

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

Luigi Marengo

Scuola Superiore S. Anna, Pisa, Italy, l.marengo@sssup.it

Corrado Pasquali

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

Abstract

We present a computational model of the interaction between incentives and the allocation of decision rights in an organization. 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.

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.

1 Introduction

One of the main differences between organizational economics and the evolutionary and capabilities approaches concerns the emphasis that is put on organizational change. Organizational economics is mainly concerned with the problem of efficient allocation of given resources and given capabilities. Key research questions within this perspective are the design of optimal incentives 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 rights, decision rights and exclusion rights (Bester 2009, Rajan and Zingales 1998). Evolutionary and capabilities based approaches¹ on the contrary are especially concerned with processes of learning and adaptation. The key research questions concern how capabilities are acquired and modified and in particular what are the organizational structures and processes that favor the generation of organizational capabilities.

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 of neoclassical economics that the role of institutions and organizations is basically to set the right incentives in order to align the individual pursuit of self interest with the promotion of collective efficiency and welfare. The capabilities view has on the contrary downplayed the role of incentives, often making the implicit assumption that individual motivation plays little or no role in the generation and accumulation of capabilities. Last but not least, the two tradition 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 more based on appreciative theory and assumes individuals with strong bounds in rationality and knowledge. Moreover sometimes it seems to depart from the methodological individualism of economic theory and assume a separate organizational dimension that cannot be reduced to

¹There 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.

the behavior of its member (more on the relations between the two theories in, e.g., Foss and Foss (2000)). Some existing attempts of bridging the gaps of the two streams of research have been made (Foss 2005, Dosi, Levinthal, and Marengo 2003, Coriat and Dosi 1998), 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 the capabilities literature does not deal in a satisfactory way of the role of incentives, delegation and power in the creation and modification of capabilities.

In this paper we make a novel attempt at bridging this gap that, we believe, makes some non trivial steps forward. 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 re-allocation 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 the agents' alternative visions that may prove very useful for collective learning. This trade-off, which is nothing but an instance of the widely discussed exploitation vs. exploration trade-off, requires a careful tuning of division of decision rights and incentives.

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 in the following section 4 we study its main properties. Results are presented in sections 5 and 6 where we discuss the behavior of the model respectively when the agent knows precisely

what she wants to get from the agents and when instead she tries to learn what are the best courses of actions. 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.

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 organizations has probably experienced such conflict: people simply have different ideas about what should be done and how it should be done and very often such different ideas can only partly, or not all, be ascribed to their self interest. Sometimes they hold diverging and motivationally strong view for the mere fact that they sincerely believe that their suggested course of action is good for the organization and attach high value to this belief.

Conflict arising from diverging interests and conflict arising from diverging vies 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.

Whereas agency conflict is inevitably a source of inefficiency and incentives are needed in order to correct for misalignment of objectives, possibly together with other devices that act rather upon intrinsic motivation, when instead conflict arises because of different views of what must be done, aligning them by providing the adequate incentives may prove both more difficult and less efficient.

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 situation in which non-routine courses of action must be envisaged and what must be done is not evident and transparent. 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' visions. In this respect, the standard solution to the problem of conflict suggested by agency model, that is aligning the agent's preferences as much as possible with the principal's, may actually prove detrimental and curb this important source of learning through diversity.

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 to be the right policies.

In the business strategy literature, the former 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 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-off under more rigorous terms.

3 The model

We consider a firm that has to take decisions on a set of n policies $P = \{p_1, p_2, \dots, p_n\}$. For simplicity we assume that each policy may take only two values $p_i \in \{0, 1\}$ and therefore the set of policies is formed by the 2^n vectors of n binary elements. We will call X this set of 2^n policy vectors and $x_i = [p_1^i, p_2^i, \dots, p_n^i]$ 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 level F , that we normalize in the interval $[0, 1]$: $F : X \mapsto [0, 1]$. This mapping determines a policy landscape whose ruggedness 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). In the analysis and the simulation exercises below we will consider a generic random assignment of performance value to each policy vector, thus assuming that the policy landscape has maximum ruggedness and complexity.

