Philippe Mongin - Bernard Walliser

INFINITE REGRESSIONS IN THE OPTIMIZING THEORY OF DECISION

It is a crucial assumption of economic theory that individuals make decisions by maximizing an objective (i.e. preference-representing) function under specified constraints. This decision-theoretic model has been subjected to strong empirical criticism on the grounds that it lacks psychological realism and does not fit well with known data on individual decision-making (e.g. Lester, 1945, Simon, 1959 and 1979). A related, methodological rather than empirical, strand of criticism emphasizes the barrenness of an approach that discards the cognitive and deliberative import of decisions as scientifically irrelevant and can more or less trivially be reconciled with any known observable behaviour (e.g. Simon, 1976, Latsis, 1976). Finally, there is a logical critique of optimization which is not as well developped in the technical literature. This paper aims at making it more precise.

If it is granted that optimizing always involves some computations and the latter are costly, it becomes obvious that decisions made in this way may not be optimal after all, since computation costs are usually not taken into account. Of course, a more sophisticated model could make computation costs part of the relevant exogeneous variables, but metaoptimization is again costly, and a further problem is created. In order to avoid logical pitfalls, optimizing decision theorists should demonstrate that there are optimizing decisions that take their *own* costs into account, or alternatively that the infinite regress which shows up in the previous argument is not a "vicious" one. The latter route is a promising one, since infinite regresses can easily be translated into the formalism of infinite sequences and a "non-vicious" regression can simply be defined as one which converges in an appropriate mathematical sense.

There are tentative models of how actual agents relate to the first two logical levels at which optimizing behaviour can be defined (Beach and Mitchell, 1979, Shugan,

B. R. Munier (ed.), Risk, Decision and Rationality, 435–457.

1980). There are also more formal, but *ad hoc* models, such as Baumol and Quandt's (1964) and, more recently, Göttinger's (1982), which convey some results and insights, but again virtually restrict themselves to the first two relevant logical levels. Winter has done much in terms of eliciting a distinctive logical objection against the optimizing theory of decision. His working emphasizes self-reference rather than infinite regress : "the optimization whose scope covers all considerations including its own costs... sounds like it may involve the logical difficulties of self-reference. To demonstrate this - to prove logically that there is no superoptimization - would require the development of a formal framework within which the statement could be interpreted. That would be an interesting project." (1975, p. 83).

Infinite regresses do not appeal to commonsense reasoning and, from Aristotle onwards, philosophers have generally believed that no sound argument should ever involve an infinite regress (recall the celebrated sentence in *Physics*, VIII, 5 : ἀνάγκη στῆναι καὶ μὴ εἰς ἄπειρον ἰέναι). Still, a well-known example of an infinite regress which turned out to be inoccuous is Zeno's paradox. The arrow, it is said, would never reach the wall ; but it does, and the mathematical model which makes the "paradox" evaporate is that of a convergent geometric series. This is a useful example, since it suggests how to handle infinite regresses formally. Contemporary economics and social sciences are replete with infinite regresses which are referred to as "paradoxes"[1] in much the same way as Zeno's was before a solution was offered.

It is worth emphasizing that the logical problem which has just been raised in the context of optimization is by no means specific to it, but could emerge from any theory of rational action whatsoever. Aristotle pointed out that any deliberation could involve one in an infinite regression (*Nicomachean Ethics*, III, 5). This may be understood as follows : for the theorist to start regressing from one decision to another, it is enough for him to be committed to this rather weak claim -- agents do or should apply to the higher decision level the very same choice mechanism that they applied to the lower level. From the same basic argument, along with the implicit premisse that *any* infinite regress is a logically vicious one, Ryle (1949) inferred

that philosophy should purge itself ofthe "intellectualist" tradition of analyzing action as relying on some prior knowledge [2] . As far as contemporary decision theory is involved, the question arises of whether or not the standard optimizing model of decision is more prone to "vicious" infinite regresses than its heterodox rivals, such as Simon's "satisficing" model.

The aim of this paper is to arrive at a formalism of the infinite regress of optimization and a definite interpretation of what it is for a decision theory to *supersede* an infinite regress. We shall also discuss and formalize higher-order optimization both as a preliminary step and for its own sake. Thus the strategy of the paper is as follows. After reviewing some useful facts about the standard optimizing model as well as the less commonly used notion of "decision procedure", we discuss metaoptimization, i.e. the optimal trade-off between low-level utilities and decision costs. This framework is further extended to higher logical levels, and applied to two particular models of infinite regress, one of which relies on the full-fledged version of the optimizing model, the other is set up in the more tractable language of decision rules. Finally, it is shown that optimizing agents cannot avoid arbitrarily choosing the logical level at which they operate, but the observer-theorist could determine whether or not there is a finite level at which the agent *should have* made his optimal decisions, and whether or not the infinite regress of optimizing decisions is superseded, i.e. convergent in the appropriate sense.

1. SOME USEFUL FACTS ABOUT OPTIMIZATION AND DECISION PROCEDURES

Consider a typical optimization problem: when faced with a state of nature $e \in E$ and a feasible set $X(e)$ which depends on e, the agent chooses his command variable x by maximizing his objective function $U(x, e)$ under the constraint $x \in X(e)$. U, X and e will be referred to as the *determining factors* of his choice. At this stage, they are assumed to be known to the agent and no computation cost is yet taken into account. If there is a unique solution for each e, the

agent's maximizing behaviour can be expressed by a function, referred to as an *optimal decision rule* R^*:

$$\underset{x}{\text{Max}}\ U(x, e) \text{ s.t. } x \in X(e),\ e \in E \rightarrow x = R^*(e)$$

The above optimization problem may be understood as referring to either a fixed or a variable e. Accordingly, we shall henceforth distinguish *local* from *global* optimization problems.

