Helsinki GSE

HELSINKI GSE DISCUSSION PAPERS 36 · 2025

# Markets, Contracts, and Hierarchies: How Bargaining Frictions Affect Governance

Tore Ellingsen
Topi Miettinen

# Helsinki GSE Discussion Papers

Helsinki GSE Discussion Papers 36 · 2025

Tore Ellingsen and Topi Miettinen:
Markets, Contracts, and Hierarchies: How Bargaining Frictions Affect Governance

ISBN 978-952-7543-35-1 (PDF)
ISSN 2954-1492

Helsinki GSE Discussion Papers:
https://www.helsinkigse.fi/discussion-papers

Helsinki Graduate School of Economics
PO BOX 21210
FI-00076 AALTO
FINLAND

Helsinki, Revised January 2026

# Markets, Contracts, and Hierarchies: How Bargaining Frictions Affect Governance*

Tore Ellingsen†          Topi Miettinen‡

January 9, 2026

**Abstract**

We develop an organizational governance model with a single buyer and endogenous upstream entry. Investments and control rights over assets and actions are immediately contractable; production is contractable after uncertainty resolves. We show the following: Supplier competition eliminates pre-entry bargaining frictions. To minimize post-entry bargaining frictions, control rights over assets and actions are always bundled. If entry is sufficiently cheap, there is frictionless post-entry competition, sometimes due to buyer sponsorship. Otherwise, only one supplier enters. There is vertical integration if the asset's expected profitability is highest in the buyer's favorite use; if not, the buyer contracts with an autonomous supplier.

**JEL Codes:** D23, D86, L22

**Keywords:** transaction cost, property rights, relationship specificity, organization, bargaining, vertical integration, governance, inefficiency, master supply agreement, managed market

*Previous title: "Markets, Contracts, or Firms? A Simple Model of Governance." Thanks to Ricardo Alonso, Bob Gibbons, and Birger Wernerfelt for detailed discussions, to Constantin Blome, Patrick Bolton and Oliver Hart for helpful comments, and to Maija Halonen for inviting us to the 2024 Theory and Applications of Contracts conference. The criticisms and suggestions of several anonymous referees and of the Editor, Alexander Wolitzky, have greatly improved the paper. We are responsible for remaining shortcomings. Financial support of the Fulbright Finland Foundation, the Jan Wallander and Tom Hedelius Foundation, the Norwegian Research Council (250506), and the Yrjö Jahnsson Foundation is gratefully acknowledged.

†Department of Economics, Stockholm School of Economics, Box 6501, S–11383 Stockholm, Sweden. E-mail: gte@hhs.se

‡Department of Finance & Economics, Hanken School of Economics, Helsinki GSE, P.O. Box 479, Fi-00101 Helsinki, Finland. E-mail: topi.miettinen@hanken.fi

# 1 Introduction

The study of vertical integration comprises two largely separate literature traditions. A literature on *organizational economics* emphasizes how contracting problems shape firms' internal governance mechanisms, whereas a literature on *industrial organization* emphasizes how technology and tastes shape market structure and firms' external conduct, leaving aside internal governance issues. However, as argued by Bresnahan and Levin (2013), the organization of an entire industry can only be properly understood if we bridge the two traditions:

> Empirical research [in the organizational economics tradition] has faced a tension between testing specific predictions derived from the theory and trying to explain broad industry patterns of integration. Most contractual theories of integration aim to isolate particular mechanisms. The models tend to focus on simple environments, for example two parties structuring a supply relationship with the surrounding market held fixed. This means that the clearest empirical analogues are marginal decisions at the firm boundary, such as whether a firm should make or buy a particular input. Moving from this level of analysis to explain the organization of an entire industry requires additional steps because *the availability of suppliers or the existence of functioning markets or contracts needs to be treated as an outcome of the analysis rather than a fixed parameter.* (Our emphasis.)

Here, we build one such theoretical bridge in the context of the make-or-buy decision, a paradigmatic problem in the organizational economics literature (Tadelis and Williamson, 2013). Concretely, we allow the entry of suppliers and the governance mechanism to be jointly determined. Spot market trade, long-term contracts, and vertical integration can all be rationalized and linked to the primitive parameters of the economic environment.

The key to obtaining a unified treatment of markets, contracts, and hierarchies is our hypothesis that both the costs of separation and the costs of integration consist primarily of *bargaining frictions*. Under vertical separation, parties sometimes fail to agree on efficient outcomes that favor the buyer; under upstream integration the buyer can always implement these outcomes, but the parties sometimes fail to agree internally on efficient outcomes that favor the supplier.[1] By contrast, the conventional assumption of Williamson (1985) and Masten (1984, 1986) is that vertical separation entails bargaining frictions whereas vertical integration instead entails costs of bureaucracy.[2] Thus, we comply with the advice of Gibbons (2005), who writes (page 222):

[1]In terms of formally modeling the trade offs between markets, contracts, and hierarchies, and giving a prominent role to bargaining costs, Wernerfelt (2015) is a closely related prior contribution; see also Wernerfelt (2016). However, Wernerfelt focuses on how integration saves bargaining costs by bundling many transactions, whereas we consider how vertical integration saves bargaining costs in some states of the world but incurs extra bargaining costs in other states. Thus, the substantive insights are quite different.

[2]For a catalog of the costs of bureaucracy that Williamson thought could be associated with vertical integration, see Chapter 6 of Williamson (1985).

It seems preferable to develop an elemental rent-seeking theory of the firm where rent-seeking provides a unified account of both the costs and the benefits of integration.

The analytical challenge is to devise a useful model of bargaining frictions. To do so, we build on Ellingsen and Miettinen (2008, 2014), a framework that rationalizes risks of disagreement in bilateral bargaining environments with complete information.[3] Miettinen and Vanberg (2025) apply the framework to multilateral bargaining with more than two irreplaceable negotiators. Here, we consider instead the case in which some negotiators are substitutes; there can be more than one supplier, yet only one supplier needs to trade with the buyer in order for the surplus to be maximal.

In this bargaining model, competition is shown to reduce the risk of disagreement. With multiple suppliers, any attempt by one supplier to improve the terms of trade risks losing the trade to the competitor. Therefore, competing suppliers find it less attractive to engage in potentially wasteful rent seeking than a monopoly supplier does.[4]

This formal analysis complements the analysis of Williamson (1971, 1975, 1985), who often emphasizes the connection between haggling costs and lack of competition. He argues that there tends to be more intense competition *before* specialized investments are made than afterwards, calling this the *fundamental transformation*. Yet, Williamson rarely discusses how competition might reduce bargaining costs. Likewise, Wernerfelt (2015), who studies the tradeoff between bargaining costs associated with a sole supplier and the costs of accessing a larger market, simplifies by exogenously setting the bargaining costs in markets to zero.[5]

Our analysis rests on three central assumptions. Two of these assumptions are shared with Masten (1986), an early formalization of Williamson's transaction-cost theory. The first of Masten's assumptions is that it is more difficult to contract on trades that take place further into the future.[6] While lack of detailed long-term contracts is not necessarily a problem, we

[3] That analysis in turn builds on Schelling (1956) and, especially, Crawford (1982). For recent contributions to this literature on conflicting commitments in bargaining, see Basak and Deb (2020), Dutta (2024), and Pei (2025).

[4] Other bargaining models might yield similar predictions. A related approach to understanding contracting failure is that parties may afraid of being taken advantage of because they might not fully grasp the offer that the other side is making. For early work along these lines, see Katz (1990). Recent contributions, building on the rational inattention model, are due to Ravid (2020), Wolitzky (2023), and Cusumano, Fabbri, and Pieroth (2024)—the latter showing that inefficiency is smaller when there is competition on the side that makes offers.

[5] A prior formalization of the fundamental transformation is Hart and Moore (2008) and Fehr, Hart, and Zehnder (2011). They propose that negotiators are more prone to accept unequal outcomes when they are arrived at through a competitive process. Our explanation does not require such a direct link from the economic environment to the negotiators' preferences. In this respect, a more closely related contribution is Milgrom and Roberts (1990), especially pages 72-75. They observe that several models of bilateral bargaining, including the Nash demand game, have multiple equilibria, some of which are inefficient, and in the context of the demand game they demonstrate that competition solves this coordination problem. Other formal models of incomplete contracts and ex post inefficiencies have highlighted related frictions, such as asymmetric information (Bajari and Tadelis, 2001, Tadelis, 2002, Loertscher and Marx, 2022, Raith, 2023) and negative reciprocity (Hart and Moore, 2008, Hart and Holmstrom, 2010), but not studied how these inefficiencies depend on supplier competition and may thus affect market structure.

[6] Similar assumptions are also maintained by, e.g., Bajari and Tadelis (2001), Tadelis (2002), Hart and Moore (2008), Hart and Holmstrom (2010) and Raith (2023). By contrast, many classical incomplete contracting models

demonstrate that it can become a problem if future negotiations involve less competition than current negotiations.[7]

Masten's second main assumption, which we also adopt, is that the buyer's profitability is uncertain ex ante.

However, our third main assumption differs sharply from Masten's. We assume that the right to decide over assets and the right to assign supplier effort are both immediately tradable, but that the future benefits associated with implementing the projects are not immediately tradable. By contrast, Masten assumes that vertical integration grants the buyer not only the decision rights, but also the rights to the supplier's future benefits—even if it turns out that the implemented project is not the project that the buyer wanted access to the asset for. Under our assumption that these benefits cannot be traded in advance (or more generally that not all of them are pledgeable), vertically integrated parties need to bargain if it turns out that the supplier's preferred use of the asset is efficient.[8]

To focus on the implications of the three central assumptions, our framework is otherwise as simple as possible. Several potential suppliers face a fixed entry cost. The entry cost is associated with the purchase of an asset whose value outside the two alternative projects is zero. The right to control the asset and the right to choose projects are both contractable when suppliers make their entry decisions, but the project choice itself is not. Only after entry will all parties learn whether the buyer's preferred project is better or worse than the suppliers' preferred project.

