

Organization and Information: Firms' Governance Choices in Rational-Expectations Equilibrium

Robert Gibbons, Richard Holden and Michael Powell*

August 23, 2012

Abstract

We analyze a rational-expectations model of price formation in an intermediate-good market under uncertainty. There is a continuum of firms, each consisting of a party who can reduce production cost and a party who can discover information about demand. Both parties can make specific investments at private cost, and there is a machine that either party can control. As in incomplete-contracting models, different governance structures (i.e., different allocations of control of the machine) create different incentives for the parties' investments. As in rational-expectations models, some parties may invest in acquiring information, which is then incorporated into the market-clearing price of the intermediate good by these parties' production decisions. The informativeness of the price mechanism affects the returns to specific investments and hence the optimal governance structure for individual firms; meanwhile, the governance choices by individual firms affect the informativeness of the price mechanism. In equilibrium, the informativeness of the price mechanism can induce *ex ante* homogeneous firms to choose heterogeneous governance structures. (JEL: D20, D23).

*Gibbons: Massachusetts Institute of Technology, rgibbons@mit.edu. Holden: University of New South Wales, richard.holden@unsw.edu.au. Powell: Kellogg School of Management, mikepowell@kellogg.northwestern.edu. We thank Daron Acemoglu, Philippe Aghion, Mathias Dewatripont, Glenn Ellison, Oliver Hart, Bengt Holmstrom, Matt Selove, Sarah Venables, Birger Wernerfelt, Oliver Williamson, the editors, three anonymous referees, and seminar participants at Chicago, Columbia, Harvard, MIT, MBS, Munich, Northwestern, Penn, Queen's, Sydney, Toronto, ULB, UNSW and USC for helpful comments. All three authors thank MIT Sloan's Program on Innovation in Markets and Organizations for financial support, Holden thanks the University of Chicago Booth School of Business for financial support, and Powell thanks the NSF.

1 Introduction

Economists have long celebrated the market’s price mechanism for its ability to aggregate and transmit information (Hayek, 1945; Grossman, 1976). This informativeness of the price mechanism raises the possibility that the market could substitute for certain information-gathering and communication activities within a firm, thereby affecting the firm’s optimal design. But as Grossman and Stiglitz (1976, 1980) pointed out, market equilibrium must be internally consistent. For example, when information is costly to acquire, market prices cannot be fully informative, otherwise no party would have an incentive to acquire information in the first place.

In quite a different tradition, scholars in business history and strategic management have long argued that it is enormously difficult for one firm to do two things exceedingly well. At a general level, Chandler (1962) famously argued that a firm’s strategy and organizational structure are inextricably linked, so it will be hard to pursue two strategies if this requires two structures. More specifically, Porter (1985: 23) noted that “Cost leadership usually implies tight control systems, overhead minimization, pursuit of scale economies, and dedication to the learning curve; these could be counterproductive for a firm attempting” a different strategy. Similarly, Roberts (2004: 255) suggested that different strategies can involve “quite different tasks, calling on different organizational capabilities and typically requiring different organizational designs to effect them.” In short, in choosing a competitive strategy and an organizational design, it often pays to focus.

In this paper we view firms and the market as institutions that shape each other: in industry equilibrium, each firm takes the informativeness of the price mechanism as an important parameter in its choice of strategy and design, but these decisions in turn effect the firm’s participation in the market and hence the informativeness of the price mechanism. We thus complement the large and growing literature on how an organization’s design affects its members’ incentives to acquire and communicate information.¹ In particular, our analysis

¹See Milgrom and Roberts (1988), Holmstrom and Tirole (1991), and Aghion and Tirole (1997) for early

shows how one firm’s optimal strategy and design depend not only on the uncertainty that firm faces but also on the strategy and design decisions that other firms make. For example, if the market price is very informative, then many firms will choose designs that sacrifice incentives for information gathering in favor of those for other activities (specifically, cost reduction), effectively free-riding on the informativeness of the price mechanism. But the Grossman-Stiglitz insight implies that not all firms can free-ride, lest the price mechanism contain no information.

To pursue these issues, we develop a rational-expectations model similar to Grossman and Stiglitz (1976, 1980) but applied to a market for an intermediate good (e.g., prices and net supply are non-negative and the players are risk-neutral). As in other rational-expectations models, the price mechanism both clears the market and conveys some information from informed to uninformed parties. The fact that the price is not perfectly informative provides the requisite incentive for some parties to pay the cost of acquiring further information. Relative to many rational-expectations models, the innovation here is the enrichment from individual investors to firms, in two respects. First, each firm chooses one of two focused strategies/designs—one that inspires a party within the firm to collect costly information, or another that inspires a different party to undertake cost reduction.

The second way that our firms differ from the individual investors in the classic rational-expectations model of an asset market is that our firms do not buy or sell an asset based on expectations of changes in that asset’s price; instead, our firms choose whether or not to buy in an input market, based on expectations of the value of final goods. This separation of the topic of information (the value of final goods) from the market in which prices convey information (the market for intermediate goods) suggests other applications of our model. For example, the uncertainty might concern whether tariff barriers will change or whether a new technology will fulfill its promise. Interestingly, however, not all sources of uncertainty will do: our rational-expectations model requires some element of common-value uncertainty

work and Alonso et al. (2008) and Rantakari (2008) for a sample of recent work; see Bolton and Dewatripont (2012) and Gibbons et al. (2012) for surveys.

(possibly partially correlated rather than perfectly common values) rather than pure private-value uncertainty. As Grossman (1981: 555) put it, in non-stochastic economies (and certain economies with pure private-value uncertainty), “No one tries to learn anything from prices [because] there is nothing for any individual to learn.” Often, however, there is something to learn from prices, such as when there is an element of common-value uncertainty.

To model our firms, we develop a simplified version of the classic incomplete-contracting approach initiated by Grossman and Hart (1986), but applied to the choice of governance structure within an organization (akin to Aghion and Tirole (1997)). To keep things simple, our incomplete-contracts model involves only a single control right (namely, who controls a machine that is necessary for production) and hence two feasible organizational designs. Regardless of who controls the machine, each party can make a specific investment, but the incentives to make these investments depend on who controls the machine. Following the incomplete-contracts approach (i.e., analyzing one firm in isolation) reveals that the optimal organizational design is determined by the marginal returns to these investments. In our model all firms are homogeneous *ex ante*, so an incomplete-contracts analysis of a single firm would prescribe that all firms choose the same organizational design. Relative to the incomplete-contracts approach, the novel component of our model is the informativeness of the price mechanism, which endogenizes the returns to the parties’ specific investments and hence creates an industry-level determinant of an individual firm’s choice of organizational design.