Among numerous applications of this well-known decision-theoretic framework, two elementary examples deserve mentions. In the standard consumer theory, x is to be interpreted as a basket of goods, e is (p, r) with p the price vector and r the individual's income, R^* is, of course, the individual's demand function :

$$\underset{x}{\text{Max}}\ U(x) \text{ s.t. } p \cdot x \leq r \rightarrow x = D(p, r)$$

In the context of monopoly, x is to be interpreted as the production level, e is a vector of exogeneous variables which influence both the inverse demand function F(x, e) and the firm's cost function C(x, e) and optimization leads again to a reaction function H :

$$\underset{x}{\text{Max}}\ F(x, e) \cdot x - C(x, e) \rightarrow x = H(e)$$

Global optimization problems admit of a specific, sometimes more convenient, rephrasing. Define a new utility function V on the set $\mathcal{R}$ of all feasible decision rules R ($\mathcal{R} \subseteq X^E$) by giving each utility U(x, e) a positive weight $\pi(e)$ and maximise on $\mathcal{R}$:

$$\underset{R \in \mathcal{R}}{\text{Max}}\ V(R) = \int_E U(R(e), e)\, d\pi(e)$$

Suppose $\mathcal{R} = X^E$; then, this maximization programme leads to the agent's optimal decision rule R^* on $\mathcal{R}$, no matter how the $\pi(e)$ are chosen, provided that *all* of them are positive

(i.e. π is a *strictly* positive measure on E). If $\mathcal{R} \neq X^E$, the choice of π does influence the optimal R^*. When π is taken to be a probability measure on E and the feasible set $\mathcal{R}$ is restricted to constant rules only, the model collapses into the usual expected utility criterion.

In actual terms, global optimization problems may be solved either by using a general (formal) algorithm which generates R^* directly or by using a specific, possibly numerical algorithm that computes R^* pointwise. As far as local optimization problems are concerned, they can be solved either by applying R^* as generated by the formal algorithm to the given $\bar{e}$ or by computing only $R^*(\bar{e})$ from a specific algorithm. In the case where X and R are finite, there is always one algorithm available for both global and local problems. It consists in computing and successively comparing the U(x, e) or the V(R) for each relevant argument of the U or V function. This is the so-called *trivial algorithm*, to be used later in this paper.

We are also in need of a formal notion of *decision procedure*. The latter turns out to be rather awkward to make precise. Let us here define it as a pair (T, D) where T stands for the procedure *type* and D is a relevant set of *determining factors*. The type variable is assumed to be chosen in a predetermined list - e.g. optimizing, satisficing and random choice. This is a fairly *ad hoc* characterization, but it is difficult to improve on in the absence of a generally agreed upon theoretical framework capable of defining procedure types independently of particular examples [3]. The determining factors are, so to speak, the agent's background knowledge ; this notion is closely related to H. Simon's concept of "decision premisses". It is, again, best understood by means of examples.

The determining factors that are consistent with the optimizing type T^o are (U, x, e) and either (U,x) or (V, $\mathcal{R}$), depending on whether the problem at hand is a local or a global optimization problem. Note that E is not made explicit in this formalism, since the feasible set X is taken to be a function on E. As far as the satisficing type T^s is concerned, we shall assume that individuals make their choice by sequentially searching in an ordered action set

$\hat{X}(e)$, stopping as soon as a $x \in \hat{X}(e)$ is seen to ensure more than a predetermined amount of (usually vector-valued) utility, that is, satisfies $U(x, e) \geqslant \underline{U}(e)$. In this case, the relevant determining factors are $\hat{X}$, U, and $\underline{U}$ and the fixed environment e, if the problem to be solved is not a global, but a local one. Concerning the random choice type T^r, the determining factors clearly are X, a probability density μ on X and possibly environment e. Once these definitions have been introduced, there is no further difficulty in defining the *procedure set* $\mathcal{P}$ as well as functions $f : \mathcal{P} \to \mathbb{R}$, which can be denoted by f(P) or f(T, D) indifferently. It should be clear from the three examples above that determining factors are not independent of the type variable [4].

2. METAOPTIMIZATION

When solving a decision problem with procedure P = (T, D), individuals incur two kinds of costs. There is a discovery cost C'(T, D), i.e. one of finding out what the relevant type and determining factors are, and a computation cost, C" (T, D), i.e. one of finding out what the solution is once T and D are given. For instance, the costs of optimization are $C'(T^o,U,X) + C"(T^o,U,X)$ or $C'(T^o,U,X,e) + C"(T^o,U,X,e)$, depending on whether the problem at hand is a global or a local one. The C', C" functions are assumed to be real valued and normalized in such a way as to be commensurable with the U and V functions.

