UNIVERSITÄT ST. GALLEN

VOLKSWIRTSCHAFTLICHE ABTEILUNG

Learning in Economics: Where Do We Stand? A Behavioral View on Learning in Theory, Practice and Experiments

Tilman Slembeck

August 1999 Discussion paper no. 9907

DEPARTMENT OF ECONOMICS UNIVERSITY OF ST. GALLEN

August 1999
Discussion Paper No. 9907

Learning in Economics: Where Do We Stand?
A Behavioral View on Learning in Theory, Practice and Experiments*

Tilman Slembeck

Abstract

This paper briefly reviews the current literature on learning in economics from a behavioral point of view. It critically compares theory with aspects of learning in real-life and with evidence from laboratory experiments, and argues that most customary approaches lack criteria for their applicability. Hence, there is a need for a theory that includes criteria when to employ which theory or which element(s) of existing theories contingent on the situation or environment in question. A discussion of several unsolved issues in economic learning stresses the fundamental role of learning conditions that have be neglected in the literature, but are accounted for in behavioral approaches such as "contingent learning".

JEL Classification: B4, C9, D8
Keywords: economic learning, behavioral economics, economic experiments, game theory, information, feedback, contingent learning.

Author's address: Tilman Slembeck
University of St.Gallen
Varnbuelstrasse 19
9000 St.Gallen / Switzerland
Tel: ++ 41 71 224 25 49
Fax: ++ 41 71 224 35 38
E-Mail: tilman.slembeck@unisg.ch

Publisher: Forschungsgemeinschaft für Nationalökonomie
an der Universität St.Gallen
Dufourstrasse 48
CH-9000 St.Gallen
Tel: ++ 41 71 224 23 00
Fax: ++ 41 71 224 26 46

---

*This paper is based on an invited talk given at the International Workshop on Evolution and Learning in Games, June 24–26, 1999, Max-Planck-Institute for Research into Economic Systems, Evolutionary Economics Unit, Jena, Germany. – I would like to thank the participants for interesting discussions and suggestions, especially Tilman Börgers, Thomas Brenner, Ido Erev, Marco Lehmann-Waffenschmidt, Dov Monderer, Rajiv Sethi, Fernando Vega-Redondo, and Nick Vriend. Special thanks go to my discusstant Ulrich Witt for his most stimulating comments.

1. Introduction

Learning is an important feature of human behavior. Most social scientists and many economists agree. Since learning is a relatively new topic in economics, there is a large variety of approaches that are not easy to classify (see FUDENBERG & LEVINE, 1998, for a comprehensive overview of learning in games). Each approach involves different methods, and – most importantly– pursues goals that may be different. Hence, one difficulty with the current learning literature is that it is often not clear what a given theory aims to explain beyond some narrowly defined theoretical or technical problem. Thus, I maintain that the literature would benefit greatly from making the aims that an approach or theory pursues more explicit, not only at a technical level, but also with respect to some phenomenon to be explained. My suspicion, however, is that much effort is directed toward improving theory for theoretical reasons, not toward explaining observable phenomena.

Thus, I am wondering what the learning literature has to offer to the average member of the average economic department. I believe that a better job should be done in explaining why learning processes are important in economics more generally, and what kinds of results learning theory has produced that are of interest to economist not specialized in this field. To do so, one has to ask what can be learned from learning in theory *and* in experiments, and how these results relate to real-world situations.

What are the contributions of learning theory to economic theory more generally? That is: why should a theorist care? – How relevant are our results to explain real phenomena and how can these results be applied? That is: why should an applied economist bother? In sum: what are the robust findings we can tell our colleagues about? Or is learning just another highly specialized field that only insiders care about?

I cannot offer any definite answers to these questions. Instead, in this paper I shall outline some current issues in economic learning, mainly from a behavioral perspective. This perspective maintains that theories about (human) behavior, –such as economic theories– can be improved by providing a stronger empirical basis through careful observation of actual behavior and by confronting theory with evidence. The aim of the behavioral (not behavioristic) approach is to broaden the view on economic behavior and go beyond the arm-chair reasoning of much economic theorizing. This is not to disregard the value or role of pure theory as a benchmark –telling us how behavior should be under mostly ideal assumptions and conditions at the limit–, but to put this theory on a stronger behavioral basis by including psychological insights collected in experiments and the field. Clearly, the aim of such effort is not only to improve theory but also to apply it to observable phenomena, and to highlight the limits it is subject to.

With regard to learning behavior, for example, the question arises in which cases or under what circumstances one can be confident that the strong assumptions made in theory do hold:

> *“The assumption that conduct is prompt and rational is in all cases a fiction. But it proves to be sufficiently near to reality, if things have had time to hammer logic into men. Where this has happened, and within the limits in which it has happened, one may rest content with this fiction and build theories upon it.“* Joseph Schumpeter (1934, 80; first published 1911)

It seems important to note that most traditional and learning theory does not include statements about the circumstances or situations in which the involved assumptions are supposed to hold. Hence, the *limits* imposed by these assumptions are usually not discussed – especially not with regard to learning behavior and learning conditions. The basic question is when and under what circumstances we can be confident that “*things have had time to hammer logic into men.*“ From a behavioral view it seems that many theories make bolder claims than can be supported by behavioral evidence, not only with respect to cognitive abilities and capacities (hence, the rationality assumption), but also with respect to situational constraints for learning. The role of these constraints, that is of learning conditions, will be the main theme of this paper.

In Section 2 the existing learning literature will briefly be summarized and commented on. Section 3 focuses on some aspects of learning “in practice“ that deserve more attention than usually attributed in theory. In Section 4 some problems of learning in experiments will be discussed, and Section 5 concludes.

2. Learning in Theory

As mentioned in the introduction, the goals one pursues when doing theory may differ substantially. There are at least two general views on (game) theory:

view A) *Game theory is for doing theory, not for playing games.*
view B) *We need a theory for playing games, not just for doing theory!*

It does not always seem clear which of these two extremes a given theory belongs to.[1] Some (possibly most) theories have purely theoretical aims, while others are inspired by empirical findings. Still others mainly focus on observable behavior and are able to organize experimental data quite well. But since the results of “horse races“ among theories in the laboratory are still mixed, it remains an open question which approach fits best for which type of situation or

[1] The first view has been attributed to Reinhard Selten, e.g. by GEOREE & HOLT (1999a) who favor the second view, which is also typical for the work of Ido Erev and Al Roth (see EREV & ROTH, 1998).

“overall“. An additional unsettled question is if and to what extent the results of “fit-theory-to-data exercises“ can be applied to a wider range of phenomena possibly outside the laboratory (i.e, the question of the external validity of results).

