Elia Piludu
ID: 667432
Submitted: 13/02/2026

Economics of Management and Innovation

DOSI-VIRGILLITO Module

Call date: 2 feb 2026

Approach Preview

I based everything I wrote on the course material (and what it pointed to, if needed). For intra-text citations, I relied on the classic author-date format, and set a bibliography (at the end) to reference the precise articles/lectures I used while answering[1]. The lectures' bibliography voice and citations are of course less formal in their typing. Regarding the remaining (personal) editorial criterias I've followed, the main ones are:

- the double apex (") is used for direct quotes, which are considered terms, expressions, or sentences directly taken from a certain source and pasted as they were in it;
- the single apex (') may emphasize a content, flag a term as 'particular', or indicate an expression to which has been given a slightly (sometimes <u>very</u> slightly) different meaning (connotation) compared to its standard denotation;
- *Italics* formatting serves to emphasize further the term, expression or sentence to which it has been applied (merely aesthetic/preference purpose).

Besides this approach introduction, I tried to stick to an impersonal tone while writing, although from time to time it felt necessary to switch to the opposite. It may lower a bit the formality of the document in some of its parts, but I saw it as viable in the end and left it like that.

[1] In-text cits are also all bookmarked, so that clicking on one leads to the precise bibliography voice it refers to.

Exam Contents

Question 1:

Compare a case in which an industry is populated by homogenous firms versus a case in which an industry is populated by heterogeneous firms. Which are the implications of heterogeneity for industry evolution and performance?

Answer 1:

The comparison between an industry populated by *homogenous* firms and one characterized by *heterogeneous* agents highlights the fundamental divide between the standard neoclassical equilibrium paradigm and the evolutionary approach to economics (Dosi, 2023). These two perspectives rest on radically different micro-foundations regarding the nature of the firm, the characteristics of technological knowledge, and the mechanisms of competition.

In the standard neoclassical framework, often grounded in the metaphor of the "*representative agent*" (Dosi, 2023), firms are treated as *homogenous* entities. They are essentially identical processors of inputs into outputs, fully described by a shared production function and facing the same set of available technologies. In this context, technology is akin to a "public good" or a *book of "recipes*" (Dosi, 2023) freely available to all. If firms are homogenous, they possess perfect information (or at least identical probability distributions regarding uncertainty) and unlimited 'computational' capabilities to maximize profits. Consequently, competition is a static state rather than a dynamic process. Since all firms are identical in their efficiency and cost structures, there is no scope for a selection mechanism to operate. Any deviation from the equilibrium is corrected by instantaneous adjustments in prices and quantities, driving profits to zero in the long run. In this theoretical world, industry evolution can only be triggered by *exogenous* shocks, such as a shift in the production function that affects all agents simultaneously, leaving no room for the *endogenous* generation of variety or the messy dynamics observed in real markets.

On the other hand, the evolutionary perspective portrays industries as populated by strictly *heterogeneous* firms. Here, firms are viewed as behavioural entities characterized by 'bounded rationality' and distinct capabilities. They do not optimize on a known bunch of possibilities but search, learn, and adapt in complex, changing environments. Their heterogeneity stems from firm-specific '*organizational routines*'[2], the recurrent interaction patterns that allow organizations to solve problems, and specific technological competencies (Dosi, 2023). As highlighted by Costa et al. (2023), these internal capabilities act as the firm's '*genotype*', which then maps into the '*phenotype*' of revealed performance. Because knowledge is often tacit, cumulative, and path-dependent, firms cannot simply copy the best practice of competitors. This

[2] Even if present in Dosi's book (2023), it is a quote pointing to an outside work so I preferred to maintain it as a concept (single apex) rather than a direct quote from the document (double apex). I used the same reasoning in other instances throughout this paper.

results in persistent asymmetries in productivity, profitability, and growth rates, even among firms operating in the same sector[3].