Computation costs can be further analyzed into subcosts, each of which is related to one formal or numerical operation performed by the individual. Assuming that determining factors are known to the agent, i.e. that discovery costs have already been incurred, there is a cost $\Gamma'(T, D)$ of finding the solution and, for any function $F : y \mapsto F(y)$, a cost $\Gamma"(F, y)$ of applying F to y as its argument. Additionally, we may define $\Gamma'''(\hat{S})$ as the cost of selecting the maximal elements of an ordered set $\hat{S}$. Of course, all these costs are not to be incurred

simultaneously. This, and the definitions above, can be clarified in terms of a local optimization problem. In this example, assuming again that discovery costs $C'(T^o, U, X, e)$ have already been incurred, the agent will have to pay $\Gamma"(T^o, U, X, e)$ only if he directly uses a specific algorithm, or $\Gamma"(T^o, V, \mathcal{R}) + \Gamma"'(R^*, e)$ (or $\Gamma'(T^o, U, X) + \Gamma"'(R^*, e)$) if he goes the roundabout way of using a general algorithm first. To take a particular case, his algorithm might be the trivial one, i.e. it might be the case that

$$\Gamma"(T^o, U, X, e) = \sum_{x \in X(e)} \Gamma"'(U, (x, e)) + \Gamma''''(\{U(x, e), x \in X(e)\})$$

Repeated decisions can lower unit costs substantially if the agent retains some of his previously gathered information and manages to use it again. For instance, if U and X are fixed, and e is variable, the cost of solving the local optimization problem for each e reduces to $\Gamma"'(R^*, e)$, providing that R^* has not been forgotten by the agent. If not only U and X, but also e are fixed, and repeated decisions are made, the total decision cost falls to zero, providing of course that the agent remembers the solution. In either case, the once-and-for-all incurred costs may or may not be divided, depending on the problem at hand, by the number of repetitions, to become part of a (decreasing) decision unit cost; they are very much like the fixed costs of the neo-classical theory of the firm. Still, this line of reasoning cannot be pursued very far if assumptions about the agent's memory (and possibly about his storage costs) are not explicitely made.

In view of the various decision costs he would incur, the agent may switch from one optimizing procedure to a less costly one, which might be an optimizing procedure with "simpler" determining factors, or he may choose a non-optimizing procedure. A metadecision of this sort relies on weighing the loss in primary utility U or V against the gain made on the C' or C" score. Call $\tilde{x}(P)$ the solution of the local optimization problem for procedure P and C(P) =

$= C'(T^o, D) + C''(T^o, D)$ the total cost of using P. The optimal trade-off is given by the programme :

$$\underset{P \in \mathcal{P}}{\text{Max}}\ U(\tilde{x}(P)) - C(P)$$

the solution of which, P^*, depends on T^o, U, C and $\mathcal{P}$. In case of a global optimization problem, denoting by $\tilde{R}(P)$ the decision rule implied by P, the optimal solution P^* can be seen to result from the following programme:

$$\underset{P \in \mathcal{P}}{\text{Max}}\ V(\tilde{R}(P)) - C(P)$$

We shall later on use the notation $R^* = \tilde{R}\ (P^*)$.

The above framework applies to the following example from Göttinger (1982). A consumer is assumed to know his utility function U which is C^∞ and to face a local maximization problem (the procedure here is taken to be (T^o, U, X, p, r)). Since the consumer finds it difficult to use U in order to find out the optimal solution, he considers resorting to "simplified" functions U^i. The latter are given to him by the i-order Taylor expansion of U around (x_0, p, r), where $x_0 \in X(p, r)$ is fixed. Not surprisingly, Göttinger assumes that the computation cost $C(i) = C(T^o, U^i, X, p, r)$ of the value of the demand function increases with i for x close enough to x_0. The agent's final decision is then seen to result from the following metaoptimization programme :

$$\underset{i}{\text{Max}}\ U(i) = U\Big(D^i\ (p, r - C(i))\Big),$$

where D^i is the demand function derived from using U^i as a proxy, and the computation cost C(i) has been subtracted from the exogeneous income r. Nothing in the model as it was stated ensures that the metaoptimization programme has a finite solution, though this is clearly the case when $C(i+1)-C(i)$ is constant or increasing and $U(i+1) - U(i)$ is decreasing with i.

Baumol and Quandt (1964) study the example of a monopolist who does not know his total cost and demand functions except for two points drawn at random on each curve. (The observer of course knows what the curves are ; they are assumed to belong to one of six classes of total cost - demand combinations). There are assumed to be four procedures P_i by means of which the agent can choose his price p_i : he might fix it at a constant value, draw it at random, optimize it under the assumption of linear total cost and demand functions, or finally optimize it under the assumption of a quadratic cost and a linear demand, the known points being each time on the assumed curves. Baumol and Quandt are satisfied with a comparison, for various e, of the results accruing from each P_i. If m is the random drawing of two price values and π the monopolist's profit as a function of p_i and e, the metaoptimization programme for a given e is here expressed by :

$$\underset{i}{\text{Max}}\ E_m\ \{\pi(p_i\ (m), e) |\ e\} - C(P_i)$$

where $C(P_i)$ stands for the cost of using procedure P_i - a value which is generally not independent of e. Il may be assumed, as in Göttinger's model, that $C(P_i)$ is increasing in i for a fixed e. Still, no special assumption is needed here to derive a (meta)optimal solution, since the procedure set $\mathcal{P}$ is assumed to be finite.

3. HIGHER-ORDER OPTIMIZATION

The weakness in the models discussed in section 2 is that they arbitrarily focus on metaoptimization, which involves decision costs of its own and should be superseded in exactly the same way as was standard optimization. We may here clarify the costs involved in metaoptimization, costs that remained implicit in the models discussed above. From this point on, logical levels are to be distinguished from each other. Then, we should rewrite the various costs discussed in the last section as C_1', C_1'', C_1 and the

procedures as $P_1 = (T_1, D_1) \in \mathcal{P}_1$. The costs incurred by a metaoptimizer can be taken to depend on $(U, X, C_1, \mathcal{P}_1)$ or $(V, C_1, \mathcal{P}_1)$ for a global problem and on $(U, X, e, C_1, \mathcal{P}_1)$ for a local one. This is in accordance with the examples as well as with the abstract metaoptimization formula in section 2. The following discussion of costs is limited to the notationally convenient case of *global* metaoptimization. As was done in the C_1 case, there is a natural distinction between discovery and computation components, respectively C_2' and C_2''.