As in Klein, Crawford, and Alchian (1978) and Grossman and Hart (1986), the terms of the ex ante contract affect the outside options of the different parties when they negotiate subsequent production decisions and transfers. For example, a supplier without any control rights has no credible threat, and thus cannot claim larger transfers than those specified in the ex ante contract. By contrast, a supplier with both asset ownership and authority over own effort—what we call an *autonomous* supplier—can credibly threaten to pursue the

assume that certain near-term decisions are never contractable. For example, in Grossman and Hart (1986), Rajan and Zingales (2002), Hart and Moore (1990), Bolton and Whinston (1993), Loertscher and Riordan (2019), ex ante investments are not contractable at all, while future trades will eventually become contractable.

[7] Williamson (1985) (p.62) links the fundamental transformation to the fact that—even when the investment is primarily in physical capital—it takes time to build the human capital that is required to make use of the investment, and this reduces substitutability between people:

> [...] there is more to idiosyncratic exchange than specialized physical capital. Human capital investments that are transaction-specific also occur. These evolve during contract execution. Specialized training and learning-by-doing economics in production operations are illustrations. Except when such investments are transferable to alternative suppliers at low cost, which is rare, the benefits can be realized only as long as the relationship between the buyer and seller is maintained.

Accordingly, we assume that there is a strict complementarity between the asset and the supplier that acquires it.

[8] Bargaining frictions inside firms have attracted little attention in economics (if we except the related work on influence costs, e.g., Milgrom, 1988 and Meyer, Milgrom, and Roberts, 1992). By contrast, the academic management literature has paid more attention to this issue, see for example Andersson, Forsgren, and Holm (2007). The perspective that subsidiaries sometimes need to bargain with headquarters in order to implement desirable projects is especially pronounced in Lunnan, Tomassen, Andersson, and Benito (2019).

own favorite project, and thus bargain for a share of the additional surplus that would be generated by the buyer's favorite project.

The main findings are simple and, we think, intuitive. Aside from *spot trading*, only three governance modes ever occur in equilibrium. (i) *Integration:* A single asset is acquired, the buyer owns the asset and has the right to decide on projects—one of the suppliers is thus employed (if the supplier is a firm, that firm becomes a division or a subsidiary). (ii) *Single sourcing.* A single asset is acquired by an autonomous supplier. This relationship is sometimes (but not always) governed by an exclusive contract, through which the buyer extracts rents. (iii) *Parallel sourcing:* At least two assets are acquired, each asset by an autonomous supplier.

In summary, suppliers are either fully autonomous or fully subordinate, but the buyer does not integrate with multiple suppliers. That is, *control rights over assets and people are always bundled, while competition only occurs between firms and not within them*.[9]

The three governance outcomes are linked to parameters as follows. When entry costs are below a threshold, all the suppliers enter, there is spot trade, and outcomes are first-best. With entry costs slightly above the threshold, the buyer subsidizes the entry of two autonomous suppliers, even if first-best would require that exactly one supplier enters. The reason is that competition between the suppliers ensures ex post efficient trade, and trading efficiency is not attainable with only one supplier—regardless of the allocation of control rights—because bilateral bargaining is never entirely efficient. Finally, if entry costs are even higher, only one asset is acquired. The governance mode is then determined by whose use of the asset is likely to generate the largest profit and which is therefore more important to protect from bargaining friction.[10]

As a broad generalization, our model supports the idea that equilibrium governance arrangements tend to be (second-best) welfare-maximizing. That is, they maximize expected benefits minus the expected sum of production costs and transaction costs. However, there is an exception. If the buyer proposes contracts, and the welfare-maximizing outcome entails a single autonomous supplier, the buyer might nevertheless prefer to encourage the entry of a second supplier—using parallel sourcing to extract supplier rents.

The model is also related to the literature that compares different sourcing strategies, typically taking vertical separation for granted but considering competition as a choice variable.[11] Recently, the issue of bargaining frictions has come to play a role in this literature too,

[9]Observe that our argument is not related to asset value depreciation, as in Holmstrom and Milgrom (1991) and Wernerfelt (2002). For more closely related argument for bundling of control rights, see Van den Steen (2010), which we briefly describe in Footnote 22.

[10]This last finding has some resemblance to the idea that vertical integration is useful to guarantee assurance of supply when there is a risk that prices are suboptimally set; see especially Carlton (1979). Under vertical separation, the risk of facing an excessive price is similar to the risk of haggling with a resistant supplier. However, unlike us, Carlton does not consider that the drawback to vertical integration might be symmetric; under integration the buyer could insist on an excessive price in return for allowing trades benefiting the supplier. Instead, Carlton considers the drawback that the buyer might have to take delivery of goods that are not wanted.

[11]This literature sometimes distinguishes parallel sourcing (which is the organized form of competition that we

e.g., in Leider and Lovejoy (2016), but we are not aware of previous formal models of single versus parallel sourcing that study endogenous entry.

Before we present the details of the model, let us fix ideas by relating the model's assumptions and implications to the famous case of General Motors (GM) and its sourcing of automobile bodies.[12] Before 1919, GM bought its bodies from several suppliers. Our explanation is that the supply of bodies was initially quite large relative to GM's needs. There was already a big industry producing horse-drawn carriage bodies—for which demand was now decreasing—and the technological step to producing bodies for automobiles was small. In other words, the early period corresponds to our low entry-cost scenario, where the buyer optimally sources from several autonomous suppliers.

Soon, the demand for GM's closed-body automobiles increased massively. To meet the new demand, it became necessary to invest in large new production plants. At this stage, in 1919, GM chose to write a ten-year contract with one of the suppliers, Fisher Body (Fisher). Fisher remained a separate firm, governed by the Fisher brothers, that also supplied bodies to other automobile producers. However, the individual Fisher brothers also signed 5-year contracts as GM employees. In our model, the need for plant expansions corresponds to an increase in entry cost, while the relative growth of GM's demand corresponds to an increase in relationship-specificity. The observed outcome is consistent with the parameter region above the entry-cost threshold and below the specificity threshold.[13]

After a few years of rapidly growing demand, there were increasing tensions between GM and Fisher over the location of new body plants (Klein, 2007). GM wanted all new body plants to be located adjacent to their assembly plants, saving time and transportation costs. Fisher sometimes preferred different locations to serve its other customers and possibly to take advantage of local cross-plant externalities. Under the long-term contract, Fisher had complete autonomy in coordinating its production, including plant location. Therefore, each new plant required a negotiation between the two firms. These negotiations typically ended with the adjacent location that GM was seeking (see Klein, 2007, Table 1). Ultimately, in 1926, after a period of tense negotiations over the location of a new mega-plant—Fisher wanted it to be located in Detroit and GM wanted it in Flint—GM took over Fisher, which became a division of GM. Our model straightforwardly explains this outcome: the demand for GM's cars became so large and so profitable that it was highly likely to be efficient to locate all new Fisher plants next to GM's assembly plants. At this point, it is better that GM has all

consider) from dual sourcing, multiple sourcing, and second-sourcing. For example, some of the dual sourcing literature studies the splitting of production of a homogeneous good, the second-sourcing literature has a hierarchy of suppliers, and multiple sourcing often refers to different suppliers selling differentiated goods rather than competing to supply the same goods. But all versions share the view that multiple suppliers engender more competition than single sourcing.

[12]The case was brought to economists' attention by Klein, Crawford, and Alchian (1978). The empirical picture was gradually enriched, with particularly important additions by Coase (2000), Freeland (2000) and Klein (2007).

[13]The model fails to rationalize that the Fisher brothers were employed by GM while separately keeping control of their Fisher assets. On the other hand, as noted by Coase (2000), the employment contracts were not considered a permanent solution, but more of a "trial marriage."

the control rights. Letting Fisher retain control rights merely increases the risk of costly bargaining frictions.

Section 2 presents the model, except for the details of the bargaining process, which are presented and analyzed in Section 3. Section 4 solves the model. Section 5 offers some final remarks.

## 2 The Model

There are four players. One player is a buyer, $B$. The other three players are ex-ante identical suppliers $S_1, S_2$, and $S_3$.[14] All players are selfish and risk neutral and each player $i \in \{B, 1, 2, 3\}$ has an outside option (non-production) that we normalize to 0. As our recurrent example, we think of the buyer as GM (through its President) and one of the suppliers as Fisher Body (through the Fisher brothers, who were jointly in charge of Fisher Body).

Production of output requires a single asset. Suppliers purchase assets in a competitive market at cost $f > 0$. The buyer can subsidize asset purchase but is unable to purchase or operate the asset. Operation requires the (full) effort of the supplier that purchased it.[15] For example, GM could fund Fisher's plant investments, but lacked the expertise (tacit knowledge and patents) of the Fisher brothers to operate body plants efficiently. Symmetrically, we assume that suppliers are unable to run the buyer's business.

There are two types of output (projects)—a regular output and a tailored output. For example, the tailored output is a special closed body design that suits one of GM's new automobile models, the regular output is a more standard design that can be sold to several other car makers.

The regular project generates a quantity of regular output $q_R$, which (for simplicity) has a known selling price $p_R$. Thus, it produces gross profits $v_R = p_R q_R$. The tailored project generates a quantity of tailored output $q_T$.

Initially, the revenues associated with the tailored project are uncertain. Eventually, the price is either $p_T^H$ or $p_T^L$. More precisely, the tailored project produces revenues (gross profits) $v_T^H = p_T^H q_T$ with probability $h$ and $v_T^L = p_T^L q_T$ with probability $(1-h)$. The realized net profits of project $j$ is denoted

$$\pi_j = v_j - f. \tag{1}$$

The expected net profits of the regular project is also $\pi_R = v_R - f$, while the expected net profits of the tailored project is

$$\bar{\pi}_T = h v_t^H + (1-h) v_T^L - f. \tag{2}$$

[14]It is straightforward to extend the model to have more than three suppliers, but it adds more complications than insights.

[15]We could think of the asset as a bundle of physical and human capital required to produce car bodies, for example, and only the purchaser obtains the relevant training required to use the asset productively; see Footnote 7.

We assume that $v_T^H > v_R > v_T^L$. In other words, the tailored project will sometimes be ex post superior to the regular project, but not always. The buyer can undertake at most one (tailored) project.

Like Masten (1986), we assume that the revenues from the tailored project initially accrue to the buyer, $B$. Similarly the revenues from the regular project initially accrue to the supplier that operates the asset. (GM obtained the revenue from selling GM cars, while Fisher Body obtained the revenue from selling standard bodies to other car manufacturers.)