Our organization is composed by a principal Π and a number of agents that may range from 1 up to n . Such agents are attributed decision rights over a subset of policies. Let $A = \{a_1, a_2, \dots, a_h\}$ with $1 \leq h \leq n$ be a set of agents and let each agent be defined by a non empty subset of policies under his control. More precisely, let $d_i \subseteq P$ a generic non empty subset of the set of policies, we call a decomposition of decision rights a partition ² of the set of policies, i.e. a set of non-empty subsets $D = \{d_1, d_2, \dots, d_k\}$ such that:

$$\bigcup_{i=1}^h d_i = X \text{ and } d_i \cap d_j = \emptyset \forall i \neq j$$

We call organizational structure a mapping of the set D onto the set A of

²Actually we could also assume that some decision rights are ambiguously allocated and 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 further investigation.

agents, i.e. a mapping that assigns each subset of policies to one and only one agent. Note that, for simplicity, we assume that the principal does not directly control any policy item.

Examples of organizational structures in our framework and assuming 4 policy items are:

- $\{a_1 \leftarrow \{p_1, p_2, p_3, p_4\}\}$, 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, 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 space of policy vectors. These different orderings reflect both their different visions of how things should be done for the organization’s best and their own personal interest. The principal is interested in the overall performance of the organization but in general does not know how to achieve this objective, i.e. her ordering over the policies space is different from nature’s. Also the agents have different orderings from nature’s, the principal’s and the other agents’. This may reflect both their ignorance, their different visions and their different interests. We will call \succeq_{Π} the principal’s ordering and \succeq_{a_i} the ordering held by agent i . We will assume that all such orderings are complete and transitive, i.e. that if $x_i \succeq_k x_j$ and $x_j \succeq_k x_l$ then $x_i \succeq_k x_l$, with $k \in \{\Pi, a_1, a_2, \dots, a_h\}$.

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 preferences, 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 preference 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 ³. 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 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) given⁴. Then the first – according to the agenda – agent may modify the policies under his control. He randomly generates new policy sub-vectors and chooses the one that, together with current status quo 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.

When all instantiations have been examined for the first agent in the agenda and he 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, the agenda is repeated until an optimum or a cycle are encountered. A (local) optimum is a policy vector for which no agent finds it convenient to modify items under his control according to the procedure outlined so far. A cycle is instead a subset of policy vectors among which the chosen policies keep circling.

In the following section we will present an computational model of this procedure and show that the final outcome, 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

³We 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 $rank(x_0) = 1$.

⁴In what follows we often find properties for all possible initial vectors.

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. We will also show that cycles may be avoided by appropriately allocating decision rights and that for the principal it is more likely and easy to obtain a desired outcome if she relies on a finer division of decision labor, i.e. if she assigns fewer policy items to each agent.

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 Paths in the policy space

In this section we analyze the properties of the paths in the space of policies that emerge out of the procedure informally outlined in the previous section and show how the outcome of these paths may be manipulated by the principal through appropriate modification of the allocation of decision among given agents. For simplicity we will begin by assuming away monetary incentives and concentrate only on the organizational structure.

Given a decomposition of decision rights $D = \{d_1, d_2, \dots, d_k\}$, we say that the policy vector x_i is a **preferred neighbor** of vector x_j with respect to agent a_h who has control of the set of policies d_h if the following three conditions hold:

1. $x_i \succeq_{a_h} x_j$
2. $x'_i = x'_j \forall \nu \notin d_h$
3. $x_i \neq x_j$

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

We call $H_i(x, a_i)$ the set of preferred neighbors of an outcome x for agent a_i .

A **path** $P(x_i, D, \alpha)$ from an outcome x_i , a decomposition D and an agenda α is a sequence, starting from x_i , of preferred neighbors for the agents in :

$$P(x_i, D, \alpha) = x_i, x_{i+1}, x_{i+2}, \dots \text{ with } x_{i+m+1} \in H(x_{i+m}, D, a_{i+m+1} \in \alpha)$$

A vector x_j is **reachable** from another vector x_i and for the decomposition D if there exist a path $P(x_i, D, \alpha)$ such that $x_j \in P(x_i, D, \alpha)$.

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.

The set of **best neighbors** $B_i(x, a_i) \subseteq H_i(x, a_i)$ of a vector x for agent a_i is the set of the most preferred outcomes in the set of neighbors for the agent:

$$B_i(x, a_i) = \{y \in H_i(x, a_i) \text{ such that } y \succ_{a_i} z \forall z \in H_i(x, a_i)\}$$

By extension from a single agent to the entire organization, we can give the following definition of the set of neighbors for an organization as:

$$H(x, A) = \bigcup_{i=1}^k H_i(x, a_i)$$

An outcome x is a **local optimum** for the decomposition D and the set of agents A if there does not exist an outcome y such that $y \in H(x, A)$.