First of all, there is a $C_2'(T_2^o, V, C_1, \mathcal{P}_1)$ to be paid in order to discover the determining factors of the metaoptimization problem. If the agent has *already* made an optimal choice, the "lower-level" piece of information V can safely be assumed to be known to him when he turns into a metaoptimizer, so that it costs him nothing. Similarly, the value of C_1 for the optimizing procedure $\hat{P}$ *was incurred* by the agent and can, again, be assumed to be *known* to him; thus, he would have to pay only for discovering the values of C_1 on the remaining domain $\mathcal{P}_1 \setminus \{\hat{P}\}$. It would be an interesting project to investigate how individuals in effect come to know a procedure set $\mathcal{P}_1$. All metaoptimizing models we are aware of rely on a completely *ad hoc* characterization of available procedures - witness Baumol and Quandt's list or Göttinger's Taylor expansion assumption in section 2. The only reasonable hypothesis is that actual agents use an extremely simple $\mathcal{P}_1$ which either is finite or is generated from a simple recurrence rule, and includes some precise, but computationally complex typical rules (usually optimizing one) along with loose, but computationally handy typical rules (possibly satisficing ones). This bias towards simple procedure sets can of course be rationalized in terms of metametaoptimization as is shown later in this section.

Computation costs $C_2''(T_2^o, V, C_1, \mathcal{P}_1)$ relate to level 1 as well as to level 2 consumptions. "On the top" is $\Gamma_2'(T_2^o, V, C_1, \mathcal{P}_1)$, which is the cost of discovering the optimal procedure $P_1^* \in \mathcal{P}_1$ (the latter is of course likely to

be a non-optimizing one). "At the bottom" is the cost of finding the (meta) optimal rule which is associated with P_1^*. That is, once $P_1^* = (T_1^*, D_1^*)$ is known, the agent must compute $\tilde{R}^*$ which is associated with the now endogeneously defined type and determinant and has cost $\Gamma_1'(T_1^*, D_1^*)$. When all of the computations are done using the trivial algorithm only, the total computation cost $\Gamma_2'(T_2^o, V, C_1, \mathcal{P}_1) + \Gamma_1'(T_1^*, D_1^*)$ is:

$$\underbrace{\sum_{P_1 \in \mathcal{P}_1}{}^{\cdot} \Gamma_1'(P_1)}_{\text{Cost of computing the rules associated with the various procedures}} + \underbrace{\sum_{P_1 \in \mathcal{P}_1}{}^{\cdot} \Gamma''\left(V, \tilde{R}(P_1)\right)}_{\text{Cost of computing the utility of the rules associated with the various procedures}}$$

$$+ \underbrace{\sum_{P_1 \in \mathcal{P}_1}{}^{\cdot} \Gamma''(C_1, P_1)}_{\text{Cost of computing the total cost associated with the various procedures}} + \underbrace{\Gamma'''\left(\left\{V\left(\tilde{R}(P_1) - C_1(P_1), P_1 \in \mathcal{P}_1\right)\right\}\right)}_{\text{Cost of comparing the net utilities of the various procedures}}$$

Second-order procedures $P_2 = (T_2, D_2)$ may be defined in a way similar to P_1; for instance, a global metaoptimization problem is described by $(T_2^o, V, C_1, \mathcal{P}_1)$. The optimizing answer to the question of where the C_1 and $\mathcal{P}_1$ components come from is roughly the following: agents would endogeneously determine the $\mathcal{P}_1$ and C_1 which make it possible to solve the lower-order (metaoptimization) problem as economically as possible. Individuals are required to go through the metametaoptimization programme :

$$\operatorname*{Max}_{P_2 \in \mathcal{P}_2} V\left(\tilde{R}\left(\tilde{P}_1(P_2)\right)\right) - C_2(P_2)$$

Here $\tilde{P}_1(P_2)$ stands for the level one procedure which results

from using P_2 as a level 2 procedure and $\tilde{R}(\tilde{P}_1(P_2))$ for the implied decision rule in $\mathcal{R}$. This obviously extends to any higher level n. Thus define a (n+1)-order maximization programme as :

$$\underset{P_n \in \mathcal{P}_n}{\text{Max}} \; V\left(\tilde{R}\left(\tilde{P}_1\left(\ldots\left(\tilde{P}_{n-1}(P_n)\right)\right)\right)\right) - C_n(P_n)$$

where the symbols have an obvious interpretation. Denote by P_n^* the solution of this programme and let $R^{n+1*} = \tilde{R}\left(\tilde{P}_1\left(\ldots\left(\tilde{P}_{n-1}(P_n^*)\right)\right)\right)$, i.e. the rule in $\mathcal{R}$ induced by using the optimal n-level procedure. Call $\hat{P}_{n+1}$ the one-step ahead optimizing procedure $(T^o, V, C_n, \mathcal{P}_n)$ such that $P_n^* = \tilde{P}_n(\hat{P}_{n+1})$.

