

Richard F. Bonner¹, Tetyana I. Mamchych²

¹ Mälardalen University, Västerås, Sweden

² Volyn State University, Lutsk, Ukraine

MAKING SENSE OF KNOWLEDGE MANAGEMENT: RATIONAL CHOICE, LEARNING, AND THEIR INTERPLAY

Abstract. We observe a learning problem, which occurs when changes in the knowledge system of a firm (learning) alter its business objectives (preference). Grounds for evaluating learning may become known only *after* the learning. Moreover, there may simply be *no ground* for comparing consequences of learning strategies, unless their resulting preferences are expressly reconciled. Routine reconciliation by monetary value is often ill founded. Logical foundations for managing knowledge in a firm are thus openly put into question. The problems of evolution of preference through learning, and the rationality of such learning, must be formally addressed. As a first step towards some answers, we briefly scan the current state of theories of learning and rational choice. Central to rational choice theory, and to decision schemes of artificial intelligence, is the exploration/exploitation problem of effective dynamical choice of action by an economic agent; here, the agent learns the environment and the consequences of its actions, but agent's taste is normally set, not subject to learning. Under this assumption, and in the presence of bayesian prior, rational learning is in principle well understood. Considerations of pragmatic flavor, however, largely question the presence of bayesian prior. Learning without bayesian prior is studied in empirical models; here, however, rational grounds are uncertain. Without the assumption of persisting preference, the potential preferences must be reconciled, as in the theories of social choice and reference-dependent preferences; however, questions of learning have not here been addressed. Thus, current theory leaves the learning problem largely open, but the tools appear to be in place, inviting investigation.

Key words. *computational learning, decision process, economic search, empirical learning, knowledge management, rational choice, rational learning,*

1. A problem

Central to any life science, of which the field of business information systems is a prime example, is the problem of choosing between acting and learning. Immediate economic action (sell, buy, employ, invest, merge, etc.) may often be delayed pending learning (researching markets, educating employees, validating information, developing IT support, etc.), so that a better choice of action may be arrived at.

With criteria of choice of action, in function of knowledge of environment, consistently set, the problem whether to act on the spot, or to postpone action pending learning, invites few ground considerations. Granted due prior knowledge of the environment, such problems are in principle solvable by dynamical programming techniques, now standard in the theory of decision processes. Thus, with business objectives of a firm explicitly set, and unaffected by potential learning, it makes sense to say that the knowledge management function in the firm should support its business function. The information systems of IKEA do help to 'furnish the world'.

We all know, however, that such assumptions are not always met. Education changes preference. Business objectives evolve with learning. The adoption of a major information system may render the very reasons for its adoption obsolete.

The logic of the problem is elementary. Assume perfect information. Suppose our initial preference is \leq , the consequence of learning x_i is X_i , and our resulting preference is \leq_i , $i = 1, 2$. The preferences \leq_i may at the outset be known, partially known, or not known at all. How can we choose x_i ? Suppose \leq_i known. Choosing by \leq is questionable, as it may leave us dissatisfied with X_i in our new identity \leq_i , unless the new preference does not differ much from the old (and how to measure dissimilarity?). Choosing by any of \leq_i makes sense if both rank X_1 and X_2 the same, say $X_1 \leq_i X_2$, $i = 1, 2$. But what if $X_1 <_1 X_2$ and $X_1 >_2 X_2$, say? Could we reconcile conflict by a 'social' preference $\leq_{1,2}$? On what grounds would we choose it? Lack of information embeds the problem in frameworks of decision processes.

Formal grounds for knowledge management in a firm call for clarification. The problem voiced, however, is not unique to business firms. It occurs in all planning situations, where the actions planned affect action evaluation criteria. One is thus lead to formally examine rational decision from the standpoint of learning, and, in the opposite direction, to examine the rational grounds for learning. We tentatively scan the state of the two fields, and their intriguing and largely open connections. The elusive nature of evolving identity calls for precision of language, which only

formal models may provide. This sets a tone of approach, complementary to disciplinary discussions and domain pragmatics, both left for parallel efforts.

The present note aims to anchor voiced questions in the standard models of rational choice and learning. It does not attempt to state them formally, let alone to answer them. Sections 3 and 4 recall the standard bayesian set-up for one-step and many-step decision problems, respectively. Section 5 recalls the non-bayesian empirical learning model. Section 6 points to work on reconciliation of preferences, as in problems of social choice. Only Section 6 addresses the logic of the learning problem; the other sections address procedural and statistical frameworks for its setting with imperfect information. The References, not all cited, collect ground material.

2. Background

Connections of rational choice and learning are indeed intricate. In essence, rational choice is a principle fixing the behavior of an agent, while learning manifests itself in changes of behavior. Experience affects the agent's taste, the taste determines the agent's choices, and the choices affect the agent's further experience. Thus, systematic changes in the behavior of a rational agent in a stable environment may be thought of as due to change of taste (the agent's perception of desirability of potential consequences of its acts), or due to the agent's improving ability to act towards maximum pleasure according to persisting taste.

