

A BEHAVIORAL APPROACH TO LEARNING IN ECONOMICS

TOWARD AN ECONOMIC THEORY OF CONTINGENT LEARNING ¹

TILMAN SLEMBECK

University of St.Gallen
Department of Economics ²

first version August 1997 ³
current version prepared for the

European Economic Association Congress

Berlin, September 1998

Abstract. In economics, adjustment of behavior has traditionally been treated as a “black box.” Recent approaches that focus on learning behavior try to model, test, and simulate specific adjustment mechanisms in specific environments (mostly in games). Results often critically depend on distinctive assumptions, and are not easy to generalize. This paper proposes a different approach that aims to allow for more general conclusions in a methodologically more compatible way. It is argued that the introduction of the main determinants of learning behavior as situational restrictions into the standard economic model may be a fruitful way to capture some important aspects of human behavior that have often been omitted in economic theory. Based on a simple model of learning behavior (learning loop), robust findings from psychology are used to explain behavior adjustment, and to identify its determinants (contingent learning). An integrative methodology is proposed where the “black box” is not opened, but instead the factors that determine what happens inside, and the limits imposed by these factors can be analyzed and used for model building. The paper concludes with testable hypotheses about learning behavior in the context of economics.

¹For many helpful comments on earlier versions of this paper I would like to thank Ken Binmore, Andreas Brandenberg, Amnon Rapoport, Al Roth, Jürgen Weiss, and the participants of the HSG Lunch Discussion (University of St.Gallen, Switzerland), the 29th Annual Meeting of the Society for Mathematical Psychology (University of North Carolina, Chapel Hill, NC), the FAS Behavioral Economics Workshop (Harvard University, Cambridge, MA), and the Economic Theory Seminar (University College London).— This research has been financially supported by the Swiss National Science Foundation (No. 8210-46574).

²Author’s address: Tilman Slembeck, University of St.Gallen, Varnbuelstrasse 19, 9000 St.Gallen, Switzerland. E-mail: tilman.slembeck@unisg.ch

³Working Paper No. 316, Department of Economics, University of Pittsburgh.

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)

1. Introduction

In recent years, economists have shown an increasing interest in learning processes. It seems to be widely acknowledged that learning is a basic feature of human behavior and therefore must be important in all economic analysis that aims to advance our understanding of that behavior. This is especially the case for a rational choice approach that tries to comprehend human behavior not only in a narrowly defined economic context, thereby contributing to social sciences more generally.

Despite the basic and longstanding acceptance of its importance (see e.g., Schumpeter above), the phenomenon of human learning has mostly been neglected in economic theory. As will be outlined in Section 2, learning has traditionally not been recognized as a truly behavioral feature but was treated as a black box of automated adjustment, and it is only since the early eighties that learning has been analyzed more extensively by economists. The shift from models that (implicitly) assume perfect adjustment of behavior to approaches that explicitly model adjustment processes originates on one hand in the *rational expectations literature* where theorists felt the need for a better understanding of how expectations are formed. The occurrence of multiple equilibria in many of these models led to theoretical efforts to model rational expectations equilibria as the result of some adjustment process. On the other hand, there was also an increasing demand from *game theorists* for a more comprehensive theory that allows to distinguish between multiple equilibria in games.

After more than one decade of research, the question arises what can be learned from this learning literature. It will be argued that most current approaches that stipulate learning mechanisms contribute to refinements of theoretical answers to theoretical questions in an often axiomatic or ad hoc way, but avoid to address truly behavioral questions. That is, they fail to enhance our understanding of learning as a basic process in economic behavior, a process that helps the individual to adjust to changing circumstances under a variety of conditions.

Section 3 discusses the methodology involved with current approaches and presents the methodology of an alternative approach: contrary to standard theory that assumes learning to be perfect – and current approaches that study specific learning mechanisms either theoretically or experimentally – an alternative methodology is to identify the *determinants* for learning

processes.⁴ Section 4 presents an approach of contingent learning (CLA) based on this idea and concludes with behavioral⁵ hypotheses about the influence of situational factors on learning.

The goal of this approach is on both levels of analysis. On the *positive* level it provides a framework that permits to account for the contingencies of a given situation with respect to learning in a way that is methodologically compatible with traditional economics. Thereby, the CLA aims to enlarge the economists' analytical instrumentarium with theory elements that allow for deeper explanations and better predictions in cases where behavioral adjustments appear to be important, but where the conditions for these adjustments cannot be assumed to be as perfect as presumed in standard theory (see next section for a discussion of these assumptions). Since these cases are likely to include processes of economic transition, privatization, the introduction (or design) of markets, as well as changes in economic policy, the CLA may also have implications on a *normative* level.

2. Learning in economics – a critical review

Virtually no economist would maintain that learning processes are unimportant for economics – especially if economics is understood as a genuinely social science that studies human behavior more generally. Indeed, it is easy to see that some sort of learning lies at the heart of the economic approach. As discussed more extensively next, traditional economics, in its most general methodological characteristic, is the analysis of changes in behavior as a response to changing restrictions or relative prices. This basic methodological concept – in conjunction with the rationality assumption – provides for the predictive power and the relative success of the rational choice approach as compared to other disciplines or approaches in social sciences. Hence, the need for some notion of the adjustment of behavior inevitably follows from this basic concept.

2.1. Standard Approach

Traditionally, economics has assumed a perfect adjustment mechanism that allows the analysis to concentrate on the steady states of some unspecified underlying process:

“Economics has tended to focus on situations in which the agent can be expected to ‘know’ or to have learned the consequences of different actions so that his observed choices reveal stable features of his underlying preferences. We use economic theory to

⁴Note that the notion of “determinants” does not mean that learning is viewed as a deterministic process. Instead, the idea is that learning is influenced in a predictable way by situational factors, where the term “situational” refers to factors that are part of the situation rather than part of the psychological makeup of the individual.

⁵Note that the term “behavioral” is used here to denote a theory (or a hypothesis) that is derived from evidence about actual human behavior. This evidence is based on both behavioristic *and* cognitive theories and empirics in psychology and economics. Hence, *behavioral* approaches in this sense may profoundly differ from *behavioristic* approaches in the tradition of Pavlov, Thorndike and Skinner.

calculate how certain variations in the situation are predicted to affect behavior, but these calculations obviously do not reflect or usefully model the adaptive process by which subjects have themselves arrived at the decision rules they use. Technically, I think of economics as studying decision rules that are steady states of some adaptive process, decision rules that are found to work over a range of situations and hence are no longer revised appreciably as more experience accumulates.” [LUCAS (1987, 218)].

As SIMON (1978, 10) has pointed out, the interest was – and mainly still is – not in the adjustment process itself, but in its outcome, i.e., in *which* decisions are made by individuals and not *how* decisions are made.⁶ Hence, in most standard equilibrium models learning is reduced to some sort of *passive* “adjustment” or “adaptation” in that individuals are thought to automatically react to changing circumstances (i.e., to changes in restrictions or relative prices) in otherwise stable environments.⁷

This reduction (in the spirit of Lucas’ above statement) implies strong assumptions about learning – in that a perfect and complete learning process is implicitly presumed – that could only be defended in an environment that provides perfect conditions for learning. Despite the fact that this kind of reasoning not only prevails in microeconomic theory, but extends to macroeconomics (e.g., to models that involve the rational expectations hypothesis which implicitly assumes a perfect and complete learning process) in the standard literature virtually no attempt is made to justify these assumptions that are so crucial for the functioning of economic theory.

Rather, the strong learning assumptions are commonly defended in the same vein as other strong assumptions in economic theory (e.g., the rationality assumption) – that is, by the traditional “as if” argument à la Friedman, or by similar arguments like “the market leads to perfect learning”.⁸ Nevertheless, the critique from various sides continues. Aside from other

⁶Note that Lucas had admitted earlier that hypothesizing about the quantitative implications of economic policies “involves imagining how agents will behave in a situation which has never been observed”, and therefore “one must have some understanding of the way agents’ decisions have been made in the past and some method of determining how these decisions would be altered by the hypothetical change in policy.” [LUCAS, 1981, 180] This statement clearly calls for the study of learning processes.

⁷WITT (1992, p. 4): “In the [neoclassical paradigm] change is always interpreted as reactive. Individual agents or economic aggregates are viewed as responding to events that affect the basis of decision making. Economic actors are portrayed as attempting to adapt optimally to new conditions imposed on them. They are not credited in any way with creating the new conditions themselves. The reason for this narrow interpretation is the very core of the neoclassical paradigm, the synthesis of optimization and the equilibrium concept. Used together, the two ingredients rule out any explanation of individual behavior other than that of adapting to changing circumstances.”