4. TWO SPECIFIC FORMALISMS

As noted earlier, the crucial difficulty in modeling metadecision problems is that of defining a meaningful set of procedures $\mathcal{P}_1$ among which agents are to choose. This section discusses two possible solutions to this problem. The first has already been dealt with briefly ; it consists in restricting $\mathcal{P}_1$ to the optimizing type and considering an *ad hoc* set of simplified determining factors. We shall here elaborate upon this approach and then move to the more radical method of discussing metaoptimization in terms of decision rules only.

Restricting oneself to the T^o type has greater potential generality than at first seems to be the case, since various non-optimizing procedures can formally be reduced to optimizing ones. For instance, following a predetermined rule $\bar{R}$ is trivially equivalent to solving the programme : for each $e \in E$

$$\underset{x \in X(e)}{\text{Min}} \; \| x - \bar{R}(e) \| \;,$$

provided, of course, a norm has been defined on the action space. Similarly, satisficing models can be rewritten as *special* optimizing models, which does not imply, of course, that *standard* optimization includes satisficing as a particular case. For given e, call $\hat{X}^n(e)$ the sequence of the first n actions of $\hat{X}(e)$ which have been considered by the agent and $W(x, e)$ the characteristic function of $\{U(x, e) \geqslant \underline{U}(e)\}$. Then, the first satisfactory action for a given e is the solution of the programme :

$$\underset{n,x}{\text{Max}} \; \frac{W(x, e)}{n} \quad \text{s.t.} \quad n \in \mathbb{N}^*, \; x \in \hat{X}^n(e)$$

(if x_N is the satisfactory solution, $\frac{W(x, e)}{n} = 0$ for $n < N$, and $= \frac{1}{n}$ for $n \geqslant N$). Note further that restriction to the optimizing type makes the list of determining factors a redundant one. To dispense with the constraint element X_1, it is enough to convert $U_1(x, e)$ into $U_1'(x, e)$ with $U_1'(x, e) = -\infty$ if $x \notin X_1(e)$, or to use characteristic functions again.

Thus, after redefining utilities in the proper way, the level (n+1) optimization programme of the cost section now reads:

$$\underset{V_n \in \mathcal{V}}{\text{Max}} \; V\left(R^*\left(V_1^*\left(\ldots\left(V_{n-1}^*(V_n)\right)\right)\right)\right) - C_n(V_n)$$

for a global problem. To illustrate this method of analyzing higher-order optimization in terms of utilities only, consider the case in which a level 2 optimizer does not need to take computation costs into account (e.g. they are similar for all utilities), but is only concerned with the discovery costs of V_1. This case is examplified by Stigler's well-known consumer search model (1961). According to this model, the consumer is assumed to know the set of shops X(e)

in which he could make a purchase, but not the price charged in each shop, so that he ignores his "true" utility function U(x, e). Nonetheless, he does have information about utility values, since it is also assumed that he knows the cost of visiting one more shop and the true distribution of prices. From the special vantage point of this paper, this model is a specialized metaoptimization case in a sense diametrically opposed to Göttinger's model (the latter has no discovery, but computation costs only).

A more radical simplification of the higher-order optimization problem can be found in Mongin (1986). At every logical level, procedures are replaced by the more tractable decision rules implied by them ; optimization and particular optimizing procedures are specified by formal conditions placed on the various decision rule sets [5]. Technically, a level n decision rule is defined as a member R_n of $\mathcal{R}_n = X^{E^n}$, $R_1 \in X^E$ being a rule in the usual sense. Applying a n-rule to a given e engenders a (n-1)-rule, and so on, finally resulting in $x(e) = R_n(e)\ldots(e)$ [n times] in the action space. A level n-procedure is characterized by two data : on the one hand, its associated top level decision rule R_n, and on the other hand, a subset A_{n-1} of $\mathcal{R}_{n-1}$, which is the maximum range onto which R_n could be mapped and that should be interpreted as the set of lower-order rules of which the agent is aware. An n-procedure is associated with the (n-1)-procedure implied by it through certain plausible set-theoretic conditions.

Optimizing procedures are specified by the requirement that the agent explore *all* of the a priori available rules, i.e. level n-optimizing for a local problem is described by a sequence :

$$\langle R_n, A_{n-1}\rangle, \langle R_n(e), A_{n-2}\rangle, \ldots, \langle R_n(e)\ldots(e), A_0\rangle$$

with $A_j = A^E_{j-1}$, $1 \leqslant j \leqslant n$, and $A_0 \subset X$. This is a rather strong assumption, since it implies that no other algorithm than the trivial one be used by the agent. There are as many particular optimizing n-procedures as there are subsets of X; using an $A_{n-1} \subset \mathcal{R}_{n-1}$ is tantamount to optimizing on a "simplified" procedure set $\mathcal{P}_{n-1}$ in the terminology of the previous sections. Discovery and computation costs are

redefined so as to be compatible with this framework. There is a cost associated with finding $\langle R_n, A_{n-1}\rangle$ and one of computing the lower-order rules for the given e. The former has a discovery as well as computation element [6], whereas the latter is a pure computation cost $\Gamma^m(R_n, e)$.

