

Alós-Ferrer, Carlos; Ritschel, Alexander

Working Paper

Multiple behavioral rules in Cournot oligopolies

Working Paper, No. 331

Provided in Cooperation with:

Department of Economics, University of Zurich

Suggested Citation: Alós-Ferrer, Carlos; Ritschel, Alexander (2020) : Multiple behavioral rules in Cournot oligopolies, Working Paper, No. 331, University of Zurich, Department of Economics, Zurich,
<https://doi.org/10.5167/uzh-172561>

This Version is available at:

<https://hdl.handle.net/10419/222355>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



**University of
Zurich**^{UZH}

University of Zurich
Department of Economics

Working Paper Series
ISSN 1664-7041 (print)
ISSN 1664-705X (online)

Working Paper No. 331

Multiple Behavioral Rules in Cournot Oligopolies

Carlos Alós-Ferrer and Alexander Ritschel

Revised version, July 2020

Multiple Behavioral Rules in Cournot Oligopolies*

Carlos Alós-Ferrer[†]
University of Zurich

Alexander Ritschel[‡]
University of Zurich

This Version: July 2020

Abstract

We study intra-individual behavioral heterogeneity in an experimental Cournot oligopoly. Previous empirical results in this setting have demonstrated convergence to competitive outcomes, in agreement with theoretical predictions assuming that players imitate successful opponents. We postulate that players sometimes rely on imitation of successful behavior, and sometimes best reply to the actions of others. Testable predictions are obtained from a model allowing for multiple behavioral rules which accounts for differences in the cognitive nature of the underlying decision processes. Those include non-trivial response time interactions depending on whether the rules share a common prescription (alignment) or not (conflict), a classification which is *ex ante* observable. The results confirm the hypotheses and support the presence of multiple behavioral rules at the individual level.

Keywords: Cournot oligopoly · Imitation · Best Reply · Multiple behavioral rules

JEL Codes: C72 · C91 · D91

1 Introduction

Imitation is pervasive in humans. It is well established that imitative behavior plays a fundamental role for human learning (Bandura, 1977). Humans actually *overimitate*, even when compared to chimpanzees (Horner and Whiten, 2005), imitate gestures of others unwittingly (Chartrand and Bargh, 1999), and imitate others even when it impairs task performance (Cracco et al., 2018). It has been argued that imitating successful others is also a common practice in markets, although the evidence is less direct. For instance, Williams and Miller (2002) conducted a cluster analysis of the decision-making styles of 1,684 executives and showed that the most

*The authors gratefully acknowledge financial support from the German Research Foundation (DFG) under project Al-1169/4, part of the Research Unit “Psychoeconomics” (FOR 1882). The experiment uses a decision task complying with standard conventions in experimental economics and was hence covered by umbrella ethical agreements in the institution where the experiment was carried out.

[†]Corresponding Author: Email: carlos.alos-ferrer@econ.uzh.ch. Zurich Center for Neuroeconomics (ZNE), Department of Economics, University of Zurich, Blümlisalpstrasse 10, 8006 Zurich, Switzerland.

[‡]Zurich Center for Neuroeconomics (ZNE), Department of Economics, University of Zurich, Blümlisalpstrasse 10, 8006 Zurich, Switzerland.

numerous group (“followers:” 36%) was described to “make decisions based on how they’ve made similar decisions in the past or how other trusted executives have made them.”

Imitative behavior would have important consequences for market outcomes. Vega-Redondo (1997) showed that, in dynamic Cournot oligopolies, if market actors imitate quantities leading to the highest profits and make occasional mistakes, the long-run market outcome corresponds to the Walrasian equilibrium, and not to the Cournot-Nash one (see also the “imitation equilibrium” of Selten and Ostmann, 2001). In this sense, imitation leads to more competitive outcomes than those predicted by standard game-theoretic concepts. This surprising result was generalized to the class of *aggregative games* by Alós-Ferrer and Ania (2005), and rests on the observation that, for payoff structures as that of a Cournot oligopoly and other games, imitative behavior systematically increases *relative payoffs*, as pointed out in earlier work by Schaffer (1988, 1989), which can destabilize the Cournot-Nash equilibrium but not the Walrasian one.

The result of Vega-Redondo (1997) motivated a series of experimental studies on Cournot oligopolies. Huck et al. (1999) and Offerman et al. (2002) conducted experiments with repeated interactions in Cournot triopolies and found more competitive outcomes when imitation is facilitated by providing information on individual actions and profits. In contrast, providing more information on the market structure, or restricting information to aggregate quantities, decreases competitiveness, as would be expected if behavior shifted toward (myopic) best-reply. Huck et al. (2000) found similar results in an experiment with differentiated-products oligopolies. In a different Cournot oligopoly experiment, Huck et al. (2002) suggest that the data would be consistent with a mixture of best reply and imitation. Convergence to Walrasian outcomes has been shown to be robust even under asymmetric costs, where the theoretical prediction does not generalize (Apesteguía et al., 2010).

Generally, experimental support for Walrasian outcomes as predicted by models of imitation seems to be stronger with larger numbers of participants per Cournot market. Huck et al. (2004) showed that experimental markets with four or more firms frequently become very competitive, while less competitive outcomes were observed for three or less firms. A few experiments have also examined Cournot oligopolies with very long horizons (1200 periods) and come to similar conclusions. Friedman et al. (2015) showed that duopolies and triopolies initially moved toward competitive outcomes, but participants gradually learned to coordinate on quantities near collusion. Oechssler et al. (2016) replicated this result for duopolies, but found quantities above the Cournot one for tetrapolies and concluded that “four remain many even with 1200 periods.”¹

¹Other experiments have examined the role of imitation in different settings. Abbink and Brandts (2008) carried out an experiment on Bertrand competition under convex costs and found evidence compatible with the presence of a heuristic based on imitation and experimentation, as modeled theoretically by Alós-Ferrer et al. (2000). Offerman and Sonnemans (1998) found that people often imitated successful others (but also learned from their own experience) in an experiment on individual investment decisions. Offerman and Schotter (2009) conducted an experiment using both individual production tasks and take-over games, where participants could sample the actions and payoffs of others. They found behavior mostly compatible with imitation even when it had detrimental payoff consequences.

While compatible with the theoretical predictions of imitation models as Vega-Redondo (1997), the experimental results above do not constitute a direct test of the actual presence of imitative behavior in Cournot oligopolies. Convergence to competitive outcomes is a joint implication of imitative behavior, a stochastic dynamics based on infrequent mistakes, and a focus on long-run outcomes. Thus, experiments showing such convergence demonstrate the (very relevant) prediction of increased competitiveness, but do not test imitative behavior in itself. In this work, we aim to provide more direct evidence on the use of imitation at the individual level, not relying on any convergence results.

A first step to investigate the presence or prevalence of imitation is to concentrate on whether actual choices are compatible with the predictions of imitative rules or not. Obviously, this is an incomplete test, as a choice predicted by an imitative rule might also be predicted by an alternative rule as e.g. best reply. However, evidence on choices compatible with a behavioral rule is still highly informative. Apesteguía et al. (2007) conducted a Cournot triopoly experiment with random re-matching each period within pre-specified groups and treatments that differed on whether players were informed about the actions and profits of their competitors, or those of other players in the experiment. They evaluate behavioral rules on the number of individual decisions compatible with the rules' predictions, and find that most subjects either repeat their previous choice, switch to the action with the highest-observed payoff, or adopt an action not observed in the last period. Strikingly, over 35% of their subjects switch to the action with the highest-observed payoffs in more than 60% of all decisions, and the percentage increases to 80% for about 10% of the players.

An alternative way to demonstrate the relevance of imitation is to concentrate on the characteristics of this behavioral rule as a heuristic, and this is the avenue we will pursue. In particular, we will use the fact that imitation can be seen as a more intuitive behavioral rule when compared to other rules as myopic best reply, in terms of their cognitive characteristics. In this sense, our work is related to Bosch-Domènech and Vriend (2003), who attempted to show the presence of imitative behavior in a Cournot oligopoly experiment. Their argument was that imitation should be more prevalent for more demanding environments because more sophisticated behavioral rules are then harder to apply. Thus, they used different treatments with different levels of complexity implemented through time limits and inconveniently-described payoff tables. The results showed no stronger reliance on imitation as complexity increased. In particular, aggregate quantities did not become closer to the Walrasian one (although they stayed above the Cournot one for all treatments). By analyzing the percentage of correct predictions at the individual level for different behavioral rules, Bosch-Domènech and Vriend (2003) find that imitation rules do not describe behavior better as complexity increases. However, the same is true for best-reply rules, which leads the authors to conclude that the more demanding treatments increased player disorientation.

A different approach following a conceptually-similar strategy was pursued by Buckert et al. (2017). These authors conducted a Cournot triopoly experiment and used treatments where subjects were either placed under time pressure or distracted by concurrent tasks, which induced

higher levels of stress according to cortisol measurements. Subjects under time pressure made a larger percentage of choices consistent with imitation.

In contrast to these works, our aim is to demonstrate that imitation coexists with other, more deliberative rules of behavior along the lines of (myopic) best reply. That is, our hypothesis is that there is *intra-individual heterogeneity* in behavior. To show this, we borrow from models analyzing multiple behavioral rules and derive predictions which would not hold if a single behavioral rule was at work. In particular, we obtain testable predictions both when imitation and best reply prescribe different actions and when they do prescribe the same action, even though in the latter case choice data cannot distinguish the rules. That is, we will not exclusively rely on the percentage of choices consistent with imitation, but rather formulate more nuanced predictions.

Specifically, we apply a “dual-process diffusion model” previously used to study decision process multiplicity in individual, non-strategic, binary decisions (Achtziger and Alós-Ferrer, 2014; Alós-Ferrer, 2018). We adapt this model to our setting (with non-binary decisions), but otherwise rely on the model structure as used in previous works. To derive testable predictions, the model develops ideas from *dual-process theories*, which postulate that the human mind is mainly influenced by two kinds of processes, called *automatic* and *controlled* (see Kahneman, 2003; Alós-Ferrer and Strack, 2014; see also Evans, 2008 and Weber and Johnson, 2009 for detailed reviews). Automatic processes are fast, unconscious, and require few cognitive resources. They capture impulsive reactions and behavior along the lines of stimulus-response schemes. Controlled processes are slow, consume cognitive resources, and are reflected upon (partly) consciously. Although there is a clear analogy between dual-process theories and the distinction between full and bounded rationality in the economic sciences, the key difference is that dual-process models assume heterogeneity within the individual.