⁸A long list of such defensive arguments with respect to the rationality assumption can be found in CONLISK (1996, pp. 683) who presents learning as an argument in favor of rationality but criticizes it since it holds only under very perfect conditions. Hence, it is the weakness of this argument and theories that make strong assumptions about learning that these conditions are not identified and the limits imposed by them are not analyzed. See also THALER (continued...)

forceful counter arguments – concerning the impossibility to justify the rationality assumption as a result of learning,⁹ and the inability of neoclassical theory to genuinely account for evolution or development¹⁰ – an important point of criticism that has been put forward (e.g., by WINTER, 1987) is that a theory that gives no rationale for its underlying strong assumptions provides no internal criteria for its applicability.

That is, neoclassical theory assumes perfect and complete learning but “*itself provides no indication as to how long it takes for adaptive processes to reach something like steady-state conditions, it provides no guidance regarding the quality of the predictions that...may be expected to provide in particular cases.*”¹¹ Furthermore, although *the limits* of the underlying assumptions may generally be acknowledged, those who defend these assumptions along the lines mentioned above “*do little to help define [these limits] and nothing to explore the important phenomena that lie beyond them.*”¹²

Generally speaking, standard economics has treated human learning as a black box process of perfect adaptation and has not attempted to explore either the conditions under which this may be justified, or the limits that are implied by the learning assumptions.

2.2. Approaches to Economic Learning

Instead of assuming a black box process, some more recent approaches have tried to explicitly model behavior adjustment by introducing statistical techniques and other mechanisms of information updating or gathering, referred to as “learning”. Three main fields of research can be distinguished. The first addresses learning in (rational) equilibrium models, and the second focuses on learning in a game theoretical context.¹³ While these two fields approach learning mainly on theoretical grounds, a third field studies actual learning behavior in experiments.

In experimental approaches to learning, researchers aim to find algorithms that mimic actual learning behavior (as observable in experiments) in order to predict that behavior in certain well-defined situations (e.g., in classes of games, see EREV & ROTH, forthcoming).

⁸(...continued)
(1991, p. 158).

⁹See ARROW (1987, 201).

¹⁰See NORTH (1994).

¹¹WINTER (1987, 248).

¹²WINTER op. cit.

¹³Another field may be distinguished that focuses on “social learning” and analyzes “*the process by which certain mechanisms in society aggregate the information of individuals*” (VIVES, 1996, 589). It draws from both, rational expectations and game theoretic models but has developed genuine concepts (e.g., “herding”, “informational cascades”).

In both theoretical fields researchers face the same problem that many models tend to have multiple equilibria.¹⁴ Thus, some “learning” mechanism is introduced in order to model adjustment paths that allow to determine or to distinguish between equilibria. The main interest lies in the analysis of the properties under which behavior converges to some equilibrium, and the stability of the equilibrium (if any is reached). The properties analyzed typically include the rationality assumption, informational assumptions, assumptions about (prior) beliefs, and, of course, the adjustment process itself. Hence, an example of a typical question is: “Does x -learning of w -rational players with y -beliefs (or y -expectations) lead to z -equilibrium?”¹⁵ With focus on the assumed “learning” process (i.e., on x) various mechanisms are discussed in the literature of which the most common ones will shortly be presented in order to characterize the underlying concept of learning.¹⁶

2.2.1. Learning and Rational Expectations

In the rational expectations literature (where z = rational expectations), the most common adjustment mechanism is Bayesian (or rational¹⁷) learning. Hence, this and other statistical methods¹⁸ are viewed as representing the human learning processes in the formation of (rational) expectations. This way of modeling may be justified by an “as if” assumption in conjunction with empirical success that would override the involved superrationality (which is assumed in that huge cognitive abilities are needed to perform the calculations). Unfortunately, econometric evidence seems not to be overwhelming, and the introduction of learning mechanisms does not appear to reduce the set of potential outcomes in a meaningful way.¹⁹ The difficulties lie both in the definition of optimal learning, and in limiting it to one mechanism (out of a large set of plausible mechanisms) that can be justified by some optimality argument. Also, common statistical techniques cannot be applied in a strict sense to the problem of individual inference in rational expectations (macroeconomic) models since to do so would require individuals to be unaware of the effects of beliefs on outcomes, whereas in these models beliefs affect outcomes and outcomes affect beliefs. Assuming that individuals are unaware of the effects of beliefs on

¹⁴Hence, the recent concern with so-called “learning” is not mainly a reaction to the mentioned criticism, but originates from the need for further theoretical development within the prevailing paradigm.

¹⁵Further questions may concern the speed of convergence, the modeling of short- or medium-term behavior (or play), and assumptions about knowledge or information.

¹⁶See KIRMAN & SALMON (1995) for a general overview of current research on learning in economic theory, and FUDENBERG & LEVINE (1997) for a comprehensive work with focus on learning in games.

¹⁷BLUME & EASLEY (1995, 13) refer to rational learning as a “*a poorly chosen euphemism for ‘Bayesian learning’*.”

¹⁸Following BULLARD (1991) at least three more approaches within rational expectations macroeconomics may be distinguished with respect to the forecast functions employed: forecast functions that (1) use only historical data, that (2) include the beliefs of others, and that (3) include frivolous variables (or “sunspots”).

¹⁹See e.g., SALMON (1995, 236) and BULLARD (1991, 57).

outcomes, however, is unsatisfactory since this would imply that individuals ignore relevant and potentially useful information.

2.2.2. Learning in Games

In the *game theoretic literature* (where z is qualified by some solution concept, e.g., the Nash equilibrium), several other mechanisms are introduced in addition to Bayesian learning to provide a better foundation for equilibrium theory, especially in settings with non-equilibrium expectations.²⁰ These models also involve various assumptions about the players' rationality, knowledge and beliefs, and about the ways players interact.²¹ While one type of games models the "learning" behavior of individual players, another type focuses on the aggregate behavior of a population (or its subgroups). To illustrate the involved concepts of learning, for each class a common adjustment mechanism will be presented below.

a) In *individual level approaches*, a common way to model learning is a quasi-Bayesian updating mechanism called "fictitious play" where players are assumed to behave as if they think they are facing an exogenous, stationary, unknown distribution of opponents' strategies.²² Hence, learning is modeled as statistical updating of information about the frequency that each strategy is played, starting from exogenous prior beliefs about the distribution of opponents' strategies. The process is backward oriented, i.e., players are anticipative only in a rather limited statistical sense, not in the sense that they think in terms of the effects of their own actions on their opponents' play, and, furthermore, players do not try to influence the future play of their opponents.²³ Since players only keep track of data about the frequency of opponents' play, they ignore data on conditional probabilities, and therefore may not recognize the emergence of cycles. Hence, the notion of human learning employed in this literature is relatively simple and myopic, and behavior is surprisingly unstrategic, if not naive. This simplification is often justified by its proponents with the "economically interesting" case of large populations of

²⁰Although it seems to be widely acknowledged that learning does not necessarily converge to any equilibrium concept (beyond the very weak notion of rationalizability), learning models are thought to suggest useful ways to evaluate and modify the traditional equilibrium concepts in that they lead to refinements of Nash equilibrium and to descriptions of long-run behavior weaker than Nash equilibrium (FUDENBERG & LEVINE, 1997).

²¹For a comprehensive introduction to learning in games and an overview of current research see FUDENBERG & LEVINE (1997). A short overview can be found in FUDENBERG & LEVINE (1998).

²²Note that there is no unique fictitious play rule, and there are several variants of fictitious play (like smooth or stochastic fictitious play) that all share the same notion of learning as described below. – Another class of individual level approaches that may be referred to as "stimulus-response" or "reinforcement" models of learning do not use the opponents' play as the object of updating, but the players' own payoffs in that strategies that yield higher payoffs become more likely to being played.

²³This is due to the assumption that the distribution of opponents' strategies is stationary which makes sense only if the system converges (at least in long-run behavior) and involves limited learning abilities by the opponents.

players where considerations of strategic interaction can be omitted, or may be defended by an “as if” argument in the case of the theory’s empirical success.

b) *Aggregate level approaches* are usually identified with evolutionary models that deal with the aggregate behavior of populations of players. The analysis typically focuses on the emergence and stability of behavioral strategies under most of the assumptions and conditions mentioned above.²⁴ Learning is commonly represented by a “*replicator dynamic*” that was originally motivated by models of biological evolution.²⁵ The replicator dynamic assumes that a population (or a fraction of it) playing a particular strategy grows in proportion to how well that strategy is doing relative to the mean population payoff.²⁶ – Since the replicator is only a mechanism of reinforcement of certain strategies, the link to learning processes is established by one of the following stories: One is the idea of “asking around” or “social learning” where players can only learn from other players in the population because they do not remember their own past experience, or because they are periodically replaced so that only new players make choices. A second story is based on the assumption that players do satisfice rather than optimize, i.e., they try to achieve only a certain aspiration level instead of a maximum. Hence, the payoff monotone dynamic²⁷ generated by the replicator it thought to imply that players choose only the best strategy “available in their neighborhood” (e.g., they compare an “inherited” strategy with one randomly observed), and do not screen all strategies for the one with maximum payoff.