5. THE AGENT'S ATTITUDE TOWARDS THE INFINITE REGRESS

Infinite regresses can be studied from two distinct vantage points, the agent's and the observer-theorist's. Let us here discuss the former. Predictions on the agent's behaviour can be made by estimating the variation of his net utility when two optimizing procedures $\hat{P}_n$ and $\hat{P}_{n-1}$ are applied in succession:

Logical level at which optimization takes place \ of procedures	n+1	n	n-1	0
n+1	$\hat{P}_{n+1}$	P^*_n	$\tilde{P}_{n-1}(P^*_n)$	R^{n+1*}
n		$\hat{P}_n$	P^*_{n-1}	R^{n*}

As a particular example consider the first three levels :

Logical level at which optimization takes place \ of procedures	3	2	1	0
3	$\hat{P}_3$	P^*_2	$\tilde{P}_1(P^*_2)$	R^{3*}
2		$\hat{P}_2$	P^*_1	R^{2*}
1			$\hat{P}_1$	R^{1*}

There are several ways of measuring the agent's net utility variation, depending on what assumptions are made

about his problem-situation. First, consider the following simple case : it is assumed that the discovery costs of any level are eliminated *and* that the agent knows the optimal procedures P^*_{n-1} and P^*_n without charge. Thus, the total cost of n and (n+1)-optimizing reduces to a computation cost depending on P^*_{n-1} and P^*_n only, i.e. $C_n(\hat{P}_n) = C^{\prime\prime}_{n-1}\left(P^*_{n-1}\right)$ and $C_{n+1}\left(\hat{P}_{n-1}\right) = C^{\prime\prime}_n\left(P^*_n\right)$. The relevant variational expression is therefore :

$$\Delta_n = V(R^{n+1*}) - V(R^{n*}) - C^{\prime\prime}_n\left(P^*_n\right) + C^{\prime\prime}_{n-1}\left(P^*_{n-1}\right)$$

$$= V(R^{n+1*}) - V(R^{n*}) + C^{\prime\prime}_{n-1}\left(P^*_{n-1}\right) - C^{\prime\prime}_{n-1}\left(\tilde{P}_{n-1}\left(P^*_n\right)\right)$$

$$+ C^{\prime\prime}_{n-1}\left(\tilde{P}_{n-1}\left(P^*_n\right)\right) - C^{\prime\prime}_n\left(P^*_n\right)$$

This expression can be seen to be nonpositive given a plausible assumption, since the first line is obviously nonpositive because of the optimality of P^*_{n-1}, while the second is nonpositive as well, if it is granted that computations on the lower-level procedure $\tilde{P}_{n-1}(P^*_n)$ are part of the computations made on the higher-level P^*_n. Hence the conclusion in this particular case : the agent should (n-1)-level optimize and not move up. For instance, when no discovery cost is involved, an agent who knows the (level 1) optimizing rule should not use the metaoptimizing rule even if it is given to him free of charge.

Second, consider the slightly more general case in which discovery costs are again assumed away *and* the agent is given P^*_{n-1}, but not P^*_n free of charge. Then, his incremental utility is :

$$\Delta'_n = V(R^{n+1*}) - V(R^{n*}) - C^{\prime\prime}_{n+1}\left(\hat{P}_{n+1}\right) + C^{\prime\prime}_{n-1}\left(P^*_{n-1}\right)$$

$$= V(R^{n+1*}) - V(R^{n*}) - C^{\prime\prime}_n\left(P^*_n\right) + C^{\prime\prime}_{n-1}\left(P^*_{n-1}\right)$$

$$+ C^{\prime\prime}_n\left(P^*_n\right) - C^{\prime\prime}_{n+1}\left(\hat{P}_{n+1}\right)$$

The first line in the last expression of Δ'_n is equal to Δ_n and therefore nonpositive, as is the second line, under the "repeated cost" assumption of the last paragraph. Hence, the agent should again stop at P^*_{n-1}. This means that in the absence of discovery costs, individuals who have already (level 1) optimized should not undertake the computation of the metaoptimizing rule. In order to reach this conclusion himself, the agent does not need to know any actual figures; it is enough for him to be able to do the abstract reasoning above.

Third, suppose that no free information is available and no previous optimization has occurred. The general problem facing the agent in this case is to find the level n at which he should operate. Still assuming away discovery costs, incremental utility should be expressed as :

$$\Delta^{\prime\prime}_n = V(R^{n+1*}) - V(R^{n*}) - C^{\prime\prime}_{n+1}\left(\hat{P}_{n+1}\right) + C^{\prime\prime}_n\left(\hat{P}_n\right)$$

$$= V(R^{n+1*}) - V(R^{n*}) + C^{\prime\prime}_n\left(P^*_n\right) + C^{\prime\prime}_{n-1}\left(P^*_{n-1}\right)$$

$$+ \left[C^{\prime\prime}_n\left(\hat{P}_n\right) - C^{\prime\prime}_{n-1}\left(P^*_{n-1}\right)\right] - \left[C^{\prime\prime}_{n+1}\left(\hat{P}_{n+1}\right) - C^{\prime\prime}_n\left(P^*_n\right)\right]$$

The first line here is again equal to Δ_n, therefore nonpositive, but the second line has no determinate sign, since it is a difference of two positive quantities (as a result of the "repeated cost" assumption). Intuitively, $\Delta^{\prime\prime}_n \gtrless 0$ depending on whether the certain loss in utility and lower-level computation costs (first line) is or is not balanced by a gain in the upper-level computation cost (second line). Thus, if the utility function is very flat and computations have a uniformly higher cost in n+1 than in n, it is not commendable to move up from n to n+1.

This conclusion may be sharpened within the context of particular models, but the crucial point is that no general result on the sign of Δ_n^* is available here. As a result, the optimal level n^* at which the agent should operate may be finite or infinite [7]. (A sufficient condition for there be a finite n^* it that Δ_n^* be positive and negative in succession.) The problem of indeterminacy has quite different consequences for the observer-theorist on the one hand, and on the other hand for the agent in the model. The latter, as opposed to the former, never knows n^* since he does not know the actual values of all the Δ_n^* (he cannot know more than all the Δ_n^* up to a finite N).