As one example of how the informativeness of the price mechanism and firms’ strategic choices interact, consider firms like Apple versus firms like Dell. Such firms participate in many of the same input markets, and broad industry trends affect demand at both kinds of firms. Although firms like Dell have very different strategy and structure from firms like Apple, the former might nonetheless learn something by observing prices in the latter’s input markets. For example, suppose Dell observed a change in the pricing or availability of “electronics manufacturing services” from firms such as Flextronics, which provide critical

outsourced manufacturing and assembly services for original electronics equipment manufacturers. Dell might then update its beliefs about Apple's production plans. Quite consistent with this scenario, both Dell and Apple are indeed large customers of Flextronics, and in July 2011 a senior Flextronics executive pleaded guilty to insider-trading charges involving Apple's production plan.² The possibility of Dell inferring information about demand for its products from the availability or price of electronics manufacturing services parallels our model, in that it is the market-clearing price of an intermediate good (or, here, service) that provides information about demand for a final good.

In summary, our model integrates two familiar approaches: rational expectations (where an imperfectly informative price mechanism both permits rational inferences by some parties and induces costly information acquisition by others) and incomplete contracts (where equilibrium investments depend on the allocation of control, and control rights are allocated to induce second-best investments). Our main results are that: (1) under mild regularity conditions an equilibrium exists; (2) *ex ante* identical firms may choose heterogeneous organizational designs; and (3) firms' choices of organizational design and the informativeness of the price mechanism interact. In fact, in our model, certain organizational designs may be sustained in market equilibrium *only* because the price system allows some firms to benefit from the information-acquisition investments of others. We also provide comparative statics on the proportion of firms that choose one organizational design or the other.

Grossman and Helpman (2002), Legros and Newman (2008) and Legros and Newman (2009) analyze other interactions between firms' governance structures and the market. These papers differ from ours in two respects. First, in modeling firms' choice of governance structures, they focus on the boundary of the firm (i.e., the integration decision) whereas we focus on the organizational design (specifically, the allocation of control within the organization). Second, we focus on the informativeness of the price mechanism, whereas

²See, for instance, <http://dealbook.nytimes.com/2011/07/05/executive-pleads-guilty-to-leaking-apple-secrets/>

they focus on different aspects of the market.³ As Grossman (1981: 555) suggests, however, such models are not useful “as a tool for thinking about how goods are allocated... when... information about the future... affects current prices.” In contrast to the aforementioned papers, our model focuses on the informative role of prices—transferring information from informed to (otherwise) uninformed parties. We therefore see our approach as complementary to these others: in economies with uncertainty, the price mechanism clears the market and communicates information and hence can affect how firms design their structures and processes to acquire and communicate information within the firm; without uncertainty, however, governance and pricing can still interact, for the reasons explained in these papers.

The remainder of the paper proceeds as follows. In Section 2 we specify and discuss the model. Section 3 analyzes the organizational-design choice of a single firm in isolation, and Section 4 analyzes the informativeness of the price mechanism, taking firms’ organizational-design choices as given. Section 5 then combines the incomplete-contracts and rational-expectations aspects of the previous two sections, analyzing the equilibrium choices of organizational designs for all the firms in the industry and hence deriving our main results. Section 6 offers an enrichment of our model in terms of firms’ choices about their boundaries and discusses how our approach relates to existing theories of firm boundaries. Section 7 discusses our model’s implications for empirical work on organizational structures and firms’ boundaries, and Section 8 concludes.

³For example, Grossman and Helpman (2002) view the market as a matching mechanism, where efficiency increases in the number of firms participating; the quality of matching determines the returns to outsourcing, which then depends on how many other firms choose outsourcing. And in Legros and Newman (2009), supply and demand determine prices, which in turn determine the return to parties’ actions and hence the parties’ optimal governance structures; meanwhile the parties’ actions in turn determine supply and demand, so governance and pricing interact.

2 The Model

2.1 Overview of the Model

We begin with an informal description of our model. There is a continuum of firms, each consisting of an “engineer” and a “marketer” who both participate in a production process that can transform one intermediate good (a “widget”) into one final good. Any firm may purchase a widget in the intermediate-good market. Each firm has a machine that can transform one widget into one final good at a cost. The engineer in a given firm has human capital that allows her to make investments that reduce the cost of operating that firm’s machine. Likewise, the marketer in a given firm has human capital that allows him to make investments that deliver information about the value of a final good.

As is standard in incomplete-contracting models, the parties’ incentives to make investments depend on the allocation of control. There are two possible organizational designs (i.e. governance structures inside the firm): marketing control and engineering control. In particular, in our model, only the party that controls the machine will have an incentive to invest. Thus, in firms where the marketer controls the machine, the marketer invests in information about the value of the final good, whereas in firms where the engineer controls the machine, the engineer invests instead in cost reduction and relies solely on the price mechanism for information about the value of the final good. Naturally, if the price mechanism is more informative, the returns to investing in information are lower so firms have a greater incentive to choose engineer control and invest instead in cost reduction. As in rational-expectations models, however, when fewer firms invest in gathering information, the price mechanism becomes less informative, thereby making marketer control more attractive. An industry equilibrium must balance these two forces. We show that, given a rational-expectations equilibrium, a unique equilibrium exists and is often interior (even though firms are identical *ex ante*). In this sense, the price mechanism induces heterogeneous behavior

among homogeneous firms.⁴

In subsection 2.5 we offer an elaboration of the basic model where rather than the uncertainty being about an existing good, it is about the value of a new product. There are now two production periods, and in order for the new product to be produced the machine must be taken “offline” and “retooled” in the first period. One can think of this “retooling” as devoting resources to innovation and new product development. In the spirit of Christensen (1997), the new product created by informed firms may be more valuable than the current product produced by uninformed firms. The cost of devoting these resources to innovation is the inability to produce the existing product in the first period. The controller of the machine now has a choice between producing the existing good in both periods, or the new good in second period but nothing in the first. We feel that this elaboration fits the Apple-Dell example (and others like it) quite well. We show that despite the additional model complexity, none of the existence results or qualitative predictions of the model is altered. But since this elaborated model adds notational complexity we perform our analysis in the body of the paper on the basic model.

2.2 Statement of the Problem

There is a unit mass of risk-neutral firms. Each firm $i \in [0, 1]$ consists of two parties, denoted E_i and M_i , and a machine that is capable of developing one intermediate good (a “widget”) into one final good at cost $c_i \sim U[\underline{c}, \bar{c}]$. The machine can be controlled by either party, but it is firm-specific (i.e., the machine is useless outside the firm) and its use is non-contractible (i.e., only the party who controls the machine can decide whether to operate it). If party E_i controls the machine, we say that the governance structure in firm i is $g_i = E$, whereas if party M_i controls the machine, we say that $g_i = M$.

Final goods have an uncertain value. Party M_i can invest at cost K_M to learn the value of a final good in the market, $v \sim U[\underline{v}, \bar{v}]$. If M_i incurs this cost, E_i observes that M_i is

⁴We label our parties “engineer” and “marketer” because their investments produce cost reductions and demand forecasts, respectively.

informed but does not herself observe v . Party E_i can invest at cost K_E in reducing the cost of operating the firm's machine. If E_i incurs this cost, M_i observes that E_i invested, so it is common knowledge that c_i is reduced to $c_i - \Delta$, where $\Delta \leq \underline{c}$. Both of these investments are non-contractible (e.g., for E_i , neither the act of investing nor the resulting cost is contractible).

