Convergence in sequential decision making with learning and costs to change

Juan Enrique Martínez-Legaz*†
Departament d'Economia i d'Història Econòmica
Universitat Autònoma de Barcelona
08193 Bellaterra
Spain

Antoine Soubeyran
Aix-Marseille University (Aix-Marseille School of Economics)
CNRS & EHESS
Château Lafarge, route des Milles
13290 Les Milles
France

Abstract

We analyze a sequential decision making process, in which at each step the decision is made in two stages. In the first stage a partially optimal action is chosen, which allows the decision maker to learn how to improve it under the new environment. We show how inertia (cost of changing) may lead the process to converge to a routine where no further changes are made. We illustrate our scheme with some economic models and an example about institutional changes.

Keywords: sequential decision making, costs to change, convergence.

1 Introduction

There are many sequential decision making processes in which, at each step, the decision is made in two stages. In the first stage the decision maker chooses a partially optimal solution, after the implementation of which he learns about the new environment and uses this learning to improve the current solution.

*Corresponding author. E-mail address: JuanEnrique.Martinez.Legaz@uab.cat. Phone: (+34)935811366. Fax: (+34)935812012. Postal address: Departament d'Economia i d'Història Econòmica; Universitat Autònoma de Barcelona; 08193 Bellaterra; Spain

†This author has been supported by the MICINN of Spain, Grant MTM2011-29064-C03-01. He is affiliated to MOVE (Markets, Organizations and Votes in Economics).

After presenting our mathematical model in the next section, we will give several examples in which this situation occurs and fit our model very well. In this paper we focus on the interlinked dynamic effects arising in this type of processes, an essential ingredient in our model being the inertia effects (costs to change). We present simple conditions under which the process converges, leading to a routinized action choice in the limit. We also discuss some real world situations in which our model applies.

The rest of the paper is organized as follows. Section 2 presents the model and illustrates it by means of some examples. Section 3 presents a convergence theorem. In Section 4 we summarize our conclusions. Our convergence theorem relies upon a new Caristi type lemma.

2 The model

We consider a decision maker who wants to maximize a utility function $U : X \longrightarrow \mathbb{R}$ over the decision space X. We assume that $X \subseteq X_1 \times X_2$, the sets X_1 and X_2 consisting of partial decisions belonging to two different categories. Each decision $x = (x_{1,}x_2)$ has two components; the first one, $x_1 \in X_1$, can be implemented at no cost, whereas to implement the second component, $x_2 \in X_2$, one incurs a cost which depends on the amount of change imposed on this variable relative to the preceding decision. We assume that, for every given $x_1 \in X_1$, the problem of maximizing $U(x_{1,}x_2)$ over the set $F(x_1) := \{x_2 \in X_2 : (x_{1,}x_2) \in X\}$ is (relatively) easy to solve. Thus, what makes the decision problem difficult is the choice of $x_1 \in X_1$, and we assume that there is a nonnegative cost $C(x_1, x_1')$ of changing from a given $x_1 \in X$ to a new $x_1' \in X$. We assume that the costs are measured in the same units as the utility function.

Notice that the decision maker may be an individual or some relatively complex organization. In the latter case, the tasks of choosing $x_1 \in X_1$ and choosing $x_2 \in X_2$ may be made by different parts within this organization.

We consider the following sequential decision making process. Given an initial $x_1 \in X_1$, the decision maker makes an optimal (relative to x_1) decision by maximizing $U(x_1, x_2)$ over $F(x_1)$. Let $x_2 \in X_2$ be an optimal solution to this maximization problem. The decision maker will search for the optimal $x_1 \in X_1$ relative to x_2 taking into account the costs to change, that is, he will maximize $U(x_1', x_2) - C(x_1, x_1')$ over the set $F^{-1}(x_2) := \{x_1' \in X_1 : (x_1', x_2) \in X\}$. He will then change x_1 accordingly, after which he will proceed to the next iteration, choosing an optimal $x_2 \in X_2$ relative to the new x_1, etc..

We devote the rest of this section to describe a few real world situations which fit into our model.

Consider first a consumer who first chooses a store x_1 from a set X_1 of available stores and then a commodity (say, a computer) x_2 from a set X_2 of commodities of the same type offered by the chosen store. The initial store may be the one closest to his home, and he will choose the optimal computer available at this store. After having decided which computer he wants to buy, he will search for the same computer in other stores, in order to find out the store that offers the best conditions (cheaper price, better technical service,...). The choice of a new store involves costs of change, including, for instance, the ones derived from traveling from the current store to the new one. Once an optimal store is chosen, the consumer will search for the optimal computer offered by this store; then he will search for the optimal store where this new computer is available, and so on so forth.