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

5 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 vision to be implemented, whatever the result.

The principal has two means of achieving this goal, she can act on the incentives 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 instead an appropriate allocation of decision rights to the managers. Let us first show, by means of a few examples, that the principal can in principle manipulate the agent's decision to a large extent 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. The following table 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

Example I: the emergence of a local optima

All the agents prefer vector $[0, 1, 1]$ to any other option, but this vector is the least preferred by the principal. This looks indeed 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 $[1, 1, 0] \preceq_{a_1} [0, 1, 0]$), then agent 2 will switch 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 $[0, 0, 1] \preceq_{a_3} [0, 0, 0]$, 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.

The same result could be obtained for instance with the organizational structure $\{a_1 \leftarrow \{p_1, p_2\}, a_3 \leftarrow \{p_3\}\}$. However, of course, with other organizational structures, for instance by keeping all 3 items together and assigning them to the one of the agents, this result could not be achieved and the principal's least preferred policy would always prevail unless further monetary incentives are put in place.

Actually the results may be stronger than this. 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 with the following table 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

Example II: 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 structure. 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 $[001] \rightarrow [000] \rightarrow [110] \rightarrow [111] \rightarrow [001]$. It is therefore a structure in which intra-organizational conflict does never settle into an equilibrium. Structure $\{a_1 \leftarrow \{p_1\}, a_2 \leftarrow \{p_2\}, a_3 \leftarrow \{p_3\}\}$ has the unique global optimum $[001]$ 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 $[000]$.

We cannot here provide more general results, but in Marengo and Settepanella (2008) 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 simply shown some example crafted in such a way as to show the possibility of manipulation of the outcome of the organizational decision processes by allocating differently 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 investigated the general properties of random populations of agents and principals. We simulated randomly generated problems with $n = 8$ policy items and up to 8 agents with randomly generated preferences. We have tested the following organizational structures with 1, 2, 4 and 8 agents⁵:

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

The following table presents the average number of local optima obtained over 1000 different randomly generated problems (with standard deviations in brackets):

Org. Structure	N. of local optima
O1	1 (0.0)
O2	10.4 (1.3)
O4	18.1 (2.7)
O8	24.7 (4.1)

Number of local optima for different organizations

(n=8, 1000 repetitions)

⁵When only a subset of the 8 agents were employed, i.e. in all organizational structure but the one designated by O8, the assignment of agents to the elements of the decomposition was also made randomly.

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 the higher possibility of finding a local optimum equal or close enough to the principal's most preferred policy profile.

So far we have shown how the principal's objective may be obtained by appropriate design of the organizational structure. The same result could be also obtained by monetary incentives, that is by paying the cost $c(r_j - r_i)$ in order to induce the agent to choose a policy profile against his own preferences. The two are clearly substitutes as shown by the following table that summarizes the results of a series of simulations. The following table reports the total costs incurred by a randomly generated principal in order to obtain the implementation of her most preferred policy vector for the four different organizational structures and for every possible initial condition. It is shown that in more fine grained structures the principal can achieve his most preferred vector at a much lower cost, because she can take advantage of the agent freely choosing a local optimum close enough to her preferred policy and having to pay only the extra cost of moving from that local optimum to her most preferred vector.

Org. Structure	Average cost for principal
O1	717
O2	1051
O4	2403
O8	15071

Average cost incurred by principal for preferred policy
(n=8, 1000 repetitions, all costs have to be multiplied by c)

All in all, the organizational structure is a substitute for incentives.

6 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. She holds an ordering of

the policy vectors that may not reflect their true relative performance and she is aware of her ignorance. Thus the principal, along with trying to have her preferred policies implemented, 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 preferences⁶. 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 implements looser incentives and organizational structures that divide less finely decision rights and leaves higher freedom to the agents to implement their own preferred policies, she may learn that some of the agents' ideas may actually perform better in the environment but on the other hand she may lose control of the organization and the latter may be finally oriented by some agents to serve their own interests.

In this section we examine this trade-off and analyze in particular how the choice of incentives on one hand and the choice of organizational structure on the other may strike different balances. We will assume a very simple learning mechanism for the principal: if at two successive moments in time t and $t + 1$ two different policy vectors x_t and x_{t+1} are implemented with $x_t \neq x_{t+1}$, the principal may check if their performance levels are in line with her preferences and swaps their positions in the ranking if they are not. On the contrary we assume that agents do not learn and keep their preferences unchanged⁷.