Both of these stories – intended to motivate the reinforcement mechanism of replicator dynamics as being a good representation of human learning – do not seem to be very satisfying: In the first, individuals are assumed to be able only to copy existing strategies in an automated manner (thereby lacking any ability for cognition, creativity, or even learning from their own experience), and the second story matches more with a “best response behavior under bounded rationality”²⁸ than with (sophisticated) human learning.

²⁴In addition to the mentioned assumptions about rationality, information, interaction, etc., and populations may be homogenous or heterogenous. See FUDENBERG & LEVINE (1996) and WEIBULL (1998) for short introductions to learning and evolution in games.

²⁵A related concept that is also widely used, is the notion of an *evolutionary stable strategy*, i.e., a strategy that is robust when it is invaded by a small population playing a different strategy.

²⁶Note that the replicator is based on basically the same idea as the individual level stimulus-response or reinforcement models mentioned above.

²⁷Following FUDENBERG & LEVINE (1997) a model that generates a “payoff monotone dynamic” is a system in which the growth rates of strategies are ordered by their expected payoff against the current population, so that strategies that yield a payoff above the mean grow faster. There are several ways to formulate such a model, and the replicator is only one particular form of a payoff monotone dynamic in which the rates of growth are proportional to payoff differences with the mean.

²⁸“Best response” here means that players simply pick the strategy with the higher payoff when observing two (or more) strategies that exist in the current population. Note that *this* concept of “best response” as it is incorporated in the replicator dynamic *itself* refers more to a maximizing choice in a static situation than to learning.

More generally, most models abstract from forms of *sophisticated* learning in that considerations of repeated play are excluded.²⁹ Whereas a first level of sophistication may be included in the way that players use statistical procedures to forecast opposing players behavior, the second level where players consider the possibility that their current play may influence the future play of their opponent is mostly omitted by assuming large numbers of (anonymous) players, or simply by assuming that the incentives to alter the future play of the opponents are small enough to be neglected. – Other models that are not limited to large populations (e.g., KALAI & LEHRER, 1993) seem to hold only under very strong assumptions (e.g., about the agents' initial beliefs).³⁰ As BLUME & EASLEY (1995) have pointed out, existing results for models in which the only intertemporal link is learning are delicate, and one may remain rather skeptical when asking what can be learned from the current literature on learning in economics (indeed, Blume and Easley are pessimistic³¹).

2.2.3. Experimental Approaches: Tests of Learning Mechanisms

When robustness and validity of models – as well as the “as if” assumption put forward in the defense of unrealistic updating mechanisms, and the strong learning assumptions – are questioned, a natural answer is to test them in experiments.

Compared to the growth of the literature on “learning” theories, the body of experimental work to test proposed learning mechanisms is still relatively small. One reason may be that most theories are not designed for testing. Instead, they contribute to refinements of theoretical answers to theoretical questions, based mainly on introspection and casual empiricism since the unstated goal of most such analyses has been to predict behavior entirely by theory.³²

Hence, when testing a specific learning mechanism, a number of problems arise. One main problem is that the mechanism cannot be tested *per se*, that is, it is always tested in connection with a model or game that involves various additional assumptions.³³ Furthermore, most models contain little or no information about the conditions under which the proposed mechanism is assumed to be a good approximation of human learning behavior. Hence, a theory that would

²⁹For a summary of literature on repeated-game play and some experimental evidence see CRAWFORD (1995a, 40), who suggests that existing results raise difficult theoretical issues because significant learning often seems to occur within a single play of the repeated game, whereas standard methods for analyzing learning would assume repeated play of the repeated game, so that theoretical analysis should be combined with plausible restrictions on behavior, in order to understand these phenomena.

³⁰See KIRMAN & SALMON (1995, 2), and CRAWFORD (1995a, 3), who finds that “*theoretical analyses (traditional or adaptive) usually yield definite predictions only under strong assumptions, which are reasonable for some applications but unrealistic and potentially misleading for many others.*” [emphasis added]

³¹See KIRMAN & SALMON (1995, 2).

³²See CRAWFORD (1995a, 2).

³³This problem of testing joint hypotheses has a long tradition in empirical economics.

specify the conditions for and limits of the proposed learning mechanisms is missing, so that experimenters must pursue a course of trial-and-error-meta-learning in order to find out which mechanism describes human learning best in which setting (e.g., game).³⁴ Since every model or game is precisely defined by numerous restrictions and assumptions that critically influence the agents' actions and the model's possible equilibria, the learning process itself is likely to be strongly influenced by these restrictions and assumptions. – Moreover, as CRAWFORD (1995a, 3) has argued in his recent review of experimental studies of strategic interaction, none of the leading theoretical frameworks for analyzing games adequately identifies the principles that govern behavior by itself.³⁵ Thus, within this methodology the quest of finding a general learning mechanism – i.e., a theory that would enrich our understanding of learning in economics under a wider range of assumptions, and in other than perfect environments – may be questioned.

A traditional way of finding *more general* evidence for a given mechanism is to adjust a model to, or calibrate a model with experimental data,³⁶ and then to test it in different other situations. Several attempts to fit models with experimental data have been made. CHEUNG & FRIEDMAN (1994) fitted experimental data to a modified version of fictitious play with some success and found a better fit than for stimulus-response type models. ROTH & EREV (1995; see also EREV & ROTH, forthcoming) have proposed a stimulus-response (or reinforcement) model whose simple learning mechanism fits aggregate level data of several experiments relatively well (especially the intermediate term behavior). RAPOPORT ET AL. (1995) found that the adaptive dynamics of a modified version of the Roth-Erev model can be relatively accurate at the individual *and* aggregate levels in fitting the experimental data of a coordination (i.e., market entry) game.³⁷ VAN HUYCK, BATTALIO & RANKIN (1996) compared several learning models and report that exponential fictitious play fits their data best, whereas NAGEL & TANG (1997) find that some simple reinforcement models outperform the standard Nash equilibrium model and the quantal response model on experimental data in a centipede game. The results of a study by EREV & RAPOPORT (1997) also favor a reinforcement model, but show that this model cannot account for the observed effects of variations in the content of feedback information.³⁸

³⁴On the problem of evaluating learning models see also RAPOPORT ET AL. (1995, 35) who emphasize the lack of alternative models that may be compared.

³⁵CRAWFORD (1995a) examines experimental evidence with respect to traditional noncooperative game theory, evolutionary game theory, and adaptive learning models.

³⁶See ARTHUR (1994, pp. 135) for a calibrated learning algorithm for automata.

³⁷In DANIEL, SEALE & RAPOPORT (1996) and RAPOPORT, SEALE & WINTER (1997) an alternative reinforcement learning model to Roth/Erev was proposed and tested for individual level behavior.

³⁸For more models that are designed to fit experimental data see for example CRAWFORD (1995b), ZAUNER (1994),
(continued...)

2.3. Comment

To summarize the status quo, it seems fair to say that, despite its immense importance for the functioning of economic theory, *theoretical* approaches to learning in economics have so far addressed learning not as a behavioral feature, but – to paraphrase Blume and Easley – as a poorly chosen euphemism for farfetched statistical techniques and passive mechanisms of information accumulation and updating under mostly perfect situational conditions that cannot be justified neither by a sound concept of human learning, nor by empirical success.

Since *experiment-based approaches* are only a more recent development, it may be premature to judge the results, but so far only few learning mechanisms that would cover a wider class of models (or games) have been proposed, and attempts to test such a mechanism are relatively rare. As the number of models and experiments to test them increases, 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 (BUSH 1963). At the same time criteria for model comparison have to be developed that are widely missing today and that involve a number of difficult problems.³⁹ 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. Even if a single relatively general mechanism can be found, what can we conclude other than that some fundamental principles of learning which are well-documented in the psychological literature indeed appear to hold in interactive settings?⁴⁰ These findings may be interesting *per se* – though more to psychologists than to economists.⁴¹

³⁸(...continued)
and MCKELVEY & PALFREY (1992).

³⁹Criteria 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.

⁴⁰These 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).

⁴¹The 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
(continued...)