Today, the literature offers a variety of approaches to learning in economics or in games. There are two branches of mainly theoretical approaches (view *A* above). The first involves *individual-based learning mechanisms*, while the second focuses on the *aggregate level*. In contrast, *experimental approaches* support view *B*, in that they aim to account for behavior observable in laboratory experiments. In the same vein, *behavioral approaches* are based on behavioral evidence and include findings that are considered robust in psychology. In addition, the contingent learning approach discussed below stresses the importance of learning conditions.

Individual Level Approaches

Approaches of individual learning model learning behavior of a representative actor, in that some rule or mechanism is introduced according to which the actor updates his beliefs when new information about the environment arrives. Typical statistical rules include *Bayesian* (or “rational“) learning or *least-squares learning* (see BLUME & EASLEY, 1995). In *imitation learning*, actors are thought to copy the behavior of others, especially behavior that is popular or appears to yield high payoffs (see e.g., SCHLAG, 1998 and 1999).[2] In models of *best reply learning*, actors behave more strategically, in that they aim to maximize their payoffs given their expectations about what others will do. In simple versions of this type of learning, known as “*fictitious play*“, actors choose their best reply to the observed frequency distribution of their opponents‘ previous actions (see e.g., YOUNG, 1998).

Aggregate Level Approaches

It is typical for aggregate level approaches that actors or players are assumed to interact repeatedly within (possibly large) populations. In *evolutionary approaches,* actors are usually programmed to behave in a given and fixed manner or to play a certain strategy.[3] Actors that use a higher-than-average-payoff strategy are at a reproductive advantage so that their frequency within the population increases over time. The most widely used model is known as the “*replicator dynamics*“ (see FUDENBERG & LEVINE, 1998, pp. 51). It assumes that the growth rate of a certain player type or strategy in a population is a linear function of its payoff relative to the average payoff. – Also at the aggregate level, there are approaches that model learning in *neural*

[2] Recent experimental studies on imitation include HUCK, NORMANN & OECHSSLER (1999) and BOSCH & VRIEND (1999).

[3] On the sub-field of evolutionary game theory, see WEIBULL (1995) and SAMUELSON (1997).

networks (e.g., RUMELHARD & MCCLELLAND, 1986) or via *genetic algorithms* (e.g., HOLLAND, 1975).

Equilibrium Selection and Equilibrium Concepts

Note that most theoretical approaches – be it at the individual or at the aggregate level – have mainly been motivated by multiple equilibria in games (or in rational expectations models). Thus, they focus on well-defined equilibrium selection processes. The driving force behind such effort is rooted in the very concept of equilibrium in that most economists perceive it as their duty to find and formally describe some state where behavior is no longer modified (and possibly no more learning occurs). The two most basic concepts are the market equilibrium and the Nash equilibrium.

In *non-strategic settings* of simple dependence like perfect markets with an atomistic structure, free information, homogeneous goods, and costless transactions adjustment or learning is relatively straightforward because there is only one observational variable, namely the market price, and only one behavioral variable, namely the quantity that is supplied or demanded. *Market equilibrium* is reached simply via adjustment of quantities if the current price is above or below the equilibrium price.

In contrast, in *strategic settings* of interdependence –typically modeled as games– adaption or learning is often not so simple since the "environment" is not only reacting to a player's behavior, it is also acting strategically and may anticipate behavior. This introduces *strategic uncertainty* about the behavior of others, which may complicate learning substantially. The most basic equilibrium concept in game theory –and possibly in economic theory– is the *Nash equilibrium*. It describes a situation in which none of the involved players (or market participants) has an incentive to change her current behavior given that all others do not change their current behaviors either (see e.g., BINMORE, 1992, for a formal definition).

The problem with this concept, however, is that it does not include any statements about how a Nash equilibrium is reached. Also, even in simple 2 x 2 matrix games with only 2 players and 2 strategies more than one Nash equilibrium may exist. Therefore most theory focuses on the following question: "Suppose that players reach a Nash equilibrium, and suppose that there are several Nash equilibria, which Nash equilibrium is most likely to be reached?" The majority of literature aims to answer this question via theoretical refinements of the Nash concept or by introducing alternative equilibrium concepts (see FUDENBERG & TIROLE, 1992, for an overview of the refinement literature), while the learning literature tries to answer it via some learning mechanism (see MAILATH, 1998, for a survey).

Logically prior to the refinement question, however, is the question "When or under what circumstances does learning lead to Nash equilibrium?" This question, again, stresses the importance of learning conditions, and it has recently been addressed by BÖRGERS (1999) who presents a list of factors that are assumed to determine whether subjects approach Nash equilibrium in experiments. This list is somewhat similar to the list of learning determinants presented in the contingent learning approach below, and Börger‘s approach seems one of the rare exceptions where theorists aim to address the basic question of learning conditions.

Aside from the last mentioned efforts, theoretical learning approaches are mostly not motivated by some behavioral or psychological criticism, and are not based on behavioral theories or empirical evidence about human learning. Hence, it appears that most of these analyses aim to predict behavior entirely by theory. Furthermore, definite predictions usually hold only under strong assumptions that may be reasonable in some ideal situations, but are unrealistic and potentially misleading for many applications. There seems to be an evident lack of knowledge about some of the very principles that govern learning behavior that requires not only theoretical progress, but also systematic observation and careful empirical work (see CRAWFORD, 1995).

Experimental Approaches

This type of approach aims to model how people behave or play games in experimental laboratories. Hence, they support view *B* mentioned above. The most often studied theory is *reinforcement learning* which basically assumes that people adopt actions or strategies that have been successful in the past. This assumption is one of the cornerstones of behavioral psychology, and it is based on the "Law of Effect" (Thorndike, 1898). Reinforcement learning has widely been studied in animals and humans by psychologists, and has attracted the attention of economists (see EREV & ROTH, 1998). This type of learning is *non-cognitive* in that individuals do not systematically reflect their behavior, but simply use any high-payoff strategy with increasing probability regardless of the situation.[4] It has therefore been labeled a "low rationality theory" (see ROTH & EREV, 1995). Reinforcement learning is also *non-strategic* because actors are assumed to consider only their own actions and outcomes, not those of others. This is a main difference to models of best reply learning, like fictitious play, where actors are assumed to know and consider the actions and payoffs of others.

Another experimentally oriented approach is the *learning direction theory*, proposed by Reinhard Selten (see SELTEN & STOECKER, 1986). It is based on the idea that after some experience, people think about what might have been a better decision last time, and then adjust their behavior in that direction (SELTEN & BUCHTA, 1999). Hence, directional learning makes only qualitative predictions, but does not attempt to predict the amount of change.