We embed these firms in a rational-expectations model of price formation in intermediate good markets. Firms may purchase widget(s) in the intermediate-good market. The supply of widgets, x , is random and inelastic. Assume $x \sim U[\underline{x}, \bar{x}]$.

Equilibrium in the market for widgets occurs at the price p that equates supply and demand (from informed and uninformed firms). In making decisions about purchasing a widget, firms that are not directly informed about v (from investments by their marketers) make rational inferences about v from the market price for widgets. Firms choose their governance structures (i.e., machine control) taking into account the information they will infer from the market price and hence the relative returns to their two parties' investments.

2.3 Timing and Assumptions

We now state the timing and assumptions of the model more precisely. We comment on these assumptions in Section 2.4. There are six periods.

[FIGURE 1 HERE]

In the first period, industry-level uncertainty is resolved: the value of a final good v is drawn from $U[\underline{v}, \bar{v}]$ and the widget supply x is drawn from $U[\underline{x}, \bar{x}]$, but neither of these variables is observed by any party.

In the second period, the parties in each firm negotiate a governance structure $g_i \in \{E, M\}$: under $g_i = E$, party E_i controls the machine that can develop one widget into one final good; under $g_i = M$, party M_i controls this machine. This negotiation of governance structure occurs via Nash bargaining.

In the third period, parties E_i and M_i simultaneously choose whether to make non-contractible investments (or not) at costs K_E and K_M , respectively. The acts of making these investments are observable but not verifiable, but the outcome of the marketer's investment (namely, learning v) is observable only to M_i , not E_i .

In the fourth period, production planning takes place, in two steps. In period 4a, the parties E_i and M_i commonly observe $c_i \sim U[\underline{c}, \bar{c}]$, the raw cost of running their machine, as well as $\delta_i \in \{0, \Delta\}$, the amount of cost reduction achieved by E_i 's specific investment. Also, M_i (but not E_i) observes $\varphi_i \in \{\emptyset, v\}$, a signal about the value v of the final good, where $\varphi_i = \emptyset$ is the uninformative signal received if party M_i has not invested K_M in period 3, and $\varphi_i = v$ is the perfectly informative signal received if K_M has been invested. We use the following notation for the parties' information sets: $s_i^M = (c_i, \delta_i, \varphi_i)$, $s_i^E = (c_i, \delta_i, \emptyset)$, and $s_i = (s_i^M, s_i^E)$. In period 4b, the market for widgets clears at price p . In particular, any firm may buy a widget ($w_i = 1$) but will not demand more than one widget because the machine can produce only one final good from one widget.

In the fifth period, production occurs: if the party in control of the machine in firm i has a widget, then he or she can run the machine to develop the widget into a final good at cost $c_i - \delta_i$. We denote the decision to produce a final good by $q_i = 1$ and the decision not to do so by $q_i = 0$. In principle, off the equilibrium path, one party might control the machine and the other have a widget, in which case the parties bargain over the widget and then the machine controller makes the production decision. We assume that cashflow rights and control rights are inextricable, so that whichever party controls the machine owns the final good (if one is produced) and receives the proceeds.

Finally, in the sixth period, final goods sell for v and payoffs are realized. The expected payoffs (before v is realized) are

$$\pi_{E_i}^{g_i} = 1_{\{g_i=E\}} 1_{\{w_i=1\}} [1_{\{q_i=1\}} (E[v|s_i^E, p(\cdot, \cdot) = p] - c_i + \delta_i) - p], \text{ and}$$

$$\pi_{M_i}^{g_i} = 1_{\{g_i=M\}} 1_{\{w_i=1\}} [1_{\{q_i=1\}} (E [v|s_i^M, p(\cdot, \cdot) = p] - c_i + \delta_i) - p]. \quad (1)$$

2.4 Discussion of the Model

Before proceeding with the analysis, we pause to comment on some of the modeling choices we have made.

First, we assume that the machine is firm-specific. This assumption allows us to focus on the market for widgets by eliminating the market for machines. By allowing both markets to operate, one could analyze whether the informativeness of one affects the other.

Second, we have only one control right (over the machine) and hence only two candidate governance structures. Our choice here is driven purely by parsimony; extending the model to allow more assets (and hence more governance structures) could allow more interesting activities within organizations than our simple model delivers.

Third, we make the strong assumption that control of the machine and receipt of cashflow from selling a final good are inextricably linked. We expect that richer models based on weaker assumptions would yield similar results (if they can be solved).

Fourth, we have binary investments in cost reduction and information acquisition (at costs K_E and K_M , respectively), rather than continuous investment opportunities. It seems straightforward to allow the probability of success (in cost reduction or information acquisition) to be an increasing function of the investment level, which in turn has convex cost.

Fifth, we assume inelastic widget supply x . This uncertain supply plays the role of noise traders, making the market price for widgets only partially informative about v , so that parties may benefit from costly acquisition of information about v .

Sixth, our assumptions that all the random variables are uniform allow us to compute a closed-form (indeed, piece-wise linear) solution for the equilibrium price function for the intermediate good. This tractability is useful in computing the returns to alternative governance structures, at the firm level, and hence the fraction of firms choosing each governance structure, at the industry level.

Seventh, as in Grossman-Stiglitz and the ensuing rational-expectations literature, our model of price formation is a reduced-form model of price-taking behavior, rather than an extensive-form model of strategic decision-making (which might allow information transmission during the price-formation process, either by the parties as described in our model or by one party who separates from his engineer and becomes something like a marketer).

2.5 Alternative Formulation: New Products

The idea that the information contained in prices can influence the governance structure of a firm does not rely critically on the uncertainty being about the demand for a final good. In this subsection, we show that the framework developed in sections 2.1 – 2.3 is equivalent to a model in which the uncertainty concerns the potential profitability of a new product.

Each dyad $i \in [0, 1]$ still consists of two parties, E_i and M_i . There is still a single (firm-specific) machine that can be controlled by either E_i or M_i , and its use is non-contractible. Now, however, production occurs over two production periods, and there are two options facing the controller of the machine. The machine can either be (1) used in the production of the *current* final good in each period or (2) taken offline for a period, retooled, and then deployed toward the production of a *new* good. The current good can be produced in both production periods if the dyad purchases a single widget at price p . In this case, the current good sells for $\frac{1}{2}v_0$ in each period and costs $\frac{1}{2}(c_i - \delta_i)$ in each period to produce. Production of the new good does not require a widget, and it yields a net benefit of 0 in the first production period and v_1 in the second production period. As a departure from sections 2.1 to 2.3, v_0 is commonly known, but v_1 is uncertain. Define $v = v_0 - v_1$ and assume that $v \sim U[\underline{v}, \bar{v}]$. Party M_i can invest K_M to learn about v_1 (and hence v). Party E_i can invest K_E to reduce the cost of producing the current good (but not the new good) by Δ . As before, the supply of widgets, x , is random and inelastic. Assume $x \sim U[\underline{x}, \bar{x}]$.