The buyer can make transfers to the suppliers or receive transfers from them. Let $t_i$ denote the transfer between the buyer and supplier $i$, with $t_i > (<) 0$ indicating that the transfer runs to (from) the supplier. Thus, the buyer's net profit is

$$\pi_B = \mathbb{1}_T v_T - \sum_{i=1}^{3} t_i, \tag{3}$$

and supplier $i$'s net profit is

$$\pi_i = t_i - \mathbb{1}_I f + \mathbb{1}_I \mathbb{1}_R v_R, \tag{4}$$

where $\mathbb{1}_I$ is an indicator function that take the value 1 if the supplier invests in an asset, $\mathbb{1}_R$ takes the value 1 if this asset is used for the regular project, and $\mathbb{1}_T$ takes the value 1 if the asset is used for the tailored project.

## 2.1 Asset specificity versus profitability

In our view, the natural definition of *asset specificity* is $q_T/q_R$. As the asset is more and more suitable for the tailored project or less and less suitable for the regular project, asset specificity tends toward infinity.[16] An alternative definition is $v_T/v_R$. However, in our setting, that ratio is stochastic, since it depends on $p_T$, and hence on market conditions that are orthogonal to the technology parameters.

The definition of specificity is not important as such, but one of our findings is that the tailored product's *demand favorability*, $E[p_T]/p_R$ will be just as decisive for governance outcomes as the asset specificity $q_T/q_R$.

## 2.2 Preview of the contracting problem

What makes the governance problem challenging are the limitations to contracting. At the time the asset is acquired, it is impossible to specify what precisely is to be meant by a tailored project. Thus, assignments of assets and effort to projects are not initially contractable but will become contractable later. Also, it is impossible to contract on revenue-sharing; no future revenues are pledgeable at Date 0.[17]

[16] Since the asset costs $f$, the unit cost associated with pursuing project $j$ is $C_j = f/q_j$. Thus, we might equivalently say that asset specificity is the inverse of the unit cost ratio $C_T/C_R$.

[17] For an early discussion of why there are limits to the pledging of future returns, see, e.g., Williamson (1975, Chapter 5, Section 1.5.3).

The detailed process of future interaction is also not fully contractable initially; parties cannot design state-contingent mechanisms that they will be playing later. Only unconditional rights to assign assets to projects (allocation of property rights) and supplier effort (employment or autonomy) can be regulated in an initial contract. For example, it is not possible to write a contract that allocates control through option contracts, as envisaged in the context of trading goods and services by Nöldeke and Schmidt (1995). The idea is that changes in control rights are large and time-consuming operations that cannot be implemented ahead of each individual transaction. For example, GM took over the control of Fisher Body once, and it remained an integrated supplier forever after.[18]

These control rights in turn affect incentives at later stages.

### 2.3 First-best benchmark and brief guide to results

If $f < v_R$, the first-best outcome is that all suppliers invest. In state $H$ one of them should undertake the tailored project and the others should undertake the regular project. In state $L$ all suppliers should undertake the regular project. Perhaps unsurprisingly, we can show that this first-best outcome is always attained in equilibrium, which entails the entry of all three suppliers, who each retain the rights to their effort as well as their assets. No long-term contract is required; spot market interaction suffices.

The interaction between the four players becomes more intricate when the investment cost $f$ is above $v_R$ but below $hv_T^H + (1-h)v_R$. In this case, the first-best outcome is that only one of the suppliers invests and that this supplier pursues the tailored project in state $H$ and the regular project in state $L$.

Whenever entry is desirable, the maximally attainable total welfare can be expressed generally as

$$W^{FB} = \pi_B + \pi_1 + \pi_2 + \pi_3 = hv_T^H + (1-h)v_R - f + \max\{2(v_R - f), 0\}. \tag{5}$$

As we shall see, none of the available governance modes are guaranteed to implement the first-best outcome when $f > v_R$. Depending on $f$ and $h$, one of three different governance modes can be second-best, including one that involves sponsored entry by two autonomous suppliers, a form of pre-arranged spot market interaction.

Intuitively, the advantage from having two suppliers entering is that competition eliminates the temptation to engage in costly rent-seeking at the expense of the trading partner, whereas the disadvantage is the duplication of capacity. Thus, the buyer sometimes strategically prefers parallel sourcing from two autonomous suppliers even when such governance does not maximize overall (second-best) welfare, because the buyer may not be able to extract all the surplus from a single supplier.

[18]This line of justification goes back at least to Simon (1951), see also Wernerfelt (2015), but we make no attempt to formalize it further here. For other (more fully formalized) reasons why contracts are incompletely responsive to the state of the world, see Tirole (2009) and Aghion et al. (2012).

When duplication is too costly, only one supplier enters, and rent-seeking at Date 3 is unavoidable. In this case, it turns out that the best form of governance is vertical integration if $E[\pi_T]$ exceeds $\pi_R$ and vertical separation otherwise, as we show in Proposition 7. This condition can equivalently be written $E[v_T] > v_R$ or $E[p_T]q_T > p_R q_R$. Hence, integration is preferable to separation when

$$\frac{E[p_T]}{p_R}\frac{q_T}{q_R} > 1. \tag{6}$$

That is, integration is preferable to separation when the product of demand favorability and asset specificity exceeds one.

Inasmuch as we compare a single external supplier to vertical integration, it turns out that this condition is also the relevant equilibrium prediction. However, since some of the parameters that affect the inequality (6) also affect equilibrium entry of multiple suppliers, the inequality does not cleanly predict the choice between vertical integration and all forms of vertical separation.

The role of relative profitability over and above asset specificity may encourage new agendas for empirical research. Extensive evidence has already demonstrated that there is more integration when the asset is more specific; see, e.g., Masten (2002) and Lafontaine and Slade (2007) for surveys. The empirical measures of asset specificity used in these studies typically reflect technology rather than demand. To the best of our knowledge, there is no corresponding body of evidence examining whether increases in downstream demand cause upstream vertical integration.

Incidentally, the hypothesis is not entirely new. Masten (1984) writes (page 405): "In general, the organizational costs associated with market exchange increase, the more specialized, profitable, and durable are the investments associated with a given transaction." However, we are not aware that it has been carefully developed or tested.[19]

### 2.4 Timing and contracting

The timing is depicted in Figure 1.

We adopt the assumption that all contracts and contract negotiations are bilateral. For example, the buyer cannot offer a grand contract whose clauses depend on the signing decisions of all the (potential) suppliers.[20]

Thus, at Date 0, $B$ bargains bilaterally with each of the three suppliers, and when a contract is accepted, the clauses of that contract enter into force irrespective of which other contracts are signed.

[19]When Masten (1986) devises a formal model, it has the feature that higher expected profitability of the buyer may not be associated with integration, because in that model the highest profits are associated with states in which the relationship-specific asset is not used for the tailored project, as the buyer has an even better project to attend to.

[20]For discussions of how to model non-cooperative contracting games, see Ellingsen and Paltseva (2016) and Watson (2024).

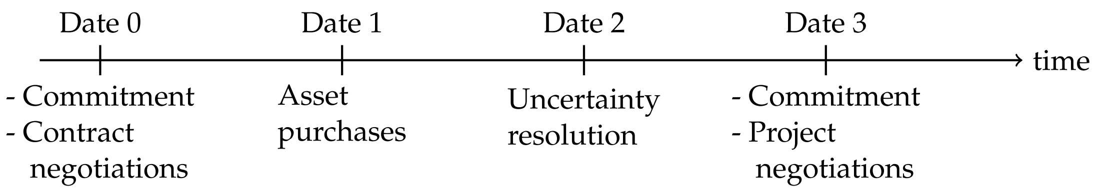


Figure 1: Timing

A crucial feature of the analysis is that players have the opportunity to commit themselves to strong bargaining positions before they meet, with the purpose of increasing their share of the gains from trade. Since the bargaining problem at Date 3 is closely similar, we postpone the description of these bargaining tactics to the next section.

Let us instead begin by describing the contents of the contracts. A Date 0 contract between $B$ and $S_i$ comprises a transfer $t_i$ that is contingent on $S_i$ purchasing an asset and that may also depend on whether, at Date 0, $S_i$ agrees to give $B$ ownership of the asset and/or authority over $S_i$. Asset ownership and authority are defined as follows.

*Ownership of an asset* entails the right to permit of forbid the assignment of the asset to any project.

*Authority over a supplier* entails the right to forbid or permit the assignment of that supplier's effort to any project.

If $B$ has *authority* over $S_i$, $B$ can therefore choose what project $S_i$ spends effort on. Note that authority over $S_i$'s effort does not extend to the use of $S_i$'s asset.[21]

Formally, the players negotiate about bilateral contracts between $B$ and each $S_i$ of the form $(t_i, o, a, e) \in \{\mathbb{R} \times \{b, s\}^2 \times \{0, 1\}\}$, where $o = b$ indicates that $B$ owns the asset, $o = s$ indicates that $S_i$ owns the asset, $a = b$ indicates that $B$ has authority over the supplier's effort (there is an employment relationship), $a = s$ indicates that $S_i$ has authority over the own effort (there is no employment relationship), $e = 1$ indicates that $B$ contracts *exclusively* with $S_i$, and $e = 0$ indicates that the contract is nonexclusive. If no contract is agreed, any trade at Date 3 will be a *spot trade*. The main role of exclusivity clauses, $e = 1$, will be to prevent the buyer from accepting multiple supplier offers—and thereby making entry unprofitable for all of them—when entry of a single supplier is optimal.

We assume that agreed contracts are binding, which is shorthand for saying that contract violations will be sufficiently harshly punished to deter the violation. This assumption is perhaps most questionable when it comes to authority over employees, who will often have a legal right to leave. While we do not model such ex post participation constraints explicitly, it is worth noting that an employed supplier cannot undertake the regular project without

[21] As Wernerfelt (2002) notes, there are many practical examples in which asset ownership is separated from authority over people's efforts. However, authority over a person sometimes implicitly carry over to the asset. For example, employed building workers sometimes own their tools, and when they are assigned to a project it is implied that they are responsible for bringing tools as well.

an asset. Thus, the outside option is indeed unattractive.[22]

As already noted, it is impossible at Date 0 to contract directly on the Date 3 project choice. Indeed, whenever Date 0 contracting gives rise to a non-exclusive contract with $S$ ownership and authority, the contract is akin to the highly incomplete Master Supply Agreements (MSAs) described by Bernstein (2015). For the most part, MSAs are not intended to be legally binding documents that tie the parties to specific terms of trade. Rather, in the words of one of Bernstein's interviewees (p.566):

> the contract is just a formalized handshake that says that your intention is to put business in here

The only legal obligation of the buyer imposed by a typical MSA is to reimburse the supplier for "reliance expenses" (p.567), i.e., for costs that the supplier would not have taken except because of the MSA. In our setting, reimbursement for reliance expenses are payments by $B$ to cover a part of the entry cost $f$ in the case that entry is otherwise unprofitable for the supplier(s).