Nonetheless, for most applications in economics, rather than knowing the particular learning algorithm that governs behavior in a specific situation, it may be even more useful to understand the determinants or conditions that influence behavior in a variety of situations, and the limits imposed by these conditions. Hence, questions related to the application of an economic learning theory or its practical implications – for example questions of market design, the occurrence of anomalies, or government interventions in markets (all in which learning may play an important role) – may be difficult to answer within the prevailing methodology. To address this type of questions more directly, a different approach that parallels the traditional methodology in economics more closely may be appropriate, as will be discussed in the next Section. [For a short comparison of current approaches with the contingent learning approach see the Appendix.]

3. Methodological Issues

A main methodological characteristic of the *standard economic approach* is its focus on the analysis of changes in patterns of behavior (in terms of quantities produced, consumed, exchanged etc.) as a reaction to changes in restrictions or relative prices. Based on a bundle of assumptions that form the core of the theory, the analysis typically begins with the specification of situation variables that predetermine stable patterns of (equilibrium) behavior. Next, restrictions or relative prices are assumed to change, thereby inducing an adjustment process that leads to new equilibrium behavior. As discussed above, the adjustment process itself has traditionally been treated as a black box, and perfect conditions for this process are assumed implicitly.

Figure 1 shows this basic scheme of economic analysis. The theory core consists of strong assumptions about the economic actors (i.e., about rationality, preferences, autonomy).⁴² Additional assumptions about the actors' risk behavior, patience, time horizon, etc. may be introduced depending on the situation being modeled. This situation is generally defined by the type and complexity of the problem,⁴³ the number of involved actors,⁴⁴ and assumptions about

⁴¹(...continued)
average individual behavior.

⁴²Traditionally, the core also involves that the situation is one of scarcity. Although scarcity is a *conditio sine qua non* for economics, and for the definition of rationality, it can be constructed by various stories in virtually any situation so that it becomes an almost empty concept (at least in terms of methodology).

⁴³The “type” of problem may be one of consumption, production, investment, etc., whereas its “complexity” denotes the number of alternatives available, or dimensions of the problem for a given type of problem (e.g., an investment problem may be modeled more or less complex depending on alternatives and dimensions of the investment).

⁴⁴This number is rather important in economic analysis because it determines the degree of dependency among actors and the mechanism of interaction (see Section 4 for more details).

information.⁴⁵ These assumptions and definitions taken together determine the starting equilibrium (A, in Figure 1) as existing, stable patterns of behavior in terms of quantities.

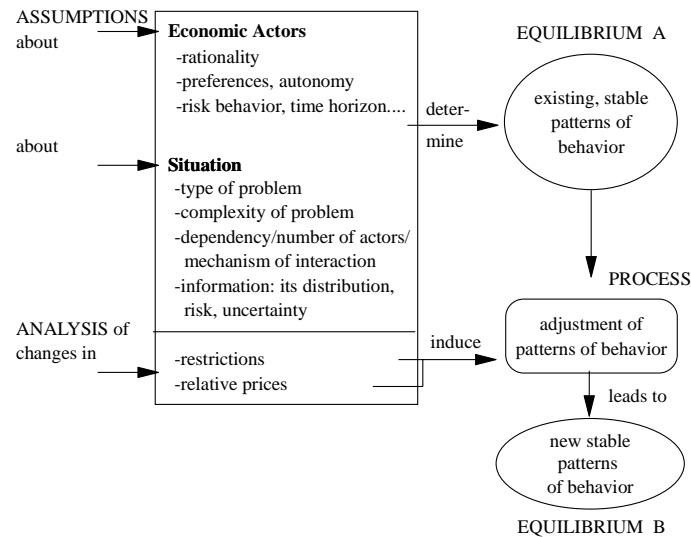


Figure 1: Basic Scheme of Economic Analysis.

The analysis then proceeds by introducing some exogenous stimulus in terms of changes in restrictions or relative prices that induces actors to adjust their patterns of behavior, eventually leading to new and again stable patterns of behavior (equilibrium B, Figure 1).

In learning models, this basic scheme is extended by introducing an adjustment mechanism that allows to distinguish between equilibria (B) for the case where several equilibria emerge. As discussed above, this is a way to “open the black box” of behavior adjustment, but one that may turn out to be problematic, especially when it comes to empirical testing due to the reasons already mentioned. Thus, an alternative way to approach learning in economics may be to “keep the black box closed”, and to identify the conditions for learning and the limits imposed by these conditions instead.⁴⁶ Henceforth, this idea will be referred to as *contingent learning*.

To acknowledge and analyze conditions for, and limits of learning as situational constraints may be viewed as similar to the introduction of cognitive constraints in models of bounded rationality put forward by Herbert Simon and others.⁴⁷ That is, whereas models of bounded

⁴⁵This includes assumptions about the distribution of information among actors, and a definition of the situation as one of certainty, risk, or uncertainty.

⁴⁶The more fundamental problem that economists face when trying to incorporate learning in economic models, may be analogous to the problem psychologists face when analyzing cognitive processes (see e.g., CATANIA [1992, 7]): Since learning (alike cognition) cannot be observed as such or directly, empirically based theories can be derived only from observable behavior, thereby leading to a variety of problems of induction.

⁴⁷For a recent survey see CONLISK (1996).

rationality relax the strong rationality assumption based on findings in cognitive science, a model of contingent learning aims to loosen the strong learning assumptions in order to provide a basis for richer economic models. Hence, contingent learning should be seen as a supplementary approach that may fruitfully be combined with existing economic theories and with bounded rationality models, not as their substitute.

In a behavioral view, the actual behavior of economic actors results from the interplay of cognitive factors on the one hand and situational factors on the other.⁴⁸ Economics has quite successfully focussed on rather extreme and specific forms of both types of factors. The mentioned problems of existing learning theory seem to suggest that analysis may not only be extended to include cognitive factors, but also additional situational factors in a behavioral – not behavioristic – sense.

As for the contingent learning approach described in the next section, standard assumptions about economic actors will mainly be maintained,⁴⁹ but the analysis of situational factors is extended by including determinants of learning processes. These determinants can be regarded as situational constraints that supplement the traditional economic restrictions. They reflect robust findings from psychology, and will be employed to account for the criticism presented in the previous section.

4. Contingent Learning

In a behavioral perspective, learning may generally be *defined* as an enduring change in behavior, or in the capacity to behave in a given fashion, which results from practice or other forms of experience.⁵⁰

The three main elements of this definition are (i) the occurrence of an actual or potential change in behavior, (ii) the persistence of this change, and (iii) the change's origin in practice or other forms of experience. Hence, learning leads to changes in behavior through interaction with the environment, and – since learning is a process – a concept of its basic dynamics is required that links the cognition and behavior of an individual to the situation.

Figure 2 presents a simple learning loop where cognition allows the individual to choose a behavioral strategy from a repertoire of strategies (that has been developed in previous interaction with the environment), and apply that strategy to a situation by some action that yields consequences which feed back to the individual. Depending on these consequences, the

⁴⁸As HEINER (1983) has argued, predictable behavior may result from the individuals' attempt to bridge what he calls the competence-difficulty-gap.

⁴⁹That is, (i) actors are assumed to be rational without any attributes, and (ii) methodological individualism is retained.

⁵⁰See eg., SCHUNK (1991).

repertoire may be modified, and a new or revised strategy may be chosen, thereby inducing an additional turn of the loop to a higher level of experience.⁵¹

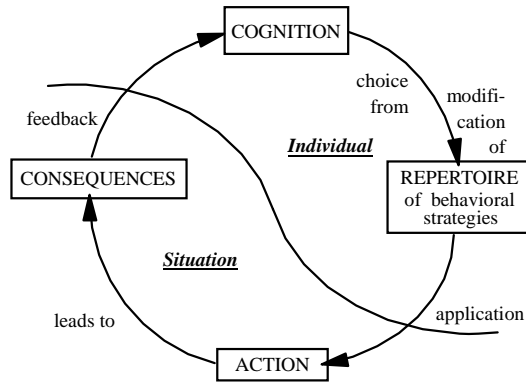


Figure 2: A Simple Learning Loop.

In order to integrate this basic concept of learning into the methodology of the economic approach as presented in the previous section, a traditional, stable pattern of equilibrium behavior (A) is assumed that is challenged by exogenous changes in economic restrictions or relative prices (see top of Figure 3). These changes induce cognitive processes, if a certain aspiration level is not satisfied – otherwise (unchanged) behavior B results.




Figure 3: An Extended Learning Loop.

