Laboratory of Economics and Management
Sant'Anna School of Advanced Studies

Piazza dei Martiri della Libertà, 33 - I-56127 PISA (Italy)
Tel. +39-050-883-341 Fax +39-050-883-344
Email: lem@sssup.it Web Page: http://lem.sssup.it

LEM
Working Paper Series

How to get what you want when you do not know what you want. A model of incentives, organizational structure and learning

Luigi Marengo[1], Corrado Pasquali[2]

[1]*LEM, Scuola Superiore Sant'Anna, Pisa, <l.marengo,g.dosi>@sssup.it*
[2]*Università di Teramo, cpasquali@unite.it*

2010/08 May 2010

How to get what you want when you do not know what you want. A model of incentives, organizational structure and learning

Luigi Marengo

LEM, Scuola Superiore Sant' Anna, Pisa, Italy, l.marengo@sssup.it

Corrado Pasquali

DSGSS, Università di Teramo, Teramo, Italy, cpasquali@unite.it

Abstract

In this paper we present a model of the interplay between learning, incentives and the allocation of decision rights in the context of a generalized agency problem. Within this context, not only actors face conflicting interests but diverging cognitive "visions" of the right course of action as well. We show that a principal may obtain the implementation of desired organizational policies by means of appropriate incentives or by means of appropriate design of the allocation of decisions, when the latter is cheaper but more complex. We also show that when the principal is uncertain about which course of action is more appropriate and wants to learn it from the environment, organizational structure and incentives interact in non-trivial ways and must be carefully tuned. When learning is not at stake, incentives and organizational structure are substitutes. When instead learning is at stake, organizational structure and incentives may complement each other and have to be fine tuned according to the complexity of the learning process and the competitive pressure which is put on fast or slow learning.

Keywords: Incentives, Organizational Structure, Learning

Acknowledgements: The support to this research by the European Union - 6th Framework Programme - DIME Network of Excellence is gratefully acknowledged.

In order to arrive at what you do not know
You must go by a way which is the way of ignorance
T. S. Eliot, East Coker, 1940

1 Introduction

Broadly speaking, the main cutting divide between organizational economics and the evolutionary and capabilities approaches should probably be looked for in the way in which firms' specificity as coordination devices is represented and analyzed. In particular, it seems that the fundamental epistemological differences between the two research streams are grounded on the adoption of different mind-sets and primitive notions.

Organizational economics takes incentives and governance as its primitives and mainly describes firms as contract- and/or hierarchy-based solutions to market failures. Its main concern is the problem of efficient allocation of given resources and given capabilities and its key research questions are the design of optimal information and incentive systems that, in a static context, boils down to a problem of optimal allocation of risk (Jensen and Meckling 1976, Fama 1980), of optimal information structures (Sah and Stiglitz 1986, Radner 1993), of optimal allocation of property rights (Hart and Moore 1990), of control, decision and exclusion rights (Bester 2009, Rajan and Zingales 1998), depending on which informational or incentive problem the analysis is focused upon.

The evolutionary and capabilities approaches[1], quite on the contrary, find their "primitives" for the analysis of the nature of economic organizations in their problem-solving features. The latter, in turn, are viewed as embedded in some form of human bounded rationality, in imperfect processes of learning and diverse mechanisms of social distribution of "cognitive labor". In this perspective, firms are conceived of as the main place for the creation and implementation of productive knowledge. Learning and adaptation are the main concerns of this research perspective and the key research questions concern how capabilities are acquired and modified and what are the organizational structures and processes that favor the generation of organizational capabilities quite independently from any issue of incentive compatibility and

[1] There are indeed important difference between evolutionary and capabilities theories of the firm (see for instance Dosi, Faillo, and Marengo (2008)), but for the present discussion these differences do not seem fundamental.

transaction governance.

A real confrontation between the two theories is somehow difficult because they largely lack a common ground. Organizational economics has strongly emphasized the role of incentives in organizations, reflecting the overall idea that the role of institutions and organizations is basically to mitigate the negative consequences of opportunism by setting the right incentives and/or the right governance structure. The capabilities view often makes the implicit assumption that individual motivation plays little or no role in the generation and accumulation of capabilities, or at least that incentive compatibility can be loosely assumed.

Last but not least, the two traditions have also important differences in methodology that make cross-communication difficult: organizational economics is deeply embedded in the neoclassical tradition of abstract analytical modeling based upon the standard toolbox of rationality and equilibrium behavior, while the capabilities view is usually oriented towards appreciative theorizing and assumes individuals with strong bounds in rationality and knowledge. Moreover, sometimes it assumes that there might well exist organizational dimensions that are not necessarily reducible to the individual behaviors of its members (more on the relations between the two theories in, e.g., Foss (2000)).

Some existing attempts of bridging the gaps of the two streams of research have been made (Langlois and Foss 1999, Dosi, Levinthal, and Marengo 2003, Coriat and Dosi 1998, Nickerson and Zenger 2004, Kaplan and Henderson 2005), but on the whole it is not unfair to say that the organizational economics literature has very little to say on learning and capabilities creation and that the capabilities literature does not deal in a satisfactory way with the role of incentives, delegation and power in the creation and modification of capabilities. In other words, at one extreme one finds a theory that to some extent censors any competence issue associated with what organizations do and how well they do it except for issues of misrepresentations of intrinsic individual abilities and adverse selection, or incentive misalignment in effort elicitation. At the other extreme, one finds a theory that censors precisely the incentive-alignment issue, in a sense pretending that all agents are perfectly benevolent cooperators as far as their individual motives are concerned. At the very same time, it focuses on the problem solving efficacy of what they do, especially in so far as what they do primarily stems from the social division of labor.

In this paper we make a novel attempt at bridging this gap that, we believe, makes some non trivial steps forward. Painted with an extremely broad brush, our contribution amounts to adding two dimensions to the the general picture. The first is a political one: in our model there exists a social function for power that amounts to the possibility given to a principal of structuring and constraining agents' adaptive search both through (re)allocating decision rights and through implementing different incentive mechanisms. On the other hand, we explore the main properties of this function of power in agency relations in which conflict arises not only from diverging interests but also from diverging views of the appropriate courses of actions or from different representations of the world. The second dimension we add is thus a cognitive one. Our main focus is on how a principal has to choose between the two forms of power - i.e. changing organizational structure and fine tuning incentives - in order to maximize his utility. As we will show, these two dimensions interact in non trivial ways depending both on the nature and the representation of the problem at hand.

We present an abstract model of the interplay between organizational structure, incentives and learning and we focus on the interaction between the allocation of decision rights and incentives when the organization is facing complex problems, i.e. problems in which the organizational behavior is the outcome of the interaction of many interdependent decisions with strong externalities, both positive and negative. We show that allocation of decision rights and incentives are largely substitutes: a principal can obtain a desired course of action by appropriate reallocation of decision rights and/or by providing appropriate monetary incentives to the agents. The former strategy, i.e. acting on the organizational design, is very powerful and less expensive and we show that in general by increasing the division of decision making rights the principal may have her policies more easily implemented.