The model considers two different behavioral rules. As an alternative to imitation, we will focus on myopic best reply, that is, payoff maximization taking current information on other agents’ behavior as given. This is a natural choice for Cournot oligopolies, as it captures one-step strategic behavior and acts as a first-order proxy of deliberative thinking. Obviously, however, this is a simplification, as it should be expected that more complex deliberative rules are used by at least some participants.

We postulate that imitation is a more automatic rule, where individuals react to a more successful action and respond by imitating this action, while myopic best reply is a more deliberative rule which involves active maximization after considering available information. This is confirmed by widespread evidence from cognitive psychology and neuroscience, which indicates that imitation learning display the characteristics of automatic processes (for a meta analysis of 226 experiments see Cracco et al., 2018). For our purposes, the key observation is that imitation, as a boundedly-rational behavioral rule, can be expected to be more automatic than rules assuming explicit payoff maximization as myopic best reply. However, there are clear differences between our approach and research in cognitive psychology. First, the paradigm we focus on is far more complex than those typically encountered in that literature. Second, the behavioral

rules we are interested in all involve cognitive aspects (as opposed to purely automatic reactions). That is, we do not postulate that imitation is an exclusively-automatic process as those studied in cognitive psychology, but merely that it is *more automatic* (or less deliberative) than myopic best reply.²

Our predictions focus on one of the most basic measures of process data, *response times*. Those are a standard tool in psychology and are now receiving increasing attention for the study of economic decisions (Moffatt, 2005; Rubinstein, 2007, 2016; Achtziger and Alós-Ferrer, 2014; Alós-Ferrer et al., 2016; Alós-Ferrer and Ritschel, 2018; Spiliopoulos and Ortmann, 2018). The key insight allowing for testable predictions is that more automatic processes are faster than more deliberative ones, and hence response times can be used as a direct source of evidence for the involvement of different decision processes. This does *not*, however, mean that one can simply classify decisions in fast and slow according to some exogenous criterion and conclude that one kind of decisions is more automatic. This would be an example of the “reverse inference” fallacy (Krajbich et al., 2015). The problem is that processes, and behavioral rules, are not directly observable. Hence, when observing a choice and its associated response time, we cannot know which process has generated them. Each process will result in a distribution of response times (and choices!). However, by exploiting the concepts of conflict and alignment among behavioral rules (i.e., whether they prescribe the same answer or different ones), our model avoids reverse inference while still allowing for specific, non-trivial predictions (on response times conditional on specific types of choices).

The model delivers four kinds of predictions. First, whenever best reply and imitation are in conflict (make different prescriptions), choices where a best reply is selected are slower than when both rules are aligned (make the same prescription). Intuitively, this is because in case of conflict best replies come almost exclusively from the slower best reply rule, while in case of alignment the faster imitation rule also contributes a significant proportion of best replies. This is analogous to the well-established Stroop effect from cognitive psychology (Stroop, 1935; see Section 2). Second, in case of conflict, best replies are slower than imitative decisions, essentially because many of the latter arise from the faster imitation rule. Third, in contrast to the case of conflict (and somewhat counterintuitively), in case of alignment best replies are *faster* than other responses. This is because, in this case, the faster imitation rule contributes a large number of apparent best replies. That is, we obtain a testable prediction even in cases where the actual choices would be uninformative in order to disentangle the rules. Fourth, there are fewer best replies in case of conflict than in case of alignment. Also, there are fewer imitative choices in case of conflict than in case of alignment. This is simply because in case of alignment both behavioral rules favor a common prescription. None of the differences above would obtain if a single behavioral rule determined behavior.

²Dual-process theories often use the labels automatic (or intuitive) and controlled (or deliberative) in a dichotomous (“dual”) way for simplicity. However, the underlying dimension (automaticity) is actually viewed as a continuum in psychology (e.g., Allport, 1954; Schneider and Shiffrin, 1977; Shiffrin and Schneider, 1977; Bargh, 1989; Cohen et al., 1990). Few processes are purely automatic or purely deliberative. In particular, many processes that are often informally described as “intuitive” are merely seen as *less* deliberative than others.

In a laboratory experiment, we find clear evidence for the predictions detailed above. The results suggest that multiple behavioral rules codetermine behavior in a complex strategic setting (Cournot oligopoly), with imitation of past success and myopic payoff maximization being the two main drivers of decisions. In particular, we explicitly reject that a single rule per individual explains behavior. This multiplicity occurs at the individual level, that is, behavioral heterogeneity starts within each single decision maker.

It is worth mentioning that our analysis involves relatively long response times (typically 10 to 15 seconds), compared to most response-time studies in cognitive psychology. This is not surprising, since the task we study (as many relevant tasks in economics) is more complex than those typically used in that field. This is, however, no obstacle for our analysis. Obviously, if the differences in response times between more intuitive and more deliberative processes are large enough, effects will still be observed even if overall response times are long, and even if none of the processes is purely automatic.³ Ultimately, whether those effects are large enough to be detected is an empirical question, which we answer in the affirmative.

The remainder of the paper is structured as follows. Section 2 presents the formal model and derives our predictions. Section 3 presents the experimental design and procedures, and describes the strategy of analysis. Section 4 discusses the results. Section 5 provides a complementary discussion on the possibility of disentangling different types of imitative decisions (imitating yourself vs. imitating others). Section 6 concludes. Proofs are relegated to the appendix. Experimental instructions and screenshots are in the Online Appendix.

2 Predictions for Multiple Behavioral Rules

This section generalizes the model of Achtziger and Alós-Ferrer (2014) and Alós-Ferrer (2018), which was restricted to binary choice, to the multiple-alternative case. The model assumes that two behavioral rules codetermine behavior, a more deliberative one and a more intuitive/impulsive one. For our purposes, we concentrate on myopic best reply and imitation.⁴

2.1 A Simple Formal Model

Consider a given decision in a Cournot oligopoly, where a decision maker has received information on the quantities produced and profits earned by all involved firms. This, together with the structure of the game, allows to compute both the action which maximizes payoffs given the actions of other players (myopic best reply) and to observe the action which has led to the largest payoffs in the last interaction (imitative choice). Suppose that only finitely many options are available (as will be the case in the experiment). Denote by X the finite

³Recall footnote 2. As an extreme example, Alós-Ferrer et al. (2016) study the responses to questions as in the Cognitive Reflection Test, with median response times typically above 30 seconds, and find differences depending on whether the questions elicited more or less intuitive responses.

⁴Achtziger and Alós-Ferrer (2014) applied the model to reinforcement and Bayesian updating in a belief updating task. Spiliopoulos (2018) used the model to study win-stay, lose-shift vs. more sophisticated (cognitive) heuristics in a repeated game played against computer algorithms. Ludwig et al. (2020) applied it to heuristic decisions in probability judgments.

set of options, with typical element $x \in X$. Let x^B denote the myopic best reply and x^I the imitative choice. Assume for simplicity that there are no ties (again, as will be the case in the experiment). We will assume below that the myopic best reply favors x^B above other options, and that the imitation rule favors x^I above other options. We speak of *conflict* if $x^B \neq x^I$, that is, imitating the best observed payoffs would not result in a best reply, and we speak of *alignment* if $x^B = x^I$.

Example 1. To fix ideas, suppose $X = \{7, 8, 9, 10\}$ and $N = 4$ players compete in a symmetric Cournot oligopoly with linear inverse demand function given by $P(Q) = 48 - Q$, $Q = x_1 + \dots + x_4$, and linear cost function $C(x_i) = 8x_i$, hence $\pi_i(x_i|x_{-i}) = (40 - Q)x_i$. Consider player 1, and suppose current play is given by $x_1 = 9$, $x_2 = x_3 = 8$, and $x_4 = 10$. Given $x_2 + x_3 + x_4 = 26$, player 1's profits $\pi_1(x_1|x_{-1}) = (14 - x_1)x_1$ are maximized at $x_1 = 7$, hence (for player 1) $x^B = 7$. However, given the current quantities, $P(Q) = 13$ and actual profits are $\pi_1 = 45$, $\pi_2 = \pi_3 = 40$, and $\pi_4 = 50$, hence player 4 reaches the largest observed profits and $x^I = x_4 = 10$. Since $x^B \neq x^I$, this is a case of conflict. Suppose that current play was given by $x_1 = 7$ and $x_2 = x_3 = x_4 = 8$ instead. Then, $\pi_1(x_1|x_{-1}) = (16 - x_1)x_1$ is maximized at $x^B = 8$. Since $P(Q) = 17$, current profits are $\pi_1 = 63$ and $\pi_2 = \pi_3 = \pi_4 = 72$, hence $x^I = 8 = x^B$, yielding a case of alignment.

Let BR and Im denote the myopic best reply and imitation rules, respectively. Which of the two rules will actually determine behavior is a stochastic event. Let $\Delta > 0$ be the probability that the actual response is selected according to imitation, and $1 - \Delta$ the probability that it is selected according to myopic best reply. Moreover, we assume that all rules are stochastic in nature, i.e., they carry an amount of noise, resulting in errors (deviations from the rule's prescription). Note that, hence, myopic best reply can select x^I and imitation can select x^B even in case of conflict, and any of them could select actions $x \neq x^B, x^I$. That is, in case of alignment ($x^B = x^I$) both behavioral rules tend to make the same prescription and in case of conflict ($x^B \neq x^I$) they would make different prescriptions in the absence of noise, but due to behavioral noise they might actually select either option in either case, or a third, different one.

Denote by P^{BR} the probability with which the myopic best reply rule indeed selects the best reply x^B , and by P^{Im} the probability with which the imitation rule selects the alternative with the highest observed payoff, x^I . That is, if $P^{BR}(x)$ and $P^{Im}(x)$ denote the probabilities with which each rule selects $x \in X$, conditional on the rule being the one which actually determines the response, then $P^{BR} = P^{BR}(x^B)$ and $P^{Im} = P^{Im}(x^I)$. Our first assumption is as follows.

(P1) For each decision situation,

$$P^{BR} > P^{BR}(x) \quad \forall x \in X, x \neq x^B \quad \text{and} \quad P^{Im} > P^{Im}(x) \quad \forall x \in X, x \neq x^I.$$

This is a minimal consistency condition which simply declares that the prescription of a rule is indeed the rule's most frequent selection, but it is a rather mild one, since for the multi-alternative case it does not even imply that the prescription is selected more than half of the time.