Let us begin by examining the role of the organizational structure in the absence of extra incentives. We concluded the previous section showing that in a static environment the principal may more cheaply obtain the desired policies by decomposing as much as possible decision rights. When instead

⁶This trade-off is a special case of the more general trade-off between exploitation and exploration in organizational learning (March 1991).

⁷An 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.

the principal also needs to learn the true structure of the environment such decomposition principle does not hold.

Having the desired policy strictly implemented, either by appropriate organizational design or by strong incentives prevents the testing of alternative policies and hinders learning. At the opposite extreme if the principal simply hires one agents and delegates to him all decisions without any decomposition, no learning will be possible as the agent will always implement his own preferred policy vector. Intermediate levels of decompositions of decision rights seem more appropriate for striking a good balance between the efficient implementation of the principal’s plans and learning and adaptation. This result is summarized in the following figure 1 where we compare the performance of the four different organizational structures already used in the previous section in randomly generated problems with learning principals and no extra incentives. The figure shows that structure **O1** does not allow any learning, structure **O8** allows very little, while better performances are achieved with structures **O4** and, especially, **O2**.




Figure 1: Learning and organizational structures

A similar trade-off emerges if the principal uses incentives rather than the organizational structure. Strong incentives do indeed induce agents to choose

the principal's preferred policy but do not allow the principal to test different policies and learn from the environment. Moreover, obviously, strong incentives are also more costly. On the contrary weak incentives leave too much freedom to agents and do not result in the implementation of the policy profiles the principal may have learnt to perform best. Intermediate levels of incentives are those more apt to promote learning and better performance.

More interesting is perhaps the case in which incentives and organizational structure may be dynamically combined together. In these cases the picture becomes more complicated, as the organizational structure also determines the speed and breadth of learning (Marengo and Dosi 2005). As already shown, finer decompositions of decision rights determine a sharp increase in the number of locally optima organizational equilibria. Using little monetary incentive the principal may easily induce the organization to climb quickly one of these local optima. However the organization is then locked-in into this local optimum and re-designs of the incentives and/or organizational structures are needed for un-locking. Lower levels of division of decision making labor on the contrary are less subject to lock-in and increase the number of policy profiles that are tested. This results in a much broader scope for learning, but such learning tend to be slow. Moreover incentives must be carefully designed, otherwise the organization will simply follow the agents' preferences.

The following propositions summarize the simulations results:

- in environments with low complexity (i.e. smooth performance landscapes with few peaks), fine decompositions of decision rights and medium or low powered incentives are more efficient, as the organization quickly climbs one of the few local optima whose location is easily learnt by the principal (because of low complexity);
- in environments with high complexity (i.e. rugged performance landscapes with many local peaks), we must distinguish cases in which the competitive pressure is low or high. With the former we indicate cases in which the organization may learn slowly and low short-term performance is tolerated provided it conduces to high future performance levels. By the latter instead we mean that low performance in the short term is promptly punished (possibly with bankruptcy):
 - with low competitive pressure, learning is fostered if decision rights are kept together under the control of one or very few agents and

incentives are medium powered. However learning tends to be slow;

- as competitive pressure increases and it becomes necessary to increase fast the performance, the division of decision rights must increase and incentives must become stronger.

All in all, competitive pressure tends to produce finer than optimal (in the long run) decompositions of decision rights and stronger than optimal (in the long run) incentives.

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 visions 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 visions 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 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 the subject of further work.

References

- 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.
- 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.
- FOSS, K., AND N. FOSS (2000): “Competence and Governance Perspectives: How Much Do They Differ? And How Does It Matter?,” in *Competence, Governance, and Entrepreneurship*, ed. by N. J. Foss, and V. Mahnke. Oxford University Press, Oxford.

- FOSS, N. J. (2005): “Cognition and Motivation in the Theory of the Firm,” *Journal des Economistes et des Etudes Humaines*, 13.
- 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.
- LEVINTHAL, D. (1997): “Adaptation on rugged landscapes,” *Management Science*, 43, 934–950.
- 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 (2008): “Social choice on complex objects: A geometric approach,” Wp 2008/28, LEM, Scuola Superiore Sant’Anna.
- 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.
- PAGE, S. E. (1996): “Two measures of difficulty,” *Economic Theory*, 8, 321–346.
- 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.