Two kinds of learning processes are thus immediately distinguished: (i) learning in absence of persisting taste such as the formation of the taste itself, perhaps subject to meta-preference, (ii) learning in presence of persisting taste, with respect to which the agent's choices successively 'improve'. The economists' models of learning [15, 17, 18, 36, 88] address mainly the latter case in the presence of bayesian prior; it may be then noted that the principles of learning within a theory of (dynamic) rational choice, are still much discussed [10, 26, 47]. Among these are batch learning problems, where the environment is first sampled, a theory is formed, and the theory is thenceforth exploited. With fixed algorithm for forming theory, the sole decision variable is here the stopping time; stopping problems are well represented in decision theory [2, 14, 19, 33, 103]. The interesting former case, (i), of evolution of preference, seems better recognized in biological fields [25], and seems to lack formal theory.

Independently of economic theory, problems of learning have been in recent decades much studied in the computational sciences, in particular in connection with neural computation [8], where no bayesian prior is normally present. This is empirical learning, the subject of Statistical Learning Theory, largely due to Vapnik [95], and of Computational Learning Theory [7, 8, 65, 97], going back to Valiant [94]. For an overview of the two schools see [96] and [4, 5], respectively, and see [87] for an abstracted mathematical account of the ground ideas.

Essentially, the concern here is the 'learnability' of classes of objects, that is, the feasibility of identifying an object within a class by random sampling in a stable but unknown environment; part of the theory is the provision of bounds on the sample size in function of parameters of accuracy of identification. Statistical foundations of empirical learning rest on the theory of empirical processes [79] and uniform central limit theorems [30]. Since Valiant's paper [94], convergence uniform in probability is also known as 'probably approximately correct' convergence, whence the now established term of PAC learning.

Fundamental connections of empirical learning and rational choice are largely uncharted. Computational Learning Theory, despite its recognized pragmatic flavor, pronounced already by Valiant [94] and in several papers on its decision-theoretic generalization [54, 55, 35], has largely avoided contact with economic theory. Statistical Learning Theory poses the learning problem as the problem of empirical optimization of a risk functional interpretable as expected utility, but it does not pursue this interpretation.

In its basic formulation, the empirical model is a batch learning model, where prior knowledge of the environment is ambiguously presented as a set of probability measures. The expected utility of learning is then ambiguous, and questions of alternative - non-rational in standard sense - decision criteria arise. Discussions of non-standard criteria are broad [1, 12, 29, 31, 62, 63, 64, 67, 86, 101, 102, 104], especially from the standpoint of computational science [26, 27],

3. Rational choice

Roughly, the notion of rational choice is an artifact designed to explain (control, coordinate, norm, predict, etc) the behavior of individuals and groups. An agent is an individual whose behavior is defined by its action, a function of choice of act. A rational agent, roughly, acts towards (uncertain) consequences of highest appeal. Formal foundations of 'rational choice under uncertainty' were laid by Savage [83] in the 1950's, establishing expected utility as a universal criterion of rational choice; see also [24, 59]. Most of subsequent economic theory subsumed Savage, essentially equating rational choice with the bayesian framework.

Recall briefly the principal setting. An act is a map $f: X \rightarrow Y$ associating with a 'state of the world' x the consequence $f(x)$ of act f in state x . Denote the set of all admissible acts by $\Phi \subset Y^X$ and let \mathcal{F} be a family of sets F of acts $f \in \Phi$. A primary element of agent's knowledge is an event, a subset E of X . The set \mathcal{E} of all admissible events usually carries an algebra structure with some completeness properties. The agent's choice function is a map $\gamma: \mathcal{E} \times \mathcal{F} \rightarrow \Phi$, which for every event $E \in \mathcal{E}$ selects in every $F \in \mathcal{F}$ an act $\gamma_E(F) \in F$; the agent's choice function defines the agent's action or behavior. The agent's preference function is a map $<$ taking events $E \in \mathcal{E}$ to preference (reflexive, transitive, and complete) relations $<_E$ on Φ ; a preference function is normally further qualified by consistency conditions.

An agent is rational with respect to a preference if it always picks acts of maximum preference, that is, if $g \prec_E \gamma_E(F)$ for all $g \in F \in \mathcal{F}$ and $E \in \mathcal{E}$. The action of a rational agent may thus be 'explained' by its preference.

Classically, under standard measurability conditions, if consequences are ordered by a real-valued utility function u on Y so that $y_1 \in Y$ is more appealing than $y_2 \in Y$ if and only if $u(y_1) \geq u(y_2)$, and if uncertainty of event is expressed by a probability measure μ on \mathcal{E} , then the conditional expected utility $\mathcal{E}_{u,\mu}(f|E) := (\int_E u \circ f d\mu) / \mu(E)$ defines a preference \prec on acts, $f \prec_E g$ if and only if $\mathcal{E}_{u,\mu}(f|E) \leq \mathcal{E}_{u,\mu}(g|E)$, $E \in \mathcal{E}$, $f, g \in \Phi$. The measure μ may here have statistical origins, or it may represent the agent's prior beliefs in the bayesian decision model. Savage's [83] orthodox justification of the bayesian approach, recall, consisted in showing that any preference \prec on a set Φ of acts subject to some compatibility conditions valid for an expected utility preference - the six Savage axioms - is indeed realized by expected utility with respect to some distribution μ on X , and some utility u on Y , both essentially determined by \prec . Thus, in standard interpretation, behavioral rationality can be cognitively explained: an agent, whose choice function γ induces a Savage preference ($f \prec_E g$ if and only if $\gamma_E(\{f, g\}) = g$), behaves as if maximizing some expected utility $\mathcal{E}_{u,\mu}(f|E)$ with μ and u essentially determined by γ .