Given that the aspiration level is not met, a behavioral strategy is chosen from the individual’s repertoire, and applied to the situation where it leads to consequences that feed back to the

⁵¹Hence, the loop may be better presented as a *spiral* where each turn represents a higher level of experience incorporated in an enlargement of the repertoire over time. The third dimension missing in Figure 2, thus, may be interpreted in terms of accumulation of knowledge, or cognitive development over time.

individual, inducing deliberation.⁵² If the feedback indicates that the chosen action yields consequences that do not restore the aspiration level, the applied strategy may be modified, or a new strategy may be chosen from the repertoire, and applied to the situation. This process continues until the aspiration level is satisfied, and a new equilibrium behavior (B) results.⁵³ Note that in Figure 3 the *situation* is characterized by five situational factors that will next be explained in more detail.

4.1. Determinants of Learning

Since, by definition, human learning evolves through the individual's interaction with the environment, it is a process that is always embedded in a situational background, and is therefore influenced by various situational factors. Some of these factors are accounted for in theoretical approaches to learning in economics, but are mostly trivialized, while others are idealized or completely omitted. These factors are nevertheless known to be important for actual learning behavior in psychology, and, thus, are likely to influence economic (learning) behavior in predictable ways.⁵⁴ For instance, existing institutions, or the actual design of a market may induce notable variations in these determinants, and thereby influence (learning) behavior and aggregate outcomes significantly. Hence, learning may be viewed as being *contingent* on the situational factors as discussed below.

Before discussing single determinants, two remarks seem appropriate. The first concerns the question in what sense learning can be assumed to be influenced by the determinants. In Section 4.2. the effects on learning will be operationalized in terms of outcomes, so that in the following discussion “easier” (or “more difficult”) learning is associated with better (or worse) outcomes. – The second remark concerns the isolation of the effects of single determinants. As always in the analysis of complex situations, the proposed effects of a single factor may be overlapping with the effects of other factors. Therefore, the following discussion of single factors has to abstract from possible cross-effects of factors.

Two groups of learning determinants (situational factors) can be distinguished: One group includes structural influences that exist *prior* to any (inter)action: (i) the complexity of the environment and the task; (ii) the number of actors involved with the situation, and (iii) the

⁵²The notion of an aspiration level is an example how bounded rationality can be combined with contingent learning, but this notion can be replaced by a maximizing assumption so that, given the exogenous changes, learning ends when utility is maximized under the new restrictions.

⁵³Again, this process may better be represented by a spiral where experience and knowledge accumulate over time, and behavior is modified until the aspiration level is restored. – For simplicity, the aspiration level is assumed to be constant, though aspiration level models in psychology suggest that the aspiration level may be affected by experience (see KLEINDORFER, KUNREUTHER & SCHOEMAKER (1993, 41) for references).

⁵⁴For basic texts on learning and behavior see MAZUR (1994), and CATANIA (1992). Formal models of learning can be found in LUCE, BUSH & GALANTER (1963).

degree of uncertainty associated with the structure of the situation. These factors are labeled *structural* determinants.

A second group of factors is also induced by the structure but is related to (inter)action: (i) content, quality, and quantity of feedback, and (ii) strategic uncertainty in situations of strategic interdependence; labeled *interaction* determinants. Table 1 summarizes the determinants and assesses their significance with respect to ideal types of situations.

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
interaction determinants			(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 Determinants for Learning

4.1.1. Complexity of and Information about the Structure

A situation (commonly modeled as an individual decision problem, a market, or a game) can be defined by its structure. The *structure* consists of the actors, the decisions they face, the information they have when making them (i.e., the actors' knowledge⁵⁵), how their decisions determine outcomes, and their preferences over outcomes. It also incorporates any repetition, correlating devices, or opportunities for communication.⁵⁶

With respect to the *complexity* of the structure, two aspects can be distinguished: (i) the complexity of the environment and (ii) the complexity of the task that has to be accomplished within an environment. The *environment* may vary in complexity in terms of the number and complexity of its elements, its dimensions, and the relations between them.⁵⁷ Within a given

⁵⁵Something is *mutual* knowledge if all actors know it, and *common* knowledge if all actors know that all actors know it, and so on ad infinitum.

⁵⁶See CRAWFORD (1995a, 3). – When behavior is influenced by additional factors, such as how the situation is presented or the social setting in which it takes place, this is sometimes called the *context*.

⁵⁷For example, a simple environment is a room with only a sheet of paper and a pen, whereas a fully equipped (continued...)

environment, the *task* may also vary in complexity depending on (i) the number of the dimensions of the task or problem, (ii) the number of combinations of dimensions, and (iii) the number of outcomes involved with particular actions or potential solutions. Hence, depending on these factors, the task may involve more or less complex decision making, calculus, etc., and, learning is fostered with decreasing complexity of task and environment.

For a given level of structural complexity, *information* about the structure available to actors may vary. In most cases, at least some of this information is given prior to any (inter)actions. Generally, the higher the level of information about the structure that is initially given, the easier is learning. But since complexity is not a simple function of the quantity of information available to the actors,⁵⁸ the effect of the *content* of structural information on learning must be analyzed for each situation separately. This analysis may be based on the assumption that the more the content of information reveals about the true nature of the structure, the easier is learning. – If actors know that they do not have all information about the structure, *structural uncertainty* results. The effects of structural (and strategic) uncertainty will be discussed in Section 4.1.3.

Note that in many situations learning about the structure involves experience, i.e., direct or indirect observation of actions and/or their consequences. Nevertheless, the underlying complexity of the structure, and information about the structure may influence learning significantly, and should therefore be analyzed separately.

In traditional economics, the complexity of the environment is often reduced to a minimum, and the task is usually a relatively simple choice among two alternatives. The effects of the complexity of the structure and information about the structure are commonly not analyzed as a separate variable, or as a determinant for learning processes.⁵⁹ – The effects of variations in structural complexity and information on learning discussed in this section, are summarized in the form of hypotheses in Section 4.3.

⁵⁷(...continued)

modern office is a relatively complex environment.

⁵⁸For example, a situation with “few information” may appear complex because many questions are left open, while a situation with “much information” may also appear to be complex because the additional information may be redundant and may obscure what is important.

⁵⁹See also HEY (1992, 95), who argues in favor of experiments with varying degrees of complexity because theories and experiments that are too much “stripped-down” fail to handle behavior in complex environments appropriately.

4.1.2. The Number of Actors: Degree of Dependency

Due to its significance in economics, one structural determinant that deserves special attention is the number of actors.⁶⁰ In economics, the mechanism of interaction among individuals (i.e., economic actors) is commonly modeled with respect to dependency and the number of actors involved: *Dependency* – in contrast to interdependency – means that there are many actors that depend on each other only via a single parameter. The *mechanism of interaction* in this situation is typically a (perfect) market that connects actors only via the market price. With a decreasing number of actors the “degree of dependency” increases in terms of increasing *interdependency*. The extreme case of interdependency is a situation with only two actors, where outcomes entirely depend on each other’s behavior. With several or few (to two) actors involved, the mechanism of interaction is often modeled as a strategic game (e.g., as a bargaining game). Hence, the analytical continuum runs from markets (dependency) to strategic games (interdependency) with decreasing number of actors (see Figure 4).⁶¹




Figure 4: Degree of Dependency.

With respect to the number of actors learning may generally be affected in the following way: Starting from a situation with many actors, and all other things equal, with a gradual decrease in the number of actors the influence of the other actors’ behavior on each other increases, so that it becomes more important for a single actor to take the behavior of others into account (monotonistic increase in dependency). Thus, more information has to be collected and processed by the actors as the behavior of single other actors becomes more relevant so that learning the other actors behavior becomes more important and difficult. However, at a certain number of actors there are so few actors (two at the margin) that learning becomes easier again (see Figure 5). In Section 4.3., a separate hypothesis is added to capture this idea.

⁶⁰Of course, in tasks of individual decision making the number of actors is always one so that the following discussion does not apply to these tasks.

⁶¹Note that perfect markets can be modeled as a n-person game in game theory, and that 2-person situations can be captured in traditional models. Thus, the degree of dependence does not imply a specific type of model, although there are clear advantages of each approach with respect to the number of actors.




Figure 5: Dependency and Learning

This line of reasoning implies, for example, that learning in perfect markets may be easiest, and may be most difficult in markets with several actors, while the difficulty of learning in two-person markets may be between these two extremes. However, as noted in the introduction to this section, in some situations this general pattern may be veiled by other structural factors, such as the availability of structural information, of knowledge, of correlating devices, of communication, and effects of collusion and reputation, so that learning (and the resulting behavior) may depend more upon these factors than on the number of involved actors. Therefore, it may not always be easy to isolate the effects of dependency on learning.

4.1.3. Uncertainty

Generally, and all other things equal, learning can be assumed to be more difficult under uncertainty than under certainty. Two types of uncertainty may affect learning: structural and strategic uncertainty.

