

TSE DIGITAL CENTER POLICY PAPER SERIES

ECONOMICS OF PROCUREMENT AND ORGANIZATIONAL DESIGN: A REVIEW OF SELECTED LITERATURE

Matthieu Bouvard, Bruno Jullien and David Martimort

September 22, 2025

Policy Paper published as Working Paper No. 25-1667

Economics of Procurement and Organizational Design: A review of selected literature¹

Matthieu Bouvard, Bruno Jullien and David Martimort

Toulouse School of Economics, 22 September 2025

1. Introduction

Procurement theory and organization economic theory provide insights into how firms should structure their relationships with suppliers and their internal decision-making processes. This review synthesizes some key theoretical insights from influential papers on procurement theory and organizational design. The focus is on how information and incentives shape organizational structure – whether through integration, delegation of decision authority, hierarchical contracting, or monitoring arrangements. With a view on the development of Open RAN, we focus on aspects of the literature that we believe are the most useful for understanding these issues.

We group the contributions under thematic categories for clarity:

- *Integration and Communication,*
- *Centralization vs. Delegation of Authority,*
- *Hierarchies and Incentives in Organizations,*
- *Procurement with Competing Suppliers,*
- *Informational Alliances.*

We then provide two supplementary sections on related themes in organization theory for completeness: *Adaptive Organizations and Task Design*, and *Collusion and Delegated Monitoring*.

2. Integration and Communication

One classic motive for vertical integration is to improve information flow and coordination between stages of production. Arrow (1975) provides an early formal argument along these lines. In Arrow's model, a downstream firm faces uncertainty about its upstream input supply. Integrating an upstream supplier into the firm can alleviate this uncertainty by improving communication of supply information. A key insight is that even starting from a competitive benchmark, the need to share information under uncertainty can drive firms toward vertical integration, with the outcome often moving the market away from perfect competition. In other words, integrating an upstream supplier is valuable because it allows the downstream decision-maker to obtain timely and accurate information (for instance,

¹ This review has benefited from the financial support of Orange in the context of the research foundation TSE-Partnership

about input availability or costs) without the distortions that arise when separate firms interact through a market. Arrow's conclusion was that the information advantages can lead integrated firms to gain market power, introducing a tendency toward imperfect competition.

Gilbert and Riordan (1995) extend the integration argument to a regulatory context with complementary products. They consider an industry requiring two complementary inputs and ask whether a regulator should treat them as one integrated monopoly or as separate entities. Under bundled supply (one integrated firm producing both components), the regulator deals with one information source; under unbundled supply (two separate firms), each component is regulated independently. The authors find an important trade-off: unbundling two complementary goods introduces an additional informational cost analogous to the classic double-marginalization problem in vertical structures. With separate suppliers, the regulator must leave each firm an information rent (compensation for private cost information), much like two monopolists each adding a price markup. This can make separate regulations inefficient, even if it might allow competition in one of the components. Gilbert and Riordan (1995) show that unless unbundling yields substantial competitive benefits in at least one stage, the extra "information cost" can outweigh those benefits. Thus, an integrated structure tends to be optimal from the regulator's perspective when information asymmetry is significant. This resonates with Arrow's point: consolidation simplifies communication (here, between the firms and the regulator) and avoids duplication of informational inefficiencies. However, if de-integration can harness competition (e.g., one component is competitively supplied), it might justify the added information costs, an important consideration in regulated network industries.

In summary, early work indicates that integration can achieve better information sharing and coordination. Arrow (1975) highlighted improved internal communication as a rationale for integration, while Gilbert and Riordan (1995) quantified the cost of separated communication channels in regulation. Both emphasize that organizational boundaries affect how information is transmitted and used: integrated parties can eliminate information frictions that separate contracting would create.

3. Centralization vs. Delegation of Authority

The work reviewed in the previous section compares the integration of multiple activities into a single firm to their separation into distinct entities. A related line of research studies the organization of production within a corporation. Indeed, even when multiple activities are hosted under a common corporate roof, the question remains of how much autonomy should be given to these different entities by the headquarters. In other words, should decision-making be retained by the top management – the principal – or delegated to better informed subordinates – the agents. In the latter case, even an integrated firm may not function as an entity where all information is aggregated at the top – the implicit assumption in Gilbert and Riordan (1995).