6. THE OBSERVER-THEORIST'S VIEWPOINT : CONVERGENCE PROBLEMS

There are two problems that the observer-theorist must address : first, is there an optimal n^*, i.e., a logical level at which the agent should, or should have, operated? And second, in what sense and when does an infinite regress of optimizing theory converge? Existence problems relating to n^* were briefly touched upon in the last section and will not be discussed any further here. We shall only mention one possible attractive interpretation of the optimal n^*. It has been emphasized that n^* is not rationally chosen by agents; hence it could prevail only by some impersonal process of the "natural selection" type.

To be specific, consider a set of agents who go through a natural section process which conforms to the following, standard principle: the survivors are those who have the highest net utility. Agents may all be optimizing at the topmost logical level n at which they operate, with n varying throughout the population; then, only n*-optimizers will remain alive (supposing, of course, n* finite). They may use any procedure whatsoever at a fixed n, n being their topmost logical level; the survivors will be those agents who have adhered to P_n^* spontaneously. In the most general case, where both procedure types and topmost logical levels

vary, only agents who have followed R^* spontaneously will survive.

There is another task that the agent could not himself undertake and is for the observer-theorist only - formally to investigate infinite regresses. The latter notion has been left implicit thusfar. Restricting ourselves to the case of global problems, define an *infinite regress of optimizing theory* as an infinite sequence of optimizing procedures $(\hat{P}_1, \ldots, \hat{P}_n, \ldots)$ generated from each other in the following way:

- if $\hat{P}_1 = (T^o, V)$, $\hat{P}_2 = \left(T^o, V, C_1, \mathcal{P}_1\right)$ with $\hat{P}_1 \in \mathcal{P}_1$ and $C_1 : \mathcal{P}_1 \to \mathbb{R}$;
- more generally, if $\hat{P}_n = \left(T^o, V, C_{n-1}, \mathcal{P}_{n-1}\right)$, $\hat{P}_{n+1} = \left(T^o, V, C_n, \mathcal{P}_n\right)$, with $\hat{P}_n \in \mathcal{P}_n$ and $C_n : \mathcal{P}_n \to \mathbb{R}$.

These set-theoretic conditions are clearly necessary for each n-level programme to be "superseded" by the (n+1)-programme, i.e. to be compared with some rivals of the same logical level in a metaoptimizing way. This definition, of course, does not ensure unicity of $\hat{P}_{n+1}$ once $\hat{P}_n$ is given ; the number of infinite regresses generated from a standard problem is in principle indefinitely large and can be narrowed only by *a priori* (non-endogenizable) restrictions on the $\mathcal{P}_n$ and C_n. Note that sequence of $\hat{P}_n$ may be replaced by sequences of $\hat{V}_n$ or $\hat{P}_n$ in the context of the simplifying approaches of section 4.

"Convergence" here may formally be defined in a variety of ways, but the intuitively relevant concept is *decision convergence*: there is a $N \in \mathbb{N}^*$ such that the R^{N*} induced by higher-order $\hat{P}_{N+1}, \ldots, \hat{P}_{N+j}, \ldots$ does not change.

Formally :

$$(\exists\, N \in \mathbb{N}^*)(\forall\, j \in \mathbb{N}^*) \qquad R^{N+j*} = R^{N*}$$

This *stationary* convergence concept is the suitable one when the $\mathcal{P}_n$ are finite; it may be generalized into an *asymptotic* convergence concept when the $\mathcal{P}_n$ are infinite topological

spaces. It is worth emphasizing that an infinite regress may decision-converge with net utility continuing to vary from one level to another (for instance, decreasing from some n onwards, which is the most plausible empirically) [8]. An infinite regress can be shown to be stationary at N if there is *a* $j \geq 1$ such that (N+j)-optimizing leads back to the N-optimizing procedure: $\tilde{P}_N\left(\tilde{P}_{N+1}\left(\ldots(\hat{P}_{N+j})\right)\right) = \hat{P}_N$. Intuitively, $\hat{P}_N$ is revealed to be both efficient and simple enough for discovery and computation purposes; if metaoptimization "warrants" optimization, there is nothing more to do. Even more trivially, decision convergence obtains when there is a N and a P_N such that: $P_N = \tilde{P}_N(P_{N+1})$ for all $P_{N+1} \in \mathcal{P}_{N+1}$. For instance, if all the procedures compared by the metaoptimizer lead to the same decision rule, there is nothing more to do. Beyond these unsurprising results, necessary and/or sufficient conditions for decision or utility convergence can only be defined within the framework of particular models.

Here is an example borrowed from Mongin (1986). As mentioned in section 4, procedures are identified with pairs $\langle R_n, A_{n-1} \rangle$ so that infinite regresses may be redefined as sequences

$$\langle R_1, A_0 \rangle, \ldots, \langle R_n, A_{n-1} \rangle, \ldots$$

the $\langle R_n, A_{n-1} \rangle$ being connected with the $\langle R_{n-1}, A_{n-2} \rangle$ through natural set-theoretic conditions. It is assumed that the agent incurs no discovery cost, but only a computation cost $\Gamma''(R_n, e)$ which is associated with his computing the lower-level $R_n(e)$ once R_n and e are given. It is also assumed that $\Gamma''(R_n, e)$ increases with Card(Rge R_n) for any given e - roughly speaking, the more values in the range, the more costly the decision rule is to operate. These hypotheses can be shown to imply a consequence which is embarrassing for the optimizing theory of decision: an infinite regress is decision-convergent only if the R_j, $j \geq 1$, are *constant* rules [9]. Intuitevely, if a complex, i.e. non constant rule R_j is optimally chosen at level j, it is

replaced by a simple, i.e. constant rule R'_j when optimizing takes place at (j+1). Once this substitution process has started, it never comes to a halt, so that the infinite regress cannot decision-converge. The only convergence case is therefore that in which the substitution process does not even *begin*, i.e. the level 1 optimal rule is a constant rule. Since this is a very particular case, the model implies a logical objection against optimization. It is one among many examples of how Winter's hint that there is no "superoptimization" can be sharpened into a formal argument against the optimizing theory of decision.

