FROM FLEXIBILITY TO INSECURITY: HOW VERTICAL SEPARATION AMPLIFIES FIRM LEVEL UNCERTAINTY

David Thesmar — HEC

Mathias Thoenig — University of Geneva and Paris School of Economics

June 30, 2007

Abstract

This paper presents a model where firms may endogenously externalize part of their production process. We start from the premise that adaptation to uncertainty cannot be contracted upon in the worker/employer relationship. Vertical separation then balances flexibility gains against hold-up costs of opportunistic behavior by outside contractors. In equilibrium, the degree of separation is shown to depend on the degree of product market competition, contractor's bargaining power, and the volatility of demand shocks. Our main result is that an increase in the degree of vertical separation amplifies the elasticity to demand shocks of firms' sales and employment. It does not, however, amplify aggregate

Acknowledgements: For insightful and numerous discussions we thank Philippe Aghion, Bernard Caillaud, Eve Caroli, Antonio Ciccone, Matthew Ellman, Dalia Marin, Niko Matoushek, Xavier Ragot, Thierry Verdier, John Van Reenen, as well as participants at seminars in University College London, Amsterdam, University College Dublin, Ecole Normale Supérieure Paris, University Pompeu Fabra, University of Southampton, Toulouse University, and CEPR-ERWITT conference. All remaining errors are, of course, our owns. Both authors are members of the CEPR.

Email addresses: Thesmar: thesmar@hec.fr; Thoenig: thoenig@ecopo.unige.ch

uncertainty. Evidence from firm level data is shown to be largely consistent with the main implications of our theory. (JEL: L16, L23, L24)

1. Introduction

It is a common view among practitioners of the firm that vertical separation improves corporate flexibility. Because outside suppliers are given stronger incentives to adapt, vertical separation allows one to follow demand shifts more closely. This paper starts from such premises in order to assess the interplay between vertical separation and firm-level uncertainty. The main result of our theoretical and empirical investigations is that any exogenous increase in vertical separation amplifies the impact of demand shocks on the product and labor markets by increasing firm-level volatility of sales and employment.

On a theoretical basis, our view of vertical separation rests on two core assumptions. The first is that firms face demand shocks that require adaptation efforts from the workforce; the second is that these shocks are to some extent not contractible upon. It is therefore difficult for firms to enforce the ex post efficient level of efforts from their workers. By leaving ownership of intermediate inputs to subcontractors, firms make these workers credible residual claimants on joint output and so provide them with the proper incentives to adapt. In sum, vertical separation increases flexibility because it enables firms to switch from ex ante to ex post production decisions. However, the bilateral nature of the subcontractor/final producer relation gives rise to a hold up problem (Grossman and

Hart 1986). In our model, firms consequently trade off the benefits of flexibility against its costs in terms of opportunism. The equilibrium level of vertical separation is shown to depend positively on the degree of product market competition and final demand uncertainty and to depend negatively on coordination and transport costs. Competition matters because, when final demand becomes sensitive to the price policy of the firm, any ex post productive inefficiency is costly.

An important consequence of this microeconomic framework is that any increase in vertical separation *amplifies* the impact of demand shocks in term of sales and employment volatility at the firm-level. The intuition is that both parts of the production process are partial complements in that increased flexibility on the subcontractor's side must be matched by a similar increase on the contractor's side. This result is at odds with conventional wisdom, which views vertical separation as a practice of large and dominant companies to transfer risk to small subcontracting firms. Here our theoretical prediction, confirmed by empirical investigation, is that the risk borne by the contracting firm also increases with the degree of separation.

We then ask whether our theory of vertical separation also predicts amplification of aggregate uncertainty. Toward this end, our microeconomic theory is embedded into a simple general equilibrium model. At equilibrium, aggregate demand shocks do not induce vertical separation: since aggregate shocks cause co-movements in the price of intermediate inputs, the flexibility gain of separation is fully counteracted by equilibrium adjustments on the input market. We show also that aggregate shocks are *not* amplified by the degree of vertical separation—only idiosyncratic demand shocks are.

Empirical evidence is provided using a balanced panel of some 2,300 French firms for the 1984–1999 period. Our theory yields two predictions: (i) demand shocks promote vertical separation; and (ii) vertical separation amplifies the firm-level elasticity of sales and employment to demand shocks. The combination of these two channels unambiguously implies a positive correlation between volatility of sales/employment, and vertical separation, two variables that we can observe in the data. In a first set of specifications, we find robust evidence of such a positive correlation. The results hold in the cross section of firms and in the panel dimension (controlling for various observables) regardless of which measure of firm-level volatility is used, variance of sales growth or of employment growth. Yet these encouraging results do not enable one to discriminate between our two theoretical channels.

The rest of our empirical analysis is focused exclusively on the *amplification channel*—namely, the causal effect of vertical separation on firm-level uncertainty. The task is difficult because we do not directly observe demand shocks and because, lacking a proper exogenous source of variation for vertical separation at the firm-level, we cannot implement truly credible instrumental variable techniques. We thus use the lagged value of vertical separation and ask if an increase in separation is followed by a rise in firm-level volatility. This requires that we use annual data. However, since we cannot compute variance of sales on an annual basis, we propose an alternative estimation technique that is exactly derived from the reduced form of our model: we look at the sensitivity of firm sales (or employment) to industry sales shocks. In passing, we propose an original and simple estimation procedure to account for fixed effects in the sensitivity of firm sales to industry shocks. This procedure is useful because standard within techniques do not work in these nonlinear specifications. Empirically, we find that sensitivity of labor demand increases for up to three years after an increase in vertical separation.[1]

From a microeconomic viewpoint, our emphasis on flexibility and adaptation takes us close to a somewhat older literature on outside contracting that includes modeling of Chandler's intuitions by Carlton (1979). According to this view, firms use vertical integration in order to secure a steady supply of inputs when the upstream market is subject to uncertainty; in contrast with our theory and empirical results, this view predicts a negative correlation between vertical separation and firm-level uncertainty.

From a macroeconomic perspective, an emerging literature documents the recent trend in the uncertainty of corporate environment within industrialized economies (see Comin and Philippon (2005) for a survey). There is concurrently increasing evidence that firms are now refocusing on their core activities by outsourcing part of their production to outside contractors. This trend in vertical separation is explained by increased competition and international specialization (Feenstra 1998; Hummels et al. 2001) and by the spread of new technological and

1. This evidence is consistent with recent results of Mulhainathan and Scharfstein (2001) from the chemical industry. The authors find that nonintegrated firms tend to have production capacity that is more sensitive to demand shocks, although they interpret this piece of evidence differently.

organizational practices (Caroli and Van Reenen 2001). A direct consequence of our theory is that those two trends are not independent. In particular, the increase in demand uncertainty may explain part of the observed changes in vertical separation. Yet our amplification channel indicates that the causality also goes in the other direction: any force that has contributed to the increase in outsourcing and vertical separation—such as competition or the evolution of transport costs—may have contributed to the increase in firm-level uncertainty. Although a combination of both effects must surely have occurred at the macro level, our firm-level empirical analysis does not allow to take a stand on this.

This paper proceeds as follows. Section 2 highlights the key features of our model and Section 3 details the most important predictions of our theory. Section 4 puts our theory to the test by analyzing French firm-level data from 1984 to 1999. Section 5 concludes and offers suggestions for further research.

2. The Baseline Story

This section presents our microeconomic mechanism. In this partial equilibrium model, demand shocks promote vertical separation. We show that vertical integration then amplifies the elasticity of firm-level sales and employment with respect to demand shocks.

2.1. Set Up

The production structure is composed of a downstream stage (the *buyer*) and an upstream stage (the *supplier*). The supplier produces one intermediate input whose quality q must be specifically tailored to the buyer's need (otherwise, the final production cannot take place). Then the buyer produces y units of output with l units of labor and the intermediate input of quality q according to a Cobb–Douglas function:

$$y = q^{\alpha} l^{1-\alpha}. \tag{1}$$

Hence, the higher is the quality q of the intermediate input, the larger is final production. The downstream firm is a monopolist facing an uncertain demand curve:

$$p \equiv \tilde{T} y^{-1/\sigma} \tag{2}$$

where y is the quantity of output sold, $\tilde{T}$ is a random variable of mean T and variance σ_T^2, and $\sigma > 1$ is the price elasticity of demand. This demand function can, for example, be derived from Dixit–Stiglitz preferences (see Section 3).

To fix ideas, think about a telecom operator launching its mobile phone service. It does not know a priori how many new customers it will garner—say, between fifty thousand and one million. Aside from buying the handsets and other pieces of equipment, the firm needs to hire employees for customer care and equipment maintenance. The firm also needs to set up software for registering and billing clients and for assisting customer care employees. The software is designed by a team of computer programmers, who have two choices.

1. They can build simple codes in parallel and then merge them. The resulting software will run well with one thousand orders per day. But because it has not been conceived in an integrated way, the software becomes slow and memory consuming with a larger number of customers. The advantage is that this code is easy to create and quick to implement.
2. They can instead write a more sophisticated code that is well optimized and hence well suited to deal with large numbers of customers. The registering function will connect well with the billing part and provide quickly the relevant information needed for adequate customer care. Implementing this software would, however, take more time and manpower.

Obviously, it is feasible to specify in advance how fast and efficiently the software can run. However, it may not be optimal for the firm to request ex ante a costly and efficient code when it might end up with fewer than one hundred thousand customers. In this illustrative example, an efficient software (high q) allows the firm to bill more customers, which would also require more employees (higher l). But such software is costly to create. On the other hand, an inefficient software would prevent the firm from billing large numbers of customers; in this case, the firm would optimally hire fewer employees.

Under uncertainty, the first-best level of output $\tilde{y}^*$, and thus the optimal input's quality level $\tilde{q}^*$, depend on $\tilde{T}$. This optimum could be achieved by writing—before the revelation of uncertainty—a contract specifying a payment to the supplier contingent on (i) the quality q and (ii) the demand shock $\tilde{T}$. Another solution would be to sign the contract contingent on q only, but after the shock $\tilde{T}$ is revealed. We rule out these possibilities by assuming contractual incompleteness in a fashion similar to Battigalli and Maggi (2002): the demand shock $\tilde{T}$ (contingencies) and the characteristics q (the actions) are observable and verifiable; contracts are signed ex ante (before information disclosure) and cannot be modified;[2] and there are costs of describing contingencies and q when writing the contract.

For simplicity, we restrict the set of contracts to the two polar forms highlighted by Battigalli and Maggi (2002): a *rigid contract* $(\bar{q}, W)$ specifies the level of quality $\bar{q}$ and the associated payment W given to the supplier; and a *discrete contract* $(\emptyset, W)$ does not specify the quality.[3]

---

2. A justification for this would be that production must occur soon after the shock is revealed and that writing a contract is time consuming.

3. More precisely, a complete contract binds the parties by specifying $\{(\tilde{T}, \tilde{q}, \tilde{W})\forall\tilde{T} > 0\}$, where $\tilde{T}$ is the state of the demand, $\tilde{q}$ is the input's quality, and $\tilde{W}$is the supplier's payoff. Here we assume that the costs of writing the contract are extremely nonlinear: they are 0 if at most either one contingency $\tilde{T}$ or one quality requirement q is described; they are infinite otherwise. Hence only three contracts can be written at a non-prohibitive cost: $(\emptyset, \emptyset, W)$, $(\emptyset, \bar{q}, W)$, or $(\bar{T}, \emptyset, W)$.

In the last of these possible contracts, transaction and input production will take place only if the state of nature $\bar{T}$ is realized (but without specifying a quality requirement). Since profits are positive for every $\tilde{T}$, this contract clearly restricts the opportunities for a positive profit from the buyer's point of view. As a result, this contract is always dominated by the other two types.

We are thus left with two polar cases: the rigid contract $(\emptyset, \bar{q}, W)$, which states that a transaction will take place for any $\tilde{T}$ but with a quality requirement $\bar{q}$ and a payment W; and the discrete contract $(\emptyset, \emptyset, W)$, which states that a transaction will take place for any $\tilde{T}$, with *no* quality requirement $\bar{q}$ and with a payment W.

Aside from contracting decisions, the buyer chooses the degree of vertical integration of his firm. We take here a standard property rights approach to vertical integration (see Grossman and Hart 1986) and so assume that the buyer cannot sell his firm to the supplier. Under vertical integration the buyer owns the intermediate input q, which is produced by the supplier at a linear cost $C^{\mathrm{IN}}(q) = cq$. Under separation, the supplier owns the intermediate input: this gives her ex post bargaining power when bringing her productive intermediate inputs to the buyer for production. However, the supplier faces γ, an extra cost of coordination (or a transportation cost), such that production of the intermediate input costs her $C^{\mathrm{SEP}}(q) = \gamma cq$ with $\gamma > 1$. This last assumption may seem ad hoc, but it clarifies the exposition of our main mechanism and it will serve also for comparative statics exercises. Our theory of vertical integration thus builds both on the property rights framework of Grossman and Hart and on a technological assumption that tends to make vertical separation always less efficient. As we shall argue, a big advantage of this modeling strategy is its simplicity when dealing with the macroequilibrium (see Section 3). For readers uncomfortable with our two-sided approach, we propose in Appendix A a slightly different version of our model based only on an incomplete contract theory (and thus where $\gamma = 1$).