At Date 1, any supplier who has agreed to acquire an asset does so and makes (or receives) the associated transfers. A supplier without a contract decides freely whether to acquire an asset or not. If $S_i$ has accepted a contract that gives $B$ both asset ownership and authority, we call this mode of governance *full integration*; if $B$ has only one of the two control rights, we call it *partial integration,* and if $B$ has neither, we call it *market*. Markets can involve either a single autonomous supplier, i.e., bilateral monopoly, or two or three autonomous suppliers—competition.

At Date 2, $v_T$ is revealed to be either $v_T^H$ or $v_T^L$.

At Date 3, project choices are made. Agents holding control over both assets and efforts now either implement their inherently preferred project ($R$ for suppliers and $T$ for $B$) or—if this project is not the most valuable overall—seek negotiations with the purpose of obtaining compensation for implementing the other project instead. Since they are not allowed to implement projects unilaterally, agents holding partial control always seek negotiations. At this stage, any contract can be written and enforced, regardless of governance mode.[23]

Exactly as at Date 0, agents can attempt to make commitments before they conduct the negotiations, and the sequence of moves and the information structure is the same, except

[22] Van den Steen (2010) argues that authority is bundled with asset ownership precisely to satisfy such ex post participation constraints.

[23] Note that the latter assumption implies that integration is compatible with contracting; concretely, the buyer can "sell" to the supplier the right to undertake the regular project. As a referee has noted, citing Williamson (1991), such a within-organization contract might not be enforceable in court, so this is one place where we clearly take a short-cut. One justification might be that an internal supplier could always find ways to credibly share a suitable fraction of the project benefits with the buyer, an issue that we briefly return to in the final section. For example, after becoming a division within GM it was presumably much easier for the Fisher brothers to credibly pledge to share with the GM owner a large part of the benefits that they might obtain from sales of bodies to other auto makers. (If anything, the problem might be to withhold these benefits in case GM fails to compensate adequately.) However, fully satisfactory treatment of contracting inside firms might require a relational contracting model of the firm, such as Baker, Gibbons, and Murphy (2002) or Halonen (2002).

the negotiation now concerns project choice rather than governance.

### 2.5 Solution concept

We confine attention to consistent equilibria, as defined by Bernheim and Ray (1989). These are subgame-perfect Nash equilibria that, in case a subgame has multiple equilibria, prioritizes efficiency in a recursive fashion. A key implication is that a consistent equilibrium cannot rely on the threat of playing an inefficient equilibrium in future subgames to sustain desired behavior early on in the game; the equilibrium must be "renegotiation-proof".

### 2.6 Governance mode

A governance mode $G$ specifies the number of active suppliers and, for each supplier whether that supplier is in control of its own effort, whether it has acquired an asset, and if so who controls that asset. The governance mode is determined by Date 0 contracting.

We say that a governance mode is optimal if it maximizes welfare (the sum of expected profits) among the feasible governance modes. Denote such an optimal mode $G^*$.

In general, there is no guarantee that governance modes that emerge in equilibrium will be optimal.

## 3 Bargaining

Bargaining potentially takes place both at Date 3 and at Date 0. Date 3 bargaining is most straightforward, since there is no subsequent strategic interaction. We therefore first study this type of "post-entry" bargaining.

### 3.1 Post-entry bargaining

If there are only two parties involved in post-entry bargaining, we rely directly on the analysis of Ellingsen and Miettinen (2008, 2014). Such two-party bargaining is relevant in the case that only one supplier has entered. However, if more than one supplier has entered, there is instead multi-party bilateral bargaining. The multi-party extension of the bargaining model is newly developed here.

Let $Q$ (for quasi-rent) denote the gains that the parties bargain over. (In Section 4, we link $Q$ to the actual payoffs of the model. The quasi-rent to be bargained over typically depends on the state and the governance mode.) Before the players engage in actual negotiation, each player $i$ makes an observable commitment, $\bar{x}_i \in [0,1] \cup \{n\}$, where $n$ denotes "no commitment". The commitment specifies *the largest share* of $Q$ that the committing player is *willing to offer* the trading partner. Making any commitment $\bar{x}_i \geq 0$ entails a small cost $c$ for player $i$, whereas $\bar{x}_i = \{n\}$ is costless.

After the commitment attempts are made, bilateral negotiations start. In the case of more than one supplier, all negotiations are conducted simultaneously. Each negotiation involves a one-shot bargaining game where proposal power is allocated randomly. With probability 1/2, the buyer makes take-it-or-leave-it (TIOLI) contract proposals, and with probability 1/2, each supplier (independently and simultaneously) makes a TIOLI contract proposal to the buyer. Denote player $i$'s negotiation offer $x_i$. The proposer observes the responder's commitment, and whether or not the responder has a loophole, before making the offer.

Commitment attempts $x_i \in [0, 1]$ translate to binding commitments as follows. With probability $\rho$ the commitment sticks, and player $i$ is constrained to make offers $x_i \leq \bar{x}_i$ and to only accept the opponent's offer if $x_j \geq 1 - \bar{x}_i$. With probability $1 - \rho$, the commitment attempt fails, and player $i$ is free to propose any $x_i$ and to accept or reject any $x_j$. Following Schelling (1956), we say that a player whose commitment attempt fails has a *loophole*. We assume that the success of commitment attempts is i.i.d. across players.

If there are several suppliers, the buyer observes all offers before responding. If the buyer makes offers, each supplier sees all offers before responding. Responders who are able to accept (independently and simultaneously) decide whether to accept or reject the relevant proposal(s). If a proposal is accepted, that proposal becomes a binding contract, and the surplus is divided accordingly.

**Two-party post-entry bargaining**

Our specification of the two-party case is encompassed by Ellingsen and Miettinen (2014).[24] Their Proposition 1 can be restated as follows.

**Proposition 1.** *Suppose $c < Q(1-\rho)^2/2$. Then, there is a unique subgame perfect Nash equilibrium outcome of the two-party bargaining game. At the commitment stage, each player $i$ commits to $\bar{x}_i = 0$. At the negotiation stage, if chosen to be proposer and being flexible, player $i$ proposes $x_i = 1 - \bar{x}_j$ if $j$'s commitment sticks and $x_i = 0$ if $j$ has a loophole. Flexible responders accept these proposals.*

We refer to their Appendix A for the full proof. However, it is easy to see that no player can profitably deviate. In particular, if player $i$ deviates to $\bar{x}_i > 0$ or $\bar{x}_i = n$, player $i$'s payoff falls when the opponent has a loophole and is selected to propose, while remaining unchanged in all other states. The flexible strategy $\bar{x}_i = n$ reduces the commitment cost $c$, but this gain is too small to be profitable when $c < Q(1-\rho)^2/2$.

The proof of uniqueness is more subtle, but a key part of the intuition is that it is never an equilibrium to commit to exactly compatible demands, $\bar{x}_i + \bar{x}_j = 1$, since the deviation to $\bar{x}_i = n$ by a player with $x_i \geq 1/2$ is profitable.

Proposition 1 has two immediate implications that pin down the two-party bargaining payoffs, which we shall be using below.

[24] They allow discounting between the commitment stage and the negotiation stage, and the surplus is normalized to 1, but these differences are immaterial.

**Corollary 1.** *(i) The two players reach agreement with probability* $\varphi = 1 - \rho^2$*. (ii) The expected gains from trade* $\varphi Q$ *are split equally.*

While this one-shot bargaining model is very simple, Ellingsen and Miettinen (2014) show that closely similar equilibrium payoffs arise in a more realistic dynamic model, where the parameter $\rho$ captures the expected duration of a stochastically depreciating commitment, the equilibrium offers become less aggressive, and the inefficiency comprises stochastic delay rather than stochastic complete impasse.

### Multi-party post-entry bargaining

Suppose now instead that at least two suppliers have entered, and that the buyer needs to contract with exactly one of them to generate the surplus $Q$. This form of competition eliminates all the inefficiency. There is agreement with probability 1 and the buyer obtains the entire surplus.

**Proposition 2.** *(i) The subgame perfect equilibrium payoffs are unique; the buyer obtains Q and each supplier obtains 0. (ii) All suppliers and the buyer pick* $\bar{x}_i = n$ *at the commitment stage.*

**Corollary 2.** *One strategy profile yielding these payoffs has the additional features: (a) At the negotiation stage, if the buyer is proposing, the buyer offers a contract entailing* $x_b^i = 0$ *to one supplier and no contract to the others. (b) If suppliers are proposing, each supplier proposes an exclusive contract yielding* $x_i = 1$ *to the buyer. (c) In either case, a single offer is accepted.*

The proof runs as follows. First, we check that no player has an incentive to deviate from the posited strategy profile. The buyer earns $Q$, so clearly cannot earn more in any equilibrium. At the negotiation stage, each supplier earns 0 from contracting regardless of what action is taken, both as a proposer and a responder. Deviating at the commitment stage is even less fruitful; it incurs the commitment cost $c$ without generating any compensating benefit at the negotiation stage. Specifically, a supplier setting $\bar{x}_i < 1$ will not be chosen by the buyer if the buyer proposes and will be beaten by a flexible competitor if suppliers propose.

Second, we prove that there are no subgame perfect equilibrium outcomes involving commitment. Let $\underline{x} < 1$ denote the lowest offer that any supplier $i$ commits to in some pure or mixed-strategy equilibrium. If this offer is not in the support of the opponents' strategies, or the support of opponents' strategies is atomless at this point, it would be profitable for supplier $i$ to save $c$ by deviating to $\bar{x}_i = n$. If instead another supplier has an atom at $\underline{x}$, then all other suppliers must pick $\underline{x}$ with positive probability (because the offer is beaten by all higher offers and will also be beaten at the negotiation stage if some supplier is flexible). Let $p$ be the probability that there is a tie between all suppliers at $\underline{x}$. Consequently, there is a supplier that gets to trade at this offer with probability at most $p/N$. But then this supplier can raise the expected revenue to $(1 - (\underline{x} + \epsilon))p$ from (at most) $(1 - \underline{x})p/N$ by raising the offer by a tiny $\epsilon$—a contradiction.