There are now seven periods:

1. Uncertainty resolution: $v \sim U[\underline{v}, \bar{v}]$ and $x \sim U[\underline{x}, \bar{x}]$ are drawn. Neither is observed.

2. Governance structure determination: E_i and M_i negotiate a governance structure $g_i \in \{E, M\}$ via Nash bargaining
3. Investment period: E_i and M_i simultaneously decide whether or not to invest at costs K_E and K_M , respectively.
4. Production planning: E_i and M_i commonly observe $c_i \sim U[\underline{c}, \bar{c}]$ and $\delta_i \in \{0, \Delta\}$. Also, M_i (but not E_i) observes $\varphi_i \in \{0, v\}$. The market for widgets clears at price p . Any firm may buy a widget at this price.
5. Production period 1: The party with control chooses either to produce or to retool the machine. Production is possible only if the dyad purchased a widget in period 4. If the party decides to produce, then one unit of the current final good is produced at cost $\frac{1}{2}(c_i - \delta_i)$ and sold into the market at price v_0 . If the party decides to retrofit the machine, then no production occurs.
6. Production period 2: If dyad i has produced in the previous period, it can produce again and receive another $\frac{1}{2}v_0 - \frac{1}{2}(c_i - \delta_i)$. If dyad i has not produced in the previous period, it can produce the new good and receive net surplus v_1 .
7. Payoffs are realized. Define q_i to be equal to 1 if production occurs in periods 5 and 6 and equal to 0 otherwise. The expected payoffs (before v is realized) are

$$\begin{aligned} \pi_{E_i}^{g_i} &= \mathbf{1}_{\{g_i=E\}} \mathbf{1}_{\{w_i=1\}} \left[\mathbf{1}_{\{q_i=1\}} \left(E \left[v \mid s_i^E, p(\cdot, \cdot) = p \right] - c_i + \delta_i \right) - p \right] \\ &\quad + \mathbf{1}_{\{g_i=E\}} E \left[v_1 \mid s_i^E, p(\cdot, \cdot) = p \right], \text{ and} \end{aligned}$$

$$\begin{aligned} \pi_{M_i}^{g_i} &= \mathbf{1}_{\{g_i=M\}} \mathbf{1}_{\{w_i=1\}} \left[\mathbf{1}_{\{q_i=1\}} \left(E \left[v \mid s_i^M, p(\cdot, \cdot) = p \right] - c_i + \delta_i \right) - p \right] \\ &\quad + \mathbf{1}_{\{g_i=M\}} E \left[v_1 \mid s_i^M, p(\cdot, \cdot) = p \right]. \end{aligned}$$

Here, s_i^E and s_i^M are defined as in section 2.3. The first line in each of the above

expressions is the same as (1). We show in the appendix that the second line in each expression does not affect any of the qualitative predictions of the model.

In light of this, as we mentioned earlier, we now return to basic model (i.e. a single production period) for the remainder of the paper.

3 Individual Firm Behavior

As a building block for our ultimate analysis, we first analyze the behavior of a single firm taking the market price p as given. Optimal behavior involves purchasing a widget only if one is going to produce. Define the gross surplus to the parties in a firm as $GS_i^{g_i} = \pi_{M_i}^{g_i} + \pi_{E_i}^{g_i}$, i.e.

$$GS_i(g_i, s_i) = 1_{\{q_i=1\}} [E[v|s_i^{g_i}, p(\cdot, \cdot) = p] - p - (c_i - \delta_i)].$$

The efficient production decision is $q_i^* = 1$ if $E_{x,v}[v|s_i^{g_i}, p] \geq p + c_i - \delta_i$, and the maximized expected gross surplus in period 4 is then

$$GS_i^*(g_i, s_i) = E_{x,v}[(v - c_i + \delta_i - p) q_i^*(g_i, s_i, p) | s_i^{g_i}, p].$$

Recall that the controller of the machine both controls the production decisions and receives the cashflows. Consequently, the other party receives zero. These payoffs determine the parties' investment incentives in period 3, as follows.

Let the subscript pair $(I, 0)$ denote the situation in which M_i invested and hence is informed about v but E_i did not invest in cost reduction. Likewise (U, Δ) , denotes the situation in which M_i did not invest but E_i did, hence reducing production costs by Δ , and

$(U, 0)$ denotes the situation in which neither invested. Now define the following:

$$\begin{aligned}\pi_{I,0} &= E_{c_i} [GS_i^* (M, s_i)] \text{ if } \varphi_i = v, \delta_i = 0, \\ \pi_{U,\Delta} &= E_{c_i} [GS_i^* (E, s_i)] \text{ if } \varphi_i = \emptyset, \delta_i = \Delta, \text{ and} \\ \pi_{U,0} &= E_{c_i} [GS_i^* (g_i, s_i)] \text{ if } \varphi_i = \emptyset, \delta_i = 0.\end{aligned}$$

Formally, these expectations are triple integrals over (c_i, x, v) space:

$$\begin{aligned}\pi_{I,0} &= \int_{\underline{v}}^{\bar{v}} \int_{\underline{x}}^{\bar{x}} \int_{\underline{c}}^{v-p(x,v)} (v - p(x, v) - c_i) dF(c_i, x, v), \\ \pi_{U,\Delta} &= \int_{\underline{v}}^{\bar{v}} \int_{\underline{x}}^{\bar{x}} \int_{\underline{c}}^{E[v|p]-p(x,v)+\Delta} (v - p(x, v) + \Delta - c_i) dF(c_i, x, v), \text{ and} \\ \pi_{U,0} &= \int_{\underline{v}}^{\bar{v}} \int_{\underline{x}}^{\bar{x}} \int_{\underline{c}}^{E[v|p]-p(x,v)} (v - p(x, v) - c_i) dF(c_i, x, v),\end{aligned}$$

where F is the joint distribution function.

Since one party's expected payoff in period 4 is independent of its investment, at most one party will invest in period 3. If E_i controls the machine ($g_i = E$), she will invest if $\pi_{U,\Delta} - K_E \geq \pi_{U,0}$. Similarly, if M_i controls the machine ($g_i = M$), he will invest if $\pi_{I,0} - K_M \geq \pi_{U,0}$. We assume that K_E and K_M are small relative to the benefits of investment, so the party that controls the machine will invest.⁵

To proceed, we need to compute the price function $p(x, v)$. This involves analyzing the behavior of other firms, as follows.

⁵This condition can be stated in terms of primitives of the model, but since this is the economic assumption we are making, we state it in this fashion.