The timing of actions is as follows:

0. *Ownership structure and contracting.* The buyer chooses the ownership structure and signs either a discrete or a rigid contract with the supplier.
1. *Information.* The demand shock $\tilde{T}$ is revealed to both parties.
2. *Production of the intermediate input.* The supplier produces the quality level

q and brings it to the buyer. The payment to the supplier depends on: (a) the contract signed ex ante and possibly on (b) bargaining if the supplier owns the intermediate input.

3. *Production of the good.* The buyer uses the intermediate input, hires labor in quantity $\tilde{l}$, and produces the good. The firm's profit with respect to demand shock and intermediate input quality is given by

$$\pi(q, \tilde{T}) = \max_{l} \ \tilde{T}^{1/\sigma}(q^{\alpha} l^{1-\alpha})^{1-1/\sigma} - wl \tag{3}$$

$$= \frac{1}{\sigma\varepsilon} \tilde{T}^{\varepsilon} q^{1-\varepsilon} \left(\frac{(\sigma-1)(1-\alpha)}{\sigma w} \right)^{\sigma\varepsilon - 1}, \tag{4}$$

where w is labor's price and $\varepsilon = 1/(1 + \alpha(\sigma - 1)) < 1$.

We next study the unconstrained first-best solution as a benchmark case and then look at the four organizational structures: integration under discrete and rigid contracts and separation under discrete and rigid contracts.

2.2. Unconstrained first-best

The unconstrained first-best solution corresponds to the case where an ex ante complete contract contingent on q and $\tilde{T}$ can be written or when contracts contingent on q can be signed after the realization of $\tilde{T}$. Here integration is always chosen owing to the additional costs γ of coordination under vertical separation. Hence, the first-best level of quality is given by maximizing buyer's surplus $\pi(\tilde{q}^{\mathrm{FB}}, \tilde{T}) - \tilde{W}^{\mathrm{FB}}$ under supplier's participation constraint $\tilde{W}^{\mathrm{FB}} - c\tilde{q}^{\mathrm{FB}} \geq 0$. This is equivalent to maximizing total surplus:

$$\Pi^{\mathrm{FB}} = E\Big[\max_{q} \ \pi(q, \tilde{T}) - cq\Big]. \tag{5}$$

One then sees easily that first-best quality and expected profit are equal to

$$\tilde{q}^{\mathrm{FB}} = \alpha R\tilde{T}/c, \tag{6}$$

$$\Pi^{\mathrm{FB}} = RE[\tilde{T}]/(\sigma-1), \tag{7}$$

respectively, where

$$R \equiv \left(\frac{\sigma-1}{\sigma}\right)^{\sigma}\left(\left(\frac{\alpha}{c}\right)^{\alpha}\left(\frac{1-\alpha}{w}\right)^{1-\alpha}\right)^{\sigma-1}.$$

Intuitively, the optimal quality $\tilde{q}^{\mathrm{FB}}$ is an increasing function of the shock $\tilde{T}$, and its elasticity with respect to $\tilde{T}$ is 1: there is full adaptation.

2.3. Integration

Under integration, the buyer owns the upstream firm and the intermediate input. Assume first that the buyer and the supplier write a discrete contract $(\emptyset, \bar{W})$ specifying a non-contingent payment W, to be made to the supplier in case of production of the intermediate input, but where the quality is not specified. In this case, the supplier receives W in period 2 for any level of quality. She therefore chooses the lowest possible quality level: $q = 0$. As a result, the buyer cannot produce at all and gets zero surplus. Thus, *under integration a discrete contract cannot sustain the production of an intermediate input with nonzero quality*.

Under a rigid contract $(\bar{q}, \bar{W})$, the payment $\bar{W}$ is made only if an intermediate input of verifiable quality $\bar{q}$ is delivered. When such a contract is signed in period 0, the buyer's expected payoffs and the supplier's participation constraint are given by $E[\pi(\bar{q}, \tilde{T})] - \bar{W}$ and $\bar{W} - c\bar{q} \geq 0$ respectively. Given this, the buyer chooses the level of quality that maximizes expected surplus:

$$\Pi^{\mathrm{IN}} = \max_q E\left[\pi(\bar{q}, \tilde{T}) - c\bar{q}\right]. \tag{8}$$

This yields the quality level and expected profit under integration as follows:

$$\bar{q} = \frac{\alpha R}{c} E[\tilde{T}^{\varepsilon}]^{1/\varepsilon}; \tag{9}$$

$$\Pi^{\mathrm{IN}} = \frac{R}{\sigma - 1} E[\tilde{T}^{\varepsilon}]^{1/\varepsilon}. \tag{10}$$

(Throughout this paper, we write e.g. $E[\tilde{x}]^{\alpha}$ to mean $(E[\tilde{x}])^{\alpha}$.) By definition, $\bar{q}$ is deterministic and fixed in advance. With respect to the unconstrained first-best, the fit with demand deteriorates and thus profits under (rigid) integration are lower than under the first-best scenario owing to this absence of flexibility. The loss of flexibility due to the rigidity of contracts under integration appears in the last term of the profit expression:

$$\frac{\Pi^{\mathrm{IN}}}{\Pi^{\mathrm{FB}}} = \underbrace{\frac{E[\tilde{T}^{\varepsilon}]^{1/\varepsilon}}{E[\tilde{T}]}}_{\text{flexibility}};$$

this is less than 1 by virtue of Jensen's inequality since $\varepsilon < 1$. For the same reason, $\bar{q}$ is smaller than $E[\tilde{q}^{FB}]$, the average quality under first-best.

Although the buyer is risk neutral, he dislikes risk because the intermediate input's quality cannot be adapted ex post. The buyer behaves with respect to uncertainty like a risk-averse agent (with a CRRA utility), whose relative risk aversion is $1/\varepsilon = 1 + \alpha(\sigma - 1)$. This risk aversion increases with α, the returns to scale of the intermediate input's quality in production function (1): this follows because only the intermediate input's quality (and not labor) is subject to contractual rigidity. Moreover the more prevalent are rigid inputs, the more sensitive are profits with respect to uncertainty. Risk aversion also increases with σ, the elasticity of demand faced by the monopolist—a measure of market power and competition. The tougher the competition, the more consumers choose a substitute if the monopolist's scale of production is ill-adapted to the demand shock.

2.4. Separation

Under separation, the supplier owns the intermediate input. This potentially gives her ex post bargaining power in the absence of a contract but at the price of an extra cost of coordination γcq. We start with discrete contracts.

Consider first the case of a discrete contract specifying a payment W but not the quality. In period 2, the supplier decides to produce a quality level $q^{\mathrm{OUT}} \geq 0$. Since she owns the intermediate input q^{OUT}, she may threaten the buyer with nondelivery; if this happens then we assume that both parties derive zero utility. Hence, both parties bargain over the total surplus $\pi(q^{\mathrm{OUT}};\tilde{T})$ and so gain respectively shares φ and $1-\varphi$ of that surplus. In period 2, the supplier produces a nonzero quality if and only if $\varphi\pi(q^{\mathrm{OUT}};\tilde{T}) - c\gamma q^{\mathrm{OUT}} > W$. Clearly this means that it is optimal[4] for the buyer to specify a contractual payment $W = 0$. Hence the supplier chooses, once uncertainty is revealed, a quality level q^{OUT} that solves

4. By initially setting $W > 0$, the buyer can eliminate the supplier's incentives when the demand shock is low—that is, when $\varphi\pi(q^{OUT};\tilde{T}) - c\gamma q^{OUT} < W$. In this case, the supplier produces an asset of zero quality and hence the buyer loses opportunity for a profit equal to $(1-\varphi)\pi(q;\tilde{T})$. This is why the optimal discrete contract under separation specifies $W = 0$.

the following maximum problem:

$$\max_{q} \quad \varphi\pi(q;\tilde{T}) - \gamma c q. \tag{11}$$

The level of quality q^{OUT} is therefore given by

$$q^{\mathrm{OUT}} = \left(\frac{\varphi}{\gamma}\right)^{1/\varepsilon}\left(\frac{\alpha R}{c}\right)\tilde{T}. \tag{12}$$

Now q^{OUT} co-varies with the shock $\tilde{T}$ with exactly the same elasticity as in the first-best case: the adaptation goal is fulfilled. The problem, however, is that q^{OUT} is always smaller than q^{FB} by a factor of $(\varphi/\gamma)^{1/\varepsilon}$. This factor captures two effects: (i) compared to first-best, the supplier under separation faces increased production costs owing to lack of coordination; and (ii) compared to first-best, the supplier reaps only a fraction φ of the overall profits and hence produces a lower quality level.

Given expression (12) for q^{OUT}, we can calculate the expected sum of the buyer's and seller's profits:

$$\Pi^{\mathrm{OUT}} = (1 + \alpha(\sigma - 1))\left(\frac{\varphi}{\gamma}\right)^{\alpha(\sigma-1)}\left(\frac{R}{\sigma - 1}\right)E[\tilde{T}].$$

The ratio of Π^{OUT} to Π^{FB} has two components

$$\frac{\Pi^{\mathrm{OUT}}}{\Pi^{\mathrm{FB}}} = \underbrace{(1 + \alpha(\sigma - 1))\,\varphi^{\alpha(\sigma-1)}}_{\text{hold up}}\ \underbrace{\gamma^{-\alpha(\sigma-1)}}_{\text{coordination}}.$$

The first term is less than 1, which makes Π^{OUT} smaller than Π^{FB}. Because the supplier does not internalize the entire gains of increasing quality, she produces on average a lower quality than the fixed level of quality chosen under a rigid contract. Hence, holdup under vertical separation increases the unit costs of production. This first term approaches unity when φ does; in this case, the holdup cost of separation disappears because the supplier obtains the entire surplus produced by her effort q. The second term also makes Π^{OUT} smaller than Π^{FB} and stems from our ad hoc extra cost of vertical separation. We interpret this term as summarizing the cost advantages of vertical integration deriving from transport or coordination inefficiencies.

Consider now the case of a rigid contract $(\bar{q}, \bar{W})$ specifying ex ante the quality $\bar{q}$ and the payment $\bar{W}$. From our previous discussion it should be clear that such a solution is not appealing under separation. Indeed, with a rigid contract, the supplier has no ex post bargaining power (because the quality and the price of the intermediate input are fixed ex ante by contract) and still faces the extra production cost γ due to coordination. The situation is similar to that under integration with a rigid contract, except now the total cost of production is larger. This clearly shows that rigid contracts under separation are always dominated by other organizational structures.

According to a strict interpretation of property rights theory à la Grossman and Hart (1986), the issue of vertical integration versus separation is linked to the presence of non-contractual dimensions of the relationship. As soon as a contract is signed, the issue of ownership structure should become irrelevant. This is not the case here given our ad-hoc assumption that separation is technologically dominated by integration: $\gamma > 1$. This simplifies the presentation but unparsimoniously uses two distinct theories of vertical integration: incomplete contract theory and technological assumption. In Appendix A we present a similar model that uses only ingredients from Grossman and Hart's theory and thus in which

$\gamma = 1$. However, the allocation of ownership still matters in the presence of a contract, because in this alternative version we assume an additional source of incompleteness in the rigid contract: suppliers face only limited liability for not following the contract terms. A supplier who owns the intermediate input may be willing to deviate from the initial contract because large punishments for doing so are ruled out. This creates a second moral hazard problem: the supplier may, or may not, comply with the terms of the contract. To make sure that she does, the buyer must pay an agency rent (in addition to compensation for effort) that is larger when the supplier owns the intermediate input. This makes vertical separation more costly than integration when rigid contracts are optimal. Conversely, as in the model presented in the main text, flexibility cannot be achieved by integration and so requires separation.

2.5. Choice of Ownership Structure: Determinants

The buyer chooses vertical separation when the expected value of doing so is no less than that under integration—that is when $\Pi^{\mathrm{OUT}} \geq \Pi^{\mathrm{IN}}$ where

$$\frac{\Pi^{\mathrm{OUT}}}{\Pi^{\mathrm{IN}}} = \underbrace{(1+\alpha(\sigma-1))\varphi^{\alpha(\sigma-1)}}_{\text{(I) hold up}} \; \underbrace{\gamma^{-\alpha(\sigma-1)}}_{\text{(II) coordination}} \; \underbrace{E[\tilde{T}] \,\big/\, E[\tilde{T}^{\varepsilon}]^{1/\varepsilon}}_{\text{(III) flexibility}}; \tag{13}$$

the effects (I), (II), and (III) have been discussed previously.

In terms of comparative statics, an increase in uncertainty increases the flexibility gain (term (III) in (13) rises) and thus pushes the firm toward more vertical separation. The intuition is that uncertainty reduces the average fit of the quality level $\bar{q}$ to demand and thus reduces profits under vertical separation. A second important consequence is that an increase in the supplier's bargaining power makes separation more likely because it reduces holdup inefficiencies under separation (term (I) in (13) becomes larger) but does not affect flexibility.

The following proposition summarizes the predicted direction of the main determinants of vertical integration in this model.

PROPOSITION 1. *Vertical separation is promoted by:*

1. *an increase in the supplier's ex post bargaining power φ; and*
2. *a mean-preserving spread of the distribution of demand shocks $\tilde{T}$.*

The proof of this proposition is straightforward. The first result follows because the term (I) in (13) is an increasing function of φ. The second result follows because $\varepsilon < 1$ and so $\tilde{T}^{\varepsilon}$ is a concave, increasing function of $\tilde{T}$. Hence, a mean-preserving spread of $\tilde{T}$ reduces term (III) in (13).