Problems. Its normative merits notwithstanding, the fundamental bayesianism as laid by Savage has been much discussed, on general [40, 47, 73, 82] and on empirical [32, 22, 37] grounds. The thrust of voiced arguments puts into doubt whether agents are in practice capable of rational behavior as prescribed by the Savage axioms, and whether such behavior could effectively be engineered from the information available to the agent. Also troubling has been the subjective nature of bayesian rationality, providing no link to empirical formation of beliefs of agents. Proposed relaxations of the Savage model include weaker axioms for preference [52, 91, 92], non-expected utility [21, 28, 41, 42, 53, 75, 84], and symbolic treatment [29]. Equilibrium models with non-expected utility are under investigation [98]. Ground notions of information as means for action are being revised [43, 44, 45, 46]. Much work is fuelled by developments in the cognitive and computational sciences, in particular in artificial intelligence, where problems often hinge on computational implementation of rationality [26, 27, 29, 86, 101, 102].

4. Rational learning

Roughly, a decision process is a sequence of choices made in function of consequences of previous choices, and rational learning is the learning in, or of, a decision process. Rationality of learning is qualified in terms of a consistent preference on the consequences of the process; this, by Savage, is formally equivalent to the presence of a utility function on the consequences, and a bayesian prior expressing uncertainty.

A decision process in discrete time is operationalized as a sequence of transactions of an agent (A) with its environment (E), both ultimately modelled as computers. The agent acts through queries and moves, the environment responds to queries with answers, and to moves with rewards. At times t_1, t_2, \dots , the agent and the environment engage in a 'cycle' of potential transactions: (i) A picks a query, puts it to E, and E responds with an answer, (ii) A picks a move, puts it to E, and E responds with a reward, (iii) A picks a meta-query (a query on its choice of move), puts it to E, and receives an answer. The execution of any query has no bearing on the answer of any later query, nor on the reward from any later action, while the execution of a move in general has. Stages (i) and (iii) are exploratory, the former exploring the environment, the latter - the untested moves. Information on untested moves is then thought to come from a 'teacher' or an 'oracle'; in a context of a game - when the environment contains other agents - it can also be thought of as 'private information'. Stage (ii) explores the reward function by sampling at moves of unknown reward, and 'exploits the environment' by picking moves of known high reward; in the context of a game, the information it provides is typically 'public'.

In general, the agent's choices of action, and their execution, all carry costs to the agent; the costs of storing and processing information are thought to be internal to the agent, the costs of interacting with the environment are then external. The agent's choices of action, in function of past experience, constitute the agent's strategy. A Savage-rational agent chooses a strategy so as to maximize the agent's expected utility of the resulting stream of costs and rewards, which a strategy generates. Note that neither the costs nor the rewards need here be monetary.

During the process, the agent's information about the environment increases through the answers to its queries. Roughly, learning about the environment takes place if this information at some time allows the agent to infer something not known from the start about the answers to further queries. This learning may or may not take place; in itself, it is in general not the objective of a rational learner. Likewise, learning about the choice of move takes place if at some time the expected reward on the choice to be made by the agent is increased by past experience. Neither this learning needs to be the objective of a rational agent; the learning cost could be too high. This learning is said to be 'supervised' if stage (iii) of the process cycle is nontrivially present; if the meta-queries in answer require a choice of move of highest reward in the current situation, the supervised learning is said to be 'by example'. Otherwise, if stage (iii) is empty, the learning is said to be 'by reinforcement'.

Thus, in principle, the notion of rational learning here depicted, pertains to information gain inadvertently happening in a decision process, which merely aims at maximizing the expected utility of its consequence; note, however, that utility could well be set to maximize information. In method, the term rational

learning is associated with bayesian sampling procedures, which start with a prior belief, and the beliefs converge to the true measure [15, 61, 81, 57, 48]. Thus, starting with a prior probability distribution over states of knowledge, the learner successively narrows down the true state; in a decision process, in particular, this state may represent both the state of the environment and the optimal decision strategy.

A decision process is thus an orthodox artifact for studying the interplay of purposeful action and procurement of information: interpretation of rational choice is operationalized. Foundational work appears to mostly concern consistency of decisions in function of evolving information [51].

5. Empirical learning

Roughly, empirical learning is learning without a bayesian prior; for a brief account see [4, 5, 96], and see [87] for mathematical refinement. The basic set-up is a trivial query game played by a Learner (L) and an Environment (E). Ex-ante knowledge shared by L and E consist of: (1) a measure space X with a class M of probability measures, (2) a measure metric space $(Y; d)$ and a class F of measurable functions on X with values in Y , and, (3) a class H of consistent learning algorithms for F . Here, a (deterministic) learning algorithm for F is a sequence h of computable functions $h_n : (X \times Y)^n \rightarrow F$, $n = 1, 2, \dots$; write $h_n(\xi)(x)$ as $h_n(x)$, $x \in X$. An algorithm h is F -consistent if it reproduces sequences of form $(x_i; f(y_i))$ with $f \in F$, that is, $h_n(x_i) = f(x_i)$ for $i \leq n$.