Structural uncertainty results when actors have incomplete information about the structure and know this. An interesting case of structural uncertainty is uncertainty about knowledge (i.e., actors don't know what/if others know). Behavior is sensitive to this type of uncertainty and learning is likely to be more difficult under this condition.⁶² From a purely theoretical point it is difficult to make predictions about differences in the effects of incomplete structural information versus structural uncertainty. This may be mainly an empirical question. In Section 4.3 the hypothesis will simply be that uncertainty hinders learning.

Strategic uncertainty arises if – with increasing dependency among actors – the behavior of other actors becomes more important so that strategic considerations play an increasing role. Therefore, learning to anticipate the behavior of others also becomes more important as the degree of dependency rises. That is, actors must form expectations about the behavior of others; often called beliefs. Strategic uncertainty refers to the extent to which agents' beliefs differ, and

⁶²For an example of structural uncertainty in knowledge see ROTH & MURNIGHAN (1982) who found that the highest frequency of disagreement in a bargaining game of asymmetric occurred when only the player with the low price knew both prices, but was not sure if his opponent knew both prices (not common knowledge condition) because the player with the low price could not credibly signal that he had only a small price.

is reduced by learning about actions of others and the observation of the outcomes of these actions. Since this learning requires (inter)actions or at least information about the outcomes of (inter)actions, strategic uncertainty is a interaction determinant (see Table 1). However, because it requires feedback, as discussed in the next section, no separate hypothesis will be specified.

4.1.4. Feedback

The probably most obvious factor that influences learning is feedback. As shown in Figures 2 and 3, feedback closes the learning loop, and is therefore essential to all learning. Note that the concept of feedback is a broad one, in that it includes all kinds of information about the connection (causalities) of behaviors (actions) and consequences (outcomes) in a given situation. Furthermore, various sources of that information may be available, including observation of the behavior of others, of the consequences of that behavior, and sometimes communication with others. Therefore, feedback information can have different contents, which in turn is important for learning as discussed below (see Table 2).

Economic theories usually assume clear, complete, and instant feedback about the consequences of the actors own actions,⁶³ though some approaches differ in the assumptions about the availability of information about actions and outcomes of other actors. The latter type of information has been found to be important in some experiments,⁶⁴ but to date analysis of the effects of variations in content, quantity and quality of feedback seems rather eclectic.⁶⁵

That effective learning “*requires accurate and immediate feedback about the relation between the situational conditions and the appropriate response*” (TVERSKY & KAHNEMAN, 1987, 90) has been acknowledged in the psychological literature,⁶⁶ and the obstacles to learning with respect to quality and quantity of feedback as “external blocks” are well known.⁶⁷ That is,

⁶³See CAMERER (1995, 608), who describes this characteristic of the economic analysis of learning as an important source of disagreement between psychologists and economists.

⁶⁴See, for example, ROTH & EREV (1995), who suggest that feedback about others’ actions may be more important for the convergence of learning processes than information about players’ own payoffs, though the latter seems to speed convergence. See also RAPOPORT ET AL. (1995, 19), who report some effect of feedback on prior beliefs, and EREV & RAPOPORT (1997, 19) who found that “public feedback” fosters learning compared to “personal feedback”.

⁶⁵On the difficulties of learning only from feedback about outcomes in a psychological view, see DOHERTY & BLAZER (1988).

⁶⁶In most real-life situations quality and quantity of feedback is much lower than commonly assumed in economic theory because (i) outcomes are delayed and not easily attributable to particular actions, (ii) 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, (iii) there is often no information about what the outcome would have been if another action had been taken, and (iv) most important decisions are unique and therefore provide little opportunity for learning, see EINHORN & HOGARTH (1978).

⁶⁷In addition to these external blocks to learning, KLEINDORFER, KUNREUTHER & SCHOEMAKER (1993, pp. 110) identify two types of *internal blocks* (*ego defenses* and *cognitive biases*) that hinder learning due to the human
(continued...)

learning is fostered the more often a certain type of situation is experienced, the more frequent and immediate feedback is within a particular situation, the better outcomes are attributable to actions (i.e., the less confounded feedback is), and the less noise exists.⁶⁸ These findings apply at least to individual decision making but leave open the role of the content of information in interactive or strategic settings.

In individual decision making tasks the *content* of feedback is relatively straightforward in that it includes only information about the outcomes of the actors own actions. In markets, and especially in games, feedback may also include information about the actions and outcomes of others. This gives the following four main variants of the content of feedback (Table 2).⁶⁹

Variants	Content of Feedback			
	actor's own actions	actor's own outcomes	others' actions	others' outcomes
<i>a</i>	x	x		
<i>b₁</i>	x	x	x	
<i>b₂</i>	x	x		x
<i>c</i>	x	x	x	x

Table 2: Contents of Feedback

With respect to learning it seems reasonable to assume that the more the content of feedback reveals about the interaction, the easier is learning. Hence, learning is fostered when the content of feedback captures “more” relevant information (as in variant *c* compared to *a*). The effect of variant *b₁* compared to *b₂* is not straightforward. In some settings it may be more valuable for an actor to know the actions of others, while in other settings knowing the outcomes of others’ actions may be more instructive (e.g., because actions can be inferred from outcomes). The discussed effects of content, quality, and quantity of feedback on learning are summarized in hypotheses in Section 4.3.

⁶⁷(...continued)

psychological makeup. These psychological obstacles to learning are not covered by the approach presented here, because the focus will be strictly on situational factors that can more easily be controlled in experiments.

⁶⁸Note that confounded and noisy feedback is not the same. See KLEINDORFER, KUNREUTHER & SCHOEMAKER (1993, 111) for examples.

⁶⁹Note that Table 2 does not apply to individual decision making, and that, for simplicity, it is assumed that the actor always knows her own actions and outcomes. Both assumptions must not necessarily hold since actors may forget their own past actions over time, and there may be situations with no feedback about the actor’s own outcomes (but the actions and/or outcomes of others). This case seems rather extreme so that it is reasonable to assume that in most situations some feedback about the outcome of the actor’s own actions is available; though it may be delayed, confounded, or noisy.

4.2. Operationalization of Learning

In order to (experimentally) test for the effects of the proposed determinants of learning on behavior, the underlying concept of learning must be operationalized. Recall that, by definition, learning is an enduring change in behavior, or in the capacity to behave in a given fashion, which results from practice or other forms of experience. Thus, if learning occurs, its effects must be detectable in terms of changes in behavior or in the capacity to behave.

This raises the question how these changes can be conceptually captured or “measured”. It seems natural – and follows from the model of a learning loop – to assume that learning tends to improve behavior in terms of outcomes at least in the long run in most cases. In this view, “better” conditions for learning are assumed to lead to improvements of actions in terms of “better” outcomes.⁷⁰ The latter may be measured in various ways depending on the specific situation. For example, changes in outcomes may be gauged by the sum of received gains, incomes or payoffs over time (similar to the concept used to measure the “winner’s curse”; see KAGEL, 1995), or in terms of quantities (consumed, produced, exchanged, etc.) in an experiment.

Another question concerns the goal and the content of learning itself. In accord with the above operationalization – and with economic theory more generally –, it appears sound to assume that the implicit goal of learning is to maximize individual rewards (or payoffs), and the content of learning, therefore, is to behave (or play) optimally with respect to (i) the given situation (or game) *and* (ii) to the behavior (or play) of others involved in that situation (or game) in order to achieve that goal.

Note that in an interactive setting where the outcome for a single individual (economic actor, player) is not only the result of his or her own behavior but is affected by the behavior of other individuals, this individual has not only to learn to behave optimally with respect to the given situation but also with respect to the (dynamic) behavior of others. Therefore, learning to behave optimally *does not* imply a fixed or specific behavior with respect to a given situation as a result of the learning process. Hence, one cannot conclude from the observation of a certain type of behavior (say, cooperation) whether learning has taken place. It is, to the contrary, a characteristic of learning that the behavioral repertoire is enlarged, and behavior therefore becomes more flexible and disperse.

For example, in the repeated Prisoner’s Dilemma game, learning may induce playing a cooperative strategy with some opponents, but a defective strategy with other opponents in order to maximize individual payoffs. Thus, when we observe a player that has continuously

⁷⁰Note that changes in the capacity to behave – as included in the above definition of learning – cannot be measured with this operationalization. This may not be a problem though, because to economic theory learning and its determinants are not interesting *per se*, but it may be important only in terms of (aggregate) outcomes.

cooperated in the repeated Prisoner's Dilemma, thereby collecting a higher overall payoff than some other player who failed to cooperate with his opponent in the same type of game, we cannot conclude from the differences in payoffs that one player has learned and the other has not. The defective player may simply have been paired with a defective player or a player that was trying to exploit him so that he himself learned to defect (see HAUK, 1997, for supporting experimental evidence).