Vertical separation or outsourcing to manage uncertainty is frequently used by businesses (see Lacity et al. (1995) for the case of IT departments). In the new economy, a well-known example of this relation is given by the success of Dell in selling and shipping personal computers to its customers (Magretta 1998). Dell is the entity that coordinates a tight network of parts makers that are kept well informed about the quantities needed and the anticipated dates of deliveries. Suppliers are not owned by Dell, but the relation that binds them to the PC assembler is tight and cooperative. To some extent, Dell managed—through a lower γ (better coordination with outside suppliers) and a larger φ (a commitment to leave a reasonable part of the surplus to suppliers)—to outsource a large part of its production process. The principal gain of such a structure, which Mr. Dell calls virtual integration, is reduced time to market and thus a better adaptation to demand changes.

Although some practitioners view vertical separation as a means of adapting to uncertainty, it must be noted that existing theories on the link between vertical separation and uncertainty predict the exact opposite. Carlton (1979) proposed a model in which vertical integration enables securing a steady supply of inputs. Hence, in Carlton's model, a larger level of uncertainty in the upstream part of production promotes integration (not separation), which then reduces uncertainty. Overall, however, the obtained reduction of uncertainty is lower than the initial increase, and the direct effect dominates: uncertainty and integration should be positively correlated. In our empirical section we find the exact opposite: a *negative* correlation between uncertainty and integration, in line with our own model.

A last comparative static is competition. In this model, competition reduces markups (the price elasticity σ of demand). How is the decision concerning vertical separation affected by an increase in competition? We must specify the distribution of $\tilde{T}$ before proceeding any further. A natural assumption in this framework is to assume that $\tilde{T}$ is log-normal (i.e., $\log \tilde{T} \sim N(0, \Sigma)$). Under that assumption,[5] the gains from separation may be written

$$\frac{\Pi^{\mathrm{OUT}}}{\Pi^{\mathrm{IN}}} = \underbrace{(1+\alpha(\sigma-1))\,\varphi^{\alpha(\sigma-1)}}_{\text{(I)}}\ \underbrace{\gamma^{-\alpha(\sigma-1)}}_{\text{(II)}}\ \underbrace{\exp\left\{\frac{\alpha(\sigma-1)}{1+\alpha(\sigma-1)}\frac{\Sigma^2}{2}\right\}}_{\text{(III)}}. \tag{14}$$

Equation (14) reveals that competition has an ambiguous net effect on the vertical separation decisions.

- The hold-up cost (I) is non-monotone in σ.
- The cost of coordination (II) is decreasing in σ; through this effect tougher competition make separation less attractive.
- On the other hand, competition makes adaptation more critical (i.e., term (III) is an increasing function of σ): As competition increases, the cost of misallocating production response to demand becomes more important and thus flexibility matters. This effect makes separation more attractive.

2.6. Choice of Ownership Structure: Consequences

A less obvious consequence of our model is that vertical separation itself increases downstream uncertainty regarding output, labor demand, and profits while pricing strategy tends to become less sensitive to demand. As we shall see, these results arise not from a simple size effect but rather because the *elasticity* of output and profits to demand shocks become higher under vertical separation.

For each organizational structure, the labor demand is given by

$$\tilde{l}^{\{\mathrm{IN,OUT}\}} = \left(\frac{(\sigma-1)(1-\alpha)}{\sigma w}\right)^{\sigma\varepsilon} q^{1-\varepsilon}\tilde{T}^{\varepsilon}, \tag{15}$$

5. Assuming small shocks yields similar results.

where q^{IN} is given by (9) and q^{OUT} by (12). From this we easily obtain the variances of labor demand under integration and separation:

$$\mathrm{Var}(\log \tilde{l}^{\mathrm{IN}}) = \varepsilon^2 \, \mathrm{Var}(\log \tilde{T}), \tag{16}$$

$$\mathrm{Var}(\log \tilde{l}^{\mathrm{OUT}}) = \mathrm{Var}(\log \tilde{T}).$$

Observe that in equations (16), the pure size effects are removed (i.e., we look at the log). Since $\varepsilon < 1$, we see that the relative variance of labor demand is unambiguously larger under vertical separation. This result follows because marginal productivity of labor is an increasing function of the other input—a property shared by standard production functions. Hence vertical separation promotes more variation in labor demand because it exhibits more variation in the other input, the intermediate input's quality q.

We may similarly derive the variance (controlling for size effect) of total output under each organizational structure as follows:

$$\mathrm{Var}\left(\log \tilde{y}^{\mathrm{IN}}\right) = \varepsilon^2 (1-\alpha)^2 \, \mathrm{Var}(\log \tilde{T}), \tag{17}$$

$$\mathrm{Var}\left(\log \tilde{y}^{\mathrm{OUT}}\right) = \mathrm{Var}(\log \tilde{T}).$$

Since $\varepsilon < 1$, it follows that overall output also fluctuates more under vertical separation. Given the demand curve (2) and the production levels, the variance for prices are

$$\mathrm{Var}\left(\log \tilde{p}^{\mathrm{IN}}\right) = (\varepsilon/\sigma)^2 (1-\alpha)^2 \, \mathrm{Var}(\log \tilde{T}),$$

$$\mathrm{Var}\left(\log \tilde{p}^{\mathrm{OUT}}\right) = 0.$$

The vertically separated firm can fully adapt its production level to demand

and thus always charges the monopoly price, which is equal to a constant markup over marginal costs (assumed to be constant here). In contrast, the vertically integrated firm cannot fully adapt its production to the shock: the demand curve moves up (or down), it cannot increase production by as much as would be possible for a vertically separated firm because the intermediate input quality $\bar{q}$ has been contractually fixed. Thus, when the demand shock exceeds its average, the firm does not increase production enough and its price rises (to a higher level than under monopoly markup)—so the firm is constrained by its production capacity. Conversely, when demand is below its average, excess production capacity makes it relatively costless to expand production above the monopoly level and the price goes down. Thus, price co-varies more with demand under vertical integration.

All things considered, vertical separation amplifies volatility in quantity and dampens the shocks to prices. The overall effect on *nominal* sales is therefore less strong than it is on real quantities (of output, but also of inputs such as labor). This difference between nominal and real behavior is present in the data, as we shall will see.

We summarize these results in our next proposition:

PROPOSITION 2. *Vertical separation amplifies the fluctuations of quantities and attenuates the volatility of prices. More precisely, at the firm level, the switch to separation induces:*

1. *an increase in the volatility of output and labor demand; and*
2. *a decrease in the volatility of prices.*

These are not size effects; but rather, they arise through increased (resp. decreased) elasticity of these quantities with respect to the demand shock.

Proposition 2 contrasts with the commonly stated intuition that vertical separation allows large companies to transfer risk to their suppliers. Here vertical separation increases risk both for supplier and buyer both. Note that this result has been obtained analytically using a Cobb–Douglas production function. How robust is it to a more general production technology? Our guess is that the theoretical argument presented here works as long as there is some degree of complementarity between the inputs q and l. Indeed, ex post fluctuations in q will trigger optimal changes in l as long as the marginal productivity of labor is an increasing function of q—that is, when $\partial^2 F/\partial q \partial l > 0$, which holds for most production functions. In line with this intuition, the results obtained in the Cobb–Douglas case carry over in a Leontief framework: $F(q, l) = \min\{q, l\}$. But they vanish if we assume perfect substitutability: $F(q, l) = q + l$.[6]

3. Aggregate versus Idiosyncratic Uncertainty

Available evidence indicates that, although idiosyncratic uncertainty seems to have increased (see Campbell et al. (2001) on stock returns and Comin and Philippon (2005) on sales and labor), aggregate uncertainty has remained stable or has decreased slightly (Stock and Watson 2002). So the next stage of analysis con-

6. Unfortunately the problem cannot be solved analytically if CES production functions are used.

sists of looking, in our model, at the differential impact of vertical separation on aggregate and firm-level uncertainty. This requires that we embed the previous microeconomic framework within a simple general equilibrium model.

3.1. The General Equilibrium Model

At the firm-level, the situation is similar to the one described in Section 2. To simplify exposition we now give the production function by equation (1) with $\alpha = 1$:

$$f(q) = q.$$

In addition, producing a quality q necessitates hiring a number q of workers at a total cost of $c(q) = wq$. We assume that hiring is made on the spot labor market by suppliers after revelation of the demand shock. Finally, the coordination costs γ are stochastically distributed among the population of buyers according to a cumulative distribution function $G(\gamma)$ with $\gamma \geq 1$.

At the aggregate level, this is a closed and static economy. The product market is broken down into a continuum of varieties $i \in [0, 1]$, and each variety is produced by a monopolist (i.e. a buyer and a supplier) in quantity y_i. Varieties are used to assemble a final good Y such that

$$Y = \left(\int_0^1 \tilde{\tau}_i^{1/\sigma} y_i^{(\sigma-1)/\sigma} \, di \right)^{\sigma/(\sigma-1)}, \tag{18}$$

where $\sigma > 1$ and $\tilde{\tau}_i > 0$ are technological shocks affecting the marginal productivity of each variety i. Shocks have both a micro and a macro component. For the sake of simplicity we assume that these shocks can be broken down into two terms: an aggregate technological shock $\tilde{A}$ and a variety-specific idiosyncratic technological shock $\tilde{\delta}_i$:

$$\tilde{\tau}_i = \tilde{A}\tilde{\delta}_i, \tag{19}$$

where $\mathrm{Cov}(\tilde{A}, \tilde{\delta}_i) = 0$ and $\mathrm{Cov}(\tilde{\delta}_i, \tilde{\delta}_k) = 0$ for $i \neq k$.

Each variety i is supplied by a monopolist. Because of the final production technology (18), the demand faced by a monopolist i is given by

$$\tilde{y}_i = \tilde{\tau}_i \frac{\tilde{D}}{\tilde{P}^{1-\sigma}} \tilde{p}_i^{-\sigma}, \tag{20}$$

where the price index is $\tilde{P} = \left(\int_0^1 \tilde{\tau}_k \tilde{p}_k^{1-\sigma}\, dk\right)^{1/(1-\sigma)}$ and $\tilde{D}$ stands for the total expenditure dedicated to the purchase of the final good. This demand function is a particular case of the one studied in the previous section (see equation (2)) with a demand shock $\tilde{T}_i$ of

$$\tilde{T}_i \equiv \tilde{D}\tilde{\tau}_i\tilde{P}^{\sigma-1}. \tag{21}$$

The demand shock faced by each monopolist has two components: an exogenous shock $\tilde{\tau}_i$ and an aggregate *endogenous* shock $\tilde{D}\tilde{P}^{\sigma-1}$. Hence we must solve for the macroeconomic equilibrium in order to figure out how the price index affects the demand shock.

3.2. Firm-level Analysis

At the level of the firm, the problem is slightly different than presented in Section 2. Given that aggregate productivity may fluctuate, it may be that $\tilde{w}$ also fluctuates. This slightly changes our method of solving the model.

First we examine the case of an integrated firm. The buyer–supplier contract is rigid, and the supplier's labor demand solves

$$\max_q \quad E\left[\tilde{T}^{1/\sigma}\right] q^{1-1/\sigma} - E[\tilde{w}]q.$$

Hence labor demand, optimal price, and expected profits are given by the following equations:

$$\begin{aligned} \tilde{l}^{\mathrm{IN}} &= \left(\frac{\sigma-1}{\sigma}\right)^{\sigma}\left(\frac{E\left[\tilde{T}^{1/\sigma}\right]}{E[\tilde{w}]}\right)^{\sigma}, \\ \tilde{p}^{\mathrm{IN}} &= \frac{\sigma}{\sigma-1}E[\tilde{w}]\frac{\tilde{T}^{1/\sigma}}{E[\tilde{T}^{1/\sigma}]}, \\ \Pi^{\mathrm{IN}} &= \frac{(\sigma-1)^{\sigma-1}}{\sigma^{\sigma}}E[\tilde{w}]^{1-\sigma}E\left[\tilde{T}^{1/\sigma}\right]^{\sigma}. \end{aligned} \tag{22}$$

Labor demand is fixed, while prices must adjust ex post in order for the monopolist to be on the demand curve.