Third, we prove by contradiction that there is no equilibrium of the negotiation game without prior commitment that yields a positive payoff for any supplier: If the buyer proposes, a slightly less generous offer is also accepted by any supplier earning a positive profit. If suppliers propose, the Bertrand logic applies.[25]

Fourth, we prove by contradiction that there is no equilibrium of the negotiation game without prior commitment in which the all suppliers earn 0 and the buyer earns $q < Q$. For any strategy profile yielding this outcome, the contract proposer could deviate to offering an exclusive contract that implements the surplus $Q$ while yielding the buyer $q + \delta$ and a supplier $Q - q - \delta$, with $\delta \in (0, Q - q)$. This contract is preferable to both parties.

## 3.2 Pre-entry bargaining

At Date 0, when bargaining regulates entry rather than project choice, the logic is largely similar. However, there are three notable differences.

First, since pre-entry bargaining occurs before any investments are made, parties potentially bargain over the full rents rather than over quasi rents. Thus, we replace $Q$ (the quasi-rent) by $W$ in many cases below.

Second, at Date 0 the final payoffs are not determined by the outcome of the bargaining, so it is not entirely obvious how to think about the commitments. This problem is especially pronounced if multiple contracts are going to be signed.

Third, there is sometimes an entirely new complication: If the parties fail to reach any agreement, the continuation game may have multiple equilibria. We start by addressing this problem, which arises even if only a single contract is signed.

### Multi-party pre-entry bargaining: One supplier sought

Unsuccessful bargaining always gives rise to a continuation game in which each of the three suppliers decides either IN or OUT. Suppose that if exactly one supplier chooses IN, that supplier earns a positive profit $u + v_R - f$ and the others earn 0. (The monopoly supplier is always able to extract some surplus from the buyer, as we shall see.) If more than one supplier chooses IN, these suppliers earn $v_R - f < 0$ while anyone choosing OUT again earns 0. (There is fierce competition, and suppliers do not recoup the cost of entry.) This continuation game has several subgame-perfect Nash equilibria. For example, it has three asymmetric pure-strategy equilibria in which exactly one supplier enters, and it has a symmetric mixed-strategy equilibrium in which each supplier enters with equal probability. The mixed strategy equilibrium is inefficient, and since there are more efficient equilibria, it is

[25] Suppose the supplier earns a positive profit $g - f$. To do so, the supplier must get the contract with positive probability—thus offering the most attractive contract from the buyer's perspective. If the supplier gets the contract with probability 1, other suppliers earn 0. Thus, one of these suppliers will find it optimal to offer a better contract. If the supplier gets the contract with probability smaller than 1, it means that another supplier offers a contract that is equally good from the buyer's perspective. In this case, the supplier can earn more, because the contract can be obtained with probability 1 by making a slightly larger transfer to the buyer.

ruled out by the consistency requirement. However, there are also many consistent equilibria and the buyer may not always be able to extract all the surplus.

For example, consider the case where the buyer is selected to make a proposal. Suppose the buyer tries to extract the surplus by offering the same terms as in Corollary 2(a). Then, the suppliers may coordinate on the following equilibrium of the continuation game: The designated supplier rejects the contract proposal, and subsequently chooses IN at Date 1, while the two other suppliers, who were not offered a contract, choose OUT. In this case, the designated supplier earns $u + v_R$ rather than $v_R$ in the continuation game, and thus $u + v_R - f$ overall.

**Proposition 3.** *When the buyer proposes, and seeks a contract with a single autonomous supplier, there are consistent equilibria in which this supplier obtains a positive surplus, $u > 0$.*

To complete the example, we only need to prove that there are rational off-equilibrium behaviors such that the buyer cannot make a deviation from the posited equilibrium to a preferred set of proposals. Suppose for example that the buyer deviates by offering no contract to the designated supplier, but a somewhat better contract to one of the other suppliers—extracting a positive but somewhat smaller transfer. In this case, that other supplier declines the offer and then enters (mimicking the behavior of the originally designated supplier), thus saving the transfer. Thus, the buyer gains nothing from the deviation. Similar constructions apply to all other buyer deviations.

Indeed, if $u + v_R - f \geq 0$, we think that $t = 0$ is a plausible outcome when the buyer proposes, and would be even more plausible in the (realistic) case when each supplier only observes the own received offer. On the other hand, the buyer will always extract all the surplus if *suppliers* make the offers.

Let $W_1 = W(G^*)$ denote the maximum (second-best) welfare (sum of expected profits) that is realized if it is optimal that exactly one supplier enters. (This expected gain clearly depends on the trading frictions at Date 3, which in turn depend on the governance mode. Again, we leave until Section 4 the computation of these payoffs.)

**Proposition 4.** *Suppose suppliers make the offers and that the welfare is largest if exactly one supplier enters. In this case, the unique equilibrium payoffs are that the buyer obtains $W_1$ and that each of the suppliers obtain 0.*

To prove the result, let $\pi_s$ denote the expected revenue that the entering supplier gets under optimal governance $G^*$. Consider the following strategy profile: Each supplier offers a contract in which they pay a transfer $t = \pi_s - f$ for being the exclusive supplier, and the buyer accepts one of these proposals. Clearly, this is an equilibrium, since no supplier can earn more than 0 by making a different contract proposal. (For example, a higher transfer entails a loss and a lower transfer implies the buyer will not sign.)

Obviously, there cannot be any equilibrium in which the total payoff is higher than the maximum welfare. Thus, it only remains to prove that there cannot be equilibria in which some supplier earns more (or the buyer earns less). This follows by the standard Bertrand logic (see Footnote 25).

**Multi-party pre-entry bargaining: Two suppliers sought**

Because competition can be desirable at Date 3, contracting at Date 0 might be seeking simultaneous agreements between the buyer and two suppliers. Let us therefore finally consider the case when realization of the maximal surplus $W(G^*)$ requires that at least two suppliers are present at Date 3.

Of course, if $f < v_R$ suppliers enter even if there is no contract with the buyer. Therefore, the buyer might as well refrain from Date 0 contracting and rely on spot market competition at Date 3; Proposition 2 establishes that the outcome is efficient and the buyer extracts all the surplus.

If instead $f > v_R$ the buyer needs to make two bilateral contracts at Date 0. Since competition at Date 3 drives suppliers' share of the surplus to 0 by Proposition 2, the buyer must compensate each of the two suppliers for $f - v_R$.

But if it is desirable for the buyer to contract with two suppliers, the value of one bilateral contract depends on whether another contract is also signed. This affects the interpretation of the commitments that players make before the contracts are proposed. We need a definition of commitment that is relative to the equilibrium the players are coordinating on.

Specifically, our assumption is that players' commitments concern shares of the surplus that their contract will create if signed, over and above the surplus that would otherwise be created, conditional on play conforming to the posited continuation equilibrium. For example, suppose the buyer makes contract offers and the offers to suppliers 1 and 2 yield both of them non-negative profits if both sign. Consider a continuation equilibrium in which these two contracts, and no other contracts are signed. If supplier 1 has made a commitment to reject surplus shares below $1 - x_1$, this commitment defines a smallest acceptable share of "the surplus that is added if supplier 1 signs this proposal given that supplier 2 signs the buyer's other proposal". In other words, if the buyer's proposal yields supplier 1 a smaller fraction of this surplus than $1 - x_1$, supplier 1 will reject the proposal despite thereby reducing his payoff. (As it turns out, such commitments will not be made along an equilibrium path, but it's still necessary to specify how they would have worked.)

**Proposition 5.** *Suppose realization of the maximal (second-best) surplus $W(G^*) > 0$ requires that at least two suppliers are present at Date 3. (i) Consistent equilibrium payoffs are unique; the buyer obtains $W(G^*)$ and each supplier obtains 0. (ii) In any equilibrium, the commitments are $\bar{x}_i = n$ for all players i.*

**Corollary 3.** *One strategy profile yielding these payoffs has the additional features: (a) If B proposes, offers are $(t_i, o, a, e) = (f - v_R, s, s, 0)$ to two of the suppliers and no contract offer to the third. (b) If* *suppliers propose, all three suppliers make the contract offers* $(f - v_R, s, s, 0)$. *(c) In either case, two offers are accepted.*

Clearly, no player has any incentive to deviate from the proposed strategy profile. The buyer gets all the available surplus. All suppliers earn 0, and a supplier that attempts to get a positive surplus through a less generous offer will have its offer rejected by the buyer.

To prove uniqueness, the steps in the proof of Proposition 2 apply, but with the additional observation that any Nash equilibrium that is (recursively) Pareto-dominated by another Nash equilibrium fails consistency.[26]

# 4 Analysis

We can now derive the model's solution recursively. (A good fraction of the analysis merely entails inserting the actual payoffs instead of the generic surplus $Q$ in the analysis above.)

## 4.1 Date 3 outcomes

The final decision node is characterized by a state of nature ($H$ or $L$) and a governance mode.

If the efficient project choice benefits a player that holds all the control rights, the player implements the project.

If instead bargaining is required to attain an efficient outcome, we use the results of Section 3 to characterize the expected payoffs associated with each governance mode.

For each relevant case, we report the ex ante expected payoffs for each player, including the cost of assets.

- *No assets:* In this trivial case, all payoffs are 0.

  In the next case, $f < v_R$.

- *Three assets. Autonomous suppliers (Spot trade)*: By Proposition 2, the outcome entails no bargaining frictions. Welfare, is

$$\begin{aligned} W^{ST} &= h\left[v_T^H - v_R\right] + 3(v_R - f) \\ &= W^*. \end{aligned} \tag{7}$$

  Suppliers each earn $v_R - f$. The buyer earns $h(v_T^H - v_R)$.

  In remaining cases, $f > v_R$.