4 Rational Expectations in the Market for Intermediate Goods

Recall that there is a unit mass of firms indexed by $i \in [0, 1]$. Who buys a widget? Let $c_M(v, p) = v - p$ be the highest cost at which a marketer who has invested in information (and hence knows v) would be prepared to produce a final good, and similarly let $c_E(p) = E[v|p] - p + \Delta$ be the highest cost at which an engineer who has invested in cost reduction (but not information) would be prepared to produce. Suppose (as we will endogenize below) that a fraction λ of firms have M control (and hence know v), whereas fraction $1 - \lambda$ have E control (and hence costs reduced by Δ). Demand for widgets is therefore

$$\lambda \frac{v - p - \underline{c}}{\bar{c} - \underline{c}} + (1 - \lambda) \frac{E[v|p(x, v) = p] + \Delta - p - \underline{c}}{\bar{c} - \underline{c}}.$$

The market-clearing price equates this demand with the supply, which recall is x , so

$$p = (1 - \lambda) E[v|p(x, v) = p] + \lambda v - (\bar{c} - \underline{c})x + (1 - \lambda)\Delta - \underline{c}.$$

The conditional expectation of v given p therefore must satisfy

$$E[v|p(\cdot, \cdot) = p] \equiv \frac{p + (\bar{c} - \underline{c})x + \underline{c} - (1 - \lambda)\Delta - \lambda v}{1 - \lambda}, \quad (2)$$

where the equivalence relation indicates that (2) must hold as an identity in x and v .

Definition 1 An *industry configuration* is a vector $\eta = (\eta_{I\Delta}, \eta_{I0}, \eta_{U\Delta}, \eta_{U0})$ consisting of the masses of dyads that are, respectively, informed and have cost reduction, informed and do not have cost reduction, uninformed and have cost reduction, and uninformed and do not have cost reduction.

Definition 2 Given an industry configuration, η , a *rational-expectations equilibrium* ("REE") is a price function $p(x, v)$ and a production allocation $\{q_i\}_{i \in [0, 1]}$ such that

1. $q_i = q_i^*(g_i, s_i, p)$ for all i , and
2. The market for widgets clears for each $(x, v) \in [\underline{x}, \bar{x}] \times [\underline{v}, \bar{v}]$.

The fact that the party who does not control the asset receives none of the cashflow implies that this party will not invest, so $\eta_{I,\Delta} = 0$. Furthermore, K_E and K_M small implies $\eta_{U,0} = 0$. Therefore $\lambda = \eta_{I,0}$ and $1 - \lambda = \eta_{U,\Delta}$. The problem of finding a rational-expectations price function in this model thus becomes one of finding a fixed point of (2). In the appendix we solve for this fixed point, showing that it is piecewise-linear over three regions of (x, v) space: a low-price region, a moderate-price region, and a high-price region. This leads to

Proposition 1 *Given an industry configuration, there exists a piecewise-linear price function with three regions that characterizes a rational-expectations equilibrium.*

We prove this proposition and derive the price function in the appendix, but to build some intuition for this result, consider the figure below, which shows the three regions of (x, v) space, R_λ^j for $j = 1, 2, 3$. The low-price region R_λ^1 begins from the lowest feasible price, p_L at (\bar{x}, \underline{v}) , and extends up to the price \bar{p} at (\bar{x}, \bar{v}) . The moderate-price region R_λ^2 then extends from price \bar{p} up to the price \underline{p} at $(\underline{x}, \underline{v})$, where the under- and over-lined notation for prices is chosen to match the (x, v) coordinates. Finally, the high-price region R_λ^3 extends from \underline{p} up to the highest feasible price, p_H at (\underline{x}, \bar{v}) .

[FIGURE 2 HERE]

Within each region, the iso-price loci are linear. In particular, solving $p^j(x, v) = p$ for v yields

$$v = -\frac{\beta_1^j}{\beta_2^j}x + \frac{p - \beta_0^j}{\beta_2^j}$$

as an iso-price line in (x, v) space. Because x and v are independent and uniform, every (x, v) point on this line is equally likely. Thus, after observing p , an informed party projects

this iso-price line onto the v -axis and concludes that the conditional distribution of v given p is uniform, with support depending on which region p is in. For example, if $p < \bar{p}$ then the lower bound on v is \underline{v} and the upper bound is some $\bar{v}(p) < \bar{v}$. Alternatively, if $\bar{p} < p < \underline{p}$ then the lower and upper bounds on v are \underline{v} and \bar{v} , so p is uninformative. Finally, if $p > \underline{p}$ then the lower bound is some $\underline{v}(p) > \underline{v}$ and the upper bound is \bar{v} .⁶

Given this uniform conditional distribution of v given p , the conditional expectation on the left-hand side of (2) is then the average of these upper and lower bounds on v . The coefficients β_0^j, β_1^j , and β_2^j can then be computed by substituting $p^j(x, v)$ for p on both sides of (2) and equating coefficients on like terms so that (2) holds as an identity. The slope of an iso-price line, $-\beta_1^j/\beta_2^j$, is decreasing in λ , meaning that in regions 1 and 3 uninformed parties can make tighter estimates of v from p when more parties are informed.

5 Industry Equilibrium

To recapitulate, Section 3 analyzed the production decision, taking $p(\cdot, \cdot)$ as exogenous, and Section 4 endogenized prices. In this section, therefore, we endogenize the governance-structure choices of each firm and define an industry equilibrium, as follows.

Definition 3 *An industry equilibrium is a set of firms of mass λ^* , a price function $p(x, v)$, and a production allocation $\{q_i\}_{i \in [0,1]}$ such that*

1. *Each firm optimally chooses g_i , with a fraction λ^* choosing $g_i = M$;*
2. *Each party optimally chooses whether or not to invest;*
3. *$q_i = q_i^*(g_i, s_i, p)$ and $w_i = w_i^*(g_i, s_i, p)$; and*
4. *The market for widgets clears for each $(x, v) \in [\underline{x}, \bar{x}] \times [\underline{v}, \bar{v}]$.*

⁶Note that in this model, but not Grossman-Stiglitz, extreme prices are very informative and intermediate prices are less informative. In fact, with the slopes of the price functions as drawn in the above figure, intermediate price are completely uninformative.

The choice in period 2 is between the two possible governance structures: $g_i = E$ or $g_i = M$. Given λ , the *ex ante* expected net surpluses from choosing the two governance structures are

$$\begin{aligned} NS^E(\lambda) &= \pi_{U,\Delta}(\lambda) - K_E, \text{ and} \\ NS^M(\lambda) &= \pi_{I,0}(\lambda) - K_M. \end{aligned}$$

In an interior equilibrium, firms must be indifferent between the two governance structures. Thus our goal is to find λ^* such that $NS^E(\lambda^*) = NS^M(\lambda^*)$ and to characterize how λ^* varies as we change the parameters of the model. For simplicity we assume that $K_E = K_M = K$. (The case where $K_E \neq K_M$ is discussed at the end of this section.) We therefore seek λ^* such that

$$\pi_{I,0}(\lambda^*) = \pi_{U,\Delta}(\lambda^*),$$

or equivalently,

$$\pi_{I,0}(\lambda^*) - \pi_{U,0}(\lambda^*) = \pi_{U,\Delta}(\lambda^*) - \pi_{U,0}(\lambda^*). \quad (3)$$

To keep notation compact, let $\sigma_v = \frac{1}{\sqrt{12}}(\bar{v} - \underline{v})$, $\sigma_x = \frac{1}{\sqrt{12}}(\bar{x} - \underline{x})$, and $\mu_x = (\bar{x} + \underline{x})/2$. We will use the following fact (which is derived in the appendix).