The *implications* of heterogeneity for industry evolution are profound. Once we acknowledge that firms differ in their 'fitness', defined by their cost structures, product qualities, and innovative abilities, competition becomes a dynamic selection process, often termed '*Schumpeterian competition*'. In this view, the market acts as a selection environment that rewards relatively more efficient or innovative firms with higher profits and market shares, while penalizing laggards with shrinkage or exit (albeit often imperfectly[4]). This mechanism is formally captured by measuring "*differential growth*" (Dosi, 2023), where the rate of change of a firm's market share is a function of its fitness relative to the industry average. If a firm's fitness exceeds the weighted average fitness of the population, its share expands; if it falls short, it contracts[5].

Consequently, heterogeneity is the very fuel of industry evolution. Following Fisher's 'fundamental principle of natural selection', the rate of evolutionary change in a population is proportional to the variance of its fitness. If there were no heterogeneity (as in the homogenous case), there would be no selection-driven evolution. The presence of diversity allows the economic system to explore new technological and organizational trajectories. However, because firms constantly learn and generate new variety, this evolution is not a smooth convergence to an optimal steady state. Instead, it generates a state of 'restless capitalism' (Dosi, 2023 [6]), characterized by continuous 'instability', entry and exit of firms, and fluctuating aggregate variables.

Regarding performance, the implications of heterogeneity are visible in the statistical properties of industrial dynamics. Unlike the Gaussian distributions predicted by models of independent, random shocks to homogenous firms (Gibrat's Law), real-world data exhibit 'fat-tailed' distributions[7] of growth rates and highly skewed size distributions (Dosi, 2023). This suggests that 'winner-take-most' dynamics are at play, where a few high-capability firms, those possessing complex and interdependent routines, achieve superior performance and dominate the industry, while a large section of lower-efficiency firms struggles to survive (Costa et al., 2023).

Concluding the whole discussion, we can say that while the assumption of homogeneity offers analytical handling, it strips the competitive process of its essential dynamic nature. Admitting heterogeneity allows us to understand industry evolution not as a movement of a representative point on a static map, but as a complex transformation of a population driven by the interplay between the generation of variety (through innovation and learning) and the reduction of variety (through market selection; Dosi, 2023).

---

[3] Further empirical evidence is given by Table 6.2 (Dosi, 2023), where autoregressive coefficients ("AR(1)" of productivity) are close to unity, indicating that laggards tend to remain laggards.

[4] Selection is rarely absolute ('weak' force), often leading to the structural coexistence of efficient and inefficient firms; it has been better evidenced in answer 3.

[5] Just as a 'rough sum-up'.

[6] To clarify further footnote number 1: this was in reality coming from a Metcalfe, 1998 work, but as I explained earlier I didn't fetch in detail every single 'indirect' quote.

[7] Lectures recall "..at least exponential (Laplace) or even fatter in their tails.", and the reported graph ("figure 9.6; growth rates distributions in different years") exhibits in fact a 'tent-shaped' figure (lectures' slides).

Question 2:

What is endogenous organization learning? How can we measure it? How does it influence economic growth? Provide theoretical reasonings and empirical examples.

Answer 2:

Endogenous organizational learning is defined as a cumulative, path-dependent process through which business firms modify their internal 'organizational routines' (the recurrent interaction patterns that allow organizations to solve problems) and their shared 'mental models' of the environment (Dosi, 2023). Unlike standard neoclassical models where learning is often reduced to Bayesian information updating or the instant acquisition of a 'blueprint', endogenous learning arises from the condition of '*bounded rationality*': agents face both 'substantive uncertainty' (imperfect knowledge of the environment) and 'procedural uncertainty' (limitations in their own problem-solving competences). Consequently, learning is not a jump to an optimal solution but a local, step-by-step search process involving the modification of existing operating procedures and the exploration of new ones.