As in the previous section, the central determinant of corporate structure is the distribution of information. But we will now focus on this distribution within the firm. For

example, agents that operate at a local level may have better information about specificities of their geographical area (consumer tastes, social norms, political interests). Similarly, agents performing technical tasks may better understand production constraints. In these different cases, a principal may ask those agents to communicate their information and then decide. However, as in Gilbert and Riordan (1995), truthful communication may not come for free. If the agent objectives are not aligned with the principal objectives and information is not easily verifiable the agent may distort communication to sway the decision in his favor. In that case, delegating the decision to a better-informed agent can be an efficient alternative to centralized decision-making.

Several models address this issue by considering communication incentives and contracting limits. Dessein (2002) compares communication or delegation when the agent is better informed but has biased preferences. Communication is not verifiable (cheap talk), while decision rights are contractible. If the principal keeps authority, the agent communicates biased information but the principal anticipates the bias and “corrects” the advice. The agent, knowing this, miscommunicates in equilibrium, making communication noisy and wasting information. Delegation can solve this: the principal lets the agent decide and commits not to override, inducing honest use of information. This holds if preference conflict is not too large. Under delegation, retaining a veto or override right reduces utility because anticipation of veto makes the agent distort again. Thus, pure delegation is optimal in moderate conflicts and involves a trade-off between control and efficient use of information.

Dilip Mookherjee and several coauthors have explored further this trade-off and provided several interesting insights on optimal design of decentralization. Taking a perspective complementary to Dessein (2002), the paper by Melumad, Mookherjee, and Reichelstein (1997), examines delegation under formal contracting with complexity constraints. Considering a principal contracting with multiple agents, they focus on whether the principal should design one grand contract (centralize all decisions) or delegate contracting authority to an intermediary manager. A key new ingredient is a limit on contract complexity – specifically a bound on the number of contingencies or messages allowed in the mechanism. Without communication or complexity limits, the Revelation Principle suggests the principal could achieve the same outcome by centralizing. Indeed, Melumad et al. note that under unrestricted contracts, delegation cannot outperform centralization and at best can only match it. However, when contracts are forced to be simple, delegation may confer an advantage. The intermediary (interpreted as a manager) can use his own private information in making decisions while the principal’s coarse contract could not fully condition upon this manager’s information. The authors identify conditions under which this flexibility gain of delegation outweighs the control loss of the principal not directly choosing actions. In other words, if the environment is complex but contracts must be simple, allowing a manager to adapt decisions to local information (which the principal cannot fully specify *ex ante*) can improve performance. Melumad et al. formally show cases where the delegated arrangement yields strictly higher expected welfare for the principal than the best achievable centralized contract. Thus, their model

echoes Dessein's conclusion in a mechanism-design context: delegating decision rights can be optimal when the principal faces limits in communicating or processing information.

Adding to this theme, work by Mookherjee and Tsumagari (2014) provides a multi-agent mechanism design perspective with explicit communication constraints. They study a principal's optimal contract with several agents who each have private information on their cost, under the assumption that communication is costly or limited (for example, messages or negotiation rounds are restricted). The striking result is that centralized decision-making is strictly dominated by decentralized arrangements when communication is constrained. In an optimal mechanism with communication frictions, the principal will decentralize decision authority to the agents and facilitate direct information exchange among them. Essentially, the agents are allowed (or incentivized) to talk to each other and coordinate, rather than passing all information through the principal. By decentralizing, the organization avoids the bottleneck of a single decision-maker trying to gather and act on all private information through limited channels. Interestingly, Mookherjee and Tsumagari find that the best mechanisms often involve gradual information revelation over multiple rounds. For example, agents might share some information in an initial round, make partial decisions, then share more, etc. This gradual, peer-to-peer communication structure stands in contrast to centralized communication and commands. The takeaway is that when it is hard to communicate, empowering agents and structuring their communication network yields better outcomes than a top-down approach. This reinforces the idea that the value of delegation grows in environments where the principal's ability to gather or process information is limited. This is conceptually similar to Dessein (2002)'s result (delegation dominates cheap-talk communication) but extended to networks of agents: hierarchy can be outperformed by a more decentralized architecture utilizing direct communication among subunits.

The question of centralized vs. delegated structures has been investigated in the context of procurement of two components by Baron and Besanko (1992). Their model explicitly allows various organizational forms: (i) integration (one entity produces both components), (ii) separation (two separate suppliers both contract directly with the principal), and (iii) delegation of contract (the principal contracts with one lead supplier who subcontracts the other component with the other supplier). Unlike integration, under separation the optimal output might not simply depend on the total cost— in some cases the principal distorts the allocation in a more complex way due to separate private information, introducing a new source of inefficiency.