Fact 1 Assume $\lambda \leq (\bar{c} - \underline{c}) \frac{\sigma_x}{\sigma_v}$. Then

$$\begin{aligned} \pi_{I,0}(\lambda) - \pi_{U,0}(\lambda) &= \frac{1}{2} \frac{\sigma_v^2}{\bar{c} - \underline{c}} \left(1 - \frac{1}{2} \frac{\lambda}{\bar{c} - \underline{c}} \frac{\sigma_v}{\sigma_x} \right) \text{ and} \\ \pi_{U,\Delta}(\lambda) - \pi_{U,0}(\lambda) &= \frac{\Delta^2}{\bar{c} - \underline{c}} \lambda - \frac{1}{2} \frac{\Delta^2}{\bar{c} - \underline{c}} + \mu_x \Delta. \end{aligned}$$

Observe that the first expression is decreasing in λ and the second is increasing in λ . This leads to the following characterization of industry equilibrium, under the regularity conditions that $\bar{c}, \underline{c}, \sigma_x, \sigma_v, \Delta > 0$ with $\underline{c} \geq \Delta$. We refer to the case where $(\bar{c} - \underline{c}) \sigma_x \geq \sigma_v$ as the *noisy outside demand* case, and in that case we obtain a closed form solution for the

proportion of firms that choose each governance structure.

Proposition 2 *An industry equilibrium exists, and there is a unique industry configuration associated with the price function characterized in Proposition 1. In the noisy outside demand case, the industry configuration associated with the industry equilibrium is as follows:*

$$\lambda^* = \frac{\sigma_v^2 + \Delta^2 - 2(\bar{c} - \underline{c})\mu_x\Delta}{\frac{\sigma_v^2}{2} \frac{\sigma_v/\sigma_x}{\bar{c}-\underline{c}} + 2\Delta^2} \quad (4)$$

if the right-hand side of (4) is in $[0, 1]$. If the right-hand side of (4) is less than 0, then $\lambda^* = 0$; if it is greater than 1, then $\lambda^* = 1$.

Proof. If $\sigma_v^2 \leq 2(\bar{c} - \underline{c})\mu_x\Delta - \Delta^2$, then $\pi_{U,0}(0) \leq \pi_{U,\Delta}(0)$ and thus, since the left-hand side of (2) is decreasing in λ , it follows that $\lambda^* = 0$. Similarly, if $\sigma_v^2 \left(1 - \frac{1}{2} \frac{1}{\bar{c}-\underline{c}} \frac{\sigma_v}{\sigma_x}\right) \geq 2(\bar{c} - \underline{c})\mu_x\Delta + \Delta^2$, then $\pi_{U,0}(1) \geq \pi_{U,\Delta}(1)$, and since the right-hand side of (2) is increasing in λ , we must have that $\lambda^* = 1$. Otherwise, we want to find λ^* such that

$$\begin{aligned} 0 &= \pi_{I,0}(\lambda^*) - \pi_{U,\Delta}(\lambda^*) \\ &= \frac{\sigma_v^2 + \Delta^2 - 2(\bar{c} - \underline{c})\mu_x\Delta}{2(\bar{c} - \underline{c})} - \frac{\lambda^*}{2(\bar{c} - \underline{c})} \left(\frac{\sigma_v/\sigma_x}{\bar{c} - \underline{c}} \frac{\sigma_v^2}{2} + 2\Delta^2 \right), \end{aligned}$$

which yields expression (4). ■

Proposition 2 is our main result, establishing that, given our rational expectations equilibrium, there exists a unique industry equilibrium and providing an explicit expression for the proportion of firms that choose each of the governance structures. As the proposition makes clear, this proportion may well be interior.⁷ Recall, however, that our firms are homogeneous *ex ante*, so an incomplete-contract style analysis (taking each firm in isolation) would prescribe that they all choose the same governance structure. In this sense, the infor-

⁷Models of industry equilibrium from IO (Ordover, Saloner, and Salop (1990)) and trade (McLaren (2000), Grossman and Helpman (2002), Antras (2003)) typically feature strategic complementarities in governance structure, and hence generically produce equilibria in which *ex ante* identical firms organize identically. One exception to this is Avenel (2008), who shows that investments in cost reduction (and hence governance structures that promote cost reduction) are strategic substitutes when firms compete Bertrand.

mativeness of the price mechanism can induce heterogeneous behaviors from homogeneous firms. To put this point differently, in this model, the price mechanism can be seen as endogenizing the parameters of the incomplete-contract model so that firms are indifferent between governance structures. In a richer model, with heterogeneous investment costs, almost every firm would have strict preferences between governance structures, with only the marginal firm being indifferent.

We are also able to perform some comparative statics. First, when the *ex ante* level of fundamental uncertainty increases (i.e., σ_v is higher), the return to investing in acquiring information increases, so λ increases. An increase in noise (i.e., σ_x is higher) has an identical effect. An increase in μ_x increases the probability of production, which disproportionately benefits *E*-control firms, decreasing λ . Finally, an increase in Δ has two effects. The first is the partial-equilibrium channel through which an increase in the benefits of choosing engineer ownership (and hence investing in cost reduction) makes engineer control relatively more appealing, reducing λ . In an industry equilibrium, however, there is also a price effect. For a fixed fraction $1 - \lambda$ of parties that invest in cost reduction, an increase in Δ makes widgets more valuable, which in turn increases demand and hence average prices. Since firms with engineer control purchase widgets over a larger region of the c_i space than do firms with marketing control, the former face this increase in average price level relatively more than do firms with marketer control, so the price effect militates towards an increase in λ . Which of these two effects dominates depends on the parameters of the model. Collecting these together we have

Proposition 3 *In the noisy outside demand case: (i) an increase in the uncertainty of either the supply of the intermediate good or the value of the final good or a decrease in the average supply of the intermediate good leads to an increase in the fraction of dyads that choose to become marketing-oriented; (ii) an increase in the level of potential cost-reduction leads to an increase in the fraction of dyads that choose to become engineering-oriented if there is sufficient uncertainty regarding the value of the final good. If this level of uncertainty*

is low, the opposite may be true.