[26] There could potentially be coordination failure between the two suppliers that need to contract with the buyer. If suppliers propose, it is a Nash equilibrium that no supplier makes an offer, since a single offer is not acceptable by the buyer. However, as this equilibrium is inefficient, it is excluded by the consistency requirement. It could potentially also be a Nash equilibrium for no supplier to sign the contract that the buyer proposes. If so, this equilibrium too would fail the consistency requirement. (As it turns out, this is not even a Nash equilibrium in our model.)

- *Two assets. Autonomous suppliers (PS)*: By Proposition 2, the outcome entails no bargaining frictions. Welfare, is

$$\begin{aligned} W^{PS} &= h\left[v_T^H + v_R\right] + 2(1-h)v_R - 2f \\ &= W^* - (f - v_R). \end{aligned} \tag{8}$$

  Suppliers each earn 0. The buyer earns $W^{PS}$.

- *One asset. Integration (I): B* owns the asset and has employed $S_i$, who acquired the asset. State $L$ then entails negotiation between $B$ and $S_i$. Corollary 1(ii) implies that $B$'s payoff is $v_T^L + \varphi(v_R - v_T^L)/2$ and $S_i$'s payoff is $\varphi(v_R - v_T^L)/2$. In state $H$, B orders the tailored project and receives $v_T^H$. The other supplier earns 0. Welfare is

$$\begin{aligned} W^{I} &= hv_T^H + (1-h)\left[\varphi(v_R - v_T^L) + v_T^L\right] - f \\ &= W^* - (1-h)(1-\varphi)(v_R - v_T^L). \end{aligned} \tag{9}$$

  The second expression, which isolates the welfare loss relative to the first-best payoff (attainable with a single asset), facilitates the comparison to other governance payoffs.

  The loss obtains because there is bargaining about the integrated supplier's value in state $L$. Bargaining is unavoidable in this state, because $B$ can threaten to refuse the supplier's request to spend effort on the regular project instead of the inferior tailored project.

- *One asset. Autonomous supplier, single sourcing (SS): $S_i$* acquired the asset, owns the asset, and is not employed. State $H$ then entails negotiation between $B$ and $S_i$. By Corollary 1(ii), $B$'s payoff is $\varphi(v_T^H - v_R)/2$ and $S_i$'s payoff is $\varphi(v_T^H - v_R)/2 + v_R$. In state $L$, $S_i$ chooses project $R$ and receives $v_R$. The other supplier has no assets and earns 0. Welfare is

$$\begin{aligned} W^{SS} &= h\left[\varphi(v_T^H - v_R) + v_R\right] + (1-h)v_R - f \\ &= W^* - h(1-\varphi)(v_T^H - v_R). \end{aligned} \tag{10}$$

  Here, the loss occurs because there is bargaining over $B$'s value in state $H$; the supplier can threaten to withhold the asset.

  The two modes of governance I and SS share the feature that one party holds all the control rights. There are two other possible modes of governing a bilateral relationship. Both of them split the rights, so that the asset owner does not have the right to assign effort and vice versa. Such a separation of the control rights is harmful for two reasons. First, the parties need to negotiate in both states of the world rather than in only one of them. Second, failed negotiations now paralyze the relationship, since no party can produce anything with the single resource that it controls. Hence, these organizational forms are payoff-dominated and will never be proposed in Date 0 bargaining.[27]

- *Two assets - one exclusive contract (EC)*: Suppose both suppliers own assets but only one of them has an exclusive contract with the buyer. Competitive pressure does not apply in state $H$ since the supplier with the exclusive contract gets to deal with the buyer. $B$ earns $\varphi(v_T^H - v_R)/2$ and the supplier with contract earns $\varphi(v_T^H - v_R)/2 + v_R$ in state $H$. $B$ earns 0 and each supplier earns $v_R$ in state $L$. Welfare is

$$\begin{aligned} W^{EC} &= h\varphi(v_T^H - v_R) + 2v_R - 2f \\ &= W^* - (f - v_R) - h(1-\varphi)(v_T^H - v_R). \end{aligned} \tag{11}$$

  Notice that this form of organization yields a lower surplus than both SS and PS: there is overinvestment in assets without any gains from competition. Hence, it will not occur in equilibrium.

- *Two assets, integration (I2):* Suppose $B$ owns two assets, each operated by different employed suppliers. In state $H$, the buyer would like to sell the right to one of the suppliers to undertake project $R$ while commanding $T$ from the other. There is competition between suppliers for the additional rent, $(v_T^H - v_R)$, thus allowing $B$ to extract all of this rent, by Proposition 2(i).

  However, in state $L$ there is no competition, as both suppliers should ideally undertake project $R$. The existence of two assets, one for each bargaining table, is common knowledge. Thus, bargaining is bilateral at each table and it is optimal for the parties to attempt commitments before bargaining starts. Whenever one of the players in a bargaining pair is flexible, that pair will earn $v_R$. If there is flexibility in one pair and inflexibility in the other, $B$ will command project $T$ from that other supplier. But if both pairs end up unable to agree, $B$ will only command project $T$ from one of them, as $B$ doesn't have the capacity to implement two projects $T$. Thus, the other supplier will be

[27] For completeness, here is a description of these two outcomes:

- *One asset. Labor integration (L1):* $S_i$ acquired the asset, owns the asset, but is employed. There is then negotiation in both states. If negotiation fails, there is no production, since no party controls the necessary resources. Welfare is

$$\begin{aligned} W^{L1} &= h\varphi v_T^H + (1-h)\varphi v_R - f \\ &= W^* - h(1-\varphi)v_T^H - (1-h)(1-\varphi)v_R. \end{aligned}$$

- *One asset. Asset integration (A1):* $B$ owns the asset, but has not employed any supplier. $S_i$ acquired the asset. As in case L1, there is negotiation in both states, and welfare is accordingly the same,

$$\begin{aligned} W^{A1} &= h\varphi v_T^H + (1-h)\varphi v_R - f \\ &= W^* - h(1-\varphi)v_T^H - (1-h)(1-\varphi)v_R. \end{aligned}$$

allowed project $R$. As a result, the welfare is

$$\begin{aligned} W^{I2} &= h\left[v_T^H + v_R\right] + (1-h)\left[\varphi^2 2v_R + (1-\varphi^2)(v_R + v_T^L)\right] - 2f \\ &= W^* - (f - v_R) - (1-h)(1-\varphi^2)(v_R - v_T^L). \end{aligned} \tag{12}$$

Compared to case PS, the loss (captured by the last term) stems from bilateral bargaining replacing competition in state $L$. Even if full integration of both suppliers avoids the splitting of control rights, the problem in state $L$ remains, exactly as with a single asset and full integration, I. It follows that PS payoff-dominates I2.

Clearly, PS also payoff-dominates all other governance modes involving two assets, since these modes have the additional drawback of splitting the control rights, thus generating greater bargaining frictions as well as paralysis when negotiations fail.

## 4.2 Optimal governance

Before we derive the equilibrium outcomes, it is useful to compare the expected welfare associated with the different governance modes. We say that the governance mode, $G^*$, is (second-best) optimal if it maximizes welfare among the available modes. Trivially, spot market interaction is optimal when entry costs are so low that $f < v_R$. The remaining results derive directly from (9)-(12).

**Proposition 6.** *Suppose $f > v_R$. The optimal governance mode, call it $G^*$, is*
*(i) PS if $W^{PS} \geq 0$ and*

$$f - v_R \leq \min\{h(1-\varphi)(v_T^H - v_R), (1-h)(1-\varphi)(v_R - v_T^L)\}; \tag{13}$$

*(ii) I if $W^I \geq 0$ and*

$$(1-h)(1-\varphi)(v_R - v_T^L) \leq \min\{h(1-\varphi)(v_T^H - v_R), f - v_R\}; \tag{14}$$

*(iii) SS if $W^{SS} \geq 0$ and*

$$h(1-\varphi)(v_T^H - v_R) \leq \min\{(1-h)(1-\varphi)(v_R - v_T^L), f - v_R\}. \tag{15}$$

*Otherwise inactivity is optimal.*

### Optimality of parallel sourcing

Let us begin by interpreting Condition (13), under which subsidized parallel sourcing is optimal. The left-hand side says the obvious, namely that subsidized competition is only desirable if the entry cost $f$ does not exceed $v_R$ by too much. The right-hand side also has one straightforward message, namely that competition is particularly desirable when the bargaining friction $1-\varphi$ is large. But the right-hand side also has a more subtle message: The value of competition depends on what the best alternative governance mode fails to achieve. If $h$ is 1, integration doesn't fail to achieve anything. Likewise, if $h$ is close to 0, single-sourcing from an autonomous supplier is optimal. Hence, the gains from competition are greatest in the intermediate case when both $I$ and $SS$ are likely to entail bargaining friction. Finally, if we keep the expected $v_T$ fixed but increase the variance, i.e., decrease $v_T^L$ and increase $v_T^H$ the right-hand side also increases. The reason is again that the bargaining friction that remains under $I$ or $SS$ matters more.

**Integration or separation?**

When the entry cost is too high to accommodate a second supplier, efficiency favors either an autonomous supplier (SS) or a fully integrated supplier (I). From Conditions (14) and (13), we see that integration is better than separation if

$$h(1-\varphi)(v_T^H - v_R) > (1-h)(1-\varphi)(v_R - v_T^L).$$

We can rearrange the inequality, using Equations (1) and (2), to get the following result.

**Proposition 7.** *If there is a single supplier, vertical integration, $I$, generates higher welfare than vertical separation, $SS$, if $\bar{\pi}_T > \pi_R$ and lower welfare when the inequality is reversed.*

Intuitively, vertical integration is optimal when—in expectation—it is more valuable to protect the buyer's preferred project from bargaining friction than to protect the supplier's preferred project. Recall from (6) that an alternative formulation of this condition is that the product of demand favorability and asset specificity exceeds one.

Figure 2 summarizes the results in the $(h, f)$ parameter space, i.e, keeping asset specificity constant and varying the probability that the buyer's project is most profitable.

[Figure 2 around here.]

Observe that for parameters in the upper area under the dashed line in Figure 2 there is underinvestment compared to the first-best allocation. Unlike in Grossman and Hart (1986), this underinvestment occurs despite investments being contractible. The reason is that future bilateral bargaining is unavoidably inefficient and not merely a source of distorted investment incentives.

## 4.3 Equilibrium governance

We are now ready to study to what extent the above constrained optimal outcomes are implemented in equilibrium.