The game proceeds as follows. E picks a sequence X_1, X_2, \dots , of i.i.d. random variables in X with distribution $\mu \in M$, and a function $f \in F$, while L picks a learning algorithm $h \in H$ and two (small) positive numbers ε and δ . L now sequentially queries E for examples $(x, f(x))$, that is, L observes the sequence $(X_1, f(X_1)), (X_2, f(X_2)), \dots$ for as long as L likes. While observing the sequence, L sequentially picks 'hypotheses' h_1, h_2, \dots in the class F consistent with f on past observations, using algorithm h .

L may in principle ex-ante compute for every positive integer k , every $f \in F$, $\mu \in M$, $h \in H$, and $\varepsilon > 0$, the probability $p_k(h; f, \mu; \varepsilon)$ that for any $m \geq k$ the expectation of the error $d(h_n(X_m), f(X_m))$ of predicting $f(x_m)$ by $h_k(x_m)$, is less than ε . L may stop observing the random sequence the first time $p_k(h; f, \mu; \varepsilon)$ is larger than $1 - \delta$ uniformly in $f \in F$, $\mu \in M$, and $h \in H$, claiming to have then learned the unknown function f with accuracy ε and confidence $1 - \delta$. If it ever happens, the stopping time $n(h; F; M; \varepsilon, \delta)$ in function of ε and δ provides a bound on the learning time of any function in the class F by any algorithm in H , in a stable environment in any state in M .

Choosing H as the set of all F -consistent algorithms, and M as the set of all probability measures on X , the stopping time becomes an intrinsic measure of 'learning complexity' of class F in a stable but unknown environment. A class F is

said to be learnable if the stopping time is finite for any ε and δ ; F is learnable in polynomial time if the stopping time is bounded by a polynomial in $1/\varepsilon$ and $1/\delta$, etc. A basic theorem for the binary case $Y = \{0, 1\}$, identifies the learnable classes as those of finite Vapnik-Chervonenkis dimension; roughly, this means that there is a bound on the size of the finite subsets of X from which extrapolation in F is always possible.

Since the 1980's, empirical learning has been much investigated. The results are of fundamental interest for the theory of computational sciences. In computer science, they establish the complexity of natural classes of functions for the so-called Oracle model of computation; they have been instrumental in the study of artificial neural networks, decision trees, etc. In cognitive science, where binary functions model concepts, they express the complexity of concept classes. In econometrics, they provide bounds for sample size.

Problems. The empirical learning model is not bayesian, hence not Savage-rational, in that it assumes no prior measure on the space of probability distributions μ , but works with a 'multi-prior', $\mu \in M$. Yet, the model is pragmatic in its formulation, perhaps more so than bayesian models, a bayesian prior being often in practice unavailable. An empirical learner L faces natural decisions, which he must in practice make, Savage-irrationally.

To begin with, when facing a prospect of learning, L must decide whether to engage in the learning at all, and if so, what strategy to adopt. L must then decide when to stop: achieving small values of the parameters of accuracy and confidence, is not a goal in itself in any practical learning task, and, further, it may be unclear which values are small enough for the task at hand.

If L learns in order to act upon the acquired knowledge, a wishfully rational L may well order the streams of costs of sampling and gains from acting - the consequences of learning - but L has no prior to evaluate expectation, and may only go by the *least* expected μ -utility, over $\mu \in M$. In economic writings, the multi-prior case has been considered by several authors, starting with Ellsberg's early paper [32], most recently in [74, 38]. Under this criterion, L may decide; for example, an estimation of L 's optimal learning time in this setting was considered in [16]. In particular, L may decide if learning is feasible; note that PAC learnability is not the same as feasibility. Recall here, for example, that some natural empirically learnable classes of functions, may not be learnable computationally [6, 66]: sample processing is a cost, not unlike that of economic search [69, 100].

Generally, connections of rational choice and empirical learning largely await clarification. For example, as observed by Doyle [26], it is not clear whether PAC learning may be viewed as rational selection of concept definitions, and how this selection relates to maximum expected utility. For another example, when is a preference or a choice function effectively learnable by example? An investigation

by Kalai et al [62, 63] suggests that preferences are easy to learn, opening an argument for explaining rationality via learnability.

6. Reference-dependence and social choice

We point to two bodies of knowledge, each concerned with problems in logic akin to the learning problem.

The theory of reference-dependent preferences, see e.g. [9, 49, 70], analyzes decisions made by individuals, not with a single preference, but with several preferences \leq_α , valid according to context $\alpha \in A$. One then investigates the conditions for such decisions to be rational, in classical or weaker sense. In Giraud [49], for example, a preference \leq on pairs $(x; \alpha)$ is constructed from $\{\leq_\alpha\}$, so that $x \leq_\alpha y$ exactly when $(x, \alpha) \leq (y, \alpha)$. A learner may then answer such questions as: 'would I rather know x as person α , or know y as person β ?