Response times are also assumed to be stochastic. Let $R^B = E[RT|BR]$ and $R^I = E[RT|Im]$ denote the *expected* response times conditional on the response being selected by the myopic best reply or the imitation rule, respectively. For simplicity, we assume that expected response times do not depend on the actually-selected response. Naturally, since imitation is thought to be more automatic, hence faster *in expected terms*, we assume

(R) $R^B > R^I$.

For some of the results below, we will further assume that

(P2) $P^{Im} > P^{BR}$,

i.e. the deliberation process behind best reply (computing the myopically optimal behavior) is noisier than the stimulus-response process behind imitation (copying the action with the largest observed payoff), while the latter is more *consistent*. This is natural since imitation is assumed to be more automatic (closer to a stimulus-response process).

A simple way to think of the model is to conceive of the imitation rule as a swift cognitive shortcut, which selects the action with the largest observed payoff quickly and very frequently, while the myopic best reply rule is a slow, deliberative process which depends on actual computations and is hence less consistent.

For the binary-choice case, the model in Achtziger and Alós-Ferrer (2014) has been given a micro-foundation in Alós-Ferrer (2018) as the *dual-process diffusion model* or DPDM. In this model, the processes are instantiated as diffusion processes as in the drift-diffusion model (DDM) of Ratcliff (1978) and Ratcliff and Rouder (1998), which has been recently further analyzed by Fudenberg et al. (2018) and is standard in cognitive psychology and neuroscience (e.g. Shadlen and Shohamy, 2016). In this model, evidence accumulation (internal to the decision maker) is captured as a diffusion process with a trend μ and two barriers. Whether the process chooses an option or the other corresponds to whether the upper or the lower barrier is hit first. The response time is the time at which the first barrier is hit. Alós-Ferrer (2018) shows that, in the DPDM, assumptions (P1), (P2), and (R) are implied if one simply assumes that the drift rate of the more automatic process is larger in absolute value than the drift rate of the more deliberative process, capturing that the former is swifter than the latter.

2.2 Predictions

Since all our formal results translate directly into experimental hypotheses, we label them accordingly for convenience (H1, H2, etc). The first testable prediction of the model concerns the comparison of conflict and alignment. Recall that, by committing *ex ante* to which behavioral rules are of interest, we can identify situations of conflict and alignment before data collection. The first prediction states that the response time of best replies must be strictly larger in situations of conflict than in situations of alignment. Since this prediction arises exclusively from process multiplicity, it essentially constitutes a “smoking gun” test on the presence of multiple processes.

Theorem 1. *Under (P1) and (R),*

(H1) the expected response time of best replies in case of conflict is strictly longer than the expected response time of best replies in case of alignment.

The intuition for Theorem 1 is as follows (all proofs are in the appendix). Independently of whether the decision problem corresponds to conflict or alignment, the best reply rule delivers the same proportion of best replies, which are relatively slow. In case of conflict, the imitation rule favors the imitative choice, which is not a best reply, and hence typically contributes relatively fewer (fast) best replies. In case of alignment, the imitation rule actually favors the best reply, and hence typically contributes relatively many (fast) best replies. Hence, one obtains faster best replies under alignment than under conflict.

It is worth noticing that the prediction of Theorem 1 corresponds to the well-known “Stroop Effect” discussed in psychology (Stroop, 1935; MacCleod, 1991), which describes a slow-down of (correct) responses when one is asked to name the color that a word is printed in but that word happens to name a different color (e.g., “Red” printed in blue) compared to when the word names the color it is printed in (e.g., the word “Red” printed in red). However, work in psychology typically assumes that this and similar response-times effects are due to central executive functions of the brain related to the detection and resolution of conflict among elementary responses, which tax cognitive resources and require time (Bargh, 1989; Baddeley et al., 2001), but enable the inhibition of automatic responses in case of conflict. The model presented here does not assume such a difference in response times and Theorem 1 holds in its absence. However, it is easy to show that all results described here hold in an extended model incorporating a “non-decision time” that is longer under conflict than under alignment. Results also hold if the value of Δ is assumed to be smaller under conflict, reflecting conflict detection (Alós-Ferrer, 2018).

The model also makes more nuanced predictions for the response times of best replies and other responses. Those amount to a non-trivial interaction between responses (best replies, imitative choices, or other alternatives) and cognitive situations (conflict or alignment). Specifically, best replies must be slower on average than imitative choices in case of conflict, but in case of alignment (where best replies are also imitative choices), they must be faster than other choices. This parallels the prediction of Achtziger and Alós-Ferrer (2014) and Alós-Ferrer (2018) that in situations with normatively correct answers errors are fast in case of conflict but slow in case of alignment. This asymmetry goes beyond simple informal statements that intuitive responses should be faster, which might hide a reverse inference fallacy (Krajbich et al., 2015), and serves as a test of the basic structure of the model. The next result gathers the predictions.

Theorem 2. *Assume (R).*

(H2) Under (P1), in case of conflict, the expected response time of best replies is larger than the expected response time of imitative choices (choosing the alternative with highest observed payoff).

(H3) Under (P2), in case of alignment, the expected response time of best replies is shorter than the expected response time of other choices.

The intuition behind Theorem 2 is as follows. The (slow) best reply rule favors the best reply alternative and the (fast) imitation rule favors the imitative choice. Those two alternatives are different in case of conflict, and hence best replies end up being on average slower in this case. In case of alignment, the two alternatives coincide but by (P2) the fast imitative process contributes more of them than the best reply rule, hence in expected terms best replies end up being on average faster. In other words, in case of alignment, the imitation rule acts as a quick and efficient shortcut to identify the best reply while the more error-prone best reply rule contributes relatively more (slow) non-best-reply answers. Hence, conditional on a best reply not being observed, it is more likely that the response is generated by the slower best reply rule.

Last, the model also makes predictions for the proportion of best replies and imitative choices comparing the cases of conflict and alignment, which we summarize in the following result.

Theorem 3. *Under (P1),*

- (H4a) the proportion of best replies is strictly smaller in case of conflict than in case of alignment, and*
- (H4b) the proportion of imitative choices is strictly smaller in case of conflict than in case of alignment (when they are also best replies).*

The intuition for Theorem 3 is immediate. In case of alignment, both behavioral rules favor the same option, in the sense of being the one selected most often. That option is simultaneously a best reply and an imitative choice. In case of conflict, the myopic best reply rule still favors best replies, but the imitation rule now favors a different option, which is imitative but not a best reply. Even though each rule might still select the option favored by the other rule in case of conflict, it does so less often. Hence, in case of alignment the common prescription obtains more often than any of the individual choices in case of conflict.

All predictions above are in terms of inequalities. It is immediate that, if only one behavioral rule was present, none of those would obtain. In our empirical analysis, (H1)–(H4b) are treated as hypotheses, hence the corresponding null hypotheses (equalities) are the ones that would result under one behavioral rule only. In this sense, confirming our predictions rests on the rejection of those null hypotheses, and thus we will interpret the results below as evidence for the multiplicity of behavioral rules.

3 The Experiment

3.1 Experimental Design and Procedures

In our experiment, participants interacted in 4-player Cournot oligopolies (tetrapolies). We made this choice because previous results have shown that convergence to Walrasian outcomes, which is compatible with imitative behavior, occurs more frequently with tetrapolies than with triopolies (Huck et al., 2004; Oechssler et al., 2016). We conducted four sessions with 32

participants each for a total of $N = 128$ (82 females; median age 22 years) at the Cologne Laboratory for Economic Research (CLER). The experiment was programmed with z-Tree (Fischbacher, 2007) and participants were recruited using ORSEE (Greiner, 2015). We excluded students majoring in economics, psychology, and business, as they might have been taught game-theoretic concepts which might influence their behavior. A session lasted around 90 minutes and average earnings were 13.59 EUR, including a show-up fee of 2.50 EUR.

Each participant competed in three different Cournot oligopolies (parts), which lasted for 17 periods each. Initially, players were matched in groups of four to play the first tetrapoly (Part 1). After 17 periods, players were rematched in new groups of four and the oligopoly payoffs (demand function) were changed (Part 2). After 17 further periods, players were rematched again and played a third oligopoly with new payoffs (Part 3). To increase the number of fully-independent observations, rematching was done within 16 pre-determined blocks of 8 participants each. Identities within a part were always anonymous and could not be traced back to previous parts. In each new part, at least two players in the group were different from the previous group. The sequence of the different oligopolies was varied across sessions.

The three parts were implemented because, in contrast to previous experiments (e.g., Huck et al., 1999; Offerman et al., 2002), we are interested in behavioral correlates of individual actions, rather than on eventual convergence. If and when convergence occurs, there is no further variance in the behavioral (choice) data, and response times become meaningless as participants mechanically repeat a fixed action. Hence, we were interested in data before convergence occurred. Thus, to maximize usable data, we implemented three parts (oligopolies) with rematching of participants, reassignment of identities, and changed payoff tables (computed with different demand functions). By the same reasoning data would be meaningless if and when collusion occurred. Rematching, working with shorter oligopolies, and changing payoff tables across parts also diminish the likelihood of collusion and increase the variance in behavioral data.

Each oligopoly was implemented through a payoff table (similarly to treatments in Bosch-Domènech and Vriend, 2003 and Apesteguía et al., 2007) derived from a linear inverse demand function of the form $P(Q) = a - Q$, where P is the price, a the saturated demand, Q the total quantity in the market, and constant marginal costs, normalized to zero. A neutral framing was used and neither firms nor quantities were mentioned. We reduced the action space to four possible actions (A , B , C , and D). To further decrease the likelihood of fast convergence, the ordering of the quantities (A to D) changed with each part, that is, in some parts the assignment of quantities to letters was increasing and in some it was decreasing.⁵ The second and third parts always had a different payoff table and a reversed ordering of the quantities with respect to the previous part. Hence, in each part, the game is given by a $4 \times 4 \times 4 \times 4$ payoff table, which by symmetry can be reduced to a 4×20 table, with four rows for the possible actions and 20 columns (labeled AAA to DDD) for the opponents' actions (independently of their identity).