Baron and Besanko show that if the principal delegates to a lead supplier, and if the principal can observe the subcontracting terms between this lead supplier and the sub-supplier, then delegation results in the same outcome as the separation case (with same information structure). The intuition is that if all contracts are transparent, the lead supplier cannot use any private information beyond what the principal could have gathered by dealing with both suppliers directly. Thus, in a world of perfect contracting adding an extra layer (delegation) cannot help the principal – it either does nothing or introduces inefficiency. In fact, as in the above, benefit of delegation would come only if information is affected or if it allows some commitment/transaction the principal alone could not achieve.

Prior to their 2014 contribution, Mookherjee and Tsumagari (2004) provided a rigorous analysis of a similar issue. In this article, they consider a one principal–two agent model where the agents could either contract directly with the principal or one agent could act as an intermediary for the other. A key friction they include is adverse selection (agents have private cost info) and possible collusion among agents. Their main result is somewhat like Baron and Besanko: delegating to one agent to subcontract the other is no better for the principal than contracting with both directly, unless specific conditions hold. A scenario of interest where delegation might help is if the two agents produce complements and the intermediary has superior information or ability to coordinate those complements. If agents produce substitutes or independent goods, introducing an intermediary just inserts an extra layer of information asymmetry which hurts the principal.

Rantakari (2008) complements these perspectives by modeling organizations where each division benefits from adapting to local factors as well as coordinating with other divisions. Information is soft and distributed across the divisions. He shows that the performance gap between centralized and decentralized authority is non-monotonic in the need for coordination: both perform equally when coordination is either trivial or paramount. Indeed, in these extreme cases, the divisions' and the headquarters' incentives are aligned to either full coordination or full adaptation to local factors. A conflict arises when each division's performance balances local adaptation and global coordination. Importantly, in asymmetric environments — for example, when a small division develops a new product alongside a large, established division — asymmetric governance structures such as partial centralization or authority to one division over the other outperform both centralization and decentralization. This helps rationalize observed practices such as “skunkworks” arrangements — semi-autonomous units given freedom to innovate while larger divisions remain centrally coordinated — that prevail in innovative firms. Moreover, when incentive conflicts across divisions can be sufficiently reduced, decentralization strictly dominates centralization, further emphasizing the adaptive value of delegated authority.

To conclude this section, consider the article by Severinov (2008), which also compares three organizational forms for producing a final good that requires two inputs: (1) centralization – a single agent produces both inputs in-house (akin to consolidation), (2) separation – the buyer contracts separately with two specialized agents, each producing one input, and (3) delegation – the buyer contracts with one primary agent who in turn subcontracts the production of the other input to a second agent. Severinov adds some new angles by focusing on the degree of complementarity or substitutability between the two inputs in the final product. He finds that this technological relationship dictates which organization is optimal. For example, if the inputs are highly complementary (meaning the value of having both is larger than the sum of individual values), then having a single responsible entity (under consolidation or delegation) might ensure better coordination of investments or qualities. If the inputs are substitutes, separate sourcing might induce competition or flexibility. Severinov shows conditions under which delegation yields the same payoff to the buyer as the multi-agent direct contracting, as in Mookherjee and

Reichelstein (2001) and Baron and Besanko (1992). Moreover, Severinov also considers which agent should be the primary contractor when delegation is used. For instance, if one input is more critical, that one might be chosen as the lead who subcontracts the other.

The Severinov paper ties together the technical nature of the production (complements vs substitutes) and informational issues. The finding that if inputs are complementary, a fragmented approach can suffer from coordination failures is reminiscent of the regulation angle in Gilbert and Riordan (1995). In such cases, a more integrated procurement (one main contractor overseeing everything as often, for example, in construction work) can ensure that complementarities are managed and there's no multiplication of informational rent. On the other hand, if inputs are independent or substitutable (say two suppliers offering alternative solutions to the same problem), competition can be harnessed by keeping them separate.

A common theme emerges from this literature. If the principal can costlessly gather and process all relevant information and write complete contracts, centralization is as good as or better than delegation. However, real organizations face communication costs, bounded rationality contract complexity, and strategic misreporting. Delegation of authority can then be an optimal organizational response that contains the costs and inefficiencies created by these frictions, while leveraging information dispersed at the local level.

The literature we have reviewed focuses on inefficiencies related to lack of verifiability and softness of information. However, the lessons should extend to other forms of communication costs. For instance, Laffont and Martimort (1998)—discussed in the section of collusion between agents—study the implications of constraints on communication for optimal delegation. In their work the constraint is one of anonymity of the mechanism (identical treatment of all agents). Delegation then relaxes the constraint and may dominate centralization when agents can secretly side-contract (collude).