The theory of social choice, see e.g. [93], analyzes collective decisions made by a group of individuals (voters) with separate preferences, whereby the collective decision should somehow maximize the desires of the individual voters. A learner might then allow all the virtual people, each of which he may potentially become due to learning, to determine his choice of learning by voting.

Visibly, these theories may formally be applied to learning problems. It is essential, however, that this is not done in isolation from statistics: a method of reconciliation of preferences may be quite bad in the worst case, but good on the average. Furthermore, the preferences \leq_α may be subject to learning. It appears that such connections are yet to be made.

7. Conclusions?

The problem of rationality of learning, which affects the learner's preference, is universal: learning often occurs by choice, while choice criteria develop through learning. Neither the logic nor the dynamics of this dependence are clear.

The problem is obviously central to any theory of management of knowledge, but we have not found it expressly posed, let alone formally treated, in subject writings. It is tacitly implied in the writings that the management of knowledge in a firm should somehow be rational. But, in what sense, exactly? How, for example, could a knowledge executive officer be held accountable, unless formal grounds for the office are in place?

It appears that formal elements for a treatment of this problem are present. We observe convergence of theories of choice in economics and procedural decision methods in artificial intelligence. We also observe formal connections of these with learning, in the bayesian and non-bayesian cases. We see reconciliation of preferences formally modelled in economics and sociology. Control of preference is formally addressed by psychology and political science. Missing perhaps, is

work towards some structure for the joint dynamics of the data observed, the theory learned, and the preference formed.

Naturally, a formal development should be paralleled by conceptual discussion and specialization to settings specific to business systems. It may be early to try envision 'the final answer'. Clearly, however, drawing on the basic sciences, knowledge management in firms will in time re-define its logical foundations, in particular to explicitly encompass such things as management of preference.

References

- [1] Adam K., *Learning While Searching for the Best Alternative*, Working paper ECO 99/4, European University Institute, Department of Economics, Florence, Italy, 1999.
- [2] Ahlswede r., Wegener i., *Search Problems*, Wiley, 1987.
- [3] Angluin D., *Queries and concept learning*, Machine Learning, vol. 2 (1988), pp. 319-342.
- [4] Angluin D., *Computational Learning Theory*, survey and selected bibliography. 24th Annual ACM STOC, (1992), pp. 351-359.
- [5] Angluin D., *A 1996 Snapshot of Computational Learning Theory*, ACM Computing Surveys, vol. 28, no. 4es (1996).
- [6] Angluin D., Kharitonov M., *When Won't Membership Queries Help?* J. Comp. Syst. Sci., vol. 50 (1995), pp. 336-355.
- [7] Anthony M., Barlett P.L., *Neural Netwbrk Learning*, Theoretical Foundations. Cambridge University Press, 1999.
- [8] Anthony M., Barlett P.L., *Computational Learning*, Cambridge University Press, 1992.
- [9] Apesteguia J., Ballester M.A., *A Theory of Reference-dependet Behavior*, (February 2004). <http://ssrn.com/abstract=575782>
- [10] Aragones E., Gilboa I., Postlewaite A., Schmeidler D., *Fact-free learning*, PIER Working Paper 03-023, Penn Institute for Economic Research, University of Pennsylvania, 2003.
- [11] Auer P., Cesa-Bianchi N., Freund Y., Schapire R.E., *The Nonstochastic Multiarmed Bandit Problem*, SIAM Journal on Computing, vol. 32, no. 1 (2002), pp. 48-77.
- [12] Barr J., Saraceno F., *A Computational Theory of the firm*, Journal of Economic Behavior and Organization, vol. 49 (2002), pp. 345-361.
- [13] Bergemann D., Hege U., *The financing of innovation: learning and stopping*, Working Paper 16, Tilburg University, Center for Economic Research, 2000.
- [14] Bertsekas D.P., *Dynamic Programming and Optimal Control*, Athena Scientific, 1995.
- [15] Blume L.E., Easley D., *Rational Expectations and Rational Learning*, Economic Theory Workshop in Honor of Roy Radner, Cornell University, 1992. <http://econwpa.wustl.edu/eprints/game/papers/9307/9307003.abs>.
- [16] Bonner R.F., Fedyszak-Koszela A., *When to Stop Learning? Bounding the stopping time in the PAC model*, Theory of Stochastic Processes, vol. 7(23), (2001), no.1-2, pp. 5-12.
- [17] Brenner T., *Modelling Learning in Economics*, Edward Elgar Publishers, 1999.
- [18] Brenner T., *Computational Techniques for Modelling Learning in Economics*. Kluwer, 1999.
- [19] Cairolì R., Dalang R.C., *Sequential Stochastic Optimization*, WileyInterscience, 1996.
- [20] Chateauneuf A., Vergnaud J-C., *Ambiguity Reduction Through new Statistical Data*, 1st Int. Symp. on Imprecise Probabilities and Their Applications, Ghent, Belgium, 1999.