The picture becomes more blurred and complicated when the principal does not know the appropriate course of action but tries to learn it from environmental feed-back. In this circumstances the principal is facing a difficult trade-off. By using efficiently the organizational structure (i.e. the allocation of decision rights) and/or the incentives, she may get her policy more efficiently implemented, but she runs the risk of curbing those agents' alternative visions that may prove very useful for collective learning. This trade-off is an instance of the widely discussed exploitation vs. exploration trade-off (March 1991), in the sense that stronger incentives increase the control of the principal who may obtain a stricter implementation of the required policies (exploitation) but for the same reason they limit the exploration of alternative policies that agents may autonomously choose. A careful tuning of the trade-off between the organization of decision rights and incentives is therefore a key issue.

The paper is organized as follows: in section 2 we describe the main issues involved in the interaction between incentives, organizational structure and learning. In section 3 we outline the model and we study its main properties. Results are presented in sections 4 and 5 where we discuss the behavior of the model respectively when the principal knows precisely what she wants to get from the agents and when instead she tries to learn what are the best courses of actions. In section 6 we analyze how incentives and organizational structures cope with the complexity problem generated by the interdependencies (externalities) among the agents. Finally, in section 7 we conclude and suggest some directions for further developments.

2 Generalized agency relations in learning and adaptation

Organizational economics usually assumes that conflict in organizations arises because individuals have diverging objectives. A typical agency model assumes a principal's utility increasing in the result, decreasing in the salary paid to the agent and indifferent with respect to the latter's effort. On the contrary, the agent's utility increases in salary, decreases in effort and is indifferent to the result. Recent literature on social preferences has challenged these conventional assumptions and has highlighted that principals and agents may indeed care about each other's utility (Fehr and Fischbacher 2002), for instance because agents tend to reciprocate generous (incentive incompatible) offers of high compensations with generous provisions of effort. But if the principal's and the agent's decisions depend on the perceived behavior and attitude of the other, then the social and organizational context in which interaction takes place may indeed influence the parties' perceptions. For instance, many researchers have pointed out that monetary incentives may have a crowding out effect against cooperation as they are perceived as a signal of lack of trust (Fehr and Gächter 2002, Frey and Jegen 2001).

In this paper we investigate a complementary issue: we assume that agents are not indifferent to the result, because they have not only preferences and interests, but also cognition, ideas, visions about what the organizational course of action should be, well beyond the mere interest in maximizing the salary net of effort costs. Without downplaying the role of diverging interests, it must be recognized that also diverging views are an important source of conflict in organizations. Everyone who has had some managing role in a business, academic or governmental organization has probably experienced such conflict: people simply have different ideas about what should be done and how it should be done. Often such different ideas can only partly, or not all, be ascribed to their self interest. Agents hold diverging and motivationally strong views for the mere fact that they sincerely believe that their intended course of action is good for the organization's interest and attach high value to this belief. Conflict arising from diverging interests and conflict arising from diverging views are often strictly intertwined: a manager of a division or department may think that more resources must be allocated to the unit she manages both because she believes to the best of her knowledge that this will serve the organization's objectives (and indeed this may prove right) and because she looks for private benefits that she may reap in terms of higher salary, power, visibility and prestige.

People do indeed hold different views of what should be done, how things should be managed, which alternative courses of actions should be followed also because, in good faith, they think differently about how the same organizational objectives could be better achieved. This source of conflict is likely to be especially relevant when non-routine decisions have to be taken, when new hard problems are being faced, when strong procedural uncertainty characterizes the current situation, when organizational and or technological change is needed, that is, in all those situations in which non-routine courses of action must be envisaged and what must be done is far from obvious and uncontroversial. In such circumstances, organizations do not have to find optimally efficient allocations of given resources, but have to design complex procedures that may provide valuable solutions to ill-structured problems (Simon 1981). However, in such situations, different visions are also a fundamental source of learning. When the principal does not know exactly what should be done, she may learn from the agents' ideas. In this respect, the standard solution to the problem of conflict suggested by agency model, that is aligning the agents' preferences with the principal's, may actually prove detrimental and curb this important source of learning through diversity.

Agency conflict is a source of inefficiency and incentives are needed in order to correct for misalignment of objectives. As well known, in the presence of information asymmetries and with risk-averse agents, perfect alignment is usually impossible and full efficiency cannot be restored (Fama 1980). However, when conflict arises because of different views of what must be done, alignment may be difficult also lacking information asymmetries because agents are concerned with organizational actions and with their individual effort as well. In such cases actions and decisions by some agents tend to produce externalities on the other agents that may be both positive (agent i chooses an action that is aligned with what agent j thinks should be done) or negative (when an agent i's action differs from what agent j would have chosen in that situation). Such negative and positive externalities are a source of complexity that contractual arrangements can hardly manage optimally (Bernholz 1997). Moreover, and more importantly, alignment may not be desirable because agents' cognitions, ideas and visions may prove partly or totally superior to the principal's. If the latter succeeds in obtaining a perfect implementation of her desired actions, she looses the opportunity of discovering better ones that agents may know.

In this paper we suppose that principals and agents hold different views of an articulated course of action that we model by way of a vector of interdependent policies. The principal has both a problem of implementation of the wanted policies and a problem of inadequacy of what she believes the right policies are. In order to solve these two problems she can act on the allocation of decisions among agents[2], and on the incentives, i.e. trying to modify the agents' preferences. We will show that in our model allocation of decision rights and incentives can either be substitute or complements, depending on whether learning is at stake or not.

In the business strategy literature, the former problem is referred to as the strategy implementation problem (Hrebiniak and Joyce 1984) and considered as a source of inefficiency. The organization is viewed as knowing an appropriate course of action but for a variety of incentive and coordination reasons is not realizing that set of policies. But an alternative view, e.g. the literature on emergent strategy (Mintzberg 1973, Burgelman 1994) suggests that the divergence between expressed strategy and actual behavior may be

[2] In a recent paper, Canice Prendergast suggests a somehow similar and complementary framework by developing a model in which the principal's main tool for alleviating incentive problems is to hire the right agent in the right position (Prendergast 2009).

a favorable circumstance. The search and discovery that results from such discrepancies may yield the identification of a superior set of actions than that which would be suggested by the conscious choice of strategy.

In the following section we outline a model that should help clarifying these trade-offs under more rigorous terms.

3 The model

We consider a firm that has to take decisions on a set of n policies $P = \{p_1, p_2, \ldots, p_n\}$. For simplicity we assume that each policy may take only two values $p_i \in \{0, 1\}$ and therefore the set of policies if formed by the 2^n vectors of n binary elements. We will call X this set of 2^n policy vectors and $x_i = [p_1^i, p_2^i, \ldots, p_n^i]$ one generic element thereof.

We concentrate on those cases in which policies interact with each other in complex ways to determine the overall organizational performance. Decisions on single policy items generate externalities, both positive and negative, on other policies. Thus the determination of the correct combination of policies is a complex task as the performance contribution of a single policy item depends upon the value taken by other policies. Complementarity and superadditivity (Milgrom and Roberts 1992) among policies are special cases.

We suppose that policy vectors have an exogenously determined performance ordering that we call "nature's" ordering and write $\succ_N$. If $x_i \succ_N x_j$ then policy vector x_i has strictly higher performance than vector x_j. Nature's ordering determines a policy landscape - i.e. the coupling of every policy vector with its performance - whose ruggedness[3] reflects the extent of interdependencies among policies and thus the complexity of the problem of finding the best performing policy vector(s) (Levinthal 1997, Page 1996, Rivkin and Siggelkow 2005)[4]. In the analysis and the simulation exercises below we will consider, unless otherwise specified, a generic complete and transitive ordering of policy vectors, without any further restriction.