In practical terms, these theories help explain why many firms push decisions down to division managers or teams closest to the market: doing so can overcome the “knowledge is power” problem by empowering those with knowledge to act on it.

4. Hierarchies and Incentives in Organizations

This section examines models with multiple layers of hierarchy and their incentive implications. A fundamental question is whether adding layers of agents (supervisors, middle managers, etc.) creates additional inefficiencies, or if an organization can be scaled up without loss. Several papers in the 1990s provided answers by extending principal-agent theory to hierarchical settings.

McAfee and McMillan (1995) coined the term “organizational diseconomies of scale” to describe how firms become less efficient as they grow larger, due to incentive and information problems. In their model, a firm's productive capability is spread among many individuals (each with private information about their own part). The central insight is that when information about the firm's operations is dispersed across layers, a hierarchy must incur costs to aggregate and act on that information. Specifically, they show that the longer information must travel up a hierarchy (the more layers between the information source

and the top decision-maker), the larger the efficiency loss. The hierarchy is forced to leave “rents” or slack to motivate information sharing and proper effort at each layer. These information rents cumulate as the number of layers increases. The result is that a firm with a very tall hierarchy might give up so much rent to its agents (to ensure truthful reporting and effort) that this firm cannot compete with smaller, leaner organizations with less dispersed information. McAfee and McMillan thus provide a formal underpinning for the notion that beyond a certain size, companies face diseconomies of scale arising from internal information problems. A long chain of command means the top doesn’t know what bottom-level workers know and extracting that knowledge is costly. They conclude that in competitive markets, overly hierarchical firms may simply not survive because their unit costs (inflated by internal incentive payments and inefficiencies) are too high. Their result, qualitatively, is a limit to firm size – not due to technology or coordination per se, but due to internal incentive costs that escalate with complexity.

Challenging this view, Melumad, Mookherjee, and Reichelstein (1995) provide one of the first building blocks for these hierarchical results. They study a simpler three-level hierarchy and demonstrate how the principal could implement the same outcome as a direct contracting benchmark by using a carefully structured mechanism (essentially, a delegated incentive contract) with the intermediate manager. One feature of their solution is that the principal sometimes has to subsidize internal transactions – for example, encourage the intermediate agent to delegate some decision to the lower agent by compensating him for any information rent paid out by the intermediate agent to the lower level agent. This idea – subsidizing outsourcing – ensures the intermediate does not become a bottleneck that distorts the incentives of the lower agent.

Mookherjee and Reichelstein (2001) generalized this logic to any number of layers and branches and present conditions under which a hierarchy can be just as efficient as a simple two-tier principal-agent relationship. They study an arbitrary hierarchy (multiple vertical layers and possibly branching at each layer) in a general adverse-selection model. The remarkable main result of their paper is an equivalence result: if certain assumptions are met, any complex hierarchy can achieve the second-best outcome (the optimal outcome compatible with incentive constraints) as if the principal had contracted directly with all agents. In short, under these conditions, adding layers does not create additional agency loss beyond what is present in a one-layer principal-agent problem. What are the assumptions? They are essentially the classical mechanism design conditions: risk neutrality of the principal and the agents, no wealth constraints (so agents can freely transfer payments), no communication costs, and crucially the ability to use top-down contracting with observability of contracts between layers. In such a setting, the principal can design a comprehensive mechanism determining communication and contracting restrictions, inducing intermediate managers to contract with their subordinates in a way that ensures information is passed and incentives are aligned throughout the chain. For example, an optimal scheme might involve the principal offering incentive schemes to her immediate subordinates that depend on what those subordinates in turn offer further down, and so on. Mookherjee and Reichelstein prove that as long as each bilateral relationship can be monitored by the principal (so there’s no secret collusion or hidden

side-contract at any link) and contracting can be appropriately contingent, no extra loss arises from decentralizing decisions through the hierarchy. A key condition for this to hold is that there is no informational “leakage” or information loss between layers –and the principal must be able to subsidize or tax all internal transactions appropriately. Thus, under some ideal contracting conditions (which are quite restrictive) a well-designed hierarchy can replicate the performance of a centralized mechanism. The authors thus provide a benchmark for efficient hierarchy, which helps understanding aspects of a hierarchy that are the source of inefficiency (e.g. if contracts are incomplete, communication is costly or if agents can collude), and additional costs it will entail.