Proof. See appendix. ■

Finally, our incomplete-contracts approach sheds new light on the functioning of the price mechanism. In particular, most partially-revealing REE models compare the benefits of acquiring information to the exogenously specified costs of acquiring information. As our model shows, however, what matters is not only these exogenous costs, K_M , but also the opportunity cost of choosing a governance structure that provides incentives to invest in information (namely, the foregone opportunity for cost reduction). To analyze these opportunity costs, consider the expression for λ^* when $K_E \neq K_M$:

$$\lambda^* = \frac{\sigma_v^2 + \Delta^2 - 2(\bar{c} - \underline{c})(\mu_x \Delta + K_M - K_E)}{\frac{\sigma_v^2}{2} \frac{\sigma_v / \sigma_x}{\bar{c} - \underline{c}} + 2\Delta^2}.$$

Note the presence of production parameters, such as Δ and K_E , which have nothing *per se* to do with market clearing or price formation. More importantly, note that comparative statics regarding the informativeness of the price mechanism, such as $\partial\lambda^*/\partial K_M$, can depend on production parameters such as Δ .

In addition to comparative statics that illustrate the potential effects of production parameters on rational-expectations equilibrium, we can also say something about how the production environment affects markets. For example, in GHP we showed that (as in Grossman and Stiglitz, 1980) market thickness depends on λ^* , with concomitant implications for economic efficiency and welfare. In this paper's setting, therefore, market thickness depends on production parameters such as Δ and K_E .

6 Markets and Hierarchies Revisited

While our main focus is on the interaction between the choice of *organizational designs* by individual firms and the informativeness of the market's price mechanism, a straightforward

extension of our model also sheds light on the interaction between the choice of individual *firms' boundaries* and the informativeness of the price mechanism. Like our analysis of organizational designs, this section shows that omitting the price mechanism from the analysis of firms' boundaries can be problematic. In particular, we find that incentives to make specific investments (which now drive firms' boundary decisions) affect the informativeness of the price mechanism and vice versa.

To extend and reinterpret our model, consider a vertical production process with three stages (1, 2, and 3) and a different asset used at each stage (A_1 , A_2 , and A_3). There are again two parties, now denoted upstream (formerly E) and downstream (formerly M). The conditions of production are such that it is optimal for the upstream party (U) to own A_1 and for the downstream party (D) to own A_3 , so there are only two governance structures of interest (namely, U owns A_2 or D owns it). Thus, the asset A_2 is analogous to the machine from our original model, but we now focus on asset ownership as determining the boundary of the firm, rather than machine control as determining organizational designs. Because upstream necessarily owns A_1 and downstream A_3 , we interpret U ownership of A_2 as forward vertical integration and D ownership as backward. Beyond this reinterpretation of governance structures in terms of firms' boundaries, all the formal aspects of the model are unchanged. Under this reinterpretation, analogs of Propositions 1 through 3 continue to hold.⁸ In particular, our characterizations of the rational-expectations equilibrium and the industry equilibrium continue to hold, as do the comparative-statics results.

We see this section's discussion as directly related to some of the classic contributions to organizational economics. For example, Coase (1937: 359) argued that "it is surely important to enquire why co-ordination is the work of the *price mechanism* in one case and of the entrepreneur in the other" (emphasis added). Similarly, Williamson's (1975) title famously emphasized "Markets" as the alternative to hierarchy. For the next quarter century, however, the literature on firms' boundaries focused on the transaction as the unit

⁸For formal statements and proofs, see an earlier working-paper version: Gibbons et al. (2010), available at <http://www.nber.org/papers/w15779.pdf>.

of analysis. In short, non-integration replaced the market in the theory of the firm.

As noted in our Introduction, beginning with Grossman and Helpman (2002), recent work has begun to bring back market interactions as a determinant of firms' integration decisions; see also Legros and Newman (2008, 2009) for more in this spirit. As we described, however, our model differs from these in our focus on the informativeness of the price mechanism.

Interestingly, our focus allows us to revisit a specific argument from *Markets and Hierarchies*, beyond the title. In the book's opening pages, Williamson summarizes Hayek's (1945) observations about information in market prices, but Williamson then argues that "prices often do not qualify as sufficient statistics and ... a substitution of internal organization (hierarchy) for market-mediated exchange often occurs on this account" (1975: 5).

To our knowledge, the extent to which market prices are sufficient statistics that can influence firms' integration decisions has not been considered since 1975. The extended model in this section allows us to analyze such influence, in two ways: at the transaction level (i.e., for a given pair of parties, U_i and D_i) and for the market as a whole.

To link our analysis to Williamson's argument that prices not being "sufficient statistics" might induce parties to abandon "market-mediated exchange," we need to be precise about these two concepts. A natural way to assess the extent to which prices are sufficient statistics in our model is the following: the equilibrium informativeness of the price system is the expected reduction in variance that is obtained by conditioning on price (see Definition 4 in the appendix). And in our model, "market-mediated exchange" also has a natural interpretation: it means relying on information about the value of the final good from the price mechanism, rather than acquiring it directly.

At the transaction level, inspecting the derivation of (3) shows that (holding all else equal) an increase in the informativeness of prices reduces the returns to choosing an integration structure that induces information acquisition. In this sense, Williamson's argument holds at the transaction level in our model.

Of course, in our model the informativeness of prices is endogenous, because every other

pair of parties will also be considering the returns to choosing different integration structures. As a result, it may or may not be true for the market as a whole that when prices are less informative, more firms are organized to induce information acquisition. In Proposition 5 (in the appendix), we show that whether or not Williamson’s argument holds for the market as a whole depends crucially on the source of the change in the informativeness of prices: a change in an exogenous variable may increase the informativeness of prices and yet also increase the returns at the transaction level to choosing an integration structure that induces information acquisition. Our model thus allows us not only to formalize Williamson’s argument at the transaction level, but also to assess its validity for the market as a whole.

7 Empirical Implications

Our model has two sets of empirical implications: across-industry and within-industry. First, there are of course the *across-industry* empirical counterparts to our model’s comparative-statics predictions. For example, holding other characteristics constant, industries with greater demand uncertainty (i.e. higher σ_v^2) should have a greater share of firms that are organized to induce information acquisition via marketing control (or, as section 6 highlights, via downstream integration). Similarly, industries that make use of intermediate inputs that are subject to larger supply shocks (i.e. higher σ_x^2) should also have more firms organized to induce information acquisition.

In our model, information that firms care about (and organize themselves to acquire) is commonly valued across firms in the industry. At the other extreme, if demand were completely idiosyncratic (i.e., the consumer valuation for the product that firm i produces is independent of that for firm j ’s product), then there would be no useful information for the price mechanism to reveal from one firm to another, and our mechanism would not yield any interactions in governance structures across firms. The degree to which this uncertainty is common-value or idiosyncratic may depend on the level of product differentiation within

an industry. An undifferentiated-good industry is likely to be characterized by common values, and thus the informativeness of the price mechanism should be more important for the interactions between the governance structures of firms in the industry. In contrast, a differentiated-goods industry may be between the common-value and idiosyncratic extremes, but probably to a lesser degree.⁹

In order to carry out this type of analysis, one would need high-quality firm-level data that (a) spans industries, (b) contains information about firms' governance-structure decisions, and (c) has industry-level proxies for, say, common-value uncertainty. For example, Antras (2003) has (a) and (b), and Syverson (2004) has (a) and (c), if we interpret product substitutability as a proxy for common-value uncertainty.