Our organization is composed by a principal Π and a number of agents that may range from 1 up to n. Each agent is attributed decision rights over

[3]In our model a policy landscape is highly rugged when modifications of one policy item determines large variations in the overall performance.

[4]Actually these papers assume a fitness function, that is (random) assignments of performance level to each policy vector, usually normalized in the interval $[0, 1]$. For our model we do not need fitness (or performance) values but simply a complete and transitive performance ordering for policy vectors.

a subset of policies. Let $A = \{a_1, a_2, \ldots, a_h\}$, with $1 \leq h \leq n$, be a set of agents and let each agent be associated to a non-empty subset of policies under his control. More precisely, let $d_i \subseteq P$ be a generic non-empty subset of the set of policies. We call a decomposition of decision rights a partition[5] of the set of policies, i.e. a set of non-empty subsets $D = \{d_1, d_2, \ldots, d_k\}$ such that:

$$\bigcup_{i=1}^{h} d_i = X \text{ with } d_i \bigcap d_j = \emptyset \text{ , } \forall i \neq j$$

We call organizational structure $\mathbf{O}$ a mapping of the set D onto the set A of agents, i.e. a mapping that assigns each subset of policies to one and only one agent, i.e. $\mathbf{O} : D \mapsto A$. Note that, for simplicity, we assume that the principal does not directly control any policy item.

Assuming for instance four policy items, the following are possible organizational structures:

- $\{a_1 \leftarrow \{p_1, p_2, p_3, p_4\}\}$, i.e. one agent has control on all four policies
- $\{a_1 \leftarrow \{p_1\}, a_2 \leftarrow \{p_2\}, a_3 \leftarrow \{p_3\}, a_4 \leftarrow \{p_4\}\}$, i.e. four agents have each control on one policy
- $\{a_1 \leftarrow \{p_1, p_2\}, a_2 \leftarrow \{p_3, p_4\}\}$, i.e. two agents have each control on two policies
- $\{a_1 \leftarrow \{p_1\}, a_2 \leftarrow \{p_2, p_3, p_4\}\}$, i.e. two agents with "asymmetric" responsibilities: one has control on the first policy item and the other on the remaining three

Finally, the organizational structure may also be characterized by an agenda $\alpha = a_{i_1}, a_{i_2}, \ldots, a_{i_h}$, that is a permutation of the set of agents defining the sequence with which agents are called to decide upon the policy items under their control.

We suppose that the principal and the agents have each an idiosyncratic ordering over the entire space X of policy vectors that may or may not correspond to nature's ordering $\succ_N$. The principal is interested in the overall

[5] Actually we could also allow for some decision rights to be ambiguously allocated, so that two or more agents are entitled to modify the same policy. This phenomenon, which is often found in real organizations, can be easily modeled in our framework but we leave it to future investigation.

performance of the organization but may not know how to achieve this objective, i.e. her ordering over the policies space may differ from nature's. Analogously, each agent has an idiosyncratic ordering of policy vectors, which, in general, differs from nature's, the principal's and the other agents', reflecting the agent's idiosyncratic vision, interest, and cognition. Moreover, such ordering concerns the entire set of policies, not only those under the control of the agent himself. We call $\succeq_{\Pi}$ the principal's ordering and $\succeq_{a_i}$ the ordering held by agent i. We assume that all such individual orderings are complete and transitive, i.e. that if $x_i \succeq_k x_j$ and $x_j \succeq_k x_l$ then $x_i \succeq_k x_l$, where k may indicate the principal or any agent.

When asked to decide upon two alternative profiles for the policies under his control, an agent will choose the one that ranks higher in his own ordering, given the current state of the other policy items that are not under his control, unless the principal gives appropriate monetary incentives to override the agent's preference and induce him to make a different choice. For the sake of simplicity we make a simple linearity assumption and suppose that the incentives needed to induce an agent to accept a policy profile that ranks lower in his ordering is proportional to the difference of the rankings of the two alternatives. Suppose, for instance, that agent a_i has to choose between two policy vectors x_i and x_j (of course the vectors may differ only in items under the agent's control) that rank respectively $rank(x_i) = r_i$ and $rank(x_j) = r_j$ with $r_i < r_j$, indicating that he prefers x_i to x_j [6]. Of course the agent would choose vector x_i and if the principal wants to reverse the choice has to pay $c(r_j - r_i)$ where c is, for simplicity, constant and equal for all agents. We could interpret c as an extra monetary incentive the principal has to give to the agent in addition to the standard compensation needed to elicit a normal level of effort, which in turn may depend upon the agent's commitment, motivation and so on.

We suppose that at the outset an initial "*status quo*" policy vector is (randomly) given[7]. Then the first – according to the agenda – agent may modify the policies under his control. He generates all the sub-vectors for the policies under his control and chooses the one that, together with current policies that are not under his control, will determine the vector he prefers, unless payments from the principal induce him to make a different choice.

[6] We use the convention that if the agent strictly prefers x_i to x_j then $r_i < r_j$ and that the agent's mostly preferred policy x_0 has rank $rank(x_0) = 1$.

[7] In what follows we usually find properties for all possible initial policy vectors.

When the first agent in the agenda has taken a decision, the value he has chosen for the policies under his control become part of the new status quo. Then the same procedure is repeated for the second, third, ..., $h-th$ agents in the agenda. Once all agents have operated on the policies under their control, we may either assume that the procedure comes to a halt or that the agenda is repeated over and over again until an optimum or a cycle are encountered. A (local) optimum is a policy vector for which no agent finds it convenient or possible[8] to modify items under his control according to the procedure outlined so far. A cycle is instead a subset of policy vectors among which agents keep cycling.

In the sequel we will investigate both stopping rules. Of course if the agenda is repeated only once cycles are ruled out and the organization will reach a decision but, we will show, there will be in general many possible outcomes. On the contrary if the agenda can be indefinitely repeated until a cycle is encountered or a local optimum is reached, we will show that cycles are usually very likely, but when cycles are not encountered the number of possible local optima is very small.

In order to be more precise, we can characterize the properties of the paths in the space of policies that emerge out of the procedure informally outlined above by providing a few definitions.

Given an organizational structure $\mathbf{O} : D \mapsto A$, we say that the policy vector x is a **preferred neighbor** of vector x for agent a_k who has control of the set of policies d_k if the following three conditions hold:

1. $x \succeq_{a_k} y$
2. $p_\nu^x = p_\nu^y \ \forall \nu \notin d_k$
3. $x \neq y$

Conditions 2 and 3 require that the two vectors differ only by policy items under the control of agent a_k. According to the definition, a preferred neighbor can be reached through the decision of a single agent.

We call $H_k(x_i, a_k)$ the set of preferred neighbors of a vector x_i for agent a_k.

[8] Impossibility may derive from the rule that the agenda can be repeated only once and therefore after the $h-th$ agent in the agenda has selected his policy item the new status quo cannot be further modified, even if some agents would like to do so.