For a more concrete feel, consider a three-tier example (principal – supervisor – worker). If the principal could dispense from the supervisor and contract directly with the worker, the principal would extract certain rents and impose certain output distortion as per the usual mechanism design discussed in sections 2 and 3. If instead the principal hires a supervisor who then hires the worker, one might expect the supervisor to also need rents to be incentivized. Mookherjee and Reichelstein show that if the principal can observe and design the contract the supervisor offers the worker and transfers, the principal can effectively charge the supervisor for any excessive rent passed to the worker, thereby preventing additional cost. However, this neat result hinges on the ability of the top layer to perfectly monitor contracts between the two other layers.

Mookherjee (2006), in a survey of decentralization and hierarchies, summarizes these points eloquently: in a frictionless contracting world (no communication or complexity costs, no collusion, commitment possible), delegation of decisions in a hierarchy need not cause any loss of control or coordination. Sufficient conditions include exactly those identified above: risk-neutrality, top-down monitoring of transactions (the principal can observe and adjust contracts among subordinates), and unlimited contract complexity. Under those conditions, delegating decisions to lower levels can be as effective as centralizing them, because the principal can design incentive schemes that pass through the layers without performance loss. However, as Mookherjee (2006) goes on to review, once we introduce real-world frictions – such as communication costs (which make centralized decision-making less flexible) or limits on contract complexity – then a trade-off arises between loss of control and flexibility as we saw in Dessein (2002) and Melumad et al. (1997). Furthermore, the possibility of collusion among agents tends to expand the range of situations where delegation is optimal. We will cover collusion in the supplement section.

In summary, the literature on hierarchies and incentives reveals an optimistic scenario and a pessimistic one. Optimistically, a hierarchy need not cause loss if everyone’s incentives are perfectly calibrated and monitored – large organizations can, in theory, be as efficient as small ones (there is no inherent “bureaucratic penalty”). However, the result that “hierarchies can be efficient” is quite fragile. It relies on the principal’s ability to overcome any incentive issues that multiple layers introduce. The pessimistic (and perhaps more realistic) view is that as soon as we acknowledge limitations – e.g., people cannot communicate freely, contracts can’t cover every scenario, or subordinate agents might collude – then layers do add extra agency costs. McAfee and McMillan’s diseconomies of scale argument underscores the latter: in practice, large firms often face overhead and

incentive problems that small firms avoid. The balance between these views can inform managers: a firm growing in scale either needs to invest in very good internal controls or accept that beyond a point, internal inefficiencies will grow. This can justify outsourcing or keeping organizational structures as flat as possible. The theoretical boundary conditions found by Mookherjee & Reichelstein (2001) serve as a benchmark – a reminder that any observed inefficiency in hierarchies comes from specific failures of their assumptions, such as communication barriers.

5. Procurement with Competing Suppliers

In many cases the buyer must choose not only the quality produced and the compensation, but also which among several suppliers will benefit from the contract. When suppliers are ranked by efficiency—one is uniformly better than the other, e.g., cost decreases with the type—the solution appears to be similar to the single supplier case but with the benefit of competition: the buyer designs a price-quality schedule and let the suppliers self-select. The price-quality schedule is obtained as the most efficient quality for each level of “virtual” costs (that is the cost adjusted for informational cost needed to induce honest reporting of costs by the suppliers). The selected supplier is the most efficient and produces the quality that is efficient for her virtual cost.

The optimal mechanism resembles an auction with quality-differentiated bids. A classical procurement auction is known as a scoring auction, which assigns a score to each pair of price and quality. For instance, in second-score auction, firms bid a score and the contract corresponding to the second score is assigned to the highest bidder.

Boone and Schottmüller (2016) extend this analysis to optimal procurement mechanisms when firms are specialized in some level of quality– for instance, one supplier might be more efficient at providing high-quality versions of a product while another is better at low-quality. The question is how a buyer (principal) should structure procurement when facing unobserved heterogeneity of suppliers and there is no ranking of the suppliers. In their model, the procuring agency values both price and quality of a good or service. Suppliers have private information on their cost functions and importantly, the identity of the lowest-cost supplier depends on the quality level as described above. Hence a firm might have the lowest cost for low-quality production, while another is cheapest for high quality. This leads to interesting consequences for mechanism design. Boone and Schottmüller find that the standard conclusions “no-distortion-at the-top” and “no-rent-at-the-bottom” do not hold when the ranking of cost depends on the level of quality. Specifically, a range of firm types (their private cost characteristics) may end up earning zero profit (no informational rent) in equilibrium, while rent may accrue to higher or lower types (or both). Intuitively, if a supplier is only competitive when quality is q , the buyer need not worry that the supplier claims another quality q' where it is not competitive and thus can design a scheme where that supplier is willing to win the contract with quality q at a price equal to cost, while another type wins at quality q' . The solution still consists in implementing a single price-quality schedule with an auction-type mechanism that selects the most desirable producer and the quality within the schedule. The price-quality schedule is obtained as the most efficient quality for each level of “virtual” costs. The complexity is

that the quality distortion compared to first best (the optimal quality under full information) may be upward or downward.