---

[4] Hence, reinforcement learning implies that people ignore information about behavior and outcomes of others, even if this information is free. Therefore, people are assumed to be unaware of whether they are trading in a market, playing a game or facing an individual decision making task (see also OCHS, 1999).

These and other approaches have been experimentally tested many times (see e.g., STAHL, 1996; DUFFY & NAGEL 1997; HO ET AL. 1998). To date, it seems fair to conclude that different learning models are able to describe experimental behavior in various contexts. Especially reinforcement learning seems to perform surprisingly well in many environments, while Bayesian learning, fictitious play, and Cournot dynamics are usually not supported by the data (see also NAGEL, 1999). However, learning appears to be highly sensitive to learning conditions and information, at least in some situations (see e.g., HUCK ET AL. 1999). None of the leading frameworks gives a fully reliable account of behavior *by itself* and *"...most behavior can be understood in terms of a synthesis of ideas from those frameworks, combined with empirical knowledge in proportions that depend in predictable ways on the environment* (CRAWFORD, 1995, 3). Hence, one would like to have a unified approach that includes criteria when to employ which theory or which element(s) of existing theories *contingent* on the situation or environment in question.

A Behavioral Approach: Contingent Learning

One attempt in this direction is to start from behavioral features documented in the psychological literature. For example, it is well known in this literature that decision making is *contingent* on the environment or situation (see HOGARTH & REDER, 1987; PAYNE, 1982). Therefore, it seems natural to assume that learning is contingent on the situation as well.

This idea has motivated the *contingent learning approach* (SLEMBECK, 1998a), who studies the role of situational characteristics as restrictions to learning. As will be discussed in Section 3, in most practical situations or applications conditions for learning are far less ideal than assumed in standard and learning theories where learning is often unconstrained with respect to these conditions. Regarding learning in experiments (see Section 4) it should be noted that also most experiments provide the same type of ideal learning conditions as assumed in theory, because many experimenters aim at "*speaking to theorists*" instead of "*searching for facts*" (ROTH, 1995, 22).[5]

The basic question which the contingent learning approach tries to answer is how learning behavior depends on learning conditions. These conditions can be understood as situational restrictions to learning, in addition to the cognitive restrictions that approaches of bounded rationality aim to account for. The learning process is modeled as a *learning loop* (or spiral) that starts with some initial behavior *A* and leads to a modified behavior *B* (see Figure 1). The

[5]Explorative experiments in areas where little or no theory is available are difficult to publish since referees feel uncomfortable when they have no sound theory to compare the results with. Thus, learning under less than ideal conditions is rarely studied per se in experiments. On the other hand, however, there are some experiments that provide "bad" learning conditions, as will be discussed below in Section 4 for the case of decision making under risk.

individual gains experience by applying behavioral strategies out of an existing repertoire to some situation, and receiving feedback about the consequences. This process continues and behavior is modified as long as some aspiration level is not satisfied or some maximal utility is not reached. The repertoire of behavioral strategies is modified and enlarged as experience accumulates.

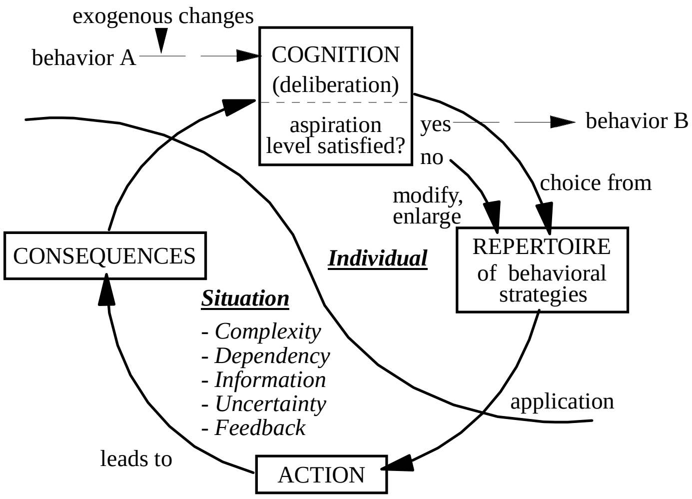

Figure 1: A Learning Loop.

The *situation* is characterized by five main learning determinants, namely by (i) the complexity of the environment and of the task, (ii) the degree of dependency among actors induced by the number of involved actors, (iii) the available information about structure of the situation, (iv) structural and strategic uncertainty, and (v) quality, quantity and content of feedback. Table 1 gives an overview of these determinants with regard to their occurrence in three types of situations, i.e. individual decision tasks, markets, and games.

type of determinant	description of determinant		type of situation		
			individual decision making	markets	games
structural determinants (prior to action)	complexity of environment and task		x	x	x
	information about structure		x	x	x
	degree of dependency			x	x
	uncertainty	structural uncertainty	x	x	x
interaction determinants		strategic uncertainty		(x)[✧]	x
	feedback information		x	x	x

[✧] In perfectly competitive markets strategic uncertainty does not exist but emerges as behavior becomes interdependent with a decreasing number of actors.

Table 1: Overview of Learning Determinants

Similar to directional learning, contingent learning makes qualitative predictions about the influence of these learning conditions with regard to the effectiveness of learning processes (see SLEMBECK, 1998a for details). The approach has been tested in a number of experiments (SLEMBECK, 1998b, 1999a, 1999b; SLEMBECK & BINMORE, 1999). The results stress the importance of structural information, structural uncertainty and feedback for adaptive behavior. In sum, behavior may differ substantially from theory predictions when the structure of the situation is not completely known or uncertain, and when feedback is deficient. Strategic uncertainty, however, seems to play a decreasing role as subjects gain experience, which is true to a lesser degree for structural uncertainty. Furthermore, qualitatively and quantitatively improved feedback can yield behavior that is closer to theory prediction in experiments that are usually run under "low" feedback conditions (see Section 4 for details).

3. Aspects of Learning in "Practice"

In real-life there exist some simple facts that make the importance of learning obvious. Firstly, people are born as babies, thus they have to learn to behave "optimally" *and* what they like. There is evidence that even adults don't have given (and fixed) preferences for each option they encounter in life. Thus, they have to learn what they like, and construct their preferences in a particular situation (see e.g., LOOMES, 1999, F37).

Time and Opportunities

Aside from the question of preferences, people live in a dynamic or evolving world where relative prices, opportunities for production and consumption, available information, and various economic restrictions (such as budgets) change all the time in the economic realm. However, in all practical situations there is *limited time and opportunity* for an individual to learn. Much of what people know, believe to know or have learnt is not acquired by personal experience but transmitted socially. Thus, most real-life learning involves social interaction that goes beyond simple imitation of successful behavior. Limited time and opportunity for individual learning means also that "*maximal learning*" –as often assumed in theory– is rarely "*optimal*" since time, cognitive effort etc. employed in learning processes are scarce resources that need to be economized on.