A **path** $P(x_i, \mathbf{O}, \alpha)$ from a vector x_i and for an organizational structure $\mathbf{O}$ and an agenda α is a sequence, starting from x_i, of preferred neighbors for the agents in the agenda:

$$P(x_i, \mathbf{O}, \alpha) = x_i, x_{i+1}, x_{i+2}, \ldots \text{ with } x_{i+m+1} \in H_{a_{i+m+1}}(x_{i+m}, a_{i+m+1} \in \alpha)$$

A vector x_j is **reachable** from another vector x_i and for the organizational structure $\mathbf{O}$ if there exist a path $P(x_i, \mathbf{O}, \alpha)$ such that $x_j \in P(x_i, \mathbf{O}, \alpha)$.

A path can end up either on a (local) optimum, i.e. a vector which does not have any preferred neighbor, or in a cycle among a set of vectors which are preferred neighbors to each other.

A vector x is a **local optimum** for the organization $\mathbf{O}$ if there does not exist a vector y such that $y \in H(x, a_k)$ for any agent a_k in the agenda.

A **cycle** is a set $X^0 = \{x_1^0, x_2^0, \ldots, x_j^0\}$ of policy vectors such that $x_1^0 \in H(x_j^0, a_{i_1})$, $x_j^0 \in H(x_{j-1}^0, a_{i_2})$, $\ldots$, $x_2^0 \in H(x_1^0, a_{i_l})$.

In the following sections we will show that paths and their outcomes, that is the (locally) optimal policy vector that is finally chosen, or the emergence of a cycle, can be highly manipulated by the principal either by changing the allocation of agents to different policies or by giving the appropriate monetary incentives. We will first examine the case in which the principal "knows what she wants" and does not modify her preferences. We will show that, in general, the principal may obtain policy vectors that are equal or very close to the ones she prefers at no or very small cost by appropriately modifying the allocation of decision rights. Incentives and organizational structure appear therefore as substitutes. Then we will consider the situation in which the principal "does not know what she wants", i.e. tries to learn from the environment which policy vectors perform better.

4 Getting what you want when you know what you want

Let us first examine the case in which the principal precisely knows the set of policies she wants to be implemented either because she has the right knowledge of the environment, i.e. her ordering over the space of policy vectors corresponds to their true performance value, or because she simply wants her preferred policy to be implemented, whatever the result.

The principal has two means of achieving this goal: she can act on the incentives and/or she can act on the organizational structure. In the former case the principal tries to align the agents's decisions to her preferences by giving the agents monetary incentives to do so. In the latter case the principal chooses an appropriate allocation of decision rights to the agents. Let us first show, by means of a few examples, that the principal can to a large extent manipulate the agent's decision and obtain a policy profile equal or very close to her preferred vector without providing extra incentives.

Consider first a very simple example in which 3 agents have a common most preferred choice, which is not the preferred option of the principal. Table 1 presents their individual preferences, ranked from the most to the least preferred outcome:

Order	**Agent1**	**Agent2**	**Agent3**	**Principal**
1st	011	011	011	000
2nd	111	000	010	101
3rd	000	001	100	111
4th	010	110	101	110
5th	100	010	000	100
6th	110	111	110	001
7th	101	101	111	010
8th	001	100	001	011

Table 1: An example of the emergence of different local optima

All the agents prefer vector $[0, 1, 1]$ to any other option, but this vector is the least preferred one by the principal. This looks indeed like a bad situation for the principal and apparently she could get better outcomes only by incurring high incentive costs, but at a closer scrutiny we notice that the principal can actually avoid such costs.

Consider for instance the organizational structure $\{a_1 \leftarrow \{p_1\}, a_2 \leftarrow \{p_2\}, a_3 \leftarrow \{p_3\}$, with agenda (a_1, a_2, a_3) and the initial status quo $[1, 1, 0]$. Agent 1 decides first and chooses to switch to 0 the policy p_1 under his control (because $[1, 1, 0] \prec_{a_1} [0, 1, 0]$), then agent 2 switches to 0 the policy p_2 under his control. The policy vector has now become $[0, 0, 0]$ and agent 3 will not further modify it because $[0, 0, 1] \prec_{a_3} [0, 0, 0]$, neither will agents 1 and 2: $[0, 0, 0]$ is a local optimum for this organizational structure and the principal can obtain it at no cost, even if it is dominated by another policy vector for all the agents.

Actually it is easy to verify that $[0, 0, 0]$ is the local optimum that the organization reaches from six out of eight initial conditions. Only for initial conditions $[0, 1, 1]$ and $[1, 1, 1]$ will the organization reach the other local optimum $[0, 1, 1]$, which is the most preferred one by all the agents.

The same result of two local optima $[0, 0, 0]$ and $[0, 1, 1]$ could be obtained for instance with the organizational structure $\{a_1 \leftarrow \{p_1, p_2\}, a_2 \leftarrow \{p_3\}\}$ and agenda (a_1, a_2). On the contrary, the organizational structures $\{a_1 \leftarrow \{p_1\}, a_2 \leftarrow \{p_2, p_3\}\}$ and, obviously, $\{a_i \leftarrow \{p_1, p_2, p_3\}\}$ $\forall i \in \{1, 2, 3\}$ possess the unique global optimum $[0, 1, 1]$.

Actually, stronger results can be shown. It is indeed possible to provide cases in which the same group of agents can generate different global optima (i.e. optima that are stably reached from any initial condition) or cycles, depending upon the organizational structure. One such example may be illustrated by table 2 that summarizes the preferences of three hypothetical agents:

Order	**Agent1**	**Agent2**	**Agent3**
1st	001	000	001
2nd	110	111	110
3rd	000	001	000
4th	010	010	010
5th	100	100	100
6th	011	011	011
7th	111	101	111
8th	101	110	101

Table 2: Emergence of cycles or different global optima

It is easy to verify that this triple of agents (note that agents 1 and 3 are identical) may either generate a cycle, or the vector [0,0,1] as unique global optimum or the vector [0,0,0] as another unique global optimum given three different organization structures. A principal could get one of these very different outcomes simply by changing the organizational structure.

Structure $\{a_1 \leftarrow \{p_1, p_2\}, a_2 \leftarrow \{p_3\}\}$ always generates the cycle $[0, 0, 1] \rightarrow [0, 0, 0] \rightarrow [1, 1, 0] \rightarrow [1, 1, 1] \rightarrow [0, 0, 1]$. It is therefore a structure in which intra-organizational conflict does never settle into an equilibrium, unless a stopping rule is provided. Structure $\{a_1 \leftarrow \{p_1\}, a_2 \leftarrow \{p_2\}, a_3 \leftarrow \{p_3\}\}$ has the unique global optimum [0,0,1] that is reached from every initial condition, whereas structure $\{a_1 \leftarrow \{p_1\}, a_2 \leftarrow \{p_2, p_3\}$ also produces a unique global optimum but a different one, i.e. vector [0,0,0].

We cannot here provide more general results, but in Marengo and Settepanella (2010) it is formally proven, by using some properties of the geometry of hyperplanes arrangements and in the slightly different context of social choice with majority voting, that any kind of cycle can always be broken by appropriate changes of what we call here organizational structure and necessary and sufficient conditions are given for any vector (e.g. the principal's most preferred policy profile) to be a global or local optimum for an appropriate organizational structure.

So far we have simply provided some examples crafted in such a way as to show the possibility of manipulation of the outcome of the organizational decision processes by differently allocating decision rights. One could wonder how general this results are and how this manipulation could complement or substitute the manipulation that may be achieved by incentives, i.e. by modifying the agents' choices through alteration of their payoff landscape.