Theoretically, this process is effectively conceptualized through the lens of "*fitness landscapes*", a framework imported from evolutionary biology into organization science (Levinthal, 1997). In this model, an organization consists of 'N' distinct attributes or policy choices. These attributes interact in complex ways, termed "*epistatic interactions*" (Levinthal, 1997) to determine the organization's performance or 'fitness'. When these interactions are high (high complexity, denoted as K in NK models; lectures' slides), the landscape becomes 'uneven', characterized by multiple local peaks rather than a single global optimum. Organizational learning is thus viewed as an adaptive walk or "hill climbing" (Levinthal, 1997) where firms search their local neighborhood for improvements. This theoretical framework explains why learning is strictly path-dependent: a firm's current position on the landscape (its history) determines which peaks it can reach. Firms that are 'tightly coupled' (high internal interdependence) may find themselves trapped on local suboptimal peaks, unable to adapt without a radical and risky disintegration of their existing structures. Complementing the structural perspective of the NK model, Nelson and Winter in their '82 work formalize learning as an evolutionary search within the space of production techniques. In this model, firms operate with specific input-output coefficients (a_1 for labor, a_2 for capital; lectures' slides). Learning is modeled as a Markovian process where the firm's future state depends on its current technological capabilities. Crucially, the model distinguishes between two latent variables driving the search: *U*, representing the distance to the current best-practice frontier (efficiency), and *V*, representing technological opportunities (the potential to extend the frontier); while *U* is bounded, *V* is theoretically unbounded (lectures' slides). This formalization is critical because it explains why learning does not exhibit diminishing returns in the long run: as firms learn, they not only close the gap (*U*) but explicitly expand the horizon of possibilities (*V*). Furthermore, Dosi et al. (2010; from lectures' slides) operationalize the economic cost of these learning processes. In this framework, learning is not a passive by-product but the result of deliberate resource allocation. Firms split their R&D budgets between Innovation (*IN*) (searching for absolute novelty) and Imitation (*IM*) (adopting existing competitors' techniques). Access to new technologies is modeled as a stochastic process (a Bernoulli draw), where the probability of success depends on the magnitude of the investment (lectures' slides). This model adds the essential layer of strategic choice to the mechanics of learning.

The measurement of endogenous learning has historically relied on the detection of 'learning curves'. Empirical regularities, famously formalized as 'Wright's Law', show that unit costs or labor inputs tend to fall as a power function of cumulative production volume ($p = \alpha X^{\square}$; Dosi, 2023). This relationship captures the gains from 'learning-by-doing', where the very act of production generates improvements in efficiency. A classic historical example cited in the literature is the production of Liberty Ships during World War II, where the labor hours required to build a vessel dropped dramatically as shipyards accumulated experience (Dosi, 2023). However, modern measurement has moved beyond simple output proxies to direct assessments of organizational capabilities. For instance, recent studies utilizing the IMCPI survey data have been able to open the "black box" (Dosi, 2023) of the firm. By applying factor analysis to granular data on management practices, researchers can cluster firms based on the complexity of their routines. Evidence from the Italian manufacturing sector shows that firms mastering more complex, interdependent routines, defined as "Complex" firms, consistently exhibit superior productivity and employment growth compared to their "Essentials" counterparts (Costa et al., 2023).

The influence of this learning on economic growth is fundamental, going well over simple efficiency gains to become the primary driver of dynamic increasing returns. Unlike physical capital, which is typically subject to diminishing marginal returns, the accumulation of organizational knowledge is cumulative and path-dependent: the learning achieved today builds the cognitive foundations for the innovations of tomorrow. This self-reinforcing process generates 'autocatalytic' growth dynamics[8], where success in problem-solving breeds further competence (Dosi, 2023). Hence, endogenous learning prevents the economic system from becoming trapped in a stationary equilibrium. By continuously exploring the fitness landscape, firms discover and open new 'technological trajectories', effectively shifting the trade-offs defined by existing techniques. Consequently, organizational learning acts as the engine of secular economic development, ensuring that the economy remains a system characterized by the continuous introduction of novelty rather than convergence to a static optimum (Dosi, 2023).

[8] It means/implies a self-reinforcing process (current knowledge fuels further learning).