The work thus shows how a buyer can combine competition and information revelation in a procurement process, at the cost of distortion of quality but also of the choice of suppliers (as the lowest cost supplier may not always obtain the contract). Boone and Schottmüller show how an auction-like mechanism can drive down rents, which may turn out to be more complex than standard procurement procedure.

6. Information Alliances

Information alliances are coalitions of independent suppliers who pool and internally verify their private information before contracting with a buyer. They sit between full mergers and arm's-length procurement, shifting the coordination locus to suppliers themselves. This internal verifiability economizes on duplicated information rents while preserving independence of suppliers.

Baron and Besanko (1999) show that such alliances, organized around a trusted verifying party, can dominate mergers, decentralized procurement, and delegation of authority. By coordinating reports, suppliers internalize misreporting externalities and raise the buyer's expected surplus, though the buyer must respect type-dependent participation constraints that alter optimal allocations.

Dequiedt and Martimort (2004) model alliances as delegated monitoring with fixed costs. A principal may pay to learn a subcontractor's type, reducing duplicated rents but incurring monitoring costs. Optimal design features discrete regimes—full consolidation, partial alliances, or arm's-length—depending on cost parameters. The framework shows why too little monitoring may arise and connects alliances to broader themes of integration, delegation, and hierarchy.

7. Conclusion

The literature on the economics of procurement emphasizes the central role of information flows and incentive alignment in shaping optimal organizational structures. Drawing on foundational insights from Arrow (1975), who emphasized the informational advantages of intra-firm coordination over market transactions, subsequent theoretical models have formalized how organizations can be designed to mitigate information loss and incentive misalignment. A fundamental principle emerges: allocating decision-making rights to those with localized information (Dessein, 2002) is often more effective than centralized decision-making, particularly when communication is costly or subject to distortion.

Research demonstrates that integration or consolidation of roles can reduce inefficiencies such as double markups or lost signals (Baron and Besanko, 1992; Gilbert and Riordan, 1995), but organizational enlargement carries risks. As McAfee and McMillan (1995) note, scaling an organization introduces new agency problems unless they are mitigated through sophisticated contracts (Mookherjee and Reichelstein, 2001). The literature resolves key paradoxes, for instance, the effectiveness of delegation in curbing strategic misrepresentation, or the value of middle managers in relieving communication

overload despite apparent cost burdens. Notably, Mookherjee and Tsumagari (2004) show intermediaries add value only when they enable coordination that contracts or communication alone cannot achieve.

Different forms of delegation versus centralization emerge depending on the prevailing constraint. For instance, while Dessein (2002) and Melumad et al. (1997) advocate delegation in the face of limited communication capacity. Similarly, hierarchical organization can either facilitate coordination or amplify hidden-action problems, depending on how contract observability and control are structured.

From an applied standpoint, these theoretical findings help clarify managerial decisions. Delegation supports flexibility and responsiveness when monitoring is limited, but requires aligned incentives (Dessein and Santos, 2006). Oversight must be structured to prevent collusion, and procurement strategies must balance coordination costs against competition benefits (Boone and Schottmüller, 2016; Severinov, 2008). Moreover, whether to consolidate or fragment contracts hinges on whether integration simplifies communication or simply adds friction.

Supplementary Section

Adaptive Organizations and Task Design

While the above models largely consider a single decision or contract, another line of research examines how organizations adapt to changing environments and how tasks should be designed between agents.

Dessein and Santos (2006) develop a theory of “adaptive organizations,” focusing on the tension between flexibility and coordination. In their setup, an organization must decide how specialized its employees’ tasks are and whether to allow employees to adjust their actions based on local information (adaptation). A highly specialized organization assigns narrow tasks to individuals, which can improve efficiency in stable environments. However, extensive specialization can become a liability when the environment changes and local information matters – because each specialist, following a predefined plan, might ignore valuable local knowledge outside their narrow scope. Dessein and Santos argue that an adaptive organization gives employees more flexibility to deviate from pre-set plans and tailor their actions to local conditions. The flip side is that, as everyone freely adapts, this may undermine coordination. The organization then optimally limits the degree of specialization of tasks but favors communication among employees. With broader autonomy, each employee can handle a wider range of contingencies while communicating to coordinate with other employees.