Under vertical separation the intermediate input's quality q is fully flexible and solves

$$\max_q \quad \varphi\tilde{T}^{1/\sigma}q^{1-(1/\sigma)} - \gamma\tilde{w},$$

which yields the following expressions for labor demand, prices, and expected profits:

$$\begin{aligned} \tilde{l}^{\mathrm{OUT}} &= \gamma\left(\frac{\sigma-1}{\sigma}\right)^{\sigma}\left(\frac{\varphi}{\gamma}\right)^{\sigma}\frac{\tilde{T}}{\tilde{w}}^{\sigma}, \\ \tilde{p}^{\mathrm{OUT}} &= \frac{\sigma}{\sigma-1}\frac{\gamma}{\varphi}\tilde{w}, \\ \Pi^{\mathrm{OUT}} &= \sigma\frac{(\sigma-1)^{\sigma-1}}{\sigma^{\sigma}}\left(\frac{\varphi}{\gamma}\right)^{\sigma-1}E\left[\frac{\tilde{T}}{\tilde{w}^{\sigma-1}}\right]. \end{aligned} \tag{23}$$

According to condition (13), the firm chooses vertical separation as soon as

$$\Pi^{\mathrm{IN}} \leq \Pi^{\mathrm{OUT}}.$$

This condition holds whenever the coordination cost γ is below a cutoff value $\bar{\gamma}$, which—using the previous condition and equations (22) and (23)—is defined via

$$\sigma\left(\frac{\varphi}{\bar{\gamma}}\right)^{\sigma-1} E\left[\frac{\tilde{T}}{\tilde{w}^{\sigma-1}}\right] = \left(\frac{1}{E[\tilde{w}]}\right)^{\sigma-1} E\left[\tilde{T}^{1/\sigma}\right]^{\sigma}.$$

By (21) we have that the demand shock is equal to $\tilde{T} = \tilde{A}\tilde{D}\tilde{P}^{\sigma-1}\tilde{\delta}$, where $\tilde{A}\tilde{D}\tilde{P}^{\sigma-1}$ and $\tilde{\delta}$ are independent processes. Hence the preceding condition may be rewritten

$$\sigma\left(\frac{\varphi}{\bar{\gamma}}\right)^{\sigma-1} = \underbrace{\frac{E[\tilde{\delta}^{1/\sigma}]^{\sigma}}{E[\tilde{\delta}]}}_{\text{(I)}} \underbrace{E\left[\frac{\tilde{A}\tilde{D}\tilde{P}^{\sigma-1}}{\tilde{w}^{\sigma-1}}\right]^{-1} \left(\frac{1}{E\tilde{w}}\right)^{\sigma-1} E\left[\left(\tilde{A}\tilde{D}\tilde{P}^{\sigma-1}\right)^{1/\sigma}\right]^{\sigma}}_{\text{(II)}}. \quad (24)$$

The decision to vertically separate thus depends on two factors: (i) idiosyncratic uncertainty (term (I) in the equation), just as in Section 2; and (ii) aggregate uncertainty (term (II)). We thus need to solve for the general equilibrium, which will tell us how $\tilde{w}$ is related to $\tilde{A}\tilde{D}\tilde{P}^{\sigma-1}$.

Finally, it is useful to remark that the degree s of vertical separation within the economy corresponds to the share of firms that choose to be vertically separated

$$s = G(\bar{\gamma}) = \int_1^{\bar{\gamma}} dG(\gamma),$$

where $\bar{\gamma}$ is given by (24).

3.3. Equilibrium Conditions

We need only write down two equilibrium conditions: the labor market clearing condition and the aggregate price determination. In these equilibrium conditions, the idiosyncratic uncertainty usually cancels out because of aggregation and the law of large numbers, but aggregate uncertainty remains.

Using equations (22) and (23), the aggregate labor demand can be easily computed and the market clearing condition can be written

$$L = \left(\frac{\sigma-1}{\sigma}\right)^{\sigma} \times \left(s\varphi\left(\frac{\varphi}{\Gamma}\right)^{\sigma-1} \frac{\tilde{A}\tilde{D}\tilde{P}^{\sigma-1}}{\tilde{w}^{\sigma}} E[\tilde{\delta}] + (1-s) \frac{E\left[(\tilde{A}\tilde{D}\tilde{P}^{\sigma-1})^{1/\sigma}\right]^{\sigma}}{E[\tilde{w}]^{\sigma}} E\left[\tilde{\delta}^{1/\sigma}\right]^{\sigma} \right), \quad (25)$$

where $\Gamma^{1-\sigma} = \int^{\bar{\gamma}} \gamma^{1-\sigma}\, dG(\gamma)/G(\bar{\gamma})$ is the average cost of coordination across vertically separated firms.

In the previous market clearing condition, all terms are deterministic with the exception of the stochastic terms $\tilde{A}\tilde{D}\tilde{P}^{\sigma-1}$ and $\tilde{w}^{\sigma}$; yet clearly the ratio of these stochastic variables must be deterministic too and equal to a scalar Λ. Indeed, we immediately obtain from condition (25) that

$$\frac{\tilde{A}\tilde{D}\tilde{P}^{\sigma-1}}{\tilde{w}^{\sigma}} = \frac{1}{L}\left(\frac{\sigma-1}{\sigma}\right)^{\sigma} \left(s\varphi\left(\frac{\varphi}{\Gamma}\right)^{\sigma-1} E[\tilde{\delta}] + (1-s)E\left[\tilde{\delta}^{1/\sigma}\right]^{\sigma} \right) \equiv \Lambda, \quad (26)$$

which is deterministic. This relationship is useful for our analysis because it allows us to rewrite condition (24) and so obtain the cutoff value $\bar{\gamma}$ below which firms choose vertical separation:

$$\sigma\left(\frac{\varphi}{\bar{\gamma}}\right)^{\sigma-1} = \frac{E\left[\tilde{\delta}^{1/\sigma}\right]^{\sigma}}{E[\tilde{\delta}]}. \quad (27)$$

Hence $s = G(\bar{\gamma})$, the degree of vertical separation within the economy, does not depend on the aggregate shock $\tilde{A}$. This is our first result.

In a similar way we can aggregate the firm-level prices as given by (22)–(23) and then use definition (20) to compute the price index; combined with (26), this gives the real wage and the total output $\tilde{Y}$:

$$\frac{\tilde{w}}{\tilde{P}} = \tilde{A}^{1/(\sigma-1)} \frac{\sigma - 1}{\sigma} \left(s \left(\frac{\varphi}{\Gamma} \right)^{\sigma-1} E[\delta] + (1-s) E \left[\tilde{\delta}^{1/\sigma} \right]^{\sigma} \right)^{1/(\sigma-1)}, \tag{28}$$

$$\tilde{Y} = \tilde{A}^{1/(\sigma-1)} L \left(s \left(\frac{\varphi}{\Gamma} \right)^{\sigma-1} E[\tilde{\delta}] + \ (1-s) E \left[\tilde{\delta}^{1/\sigma} \right]^{\sigma} \right)^{\sigma/(\sigma-1)}$$
$$\times \left(s\varphi \left(\frac{\varphi}{\Gamma} \right)^{\sigma-1} E \left[\tilde{\delta} \right] + (1-s) E \left[\tilde{\delta}^{1/\sigma} \right]^{\sigma} \right)^{-1}. \tag{29}$$

The macroeconomic equilibrium is fully described by equations (27)–(29). We see that the elasticities of real wage $\tilde{w}/\tilde{P}$ and total output $\tilde{Y}$ with respect to $\tilde{A}$ do not depend on s, the aggregate degree of vertical separation; that is, Var(log $\tilde{w}/\tilde{P}$) and Var(log $\tilde{Y}$) do not depend on s. This is our second result.

Finally, by using equations (21), (22), (23), and (26), it is straightforward to compute firm-level labor demands under integration and separation:

$$\tilde{l}^{\mathrm{IN}} = \left(\frac{\sigma - 1}{\sigma} \right)^{\sigma} \Lambda E \left[\tilde{\delta}^{1/\sigma} \right]^{\sigma},$$
$$\tilde{l}^{\mathrm{OUT}} = \gamma \left(\frac{\sigma - 1}{\sigma} \right)^{\sigma} \left(\frac{\varphi}{\gamma} \right)^{\sigma} \Lambda \tilde{\delta}.$$

Here Λ is given by (26) and $\tilde{\delta}$ is the idiosyncratic shock. Hence the elasticity of labor demand at the firm-level (i.e., Var(log $\tilde{l}$)) depends only on an idiosyncratic uncertainty and not on an aggregate one. This is our third result. Observe that a similar analysis could be performed in terms of cash flows and output.

We are now in a position to establish the following proposition.

PROPOSITION 3. *There exists an unique equilibrium with the following properties:*

1. *Aggregate uncertainty does* not *affect the decision of vertical separation:*

The cutoff $\bar{\gamma}$ and the share s do not depend on $\tilde{A}$.

2. *Aggregate uncertainty is* not *amplified by the degree of vertical separation: The elasticities of total output $\tilde{Y}$ and real wages $\tilde{w}/\tilde{P}$ with respect to $\tilde{A}$ do not depend on s, the fraction of vertically separated firms.*
3. *Only the idiosyncratic part of demand shocks is amplified by vertical separation at the firm level; the elasticities of cash flows, labor demands, and sales depend on $\tilde{\delta}$ but not on $\tilde{A}$.*

Vertical separation does not allow one to hedge against aggregate risk. The economic intuition is as follows. The share s of vertically separated firms compete ex post with each other for labor from a given pool of workers that is equal to the total supply L minus the predetermined labor demand from the share $(1-s)$ of vertically integrated firms. Hence, the total labor supply available to such firms is already fixed once the shocks are known. Assume for simplicity that there is no idiosyncratic uncertainty. When demand $\tilde{A}$ is high, all vertically separated firms simultaneously ask for more labor (in the form of intermediate input quality $\tilde{q}$). But since the residual labor supply available for these firms is fixed, no single firm receives more labor than another and so the shock $\tilde{A}$ cannot be accommodated. Hence, because the real wage co-moves with aggregate uncertainty $\tilde{A}$ (see equation (28)), the gains of flexibility due to vertical separation are completely canceled out by the ex post scarcity of labor. A direct consequence is that the *elasticity* of aggregate production with respect to aggregate demand shocks does not depend on the degree of vertical separation s. Equation (29), however, shows that the *level* of production depends in an ambiguous way on s. The reason is that a larger s means more vertical separation and flexibility; this induces an ex post better allocation of labor across firms, which is good for aggregate efficiency. On the other hand, the unit cost of production is higher on average owing to hold up inefficiencies and coordination costs.

4. Empirical Evidence

4.1. Data and Measurement Issues

Throughout this empirical section, we use two French firm-level data sets. Our main source of information, the *Bénéfices Réels Normaux*, available from the French statistical office (INSEE), is a fiscal database that is used to levy the corporate income tax an all French firms whose total sales exceed some 750,000 euros. We focus here on firms (a) that are recorded in the database for at least three of the years from 1984 to 1999 and (b) that have at least 500 employees or 30 million euros in sales. We then follow these firms over the entire period and focus on those that are continually present from 1984 until 1999. This leaves us with a balanced panel of 2,315 firms (37,040 observations). Note, however, that our empirical results are robust to including smaller firms and to including firms that leave or enter the data set. At the firm-level, the tax files provide us with accounting information (balance sheet and income statement) as well as total employment and type of industry (using an industry classification similar to the 4-digit SIC codes of the U.S. statistical system). These data are extremely noisy (mostly because of errors induced during data entry), so we windsorize most variables to the top and bottom 5%.

Our second source of information is ACEMO, a quarterly survey conducted between 1987 and 1997 by the French Ministry of Labor on some 12,000 firms. This data set provides us with figures on quarterly employment, short-term contracts, part-time workers and interim employment, and from this we extract employment based measures of uncertainty.

We investigate here the causes and consequences of vertical separation, relying on the helpful specificity of French corporate accounting norms. In French accounts, costs of goods sold are routinely broken down into personal expenses and “intermediate consumption”. Intermediate consumption is the total cost of raw materials and semi-finished goods that the firm has purchased. Put differently, intermediate consumption is the difference between sales and value added. In this section, we measure vertical separation as *the ratio of the firm’s intermediate consumption to sales*, which is equivalent to 1 minus the ratio of value added to sales.

This measure unfortunately includes both raw materials and semi-finished goods, so some of the cross-sectional variation in our measure of separation may in fact be simply due to industry specificities. For instance, taking our measure at face value would suggest that both concrete factories and car manufacturers make a considerable use of vertical separation. Yet concrete factories use mostly raw materials, whereas the automobile industry uses mostly semi-finished goods. The second drawback of this measure is that it is not based on quantities but rather on values (which are part of the firm's income statement). Hence, the ratio intermediate of consumption to sales might fluctuate owing to fluctuations in the relative price of inputs and goods sold.

However, we believe that most of these limitations are accounted for by including industry and year effects in all of our specifications. Indeed, we focus hereafter on the difference between the firm's degree of separation and the industry's average degree of separation. Moreover, fluctuations in relative prices should be captured by year dummies as well as by industry dummies interacted with year effects. These controls appear in most of our specifications. Finally, we remark that the existing (and small) empirical literature on outsourcing does not do much better. Following Feenstra (1998), a number of recent empirical studies on vertical integration have used similar measures (Acemoglu et al. 2004; Macchiavello 2006). The main reason is that quantities of inputs and output are usually not reported in financial statements, which are the main source of information used to compile databases. In addition, there is no large data set on contracting relations between firms and their suppliers.

4.2. Trends in Vertical Separation and Volatility

Our theory relates vertical separation, firm-level output, and volatility of labor demand. Do recent macroeconomic trends support the claims of such a theory? Using French National Accounts, we are in a position to compute the ratio of intermediate input consumption to total output. Results are depicted in Figure 1.

[Figure 1 here]

Thus measured, economy wide vertical separation increases throughout the 1980s and 1990s: from 85% to almost 100% of Gross Domestic Product (GDP). Around this overall trend, our aggregate measure of vertical separation exhibits cyclical movement: increasing during booms (e.g. the late 1980s and late 1990s), and decreasing during slowdowns. This aggregate evolution, however, conceals a considerable intra-industry heterogeneity. To get an idea of composition effects, consider the industry breakdown shown in Table 1. The most impressive figures are seen in the automobile industry, where the share of intermediary inputs in total sales has increased by 14.9 percentage points between 1978 and 2000. Then come consumer goods and food processing, where average reliance on intermediate input has increased by some 9.5 percentage points. This figure has decreased in three industries only (energy, equipment and agriculture), but most of the others have experienced a sizable increase. The timing of these changes varies across industries. In most manufacturing industries as well as in financial services and transportation most of the increase occurred in the 1980s. A moderate acceleration in vertical separation in the 1990s is evident in real estate, trade and construction.