⁵Payoff table 1: $P(Q) = 150 - Q$, $A = 37.5$, $B = 33.25$, $C = 30$, $D = 18.75$ (or reversed); Payoff table 2: $P(Q) = 175 - Q$, $A = 43.75$, $B = 38.875$, $C = 35$, $D = 21.875$ (or reversed); Payoff table 3: $P(Q) = 200 - Q$, $A = 50$, $B = 44.5$, $C = 40$, $D = 25$ (or reversed).

We discretized the action space to make the postulated behavioral rules (myopic best reply and imitation) both feasible and comparable. A continuous- or large-action space would have turned myopic best reply into an abstract maximization problem, while imitation would remain a discrete, intuitive rule. By choosing a discrete setup we go against our hypotheses and reduce the conceptual distance between the two behavioral rules.

Payoffs were expressed in points (rounded to the nearest integer), with an exchange rate of 18 Eurocents per 1000 points. The payoff table was permanently visible in the upper part of the screen during the corresponding part of the experiment. Example screenshots and instructions are presented in the (Online) Appendix.⁶ The points achieved in all 51 rounds were accumulated and paid at the end of the experiment. Following the standard procedure in Cournot oligopoly experiments (e.g., Huck et al., 1999; Offerman et al., 2002; Apesteguía et al., 2007, 2010), all decisions were paid. There are two reasons for this choice. First, this is a dynamic setting with feedback in which the repeated rounds are not independent of each other. Second, the possibility of imitation was an essential aspect of the design. We wanted to emphasize that the other players did in fact earn a certain amount. An alternative would have been to pay only one round at the end of the experiment, but this would raise the concern that imitation might not be triggered.

In order to focus on the interaction between myopic best reply and imitation, we highlighted the information required to implement both rules. Myopic best reply implements maximization within the column corresponding to the actual actions of the opponents in the previous period. For all rounds except the first one within each part, that column was highlighted. Thus, determining a myopic best reply required comparing four numbers only. For each round except the first, participants were also given feedback on the actions and profits of the group members in the previous period, making imitation feasible. As a robustness control, to make sure that presentation effects were minimized, we included two treatments which differed only on how that information was presented. In Treatment *FullInfo*, the choices and points of all other group members were presented in separate boxes, in addition to a box displaying the own choice and received points, and the box with the highest point amount was highlighted (as in the more demanding treatments in Bosch-Domènech and Vriend, 2003). This design choice was made to eliminate mechanical differences between imitation and myopic best reply and make both behavioral rules equally salient. Please note that, since we assume that imitation is a more automatic behavioral rule than best reply (and our predictions crucially hinge on this assumption), this choice works *against* our hypotheses, since the only remaining differences among the rules are of cognitive nature. In Treatment *BestOnly* only the own choice and points plus an additional (highlighted) box were shown, with the latter displaying the choice with the largest amount of points in the previous round (and the corresponding points). Note that in the FullInfo treatment both imitation and myopic best reply involve comparing four numerical quantities, making the mechanical aspects of the rules as comparable as possible. In contrast,

⁶Participants were asked to make a decision within 30 seconds. After that time, a request to make the decision appeared in a screen's corner. Only 162 out of 6,144 decisions (2.64%) were made after 30 seconds.

Table 1: Overview of Prescribed Actions

	AAA		AAB		AAC		AAD		ABB		ABC		ABD		ACC		ACD		ADD	
	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im
A	D	A	D	A	D	A	C	A	D	A	C	A	C	A	C	A	B	A	A	A
B	D	A	D	A	D	A	C	A	D	A	C	A	C	A	C	A	B	A	A	A
C	D	A	D	A	D	A	C	A	D	A	C	A	C	A	C	A	B	A	A	A
D	D	A	D	A	D	A	C	A	D	A	C	A	C	A	C	A	B	A	A	A

	BBB		BBC		BBD		BCC		BCD		BDD		CCC		CCD		CDD		DDD	
	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im	BR	Im
A	C	A	C	A	B	A	C	A	B	A	A	A	C	A	A	A	A	A	A	A
B	C	B	C	B	B	B	C	B	B	B	A	B	C	B	A	B	A	B	A	B
C	C	B	C	B	B	B	C	B	B	B	A	B	C	C	A	C	A	C	A	C
D	C	B	C	B	B	B	C	B	B	B	A	B	C	C	A	C	A	C	A	D

Notes: Overview of prescribed actions for each behavioral rule depending on last period's outcome. Cell entries describe the action prescribed by myopic best reply (BR) and imitation (Im) when the player previously chose the action given in the row and the opponents chose the actions given in the column. Shaded entries for BR and Im indicate that the two rules are aligned.

the BestOnly treatment closely reproduces the idea of “imitate the best” as described, e.g., by Vega-Redondo (1997), i.e. choosing the action with the highest profit in the previous round. The treatments were implemented between subjects, with half the subjects in each treatment in every session. As we will see below, results are not affected by the differences in information presentation.

3.2 Classification of Decisions and Strategy of Analysis

The data set of our experiment consists of $128 \times 48 = 6,144$ observations. The first decision within each part is always excluded since for that period there is no feedback concerning previous actions and the behavioral rules considered make no prescriptions.

Given the previous actions of all four players, the identification of the prescriptions of the different behavioral rules is straightforward. Table 1 displays the prescriptions of myopic best reply and imitation in the experiment, for the case of decreasing assignment of quantities to letters.⁷ Those prescriptions were identical for all three payoff tables. That is, the table shows the prescription of each behavioral rule when a specific combination of one's own choice (row) and the choice of the other players (column) occurred in the previous round. Whenever myopic best reply is in alignment with imitation (that is, both prescribe the same action), the corresponding cells are shaded in gray. Hence, unshaded entries indicate conflict between myopic best reply and imitation.

⁷For the analysis of the data, the case of increasing assignment of quantities was simply recoded.

Given Table 1, for periods 2–17 within each part, we can classify each actual decision of each participant according to the prescriptions (favored options) of the myopic best reply and imitation. As was to be expected, the majority of the 6,144 decisions were made in conflict situations (5,010; 81.54 percent). Of those, 26.57, 31.64, and 41.80 percent were myopic best replies, imitative decisions, or other choices, respectively. However, there were enough decisions made in case of alignment (1,134; 18.46 percent) to enable a meaningful analysis. Of those, 34.13 percent were myopic best replies (and hence also imitative). Note that our model explicitly accounts for behavioral noise, that is, we do not expect that all choices are those favored by either best reply or imitation. For instance, in conflict situations there are always two alternatives which are favored by neither myopic best reply nor imitation; since behavioral rules are stochastic, those can actually be selected by either rule.

It is natural to ask whether all subjects relied on best reply and imitation with similar intensity, or whether there was heterogeneity among subjects (i.e., subjects relying predominantly on only one of the rules). Figure 1 displays a histogram of the (log-transformed) individual ratios of myopic best replies vs. imitative choices, for the case of conflict (obviously, the exercise is not feasible in case of alignment, since then choice data cannot disentangle the rules). That is, 0 corresponds to subjects who chose an equal proportion of myopic best replies and imitative choices in conflict situations. Values below 0 correspond to subjects with a larger proportion of imitative decisions than myopic best replies, and vice versa for values above 0. The two dashed lines represent the 1-to-3 and 3-to-1 ratios. Around 78% of subjects lie between these ratios, indicating that most subjects do not predominantly rely on only one rule, in agreement with our assumptions. The figure also reveals a slight left-skewness, which suggest a somewhat higher reliance on imitation than on best reply and agrees with the descriptive statistics reported above.

In order to test our hypotheses, we will initially conduct non-parametric tests. For instance, we can test whether decisions compatible with one behavioral rule are faster than those compatible with another decision rule, conditional, e.g., on conflict among the rules (Hypotheses H2, H3). To do so, we look at all situations where the two rules conflict and build two sets of decisions for each individual, those where the prescription of the first rule was followed, and those where the prescription of the second rule was followed. Then we apply the appropriate test (in this case, a Wilcoxon Signed-Rank test).

For the analysis we consider the matching block the appropriate unit of observation, i.e. observations of all subjects who interacted anonymously with each other throughout all 3 parts are pooled into one observation. Since participants were separated into $N = 16$ different blocks (8 in each treatment) and were rematched only within those blocks, this guarantees completely independent observations. For each block, we compute the relative frequencies of choices and the average response times when following a given behavioral rule, conditional on conflict or alignment of myopic best reply and imitation.⁸

⁸A case can be made for individual observations as the appropriate unit of analysis. Following the logic of stochastic evolutionary models (Blume, 1993; Kandori et al., 1993; Vega-Redondo, 1997; Alós-Ferrer and Ania, 2005), behavioral rules have a Markovian structure, i.e. they are mappings from information (outputs and profits




Figure 1: Myopic Best Reply vs. Imitation, Log-transformed Individual Ratios. Frequency distribution of log-transformed myopic best reply/imitation ratios per subject in case of conflict and fitted normal density function. The left and right dashed lines represent the 1-to-3 and 3-to-1 ratios, respectively.

Before proceeding to the main analysis, we comment on the informational treatments. Those served as a robustness check to ensure that mere presentational effects, as salience of the maximum observed payoffs, did not significantly affect response times or drive behavior toward imitation. A block of 8 participants made on average 125.13 imitation decisions in the BestOnly treatment and 121.38 in the FullInfo treatment, which was not significantly different according to a Mann-Whitney-Wilcoxon (MWW) test ($N = 16$, $z = 0.735$, $p = .4622$). The average response time of imitation decisions was 10.38 s in the BestOnly treatment and 10.36 s in the FullInfo treatment (MWW, $N = 16$, $z = 0.210$, $p = .8336$). There were also no differences for myopic best replies. Hence, for the remainder of the analysis we will pool the data of both treatments.

in last period) to actions. Under this assumption, how exactly the input of the behavioral rule is generated is irrelevant. Hence, the fact that participants were part of tetrapolies which themselves were subgroups of certain blocks plays no role, for we are testing relative frequencies and response times which are generated *after* observation of the input, and tests condition on the relevant categories of inputs. Our conclusions were unchanged when conducting tests at the individual level (considering only those average response times with at least two observations per individual).




Figure 2: Average Response Times and Choice Frequencies. Left-hand side: Average response times of imitative choices (Im), myopic best replies (BR), and other choices conditional on conflict and alignment. Right-hand side: Relative frequency of imitative choices and myopic best replies conditional on conflict and alignment. Stars indicate the significance of Wilcoxon Signed-Rank tests. * $p < .1$, ** $p < .05$, *** $p < .01$.