One implication from Dessein and Santos is that “improvements in communication technology may reduce specialization by pushing organizations to become more adaptive”. When communication is cheap, firms optimally broaden tasks so that local adjustments can be made by individuals, with communication ensuring these adjustments are mutually

consistent. When communication is costly, organizations rely on communication to handle exceptions, enabling adaptation during exceptional times only.

The insights from this work complement the incentive-focused literature. While Dessein (2002) showed how giving agents an authority avoids communication noise due to incentive conflict, Dessein and Santos (2006) show how giving agents flexible tasks avoids the need for constant formal communication in a changing environment. Both deal with making better use of local knowledge: one by overcoming strategic miscommunication, the other by overcoming rigid task definitions. For instance, if frontline workers have knowledge of new opportunities, the organization should consider giving them autonomy, while other employees with less private knowledge may be siloed into narrow duties.

Lastly, Mookherjee and Tsumagari (2014) (discussed earlier under centralization vs delegation) general communication-constrained mechanism design sheds light on procurement with multiple suppliers who each know their costs. The finding that direct communication among agents is optimal could be applied to a supplier network: rather than the buyer being the hub of all information, it can be better to let suppliers communicate and coordinate some decisions among themselves, under proper incentives. This essentially suggests a flexible network organization rather than a strict hierarchy. For example, in a large construction project, instead of the client (principal) micromanaging interactions between different subcontractors, it could be more efficient to let subcontractors work out interface issues directly and only hold them accountable for final outcomes. By decentralizing the authority (each subcontractor makes certain decisions) and enabling inter-supplier communication, the project can benefit from their combined information. If communication constraints are the binding issue, empowering the network might outperform a centralized command. Of course, this assumes the incentives are aligned such that suppliers will truthfully share information with each other – which might need careful contract design.

In summary, adaptive organization theory stresses designing tasks and authority to handle uncertainty. Limitations on specialization and investments in communication are tools to maintain coordination while empowering adaptation. This stands in contrast to a purely centralized or rigid organization where coordination is strong, but adaptation is poor. It adds another dimension to delegation: not just who decides, but how jobs are defined.

Collusion and Delegated Monitoring

In multi-layer organizations, collusion is a central concern. Collusion refers to secret side-agreements between agents at different levels (or in different parts) of the hierarchy that undermine the principal's intent. For example, a supervisor and a worker might collude to misreport the worker's output or cost, then share the illicit gain. Several papers explore how the risk of collusion affects the optimal organizational form – whether the principal should empower an intermediary (delegated monitoring) or deal directly with agents (centralized monitoring).

Laffont and Martimort (1998) examine a three-party model: a principal, and two productive agents. The agents can engage in side-contracting (collusion) on reports to mislead the principal. Laffont and Martimort compare two archetypal organizations: centralization (no supervisor; the principal directly contracts with the agents) and decentralization (the principal delegates oversight to a supervisor who interacts with the agent). In a world of complete contracting (no communication limits and the principal can write an extensive contract with both parties), they find that collusion is not an issue under centralization – the principal can essentially design a contract that leaves no room for a beneficial collusive side contract. In fact, with full communication and commitment, centralization and delegation perform equally well. However, if there are limits on communication, the conclusion changes. When the principal cannot fully observe or dictate the interactions between the agents then collusion becomes possible under a centralized scheme. In their model, the limit is due to aggregation of reports that hides individual reports and the identity of senders, forcing equal treatment of agents (referred to as anonymity). This leaves scope for manipulation of aggregate reports that cannot be prevented by playing one agent against the other. In such cases, Laffont and Martimort show that delegation (decentralization) mitigates the collusion problem and dominates centralization. By design, a delegated hierarchy is asymmetric—the supervisor and agent have different roles and bargaining powers—which can break the symmetry that made collusion easier under centralization. In a delegated setting the principal might give the supervisor certain discretionary power or a reward scheme that makes it harder for the agent to bribe him. In summary, with unrestricted contracts, centralization performs well (collusion can be neutralized), but with information control limitations, delegation (having a hierarchical supervisor) might improve welfare because it changes the game of side-contracting in the principal's favor.