Question 3:

To what extent selection is a mechanism working in contemporary capitalism? Provide theoretical reasonings and empirical examples.

Answer 3:

The extent to which *selection* operates as a functional mechanism in contemporary capitalism is a subject of profound divergence between abstract evolutionary models and empirical reality. In theoretical terms, selection is the fundamental force that should 'cleanse' the market of inefficiency. Standard evolutionary models, such as the canonical 'replicator dynamics', postulate a rigorous environment where firms with higher 'fitness' (typically proxied by profitability or productivity) systematically expand their market shares, driving the exit of less efficient competitors (Dosi, 2023). Under this lens, one would expect a strong convergence toward best practices and a rapid punishment of underperformance. However, extensive empirical analysis reveals that market selection is, in reality, a remarkably "weak" force (Dosi, 2023), severely constrained by structural frictions and, among others, the internal *political* nature of business organizations[9].

To rigorously assess the *extent* of selection, the available literature utilizes a methodological framework known as '*evolutionary accounting*'. This technique decomposes aggregate productivity growth into two distinct components: the '*within*' effect, which captures the internal *learning* and innovation of individual firms, and the '*between*' effect, which captures the *reallocation* of market shares from low-productivity to high-productivity producers[10]. It is represented by the following (edited) equation (lectures' slides):

$$\Delta\Pi_t = \underbrace{\sum_i s_i(t-1)\Delta\Pi_i(t)}_{\text{Within Effect}} + \underbrace{\sum_i \Pi_i(t-1)\Delta s_i(t)}_{\text{Between Effect}} + \underbrace{\sum_e s_e(t)\Pi_e(t) - \sum_f s_f(t-1)\Pi_f(t-1)}_{\text{Entry \& Exit}}$$

If selection were the dominant engine of capitalism, the 'between' effect should account for the majority of industry growth. However, empirical evidence consistently shows the opposite: the 'within' component typically dwarfs the 'between' component, confirming that the displacement of inefficient firms is a slow-paced process (Dosi, 2023). This frailty of the selection mechanism is further evidenced by the persistence of 'fat-tailed' distributions of firm performance, as already evidenced in the first answer. In nearly every sector, highly efficient frontier firms coexist for decades with followers that are significantly less productive. If the market were a sharp selector, these massive differentials would be rapidly arbitraged away: their persistence proves that the market tolerates a significant degree of inefficiency.

[9] While the claim is generally verified (selection as a 'weak' force, preserving heterogeneity), Selection can be more active in specific contexts like industrial catch-up (see China case, Slide 66; lectures' slides), where the support of the distribution shrinks

[10] This is what actually represents the mechanism of 'selection' (Dosi, 2023).

The primary theoretical explanation for this 'selection failure' lies in the complex nature of the firm itself. As argued by Dosi et al. (2021), organizations cannot be reduced to simple efficiency-seeking devices. Instead, they are *dual-natured* entities: they are problem-solving systems based on 'organizational capabilities', but they are simultaneously 'structures of power'. The internal exercise of authority often prioritizes the reproduction of *hierarchical command structures* over pure productive efficiency. This occurs because the division of labor is frequently driven by a "divide-and-conquer principle" (Dosi et al. 2021), where decision rights are partitioned to maximize control rather than to align with the optimal distribution of knowledge.

This '*political*' dimension introduces severe inertia into the selection process through specific internal mechanisms. First, the exercise of "*veto power*" (Dosi et al. 2021) by managers, often used to block 'bottom-up' proposals from workers, stifles exploration and locks the organization into "suboptimal equilibria" (Dosi et al. 2021), preventing the adoption of fitter routines. Second, the hierarchy enforces *conformity*: managers tend to align their decisions with the principal's views to preserve their own "*span of control*" (Dosi et al. 2021), explicitly trading off organizational efficiency for '*obedience*'. Finally, hierarchies serve an '*extractive*' function: they are designed not just for coordination but to "rule over the extraction and appropriation of the net product" (Dosi et al. 2021). This allows top managers to engage in rent-seeking behaviors where remuneration is decoupled from productivity and grounded instead in shifting 'power relations'. In this context, the survival of an organization is effectively decoupled from its actual operational fitness.