4 Results

Figure 2 illustrates the tests of all our predictions. Average response times are shown on the left-hand side, and choice frequencies on the right-hand side. We now discuss all predictions as depicted in the figure, reporting the corresponding non-parametric tests (a regression analysis is discussed below). Note that our hypotheses yield specific directional predictions which would allow us to rely on one-sided p -values. However, we will conservatively report two-sided p -values.

Prediction (H1) serves as a first test of the presence of several, distinct behavioral rules. Myopic best replies, the prescription of the more deliberative behavioral rule, should be slower in case of conflict with imitation than in case of alignment. This corresponds to the comparison between the average response times of best replies in conflict and in alignment in Figure 2. The prediction is confirmed by the data: myopic best replies are slower in conflict (mean 12.38 s) than in alignment (mean 10.46 s), with the differences being highly significant according to a Wilcoxon-Signed-Rank (WSR) test ($N = 16$, $z = 2.947$, $p = .0032$).

Predictions (H2) and (H3) constitute a test of the nature of the involved processes and of the dual-process structure of the interaction. Essentially, myopic best replies should be relatively slow in case of conflict but relatively fast in case of alignment. Specifically, (H2) states that myopic best replies are slower than imitation decisions in conflict situations. As predicted, myopic best reply decisions are slower (average 12.38 s) than imitative choices (average 10.36 s) when the processes make different prescriptions, confirming the relatively more automatic nature of imitation decisions (compare the two left-most bars in the left-hand side of Figure 2). The difference is highly significant according to a WSR test ($N = 16$, $z = 3.361$, $p = .0008$). (H3) states that in case of alignment, myopic best replies (which are also imitative in this case) should be faster than other decisions. As predicted, myopic best replies (average 10.46 s) are significantly faster than other decisions (average 13.63 s; WSR, $N = 16$, $z = -3.258$, $p = .0011$).

The remaining two hypotheses concern relative choice frequencies. (H4a) states that myopic best replies should be less frequent under conflict than under alignment (when they are also imitative choices). This is illustrated in the right-hand side of Figure 2. In case of conflict, participants chose myopic best replies, on average, 26.57 percent of the time (average of individual averages), compared to 34.27 percent in case of alignment. The difference is highly significant (WSR test, $N = 16$, $z = -2.947$, $p = .0032$). (H4b) states that, in contrast, imitative decisions should be less frequent under conflict than under alignment (when they are also best replies). This is indeed the case, with an average of 31.59 percent of imitative decisions in case of conflict and 34.27 percent in case of alignment, although the difference is not significant with our two-tailed tests (WSR test, $N = 16$, $z = -1.344$, $p = .1788$). We remark, however, that one group successfully colluded during the last part of the experiment. When excluding the corresponding block observation, we observe less imitative decisions in conflict (average 30.97 percent) than in alignment (34.64 percent; two-tailed WSR test, $N = 15$, $z = -1.874$, $p = .0609$). All previous conclusions regarding (H1-H4a) remain unchanged when excluding the block containing the colluding group.

In summary, simple non-parametric tests already confirm our predictions. Hence, our experimental evidence is compatible with the interpretation that multiple behavioral rules, i.e. myopic best reply and imitation, codetermine behavior in complex Cournot oligopolies. We view this as a demonstration that complex economic decisions result from the interaction of multiple behavioral rules within individual economic agents.

We now turn to a more detailed regression analysis. Our data forms a perfectly-balanced panel with 48 decisions for each of the 128 participants (total $N = 128 \times 48 = 6,144$). Table 2 reports random effects panel regressions on log-transformed response times.⁹ We (conservatively) cluster standard errors at the block level. Since our hypotheses hinge on the distinction between conflict and alignment, it is important to introduce the appropriate categories in the analysis. The Conflict dummy takes the value 1 when the decision corresponds to a case of conflict between myopic best reply and imitation. To avoid having to rely on *post hoc* tests, we

⁹Response times are naturally bounded below by zero and usually present a skewed, non-normal distribution. To account for these features it is common practice to use a logarithmic transformation (Fischbacher et al., 2013; Achtziger and Alós-Ferrer, 2014).

Table 2: Random Effects Panel Regressions on (log) Response Times.

ln(ResponseTime)	Model 1	Model 2	Model 3
Conflict	0.1154*** (0.0360)	0.1418*** (0.0327)	0.1431*** (0.0329)
Imitation-Conflict	-0.1809*** (0.0307)	-0.1510*** (0.0255)	-0.1517*** (0.0255)
Other	0.1794*** (0.0424)	0.1735*** (0.0369)	0.1748*** (0.0370)
Other×Conflict	-0.1808*** (0.0406)	-0.1868*** (0.0354)	-0.1879*** (0.0355)
FullInfo Treatment	-0.0210 (0.0559)	-0.0366 (0.0570)	-0.0308 (0.0517)
Collusion		-0.2405*** (0.0525)	-0.2005*** (0.0513)
Constant	2.2418*** (0.0646)	2.6385*** (0.0579)	2.3426*** (0.1495)
Controls	No	Yes	Yes
Demographics	No	No	Yes
R^2	0.0313	0.1327	0.1587

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

further include the dummy Imitation-Conflict which only considers cases where the imitative choice was selected in conflict situations.¹⁰ Last, the dummy Other takes the value 1 for choices which are neither imitative nor best replies. Thus, the interaction Other×Conflict indicates choices which are neither imitative nor myopic best replies in case of conflict. Note that the reference group consists of decisions in case of alignment where the myopic best reply (which is also an imitative choice in this case) was selected.

With this choice of dummies, all our response-times hypotheses can be tested directly in the regressions. Model 1 in Table 2 tests for the basic effects. Models 2 and 3 add further Controls and Demographics¹¹ to show that the results are robust. All models include a treatment dummy for the presentation variants, which is never significant. Models 2 and 3 also add a Collusion dummy, taking the value 1 when the subject colluded with other subjects in a Cournot oligopoly. That coefficient is negative and highly significant, showing that the individuals who colluded were, unsurprisingly, fast. The inclusion of that dummy, however, does not affect other results.

(H1) states that best replies should be slower in case of conflict than in case of alignment. The comparison corresponds to the coefficient for Conflict, which is indeed positive and highly significant ($p = .0013$ in Model 1, $p < .0001$ in Models 2 and 3). (H2) predicts that myopic

¹⁰That is, the dummy takes the value 1 for imitative choices in case of conflict, and zero otherwise. Note that, since in case of alignment imitative choices are also best replies, this does not correspond to an interaction in the usual sense of the word.

¹¹Controls consist of a measure for normalized rounds, part 2 and part 3 dummies, and two payoff table dummies for possible medium or high payoffs. Demographics consist of age, gender, and an indicator capturing whether the subjects reported attending a game theory class.

Table 3: Panel Probit Regression Models for Myopic Best Reply.

Myopic Best Reply	Model 1	Model 2	Model 3
Conflict	−0.2183*** (0.0567)	−0.2163*** (0.0573)	−0.2158*** (0.0574)
FullInfo Treatment	−0.0159 (0.0413)	−0.0258 (0.0409)	−0.0294 (0.0461)
Collusion		−0.1815*** (0.0125)	−0.1573*** (0.0330)
Constant	−0.4092*** (0.0639)	−0.3704*** (0.0620)	−0.4225** (0.1862)
Controls	No	Yes	Yes
Demographics	No	No	Yes
Log Pseudolikelihood	−3612.3172	−3603.9443	−3603.3104
AME(Conflict)	−0.0728*** (0.0193)	−0.0719*** (0.0194)	−0.0717*** (0.0194)

Standard errors, clustered by 16 matching blocks, in parentheses.

AME=Average Marginal Effect. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

best replies should be slower than imitative choices, a comparison captured by the coefficient for the dummy Imitation-Conflict. The prediction is borne by the data, with the coefficient being negative and highly significant ($p < .0001$ in all models). (H3) predicts that, in case of alignment, best replies should be faster than other responses. The comparison reduces to the coefficient for the Other dummy, which is highly-significant and positive as expected ($p < .0001$ in all models).

The regressions also allow us to examine a number of exploratory questions. The linear combination of the coefficients Other and Other×Conflict is not significantly different from zero, i.e. in case of conflict we find no differences in response times between myopic best replies and other kinds of non-imitative decisions. In contrast, a linear combination test reveals that imitative decisions are significantly faster than other kinds of non-best-replies in conflict situations ($p < .0001$ in all models). This suggests that this latter category might include choices reflecting higher-level deliberation processes or more complex behavioral rules, as e.g. level- k considerations (best-replying to the anticipated best reply of others, etc; see Alós-Ferrer and Buckenmaier, 2018).

Tables 3 and 4 provide probit panel regressions with myopic best replies and imitative choices as dependent variables, respectively. Standard errors are again clustered at the block level. The independent variables are the Conflict dummy, a treatment dummy, a Collusion dummy, and further Controls and Demographics as in the previous regression models.

Table 3 allows us to parametrically test for Hypothesis (H4a), i.e. the prediction that myopic best replies are less likely under conflict than under alignment. This is confirmed by the negative and highly significant Conflict dummy, which is robust to the addition of Controls and Demographics (Model 1, $p = .0001$; Models 2 and 3, $p = .0002$). Analogously, Table 4 allows us to test for Hypothesis (H4b), i.e. the prediction that imitative choices are also less likely under

Table 4: Panel Probit Regression Models for Imitation.

Imitation	Model 1	Model 2	Model 3
Conflict	−0.0542 (0.0543)	−0.0754 (0.0470)	−0.0749* (0.0449)
FullInfo Treatment	−0.0159 (0.0752)	0.0282 (0.0802)	0.0120 (0.0739)
Collusion		0.6844*** (0.2523)	0.8105*** (0.1552)
Constant	−0.4557*** (0.0794)	−0.7209*** (0.1323)	−1.2212*** (0.1187)
Controls	No	Yes	Yes
Demographics	No	No	Yes
Log Pseudolikelihood	−3622.6196	−3601.3666	−3596.3697
AME(Conflict)	−0.0190 (0.0191)	−0.0261 (0.0166)	−0.0258* (0.0157)

Standard errors, clustered by 16 matching blocks, in parentheses

AME=Average Marginal Effect. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

conflict than under alignment. Although present in the data, this trend is clearly less strong than other predictions. The Conflict dummy is not significant in Model 1 ($p = .3184$), which does not control for collusion. The coefficient still misses significance in Model 2 ($p = .1087$) and becomes only weakly significant in Model 3 ($p = .0951$), after adding Controls, Demographics, and the Collusion dummy. As an additional exercise, we also ran an ordered probit panel regression (similar to Model 4 of Huck et al., 1999), which confirms our previous results and shows that both myopic best reply and imitation have a significant impact on decisions, with the latter having a larger relative weight.