In order to answer this question we investigate the general properties of random populations of agents and principals. We simulate randomly generated problems with $n = 8$ policy items and up to eight agents with randomly generated preferences. We test the following organizational structures with 1, 2, 4 and 8 agents[9]:

- **O1**: $a_1 \leftarrow \{1, 2, 3, 4, 5, 6, 7, 8\}$
- **O2**: $a_1 \leftarrow \{1, 2, 3, 4\}, a_2 \leftarrow \{5, 6, 7, 8\}$ with agenda $\alpha = a_1, a_2$
- **O4**: $a_1 \leftarrow \{1, 2\}, a_2 \leftarrow \{3, 4\}, a_3 \leftarrow \{5, 6\}, a_4 \leftarrow \{7, 8\}$ with agenda $\alpha = a_1, a_2, a_3, a_4$
- **O8**: $a_1 \leftarrow \{1\}, a_2 \leftarrow \{2\}, a_3 \leftarrow \{3\}, a_4 \leftarrow \{4\}, a_5 \leftarrow \{5\}, a_6 \leftarrow \{6\}, a_7 \leftarrow \{7\}, a_8 \leftarrow \{8\}$ with agenda $\alpha = a_1, a_2, \ldots, a_8$

In the sequel we will study the properties of decision making in randomly generated policy landscapes (that is nature's preferences). In each case we will study the outcome for every initial status quo and we will repeat the exercise for 1000 different randomly generated problems.

We first consider the case in which the agenda may be endlessly repeated until a local optimum or a cycle are encountered. Under such rule, cycles are

[9]When only a subset of the eight agents are employed, i.e. in all organizational structure but the one designated by **O8**, the assignment of agents to the elements of the decomposition is also made randomly.

very frequent when decision rights are highly partitioned as in organization **O8**, they become less frequent with coarser partitions and disappear when all decisions are delegated to a single agent. When cycles do not appear, the number of local optima is always small. Table 3 summarizes these results by presenting the average number of cycles (with standard deviations in brackets) and the share of initial conditions leading to a cycle obtained over 1000 different randomly generated problems for the four organizational structures. For instance, the first line tells that with organizational structure **O8**, 78% of the 256,000 simulated paths (256 initial status quo times 1000 repetitions with different randomly generated agents) lead to a cycle. When cycles are not encountered, paths may lead on average to 2.78 different local optima. Of course with structure **O1** simulated paths always end on the only agent's most preferred policy vector.

Org. Structure	**No. of local optima**	**Share of cycles**
O8	2.78 (1.22)	0.78
O4	1.89 (0.98)	0.74
O2	1.03 (0.45)	0.58
O1	1.00 (0.00)	0.00

Table 3: Number of local optima for different organizations
(n=8, 1000 repetitions, standard deviation in brackets)

Of course in all cases organizational outcomes are on average far both from the principal's most preferred and from the best performers according to nature's policy vectors, as we did not introduce any mechanisms for aligning them. If we introduce incentives it should be possible to align organizational outcomes with the principal's objectives. Indeed this obviously happens: as incentives grow, also control of organizational policy by the principal increases. Table 4 shows the increase of control for organizational structure **O8** as incentives increase. Control is measured by the average distance (in terms of difference between ranks) between the realized policy and the principal's most preferred one (0 meaning full control), whereas incentives are measured by the maximum sum the principal is willing to pay each agent for aligning his choice to her preferences (255 being the maximum amount for always inducing any agent having to choose between two policies to select the one preferred by the principal). Note that, in general, when decision rights are highly partitioned like in **O8**, full control cannot be achieved because of interdependencies (externalities) among agents: each agent can be induced to choose the policies the principal prefers but only within the policies under his control and given the current status quo of the policies outside his control. Because of externalities, this procedure might never generate, and therefore select, the policy vector the principal ranks highest.

Incentive level	**Average Control**	**No. of local optima**	**Share of cycles**
0	132.78 (21.22)	2.78 (1.22)	0.78
10	111.89 (16.98)	7.21 (2.98)	0.80
100	14.90 (8.45)	13.83 (3.26)	0.21
255	8.29 (1.78)	48.03 (4.17)	0.08

Table 4: Incentives, control, local optima and cycles for organization O8
(n=8, 1000 repetitions, standard deviation in brackets)

Table 4 also shows that stronger incentives have another interesting and non obvious effect: they sharply decrease the likelihood of cycles and on the other hand increase the number of local optima. Incentives tend to prevent cycles and at the same time they increase the manipulability of decisions: as the number of local optima increases, the principal may more easily induce agents to select autonomously a local optima close to her own most preferred policy vector.

Very similar results are obtained for organizational structures **O4** and **O2**, where with the highest incentives (255) average control is, respectively, 3.75 (standard deviation 0.85) and 0.47 (standard deviation 0.15), the number of local optima is 32.73 and 12.13 and the percentage of cycles is 4.7% and 1.7%. Organizational structure **O1** instead always presents only 1 optimum and no cycles. With strong incentives, full control (average 0.0 and standard deviation 0.0) is always achieved because with only one agent externalities do not exist.

These results have been obtained assuming that the agenda is repeated over and over until a local optimum or a cycle are reached, and we saw that cycles are in general very likely especially with highly decentralized structures. Of course an easy way to prevent cycles from occurring is to forbid the reiteration of the agenda: once all agents have taken their decision according to the order stated by the agenda the procedure comes to a halt and no further modifications to the policy vector are allowed.

This procedure produces very different results. Table 5 shows the results of simulations in which each agent is allowed to decide only once and in the order stated by the agenda and incentives are set to 0. Of course cycles cannot emerge with such a halting rule and simulations show that decision processes can end up in about 42 different policy vectors (not necessarily local optima, as the process is truncated) for organization **O8**, 28 for **O4**, 10 for **O2** and, obviously, only 1 for **O1**.

Org. Structure	**N. of different final policy vectors**
O8	41.93 (3.14)
O4	27.73 (2.45)
O2	10.30 (1.22)
O1	1 (0.0)

Table 5: Number of local optima without agenda reiteration and without incentives

(n=8, 1000 repetitions, standard deviation in brackets)

The table clearly shows the source of a possible "divide and conquer" strategy by the principal: by partitioning more finely decision rights and hiring more agents, each of them with responsibility on only very few policies, the principal can more easily and cheaply manipulate the organization's decision. The table shows the sharp increase in the number of local optima that can be obtained with more fine grained organizational structures and therefore higher possibility of finding a local optimum equal or close enough to the principal's most preferred policy profile. By exploiting this feature, the principal has the possibility of getting high levels of control and performance without using any extra monetary incentive. The following table 6 provides evidence in this direction. The table presents averages and standard deviations of the best control and performance achieved in each simulated problem. By best control we mean the difference between the rank, in the principal's preference ordering, of the finally implemented policy vector and the rank of the principal's most preferred vector (which is always 1, by construction). By best performance instead we mean the difference between the rank, in nature's preference ordering, of the finally implemented policy vector and the rank of nature's most preferred vector (which is always 1, by construction). In all the simulations summarized in the table incentives have been set to zero.