Availability of Feedback

One important aspect of limited opportunities in real-life learning is the availability of feedback. It has long been recognized in the psychology literature that effective learning takes place only under certain conditions in that it requires accurate and immediate feedback about the relation between the situational conditions and the appropriate response (see EINHORN & HOGARTH, 1978; TVERSKY & KAHNEMAN, 1987, 90).

In most practical situations, however, quality and quantity of feedback – e.g., to managers, entrepreneurs and politicians – is much lower than commonly assumed in economic theory because:

- outcomes are delayed and not easily attributable to particular actions
- variability in the environment may degrade the reliability of the feedback, especially where outcomes of low probability are involved, or/and when feedback is delayed
- there is often no information about what the outcome would have been if another decision had been taken
- most important decisions are unique and therefore provide little opportunity for learning.

Do Markets Heal?

The above list reflects a view that focuses on the individual, which is typical for psychologists. Economists, on the other hand, tend to focus on markets or *aggregate outcomes* instead of *individual behavior*. Economists have argued that evolutionary or market forces select "fit" actors. Hence, it is assumed that (competitive) markets provide an environment that involves incentives strong enough for actors to learn optimal behavior on average in the long run. This is the same type of (evolutionary) argument that has long been used to defend the rationality assumption in that markets are thought to "heal" the cognitive imperfections of actors through evolutionary forces, leading actors to behave *as if* they were rational.

Empirical evidence on this issue is mixed. There is evidence that "anomalous" behavior can survive for some (possibly long) time in real markets (see the large literature on "bubbles" and "herding" in markets, e.g. AVERY & ZEMSKY, 1998). – There is similar evidence from the laboratory in that some anomalies are overcome by learning in market environments while others are not: "*The data suggest the market glass is both half-full of deviations and half-empty because some deviations were drained away by learning*" (CAMERER, 1995, 675).

My own conclusion is that markets produce the types of learning assumed in the traditional and in the learning literature only under very ideal conditions (such as perfect competition and free information). These conditions are rarely met in real-life and even in experimental markets that are slightly realistic (see SUNDER, 1995, for an overview).

Furthermore, it may take a long time for markets to converge to some equilibrium even under ideal conditions – especially when the actors' initial beliefs are not coordinated. Most current theory lacks criteria for when the long run behavior it predicts can be assumed to occur. The practical question is how long the long run is, and whether the long run matters with respect to a given question or phenomenon.

Therefore, when applying learning approaches to real phenomena one may question the value of theories that make predictions for learning behavior or outcomes (equilibria) only in the long run. For applied work it would obviously be most useful to know something about short and medium run behavior, and about the approximate length of these runs.

Repeated vs. One-Shot Situations

An additional aspect of “learning in practice“ is that life is on one hand a repeated game, but on the other hand not all is repeated.

The *repeated play aspect* not only allows the individual to learn effectively in some situations (like daily shopping), to form rules of thumb, and to employ heuristics that have proven to be useful over time. It also introduces psychological or behavioral features that traditional economic theory does not account for – like concerns for fairness, reciprocity, manners, or other social norms that help real people to cooperate, coordinate or overcome social dilemmas in life. These additional features challenge theory, and they show up in experiments where they sometimes make the interpretation of observed behavior difficult.

The *one-shot aspect* has already been mentioned with respect to the problem of insufficient feedback in practical situations. Moreover, there is a general case for studying one-shot situations or games because they may sometimes be more realistic (see e.g., LOOMES, 1999). GEOREE & HOLT (1999b) argue that one-shot games are especially appealing because they allow us to abstract away from issues of learning and attempts to manipulate others‘ beliefs, behavior, or preferences (p. 29). Furthermore, psychological research suggests that *transfer* of learning across situations is (at least sometimes) surprisingly weak (LOEWENSTEIN, 1999, F29). Thus, people may find it difficult to transfer their experience from one situation to a similar or even formally identical situation so that potentially many situations may appear one-shot to them. On the other hand, people are confronted with new situations all the time, and two situations are never completely identical in life. Therefore, we may expect people to be prepared for dealing with novelty and to transfer their knowledge and beliefs between situations. Hence, for applications to novel or one-shot situations, like international politics, election campaigns or legal disputes, it seems important to have a theory about how people transfer their experience.

Imitation is one way to deal with novelty whenever the behavior of others can be observed. However, theories that assume simple copying of successful behavior may be overly optimistic for many real-life settings, since copying may be difficult or impossible in practice. For example, it seems quite demanding to imitate or copy successful companies, entrepreneurs or sportsmen like Michael Jordan. Therefore, imitation is regarded to be a challenging task by most social scientists.

Overall, there appear to be good reasons for studying both, repeated and one-shot situations or games. Economic situations may be one-shot in some (possibly important) cases, but people are also equipped with repeated play impulses from everyday life that have proven helpful in other cases (GALE ET AL.,1995). They carry these impulses with them and bring them also to the experimental laboratory (see next section).

Furthermore, real people not only adapt to given circumstances: they create or invent new situations, they change restrictions actively, they modify the rules of games, and they (re)design institutions. All these features of human behavior are obviously important to economics, but accounting for them is a very demanding task for theory. A more modest first goal would be to make clearer statements about what type of situation (e.g., repeated or one-shot) one has in mind when setting up a theory, and to explicitly acknowledge the theoretical and practical limits imposed by this choice. Hence, we need distinct criteria that tell us under which circumstances a given learning theory is thought to apply since many theories implicitly assume learning conditions that are not discussed in theory and rarely met in practice.

4. Aspects of Learning in Experiments

In experiments, the phenomenon of learning has several facets because learning is sometimes a nuisance factor, while it is the focus of research at other times (see e.g., FRIEDMAN & SUNDER, 1994).

Learning as a Nuisance Factor in Experiments

On one hand, experimenters aim to make experimental subjects familiar with a new environment and situation. They want to make sure that subjects understand the situation “correctly“ in the sense that subjects do consider all the given information. This information is typically defined by the model or theory under study so that not considering all information would bias the experiment. Practice rounds are introduced for this purpose. The fact that subjects have to make themselves familiar with the situation is a nuisance in this case, and behavior in practice rounds is typically not reported.

In the light of the arguments in favor of studying one-shot situations presented in the last section, however, not reporting initial behavior may be regarded as a waste of potentially useful information. It may in fact often be the case that experimenters have some specific sort of experimental outcome in mind before the experiment starts – an expectation that may be fueled by large series of pilot-experiments – so that experimenters have to wait until subjects finally “get it right“ before starting the “real“ experiment. This is the case when a static theory is being tested.