Therefore, the complexity of the economic environment creates 'protective niches' that shield inefficient firms from direct competition (Dosi, 2023). In markets characterized by product differentiation, informational imperfections, and customer lock-in, competition rarely occurs on a single, transparent variable like price. A firm may survive despite high costs because it serves a specific local demand or because its consumers face high switching costs. Consequently, the correlation between a firm's productivity and its growth is positive but remarkably low, confirming that the market is a 'noisy' selector (Dosi, 2023).

Selection is then not a precise optimizer but a *rough* '*filter*'. It operates effectively only in cutting off the extreme 'lower tail' of absolutely non-viable firms, yet it fails to enforce the dominance of the best. The *result* is a system of '*restless capitalism*' where efficiency and inefficiency coexist, and where survival is determined by a complex blend of technological competence, institutional power, and market friction.

Question 4:

Are intellectual property rights beneficial? If you were a regulator, which type of policies would you introduce to regulate a sector highly exposed to appropriability? Motivate with historical and empirical examples.

Answer 4:

Asking whether Intellectual Property Rights (IPR) are *beneficial* does not lead to a simple, binary answer; rather, it represents a complex trade-off between the incentive to *innovate* and the *diffusion of knowledge*. The classic justification relies on the definition of information as a 'public good', non-rivalrous and non-excludable. In this view, without the monopoly granted by a patent, competitors would free-ride on the innovator's investment, leading to a market failure where the incentive to innovate vanishes (Arrow, 1962). However, from an evolutionary perspective, this argument is often overstated. Information is rarely a pure public good: it is rather a 'latent public good' that requires significant 'learning capacity' to be replicated (Dosi, 2023). Thus, the 'beneficial' nature of IPRs is highly *context-dependent*, varying dramatically across technological regimes.

Taking an empirical perspective, the role of patents differs by sector. As established by Pavitt (1984), 'supplier-dominated' firms rely largely on professional skills, trademarks, and advertising, while 'scale-intensive' firms rely on secrecy, know-how, and technical lags. Conversely, in 'science-based' sectors like pharmaceuticals, where knowledge is codified and imitation costs are low, patents are critical. However, recent empirical studies reveal a more strategic, and potentially troubling, use of patents in high-tech industries. In the semiconductor industry, Hall and Ziedonis (2001) observe a ('revisited', as they state) "patent paradox": firms have increased their patenting rates aggressively despite survey evidence suggesting that patents are ineffective at preventing imitation. In this sector, characterized by "cumulative innovation" (Hall & Ziedonis, 2001), patents are used less for appropriating returns on specific inventions and more as "bargaining chips" (Hall & Ziedonis, 2001) to negotiate cross-licensing agreements and prevent 'holdup' by rivals. While this legal thickening has spawned "patent portfolio races" (Hall & Ziedonis, 2001) among manufacturers, it has also facilitated the entry of specialized "design firms" (Hall & Ziedonis, 2001) that rely on patents to secure venture capital.

Therefore, for the "if you were a regulator" part of the question, if I had such a role tasked with overseeing a sector highly exposed to *appropriability* (specifically pointing to one involving *upstream* scientific knowledge, like biotechnology), I would introduce a policy framework strictly designed to prevent the privatization of the '*scientific commons*'. As argued by Orsi et al. (2006), the recent extension of patentability to "upstream research" results (such as gene sequences) breaks with the traditional economic rationale of the patent system. It shifts from protecting specific commercial applications to granting "prospecting rights" (Orsi et al., 2006) over basic knowledge, creating an "anticommons tragedy" where the fragmentation of rights blocks future research (Orsi et al., 2006).