In summary, the regression models confirm our non-parametric analysis while controlling for other features. Taken together, the analyses above provide strong evidence for Hypotheses (H1), (H2), (H3), and (H4a), and weak evidence for Hypothesis (H4b). Hence, we conclude that our experiment supports our model, suggesting that interacting behavioral rules with qualitatively different properties codetermine behavior in complex economic decisions.

Remark 1. The authors have recently run a number of experiments on an unrelated research question, namely the effects (or lack thereof) of cognitive load. One of the experiments in that project (manuscript in progress) replicates the experiment described here while manipulating cognitive load across two treatments with different subjects. Each of the two treatments (no load and load, $N = 72$ each) could be considered as a replication of the experiment in this paper. All results described above are confirmed in both treatments. Details are available upon request.

5 Imitating Yourself vs. Imitating Others

There are two qualitatively different kinds of imitative decisions. In some cases, it might happen that the decision maker’s own profits in the last period were the largest observed ones. In this case, imitation actually corresponds to “imitate yourself,” while other imitative decisions are of the type “imitate others.” Decisions where a player imitates him- or herself can also be conceived of as obeying *positive reinforcement*, which prescribes to repeat the previous choice if the player has “won,” that is, obtained the maximum observed profits. This corresponds to a simple “win-stay” version of *reinforcement learning*, i.e. the tendency to repeat what has worked in the past without paying attention to whether the conditions in which past actions were successful have changed. Reinforcement is particularly important for economics, as it captures the empirically-relevant focus on *past performance*, whose consequences are well-documented (e.g., outcome bias; Baron and Hershey, 1988; Dillon and Tinsley, 2008). Evidence from neuroscience has shown that reinforcement learning is associated with extremely fast and unconscious brain responses (e.g., Schultz, 1998; Holroyd and Coles, 2002). In an explicitly economic context, Achtziger and Alós-Ferrer (2014) showed that a simple reinforcement heuristic corresponds to a highly automatic process which competes with more deliberative rules when feedback comes in a win-loss frame.

Since reinforcement is generally considered to be rather automatic, we hypothesize that “imitating yourself” should be associated with shorter response times than “imitating others.” Note that imitating yourself and imitating others are never simultaneously active processes, but rather constitute a partition of imitative decisions and hence the prediction of faster response times is straightforward: the process favoring imitation is faster in one case than in the other, while the competing myopic best reply rule remains fixed.

Figure 3 displays the response times of decisions where participants imitated themselves or others. For completeness, we disentangle the comparison according to whether imitation (or positive reinforcement) was in conflict or in alignment with myopic best reply. Imitating-yourself decisions in case of conflict were significantly faster than imitating-others decisions (average 9.48 s vs. 11.86 s, WSR, $N = 16$, $z = -3.258$, $p = .0011$). In case of alignment, imitating-yourself decisions were also significantly faster than imitating-others decisions (average 9.62 s vs. 13.71 s; $N = 16$, $z = -2.947$, $p = .0032$). Thus, we confirm that the imitation behavioral rule that we consider might be supported by a composite process which, in some cases, reflects positive reinforcement. This is of independent interest, but does not change our previous conclusions.

Decisions following “imitating yourself” (or positive reinforcement) imply upholding the previously-selected action. Hence, they are aligned with a further, particularly simple behavioral rule: *decision inertia*, i.e. the tendency to repeat previous behavior independently of any feedback. This raises the natural question of whether the driver of the effects above is actually this simple but more general rule, i.e. whether inertia results in clear effects beyond situations where the decision maker has obtained the largest profits. Previous work (Alós-Ferrer et al., 2016) has compared decision inertia with reinforcement in the belief-updating task of Charness and Levin (2005) and Achtziger and Alós-Ferrer (2014), and found that inertia does cause asym-




Figure 3: Average Response Times of Imitating-yourself and Imitating-others Decisions. Stars indicate Wilcoxon Signed-Rank tests. ** $p < .05$ and *** $p < .01$.

metries in error rates, but this behavioral rule seems weaker than reinforcement and is typically washed away by it. Apesteguía et al. (2007) found that, in their Cournot-oligopoly experiment, subjects repeated their previous choice 12% to 23% of the time (depending on treatment) when they observed a higher-payoff strategy. To see whether inertia is behaviorally relevant in our paradigm, we examined it in the cases where it is *not* aligned with imitation, since in case of alignment with imitation we obtain positive reinforcement. To avoid confusion, however, we reserve the words “alignment” and “conflict” for the confluence or not of myopic best reply and imitation. In case of conflict between myopic best reply and imitation, we will compare “stay” myopic best replies (as prescribed by inertia) with “shift” myopic best replies. In case of alignment between myopic best reply and imitation, we test within other kinds of decisions not following the common prescription of myopic best reply and imitation.

Figure 4 depicts the response times of decisions in line with inertia (“stay” decisions) and those opposed to it (“shift” decisions). For conflict, the comparison is between “stay” and “shift” best replies. There were, however, no differences in the response times of these two kinds of decisions (stay, average 12.19 s; shift, 12.51 s; WSR, $N = 16$, $z = -0.776$, $p = .4380$). Hence, whenever myopic best reply and imitation conflict, there is no evidence of involvement of inertia (beyond the possible confluence with imitation), and in particular the effects of reinforcement described above are unlikely to be due to a more general process reflecting pure inertia.

For alignment (between imitation and myopic best reply), we compare all *non*-best replies of the “stay” and “shift” forms. Such stay (inertia) decisions were significantly faster than the




Figure 4: Average Response Times of Stay and Shift Decisions. Left: Comparison of best reply decisions in case of conflict. Right: Comparison of non-best replies in case of alignment. Stars indicate Wilcoxon Signed-Rank tests. ** $p < .05$ and *** $p < .01$.

comparable shift decisions (stay, average 12.13 s; shift, 14.48 s; WSR, $N = 16$, $z = -2.741$, $p = .0061$). In this case, best replies coincide with imitative decisions, that is, the “other” decisions we examine are not imitative, and in particular can *not* follow from positive reinforcement. Although this is speculative, this result suggest that shift decisions in this case might include choices derived from higher-order reasoning or more complex behavioral rules. This would be consistent with the long response times of “other” decisions under conflict discussed in the regression analyses in Section 4.

6 Conclusion

In a Cournot oligopoly experiment designed to maximize behavioral variance (as opposed to convergence), we find clear evidence in favor of the presence of multiple behavioral rules, one of them being imitation of successful, observed behavior. This is in line with previous experimental evidence on convergence to Walrasian outcomes in Cournot oligopolies (Huck et al., 1999, 2004; Offerman et al., 2002; Apesteguía et al., 2010), which is a prediction of models assuming imitative behavior (Vega-Redondo, 1997; Alós-Ferrer and Ania, 2005).

We rely on a simple formal model where each decision maker might follow either imitation or myopic best reply. The model makes a number of predictions which allow us to test for the multiplicity of behavioral rules, in the sense that none of the predictions would hold if only one

rule was present. This is possible because the predictions rely on explicit characteristics of the rules, in terms of their cognitive requirements. First, they rest on the *ex ante* classification of decisions in *conflict* or *alignment* according to the pre-specified rules. Second, they concern both choices *and* response times, the latter being a direct correlate of the postulated characteristics of the brain decision processes underlying the behavioral rules.

We find a number of “smoking guns,” all predicted by our model: best replies are slower under conflict with imitation than under alignment (generalizing the Stroop effect from cognitive psychology), they are slower than imitative decisions under conflict but faster than other decisions under alignment, and both best replies and imitative decisions are less frequent under conflict than under alignment. The evidence is striking and systematic, and, since it is based on process data, speaks in favor of a *literal* multiplicity of competing behavioral rules in economic decision making.

More generally, our model and empirical evidence support the view that economic decision making, even in strategic settings, might sometimes be better explained by integrating different views of behavior, instead of either assuming fully-rational optimization or boundedly-rational impulse-response behavior only. Multiple behavioral rules are more than a convenient metaphor or an *as if* model, and the analysis of human decisions can be improved by viewing them as the result of the interaction of different behavioral rules and decision processes in the human brain.

References

- Abbink, K. and J. Brandts (2008). 24. Pricing in Bertrand Competition with Increasing Marginal Costs. *Games and Economic Behavior* 63(1), 1–31.
- Achtziger, A. and C. Alós-Ferrer (2014). Fast or Rational? A Response-Times Study of Bayesian Updating. *Management Science* 60(4), 923–938.
- Allport, G. (1954). *The Nature of Prejudice*. Reading, Massachusetts: Addison-Wesley.
- Alós-Ferrer, C. (2018). A Dual-Process Diffusion Model. *Journal of Behavioral Decision Making* 31(2), 203–218.
- Alós-Ferrer, C. and A. B. Ania (2005). The Evolutionary Stability of Perfectly Competitive Behavior. *Economic Theory* 26, 179–197.
- Alós-Ferrer, C., A. B. Ania, and K. R. Schenk-Hoppé (2000). An Evolutionary Model of Bertrand Oligopoly. *Games and Economic Behavior* 33, 1–19.
- Alós-Ferrer, C. and J. Buckenmaier (2018). Cognitive Sophistication and Deliberation Times. Working Paper, University of Zurich.
- Alós-Ferrer, C., M. Garagnani, and S. Hügelschäfer (2016). Cognitive Reflection, Decision Biases, and Response Times. *Frontiers in Psychology* 7 (1402), 1–21.