Learning as an Object of Experimental Study

On the other hand, there are experiments that explicitly study learning behavior. Much of the literature on experimentally based approaches mentioned in Section 2 falls into this class. In these experiments a one-shot game or situation is typically repeated several times. In individual decision making tasks repetitions are straightforward. In market or public good experiments, for instance, learning is studied by having the same subjects interacting repeatedly. In game experiments, however, there are two different procedures. While in most game experiments subjects are matched newly each round, in other experiments subjects play the same stage-game repeatedly with the same partner(s) or opponent(s). Random or rotation re-matching methods are employed to reduce potential supergame strategies or reputation effects in most game experiments, e.g. in bargaining. The reason is that theory often has not much to say about such effects as accounting for them would make the analysis considerably more difficult. Thus, experimental studies of learning and supergame behavior or the role of reputations are surprisingly rare, although both phenomena may be important in many real-life instances such as repeated bargaining – see e.g., SLEMBECK, 1999b, for the first ultimatum game experiment with fixed opponents where bargaining behavior is drastically different from what has been observed under random or rotation matching schemes.

An unsettled problem of experimental learning studies concerns the question whether short, medium, long, or ultra-long run behavior should be the focus of attention. ROTH & EREV (1995) have forcefully argued in favor of the importance of medium run or “intermediate term“ behavior, but to date not everyone seems to be convinced of the argument. After all, this is a question of the phenomenon under study or the application one has in mind. However, the literature would benefit from making this point more explicit, since in the long run that theoretical models often need for behavior to converge (asymptotically) to some equilibrium, experimental subjects – and real-life actors – may long be dead.

“Stationary Replications“

Repeating an identical task or game several times in the laboratory has been criticized by LOEWENSTEIN (1999) who labels this practice “stationary replication“. He finds that *„stationary replication is simply not a common feature of economic life..*“ and *„...it is by no means clear that behaviour at the end of repetitions is more representative of actual economic behaviour than behaviour at the beginning*“ (p. F28). A somewhat contrary view is taken by HERTWIG & ORTMANN (1998) who argue that stationary replication is a strength of experimental practices in economics as compared to psychology.

Again, it is a question of whether we are interested in initial behavior or in the evolution of behavior over time. Whenever we observe behavior to converge or stabilize over time this does not necessarily mean that his behavior is more representative for actual economic behavior. Yet, it may sometimes still be useful to study medium or long run behavior in the lab, because this allows us to analyze and identify the conditions under which learning leads to behavioral convergence or stabilization. Thus, I would maintain that stationary replication is a useful tool for analyzing learning in experiments as long as one keeps in mind the limitations mentioned by Loewenstein when interpreting experimental results and drawing conclusions for behavior outside the laboratory.

Binmore‘s Criteria

Several aspects of learning in experiments have been summarized recently by BINMORE (1999, F17) who stipulates three criteria or conditions under which one can expect economic theory to predict well:

– make the task „reasonably“ simple
– provide „adequate“ incentives
– allow „sufficient“ time for learning

All three rules have to do with learning in that a reasonably simple task is easier to learn, adequate incentives are needed to motivate sufficient learning efforts, and sufficient time for learning must be allowed. BINMORE (1999, F18) admits that *„by interpreting the criteria severely enough, we can almost guarantee that an optimising theory will work. By interpreting them loosely enough, we can almost guarantee that an optimising theory will fail*.“

The question, then, is where the boundary between these two extremes lies. In Binmore‘s view it is not interesting to ask whether economic theory works or not – he already knows that sometimes it does and sometimes it does not: “*The task is to classify economic environments into those where the theory works and those were it does not. In the shady area in between, the data obtained from the classification effort will then serve to refine the theories that do not quite make it*“ (BINMORE, 1999, F18). In this view, the role of learning conditions becomes apparent in that behavior can be manipulated by changing these conditions – at least in the laboratory. For applications, again, the crucial question arises when a given set of conditions is supposed to hold and whether conditions can be manipulated by real-life actors (such as politicians) or by “designers“ in applied market or mechanism design.

Anomalies and Learning in Individual Decision Making

Classification, however, is not always easy as experimental evidence on “Wason‘s 4-card-problem“ (see LOEWENSTEIN, 1999, F30)[6] and the “3-doors-anomaly“ (see FRIEDMAN, 1998;

[6] LOEWENSTEIN (1999, F30): “*There are four cards, each with a letter on one side and a number on the other. The exposed faces read ‘X‘, ‘Y‘, ‘1', and ‘2'. Subjects are asked which cards would need to be turned over to test the rule: ‘If there is an X on one side there is a 2 on the other.‘ Very few subjects give the right answer, which is X and 1. However, when the problem is put into a more familiar context (for example, there are 4 children from two different towns and two school districts and the rule is ‘If a child lives in Concord he goes to Concord High‘), a much higher fraction of subjects give the right answer. Subjects seem like zero intelligence agents when they are placed in the unfamiliar and abstract context of an experiment, even if they function quite adequately in familiar settings.*“

PAGE, 1998)[7] suggests. Both types of experiments involve an individual decision task that is reasonably simple to understand, and I assume that sufficient incentives were provided in the experiments. While subjects were able to solve the “4-card-problem“ correctly when the task was put in a real-life context instead of an abstract situation, the 3-doors-anomaly seems to persist even when subjects are helped in their learning. Maybe the reason is that subjects did not receive appropriate feedback in the experiments, or that there exists no real-life analogy to the latter problem so that real people don‘t have a chance to “get it right“.

As for the case of markets, some “anomalous“ behaviors in individual decision making tasks seem to vanish while others appear to persist when subjects learn. Therefore, it remains to be investigated more thoroughly under what conditions, learning leads to “more rational“ behavior, i.e. to behavior that is more in accordance with economic theory. Chances are that quality, quantity and content of feedback will prove to be among the strongest learning determinants. This may appear a too simple and too straightforward prediction. As will be discussed in the next paragraph, however, the role of feedback has been neglected in some important ways.