- [21] Chateauneuf A., Eichberger J., Grant S., *Choice Under Uncertainty With the Best and Worst in Mind: neo-additive capacities*, Working paper, 2002.
- [22] Cohen M., Gilboa I., Jaffray J.Y., Schmeidler D., *An Experimental Study of Updating Ambiguous Beliefs*, 1st Int. Symp. on Imprecise Probabilities and Their Applications, Ghent, Belgium, 1999.
- [23] Cover T.M., Thomas J.A., *Elements of Information Theory*, Wiley, 1991.
- [24] Debreu G., *Theory of Value: an Axiomatic Analysis of Economic Equilibrium*, Wiley, 1959.
- [25] DeSarbo W.S., Fong D.K.H., Liechty J., Coupland J.C., *Evolutionary Preference/utility Functions*, Psychometrika, in press (2004).
- [26] Doyle J., *Rationality and its Roles in Reasoning*, Computational Intelligence, vol. 8, no. 2 (1992), pp. 376-409.
- [27] Doyle J., Dean T., et al. *Strategic Directions in Artificial Intelligence*, ACM Computing Surveys, vol. 28, no. 4 (1996), pp. 653-670.
- [28] Dubois D., et al. *On the use of the Discrete Sugeno Integral in Decision-making*, Int. Journal of Uncertainty, Fuzziness, and Knowledge-Based Systems, vol. 9, no. 5 (2001), pp. 539-561.
- [29] Dubois D., et al. *Qualitative Decision Theory: from Savage's Axioms to Nonmonotonic Reasoning*, Journal of the ACM, vol. 49, no. 4 (2002), pp. 455-495.
- [30] Dudley R.M., *Uniform Central Limit Theorems*, Cambridge University Press, 1999.
- [31] Eboli M., *Two Models of Information Costs Based on Computational Complexity*, Computational Economics, vol. 21 (2003), pp. 87-105.
- [32] Ellsberg D., *Risk, Ambiguity, and the Savage Axioms*, Quarterly Journal of Economics, vol. 75 (1961), pp. 645-669.
- [33] Ferguson T.S., *Optimal Stopping and Applications*, <http://www.math.ucla.edu/tom/Stopping/Contents.html>.
- [34] Foster J.E., Mitra T., *Ranking Investment Projects*, Economic Theory, vol. 22 (2003), pp. 469-494.
- [35] Freund Y., Schapire R.E., *A Decision-theoretic Generalization of on-line Learning and an Application to Boosting*, Journal of Computer and System Sciences, vol. 55, no. 1, (1997), pp. 119-139.
- [36] Fudenberg D., Levine D.K., *The Theory of Learning in Games*, MIT Press, 1998.
- [37] Fugikawa T., Oda S.H., *Search and Choice Under Uncertainty*, ESA International 2003. <http://www.peel.pitt.edu/esa2003/participants.html>.
- [38] Gajdos T., Tallon J-M., Vergnaud J-C., *Decision Making with Imprecise Probabilistic Information*, Journal of Mathematical Economics, (2004), in press.
- [39] Ganter B., Wille R., *Formal Concept Analysis: Mathematical Foundations*, Springer, 1999.
- [40] Gigerenzer G., Selten R., *Bounded Rationality: An Adaptive Toolbox*, MIT Press, 2002.
- [41] Gilboa I., Schmeidler D., *Maxmin Expected Utility with non-unique Prior*, Journal of Mathematical Economics, Vol. 18 (1993), pp. 141-153.
- [42] Gilboa I., Schmeidler D., *A Theory of Case-Based Decisions*, Cambridge University Press, 2001.
- [43] Gilboa I., Schmeidler D., *Inductive Inference: an Axiomatic Approach*, Cowles Foundation Discussion Paper 1339, Yale University, 2001.
- [44] Gilboa I., Schmeidler D., *Cognitive Foundations of Probability*, Cowles Foundation Discussion Paper 1340, Yale University, 2001.
- [45] Gilboa I., Schmeidler D., *Subjective Distribution*, Cowles Foundation Discussion Paper 1341, Yale University, 2001.
- [46] Gilboa I., Schmeidler D., *A Derivation of Expected Utility Maximization in the Context of a Game*, Cowles Foundation Discussion Paper 1342, Yale University, 2001.