We can, however, observe many players under different learning conditions and look for systematic differences in outcomes or payoffs. If we find such systematic differences by comparing behavioral outcomes under varying learning conditions, we have evidence for or against the contingent learning hypotheses as described in the next section. This comparative-static method allows to analyze the effects of learning conditions in a systematic way, but it is a priori not clear which conditions yield which outcomes in which situations or games. To establish this evidence, controlled experiments are needed, and it may very well be that some hypotheses are falsified in some cases since they are deduced mainly from behavior observed in individual decision making tasks, not in interactive tasks or games.

Despite the difficulty of defining the content of learning and operationalizing its effect on observable outcomes in general terms, learning may generally be assumed to yield a best response behavior that maximizes individual outcomes (e.g. payoffs) with respect to the situation *and* the interaction with others.⁷¹

Moreover, the operationalization of learning, as proposed above, may require different experimental methods depending on the type of situation or game under investigation. Whereas the content of learning is relatively straightforward in games of pure cooperation or coordination (namely, learning to coordinate or cooperate), in some games it is more difficult to operationalize the effects of learning on outcomes. A typical example involves *zero-sum games* where not all of the players can increase their payoffs simultaneously, even if they all learn.

Generally, there exist two main methods to capture the effects of learning – and the effects of changes in learning determinants. The traditional way is to look at the *convergence* of behavior to some equilibrium or stable behavioral pattern over time. The effects of learning processes (on outcomes) and learning determinants (on learning processes) may then be assessed in terms of the *characteristics and robustness of the equilibrium* selected by the learning process,⁷² and by the *speed of convergence*. A possible assumption for the operationalization would be that behavior converges faster and is more stable under favorable learning conditions than under unfavorable conditions. The advantage of this method is that it

⁷¹Note that this formulation omits any notion of intrinsic value of behavior itself, e.g., the fun of playing a game or being part of a situation/experiment, as can sometimes be observed in experiments.

⁷²Equilibria may be characterized by game theoretic concepts or by the Pareto-criterion (where possible).

is conceptionally lucid, though no straightforward criteria exist to determine when behavior can be assumed to have settled. Also, normative implications are not easy to deduce.

An additional way to test learning effects is to vary learning determinants asymmetrically, so that some players in a zero-sum game face learning conditions that are assumed to be favorable for learning while other players face conditions that are thought to hinder learning. The general prediction is that players facing unfavorable learning conditions tend to perform relatively worse than their opponents in terms of accumulated payouts in a comparative-static analysis. The effects of most contingent learning determinants (e.g., structural information, uncertainty, feedback) can be tested in zero-sum games if introduced asymmetrically, whereas some cannot (e.g., the number of actors). The advantage of this method is that it directly maps learning conditions to the outcomes of individual behavior.

4.3. Hypotheses about Determinants of Learning Processes

Following the operationalization, the effects of the proposed determinants on learning are summarized in the hypotheses below in terms of quality of learning, where “better learning” is associated with “better outcomes”. Also, it is possible to use the commonly used measure of convergence to some equilibrium where “better learning” is, eg., associated with faster convergence.

Determinant		Hypotheses (ceteris paribus)
Com- plexity	Environment	• The lower the complexity of environment, the easier is learning.
	Task	• The fewer actions are available, the easier is learning. • The simpler the available actions are, the easier is learning.
Structural Information		• The more (initial) information reveals about the structure, the easier is learning.
Degree of Dependency		• With decreasing degree of dependency (i.e., an increasing number of involved individuals) learning initially becomes more difficult, but becomes easier when the number of involved individuals is increased beyond a certain number [see Fig. 5]
Uncertainty		• The more certain a situation appears, the easier is learning.
Feed- back	Quantity	• The more often a certain situation is experienced, the easier is learning with respect to that situation (absolute frequency of feedback). • The more often feedback occurs within a given situation, the easier is learning with respect to that situation (relative frequency of feedback).
	Quality	• The faster (or the less delayed) feedback is available relative to actions or outcomes, the easier is learning. • The better feedback is attributable (or the less confounded it is), the easier is learning. • The less feedback is disturbed by noise, the easier is learning.
	Content	• The more the content of feedback reveals about interaction (actions and outcomes), the easier is learning [see Table 2]

Table 3: Synopsis of Contingent Learning Hypotheses

5. Discussion

The contingent learning approach (CLA) aims to identify the factors that influence learning in order to explain and predict adaptive behavior. Every situation has its specifics, and there are always several factors that influence learning simultaneously, though not necessarily in the same direction. That is, the effects summarized in the above hypotheses may support or offset each other so that the net effect of the relevant determinants in a given situation is not always straightforward, but must be analyzed carefully. Two implications follow:

First, it may sometimes not be difficult to find or make up examples where single hypotheses of the CLA are expected to be falsified, because other learning determinants, or additional factors (such as the notion of fairness in some games) are dominant for the behavior in that example. Therefore, it is desirable to test the CLA hypotheses under *ceteris paribus* conditions. That these conditions can be identified, and that the determinants and effects can theoretically be distinguished may be a methodological advantage of the CLA. However, in some cases it may be difficult to provide a strict *ceteris paribus* analysis or test since determinants and effects may be interwoven, and outside factors may be important (such as framing effects).

Second, it therefore follows that the theorist's task is to untangle the net effect of the determinants on learning, and analyze them as commonly done in comparative statics. As in standard economics, a theorist who wants to model a situation where learning is involved may carefully choose and include the determinants for learning that appear to be relevant in that situation. Hence, the CLA aims to enrich economic theory by providing useful tools for model building, but does not make the art of doing so obsolete.

It follows from this short discussion that the CLA hypotheses cannot be assumed to be generic statements that hold in every instance. Rather, they reflect general tendencies that are likely to influence learning in the direction indicated by the hypotheses in many situations. The range of economically relevant situations in which they allow to explain and predict behavior remains to be explored, however.

6. Concluding Remarks

The contingent learning approach outlined in this paper is the first step in a larger research program. It aims to develop a theory that enriches the standard economic approach in a methodologically compatible way by introducing theoretical elements and evidence from other fields with respect to learning behavior in economically relevant situations. The approach therefore includes determinants in the form of situational constraints that human learning can be assumed to be contingent upon. Hence, in addition and in analogy to the *economic* restrictions in traditional economic theory – and to the *cognitive* restrictions in approaches of bounded rationality – the CLA introduces *learning* restrictions that account for the interplay between the

individual and the situation. In other words, whereas standard economics focusses on a specific type of restrictions (i.e., on endowments, budget constraints, relative prices), and the bounded rationality approach focusses on cognitive restrictions, the CLA focusses on the interplay between situational and cognitive factors, thereby bridging part of the gap between traditional economics and psychology. This additional perspective may be especially important if economics is understood as a behavioral social science whose strength lies – due to the prevailing methodological paradigm – in the explanation and prediction of *changes in individual behavior* as a reaction to *changes in restrictions and relative prices*.

The development of such a theory is the first step. It is followed by testing the hypotheses that are deduced from theory. Since the CLA hypotheses presented above are quite general and –some may argue– straightforward, it seems to be a good research strategy to follow Popper’s advice in trying to find examples where the hypotheses can be falsified, instead of finding supporting evidence, as many researches do in this field today. This enterprise is on its way and a first series of experiments has been conducted (SLEMBECK, 1998a, 1998b). The results are encouraging in that it appears to be difficult to falsify at least some of the hypotheses. It seems that some cases that appear to present counter-examples to the CLA hypotheses at first sight, show to be no such examples once overlapping effects of learning and other variables are untangled. But it has to be admitted that it is difficult to create experiments that provide tests that are truly *ceteris paribus*.

The most important and most challenging task, however, is the application of the theory to a wider range of economically relevant phenomena. Hence, the real test for the power of an economic theory enriched by contingent learning lies within its application, and may emerge in competition with other approaches that aim to explain the same phenomena.⁷³

⁷³For a general comparison of the CLA to current theoretical and experimental approaches see the appendix.

APPENDIX: A Comparison of Learning Approaches

Given that all approaches to learning in economics aim to explain or predict behavior eventually, approaches differ in their primary goal and in the method they employ.