My regulatory strategy would rest on three pillars. First, I would rigorously enforce the criterion of utility (or industrial application) to exclude 'upstream' discoveries from patentability. As noted by Orsi et al. (2006), the lowering of the utility threshold allows patents to become ‘hunting licenses’ for future applications rather than rewards for actual inventions. The granting of patents on the BRCA1 and BRCA2 genes to Myriad Genetics illustrates the danger: by patenting the gene itself rather than a specific therapeutic application, the firm was able to monopolize all diagnostic testing, effectively “stifling life-saving innovations further downstream” (a direct quote from Heller and Eisenberg, 1998, p. 698, in Orsi et al., 2006). A regulator must ensure that patents reward specific industrial applications, not the ‘basic tools’ of science.

Second, I would implement compulsory licensing regimes to prevent ‘holdup’ in cumulative technological trajectories. A warning from history is the Polaroid vs. Kodak (1986) case described by Hall and Ziedonis (2001). There, a patent injunction forced Kodak to shut down its instant-camera production entirely, costing billions and removing a competitor from the market . To avoid such ‘detrimental’ shutdowns in complex industries , regulators must ensure access to foundational technologies, unlike the aggressive ‘holdup’ strategies seen in the semiconductor sector.

Third, I would differentiate patent breadth based, among others, on the Pavitt Taxonomy (Pavitt, 1984). In sectors where innovation is cumulative and rapid (like software or semiconductors), broad patents lead to the “thicket” problems described by Hall and Ziedonis (2001). Here, I would favor narrow, shorter protection to encourage ‘*inventing around*’ and *rapid diffusion*, while reserving broader protection only for discrete, high-cost discoveries in sectors like the chemical/pharma ones.

So, to wrap up: IPRs are not an intrinsic moral right but a *utilitarian social contract*. They are beneficial only when they solve the appropriability problem without preventing the cumulative generation of knowledge.

Question 5:

Is the Marseille fish market coherent with the outcomes of neoclassical markets? Which are the similarities? Which are the differences? Which are the implications of the findings?

Answer 5:

The analysis of the *Marseille fish market* represents a canonical 'crucial experiment' in evolutionary economics, serving to test the validity of standard equilibrium theories in a real-world setting (Dosi, 2023). Theoretically, this market possesses all the structural characteristics of a perfect competition model: the goods are highly perishable[11] (fresh fish), ensuring that the market must clear daily; there are numerous buyers (500) and sellers (50; lectures' slides); the physical location is centralized, which should theoretically minimize information asymmetries. Under the strict neoclassical framework, specifically the Walrasian model, these conditions should guarantee the emergence of the 'Law of One Price' (Dosi, 2023). This law postulates that, in an efficient market, identical goods must sell for the same price at the same time, as arbitrage would rapidly eliminate any differentials. Therefore, the market should be characterized by '*anonymous*' trading, where the identity of the transacting partner is irrelevant to the price formation process (Dosi, 2023).

However, the empirical findings reveal that the functioning of the Marseille fish market is largely incoherent with these 'neoclassical predictions'. The most striking difference is the persistent and pervasive violation of the Law of One Price. As noted in the literature, the same type of fish is routinely sold at significantly different prices to different buyers at the exact same moment: prices were in fact "not posted", and "all transactions were pair-wise" (lectures' slides). That makes any assumption of perfect transparency fall, as the identity of the transacting partner becomes relevant to the price formation process. This price 'dispersion' is not a temporary disequilibrium phenomenon but a stable structural feature of the market, and Dosi (2023) further highlights that this dispersion arises because the market is then *not* structured around anonymous spot transactions, but rather around dense networks of specific "*trading relationships*" (Dosi, 2023).

This introduces the central concept differentiating the two views: *loyalty*.

In a 'neoclassical world', buyers should constantly scan the market to minimize costs, and switch sellers instantly if a lower price is available. In the Marseille reality, buyers exhibit strong loyalty to specific sellers, repeatedly transacting with them despite price differentials[12]. This behavior is rationalized not by forward-looking optimization, but by a mechanism of *reinforcement learning* (specifically the "logit rule"; lectures' slides). Rather than anticipating future value to secure supply, agents update the probability of visiting a seller based on *cumulated past profits*, simply repeating actions that proved successful in the past. The market is thus defined by "*interaction structures*" (Dosi, 2023) rather than a unified clearing mechanism.