- [47] Gilboa I., Postlewaite A., Schmeidler D., *Rationality of Belief. Or: why Bayesianism is Neither Necessary nor sufficient for Rationality*, Working Paper 04-011, Penn Institute for Economic Research, University of Pennsylvania, 2004.
- [48] Gilli M., *A General Approach to Rational Learning in Games*, Bulletin of Economic Research, vol. 53, no. 4 (2001), pp. 275-303.
- [49] Giraud R., *Reference-dependent Preferences: Rationality, Mechanism and Welfare Implications*, RUD (Risk, Uncertainty and Decisions) 2004 Conference, Zell Center for Risk Research, Kellogg School of Management, Northwestern University, Illinois, 2004
<http://www.kellogg.northwestern.edu/research/risk/rud/papers/giraud.pdf>
- [50] Gold E.M., *Language Identification in the Limit*, Information and Control, v. 10 (1967), pp. 447-474.
- [51] Grant S., Kajii A., Polak B., *Preference for Information and Dynamic Consistency*, Cowles Foundation Discussion Paper 1208, Yale University, 1999.
- [52] Grant S., Kajii A., Polak B., *Decomposable Choice Under Uncertainty*, Cowles Foundation Discussion Paper 1207, Yale University, 1999.
- [53] Halpern J.Y., *Conditional Plausibility Measures and Bayesian Networks*, Journal of Artificial Intelligence Research, vol. 14 (2001), pp. 359-389.
- [54] Haussler D., *Decision Theoretic Generalizations of the PAC Model for Neural Nets and Other Learning Applications*, Information and Computation, vol. 100 (1992), pp. 78-150.
- [55] Haussler D., Kearns M.J., Littlestone N., Warmuth M.K., *Equivalence of Models for Polynomial Learnability*, Information and Computation, vol. 95 (1991), pp. 129-161.
- [56] Haussler D., Littlestone N., Warmuth M.K., *Predicting (0; 1) Functions on Randomly Drawn Points*, Information and Computation, vol. 115 (1994), pp. 284-293.
- [57] Heinemann M., *Efficiency of Rational Learning Under Asymmetric Information*, Economica, vol. 70 (2003), pp. 1-15.
- [58] Jain S., Osherson D., Royer J.S., Sharma A., *Systems that learn*, 2nd edition. MIT Press, 1999.
- [59] Jeffrey R.C., *The Logic of Decision*, 2nd edition, University of Chicago Press, 1983.
- [60] Kahneman D., Tversky A., *Choices, Values, and Frames*, Cambridge University Press, 2000.
- [61] Kalai G., Lehrer E., *Weak and Strong Merging of Opinions*, Journal of Mathematical Economics, vol. 23, (1994), pp. 73-100.
- [62] Kalai G., *Learnability and Rationality of Choice*, Journal of Economic Theory, vol. 113, no. 1 (2003), pp. 104-117.
- [63] Kalai G., Rubinstein A., Spiegel R., *Rationalizing Choice Functions by Multiple Rationals*, Econometrica, vol. 70 (2002), pp. 2481-2488.
- [64] Kalai E., Solan E., *Randomization and Simplification in Dynamic Decisionmaking*, Journal of Economic Theory, vol. 111 (2003), pp. 251-264.
- [65] Kearns M.J., Vazirani U.V., *An Introduction to Computational Learning Theory*, MIT Press, 1994.
- [66] Kearns M.J., Valiant L., *Cryptographic Limitations on Learning Boolean Formulae and finite Automata*, Journal of the ACM, vol. 41, no. 1 (1994), pp. 67-95.
- [67] Khargon R., Roth D., *Learning to Reason*, Journal of the ACM, vol. 44, no. 5 (1997), pp. 697-725.
- [68] Kiefer J., *Sequential Minimax Search for a Maximum*, Proc. American Mathematical Society, vol. 4, no. 2 (1953), pp. 502-506.
- [69] Kohn M., Shavell S., *The Theory of Search*, Journal of Economic Theory, vol. 4, no. 2 (1974), pp. 593-123.
- [70] Koszegi B., Rabin M., *A Model of Reference-dependent Preferences*, Working Paper E04-337, Department of Economics, University of California, Berkeley, 2004.
<http://repositories.cdlib.org/iber/econ/E04-337>