Risky Behavior and Learning with Feedback

While most experiments provide rather ideal conditions for learning, the issue of learning has rarely been addressed in one essential area of experimental research, namely in individual decision making under risk. Dozens of experiments have been run to test expected utility theory (EUT) and various variants of preference functionals under risk (see e.g. HEY & ORME, 1994). From this literature STARMER (1999, F8) has recently concluded: “*One thing we have learned for sure...is that EUT is descriptively false. Mountains of experimental evidence reveal systematic violations of the axioms of EUT, and the more we look, the more we find.*“ Similarly, LOOMES (1999, F37) finds that “*the bulk of violations of EUT are broadly robust, in that they can be reproduced with a variety of displays and using familiar sums of money and straightforward, easily imaginable probabilities. Nor can these effects be explained away simply in terms of lack of incentives.*“

[7]The situation for the three-doors-anomaly is taken from Monty Hall‘s game show “Let‘s make a deal“ where the final guest is presented with three doors to choose from. One door hides the “grand prize“ while the others contain worthless prizes. After the guest has chosen a door, the show master opens one of the other two doors, revealing a worthless prize, and asks the guest whether he would like to stay with his original choice or switch to the other unopened door. While most candidates (and experimental subjects) are reported to stay with their original choices, switching would improve their chances of winning the grand prize from 1/3 to 2/3. FRIEDMAN (1998) has investigated learning by providing experimental subjects with (i) intense monetary incentives, (ii) a written track record, (iii) written advice, and (iv) comparative results of other subjects results when the task was repeated 10 to 15 times; but only about a third of his subjects learned to switch. My feeling is that many more rounds and stronger feedback would be needed for subjects to learn that switching is advantageous. PAGE (1998) tried another version were more doors were involved, but only when the choice set was reduced from 100 doors before the subject‘s choice to 2 doors after picking a door a substantial number of subjects (about 80%) learned to switch.

How about learning in this context? How about looking not only at incentives but at learning conditions such as sufficient time and feedback? In a new experiment by SLEMBECK & BINMORE (1999) evidence collected in three countries suggests that allowing subjects to learn the consequences of risky choices by giving them appropriate feedback and allowing them sufficient time to revise their decisions yields behavior that is much closer to EUT than observed before. (In the final round of the experiment 85 to 90 percent of the choices involved the expected value maximizing option). Hence, there appear to exist conditions under which EUT can be confirmed.[8]

The question remains how relevant such a result is. On one hand it sheds additional light on the importance of learning conditions as stressed by the contingent learning approach. On the other hand it demonstrates that it is indeed possible to construct conditions in which theory works well with real people. Finding such conditions is a breakthrough in that it provides an empirical benchmark to compare theory and practice with. Hence, we at least know one set of conditions where theory appears to work well.

An additional question is, of course, the *practical relevance* of such a finding. As discussed above, many important practical decisions may lack the kind of learning opportunities provided in our experiment. It is therefore a matter of application to determine if the involved learning conditions can be assumed to hold. That is, for each application one has to think carefully about the learning conditions before applying theory. It seems nevertheless very useful to know that such conditions do exist at least in the laboratory.

Experimental Tests of Learning Theories

When looking at experimental evidence from tests of learning theories, one is currently left with a somewhat empty feeling. As mentioned, evidence from comparisons between learning models in experiments is inconclusive. Currently, various authors are collecting evidence in favor of one or the other approach (see EREV & ROTH, 1998; EREV & RAPOPORT, 1998; NAGEL & TANG, 1997; VAN HUYCK ET AL., 1996; RAPOPORT ET AL., 1995).

As the number of models and experiments to test them increases, learning mechanisms can be expected to improve their ability to capture observed learning behavior, especially when the performance of models is assessed not only through hypothesis testing but through competitive tests among alternative models with respect to the same set of data. At the same time criteria for

---

[8] In a different design BARRON & EREV (1999) have recently investigated the effect of feedback on repeated decision making under uncertainty and risk, and compared the results with predictions of *prospect theory* (KAHNEMAN & TVERSKY, 1979) and *reinforcement learning theory*. They argue that many important economic decisions are made in a dynamic environment with feedback, and find that prospect theory predicts initial behavior well, while long run behavior is more in accordance with reinforcement learning. This finding seems to be driven by the effect of feedback.

model comparison have to be developed that are widely missing today and that involve a number of difficult problems (see BUSH, 1963, for a discussion in psychology).[9] Nevertheless, the issue of model comparison seems crucial for the progress of this branch of research.

However, even if the problems of model comparison can be solved successfully, the question remains what results can be expected from this type of research. I suspect that one will find a general model of learning that applies to a variety of situations very well. Most likely there will be a model that "fits best on average" over a certain range of games or situations, but there will be many more models that fit better to specific situations or games since they are adjusted, modified or calibrated. And even if a single relatively general mechanism can be found, what can we conclude other then that some fundamental principles of learning which are well-documented in the psychological literature indeed appear to hold in interactive settings like markets or games?[10] These findings may be interesting *per se* – though more to psychologists than to economists.[11] What economics needs today is a theory that allows us to combine findings from learning theory and empirical evidence on learning and connect them to situations that are economically relevant outside the laboratory.

5. Concluding Remarks

The discussion of the experimental learning literature presented in the last section may be overly pessimistic. Indeed, if there is a basic message from this literature it is that there is a tendency for people to adapt to the (economic) incentive structure in the long run in most settings. This is good news since it is what economists have assumed, or at least hoped for, all the long. There is a reservation, however, in that the message only holds provided that people have sufficient time, incentives, information and feedback for learning. Hence, as proposed by the contingent learning approach, learning much depends on the specific learning conditions, and further progress is likely to be made by studying these conditions more deeply.

One of the rather robust experimental finding is that *learning is relatively slow* in many cases. Thus, one may not conclude that “people do not learn“ when they are given only a few trials,

---

9Criteria for model comparison must be stipulated and tradeoffs among them must be discussed. These criteria include, for example, the goodness-of-fit, the number of free parameters (parsimony), the economic interpretability of these parameters, and the validation of parameters in experiments.

10These principles are likely to include the “Law of Effect” (Thorndike, 1898), the “Power Law of Practice” (Blackburn, 1936), reference points effects, and the effects of negative versus positive reinforcements (see DANIEL, SEALE & RAPOPORT, 1997, 18).

11The goal of at least some of these learning models is to account for individual, not aggregate behavior. To economists individual behavior has traditionally been of interest only in so far as something can be learned about aggregate outcomes. That is, if individual behavior differs substantially among subjects (and from theory), but aggregate behavior or outcomes are in harmony with the theory (as the results of RAPOPORT ET AL. (1995) suggest), most economists may be reluctant to these findings since economic theory is one of aggregate outcomes, based on average individual behavior.

possibly with low feedback, and behavioral anomalies may be an artifact of the experiment that is diminished as subjects learn. Of course, it is a question of whether we are interested in initial or long run behavior, but "anomalies" should not come as a surprise when people are put into a completely new situation such as an experiment. Thus, there is a case for "practice rounds" in many experiments. However, the number of these rounds as well as the behavior observed in these rounds should be reported and analyzed more carefully, even if behavior appears to be "noisy" (see SLEMBECK & BRENNER, 1999, for a theory of and experimental evidence on noisy decision making).