As another example, consider an individual who has a job. Given this job and its associated income, the individual chooses an optimal standard of living. Once this choice is made, he searches for the most convenient job that is compatible with his chosen way of life (a job closer to his home, a less demanding job,...). Changing the job entails some costs, like, for instance, adaptation costs. After changing his job, the individual aims at improving his standard of living relative to the new job, etc..

Another illustration comes from production theory. Let X_1 be a collection of technologies or production sets. We define the set X_2 as the union of all technologies, so that an element $x_2 \in X_2$ is a feasible production plan $(u, v) \in x_1$ for some $x_1 \in X_1$. The decision space is $X := \{(x_1, (u, v)) \in X_1 \times X_2 : (u, v) \in x_1\}$. Given a technology $x_1 \in X_1$, the producer chooses a production plan $(u, v) \in x_1$ so as to maximize its net profit over x_1. After this choice is made, he aims at improving the technology while keeping the chosen production plan feasible. Choosing a new technology is an innovation choice involving some costs of change. But this innovation allows for new production plans, so the producer then chooses the optimal production plan compatible with this new technology. This procedure can repeat indefinitely.

Our last example is about the evolution of institutions. In this example the decision maker is a pair consisting of an institution and a set of individuals. The institution and the individuals share a common goal, specified by a utility function. The institution sets an environment x_1 (e.g., a rule), under which the individuals choose an optimal profile of actions x_2. After observing the chosen profile of actions, the institution improves the environment subject to the constraint that the chosen profile of actions be still feasible, taking into account the costs it will incur by changing the environment. Once a new environment is established, the individuals adapt to the new situation by choosing an optimal profile of actions compatible with the new environment. We thus have an iterative process fitting our scheme. For references on institutional change, see [4, 3].

3 Convergence of the iterative process

Let as assume that the set X_1 has a metric space structure (X_1, d), the set X_2 is a compact topological space, the decision space X is a subset of $X_1 \times X_2$, the cost to change function $C : X_1 \times X_1 \longrightarrow \mathbb{R}_+$ is continuous and satisfies $C(x_1, x_1) = 0$ for all $x_1 \in X_1$, and the utility function $U : X \longrightarrow \mathbb{R}$ is continous and bounded from above.

The following technical assumptions will also be needed:

1) There exists a subadditive, non decreasing, continuous function $\phi : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$ satisfying $\phi(0) = 0$, $\phi(t) > 0$ for $t > 0$, and $\phi[d(x_1, x_1')] \leq C(x_1, x_1')$ for all $x_1, x_1' \in X_1$. Notice that $\phi \circ d$ is a metric inducing a topology equivalent to that induced by d.

2) The correspondences $X_1 \ni x_1 \rightrightarrows F(x_1) \subseteq X_2$ and $X_2 \ni x_2 \rightrightarrows F^{-1}(x_2) \subseteq X_1$ defined by $F(x_1) := \{x_2 \in X_2 : (x_1, x_2) \in X\}$ and $F^{-1}(x_2) := \{x_1 \in X_1 : (x_1, x_2) \in X\}$, respectively, are compact-valued and continuous. Notice that this continuity assumption holds, for instance, if $x_1 \in X_1$ and $x_2 \in X_2$ are independent decisions, that is, if every two such partial decisions are compatible. Indeed, this amounts to saying that $X = X_1 \times X_2$, so that $F(x_1) = X_2$ for all $x_1 \in X_1$ and $F^{-1}(x_2) = X_1$ for all $x_2 \in X_2$.

3) For every $x_1 \in X_1$, the function $U(x_1, \cdot)$ has a unique maximizer $Rx_1 := \arg\max_{x_2 \in F(x_1)} U(x_1, x_2)$. This hypothesis holds, for example, if X_2 is a convex subset of a topological vector space, for every $x_1 \in X_1$ the set $F(x_1)$ is convex, and U is strictly quasiconcave in its second argument. By Berge's maximum theorem [1], the mapping $R : X_1 \longrightarrow X_2$ is continuous.

Define $G : X_1 \longrightarrow \mathbb{R}$ by

$$G(x_1) := \max_{x_1' \in F^{-1}(Rx_1)} \{U(x_1', Rx_1) - C(x_1, x_1')\}. \tag{1}$$

This function assigns to each partial decision $x_1 \in X_1$ the maximum payoff the decision maker can get by changing it while keeping $x_2 = Rx_1$ unchanged, net of the cost to change.