**Proposition 8.** *(i) If the suppliers make the Date 0 proposals, equilibrium governance coincides with optimal governance, $G^*$. (ii) If the buyer makes the Date 0 proposals, optimum and equilibrium do not always coincide. Specifically, there may exist a consistent equilibrium with parallel sourcing although single-sourcing from an autonomous supplier would generate higher welfare.*

Part (i) of Proposition 8 follows directly from Propositions 4 and 5. In short, the Bertrand logic applies whenever suppliers make the proposals.

We stress that in case (i), the transfer ensures that the supplier earns zero profit, i.e., it is

$$t = -v_R - \frac{h\varphi(v_T^H - v_R)}{2} + f. \tag{16}$$

The transfer can be either negative, in which case the supplier pays for having the exclusive right to supply, or positive, in which case the buyer subsidizes the entry cost and exclusivity does not really need to be specified.

The intuition behind part (ii) of Proposition 8 is that buyer may be incapable of capturing all the surplus when the buyer makes the proposals and ideally would want a single autonomous supplier to enter—see the example in Proposition 3. Thus, the buyer might be able to do better by inducing two suppliers to enter instead of only one; even if total surplus is smaller, the buyer obtains a larger share of it.

Formally, suppose $f > v_R$. In equilibria with a single autonomous entrant and $t = 0$ the buyer obtains the payoff $h\varphi(v_T^H - v_R)/2$ at Date 3. Note that there is entry only if the single supplier is profitable, i.e., if

$$v_R + \frac{h\varphi(v_T^H - v_R)}{2} > f. \tag{17}$$

If the buyer instead chooses parallel sourcing, by offering nonexclusive contracts to two suppliers, it obtains the payoff $h(v_T^H - v_R)$ at Date 3, but must compensate each of the two suppliers for entering, $(f - v_R)$. Thus, the buyer prefers parallel sourcing if

$$h(v_T^H - v_R)(1 - \varphi/2) > 2(f - v_R). \tag{18}$$

This is a larger set of parameters than the set for which parallel sourcing is optimal. As an illustration, consider the example in Figure 2. With those parameter values, single sourcing yields a higher social surplus than parallel sourcing as long as $f > 2 + 2h/5$, but the buyer is better off with single sourcing only if $f > 2 + 3h/5$. (When single-sourcing is best for the buyer, despite the buyer proposing, the outcome is a Master Supply Agreement, where the buyer does not extract any rent at the date of contracting.)

Before we end this section, let us note that Proposition 7 carries over from the normative domain to the positive domain. That is, our model predicts that firms that rely on a single supplier will be vertically integrated with this supplier whenever the profitability of the buyer's preferred project is higher than that of the seller's preferred project. Lower costs of producing the tailored output (high asset specificity) is only one of the conditions favoring this outcome. High demand for the tailored output is the other.

## 5 Final remarks

Like Oliver Williamson, we take the view that arm's-length trade in thin input markets are hampered by bargaining frictions. With a single vertically separated supplier, our model predicts that there is always bargaining friction when the buyer has the most profitable project, i.e., whenever there is a transaction between the two parties. However, unlike Williamson, we posit that similar bargaining frictions can also occur under vertical integration. The difference is that internal bargaining happens only for a (potentially small) fraction of the transactions: Whenever the buyer has the most profitable project, there is no friction, as the buyer can implement the favored project without the supplier's agreement.

Thus, the notion that there is relatively little bargaining under vertical integration is an outcome of the model—vertical integration occurs precisely when it is unlikely that the bargaining option would have been valuable. By contrast, under vertical separation all interactions between the buyer and the supplier require bargaining. Of course, this greater proportion of bargaining interactions in markets does not imply that bargaining is also more frequent in *absolute* terms. To the contrary, vertical separation is chosen when interaction between the two parties is relatively unlikely.

When there is large uncertainty about who will have the more profitable project, dealing with a single supplier entails a high risk of bargaining friction both under integration and under separation. To eliminate the bargaining friction, and sometimes also to capture a larger fraction of the surplus, the buyer might then offer to sponsor the entry of more than one supplier.

Despite its simplicity, the model generates additional testable implications. For example, it suggests that vertical integration in settings with high entry costs is likely not only when asset specificity is high (a supply-side reason), but also when the buyer's preferred use of the asset is valuable for demand-side reasons.

That said, the model mutes many potentially relevant mechanisms. Let us briefly mention three of them. First, and perhaps most controversially, we have assumed that agents are not allowed to trade any claims to their returns. If it were possible, without any transaction costs, for an $S$ to sell all future returns from regular projects to $B$, full integration would produce an efficient outcome. The reason is that $B$ would voluntarily pick the regular project in State $L$; there would no longer be a need to haggle in this state. Such comprehensive advantages to integration invites the question: Why are not all firms integrated? Like Williamson, Masten (1986) answers that integration must impose costs of bureaucracy. We think the truth is somewhere between Masten's assumption and ours. It is possible to guarantee the buyer some of the returns from the supplier's favorite project, but it is difficult to pledge all these returns (see Footnote 17). A natural next step is therefore to consider more realistically incomplete financial contracts and how they interact with other control rights.

Second, we assume that information is symmetric. It might be more realistic to assume that the buyer has some private information about the value of the tailored project, $v_T$, especially if this involves the use of new technologies. In this case, contract design becomes a signaling problem and equilibrium governance might well be distorted as a result.

Third, we ignore the role of trust. Since there is no possibility for building trust over time, the model fails to capture how contracting interacts with relationship-building between and within organizations, as discussed by Bernstein (2015), Hadfield and Bozovic (2016), and Frydlinger and Hart (2024), for example. All of the above three complications are related to the issue of complexity, which our model also fails to represent.[28]

One reason to pursue some of these extensions is that they would enable us to get a better grasp on governance dynamics, which are likely to respond to changes in information and other features of the contracting environment. Some firms start out within a small niche and gradually integrate with suppliers or customers. Other firms are vertically integrated to begin with but subsequently shed business lines to focus on core activities. Indeed, the type of vertically integrated production that was common for much of the twentieth century and that Chandler (1977) and Chandler (1990) eulogized as late as 1990 was condemned soon thereafter, by Stuckey and White (1993) and Quinn and Hilmer (1994) among others. In the present model, three rationales for such a reversal from vertical integration to vertical separation is (i) a reduction of assets' sunk costs, (ii) a reduction of assets' relationship-specificity, or (iii) a reduction in the demand for buyers' preferred projects, perhaps due to tougher downstream competition.[29] Our hunch is that changes to contractibility, possibly along the lines suggested by the above extensions, might have been important too.

[28]A referee suggests that dynamics seem essential for understanding the nature of firms: "Unlike contracts, we do not see firms used for a single transaction." Our defense is that the allocation of decision rights may matter the most in connection with a few very large decisions, such as the location of the Fisher Body plant in Flint. If those decisions are in fact driving governance choices, models focusing on a few decisions might do a satisfactory job. On the other hand, our relatively static framework clearly comes with costs. For example, it is not obvious how contracts are to be enforced inside a firm. Also, our model is unable to address the distinction between transactions that are likely to be frequently repeated and transactions that are not—a central point of the infinite-horizon model of Wernerfelt (2015).

[29]Here, we have not modeled the competition that the buyer might face. We have also ignored whether vertical integration might affect competitors' input access—the problem of foreclosure; see Hart and Tirole (1990).

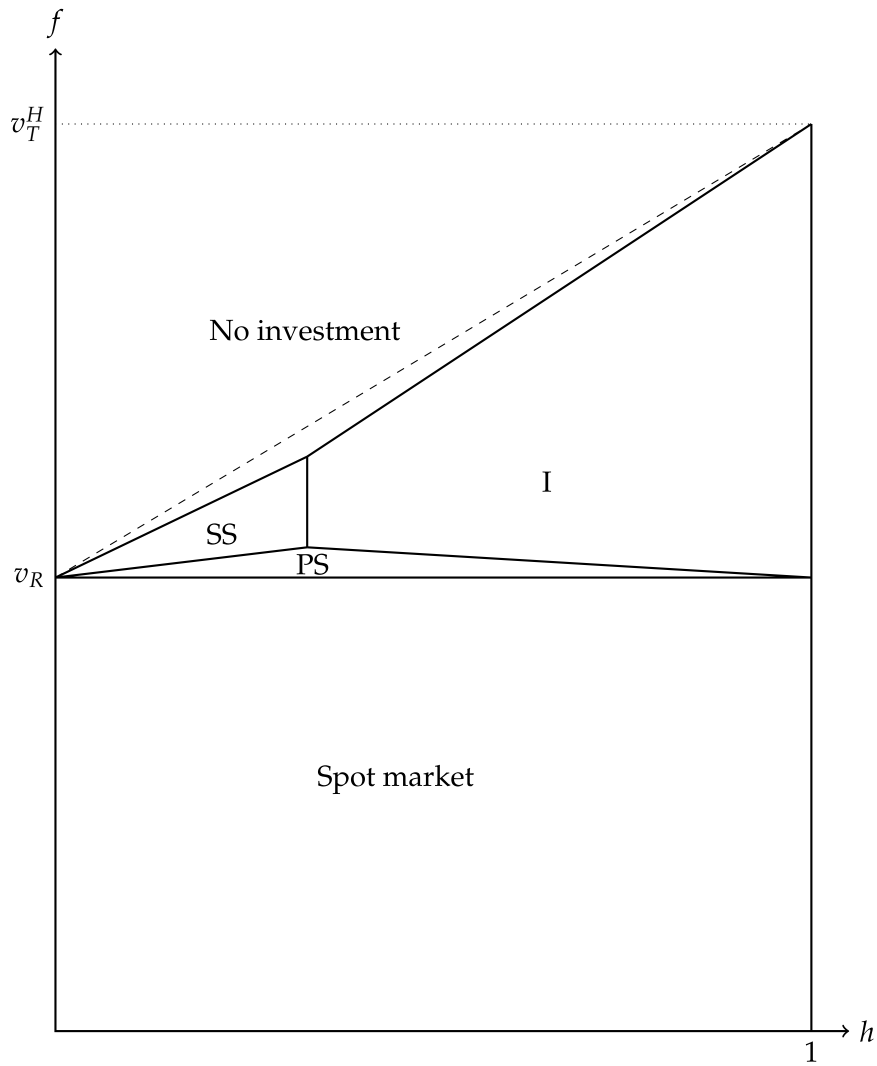


Figure 2: Second-best governance outcomes

The horizontal axis measures the probability $h$ that the tailored project will be optimal. The vertical axis measures the sunk cost. Parameter regions are labeled by the (uniquely) second-best governance choice $G^*$. $SS$ denotes single sourcing, $I$ denotes integration, $PS$ denotes parallel sourcing. No investment is first-best above the dashed line. The figure is drawn for the case when the suppliers are proposing contracts and the parameter values $v_T^H = 4, v_R = 2, v_T^L = 1, \varphi = 4/5$; if the buyer proposes, the SS region is smaller and the PS and I regions larger.