Org. Structure	Best control	Best performance
O8	4.71 (6.63)	4.82 (5.43)
O4	8.01 (10.43)	8.25 (8.48)
O2	21.47 (22.67)	22.08 (21.03)
O1	127.57 (75.98)	128.53 (75.14)

Table 6: Best control and best performance, without incentives
(n=8, 1000 repetitions, standard deviation in brackets)

Table 6 shows that high control and/or high performance *can* be achieved at zero incentives in organizational structures where decisions are highly partitioned, whereas if all decisions are delegated to one single agent best control and performance are random. Note that average control and average performance are the same (around 128.5, that is the median rank) for all structures, but best control and best performance are very different. This implies that, whereas in organization **O1** the principal can only use incentives in order to get high control and high performance, in organizations with finer partitions of decision rights, and in particular in **O8**, the principal has the possibility of achieving high control and performance by acting on the distribution of decisions and on the initial status quo, without providing any extra incentive. All in all, in this case organizational structure and incentives are largely substitutes.

5 Adaptively learning principal

Let us now turn to the more interesting and realistic case in which the principal does not know the "right" model of the world and is aware of such ignorance. She holds an ordering of the policy vectors that does not correspond to their true relative performance, i.e. the principal's ordering is different from nature's. Thus the principal tries to learn the correct ordering by a simple trial-and-error mechanism that will be explained below. When learning is in place, along with trying to have her preferred policies implemented, the principal also tries to sample the performance value of different policy vectors in order to adaptively learn from the environmental feed-backs and avoid lock-in into inferior policies. This determines a complex trade-off between aligning the agents' decisions to the principal's preferences or letting agents more free to choose policies according to their own idiosyncratic preferences. If, by means of appropriate incentive and/or organizational structures, the principal optimizes such alignment she will have her preferred policies efficiently implemented, but agents who may hold better models of the environment and could implement policies with higher performance may be forced into the straightjacket of the principal's vision. On the other hand, if the principal, by means of looser incentives and/or appropriate organizational structures, leaves higher freedom to the agents of implementing their own preferred policies, she may learn that some of the agents' ideas may actually perform better in the environment. However she may loose control of the organization and the latter may be finally oriented by some agents to serve their own views and interests.

In this section we examine this trade-off and analyze in particular how the choice of incentives and the choice of organizational structure interact together in striking a balance in this trade-off. We will assume a very simple learning mechanism for the principal: if at two successive moments in time t and $t+\tau$ two different policy vectors x_t and $x_{t+\tau}$ are implemented with $x_t \neq x_{t+\tau}$, the principal may check if their performance levels (nature's ordering) are in line with her preferences and swap their positions in her ranking if they are not. On the contrary we assume that agents do not learn and keep their preferences unchanged[10].

[10] An extension of the present model by allowing that also agents are exposed to environmental signal (possibly mediated by the principal) and adaptively learn will be the object of future research.

We measure learning with the dynamics of Spearman's rank correlation between nature's and the principal's orderings of policy vectors. A Spearman coefficient equal to 0 means that the two rankings have no correlation, a coefficient equal to 1 that their correlation is maximal.

Let us begin the analysis by assuming that the agenda can be reiterated until an optimum or a cycle are met. We noticed in the previous section that when the agenda is reiterated cycles are frequently encountered, especially in structures with high partitioning of decisions. When the organization enters in a cycle it is impossible to give a precise definition of learning, as the policy vector does not stabilize. Thus we will consider only the cases in which a local optimum, rather than a cycle, is the final outcome.

With this *caveat*, table 7 summarizes the main results, showing the final Spearman coefficient after decision and learning has taken place starting from every possible initial *status quo* vector. Since initially all agents, the principal and nature are randomly generated, Spearman's coefficient at the outset is very close to 0 (0.003 with standard deviation 0.063 in this bunch of simulations).

Org. Structure	Incentive Level	Final Spearman coeff.
O8	0	0.29 (0.26)
	10	0.31 (0.26)
	100	0.67 (0.08)
	255	0.43 (0.06)
O4	0	0.16 (0.15)
	10	0.17 (0.16)
	100	0.32 (0.08)
	255	0.14 (0.06)
O2	0	0.03 (0.07)
	10	0.04 (0.07)
	100	0.06 (0.06)
	255	0.01 (0.06)
O1	0	0.00 (0.06)
	10	0.00 (0.06)
	100	0.00 (0.06)
	255	0.00 (0.06)

Table 7: Learning, organization, and incentives. With agenda reiteration
(n=8, 1000 repetitions, standard deviation in brackets)

Table 7 shows two main results: first that learning is higher in the organizational structure **O8** and, second, that its relationship with incentives tends to be of an inverted U-shape kind. As to the former result, we noticed in the previous section that the number of local optima is highest with organizational structure **O8** and therefore also the sampling of different policy vectors is highest with such a structure. At the opposite side, with structure **O1** there is always only one optimum, regardless the level of incentives, and therefore there cannot be any sampling and any learning at all. It is worth stressing again that these results concern only cases in which an optimum is reached, and in organization **O8** most of the time a cycle is instead encountered.

As to the relationship with incentives, we observed in the previous section that under agenda reiteration stronger incentive produce more local optima and this is reflected by higher learning. However, when incentives are very high such an effect is offset by stronger control that induces agents to actually select only very few of such local optima. Thus an inverted U-shape of the function which maps the level of incentives into learning. Higher incentives also seem to make learning more predictable, as witnessed by the lower standard deviations.

Let us now turn to learning without agenda reiteration. We saw in the previous section that in such a setting the number of actually implemented policy vectors tends to be much higher than in the case with agenda reiteration, except for organizational structure **O1** in which there always is only one vector implemented. This translates into much higher learning than in the case with agenda reiteration, as shown by Table 8. The effect of incentives is instead analogous to the one of Table 7, though the maximum level of learning seems to be reached for lower levels of incentives.

Org. Structure	Incentive Level	Final Spearman coeff.
O8	0	0.79 (0.04)
	10	0.80 (0.4)
	100	0.73 (0.05)
	255	0.45 (0.06)
O4	0	0.44 (0.06)
	10	0.45 (0.06)
	100	0.36 (0.06)
	255	0.15 (0.06)
O2	0	0.08 (0.06)
	10	0.09 (0.06)
	100	0.03 (0.06)
	255	0.01 (0.06)
O1	0	0.00 (0.06)
	10	0.00 (0.06)
	100	0.00 (0.06)
	255	0.00 (0.06)

Table 8: Learning, organization, and incentives. Without agenda reiteration

(n=8, 1000 repetitions, standard deviation in brackets)

Table 7 and 8 show the average final outcomes of the learning processes of 1000 different randomly generated problems. If we observe a single learning process results are confirmed. The following figures 1 and 2 show one typical learning process for the different organizational structures respectively at 0 and maximum (255) levels of incentives. Figures show that the learning performance of structure **O8** is steadily higher. Similar dynamics appear at all incentive levels.