- Alós-Ferrer, C., D.-G. Granić, J. Kern, and A. K. Wagner (2016). Preference Reversals: Time and Again. *Journal of Risk and Uncertainty* 52(1), 65–97.
- Alós-Ferrer, C., S. Hügelshäfer, and J. Li (2016). Inertia and Decision Making. *Frontiers in Psychology* 7 (169), 1–9.
- Alós-Ferrer, C. and A. Ritschel (2018). The Reinforcement Heuristic in Normal Form Games. *Journal of Economic Behavior and Organization* 152, 224–234.
- Alós-Ferrer, C. and F. Strack (2014). From Dual Processes to Multiple Selves: Implications for Economic Behavior. *Journal of Economic Psychology* 41, 1–11.
- Apesteguía, J., S. Huck, and J. Oechssler (2007). Imitation—Theory and Experimental Evidence. *Journal of Economic Theory* 136, 217–235.
- Apesteguía, J., S. Huck, J. Oechssler, and S. Weidenholzer (2010). Imitation and The Evolution of Walrasian Behavior: Theoretically Fragile But Behaviorally Robust. *Journal of Economic Theory* 145, 1603–1617.
- Baddeley, A. D., D. Chincotta, and A. Adlam (2001). Working Memory and the Control of Action: Evidence From Task Switching. *Journal of Experimental Psychology: General* 130(4), 641–657.
- Bandura, A. (1977). *Social Learning Theory*. Englewood Cliffs, NJ: Prentice Hall.
- Bargh, J. A. (1989). Conditional Automaticity: Varieties of Automatic Influences in Social Perception and Cognition. In J. S. Uleman and J. A. Bargh (Eds.), *Unintended Thought*, pp. 3–51. Guilford.
- Baron, J. and J. C. Hershey (1988). Outcome Bias in Decision Evaluation. *Journal of Personality and Social Psychology* 54(4), 569–579.
- Blume, L. (1993). The Statistical Mechanics of Strategic Interaction. *Games and Economic Behavior* 5, 387–424.
- Bosch-Domènech, A. and N. J. Vriend (2003). Imitation of Successful Behaviour in Cournot Markets. *Economic Journal* 113(487), 495–524.
- Buckert, M., J. Oechssler, and C. Schwieren (2017). Imitation under Stress. *Journal of Economic Behavior and Organization* 139, 252–266.
- Charness, G. and D. Levin (2005). When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect. *American Economic Review* 95(4), 1300–1309.
- Chartrand, T. L. and J. A. Bargh (1999). The Chameleon Effect: The Perception-Behavior Link and Social Interaction. *Journal of Personality and Social Psychology* 76(6), 893–910.

- Cohen, J. D., K. Dunbar, and J. L. McClelland (1990). On the Control of Automatic Processes: A Parallel Distributed Processing Account of the Stroop Effect. *Psychological Review* 97(3), 332–361.
- Cracco, E., L. Bardi, C. Desmet, O. Genschow, D. Rigoni, L. De Coster, I. Radkova, E. Deschrijver, and M. Brass (2018). Automatic Imitation: A Meta-Analysis. *Psychological Bulletin* 144(5), 453–500.
- Dillon, R. L. and C. H. Tinsley (2008). How Near-Misses Influence Decision Making Under Risk: A Missed Opportunity for Learning. *Management Science* 54(8), 1425–1440.
- Evans, J. S. (2008). Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition. *Annual Review of Psychology* 59, 255–278.
- Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics* 10(2), 171–178.
- Fischbacher, U., R. Hertwig, and A. Bruhin (2013). How to Model Heterogeneity in Costly Punishment: Insights from Responders’ Response Times. *Journal of Behavioral Decision Making* 26(5), 462–476.
- Friedman, D., S. Huck, R. Oprea, and S. Weidenholzer (2015). From Imitation to Collusion: Long-Run Learning in a Low-Information Environment. *Journal of Economic Theory* 155, 185–205.
- Fudenberg, D., P. Strack, and T. Strzalecki (2018). Speed, Accuracy, and the Optimal Timing of Choices. *American Economic Review* 108(12), 3651–3684.
- Greiner, B. (2015). Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE. *Journal of the Economic Science Association* 1, 114–125.
- Holroyd, C. B. and M. G. Coles (2002). The Neural Basis of Human Error Processing: Reinforcement Learning, Dopamine, and the Error-Related Negativity. *Psychological Review* 109, 679–709.
- Horner, V. and A. Whiten (2005). Causal Knowledge and Imitation/Emulation Switching in Chimpanzees (*Pan Troglodytes*) and Children (*Homo Sapiens*). *Animal Cognition* 8, 164–181.
- Huck, S., H.-T. Normann, and J. Oechssler (1999). Learning in Cournot Oligopoly - An Experiment. *Economic Journal* 109, C80–C95.
- Huck, S., H.-T. Normann, and J. Oechssler (2000). Does Information about Competitors’ Actions Increase or Decrease Competition in Experimental Oligopoly Markets? *International Journal of Industrial Organization* 18(1), 39–57.
- Huck, S., H.-T. Normann, and J. Oechssler (2002). Stability of the Cournot Process — Experimental Evidence. *International Journal of Game Theory* 31, 123–136.

- Huck, S., H.-T. Normann, and J. Oechssler (2004). Two Are Few and Four Are Many: Number Effects in Experimental Oligopolies. *Journal of Economic Behavior and Organization* 53(4), 435–446.
- Kahneman, D. (2003). Maps of Bounded Rationality: Psychology for Behavioral Economics. *American Economic Review* 93(5), 1449–1475.
- Kandori, M., G. J. Mailath, and R. Rob (1993). Learning, Mutation, and Long Run Equilibria in Games. *Econometrica* 61, 29–56.
- Krajbich, I., B. Bartling, T. Hare, and E. Fehr (2015). Rethinking Fast and Slow Based on a Critique of Reaction-Time Reverse Inference. *Nature Communications* 6(7455), 1–9.
- Ludwig, J., F. K. Ahrens, and A. Achtziger (2020). Errors, Fast and Slow: An Analysis of Response Times in Probability Judgments. *Thinking and Reasoning* forthcoming.
- MacCleod, C. M. (1991). Half a Century of Research on the Stroop Effect: An Integrative Review. *Psychological Bulletin* 109(2), 163–203.
- Moffatt, P. G. (2005). Stochastic Choice and the Allocation of Cognitive Effort. *Experimental Economics* 8(4), 369–388.
- Oechssler, J., A. Roomets, and S. Roth (2016). From Imitation to Collusion: A Replication. *Journal of the Economic Science Association*, 1–9.
- Offerman, T., J. Potters, and J. Sonnemans (2002). Imitation and Belief Learning in an Oligopoly Experiment. *Review of Economic Studies* 69(4), 973–997.
- Offerman, T. and A. Schotter (2009). Imitation and Luck: An Experimental Study on Social Sampling. *Games and Economic Behavior* 65(2), 461–502.
- Offerman, T. and J. Sonnemans (1998). Learning by Experience and Learning by Imitating Successful Others. *Journal of Economic Behavior and Organization* 34, 559–575.
- Ratcliff, R. (1978). A Theory of Memory Retrieval. *Psychological Review* 85, 59–108.
- Ratcliff, R. and J. N. Rouder (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science* 9(5), 347–356.
- Rubinstein, A. (2007). Instinctive and Cognitive Reasoning: A Study of Response Times. *Economic Journal* 117(523), 1243–1259.
- Rubinstein, A. (2016). A Typology of Players: Between Instinctive and Contemplative. *Quarterly Journal of Economics* 131(2), 859–890.
- Schaffer, M. E. (1988). Evolutionarily Stable strategies for a Finite Population and a Variable Contest Size. *Journal of Theoretical Biology* 132, 469–478.

- Schaffer, M. E. (1989). Are Profit-Maximisers the Best Survivors? *Journal of Economic Behavior and Organization* 12, 29–45.
- Schneider, W. and R. M. Shiffrin (1977). Controlled and Automatic Human Information Processing: I. Detection, Search, and Attention. *Psychological Review* 84(1), 1–66.
- Schultz, W. (1998). Predictive Reward Signal of Dopamine Neurons. *Journal of Neurophysiology* 80(1), 1–27.
- Selten, R. and A. Ostmann (2001). Imitation Equilibrium. *Homo Oeconomicus* 18, 111–149.
- Shadlen, M. N. and D. Shohamy (2016). Decision Making and Sequential Sampling from Memory. *Neuron* 90, 927–939.
- Shiffrin, R. M. and W. Schneider (1977). Controlled and Automatic Human Information Processing: II. Perceptual Learning, Automatic Attending and a General Theory. *Psychological Review* 84(2), 127–190.
- Spiliopoulos, L. (2018). The Determinants of Response Time in a Repeated Constant-sum Game: A Robust Bayesian Hierarchical Dual-Process Model. *Cognition* 172, 107–123.
- Spiliopoulos, L. and A. Ortmann (2018). The BCD of Response Time Analysis in Experimental Economics. *Experimental Economics* 21(2), 383–433.
- Stroop, J. R. (1935). Studies of Interference in Serial Verbal Reactions. *Journal of Experimental Psychology* 35, 643–662.
- Vega-Redondo, F. (1997). The Evolution of Walrasian Behavior. *Econometrica* 65(2), 375–384.
- Weber, E. U. and E. J. Johnson (2009). Mindful Judgment and Decision Making. *Annual Review of Psychology* 60, 53–85.
- Williams, G. A. and R. B. Miller (2002). Change the Way You Persuade. *Harvard Business Review* 80(5), 64–73.

A Proofs

Proof of Theorem 1. The expected response time of best replies in case of alignment ($x^B = x^I$) is

$$E(RT|x^B, \text{Alignment}) = \frac{(1 - \Delta)P^{BR}R^B + \Delta P^{Im}R^I}{(1 - \Delta)P^{BR} + \Delta P^{Im}}$$

and the expected response time of best replies in case of conflict ($x^B \neq x^I$) is

$$E(RT|x^B, \text{Conflict}) = \frac{(1 - \Delta)P^{BR}R^B + \Delta P_B^{Im}R^I}{(1 - \Delta)P^{BR} + \Delta P_B^{Im}}.$$

Then, $E(RT|x^B, \text{Conflict}) > E(RT|x^B, \text{Alignment})$ holds if and only if

$$(P^{Im} - P_B^{Im})R^B > (P^{Im} - P_B^{Im})R^I$$

which holds by (P1) and (R). \square

Proof of Theorem 2. (H2) The expected response time of best replies in case of conflict ($x^B \neq x^I$) is as given in the proof of Theorem 1, and the expected response time of imitative answers is

$$E(RT|x^I, \text{Conflict}) = \frac{(1 - \Delta)P_I^{BR}R^B + \Delta P^{Im}R^I}{(1 - \Delta)P_I^{BR} + \Delta P^{Im}}$$

where P_I^{BR} denotes the probability with which the best reply rule selects an imitative answer when it does not coincide with the prescription of imitation, i.e. $P_I^{BR} = P^{BR}(x^I)$ when $x^B \neq x^I$. Then, $E(RT|x^B, \text{Conflict}) > E(RT|x^I, \text{Conflict})$ if and only if

$$(P^{BR}P^{Im} - P_I^{BR}P_B^{Im})(R^B - R^I) > 0$$

which holds by (R) ($R^B > R^I$) and (P1) (which implies $P^{BR}P^{Im} > P_I^{BR}P_B^{Im}$).