[Table 1 here]

To document the rise in sales and employment volatility, we exploit our firm-level data set. In Figure 2 we use the methodology of Comin and Philippon (2005) to compute, for each 4-digit industry, its cross-sectional variance of firms' sales growth; we then average those variances across industries by weighting each industry by its total sales. The same methodology is applied to employment. Hence we are left with two yearly measures of cross-sectional volatility of sales and employment growth (the cross-sectional confidence intervals are also plotted). Sales variance increases markedly starting in the mid-1980s, from 23% to 36%, while variance of employment increases from 15% to 24%.

[Figure 2 here]

Those findings are confirmed by Givord and Maurin (2004), who use French data to show that the probability of job loss increased between the 1980s and the 1990s even after proper accounting for macroeconomic fluctuations. This trend toward increased job turnover and sales volatility has also been documented by Comin and Philippon (2005) for the United States and by Neumark (2000) for several developed economies.

4.3. Evidence from Correlations

At the microeconomic level, our theory has a two-stage prediction:

1. demand shocks increase vertical separation;
2. vertical separation, in turn, amplifies the elasticity of employment to demand shocks.

Consequently—and in contrast to existing models that link vertical separation and uncertainty—our theory predicts a positive correlation between vertical separation and employment volatility. Observe that the model makes the same prediction for output quantities, which we cannot measure at the firm-level. Instead, accounting data provide us with *nominal* sales (i.e., output quantity times unit price). Our model suggests that vertically separated firms smooth demand fluctuations via their pricing policy. As a result, the correlation between separation and the variance in nominal sales should be weaker than the correlation between separation and employment fluctuation.

We estimate at the firm-level the following OLS specification:

$$\mathrm{Var}_t(Y_i) = \beta\,\mathrm{SEPARATION}_{it} + \gamma\,\mathrm{CONTROLS}_{it} + \varepsilon_{it}. \tag{30}$$

The dependent variable $\mathrm{Var}(Y_i)$ is the firm-level variance of either sales growth or employment growth as computed over the various subperiods of 1984–1999. The main explanatory variable, $\mathrm{SEPARATION}_{it}$, is our measure of vertical separation and is computed as intermediate consumptions divided by total sales, i.e. is averaged over the corresponding subperiods. If our assumptions are correct then the variance should be increasing in separation: $\beta > 0$. Estimates for various specifications of equation (30) are shown in Table 2, where we focus on the variance of sales growth as the dependent variable.

[Table 2 here]

In columns 1 and 2, we explore the pure cross-sectional correlation. For each firm, the variance of sales growth is computed over the 1984–1999 period and then the measure of separation is averaged for that period. With this methodology there is one observation per firm. The cross section specification includes 2-digit industry fixed effects and controls for firm's average level of sales. As expected, the coefficient of vertical separation is positive and significant at the 1%

level. Note that the effect estimated in column 1 is quite large. Indeed, an increase of one standard deviation in vertical separation translates into an increase of 2% in sales variance; this corresponds to one fourth of the cross-sectional standard deviation in sales variance.[7] Column 2 breaks the sample down further into two subperiods and includes two-digit industry fixed effects interacted with period dummies. Over 16 years, many factors have changed (production technologies, international and domestic competition, firm organizations) that could have affected both the propensity to outsource and the propensity of firms to explore uncertain markets. These dummies allow us to control for long-run changes in volatility related to industry-wide diffusion of competition and technology adoption. The estimated correlation between vertical separation and sales variance is unchanged (0.14 instead of 0.12).

Column 3 presents fixed-effect estimates. To obtain a time dimension, we compute the average level of separation and the variance of sales over two subperiods: 1984–1991 and 1992–1999. This allows us to include firm effects in equation (30) and thereby control for unobserved (though fixed) factors—firm organization, or the specifics of the market—that could simultaneously increase the likelihood of vertical separation and the variance of sales. Fixed unobserved heterogeneity is, however, not likely to be the only source of bias. As in column 2, we also include industry–period interactions to capture industry-wide trends.

---

7. Regarding the subsample of specification 1, the variance of sales has a mean of 16% and a standard deviation of 13% while vertical separation has a mean of 61% and a variance of 18%.

It turns out that the parameter estimate obtained in column 3 is half the size of the estimate of column 2 and statistically insignificant. We then repeat the analysis of columns 2 and 3 using three subperiods (1984–1989, 1989–1994, 1995–1999), instead of two. Again, the cross-sectional correlation is strong and significant while the fixed-effects estimate is both small and insignificant.

One potential problem is the noise surrounding our measure of sales variance: the equations estimated in columns 2 and 3 use variances computed with seven data points per firm. In columns 4 and 5, we use five data points! This may explain why our correlation is not robust with respect to introducing fixed effects. One other possibility is that firm sales is a *nominal* measure of output, whereas our model's strongest predictions concern *real* quantities (of output and of inputs such as employment). This is because firms that utilize vertical separation can attenuate the fluctuations of real output with opposite movements in prices.

The estimates reported in Table 3 accounts for these issues, by (i) looking at the variance of employment instead of nominal sales and (ii) finding alternative and more reliable short-term estimates of employment volatility. As in Table 2, columns 1–3 simply compute employment volatility as the variance of employment growth taken over various subperiods. Columns 4 and 5 provide alternative estimates of volatility: column 4 uses quarterly data to compute the variance of employment for each year; column 5 uses the share of workers employed under interim contracts as a proxy for volatility.

[Table 3 here]

In column 1 of Table 3, we explore cross-sectional correlation. The variance of employment is computed for the entire 1984–1999 period and industry-level fixed effects are included. Vertical separation is averaged over the period. The correlation between employment variance and vertical separation is statistically significant. The magnitude of the effect is similar to what we see for sales: here, an increase of one standard deviation in vertical separation translates into an increase of 3% in employment variance; this corresponds to one fifth of the cross-sectional standard deviation in employment variance.

In column 2, the variance is computed over two separate subperiods (1984–1991 and 1992–1999). We include industry dummies interacted with period dummies to account for industry-specific trends; results are similar in terms of size and statistical significance to those reported in column 1. Column 3 adds firm fixed effects to the specification in column 2. As in Table 2, the estimated correlation becomes smaller and insignificant.

In columns 4 and 5 we use information from the ACEMO survey (available from 1987 to 1997) to compute annual measures of employment volatility. These specifications include both firm fixed effects and year dummies. To account for industry-specific trends, they also include industry dummies interacted with three period dummies (1984–1989, 1990–1994, 1995–1999). Column 4 uses firm-level quarterly employment data. The dependent variable is then the variance of employment using the four figures available each year. Column 5 uses the share of workers hired under interim contracts, which is also available each year. This proxy assumes that "flexible" workers (such as those hired under fixed-term contracts) are less productive (because they expend less effort, as in Saint Paul, 1995), although their jobs are easier to terminate than long-term contracts. In this case, the share of interim workers is a measure of the relative return of flexibility, and it increases when the firm wishes to adjust employment more. Conditional on that model being true, the share of “flexible” workers should be a good proxy for expected employment variance. Correlations estimated using these two annual measures of employment volatility are significant, from both an economic and a statistical viewpoint.

4.4. Foundations for an Alternative Estimation Strategy

Section 4.3 highlights that vertical separation and firm-level uncertainty tend to move together—both in the cross section of firms and in the panel dimension, i.e. controlling for firm fixed effects. This lends credence to our theory because the positive correlation is one of its predictions. In the rest of our empirical investigation we address a more causal relationship by focusing on the amplification mechanism, which predicts that any exogenous increase in vertical separation should induce an amplification of firm-level elasticity of sales/employment to demand shocks.

The ideal empirical strategy would involve an exogenous source of variation in vertical separation that is *not* related to firm-level uncertainty. Unfortunately, we do not have such an instrument and so our only course is to use lagged separation. More specifically, we ask whether a past increase in vertical separation predicts a future rise in firm uncertainty. This approach, akin to a Granger causality test in times-series econometrics, cannot itself rule out reverse causality. For instance, it could well be that firms raised their separation level in the past because they expected uncertainty to increase in the future. We address this concern by looking at the effect of a change in separation lagged by two or three years.

However, it is difficult to use this lag approach with the specifications of Section 4.3, because computing variances requires that we collapse our data set into two subperiods at most. We propose here an alternative methodology that allows us to compute an *annual* measure of firm-level volatility. The intuition—as may be derived directly from our theoretical model—is to look at the sensitivity of the firm's activity to industry demand shocks: we ask if a past increase in separation predicts a future increase in the sensitivity of firm activity to industry sales.

This section explicitly derives from our theoretical model a reduced form that we put to the empirical test in Section 4.5. We extend our main model to a case where firms must perform a continuum of tasks—instead of just one—in order to produce their final output. In this simple extension, the production function becomes

$$\ln y = \int_0^1 \ln(q_j)\, dj,$$

where tasks are indexed by $j \in [0, 1]$. Each task j requires only a machine of quality q_j. All tasks contribute to the production of final output through a Cobb–Douglas technology. This extension is useful for bringing the model to the data because it allows us to think of vertical separation as a continuous decision. In the data, our measure of separation is continuous and ranges between 0 and 1. In this model, we will assume that firms can choose to outsource machine production q_j

for as many tasks as they wish. To ensure that firms choose interior solutions, we simply assume that some tasks are much more difficult to outsource than others: without loss of generality we assume that γ_j, the coordination cost of vertical separation task j, increases "enough" in j. The rest of the model is identical. In particular, firms operate as monopolies when the demand curve is subject to a demand shock $\tilde{T}$.

In such a setup, firms choose to outsource all tasks for which coordination costs are low enough—that is, those tasks in some interval $[0, r]$. Is is easy to calculate firm sales as a function of the demand shock (see e.g. Thesmar and Thoenig 2003):

$$py = \lambda \tilde{T}^{\varepsilon(r)}, \tag{31}$$

where $\varepsilon(r)$ is an increasing function of r, the fraction of outsourced tasks. As a result, the activity of firms make intensive use of vertical separation tends to be more elastic to demand shocks. This result is conceptually identical to that of Section 2.6. The only formal difference is that breaking down production into tasks allows us to look at *continuous* vertical separation decisions, which more nearly resemble to what we observe in the data.

To make the model empirically operational, we need a proxy for $\tilde{T}$. We consider industry shocks and assume that $\tilde{T}$ can be broken down into a firm-specific shock $\tilde{\delta}$ and a common industry component $\tilde{A}$. Let $\tilde{T} \equiv (1 + \tilde{\delta})(1 + \tilde{A})$, where $\tilde{A}$ and $\tilde{\delta}$ are small and distributed around 0. Using a first-order Taylor expansion, equation (31) can be written

$$py = \lambda \left(1 + \varepsilon(r)\tilde{\delta} + \varepsilon(r)\tilde{A}\right). \tag{32}$$

We aggregate this equation at the industry level, and we assume there are enough firms in the industry for the idiosyncratic terms $\tilde{\delta}$ to cancel out. Industry sales Y can then be written

$$PY = \lambda N \left(1 + \overline{\varepsilon(r)}\tilde{A}\right), \tag{33}$$

where N is the number of firms in the industry and $\overline{\varepsilon(r)}$ is the industry average of $\varepsilon(r)$.

Next, we log linearize (32) and (33) and then rearrange to obtain

$$\log py = \underbrace{f(r)}_{\text{level effect of } r} + \underbrace{\frac{\varepsilon(r)}{\overline{\varepsilon(r)}}}_{\substack{\text{sensitivity to} \\ \text{industry sales}}} \underbrace{\log(PY)}_{\text{industry sales}} + \underbrace{\varepsilon(r)\tilde{\delta}}_{\text{error term}}, \tag{34}$$

where $f(r)$ is a function of r. This equation is intuitive. Firms that are less vertically integrated are more sensitive to demand shocks and so in particular are more sensitive to the common industry shocks, which are well approximated in the data by total industry sales. Hence the data should confirm that firms outsourcing more than their peers are more sensitive to industry shocks.

Our empirical equation is a linearized version of (34). For firm i operating in industry s at date t, we write

$$\log \text{SALES}_{it} = \alpha_i + \gamma_t + \beta_{it} \log \text{SALES}_{st} + \varepsilon_{it}, \tag{35}$$

$$\text{with } \beta_{it} = \beta_0 + \beta_1 \text{XS\_SEPA}_{it}, \text{ and}$$

$$\text{XS\_SEPA}_{it} = \text{SEPA}_{it} - \overline{\text{SEPA}}_{st},$$

where SALES_{it} and SALES_{st} denote firm and total industry sales and where β_{it} is the sensitivity of firm i to industry sales. The variable $\overline{\text{SEPA}}_{st}$ stands for the average level of separation in the industry, and XS_SEPA_{it} is the excess separation—that is, the difference between firm and industry level of separation. We also include a firm fixed effect α_i. Equation (35) simply states that firms with above-average vertical separation should be more sensitive to industry shocks.