Furthermore, learning appears to be rather insensitive to absolute payoff magnitudes. Although there exist no experimental studies with extremely high payoffs, data from experiments where payoffs were linearly transformed (within the range of positive values) and data from the field (e.g., financial markets) suggest that absolute payoff magnitudes are less important, while learning behavior is sensitive to relative payoffs and to payoff variance (in that a higher variance slows learning).

The basic finding that learning is often slow and sensitive to learning conditions, such as feedback, has several implications for economics. Generally, economic actors cannot be assumed to adapt instantaneously to new circumstances as traditional theory suggests. Real people need time and feedback for learning, which is important not only in traditional economic settings where relative prices, budgets, technology etc. may change, but especially when new markets or other interaction mechanisms are implemented. Learning to behave "optimally" may take more time than assumed by designers, and it seems crucial to implement straightforward feedback mechanisms. The latter is even more important when entire economies are in transition, as e.g. in Eastern Europe, so that feedback is delayed and confounded for some possibly longer period.[12] In such cases adaption is also slowed by lack of information about the rules and structure of the emerging new economy, by uncertainty about possible further changes in the future, and by (strategic) uncertainty about how others will behave in the new environment. This stresses the general question of how *institutions* matter for learning and economic evolution, which I have not discussed in this paper (see e.g. YOUNG, 1998, for a basic approach).

Some further implications concern the *effects of economic policy* in that policies may take time to show effects not only because of delays in the political and economic system, but also due to slow individual learning, especially when learning conditions are unfavorable. That is, whenever a new policy is designed and people are expected to adjust or adapt to it, learning theory suggest, i.e., that the policy involves direct feedback mechanisms and predictable rules since actors will already be busy dealing with strategic uncertainty. A policy should also not decrease transparency

---

[12]For instance, learning to be an entrepreneur may take time after entrepreneurship has been banned in some countries for many decades.

or increase uncertainty because this would hinder learning. These are insights that some policy advisers have shared for a long time. However, learning theories provide a new microeconomic foundation for applied economic policy on a behavioral level.

References:

- **Avery, C. & Zemsky, P.** (1998): *Multidimensional Uncertainty and Herd Behavior in Financial Markets*, American Economic Review, 88(4), 724–748.
- **Barron, G. & Erev, I.** (1999): *The Effect of Feedback on Decision Making Under Uncertainty and Risk - Loss Aversion Without Reflection and the Reversed Certainty (Allais) Effect*, mimeo, Technion / Haifa.
- **Binmore, K.G.** (1999): *Why Experiments in Economics?*, Economic Journal, Vol.109 No.453, F16-F24.
- **Binmore, K.G.** (1992): *Fun and Games : A Text on Game Theory*, Heath Lexington, MA
- **Blume, L.E. & Easley, D.** (1995): *What Has the Rational Learning Literature Taught Us?*, in: Kirman, A.P. & Salmon, M. (Eds.): Learning and Rationality in Economics, Oxford / Cambridge, 12–39.
- **Börgers, T.** (1999): *When Does Learning Lead to Nash Equilibrium?*, mimeo, University College London.
- **Bosch-Domènech, A. & Vriend, N.J.** (1999): *Imitation of Successful Behavior in Cournot Markets*, mimeo, Pompeu Fabra and Queen Mary & Westfield College.
- **Bush, R.R.** (1963): *Estimation and Evaluation*, in: Luce, R.D., Bush, R.R. & Galanter, E. (Eds.) (1963): Handbook of Mathematical Psychology, New York etc., Wiley & Sons., 429–469.
- **Camerer, C**. (1995): *Individual Decision Making*, in: Kagel, J.H. & Roth, A.E. (Eds.): Handbook of Experimental Economics, Princeton, Princeton University Press, 587-703.
- **Crawford, V.P.** (1995): *Theory and Experiment in the Analysis of Strategic Behavior*, Discussion Paper 95-37, University of California, San Diego, published in: Kreps, D. & Wallis, K. (Eds.): Advances in Economics and Econometrics: Theory and Applications, 7th World Congress Econometric Society Monographs No. 27, Cambridge & New York, Cambridge University Press (1997).
- **Daniel, T.E., Seale, D.A. & Rapoport, A.** (1996): *Strategic Play and Adaptive Learning in the Sealed Bid Bargaining Mechanism*, University of Alberta, mimeo, April 19, 1996.
- **Duffy, J. & Nagel, R.** (1997): *On the Robustness of Behavior in Experimental Beauty Contest Games*, Economic Journal, 107, 1684–1700.
- **Einhorn, H.J. & Hogarth, R.M.** (1978): *Confidence in judgement: Persistence in the illusion of validity*, Psych. Review, Vol. 85, No. 5, 395–416.
- **Erev, I. & Rapoport, A.** (1998): *Coordination, "Magic", and Reinforcement Learning in a Market Entry Game*, Games and Economic Behavior, Vol. 23 No. 2, May 1998, 146–175.
- **Erev, I. & Roth, A.** (1998): *Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria*, American Economic Review, Vol. 88 No.4, 848-881.
- **Friedman, D.** (1998): *Monty Hall's Three Doors: Construction and Deconstruction of a Choice Anomaly*, American Economic Review, Vol. 88 No. 4, 933-946.
- **Friedman, D. & Sunder, S.** (1994): *Experimental Methods - A Primer for Economists*, Cambridge, Cambridge University Press.
- **Fudenberg, D. & Levine, D.K.** (1998): *The Theory of Learning in Games*, MIT Press, Cambridge MA and London.
- **Fudenberg, D. & Tirole, J.** (1991): *Game Theory*, Cambridge / London, MIT Press.
- **Gale, J., Binmore, K.G. and Samuelson, L.** (1995): *Learning To Be Imperfect: The Ultimatum Game*, Games and Economic Behaviour, 8, 56–90.
- **Georee, J.K. & Holt, C.** (1999a): *Stochastic Game Theory: For Playing Games, Not Just for Doing Theory*, mimeo, Department of Economics, University of Virginia.
- **Georee, J.K. & Holt, C.** (1999b): *Ten Little Treasures of Game Theory and Ten Intuitive Contradictions*, mimeo, Department of Economics, University of Virginia, February 1999.
- **Hertwig, R. & Ortmann, A.** (1998): *Experimental Practices in Economics - A Methodological Challenge for Psychologists*, mimeo, Max-Plank-Institute for Human Development, Berlin.
- **Hey, J. & Orme, C.** (1994): *Investigating Generalizations of Expected Utility Theory Using Experimental Data*, Econometrica 62(6), 1292–1326.
- **Ho, T.-H., Camerer, C. & Weigelt, K.**(1998): *Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests'*, American Economic Review, 88(4), 947–969.
- **Hogarth, R.M. & Reder, M.W.** (1987): *Perspectives from Economics and Psychology*, in: Hogarth, R.M., Reder, M.W. (Eds.): Rational Choice – The Contrast between Economics and Psychology, Chicago & London, University of Chicago Press, 1–23.
- **Holland, J.H.** (1975): *Adaption in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor.

