	—	—
# of obs.	3186	3186	3186	3186	3186	3186

Note: All coefficients and standard errors multiplied by 10^3 . Standard errors in parentheses. ***denotes significance at the 1% level, ** denotes significance at the 5% level.

Table 13: Estimating the likelihood that subjects follow IBM in treatment ROLE.

ROLE	linear, random effects		linear, fixed effects		probit, random effects marginal effects only	
constant	127*** (41.2)	146*** (41.9)	122*** (39.7)	148*** (40.8)	—	—
own payoff	-.001 (.038)	.004 (.038)	-.002 (.038)	-.009 (.038)	-.015 (.04)	-.011 (.04)
payoff diff.	.248*** (.038)	.246*** (.038)	.249*** (.038)	.248*** (.038)	.234*** (.04)	.233*** (.04)
relative propensity	—	-90.1* (47.7)	—	-126*** (48.6)	—	-113** (50.7)
R^2	.038	.038	.038	.037	—	—
# of obs.	2079	2079	2079	2079	2079	2079

Note: All coefficients and standard errors multiplied by 10^3 . Standard errors in parentheses. ***denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level.

Table 14: Estimating the likelihood that subjects follow IBM in treatment GROUP.

GROUP	linear, random effects		linear, fixed effects		probit, random effects marginal effects only	
constant	145*** (22.5)	113*** (25.4)	116*** (21.2)	111*** (24.9)	—	—
own payoff	-.043 (.030)	-.058* (.030)	-.013 (.031)	-.015 (.030)	-.069* (.04)	-.077** (.04)
payoff diff.	.551*** (.045)	.586*** (.047)	.577*** (.045)	.582*** (.047)	.546*** (.05)	.565*** (.05)
relative propensity	—	131*** (49.1)	—	20.5 (56.5)	—	98.5* (55.3)
R^2	.080	.087	.078	.080	—	—
# of obs.	1644	1644	1644	1644	1644	1644

Note: All coefficients and standard errors multiplied by 10^3 . Standard errors in parentheses. ***denotes significance at the 1% level ** denotes significance at the 5% level. * denotes significance at the 10% level.

Table 15: Estimating the likelihood that subjects follow IBM in treatment FULL

FULL	linear, random effects		linear, fixed effects		probit, random effects marginal effects only	
constant	164*** (43.6)	166*** (45.3)	154*** (40.5)	159*** (45.3)	—	—
own payoff	.056 (.038)	.056 (.039)	.059 (.038)	.059 (.039)	-.057 (.04)	-.057 (.04)
payoff diff.	.156*** (.040)	.156*** (.040)	.152*** (.040)	.151*** (.040)	.165*** (.04)	.164*** (.04)
relative propensity	—	-13.4 (61.6)	—	-22.1 (62.4)	—	-15.2 (65.2)
R^2	.009	.009	.009	.009	—	—
# of obs.	1920	1920	1920	1920	3186	3186

Note: All coefficients and standard errors multiplied by 10^3 . Standard errors in parentheses. ***denotes significance at the 1% level.