In order to fully exploit the panel dimension of our annual data set, we include both firm fixed effects and industry-specific time trends in the specification of β_{it} and we consider lags (by k years) instead of contemporaneous values of vertical separation. Thus we finally estimate, instead of (35), the following equation:

$$\log \text{SALES}_{it} = \alpha_i + \gamma_t + \beta_{it} \log \text{SALES}_{st} + \varepsilon_{it}, \tag{36}$$

$$\text{with } \beta_{it} = \beta_i + \beta_s t + \beta_1 \text{XS\_SEPA}_{i,t-k},$$

where β_i is the firm fixed tendency to be sensitive to industry sales (and is thus distinct from α_i) and β_s is the industry-wide trend on sensitivity to sales. As in Section 4.3, β_i is included to capture the impact of the diffusion of new organization, competition, or technology on firm-level volatility.

Unfortunately, equation (36) is difficult to estimate directly. Fixed effects in sensitivity β_i add a large number of variables to the right-hand side (one per firm, or more than 2,000). Traditional methods with linear equations in panel data circumvent this difficulty by differentiating equations for a given firm. This trick allows one to take out fixed effects (this is the usual "within" estimator in panel data) but does not work when equations are non-linear. In Appendix

B we propose a method designed to account for fixed effects β_i. Our appendix calculations lead to the following equation, which combines the equations in (36) as follows:

$$\begin{aligned}\log \mathrm{SALES}_{it} = \gamma_t + \alpha_i' + \left(\overline{\log \mathrm{SALES}_i} \div \overline{\log \mathrm{SALES}_s}\right) \log \mathrm{SALES}_{st} \\ + \beta_1 \left(\mathrm{XS\_SEPA}_{i,t-k} - \overline{\mathrm{XS\_SEPA}_i}\right) \log \mathrm{SALES}_{st} + \varepsilon_{it},\end{aligned} \tag{37}$$

where the overline means "average across years t within a given firm i" or "within a given industry s". This equation bears a striking similarity to standard "within" estimators for panel data. In order to account for fixed sensitivity to industry sales, we need to perform two transformations of the standard equation. First, we must include the first term of the RHS of (37), which may be interpreted as the average firm reaction to industry shocks over the period. Second, instead of focusing on excess separation, we need to include the difference between the current year's excess separation and its average value.

4.5. Does Vertical Separation Amplify Uncertainty?

The virtue of the model just described is that it allows us to look at lags in vertical separation. Before estimating reactions to industry shocks, we must check that vertical separation is indeed persistent at the firm-level. We do this in Table 4, where the following equation is estimated:

$$\mathrm{SEPA}_{it} = a_i + b\,\mathrm{SEPA}_{i,t-k} + \varepsilon_{it}, \tag{38}$$

for $k = 1, 2, 3$. Including fixed effects is important because they allow to ask whether or not vertical separation is persistent or not *at the firm-level*. Asking if

$b > 0$ amounts to checking whether, for any given firm, sudden increases in the level of separation persist or diminish. Strictly speaking, standard within estimators of such autoregressive models are inconsistent because they are polluted by the fixed effects a_i. Obtaining an unbiased estimate of b would require finding instruments for lags and using a GMM procedure. Yet, under the null hypothesis that $b = 0$ (no persistence at all), within estimates of b are unbiased—the model is not autoregressive anymore. The t-statistic should in this case be low. Thus, within estimates of the equation (38) already tell us something about persistence by allowing us to reject the hypothesis that there is none.

[Table 4 here]

Table 4 presents the within estimates of equation (38). As it turns out, the coefficient b on separation lagged by one, two, or three years is always positive and highly significant. This does not tell us much about their real size or statistical significance, but it does suggest that we can safely reject the null that there is no persistence. On the contrary, increasing vertical separation yields significant effects for at least three years. However, persistence then fades quickly as we take longer lags.

Let us now turn to the analysis of the effect of industry shocks on firms' nominal sales. Before including either lags or fixed effects in sensitivity, we first report results of equation (36) when taking the contemporary value of vertical separation ($k = 0$) and assuming no fixed effects. The latter means that β_i becomes constant across firms; we denote the common value by β_0. This yields

$$\log \mathrm{SALES}_{it} = \alpha_i + \gamma_t + \beta_1 \,\mathrm{XS\_SEPA}_{it} \times \log \mathrm{SALES}_{st} + \beta_0 \log \mathrm{SALES}_{st} + \beta_s t \log \mathrm{SALES}_{st} + \varepsilon_{it}.$$

The parameter of interest is β_1. This first specification explores the cross-sectional correlation between the sensitivity to sales shocks and vertical separation. It takes into account industry-specific trends but not the fixed propensity of firms to adapt to industry shocks.

[Table 5 here]

Estimates are reported in table 5, columns 1-4. It is reassuring that, in the cross section, firms with a high degree of contemporary vertical separation tend to display higher sensitivity to industry shocks. The estimate is positive and statistically significant at the 1% level, although the effect's order of magnitude is moderate: an increase in vertical separation of one standard deviation (+0.15) is associated with an increase in the sensitivity to sales shocks of 0.02. This is small when compared to the mean sensitivity of firm sales to industry shocks, which is equal to 0.4.[8] However, we do not know how much the *variation* in this sensitivity (which is not directly measurable) is explained by vertical separation. Column 2 in Table 5 presents a version of equation (35) without fixed effects ($\beta_i = 0$) but using separation lagged by one year ($k = 1$). Compared to the specification of

8. This figure is obtained by regressing firm sales on industry sales with firm fixed effects but without controlling for industry dummies × industry sales (non-reported regressions).

column 1, the coefficient has the same order of magnitude and is still statistically significant at 1%. We then progressively lose significance as we use longer lags. As shown in columns 3 and 4, the parameter drops from 0.13 to 0.10 (resp. 0.08) as we use vertical separation propensity lagged by two (resp. three) years.

Columns 5–7 include the controls for fixed effects in the sensitivity to sales shocks, as in equation (37), and lag separation by 1 - 3 years respectively. Going from column 2 to 5 simply controls for fixed effects of elasticity (β_i) while focusing on separation lagged one year. Accounting for fixed effects, the parameter for separation remains significant at 5%; the order of magnitude is similar for both cross-sectional and fixed-effect estimates. This is a clear gain in comparison with the approach taken in Table 2. Our annual sales shock response co-moves with vertical separation even after accounting for firm fixed effects. Unfortunately, we lose magnitude and significance with increasingly longer lags, as shown in columns 6 and 7.

[Table 6 here]

Yet our theory predicts that the effect of vertical separation on volatility should be much stronger with respect to employment. Therefore, in the spirit of Table 3, Table 6 reports at the sensitivity of *employment* to industry sales shocks by presenting estimates of equations (35) and (36) with log(employment), instead of log(sales), as the LHS variable. Results are much stronger statistically and economically significant: an increase in outsourcing of one standard deviation (+0.15) leads to an increase in employment's sensitivity to shocks of 0.03. This should be compared with the average employment to shock sensitivity in the sample, which is 0.20.[9] Thus, employment is more sensitive to sales shocks when the firm outsources a bigger fraction of its production process, compared to its industry average. The effect is strong and statistically significant (always at 1%) whether we control for firm fixed effects (columns 5–7) or not (columns 1–4) and even when we investigate the effect of long-lagged (up to three years) changes in vertical separation. This is evidence in favor of the causal effect of separation on firm volatility: three years after an increase in separation, firm employment volatility is still increasing.

5. Leads for Further Research

Although the argument we develop in this paper is simple, we believe it yields original insights in the interplay between organizational structure of the firm and uncertainty on the product and labor markets. Our main result, established in a theoretical and empirical manner, is that vertical separation increases firm-level volatility of sales and employment both for the contracting firms and their outside contractors. This result is at odds with conventional wisdom, which views separation as a practice of large and dominant companies for transferring risk to small subcontracting firms.

---

9. As with the case of firm sales, this figure is obtained by regressing now firm employment on industry sales with firm fixed effects but without controlling for industry dummies × industry sales (non-reported regressions).

We believe that the mechanism studied in this paper delivers several promising leads for further research. For example, we have not modeled imperfections that may arise from frictions on the labor market. Given such frictions, vertical separation would tend to increase the mismatch of workers to vacancies and thereby increase unemployment. Another vein of research would be to look at the consequences of vertical separation in terms of asset pricing. Since separation enhances firms' adaptation to demand shocks, this should result in a change of the risk profile of contracting firms. Such a channel would bring new insights regarding the links between a firm's organization and its financial returns.

Appendix A: An Alternative Modeling of Gamma

In the main text, when a rigid contract is signed, integration is always preferred to separation. This is not a result of our model; it is a direct implication of our ad hoc assumption that vertical separation induces an extra cost γ. Hence, given this "technological"cost difference, the allocation of ownership is not irrelevant even in the presence of a contract. The role of this appendix is to explain why ownership allocation still matters even when a contract is signed. In the main text, ownership allocation does not really matter because the contract implicitly specifies a heavy punishment for the supplier who does not comply. Hence she always does, and thence receives the fixed payment W, regardless of ownership allocation.

To find a way around this problem, we depart now from the model presented in the main text by making the contract less binding. The additional feature here is that the supplier is *limitedly liable*. Under these circumstances, any deviation from the contract cannot be punished by a utility flow of less than zero. The contract is thus less binding for the supplier and so the allocation of ownership matters once again. The intuition is as follows. A limitedly liable supplier has more bargaining power when she chooses not to comply with the terms of the contract. When she did not own the intermediate input, there was nothing to bargain over and so she honored the contract.

Now, if the supplier does own the intermediate input, then she may threaten to withdraw it from production. The buyer cannot threaten to give her less than zero utility and therefore has incentives to accept renegotiation. Hence, this additional incompleteness adds another moral hazard problem: compliance or not with the terms of the contract. To foster compliance, the buyer must be generous under the rigid contract (provide a high W) and thus leave an agency rent to the owner-supplier. Hence, *separation can be costlier than integration even when a contract is signed*, because the supplier must be allowed an agency rent as inducement for honoring the terms of the contract. As is frequent in contract theory, this agency rent arises because the agent (supplier) cannot be sufficiently punished for lack of effort (here, breaching the contract).

A.1. The Modified Framework

Let us investigate more formally the consequences of this assumption depending on the ownership allocation. Under integration, little is changed. Since the supplier does not own the intermediate input, she receives $\bar{W} - c\bar{q}$ when she complies with the terms of the contract; when she does not, she receives nothing. Given that $\bar{W} - c\bar{q} > 0$ to satisfy the supplier's participation constraint, she always prefers to comply than to breach the contract. As in the main text, the buyer's expected profit is thus given by

$$\Pi^{\text{IN}} = \max_{\bar{q}} \ E\left[\pi(\bar{q};\tilde{T}) - c\bar{q}\right] \tag{A.1}$$

$$= \frac{R}{\sigma - 1} E[T^{\varepsilon}]^{1/\varepsilon}. \tag{A.2}$$

As in the main text, if no contract is signed under vertical integration, then no intermediate input is provided ($q = 0$).

Under vertical separation, the contract $(\bar{q}, W)$ is not always executed by the supplier at period 2. On one hand, the supplier may observe the terms of the contract and she receives a surplus:

$$V^{\text{honor}} = W - c\bar{q}.$$

On the other hand, a supplier who produces a quality different than the contractual one, $q \neq \bar{q}$, faces no punishment fee. Given that she owns the intermediate input, the supplier can threaten the buyer not to deliver it, in which case both parties get zero (no production, and no punishment for the supplier). Hence, after bargaining with the buyer over the surplus, the supplier with intermediate input q receives $\varphi\pi(q;\tilde{T})$. Given that the supplier can choose q after observing $\tilde{T}$, she maximizes $\varphi\pi(q^{\text{OUT}};\tilde{T}) - cq^{\text{OUT}}$, where $\pi(\cdot)$ is given by (3). Hence, a noncomplying supplier gets

$$V^{\text{breach}}(\tilde{T}) = \frac{R}{\sigma - 1}\varphi^{1+\alpha(\sigma-1)}\tilde{T}.$$

As a consequence the supplier honors the contract whenever $V^{\text{honor.}} > V^{\text{breach}}$, which holds if and only if

$$\tilde{T} \leq T^* \equiv \frac{W - c\bar{q}}{(R/(\sigma - 1))\varphi^{1+\alpha(\sigma-1)}}. \tag{A.3}$$

As it turns out, the supplier complies with the terms of the contract as long as demand is below a certain cutoff T^*. When demand exceeds this cutoff, the surplus created by choosing a well-suited q becomes large - above the surplus provided by the contract. Therefore, the cutoff T^* is increasing in $W - c\bar{q}$, the utility left to the supplier when she observes the contract. Hence, a buyer who wants the supplier to follow the contract as often as possible (T^* high) must leave her a larger utility $W - c\bar{q}$ when she does so. This is a standard incentive compatibility constraint: the non-contractible effort is "compliance". Given that compliance cannot be enforced via heavy punishment, an agency rent must be left to the supplier as an inducement to comply. Her participation constraint ($W - c\bar{q} \geq 0$) is no longer binding, but now the incentive compatibility constraint is.