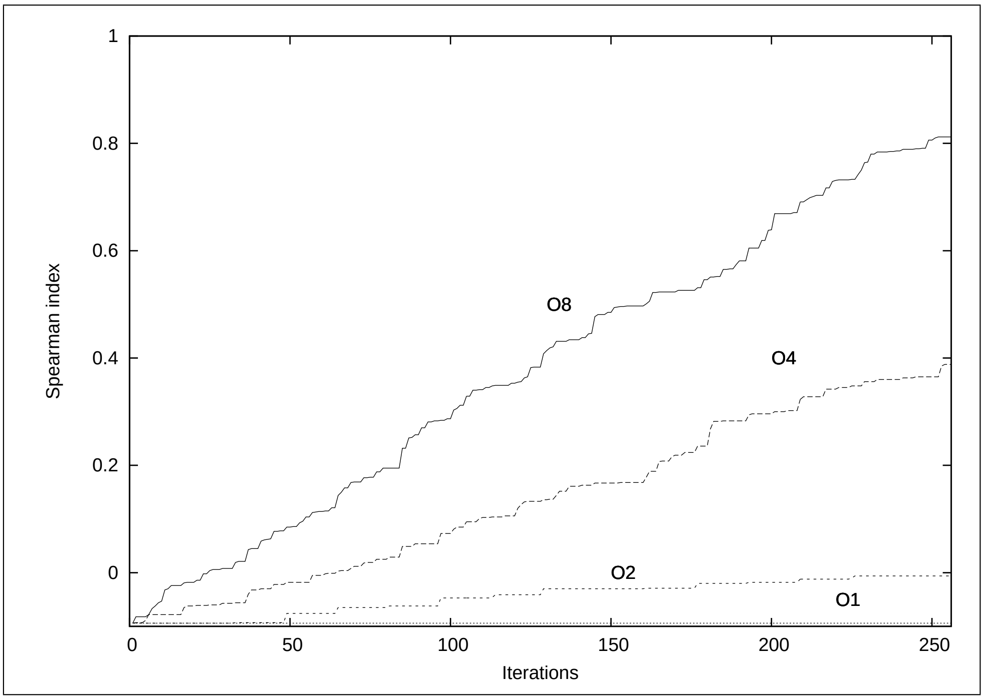

Figure 1: Learning and organizational structures with incentive level 0

To summarize the results obtained in this and the previous section, we could say that organizational structures in which the decisions are finely partitioned show an advantage for the principal in terms of higher manipulability and higher opportunities for achieving high levels of control and performance without relying too much on costly incentives. But such structures also present a dynamic advantage in terms of learning, especially when combined with a medium level of incentives.

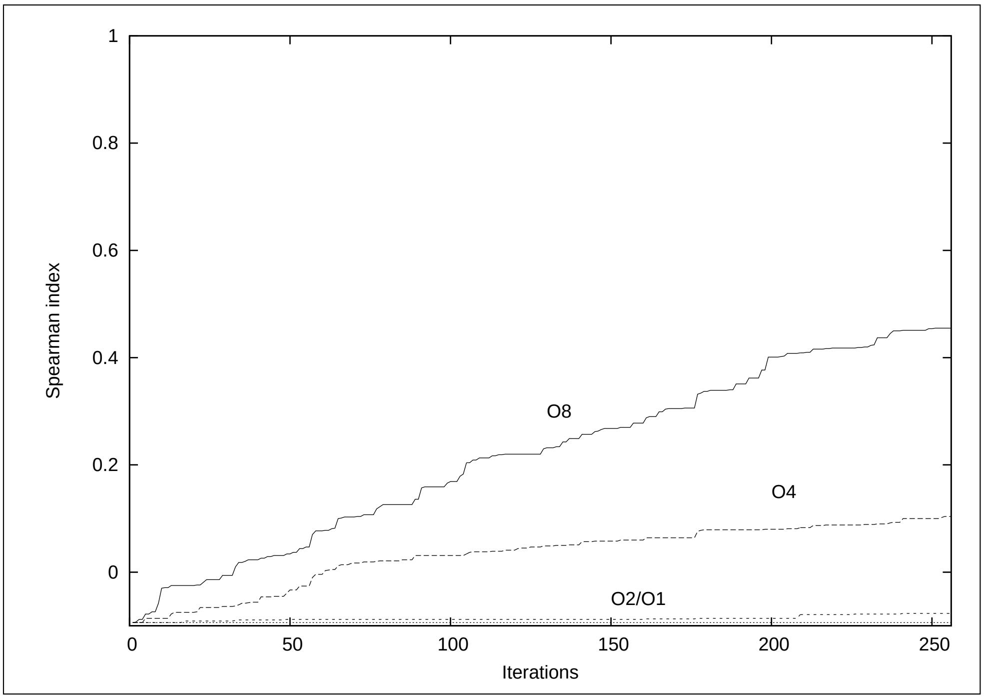

Figure 2: Learning and organizational structures with maximum incentive level

6 Externalities and the complexity of the organizational landscape

A recent stream of research has investigated how organizations can adapt and learn in complex environments in which the performance of the organization is the outcome of the interaction among organizational traits (Levinthal 1997, Levinthal and Warglien 1999, Gavetti and Levinthal 2000, Marengo and Dosi 2005, Rivkin and Siggelkow 2005). This literature has shown that when such interactions are widespread, non-linear, and imperfectly known, organizational processes of learning and adaptation take place in performance landscapes characterized by multiple local optima, and, therefore, by sub-optimality, path-dependency, and high sensitivity to small environmental perturbations. However this literature assumes that such interdependencies are exogenously determined by the nature of the "problem" the organization faces or of the"technology" (in the broad sense) it employs.

In this paper we have added to this "cognitive" source of complexity a political one, that is the complexity arising from the interdependencies among agents. In our model there is a sort of internal and political organizational landscape, whose ruggedness or smoothness is quite independent from the ruggedness or smoothness of the exogenous performance landscape the organization faces.

So far we have reported results of simulations in which agents are all randomly generated without any restriction but transitivity of their individual preferences. This determines environments of maximal complexity, in the sense that externalities tend to involve all policies and all agents. In this section we briefly analyze how different organizational structures and incentives perform when such complexity varies in intensity [11].

We have already analyzed in the previous sections how interdependencies among agents generate intra-organizational decision landscapes with multiple local optima and/or intransitive cycles. If on the other hand such externalities do not exist, cycles and multiple local optima do not appear and control by the principal becomes easier and cheaper.

Suppose for instance that in organizational structure **O8** each agent is concerned only with the policy under his control: agent a_i has random preferences between 0 and 1 for policy p_i (for every $i = 1, 2, \ldots, 8$) but is indifferent on the choices of the other policies. If no incentives are provided the organization settles into the unique optimum in which each agent chooses his own preferred value for the policy under his control, but at a minimum cost the principal can induce each agent whose choice for policy p_i differs from the principal preferred value to switch to the other value. At an average cost of $\frac{c \cdot n}{2}$ the principal can obtain her own most preferred policy vector as the unique organizational outcome.

A similar result, although with higher incentive costs, can be found in all cases in which the distribution of externalities and the distribution of decision rights coincide: if an agent is concerned with a subset of policies he should be allocated decision rights on those policies in order to minimize the cost of control.

However, whenever externalities and decision rights are perfectly aligned and incentives are set to optimize control, the organization can experience only a unique optimal outcome. Thus learning – in our strictly adaptive model – becomes impossible. In order to allow learning to take place the prin-

[11] Indeed there exist already examples of models that concentrate on the complexity of the intra-organizational decision making processes, see for instance Burton and Obel (1980) or Radner (1993)

cipal must set incentives to a lower level (thus losing some control) and/or choosing an allocation of decision rights that is finer than the scope of externalities[12]. Just to give an example, if each agent is concerned with two policy issues but is allocated only one of them (thus generating externalities between couples of agents), a final Spearman coefficient of 0.59 is on average achieved between the principal's and nature's preferences at zero incentives, and 0.46 with incentives higher or equal to $c \cdot 2$.