	Theoretical Approaches	Experimental Approaches	Contingent Learning Approach
Primary Goal	-reduce number of equilibria -model out-of-equilibrium beliefs / behavior	-predict how a certain class of games is played over time	-predict adaptive behavior contingent on situational factors / constraints
Method	-introspection -stipulate mechanism -determine equilibria (prescribe behavior)	-track behavior -build model/algorithm (motivated by psych. literature) -calibrate/test algorithm -demonstrate fitness/generality -predict behavior (dynamic, qualitative)	- identify relevant constraints - test effects of variations of constraints on behavior - predict behavioral (comparative static, qualitative)
Learning Mechanism	specified	specified	unspecified
Result	refined equilibrium theory: learning mechanism as part of some model or game	learning algorithms specific to some classes of games	theory enriched by situational constraints (in analogy to economic and cognitive constraints)
Type of Contribution	mainly normative	mainly positive	normative and positive
Main Level of Contribution	explanation	description/prediction	prediction
Policy Implications	low	medium	higher

Table 4: A Comparison of Learning Approaches

References

- Arrow, K.J.** (1987): *Rationality of Self and Others in an Economic System*, in: Hogarth, R.M., Reder, M.W. (Eds.): *Rational Choice – The Contrast between Economics and Psychology*, 201–215.
- Arrow, K.J.** (1994): *Methodological Individualism and Social Knowledge*, American Economic Association Papers and Proceedings, May 1994, 1–9.
- Arthur, W.B.** (1994): *Increasing Returns and Path Dependence in the Economy*, Michigan, University of Michigan Press.
- 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.

- Bullard, J.B.** (1991): *Learning, Rational Expectations and Policy: A Summary of Recent Research*, Federal Reserve Bank of St.Louis, Vol. 73/1, 50–60.
- 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.F.** (1995): *Individual Decision Making*, in: Kagel, J.H. & Roth, A.E. (Eds.): *The Handbook of Experimental Economics*, 588–703.
- Camerer, C.F., Loewenstein, G. & Weber, M.** (1989): *The Curse of Knowledge in Economic Settings: An Experimental Analysis*, *Journal of Political Economy*, No. 97, 1232–1254.
- Catania, A.C.** (1992): *Learning*, Englewood Cliffs, N.J., Prentice-Hall Int. Inc.
- Cheung, Y. & Friedman, D.** (1994): *Learning through Reinforcement on the Replicator Dynamics*, Santa Cruz.
- Conlisk, J.** (1996): *Why Bounded Rationality?*, *Journal of Economic Literature*, Vol. XXXIV, 669–700.
- Crawford, V.P.** (1995a): *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 University Press, Cambridge (UK) & New York (US), 1997.
- Crawford, V.P.** (1995b): *Adaptive Dynamics in Coordination Games*, *Econometrica*, Vol. 63/1, 103–143.
- 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.
- Doherty, M.E. & Blazer, W.K.** (1988): *Cognitive Feedback*, in: Brehmer, B. & Joyce, C.R.B. (Eds.): *Human Judgement – The SJT View*, Amsterdam, North-Holland.
- Einhorn, H.J. & Hogarth, R.M.** (1978): *Confidence in judgement: Persistence in the illusion of validity*, *Psychological Review*, Vol. 85, No. 5, 395–416.
- Erev, I. & Rapoport, A.** (1997): *Coordination, “Magic”, and Reinforcement Learning in a Market Entry Game*, Technion, Haifa, mimeo.
- Erev, I. & Roth, A.E.** (forthcoming): *Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria*, forthcoming in *American Economic Review*.
- Fudenberg, D. & Levine, D.K.** (1996): *Learning and Evolution in Games*, paper based on a plenary talk presented at the 1996 Summer Meetings of the Econometric Society (June 26, 1996; revised July 30, 1996), @ <http://levine.sscnet.ucla.edu/Papers/Essay/Essay7.htm>.
- Fudenberg, D. & Levine, D.K.** (1997): *The Theory of Learning in Games*, MIT Press, Cambridge.
- Fudenberg, D. & Levine, D.K.** (1998): *Learning in Games*, *European Economic Review*, 42 (3-5), 631–639.
- Hauk, E.** (1997): *Discriminating to Learn to Discriminate*, Working Paper No. 97/61 (September), Departamento de Economía, Universidad Carlos II de Madrid.
- Heiner, R.A.** (1983): *The Origin of Predictable Behavior*, *American Economic Review* 73, 560–595.
- Hey, J.D.** (1992): *Experiments in Economics – and Psychology*, in: Lea, S.E.G, Webley, P. & Young, B.M. (Eds.): *New Directions in Economic Psychology*, Aldershot/Hants, Edw. Elgar, 85–98.
- Kagel, J.H.** (1995): *Auctions: A Survey of Experimental Research*, in: Kagel, J.H. & Roth, A.E. (Eds.): *The Handbook of Experimental Economics*, 501–585.
- Kalai, E. & Lehrer, E.** (1993): *Rational Learning Leads to Nash Equilibrium*, *Econometrica*, 61, 1019–1045.
- Kirman, A. & Salmon, M.** (Eds.) (1995): *Learning and Rationality in Economics*, Oxford/ Cambridge, Blackwell.
- Kleindorfer, P.R., Kunreuther, H.C. & Schoemaker, P.J.H.** (1993): *Decision Sciences – An Integrative Perspective*, Cambridge.

- Lucas, R.E.** (1981): *Studies in Business-Cycle Theory*, Cambridge, MIT Press.
- Lucas, R.E.** (1987): *Adaptive Behavior and Economic Theory*, in: Hogarth, R.M., Reder, M.W. (Eds.): *Rational Choice – The Contrast between Economics and Psychology*, 217–242.
- Luce, R.D., BUSH, R.R. & GALANTER, E.** (Eds.) (1963): *Handbook of Mathematical Psychology*, New York etc., Wiley & Sons.
- Mazur, J.E.** (1994): *Learning and Behavior*, New Jersey, Englewood Cliffs, Prentice Hall.
- McKelvey, R. & Palfrey, T.** (1992): *An Experimental Study of the Centipede Game*, *Econometrica*, Vol. 60, 803–836.
- Nagel, R. & Tang, F.F.** (1997): *Experimental Results on the Centipede Game in Normal Form: An Investigation on Learning*, Universitat Pompeu Fabra, Barcelona, mimeo.
- North, D.C.** (1994): *Economic Performance Through Time*, *American Economic Review*, Vol. 84, No. 3 (June 1994), 359–368.
- 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.
- Rapoport, A., Seale, D.A. & Winter, E.** (1997): *Coordination and Learning Behavior in Large Groups With Asymmetric Players*, University of Arizona, mimeo, February 25, 1997.
- 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.
- Roth, A.E. & Murnighan, J.K.** (1982): *The Role of Information in Bargaining: An Experimental Study*, *Econometrica*, No. 50, 1123–1142.
- Salmon, M.** (1995): *Bounded Rationality and Learning: Procedural Learning*, in: Kirman, A. & Salmon, M. (Eds.): *Learning and Rationality in Economics*, 236–275.
- Saunders, S.** (1995): *Experimental Asset Markets*, in: Kagel, J.H. & Roth, A.E. (Eds.): *The Handbook of Experimental Economics*, 445–500.
- Schelling, T.C.** (1960): *The Strategy of Conflict*, Cambridge & London, Harvard University Press.
- Schunk, D.H.** (1991): *Learning Theories – An Educational Perspective*, New York.
- Simon, H.A.** (1978): *Rationality as Process and as Product of Thought*, *American Economic Review, Papers and Proceedings*, Vol.68., No.2, 1–16.
- Slembeck, T.** (1997): *The Formation of Economic Policy: A Cognitive-Evolutionary Approach to Policy-Making*, *Constitutional Political Economy* (8), 225–254.
- Slembeck, T.** (1998a): *As If Playing Fair – Experimental Evidence on the Role of Information in Ultimatum Bargaining*, mimeo, June 1998, University College London.
- Slembeck, T.** (1998b): *Reputations and Fairness in Bargaining – Experimental Evidence from a Repeated Ultimatum Game With Fixed Opponents*, mimeo, August 1998, University of St.Gallen.
- Thaler, R.H.** (1991): *Quasi Rational Economics*, New York, Sage.
- 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, U of Chicago Press.
- Van Huyck, J., Battalio, R. & Rankin, F.** (1996): *On the Evolution of Convention: Evidence from Coordination Games*, Texas A&M, mimeo.
- Vives, X.** (1996): *Social Learning and Rational Expectations*, *European Economic Review*, Vol. 40, 589–601.
- Weibull, J.W.** (1998): *Evolution, Rationality and Equilibrium in Games*, *European Economic Review*, 42 (3-5), 641–649.

Winter, S.G. (1987): *Comments on Arrow and on Lucas*, in: Hogarth, R.M., Reder, M.W.: Rational Choice – The Contrast between Economics and Psychology, 243–250.

Witt, U. (1992): *Evolution as the Theme of New Heterodoxy in Economics*, in: Witt, U. (Ed.): Explaining Process and Change – Approaches to Evolutionary Economics, S. 3–20, Ann Arbor, Uni of Michigan Press.

Zauner, K (1994): *A Reconsideration of the Centipede Game Experiments*, mimeo, University of California San Diego.