[11] Fish naturally decay, that's why 'perishable'. It pertains to the category of goods which lose their value if not sold immediately (or within a limited time span).

[12] Note that this loyalty is not universal but, simplifying the (heterogeneous) picture, it follows more of a *bimodal* distribution: specifically, large and frequent buyers tend to be loyal to single sellers, whereas smaller, less frequent buyers tend to 'shop around' and visit multiple sellers (lectures' slides).

Despite these deep micro-level differences, there are similarities regarding aggregate outcomes. The market does, in fact, 'clear': the fish is sold, and waste is minimized, suggesting that the system is efficient in terms of allocation even if it fails the price convergence test (Dosi, 2023). However, the deduction of these aggregate regularities follows a distinct logic compared to neoclassical theory. While standard models assume that aggregate outcomes are simply the *sum* of representative, optimizing agents, the empirical evidence here reveals a clear contrast: at the individual level, the relationship between price and quantity appears almost *random*. It is only when these heterogeneous individual behaviors are *aggregated* that a smooth, monotonic demand curve appears. Consequently, while the 'law of supply and demand' holds as a macroscopic regularity, responding predictably to supply shocks, it must be understood strictly as an '*emergent property*' of the system's *aggregation*, rather than a behavior capable of being derived from the optimization calculus of any single individual (Dosi, 2023; lectures' slides).

The implications of these findings are far-reaching for economic theory. First, they dismantle the 'representative agent' hypothesis. The stability of the aggregate market does not require identical, maximizing individuals; rather, *order* is an 'emergent property' arising from the self-organization of heterogeneous agents (Dosi, 2023). Second, they demonstrate that markets are *social institutions* embedded in trust and loyalty, not merely frictionless mechanisms for price discovery. The Marseille fish market proves then that 'who you know' matters as much as 'what you pay', roughly speaking, and that directly challenges the core neoclassical assumption of *market anonymity*.

Bibliography

- Arrow, K. (1962). *Economic Welfare and the Allocation of Resources for Inventions*. In Nelson, R. (ed.), *The Rate and Direction of Inventive Activity*. Princeton University Press.
- Costa, S., De Santis, S., Dosi, G., Monducci, R., Sbardella, A., & Virgillito, M. E. (2023). *From organizational capabilities to corporate performances: at the roots of productivity slowdown*. Industrial and Corporate Change, 32(4), 1-28.
- Dosi, G. (2023). *The Foundations of Complex Evolving Economies. Part One: Innovation, Organization, and Industrial Dynamics*. Oxford University Press.
- Dosi, G., Marengo, L., & Virgillito, M. E. (2021). *Hierarchies, Knowledge, and Power Inside Organizations*. Strategy Science, 6(4), 371-384.
- Hall, B. H., & Ziedonis, R. H. (2001). *The patent paradox revisited: an empirical study of patenting in the US semiconductor industry, 1979-1995*. The RAND Journal of Economics, 32(1), 101-128.
- Levinthal, D. A. (1997). *Adaptation on Rugged Landscapes*. Management Science, 43(7), 934-950.
- Orsi, F., Sevilla, C., & Coriat, B. (2006). *Upstream patents and public health: The case of genetic testing for breast cancer*. In Mazzucato, M., & Dosi, G. (eds.), *Knowledge Accumulation and Industry Evolution: The Case of Pharma-Biotech*. Cambridge University Press.
- Pavitt, K. (1984). *Sectoral patterns of technical change: Towards a taxonomy and a theory*. Research Policy, 13(6), 343-373.
- Virgillito, M. E. (2025/26 academic year). *Economics and Management of Innovation: The Foundations of Complex Evolving Economies*. Lecture Slides (File: Slides .pdf).