## References


Aghion, P., D. Fudenberg, R. Holden, T. Kunimoto, and O. Tercieux (2012). Subgame-perfect implementation under information perturbations. *Quarterly Journal of Economics 127*(4), 1843–1881.

Andersson, U., M. Forsgren, and U. Holm (2007). Balancing subsidiary influence in the federative mnc: a business network view. *Journal of International Business Studies 38*.

Bajari, P. and S. Tadelis (2001). Incentives versus transaction costs: A theory of procurement contracts. *Rand Journal of Economics 32*(3), 387–407.

Baker, G., R. Gibbons, and K. J. Murphy (2002). Relational contracts and the theory of the firm. *Quarterly Journal of Economics 117*(1), 39–83.

Basak, D. and J. Deb (2020). Gambling over public opinion. *American Economic Review 110*(11), 3492–3521.

Bernheim, B. D. and D. Ray (1989). Collective dynamic consistency in repeated games. *Games and economic behavior 1*(4), 295–326.

Bernstein, L. (2015). Beyond relational contracts: Social capital and network governance in procurement contracts. *Journal of Legal Analysis 7*(2), 561–621.

Bolton, P. and M. D. Whinston (1993). Incomplete contracts, vertical integration, and supply assurance. *Review of Economic Studies 60*(1), 121–148.

Bresnahan, T. and J. Levin (2013). Vertical integration and market structure. In R. Gibbons and J. Roberts (Eds.), *Handbook of Organizational Economics*, Chapter 21, pp. 853–890. Princeton: Princeton University Press.

Carlton, Dennis, W. (1979). Vertical integration in competitive markets under uncertainty. *Journal of Industrial Economics 27*(3), 189–209.

Chandler, A. D. (1977). *The: The Managerial Revolution in American Business*. Belknap Press.

Chandler, A. D. (1990). *Scale and Scope: The Dynamics of Industrial Capitalism*. Harvard University Press.

Coase, R. (2000). The acquisition of Fisher Body by General Motors. *Journal of Law and Economics 43*(1), 15–31.

Crawford, V. P. (1982). A theory of disagreement in bargaining. *Econometrica 50*(3), 607–637.

Cusumano, C. M., F. Fabbri, and F. Pieroth (2024). Competing to commit: Markets with rational inattention. *American Economic Review 114*(1), 285–306.

Dutta, R. (2024). Bargaining as a struggle between competing attempts at commitment. *Review of Economic Studies 91*(5), 2771–2805.

Ellingsen, T. and T. Miettinen (2008). Commitment and conflict in bilateral bargaining. *American Economic Review 98*(4), 1629–35.

Ellingsen, T. and T. Miettinen (2014). Tough negotiations: Bilateral bargaining with durable commitments. *Games and Economic Behavior 87*, 353–366.

Ellingsen, T. and E. Paltseva (2016). Confining the Coase theorem: Contracting, ownership, and free-riding. *Review of Economic Studies 83*(2), 547–586.

Fehr, E., O. Hart, and C. Zehnder (2011). Contracts as reference points: Experimental evidence. *American Economic Review 101*(2), 493–525.

Freeland, R. (2000). Creating holdup through vertical integration: Fisher Body revisited. *Journal of Law and Economics 43*(1), 33–66.

Frydlinger, D. and O. D. Hart (2024). Overcoming contractual incompleteness: The role of guiding principles. *Journal of Law, Economics and Organization 40*(3), 625–647.

Gibbons, R. (2005). Four formal(izable) theories of the firm? *Journal of Economic Behavior & Organization 58*(2), 200–245.

Grossman, S. J. and O. D. Hart (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy 94*(4), 691–719.

Hadfield, G. K. and I. Bozovic (2016). Scaffolding: Using formal contracts to support informal relations in support of innovation. *Wis. L. Rev.*, 981.

Halonen, M. (2002). Reputation and the allocation of ownership. *Economic Journal 112*(481), 539–558.

Hart, O. and B. Holmstrom (2010). A theory of firm scope. *Quarterly Journal of Economics 125*(2), 483–513.

Hart, O. and J. Moore (1990). Property rights and the theory of the firm. *Journal of Political Economy 98*(1), 1119–1158.

Hart, O. and J. Moore (2008). Contracts as reference points. *Quarterly Journal of Economics 123*(1), 1–48.

Hart, O. D. and J. Tirole (1990). Vertical integration and market foreclosure. *Brookings Papers on Economic Activity: Microeconomics*, 205–286.

Holmstrom, B. and P. Milgrom (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization 7*(S), 24–52.

Katz, A. (1990). Your terms or mine? the duty to read the fine print in contracts. *Rand Journal of Economics 21*(4), 518–537.

Klein, B. (2007). The economic lessons of Fisher Body–General Motors. *International Journal of the Economics of Business 14*(1), 1–36.

Klein, B., R. Crawford, and A. Alchian (1978). Vertical integration, appropriable rents, and the competitive contracting process. *Journal of Law and Economics 21*(2), 297–326.

Lafontaine, F. and M. Slade (2007). Vertical integration and firm boundaries: The evidence. *Journal of Economic Literature 45*(3), 629–685.

Leider, S. and W. S. Lovejoy (2016). Bargaining in supply chains. *Journal of Law and Economics 62*(10), 3039–3058.

Loertscher, S. and L. M. Marx (2022). Incomplete information bargaining with applications to mergers, investment, and vertical integration. *American Economic Review 112*(2), 616–649.

Loertscher, S. and M. H. Riordan (2019). Make and buy: Outsourcing, vertical integration, and cost reduction. *American Economic Journal: Microeconomics 11*(1), 105–123.

Lunnan, R., S. Tomassen, U. Andersson, and G. R. Benito (2019). Dealing with headquarters in the multinational corporation: a subsidiary perspective on organizing costs. *Journal of Organization Design 8*(1), 1–24.

Masten, S. E. (1984). The organization of production: Evidence from the aerospace industry. *Journal of Law and Economics 27*(2), 403–417.

Masten, S. E. (1986). Institutional choice and the organization of production: The make-or-buy decision. *Journal of Institutional and Theoretical Economics 142*(3), 493–509.

Masten, S. E. (2002). Modern evidence on the firm. *American Economic Review, Papers and Proceedings 92*(2), 428–432.

Meyer, M., P. Milgrom, and J. Roberts (1992). Organizational prospects, influence costs, and ownership changes. *Journal of Econoics and Management Strategy 1*(1), 9–35.

Miettinen, T. and C. Vanberg (2025). Commitment and conflict in unanimity bargaining. *American Economic Journal: Microeconomics*.

Milgrom, P. (1988). Employment contracts, influence activities, and efficient organization design. *Journal of Political Economy 96*(1), 42–60.

Milgrom, P. and J. Roberts (1990). Bargaining costs, influence costs, and the organization of economic activity. In J. E. Alt and K. Shepsle (Eds.), *Perspectives on Positive Political Economy*, Chapter 3, pp. 57–89. Cambridge: Cambridge University Press.

Nöldeke, G. and K. M. Schmidt (1995). Option contracts and renegotiation: A solution to the hold-up problem. *Rand Journal of Economics 7*(1), 163–179.

Pei, H. (2025). Commitment and agreement in repeated bargaining. *Manuscript*.

Quinn, J. B. and F. G. Hilmer (1994). Strategic outsourcing. *MIT Sloan Management Review 35*(4), 43–55.

Raith, M. (2023). Employee or contractor? A theory of the labor boundaries of the firm. *SSRN* (4517923).

Rajan, R. G. and L. Zingales (2002). Power in a theory of the firm. *Quarterly Journal of Economics 113*(2), 387–432.

Ravid, D. (2020). Ultimatum bargaining with rational inattention. *American Economic Review 110*(9), 2948–2963.

Schelling, T. C. (1956). An essay on bargaining. *American Economic Review 46*(3), 281–306.

Simon, H. A. (1951). A formal theory of the employment relationship. *Econometrica 19*(3), 293–305.

Stuckey, J. and D. White (1993). When and when not to vertically integrate. *McKinsey Quarterly*, 3–3.

Tadelis, S. (2002). Complexity, flexibility, and the make-or-buy decision. *American Economic Review, Papers and Proceedings 92*(2), 433—-437.

Tadelis, S. and O. Williamson (2013). Transaction cost economics. In R. Gibbons and J. Roberts (Eds.), *Handbook of Organizational Economics*, Chapter 4, pp. 159–192. Princeton: Princeton University Press.

Van den Steen, E. (2010). Interpersonal authority in a theory of the firm. *American Economic Review 100*(1), 466–490.

Watson, J. (2024). Contractual chains. *Econometrica 92*(5), 1735–1774.

Wernerfelt, B. (2002). Why should the boss own the assets? *Journal of Economics and Management Strategy 11*(3), 473–485.

Wernerfelt, B. (2015). The comparative advantages of firms, markets and contracts: A unified theory. *Economica 82*(326), 350–367.

Wernerfelt, B. (2016). *Adaptation, Specialization, and the Theory of the Firm*. Cambridge University Press, Cambridge UK.

Williamson, O. E. (1971). The vertical integration of production: Market failure considerations. *American Economic Review 61*(2), 112–123.

Williamson, O. E. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. Macmillan, New York.

Williamson, O. E. (1985). *Economic Institutions of Capitalism*. Macmillan, New York.

Williamson, O. E. (1991). Comparative economic organization: The analysis of discrete structural alternatives. *Administrative Science Quarterly 36*(2), 233–261.

Wolitzky, A. (2023). Unobserved-offers bargaining. *American Economic Review 113*(1), 136–173.