- [71] Marchand M., Shawe-Taylor J., *The Set Covering Machine*, Journal of Machine Learning Research, vol. 3, (2002), pp. 723-746.
- [72] McAllester D.A., *Some PAC-Bayesian Theorems*, COLT 98, ACM (1998), pp. 230-234.
- [73] McClennen E.F., *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge University Press, 1990.
- [74] Mukerji S., Tallon J-M., *Ellsbergs two-color Experiment, Portfolio Inertia and Ambiguity*, Journal of Mathematical Economics, Vol. 39, Issues 3-4, June (2003), pp. 299-316.
- [75] Mukerji S., Tallon J-M., *An Overview of Economic Applications of David Schmeidler's Models of Decision Making Under Uncertainty*, Chapter 13 in: I. Gilboa (ed.). *Uncertainty in Economic Theory: A collection of essays in honor of David Schmeidler's 65th birthday*. Routledge Publishers, 2004.
- [76] Nakhaeizadeh G., Taylor C.C., *Machine Learning and Statistics: the Interface*, Wiley, 1997.
- [77] Orbay H., *Information Processing Hierarchies*, Journal of Economic Theory, vol. 105 (2002), pp. 370-407.
- [78] Ribeiro C., *Reinforcement Learning Agents*, Artificial Intelligence Review, vol. 17 (2002), pp. 223-250.
- [79] Pollard D., *Empirical Processes: Theory and Applications*, NSF-CBMS Regional Conference Series in Probability and Statistics, vol. 2, Institute of Mathematical Statistics, 1990.
- [80] Poznyak A.S., Najim K., *Learning Automata and Stochastic Optimization*, Springer, 1997.
- [81] Sandroni A., Smorodinsky R., *The Speed of Rational Learning*, International Journal of Game Theory, vol. 28 (1999), pp. 199-210.
- [82] Savage L.J., *Bounded Rationality in Macroeconomics: The Arne Ryde Memorial Lectures*, Oxford University Press, 1994.
- [83] Savage L.J., *The Foundations of Statistics*, Wiley, 1954.
- [84] Schmeidler D., *Subjective Probability and Expected Utility Without Additivity*, Econometrica, vol. 57, No. 3, (1989), pp. 571-587.
- [85] Scholkopf B., Burges C.J.C., Smola A.J., *Advances in Kernel Methods: Support Vector Learning*, MIT Press, 1998.
- [86] Shoham Y., *The Open Scientific Borders of AI, and the case of Economics*, ACM Computing Surveys, vol. 28, no. 4es (1996).
- [87] Cucker F., Smale S., *On the Mathematical Foundations of Learning*, Bulletin (New Series) of the American Mathematical Society, Vol. 39, No. 1 (2001), pp. 1-49.
- [88] Sobel J., *Economists' Models of Learning*, Journal of Economic Theory, vol. 94, (2000), pp. 241-261.
- [89] Sutton R., Barto A., *Reinforcement Learning*, MIT Press, 1999.
- [90] Szepesvari C., Littman M.L., *A Unified Analysis of Value-Function-Based Reinforcement-Learning Algorithms*, Neural Computation, vol. 8, (1999), pp. 217-259.
- [91] Tallon J-M., Vergnaud J-C., *Choice Axioms for Positive Value of Information*, Working Paper, 2002. <http://eurequa.univ-paris1.fr/membres/tallon/tallon.htm>
- [92] Tallon J-M., Vergnaud J-C., *Beliefs and Dynamic Consistency*, Working Paper, 2003. <http://eurequa.univ-paris1.fr/membres/tallon/tallon.htm>.
- [93] Taylor A., *Mathematics and Politics: Strategy, Voting, Power and Proof*, Springer-Verlag, 1995.
- [94] Valiant L.G., *A Theory of the Learnable*, Comm. ACM, vol. 27, no. 11, (1984), pp. 1134-114.
- [95] Vapnik V.N., *Statistical Learning Theory*, Wiley Interscience, 1998.
- [96] Vapnik V.N., *The Nature of Statistical Learning Theory*, Springer, 1995.
- [97] Vidyasagar M., *A Theory of Learning and Generalization*, Springer, 1997.
- [98] Wakai K., *Aggregation of Agents with Multiple Priors and Homogeneous Equilibrium Behavior*, Working Paper, Yale University, 2001.
- [99] Wang T., *A Class of Multi-prior Preferences*, Working Paper, University of British Columbia, 2003.

- [100] Weitzman M., *Optimal Search for the Best Alternative*, *Econometrica*, vol. 47, no. 3, (1979), pp. 641-654.
- [101] Wellman M., *The Economic Approach to Artificial Intelligence*, *ACM Computing Surveys*, vol. 27, no. 3 (1995), pp. 360-362.
- [102] Wellman M., *Rationality in Decision Machines*, AAAI Fall Symposium on Rational Agency, November 1995.
- [103] Whittle P., *Optimization Over Time*, Wiley, 1982.
- [104] Zilberstein S., *Resource-bounded Reasoning in Intelligent Systems*, *ACM Computing Surveys*, vol. 28, no. 4es, 1996.

NADAWANIE SENSU ZARZĄDZANIU WIEDZĄ: RACJONALNY WYBÓR, UCZENIE I WSPÓLZALEŻNOŚCI

Streszczenie

W artykule rozważany jest problem uczenia, który obejmuje zmiany systemu wiedzy firmy wpływające na modyfikację działań tej firmy a zarazem jej preferencje. Podstawy tych zmian mogą zostać rozpoznane dopiero po uczeniu. Co więcej, podstaw dla porównania strategii uczenia może po prostu nie być, o ile nie została zapewniona jawna zgodność preferencji wynikających z różnych strategii. Procedury odwołujące się do kryterium wartości pieniężnej mogą być bezpodstawne. Logiczne podstawy zarządzania wiedzą w przedsiębiorstwie stają w ten sposób pod znakiem zapytania. Formalnego zbadania wymagają problemy ewolucji preferencji wynikających z uczenia i racjonalność takiego uczenia. Jako pierwszy krok w tym kierunku zaproponowano krótki przegląd obecnych teorii uczenia i racjonalnego wyboru. Centralnym problemem teorii racjonalnego wyboru i schematów decyzyjnych sztucznej inteligencji jest zagadnienie efektywnego i dynamicznego wyboru akcji przez „ekonomicznego” agenta, który „uczy się” środowiska i konsekwencji swoich działań – przy czym jest on standardowo konfigurowany niekonieczne pod kątem uczenia. Przy takich założeniach i z uwzględnieniem np. formalizmu Bayesa (który nie zawsze jest brany pod uwagę) - racjonalne procesy uczenia są w zasadzie dobrze rozpoznane. Uczenie bez uwarunkowań założeń Bayesa jest badane w modelach empirycznych; ich racjonalne podstawy nie są jednak jasne. Przy założeniu zmiennych preferencji agenta w funkcji uczenia – zgodność potencjalnych preferencji powinna być formalnie zapewniona z racjonalnego punktu widzenia – na wzór teorii wyborów społecznych i względnych preferencji (co nie jest jeszcze rozważane w uczeniu). Podsumowując, wydaje się, że współczesna teoria traktuje problem uczenia w sposób bardzo otwarty natomiast rozwijane są narzędzia pobudzające do badań omawianego zagadnienia.