A.2. Buyer's Choice

We now turn to the principal's (buyer's) problem. When $\tilde{T} \leq T^*$, the supplier follows the contract; she delivers $\bar{q}$ and is paid W, and the buyer collects $\pi(\bar{q}; \tilde{T}) - W$. When $\tilde{T} > T^*$, the supplier chooses q^{OUT} and bargains with the buyer over the surplus $\pi(q^{\text{OUT}}; \tilde{T})$. As a result, the buyer obtains $(1 - \varphi)\pi(q^{\text{OUT}}; \tilde{T})$. Hence, the expected surplus of the buyer is given by

$$\Pi^{\mathrm{OUT}}(\bar{q}, W, T^*) = P(\tilde{T} < T^*)\, E\left[\pi(\bar{q}; \tilde{T}) - W \,\middle|\, \tilde{T} < T^*\right]$$
$$+ P(\tilde{T} > T^*)\, E\left[(1-\varphi)\pi(q^{\mathrm{OUT}}; \tilde{T}) \,\middle|\, \tilde{T} \geq T^*\right].$$

The buyer designs a contract $(\bar{q}, W)$ that will maximize this expected surplus. Here T^* is given through the incentive compatibility constraint (A.3) by

$$W - c\bar{q} = T^* \frac{R}{\sigma - 1} \varphi^{1+\alpha(\sigma-1)}. \tag{A.4}$$

Hence, a choice of W corresponds to a choice of T^* for a given $\bar{q}$. We thus rewrite the buyer's problem as the maximization, with respect to $\bar{q}$ and T^*, of

$$\Pi^{\mathrm{OUT}}(\bar{q}, T^*) = P(\tilde{T} < T^*)\, E\left[\pi(\bar{q}; \tilde{T}) - c\bar{q} - T^* \frac{R}{\sigma - 1} \varphi^{1+\alpha(\sigma-1)} \,\middle|\, \tilde{T} < T^*\right]$$
$$+ P(\tilde{T} \geq T^*)\,(1-\varphi)\, E\left[\pi(q^{\mathrm{OUT}}; \tilde{T}) \,\middle|\, \tilde{T} \geq T^*\right].$$

This problem can be simplified by recognizing that

$$q^{\mathrm{OUT}} = \arg\max_{q} \;\; \varphi\pi(q; \tilde{T}) - cq$$

and then taking the maximum of Π^{OUT} with respect to $\bar{q}$. Hence, the buyer's problem is to maximize

$$\begin{aligned} \Pi^{\mathrm{OUT}}(T^*) \;&=\; \frac{R}{\sigma - 1}\Bigg(H(T^*)\, E[\tilde{T}^{\varepsilon} | \tilde{T} < T^*]^{1/\varepsilon} \\ &+ (1 - H(T^*)) \frac{1-\varphi}{\varepsilon} \varphi^{(1-\varepsilon)/\varepsilon}\, E[\tilde{T} | \tilde{T} > T^*] - H(T^*) T^* \varphi^{1/\varepsilon} \Bigg) \end{aligned} \tag{A.5}$$

with respect to T^*, where $H(\cdot)$ is the cumulative distribution function of $\tilde{T}$ on the range $[0, +\infty)$. The first term corresponds to the expected profit when the contract is followed by the supplier. The second term corresponds to the expected profit when the supplier breaches the contract. In this case, expected profits may be larger as a result of better adaptation—expectation of T, not of T^{ε}—but there is a holdup cost because the buyer does not garner all the profits and the supplier does not make enough effort. The last term, $H(T^*)\,T^*\,\varphi^{1/\varepsilon}$, is the interesting one and stems from the additional agency problem that we have introduced here.

The more compliance the buyer wants, the higher is T^* and the more generous the buyer must be under compliance—given that he cannot punish the supplier. This "carrot" becomes more expensive and reduces the buyer's profits as the moral hazard problem (higher values of T tempt the supplier to breach) becomes more prominent. In sum, the optimal T^* depends on the relative strength of three effects: as in the main text, it is costly in terms of adaptation but shelters the buyer from holdup problems; the third effect is a cost in term of incentive compatibility.

A.3. Consequences

It is important to note that, if the optimal T^* were very large, then the profit under vertical separation would be nearly

$$\Pi^{\mathrm{OUT}} \approx \frac{R}{\sigma - 1}\left(E[T^{\varepsilon}]^{1/\varepsilon} - T^*\varphi^{1/\varepsilon}\right),$$

which is much smaller than the profit under vertical integration (see equation (A.2)). *When rigid contracts are optimal, they work better under integration* because integration prevents the supplier from breach the contract. However, if a certain degree of flexibility is required (lower T^*), then the only way to adapt is by accepting a certain degree of opportunism: vertical integration is needed.

Another direct implication of (A.5) is that, *even under vertical separation, some degree of rigidity is always optimal.* Indeed, the first-order derivative of Π^{OUT} with respect to T^* is given by

$$\frac{d\Pi^{\text{OUT}}}{dT^*} = \frac{R}{\sigma - 1}\left(H'(0)\left(1 - \frac{1-\varphi}{\varepsilon}\varphi^{(1-\varepsilon)/\varepsilon} - \varphi^{1/\varepsilon}\right)\right) > 0,$$

which is always positive. On a tiny interval $[0, \eta]$, the flexibility loss is small and is outweighed by the reduction in holdup costs.

All in all, this model replicates most of the features described in the main text. First, the degree of vertical separation can be solved separately from the general equilibrium analysis; the trade-off involves the benefits of adaptation on the one hand and holdup costs on the other. The new feature here is an additional cost of vertical separation: in this regime, limited liability renders rigid contracts costly to enforce. Second, the optimal degree of vertical integration depends on the importance of the holdup problem, the distribution of demand shocks, and the degree σ of competition. Finally, as in the main text, vertical separation amplifies firm-level uncertainty.

Appendix B: Reaction to Sales Shocks with Fixed Effects

Equation (35) assumes that there is no unobserved heterogeneity governing the sensitivity of sales to industry shocks. From this viewpoint, such a specification is less demanding than our first set of regressions described in Section 4.2 (Tables 2 and 3). In addition, given that we plan to use lags of vertical separation at the firm-level, we want this lagged measure of vertical separation to be identified from the time series and not from the cross section. This is why we need to control for unobserved heterogeneity in elasticity to industry shocks.

For expositional purposes, we abstract here from industry trends $\beta_s t$ in equation (36). We write

$$\begin{aligned}\log \mathrm{SALES}_{it} &= \alpha_i + \gamma_t + \beta_i \log \mathrm{SALES}_{st} \\ &\quad + \beta_1 \left(\mathrm{XS\_SEPA}_{i,t-k} \times \log \mathrm{SALES}_{st} \right) + \hat{\varepsilon}_{it},\end{aligned} \tag{B.1}$$

where $\mathrm{XS\_SEPA}_{i,t-k}$ stands for excess separation: the difference between the firm's degree of separation and the average degree of separation of its industry.

First we average equation (B.1) across years at the firm-level (there $\overline{X}_i$ denotes the mean of X_i across years t for a given firm i):

$$\overline{\log \mathrm{SALES}_i} = \alpha_i + \overline{\gamma} + \beta_i \overline{\log \mathrm{SALES}_s} + \beta_1 \overline{\mathrm{XS\_SEPA}_{i,t-k} \times \log \mathrm{SALES}_{st}} + \overline{\varepsilon}_i.$$

Then we deduce an estimate for β_i,

$$\begin{aligned}\beta_i =& \frac{\overline{\log \mathrm{SALES}_i}}{\overline{\log \mathrm{SALES}_s}} - \beta_1 \frac{\overline{\mathrm{XS\_SEPA}_{i,t-k} \ \log \mathrm{SALES}_{st}}}{\overline{\log \mathrm{SALES}_s}} \\ & - \frac{\alpha_i}{\overline{\log \mathrm{SALES}_s}} - \frac{\overline{\gamma}}{\overline{\log \mathrm{SALES}_s}} - \frac{\overline{\varepsilon}_i}{\overline{\log \mathrm{SALES}_s}},\end{aligned}$$

which we plug back into the initial specification (B.1) to derive

$$\begin{aligned}\log \mathrm{SALES}_{it} =& \frac{\overline{\log \mathrm{SALES}_i}}{\overline{\log \mathrm{SALES}_s}} \log \mathrm{SALES}_{st} \\ & + \beta_1 \left(\mathrm{XS\_SEPA}_{i,t-k} - \frac{\overline{\mathrm{XS\_SEPA}_{i,t-k} \ \log \mathrm{SALES}_{st}}}{\overline{\log \mathrm{SALES}_s}} \right) \log \mathrm{SALES}_{st} \\ & + \alpha_i \left(1 - \frac{\log \mathrm{SALES}_{st}}{\overline{\log \mathrm{SALES}_s}} \right) + \gamma_t - \overline{\gamma} \frac{\log \mathrm{SALES}_{st}}{\overline{\log \mathrm{SALES}_s}} + \varepsilon_{it} - \overline{\varepsilon}_i \frac{\log \mathrm{SALES}_{st}}{\overline{\log \mathrm{SALES}_s}}.\end{aligned}$$

To further simplify this expression, we now assume that $\text{XS\_SEPA}_{i,t-k}$ and $\log \text{SALES}_{st}$ are uncorrelated *within* firm i (conditional on i):

$$\text{XS\_SEPA}_{i,t-k} \perp \log \text{SALES}_{st} \;\Big|\; i.$$

This is a reasonable assumption if k is large enough. In any case, if we want to interpret the correlation between lagged separation and sales shocks in a causal way then we must assume that firms do not outsource because they anticipate future industry shocks will be larger.

We also make the approximation that industry sales fluctuate moderately about their means:

$$\frac{\log \text{SALES}_{st}}{\overline{\log \text{SALES}_s}} \approx 1.$$

This is more debatable, but serves to simplify the estimation procedure.

Using these two assumptions, we can rewrite the econometric equation as

$$\begin{aligned}\log \text{SALES}_{it} = & \frac{\overline{\log \text{SALES}_i}}{\overline{\log \text{SALES}_s}} \log \text{SALES}_{st} + \gamma_t - \overline{\gamma} - \overline{\varepsilon}_i \\ & + \gamma_1 \left(\text{XS\_SEPA}_{i,t-k} - \overline{\text{XS\_SEPA}_i}\right) \log \text{SALES}_{st} + \varepsilon_{it}, \qquad \text{(B.2)}\end{aligned}$$

which is easily interpreted. Two things are needed to control for fixed effects in sensitivity. First, we must control for the "average sensitivity to industry shocks" (the first term on the RHS of (B.2)). Second, instead of directly interacting excess separation with industry sales, we should focus on the spread between separation and in its firm-level average. The average error term $\hat{\varepsilon}_{it}$ requires the inclusion of firm fixed effects in the estimation process. Observe the similarities between model (B.2) and a standard within estimator for panel data.

References

Acemoglu Daron, Aghion Philippe, Griffith Rachel and Zilibotti Fabrizio, [2004], "Vertical Integration and Technology: Theory and Evidence", NBER Working Paper #10997.

Battigalli Pierpaolo and Maggi Giovanni, [2002], ”Rigidity, Discretion and the Costs of Writing Contracts,” *American Economic Review*, vol. 92, pp. 798-817

Campbell John, Lettau Martin, Malkiel Burton, Xu Yexiao, [2001], ”Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk”, *Journal of Finance*, vol 56, pp 1-43.

Carlton, Dennis [1979], ”Vertical Integration in Competitive Markets under Uncertainty”, *The Journal of Industrial Economics*, Vol 27, pp 189-208

Caroli, Eve and Van Reenen, John [2001], ”Skill Biased Organizational Change? Evidence from a panel of British and French Establishments”, *Quarterly Journal of Economics*, vol 116, pp. 1449-92

Chandler, Alfred [2004], “Scale and Scope: The Dynamics of Industrial Capitalism”, Belknap Press

Comin Diego and Philippon, Thomas [2005], "The Rise in Firm-Level Volatility: Causes and Consequences", *NBER Macroannuals 2005*

Feenstra Robert, [1998], ”Integration of Trade and Disintegration of Production in the Global Economy”, *Journal of Economic Perspective*, Fall 1998, pp 31-50

Givord, Pauline and Maurin, Eric [2004], ”Changes in Job Stability and Their Causes : An Empirical Analysis for France 1982 - 2002”, *European Economic Review*, vol 48, pp 595-615

Grossman, Sanford, and Hart, Oliver [1986], ”The Costs and Benefits of Ownership : A Theory of Vertical and Lateral Integration”, *Journal of Political Economy*, vol 94, pp 691-714

Hummels, David, Ishii June and Yi Kei-Mu, [2001], "The nature and growth of vertical specialization in world trade", *Journal of International Economics*, vol 54, pp 75-96.

Lacity, Mary,Willcocks, Leslie and Feeny David [1995] ,”IT Outsourcing: Maximize Flexibility and Control ”, *Harvard Business Review*, vol 73, pp 84-93

Macchiavello, Rocco, [2006], "Contractual Imperfections, Credit Markets, and Vertical Integration: Theory and Cross-Country Evidence", mimeo LSE.