4) For every $x_1 \in X_1$, the maximization problem in (1) has a unique solution $Tx_1 \in X_1$. This hypothesis holds, for instance, if X_1 is a normed vector space, for every $x_2 \in X_2$ the set $F^{-1}(x_2)$ is convex, U is strictly quasiconcave in its first argument, and C is strictly convex in its second argument. By Berge's maximum theorem, $T : X \longrightarrow X$ is continuous.

The following lemma is similar to the Caristi fixed point theorem [2], but our assumptions are different: Instead of imposing a semicontinuity assumption on the function G we just assume it to be bounded from above, but unlike in the case of Caristi theorem we require the mapping T to be continuous.

Lemma 1 *Let (Y, d) be a complete metric space and $G : Y \longrightarrow \mathbb{R}$ be a bounded from above function. If $T : Y \longrightarrow Y$ is a continuous mapping such that $d(y, Ty) \leq G(Ty) - G(y)$ for each $y \in Y$, then, for any $y_0 \in Y$, the sequence $\{T^n y_0\}$ converges to a fixed point $\overline{y}$ of T.*

Proof. Since $G(T^{n+1}y) \geq G(T^n y) + d(T^n y, T^{n+1}y) \geq G(T^n y)$, the sequence $\{G(T^n y)\}$ is nondecreasing. As it is bounded from above, it is also convergent. For any $m, p \in \mathbb{N}$, one has

$$\begin{aligned} d(T^m y, T^{m+p} y) &\leq \Sigma_{i=0}^{p-1} d(T^{m+i} y, T^{m+i+1} y) \\ &\leq \Sigma_{i=0}^{p-1} \left[G(T^{m+i+1} y) - G(T^{m+i} y)\right] \\ &= G(T^{m+p} y) - G(T^m y). \end{aligned}$$

This proves that $\{T^n y\}$ is a Cauchy sequence. By the completeness of Y, this sequence converges to some point $\overline{y} \in Y$, which, as T is continuous, is a fixed point of T. ■

Theorem 2 *For every $x_1^0 \in X_1$, the sequence $\left\{\left(T^n x_1^0, RT^n x_1^0\right)\right\}$ converges to $(\overline{x}_1, R\overline{x}_1)$ for some fixed point $\overline{x}_1 \in X_1$ of T.*

Proof. Since U is bounded from above and C is non negative, G is bounded from above, too. For each $x_1 \in X_1$, using that $Tx_1 \in F^{-1}(Rx_1)$ and $RTx_1 \in F(Tx_1)$, that is, $Rx_1 \in F(Tx_1)$ and $Tx_1 \in F^{-1}(RTx_1)$, respectively, and that $C(Tx_1, Tx_1) = 0$, one sees that

$$\begin{aligned}
\phi(d(x_1, Tx_1)) &\leq C(x_1, Tx_1) + \max_{x_2 \in F(Tx_1)} U(Tx_1, x_2) - U(Tx_1, Rx_1) \\
&= U(Tx_1, RTx_1) - G(x_1) \\
&\leq \max_{x_1' \in F^{-1}(RTx_1)} \{U(x_1', RTx_1) - C(Tx_1, x_1')\} - G(x_1) \\
&= G(Tx_1) - G(x_1).
\end{aligned}$$

Hence, by Lemma 1 applied to the metric space $(X_1, \phi \circ d)$, the sequence $\left\{T^n x_1^0\right\}$ converges to some fixed point $\overline{x}_1 \in X_1$ of T. Since R is continuous, it follows that the sequence $\left\{RT^n x_1^0\right\}$converges to $R\overline{x}_1$. Consequently, $\left\{\left(T^n x_1^0, RT^n x_1^0\right)\right\}$converges to $(\overline{x}_1, R\overline{x}_1)$. ■

4 Conclusion

We have examined a sequential two stages decision making process where the decision maker makes first a costly partial decision and then completes it in a costless optimal way. Our model takes into account inertia, i.e. costs to change, and is shown to converge to a stable decision under suitable assumptions. The main result essentially shows that high costs to change or, more specifically, costs which increase with the distance between partial decisions are enough to have convergence.

References

[1] Berge, C. (1959) *Espaces topologiques: Fonctions multivoques*, Paris: Dunod.

[2] Caristi J. (1976) "Fixed Point Theorems for Mappings Satisfying Inwardness Conditions," *Transactions of the American Mathematical Society*, 215, 241–251.

[3] Kingston C., Caballero G. (2009) "Comparing Theories of Institutional Changes," *Journal of Institutional Economics*, 5, 151-180.

[4] Roland, G. (2004) "Understanding Institutional Change: Fast-moving and Slow-moving Institutions," *Studies in Comparative International Development* 38, no. 4, 109-131.