- **Huck, S., Normann, H.T. & Oechssler, J.** (1999): *Learning in Cournot Oligopoly - An Experiment*, Economic Journal, 109 (February), C80–C95.
- **Loewenstein, G.** (1999): *Experimental Economics from the Vantage-Point of Behavioural Economics*, Economic Journal, 109 (February), F25-F34.
- **Loomes, G.** (1999): *Some Lessons from Past Experiments and Some Challenges for the Future*, Economic Journal, Vol.109 No.453, F35-F45.
- **Mailath, G.** (1998): *Do People Play Nash Equilibrium? Lessons from Evolutionary Game Theory*, Journal of Economic Literature, 36, 1347–1374.
- **Nagel, R.** (1999): *A Survey on Experimental Beauty Contest Games - Bounded Rationality and Learning*, in: Bodescu, D.V., Erev, I. & Zwick, R. (Eds.): Games and Human Behavior, Mahwah (NJ), Lawrence Erlbaum Associates Inc., 105–142.
- **Nagel, R. & Tang, F.F** (1997): *An Experimental Study on the Centipede Game in Normal Form - An Investigation on Learning*, Hebrew University, Working Paper no.149, to appear in a special issue in the Journal of Mathematical Psychology, ed. by G. Bolton.
- **Ochs, J.** (1999): *Coordination in Market Entry Games*, in: Bodescu, D.V., Erev, I. & Zwick, R. (Eds.): Games and Human Behavior, Mahwah (NJ), Lawrence Erlbaum Associates Inc., 143–172.
- **Page, S.E.** (1998): *Let‘s Make a Deal, Economic Letters*, 61, 175-180.
- **Payne, J.W.** (1982): *Contingent Decision Behavior*, Psychological Bulletin, Vol.92, No.2, 382–402.
- **Rapoport, A., Seale, D.A., Erev, I., Sundali, J.A.** (1995): *Equilibrium Play in Large Group Market Entry Games*, University of Arizona, mimeo, December 31, 1995.
- **Roth, A.E.** (1995): *Introduction*, in: Kagel, J.H. & Roth, A.E. (Eds.): Handbook of Experimental Economics, Princeton, 3–109.
- **Roth, A.E. & Erev, I.** (1995): *Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term*, Games and Economic Behavior, 8, 164–212.
- **Rumelhard, D. & McClelland, J.** (1986)*: Parallel Distributed Processing - Explorations in the Microstructure of Cognition*, Cambridge, MIT Press.
- **Samuelson, L.** (1997): *Evolutionary Games and Equilibrium Selection*, Cambridge MA and London, MIT Press.
- **Schlag, K.** (1998): *Why Imitate, and If So How? - A Bounded Rational Approach to Multi-Armed Bandits*, Journal of Economic Theory, 78(1), 130–156.
- **Schlag, K.** (1999): *Which one should I imitate?*, Journal of Mathematical Economics 31(4), 493–522.
- **Schumpeter, J.A.** (1934): *The Theory of Economic Development*, Cambridge, Harvard University Press.
- **Selten, R. & Stoecker, R.** (1986): *End Behavior in Sequences of Finite Prisoner‘s Dilemma Supergames - A Learning Theory Approach*, Journal of Economic Behavior, 47–70.
- **Selten, R. & Buchta, J.** (1999): *Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions*, in: Bodescu, D.V., Erev, I. & Zwick, R. (Eds.): Games and Human Behavior, Mahwah (NJ), Lawrence Erlbaum Associates Inc., 79–102.
- **Slembeck, T.** (1998a): *A Behavioral Approach to Learning in Economics – Towards an Economic Theory of Contingent Learning*, paper presented at the European Economic Association Annual Meeting, Berlin, September 1998, previous version: Working Paper No. 316, Department of Economics, University of Pittsburgh, August 1997.
- **Slembeck, T.** (1998b): *As If Playing Fair – Experimental Evidence on the Role of Information in Ultimatum Bargaining*, mimeo, University College London.
- **Slembeck, T.** (1999a): *Low Information Games – Experimental Evidence on Learning in Ultimatum Bargaining*, Discussion Paper No. 9903, Department of Economics, University of St.Gallen.
- **Slembeck, T.** (1999b): *Reputations and Fairness in Bargaining – Experimental Evidence from a Repeated Ultimatum Game with Fixed Opponents*, Discussion Paper No. 9904, Deptment of Economics, University of St.Gallen.
- **Slembeck, T. & Binmore, K.G.** (1999): *Learning Preferences Over Risky Prospects – Experimental Evidence from Three Countries*, University of St.Gallen and University College London, first draft August 1999.
- **Slembeck, T. & Brenner, T.** (1999): *Noisy Decision Makers - On Errors, Noise and Inconsistencies in Economic Behaviour*, mimeo, University of St.Gallen and Max-Planck-Institute Jena.
- **Stahl, D.O.** (1996): *Rule Learning in a Guessing Game*, Games and Economic Behavior, 16, 303–330.
- **Starmer, C.** (1999): *Experimental Economics: Hard Science or Wasteful Tinkering?*, Economic Journal, Vol.109 No.453, F5–F15.
- **Sunder, S.** (1995): *Experimental Asset Markets: A Survey*, in: Kagel, J.H. & Roth, A.E. (Eds.): Handbook of Experimental Economics, Princeton, Princeton University Press, 445-500.
- **Tversky, A. & Kahneman, D.** (1987): *Rational Choice and the Framing of Decisions*, in: Hogarth, R.M. & Reder, M.W. (Eds.): Rational Choice – The Contrast between Economics and Psychology, 67–94, Chicago/London, University of Chicago Press.
- **Van Huyck, J., Battalio, R. & Rankin, F.** (1996): *On the Evolution of Convention: Evidence from Coordination Games*, Texas A&M, mimeo.

- **Weibull, J.W.** (1995): *Evolutionary Game Theory*, Cambridge MA and London, MIT Press.
- **Young, P.H.** (1998): *Individual Strategy and Social Structure***-** *An Evolutionary Theory of Institutions*, Princeton, Princeton University Press.