Faure-Grimaud, Laffont, and Martimort (2003) investigate further the collusion issue when information is soft. “Soft information” captures information that cannot be verified by outsiders (like an unverifiable report or opinion) – for instance, a supervisor's assessment of an agent's performance might be soft. They consider an organization with a supervisor and agent where the supervisor's observation is non-contractible (soft), but the supervisor can communicate it to the principal or collude with the agent to suppress or distort it. One might fear that if the supervisor's signals aren't verifiable, having a supervisor is useless or even harmful (since the supervisor could collude with the agent and lie without fear of external check). However, Faure-Grimaud et al. show that supervision with soft information can still create value in the presence of collusion. The principal can design schemes (e.g., bonus payments, whistleblowing rewards, or discretionary authority) such that even though the info is soft, the supervisor has incentives to report truthfully in many cases. In particular, they find conditions for the outcome under a soft-information supervisory scheme to be equivalent to what could be achieved if information were hard/verifiable. This is done by carefully balancing the incentives so that any attempt by the agent and supervisor to misreport would leave them no better off (or one of them would refuse to participate). More generally they show that even if a supervisor's information is soft and collusion may happen, introducing a supervisor can enlarge the set of incentive schemes

available to the principal. Collusion can be curbed by creating an internal conflict of interest. The authors find that multiple organizational forms can achieve the same second-best outcome – for instance, a form where the supervisor is purely an advisor versus one where the supervisor has some decision authority might be equivalent if properly designed. Recall that Mookherjee and Tsumagari (2004) found that, with no limits on communication, if agents can collude, delegating production decisions to one agent who subcontracts is generally not better for the principal than contracting with both directly. They argue that unless products are complements and there are specific mitigating factors (the intermediary having superior information), the principal loses out by such delegation. Laffont and Martimort (1998) show that delegation may be optimal when collusion is at stake if there is imperfect communication. One issue here may be that the conclusions depend on the detailed assumptions, in particular on the way collusion is implemented.

References

- Arrow, K. J. (1975). Vertical integration and communication. *The Bell Journal of Economics*, 173-183.
- Baron, D. P., and Besanko, D. (1992). Information, control, and organizational structure. *Journal of Economics and Management Strategy*, 1(2), 237-275.
- Baron, D. P., & Besanko, D. (1999). Informational alliances. *Review of Economic Studies*, 66(4), 743-768.
- Boone, J., & Schottmüller, C. (2016). Procurement with specialized firms. *The RAND journal of economics*, 47(3), 661-687.
- Dequiedt, V., & Martimort, D. (2004). Delegated monitoring versus arm's-length contracting. *International Journal of Industrial Organization*, 22(7), 951-981.
- Dessein, W. (2002). Authority and communication in organizations. *The Review of Economic Studies*, 69(4), 811-838.
- Dessein, W., & Santos, T. (2006). Adaptive organizations. *Journal of political Economy*, 114(5), 956-995.
- Faure-Grimaud, A., Laffont, J. J., and Martimort, D. (2003). Collusion, delegation and supervision with soft information. *The Review of Economic Studies*, 70(2), 253-279.
- Gilbert, R. J., and Riordan, M. H. (1995). Regulating complementary products: A comparative institutional analysis. *The RAND Journal of Economics*, 243-256.
- Laffont, J. J., and Martimort, D. (1998). Collusion and delegation. *The RAND Journal of Economics*, 280-305. McAfee, R. P., and McMillan, J. (1995). Organizational diseconomies of scale. *Journal of Economics and Management Strategy*, 4(3), 399-426.

- McAfee, R. P., & McMillan, J. (1995). Organizational diseconomies of scale. *Journal of Economics & Management Strategy*, 4(3), 399-426.
- Melumad, N. D., Mookherjee, D., and Reichelstein, S. (1995). Hierarchical decentralization of incentive contracts. *The RAND Journal of Economics*, 654-672.
- Melumad, N., Mookherjee, D., & Reichelstein, S. (1997). Contract complexity, incentives, and the value of delegation. *Journal of Economics & Management Strategy*, 6(1), 257-289.
- Mookherjee, D. (2006). Decentralization, hierarchies, and incentives: A mechanism design perspective. *Journal of Economic Literature*, 44(2), 367-390.
- Mookherjee, D., & Reichelstein, S. (2001). Incentives and coordination in hierarchies. *The BE Journal of Theoretical Economics*, 1(1), 20011005.
- Mookherjee, D., and Tsumagari, M. (2014). Mechanism design with communication constraints. *Journal of Political Economy*, 122(5), 1094-1129.
- Rantakari, H. (2008). Governing adaptation. *The Review of Economic Studies*, 75(4), 1257-1285.
- Severinov, S. (2008). The value of information and optimal organization. *The RAND Journal of Economics*, 39(1), 238-265.