Thus, in our model, the organizational design principle of internalizing externalities which is one of the main prescriptions of transaction costs economics, is indeed justified in terms of control optimization [13] but not in terms of adaptive learning. In order to increase the level of exploration and foster adaptive learning, externalities should not be entirely internalized within separated decision units.

7 Conclusions and directions for further research

In this paper we have introduced a model that studies the interplay between learning, incentives and allocation of decision rights (the organizational structure) in a generalized agency problem whereby principals and agents have diverging views of the right courses of action for the organization, rather than simply conflicting interests.

Our main results could be summarized as follows. When learning is not at stake, incentives and organizational structure are substitutes. Diverging views among the principal and the agents may be to a large extent diluted by careful organizational design and incentives may be used as secondary devices. Somehow our model tends to support the idea that rules and organization may be more important than incentives in order to align individual behaviors to a common goal.

When instead learning is at stake, organizational structure and incentives may complement each other and have to be fine tuned according to the complexity of the learning process and the competitive pressure which is put

[12] A similar argument can be found in Cohen (1984) who argues that some degree of conflict may be a fundamental source for organizational learning.

[13] We have already remarked, however, that if externalities are very diffused, e.g. every agent's utility depends on all policies, achieving perfect control may require very costly incentives.

on fast or slow learning.

The model is rather rich and only a subset of possible research questions have been examined in the present paper. Among the possible lines of further research is the introduction of some learning process also for the agents, possibly with partial environmental feed-back only on the policies under their control. One should also consider the costs of hiring agents that are likely to depend on their span of control. Agents that are given responsibility of larger sets of policies are likely to be more costly, whereas in the present paper such costs have not been considered.

Finally, it would be interesting to model the organizational structure itself as subject to learning. The allocation of decision rights could be modified adaptively, for instance by taking one policy item out of the control of one agent and giving it to the control of another randomly selected agent. This would introduce a new learning process, certainly slower (the space of organizational structures is larger than the space of policies) but that could interact in non trivial ways with the learning of policy profiles. This will be the subject of future work.

References

Bernholz, P. (1997): "Property rights, contracts, cyclical social preferences and the Coase theorem: A synthesis," *European Journal of Political Economy*, 13, 419–442.

Bester, H. (2009): "Externalities, Communication and the Allocation of Decision Rights," *Economic Theory*, 41, 269–296.

Burgelman, R. A. (1994): "Fading Memories: A Process Theory of Strategic Business Exit in Dynamic Environments," *Administrative Science Quarterly*, 39, 24–57.

Burton, R., and B. Obel (1980): "A Computer Simulation Test of the M-form Hypothesis," *Administrative Science Quarterly*, 25, 457–466.

Cohen, M. (1984): "Conflict and Complexity: Goal Diversity and Organizational Search Effectiveness," *American Political Science Review*, 78, 435–454.

Coriat, B., and G. Dosi (1998): "Learning how to govern and learning how to solve problems," in *The Dynamic Firm. The role of technology, strategy, organization and regions*, ed. by A. D. Chandler, P. Hagstrom, and O. Solvell. Oxford University Press, Oxford.

Dosi, G., M. Faillo, and L. Marengo (2008): "Organizational Capabilities, Patterns of Knowledge Accumulation and Governance Structures in Business Firms: An Introduction," *Organization Studies*, 29, 1165–1185.

Dosi, G., D. Levinthal, and L. Marengo (2003): "Bridging contested terrain: Linking incentive-based and learning perspectives on organizational evolution," *Industrial and Corporate Change*, 12, 413–436.

Fama, E. (1980): "Agency problem and the theory of the firm," *Journal of Political Economy*, 88, 288–307.

Fehr, E., and U. Fischbacher (2002): "Why Social Preferences Matter. The Impact of Non-selfish Motives on Competition, Cooperation and Incentives," *Economic Journal*, 112, C1–C33.

Fehr, E., and S. Gächter (2002): "Do Incentive Contracts Undermine Voluntary Cooperation?," Working paper 34", Institute for Empirical Research in Economics, University of Zurich, www.iew.uzh.ch/wp/iewwp034.pdf.

Frey, B., and R. Jegen (2001): "Motivation crowding theory," *Journal of Economic Surveys*, 15, 589–611.

Gavetti, G., and D. Levinthal (2000): "Looking Forward and Looking Backward: Cognitive and Experimental Search," *Administrative Science Quarterly*, 45, 113–137.

Hart, O., and J. Moore (1990): "Property Rights and the Nature of the Firm," *Journal of Political Economy*, 98, 1119–1158.

Hrebiniak, L. G., and W. F. Joyce (1984): *Implementing Strategy*. MacMillan, London.

Jensen, M. C., and W. H. Meckling (1976): "Theory of the firm: Managerial behavior, agency costs, and the ownership structure," *Journal of Financial Economics*, 3, 305–360.

Kaplan, S., and R. Henderson (2005): "Inertia and Incentives: Bridging Organizational Economics and Organizational Theory," *Organization Science*, 16, 509–521.

Langlois, R., and N. Foss (1999): "Capabilities and Governance: the Rebirth of Production in the Theory of Economic Organization," *Kyklos*, 52, 201–218.

Levinthal, D. (1997): "Adaptation on rugged landscapes," *Management Science*, 43, 934–950.

Levinthal, D., and M. Warglien (1999): "Landscape design: Designing for local action in complex worlds," *Organization Science*, 10, 342–357.

March, J. G. (1991): "Exploration and Exploitation in Organizational Learning," *Organization Science*, 2, 71–87.

Marengo, L., and G. Dosi (2005): "Division of Labor, Organizational Coordination and Market Mechanisms in Collective Problem-Solving," *Journal of Economic Behavior and Organization*, 58, 303–326.

Marengo, L., and S. Settepanella (2010): "Social choice on complex objects," Wp 2010/02, LEM, Scuola Superiore Sant'Anna, www.lem.sssup.it/WPLem/files/2010-02.pdf.

Milgrom, P., and J. Roberts (1992): *Economics, Organization and Management*. Prentice Hall, Englewood Cliffs, NJ.

Mintzberg, H. (1973): *The Nature of Managerial Work*. Harper and Row, New York.

Nickerson, J. A., and T. R. Zenger (2004): "A knowledge based theory of the firm – The problem-solving perspective," *Organization Science*, 15, 617–632.

Page, S. E. (1996): "Two measures of difficulty," *Economic Theory*, 8, 321–346.

Prendergast, C. (2009): "Contracts and conflict in organizations," unpublished manuscript, GSB, University of Chicago, http://faculty.chicagobooth.edu/canice.prendergast/research/ContractsAndConflictsInOrgsUnpublished.pdf.

RADNER, R. (1993): "The Organization of Decentralized Information Processing," *Econometrica*, 61, 1109–1146.

RAJAN, R. G., AND L. ZINGALES (1998): "Power in a Theory of the Firm," *Quarterly Journal of Economics*, 113, 387–432.

RIVKIN, J., AND N. SIGGELKOW (2005): "Speed and Search: Designing Organizations for Turbulence and Complexity," *Organization Science*, 16, 101–122.

SAH, R., AND J. STIGLITZ (1986): "The Architecture of Economic Systems: Hierarchies and Polyarchies," *American Economic Review*, 76, 716–727.

SIMON, H. A. (1981): *The sciences of the artificial.* MIT Press, Cambridge, MA, 2nd edn.