(H3) The expected response time of best replies in case of alignment ($x^B = x^I$) is as given in the proof of Theorem 1, and the expected response time of other answers is

$$E(RT|x \neq x^B, \text{Alignment}) = \frac{(1 - \Delta)(1 - P^{BR})R^B + \Delta(1 - P^{Im})R^I}{(1 - \Delta)(1 - P^{BR}) + \Delta(1 - P^{Im})}.$$

Then, $E(RT|x^B, \text{Alignment}) < E(RT|x \neq x^B, \text{Alignment})$ if and only if

$$((1 - P^{BR})P^{Im} - P^{BR}(1 - P^{Im}))(R^B - R^I) > 0.$$

Since $R^B > R^I$ holds by (R), the result holds if $(1 - P^{BR})P^{Im} > P^{BR}(1 - P^{Im})$, which is equivalent to $P^{Im} > P^{BR}$. The latter holds by (P2). \square

Proof of Theorem 3. (H4a) The proportion of best replies in case of alignment ($x^B = x^I$) is $P(BR|\text{Alignment}) = (1 - \Delta)P^{BR} + \Delta P^{Im}$, and the proportion of best replies in case of conflict ($x^B \neq x^I$) is $P(BR|\text{Conflict}) = (1 - \Delta)P^{BR} + \Delta P_B^{Im}$, where P_B^{Im} denotes the probability with

which the imitation rule selects a best reply when it does not coincide with the prescription of imitation, i.e. $P_B^{Im} = P^{Im}(x^B)$ when $x^B \neq x^I$. Then, $P(BR|Alignment) > P(BR|Conflict)$ if and only if $P^{Im} > P_B^{Im}$, which holds by (P1).

(H4b) is analogous to (H4a). □

B Translated Instructions

General Instructions

The experiment consists of **three parts with 17 rounds each** in which **you and three other participants** make decisions. After the completion of these three parts, a questionnaire will follow. In each of the three parts you will **earn points**. How many points you earn depends on your decisions and the decisions of the players in your group. All points you earned each round will be added up at the end of the experiment and exchanged into Euros. The exchange rate is:

$$1000 \text{ points} = 18 \text{ Eurocents.}$$

Independently of your decisions, you will receive 2.50 EUR for your participation. The total amount will be paid **in cash** and **anonymously** at the end of the experiment.

On the following pages you will receive all further information which you need for the experiment. Among other things the sequence of the experiment will be explained in detail. Once you have finished reading the instructions, please proceed to answer the control questions on the screen.

Instructions of the Experiment

General Sequence: The experiment is divided into **three parts**. The **procedure is the same** for each part. Only the payoff table (which will be discussed later in more detail) and the composition of the groups change with each part. One part consists of **17 rounds**. At the beginning of each part, participants will be divided into groups. One **group** consists of **4 players** (you included) and stays the same for the duration of a part. That means, that you always interact with the same players during one part. In every new part two of players will be replaced and therefore the composition of the group **changes**. That means that in a new part you do not interact with the same players as in the previous part. In every round you have to decide among four options, *A*, *B*, *C*, or *D*. How many points you earn in one round depends both on your choice and on the choices of the other three group members.

Payoff Tables: The payoff tables are an important component of the experiment. They show you all possible payoffs depending on your choice and the choice of the other three group members. The rows represent your choice and the columns represent the joint choice of the other group members. The appropriate cell entry is the amount of points you would receive if this combination of choices occurs. Please note that for your payoff it is irrelevant which of the other group members made which choice. That means that if the other group members choose

C , A , and B , respectively, this has the same effect on your payoff than if they choose A , B , and C . For a better overview, columns are ordered alphabetically.

Figures B.1-B.3 depict examples of such payoff tables. Please note that in the experiment other payoff tables will be used. **Important note:** The payoff table will not change during a part. The same payoff table applies to all group members.

Your Decision: In each round you have to choose one of the four options, A , B , C , or D . You have 30 seconds to make your choice. You make a choice by clicking on the appropriate button on the screen. During your choice the payoff table of the current part will be shown. The next round begins as soon as all participants made their choice. In every round the result of the previous round is shown (except for the first round of every part).

Sequence of Decisions in a Round in Detail: The **payoff table** will be shown at the beginning of each part so you can familiarize yourself with it (see Figure B.1¹²). The table will be kept on the screen during the experiment at all times – you **do not** have to memorize or copy the table. After you have familiarized yourself with the table click “continue”. The **decision phase** will start as soon as all participants are ready.

Round 1 of 51																				
	AAA	AAB	AAC	AAD	ABB	ABC	ABD	ACC	ACD	ADD	BBB	BBC	BBD	BCC	BCD	BDD	CCC	CCD	CDD	DDD
A	1000	1200	750	825	1400	950	1025	500	575	650	1600	1150	1225	700	775	850	250	325	400	475
B	816	952	646	697	1088	782	833	476	527	578	1224	918	969	612	663	714	306	357	408	459
C	1050	1330	700	805	1610	980	1085	350	455	560	1890	1260	1365	630	735	840	0	105	210	315
D	1056	1312	736	832	1568	992	1088	416	512	608	1824	1248	1344	672	768	864	96	192	288	384

For the next 17 rounds (1-17) this will be the payoff table.

Click "continue" so the experiment can continue.

continue

Figure B.1: Beginning of a part.

Now you can choose among **four options**, A , B , C , and D . To make a choice click on the appropriate button of your choice (see Figure B.2).

Starting in round two of a part, the **results from the previous round** will be shown (see Figure B.3). In the first column, “Result,” you see the choices of all four players in the group. In the example figure it was “ B , D , B , C .” The first letter (“ B ”) always represents your own choice whereas the following three letters (“ D , B , C ”) represent the choice of the other three

¹²The payoff tables in the instructions were just illustrative. They were computed using $P(Q) = 140 - Q$, $A = 25$, $B = 17$, $C = 35$, $D = 33$.

Round 1 of 51																				
	AAA	AAB	AAC	AAD	ABB	ABC	ABD	ACC	ACD	ADD	BBB	BBC	BBD	BCC	BCD	BDD	CCC	CCD	CDD	DDD
A	1000	1200	750	825	1400	950	1025	500	575	650	1600	1150	1225	700	775	850	250	325	400	475
B	816	952	646	697	1088	782	833	476	527	578	1224	918	969	612	663	714	306	357	408	459
C	1050	1330	700	805	1610	980	1085	350	455	560	1890	1260	1365	630	735	840	0	105	210	315
D	1056	1312	736	832	1568	992	1088	416	512	608	1824	1248	1344	672	768	864	96	192	288	384

Please make a choice:

A

B

C

D

Figure B.2: Decision in the first round of a part.

members of your group. The position of a letter (choice of a group member) is always assigned to a specific group member and stays the same during a given part. In the example the “left” player chose *D*, the “middle” player chose *B*, and the “right” player chose *C*.

Round 2 of 51																				
	AAA	AAB	AAC	AAD	ABB	ABC	ABD	ACC	ACD	ADD	BBB	BBC	BBD	BCC	BCD	BDD	CCC	CCD	CDD	DDD
A	1000	1200	750	825	1400	950	1025	500	575	650	1600	1150	1225	700	775	850	250	325	400	475
B	816	952	646	697	1088	782	833	476	527	578	1224	918	969	612	663	714	306	357	408	459
C	1050	1330	700	805	1610	980	1085	350	455	560	1890	1260	1365	630	735	840	0	105	210	315
D	1056	1312	736	832	1568	992	1088	416	512	608	1824	1248	1344	672	768	864	96	192	288	384

Results from the previous round		
Results	Your Choice and Points	Choice with the most Points
Choice of all Players: B, D, B, C	Choice: B Points: 663	Choice: C Points: 1365

Please make a choice:

A

B

C

D

Figure B.3: *BestOnly treatment*: New choice and the result of the previous round starting from the second round.

In the second column, “Your Choice and Points,” you will see your own choice and the points you earned in the previous round. In the example in Figure B.3 you can see in the payoff

Round 2 of 51																				
	AAA	AAB	AAC	AAD	ABB	ABC	ABD	ACC	ACD	ADD	BBB	BBC	BBD	BCC	B C D	BDD	CCC	CCD	CDD	DDD
A	1000	1200	750	825	1400	950	1025	500	575	650	1600	1150	1225	700	775	850	250	325	400	475
B	816	952	646	697	1088	782	833	476	527	578	1224	918	969	612	663	714	306	357	408	459
C	1050	1330	700	805	1610	980	1085	350	455	560	1890	1260	1365	630	735	840	0	105	210	315
D	1056	1312	736	832	1568	992	1088	416	512	608	1824	1248	1344	672	768	864	96	192	288	384
Results from the previous round																				
Results		Your Choice and Points				Choice and Points of the Group Members														
Choice of all Players: B, D, B, C		Choice: B Points: 663				Choice: D Points: 1248					Choice: B Points: 663					Choice: C Points: 1365				
Please make a choice:						A		B		C		D								

Figure B.3: *FullInfo treatment*: New choice and the result of the previous round starting from the second round.

table that you earned 663 points because you chose *B* (row “*B*” in the table) and the other group members chose *D*, *B*, *C* (column “*B C D*”). In the table, the column representing the choice combination of the group members is highlighted in yellow. The columns are ordered alphabetically for a better overview.

[*BestOnly treatment*:] The last column, “Choice with the most Points,” shows the choice which earned the most points in the previous round and the points earned with that choice. This information is always highlighted in yellow.

[*FullInfo treatment*:] The last column, “Choice and Points of the Other Group Members,” shows the choices and how many points the other group members earned in the previous round. The ordering of the group members is the same as in the first column, “Result” (“left” player – *D*, “middle” player – *B*, and “right” player – *C*). The choice and points of the player who earned the most points in the previous round is highlighted in yellow. The column “Your Choice and Points” is also highlighted in yellow if you earned the most points in the previous round. In case of a tie the choice and points of multiple players will be highlighted.

Do you have any questions? If so, please raise your hand and wait quietly.