Ecole Normale Supérieure and *Ecole Nationale des Ponts et Chaussées, Paris.*

NOTES

[1] E.g. the paradox of the origin of institutions (collectively to decide upon a social rule requires that a collective decision rule be defined first); the paradox of mutually referring expectations (I expect that you expect that I expect, etc) which has emerged from game theory literature (Walliser, 1985); the paradox of information search - in order to find out which information is available, individuals should rely on preexisting information - which is closely related to the infinite regresses discussed in this paper (Winter, 1971)

[2] Following the "intellectualist legend", "our intellectual planning process must inherit its title to shrewdness from yet another interior process of planning to plan, and this process could in turn be silly or shrewd. The regress is infinite and this reduces to absurdity the theory..." (Ryle, 1949, p. 31-32).

[3] Note in particular that the logical relationship between satisficing and optimizing models is still very unclear, as testified to by conflicting views on the subject (for a brief review, see Mongin, 1984).

[4] One might be tempted to dispense with the notion of procedure altogether and analyze decisions in terms of their observable consequences only, since applying a (T, D) is

tantamount to defining a correspondence from E into X. This would be convenient for some purposes, but is also potentially confusing, for there is no one-to-one mapping between procedures and their observable consequences. In order to keep notations handy, it is assumed in most of this paper that the *correspondence* E → X associated with any (T, D) collapses into a *single-valued function*.

[5] In this way, procedures are not dispensed with, since the use of a particular (T, D) reflects on the decision rule sets which are compatible with it.

[6] In order to find R_n and A_{n-1}, the agent should *discover* the implicit determining factors and then *compute* R_n and the members of A_{n-1}.

[7] If net utility fluctuates without limit, n^* may be taken to be the highest adherence value of the net utilities sequence.

[8] The definitions above are compatible with local as well as global optimizing problems. Considering the former case, it may happen that $(\hat{P}_1, \ldots, \hat{P}_n, \ldots)$ decision-converges for *some*, but *not all* $e \in E$. Decision convergence in the latter case is nothing but decision convergence for all e; note that this concept does not ensure uniformity (the lowest "stable" level N is variable with e) and even stronger definitions might eventually be needed.

[9] Assuming that there is a unique R_1 which solves the global level 1 optimization problem.

REFERENCES

Aristotle's 'Physics', ed. by W.D. Ross, Oxford, Clarendon Press, 1936;

The 'Nicomachean Ethics' of Aristotle, ed. by J.S. Brewer, Oxford, 1836.

Baumol, W.J., and R.E. Quandt (1964), 'Rules of Thumb and Optimally Imperfect Decisions', *American Economic Review*, 54, p. 23-46.

Beach, L.R., and Mitchell T.R. (1978) : 'A Contingency Model for the Selection of Various Strategies', *Academy of Management Review*, 3.

Göttinger, H.W. (1982), 'Computational Costs and Bounded Rationality', in W. Stegmüller, W. Balzer and W. Spohn eds., *Philosophy of Economics*, Berlin, Springer.

Latsis, S. (1976), 'A Research Programme in Economics', in S. Latsis ed., 1976, p. 1-41.

Latsis, S. (1976) (ed), *Method and Appraisal in Economics*, Cambridge, Cambridge University Press.

Lester, R.A. (1946), 'Shortcomings of Marginal Analysis for Wage Employment Problems', *American Economic Review*, 36, p. 63-82.

Mongin, P. (1984), 'Modèle rationnel ou modèle économique de la rationalité?', *Revue économique*, vol. 35, p.9-63.

Mongin, P.(1986), 'Is it Logically Impossible to Optimize?', mimeo, Laboratoire d'économie politique, Ecole Normale Supérieure, Paris.

Ryle, G. (1949), *The Concept of Mind*, London, Hutchinson.

Shugan, S.M. (1980), 'The Cost of Thinking', *Journal of Consumer Research*, 7, p. 99-111.

Simon, H.A. (1959), 'Theories of Decision-Making in Economics and Behavioral Science', *American Economic Review*, 49, p. 253-283.

Simon, H.A. (1976), 'From Substantive to Procedural Rationality', in S. Latsis ed., 1976, p. 129-148.

Simon, H.A. (1979), 'Rational Decision Making in Business Organizations', *American Economic Review*, 69, p.494--513.

Stigler, G.J. (1961), 'The Economics of Information', *Journal of Political Economy*, 69, p. 213-225.

Walliser, B. (1985), *Anticipations, équilibres et rationalité économique*, Paris, Calmann-Lévy.

Winter, S. (1971), 'Satisficing, Selection and the Innovating Remnant', *Quarterly Journal of Economics*, 85, p. 237-261.

Winter, S. (1975), 'Optimization and Evolution in the Theory of the Firm', in R.H. Day and T. Groves eds, *Adaptive Economic Models*, New York, Academic Press, p. 73-118.