Magretta, Joan [1998], ”The Power of Vertical Integration: An Interview with Michael Dell”, *Harvard Business Review*, vol 76, pp 72-85

Mulhainathan, Sendhil, and Scharfstein, David [2001], ”Do Firm Boundaries Matter ?”, *American Economic Association Papers and Proceeding*, Vol 91, pp 195-199

Neumark, David [2000], ”Changes in Job Stability and Job Security: A Collective Effort to Untangle, Reconcile and Interpret the Evidence”, NBER WP No. 7472

Saint Paul, Gilles [1995], *Dual Labor markets: A Macroeconomics Perspective*, MIT Press

Stock James and Watson Mark, [2002], ”Has the Business Cycle Changed and Why?”, NBER Working Paper 9127

Thesmar, David and Thoenig, Mathias [2003], “Why Is A Flexible World More Insecure ? How Outsourcing Amplifies Uncertainty”, CEPR No. 3406

Table 1: Trends in Vertical Separation by Industry

	1978-2000	1978-1990	1990-2000
Agriculture	-3,8	-4,6	0,8
Food processing	9,5	7,3	2,2
Consumer goods	9,7	8,7	1,0
Automobile	14,9	7,3	7,6
Equipment goods	-3,8	-4,2	0,4
Intermediate goods	6,8	9,2	-2,4
Energy	-17,3	-15,6	-1,7
Construction	3,5	1,3	2,2
Trade and repair	0,2	-0,8	1,0
Transportation	3,4	2,4	1,0
Financial services	7,6	7,9	-0,4
Real estate	0,6	-0,2	0,9
Corporate services	-2,6	0,0	-2,6
Houshold services	4,2	3,9	0,4
Education, Health, Social services	0,3	0,0	0,3
Average	2,2	1,5	0,7
Overall economy	2,9	1,3	1,6

Source: INSEE.

Note: Between 1978 and 2000, total purchase of intermediate inputs by firms in the food processing industry has increased by 7.3 points of total production.

Table 2: Variance of Sales Growth and Outsourcing

	Model (1)	(2)	(3)	(4)	(5)
Vertical separation	0.118***	0.138***	0,065	0.094***	-0,115
	[0.031]	[0.027]	[0.180]	[0.026]	[0.113]
log(Sales)	0,003	-0,004	-0,024	-0.008***	-0.046***
	[0.003]	[0.002]	[0.023]	[0.002]	[0.017]
Constant	0,046	0,054	0,386	0.104***	0.693***
	[0.048]	[0.049]	[0.371]	[0.040]	[0.254]
Industry FE	yes	no	no	no	no
Firm FE	no	no	yes	no	yes
Period Effects	no	yes	yes	yes	yes
Period Effects x Industry FE	no	yes	yes	yes	yes
Subperiods	1984-1999	1984-1991 1992-1999	1984-1991 1992-1999	1984-1989 1990-1994 1995-1999	1984-1989 1990-1994 1995-1999
Observations	2315	4630	4630	6944	6944
R^2	0,15	0,16	0,7	0,16	0,56

Notes: The sample is restricted to firms present throughout 1984 - 1999. The dependent variable is the variance of the sales growth rate. This variance is computed on various subperiods. The degree of vertical separation is measured as the share of intermediate inputs in total sales. Industry fixed effects (FE) are at the 2-digit level. All regressions exclude top and bottom 5% values of dependent and independent variables. Heteroskedasticity-robust standard errors are in brackets.

* significant at 10% ** significant at 5% *** significant at 1%

Table 3: Variance of Employment Growth and Outsourcing

	Model				
	(1)	(2)	(3)	(4)	(5)
Vertical separation	0.196***	0.153***	0,059	0.058**	0.090***
	[0.049]	[0.032]	[0.185]	[0.029]	[0.016]
log(employment)	0.011**	0	-0,024	0,012	0.032***
	[0.005]	[0.003]	[0.024]	[0.009]	[0.005]
Constant	-0.126**	-0,026	0,501	-0,031	-0.193***
	[0.058]	[0.054]	[0.437]	[0.060]	[0.039]
Industry FE	yes	no	no	no	no
Firm FE	no	no	yes	yes	yes
Period Effects	no	yes	yes	yes	yes
Period Effects x Industry FE	no	yes	yes	yes	yes
Subperiods	1984-1999	1984-1991 1992-1999	1984-1991 1992-1999	annual data	annual data
Observations	2306	4600	4600	11312	11312
R^2	0,08	0,15	0,68	0,44	0,61

Notes: The sample is restricted to firms present throughout 1984 - 1999. The dependent variable is the firm-level variance of employment growth as computed on various subperiods. In column (3) the dependent variable is the *within year* variance of employment computed on a quarterly basis; in column (4), the dependent variable is the firm-level share of interim workers on total employment. The degree of vertical separation is measured as the share of intermediate inputs in total sales. All regressions exclude top and bottom 5% values of dependent and independent variables. Heteroskedasticity-robust standard errors are in brackets.

* significant at 10% ** significant at 5% *** significant at 1%

Table 4: Serial Correlation of Our Measure of Outsourcing

	Model			
	(1)	(2)	(3)	(4)
(Vertical Separation)$_{t-1}$	0,51	-	-	0,43
	[0,01]			[0,01]
(Vertical Separation)$_{t-2}$	-	0,32	-	0,10
		[0,01]		[0,01]
(Vertical Separation)$_{t-3}$	-	-	0,19	0,02
			[0,01]	[0,01]
Firm Fixed Effects	yes	yes	yes	yes
Year Fixed Effects	yes	yes	yes	yes
Observations	86013	80194	74408	68054
R^2	0,94	0,94	0,94	0,95

Notes: The sample is restricted to firms present throughout 1984 - 1999. The dependent variable is the contemporary degree of vertical separation as measured by the share of intermediate inputs in total sales. Standard errors (in brackets) assume that the various error terms of a given firm are correlated with each other.

Table 5: Past Outsourcing and the Sensitivity of Firm Sales to Industry Sales (IS)

	Model						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Excess Separation x Log Indus. Sales	0.137**	0.134***	0.096**	0,075	0.107**	0.082*	0,06
	[0.057]	[0.049]	[0.048]	[0.046]	[0.051]	[0.049]	[0.048]
Log(IS)	2.693***	2.791***	2.826***	2.565***	2.364***	2.459***	2.244***
	[0.230]	[0.257]	[0.282]	[0.334]	[0.336]	[0.358]	[0.406]
Excess Separation	-1.885*	-2.155**	-1.570*	-1,245	-1.661*	-1,312	-0,967
	[1.057]	[0.911]	[0.884]	[0.861]	[0.946]	[0.910]	[0.895]
Mean Reaction to IS	-	-	-	-	0.558*	0,487	0,418
					[0.288]	[0.298]	[0.303]
Firm FE	yes	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes	yes
Industry FE x Log(IS)	yes	yes	yes	yes	yes	yes	yes
Industry FE x Trend	yes	yes	yes	yes	yes	yes	yes
Industry FE x Trend x Log(IS)	yes	yes	yes	yes	yes	yes	yes
N of lags for Excess Separation	0	1	2	3	1	2	3
Observations	32013	29817	27919	25994	29817	27919	25994
R^2	0,89	0,9	0,9	0,91	0,9	0,9	0,91

Notes: The sample is restricted to firms continuously present in the data during 1984 - 1999. The dependent variable is the firm level log of sales. The main explicative variable is the interaction term between the excess degree of vertical separation and the log of industry sales (IS). All regressions control for industry FE and industry-specific trends in the sensitivity to industrywide shocks. Models (1) and (2) do not include the control for firm-level unobservable heterogeneity in the sensitivity to industry shocks; models (3) - (5) include such a control ("Mean reaction to industry sales"). Model (1) uses a contemporary measure of outsourcing; models (2) and (3) use a measure of outsourcing lagged by one year, model (4) a two year lag and model (5) a three year lag. Heteroskedasticity-robust standard errors are in brackets.

* significant at 10% ** significant at 5% *** significant at 1%

Table 6: Past Outsourcing and the Sensitivity of Firm Employment to Industry Sales (IS)

	Model						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Excess Separation x log(IS)	0.208***	0.160***	0.166***	0.172***	0.155**	0.169***	0.179***
	[0.070]	[0.062]	[0.057]	[0.056]	[0.063]	[0.058]	[0.057]
log(IS)	0.579**	1.108***	1.387***	1.307***	1.048***	1.317***	1.282***
	[0.260]	[0.293]	[0.322]	[0.358]	[0.346]	[0.368]	[0.405]
Excess Separation	-4.635***	-3.598***	-3.542***	-3.505***	-3.504***	-3.597***	-3.638***
	[1.305]	[1.157]	[1.065]	[1.053]	[1.188]	[1.098]	[1.071]
Mean Reaction to IS	-	-	-	-	0,189	0,22	0,074
					[0.522]	[0.541]	[0.548]
Firm FE	yes	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes	yes
Industry FE x Log(IS)	yes	yes	yes	yes	yes	yes	yes
Industry FE x Trend	yes	yes	yes	yes	yes	yes	yes
Industry FE x Trend x Log(IS)	yes	yes	yes	yes	yes	yes	yes
N of lags for Excess Separation	0	1	2	3	1	2	3
Observations	31816	29631	27742	25830	29631	27742	25830
R^2	0,93	0,93	0,94	0,94	0,93	0,94	0,94

Notes: The sample is restricted to firms continuously present in the data during 1984 - 1999. The dependent variable is the firm-level log of employment. The main explicative variable is the interaction term between the excess degree of vertical separation and the log of industry sales (IS). All regressions control for industry FE and industry-specific trends in the sensitivity to industrywide shocks. Models (1) and (2) do not include the control for firm-level unobservable heterogeneity in the sensitivity to industry shocks; models (3) - (5) include such a control ("Mean reaction to industry sales"). Model (1) uses a contemporary measure of outsourcing; models (2) and (3) use a measure of outsourcing lagged by one year, model (4) a two-year lag and model (5) a three-year lag. Heteroskedasticity-robust standard errors are in brackets.

* significant at 10% ** significant at 5% *** significant at 1%

Graphics

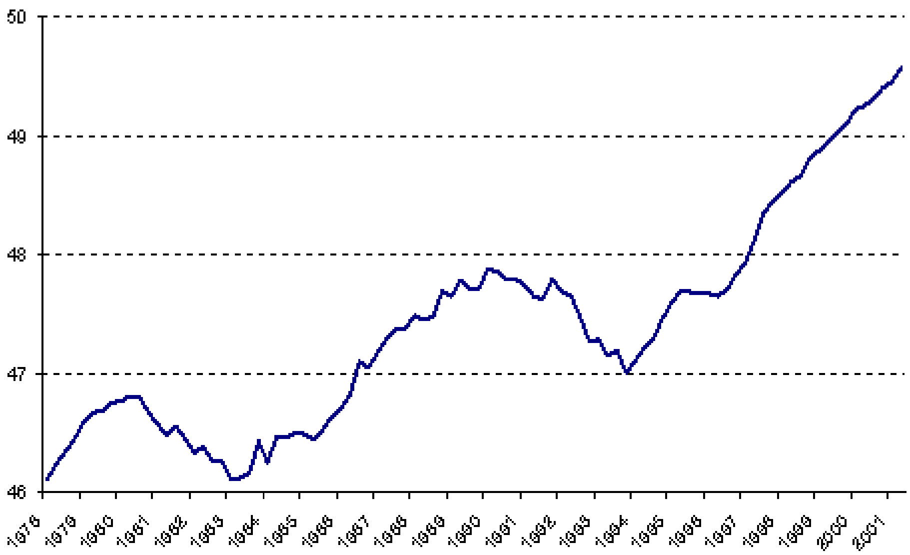

Figure 1: Consumption of Intermediate Inputs as a Fraction of Total GDP (Source: French National Accounts)

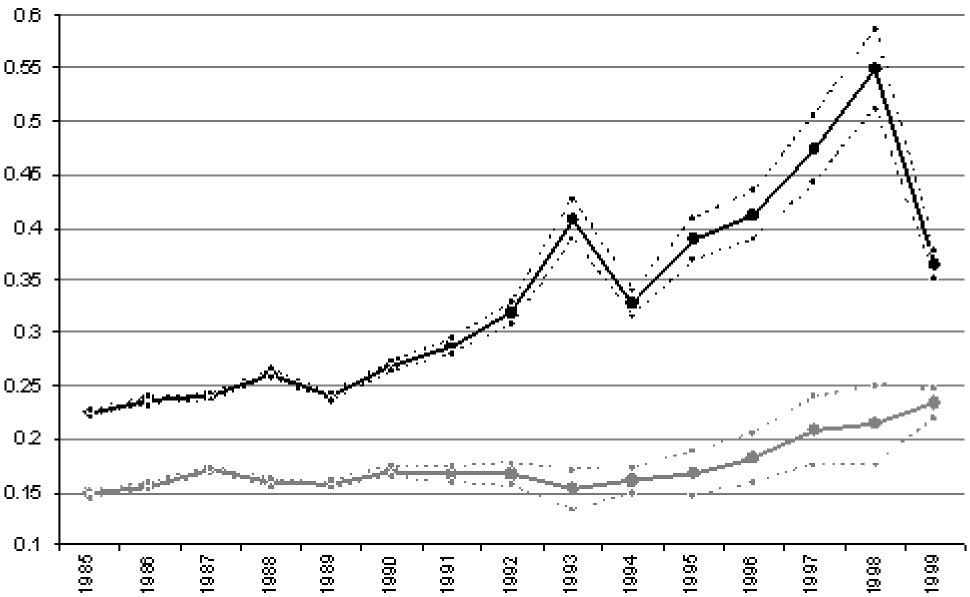

Figure 2: Cross-sectional Variance of Sales Growth (black) and Variance of Employment Growth (Grey)