Turning to what we refer to as within-industry analysis, where dependent variables are at the firm level and the analysis either focuses on a particular industry or contains industry controls, a common approach is to regress a measure of a dyad's governance structure on dyad- or transaction-level characteristics. Most of the recent empirical work on internal organization (for example, Bresnahan et al. (2002), Acemoglu et al. (2007), and Bloom, Sadun and van Reenen (2009)) and on firm boundaries (for example, Joskow (1985), Baker and Hubbard (2003), Forbes and Lederman (2009)) falls in this category.

Before proceeding, note that in our model, at the time dyads make governance-structure choices, there are no characteristics that vary at the dyad level. One goal of this homogeneity assumption was to highlight the idea that, even if firms are homogeneous ex ante, the organization of such firms could optimally be heterogeneous. In any real-world application, however, firms are likely to be heterogeneous at the time they decide upon their governance structure. This can be easily incorporated into our framework by allowing for heterogeneity in investment costs. That is, let $k_i = K_i^M - K_i^E$, with $k_i \sim U[\underline{k}, \bar{k}]$. There will then be a

⁹We have analyzed an elaboration of our model in which firm i cares about consumer valuation v_i , which is equal to a common-value component v with probability $\sqrt{\zeta}$ and an idiosyncratic component ν_i with probability $1 - \sqrt{\zeta}$. v and ν_i are uniformly distributed on \underline{v} and \bar{v} , but the ν_i values are independent across firms. At the time of production, firm i does not know whether $v_i = v$ or $v_i = \nu_i$. Under this specification, the ex ante correlation in consumer valuations is ζ and the informativeness of the price mechanism is increasing in ζ .

cutoff value k^* such that firms with $k_i < k^*$ will choose M -control and those with $k_i > k^*$ will choose E -control, with λ^* in (4) equal to $(k^* - \underline{k}) / (\bar{k} - \underline{k})$. In this model, only the marginal firm is indifferent.

In our model, one dyad’s governance structure also depends on the governance-structure choices of others. As a result, a regression of one dyad’s governance structure on its dyad-level characteristics will be biased. To see this, let $g_i = 1$ if dyad i is marketing-controlled (if the analysis is of internal organization) or if dyad i is downstream-integrated (if the analysis is of firm boundaries), and let $g_i = 0$ otherwise. Let X_i be the dyad-level characteristics that are usually included in governance structure regressions (such as the level of appropriable quasi-rents, transaction complexity, etc.), and suppose the industry we are analyzing has n dyads. Define ρ to be the correlation between dyad-level characteristics across dyads (i.e. $\rho = \text{Corr}(X_i, X_j)$) and assume that $\text{Var}(X_i)$ is common across dyads. Finally, denote by $\bar{g}_{-i} = \frac{1}{n-1} \sum_{j \neq i} g_j$ the industry-average governance structure not including dyad i . If the regression of g_i on X_i does not also include a measure of the governance structures of other dyads on the right-hand side, estimates of the coefficient on X_i will be biased.

In particular, using the terminology of Angrist and Pischke (2009), if we call $g_i = bX_i + \tau_i$ the “short regression” and $g_i = \beta X_i + \gamma \bar{g}_{-i} + \varepsilon_i$ the “long regression”, then the bias is given by the following proposition (which is proved in the appendix).

Proposition 4 *Suppose we estimate the short regression when the true model is given by the long regression. Then the bias of the estimated coefficient is given by*

$$E \left[\hat{b} \mid X \right] - \beta = \underbrace{\frac{\gamma}{1-\gamma} \rho \beta}_{\text{omitted variables}} + \underbrace{\frac{\gamma}{1-\gamma} \frac{\gamma}{n-1+\gamma} (1-\rho) \beta}_{\text{reverse causality}},$$

where n is the number of firms in the industry. As $n \rightarrow \infty$, the bias approaches $\frac{\gamma}{1-\gamma} \rho \beta$.

As the proposition shows, the bias in the short regression is a combination of two biases: (1) an omitted-variable bias that results from failure to include \bar{g}_{-i} in the regression and (2)

a reverse-causality bias that results from the fact that, if governance structures interact, g_i also affects g_j for all $j \neq i$. The latter bias goes away if firms are atomistic (which here can be approximated by taking n to infinity), as in most models of industry equilibrium. Our model predicts that $\gamma < 0$, which implies that the omitted-variable bias biases estimates of β towards zero. Further, this bias is greater the greater is ρ . Different determinants of vertical integration identified in the literature (e.g., uncertainty, transaction frequency, appropriable quasi-rents, importance of ex ante investments) may differ in their correlation across firms, so estimates of their effects on vertical integration may be differentially biased toward zero. Alternatively, if γ were found to be positive (and is necessarily less than one), then such regressions would be biased away from zero.

Other models characterizing governance structures in industry equilibrium may have different predictions. For example, Grossman and Helpman's (2002) model exhibits strategic complementarities in outsourcing decisions, and thus their model would predict that $\gamma > 0$. The model of Legros and Newman (2009) can predict both $\gamma > 0$ and $\gamma < 0$, depending on aggregate demand and the distribution of firm productivity, but since the interaction in governance structures acts only through the equilibrium price level, if one were to control for the market price in their model, they would predict $\gamma = 0$.

This discussion suggests two potential avenues for future empirical work. First, in estimating the magnitude of the classical determinants of governance structures, it would be interesting to include industry-average governance structure to eliminate the omitted-variable bias described above. Secondly, it would be useful to estimate, in a variety of contexts, the causal impact of industry-average governance structures on individual governance structures. This would require instruments for (a subset of) the governance structures of other dyads within an industry to estimate the sign and magnitude of γ ; for recent work along these lines, see Forbes and Lederman (2010).

Finally, it is interesting to note that in every empirical study mentioned above, there is significant variation in the governance-structure variable at the industry level. That is, it

is almost always the case that within an industry, there are some firms that are organized one way and other firms that are organized another. While this could potentially be due to measurement error in industry classification (i.e., it could be an aggregation problem), we take the view that this is an empirical fact to be explained. One potential explanation for this, of course, is that firms differ in their *ex ante* characteristics, and thus of course some firms organize one way and others organize differently. Another view, and one that is consistent with our model, is that industry-equilibrium effects provide forces toward heterogeneity in governance structure. This is true in models that generate a negative γ but not in models that generate a positive γ . Disentangling whether heterogeneity in governance structure is due to equilibrium effects or underlying heterogeneity in firm characteristics is an interesting empirical question, and one that could be informed by estimates of γ .

8 Conclusion

We view firms and the market not only as alternative ways of organizing economic activity, but also as institutions that interact and shape each other. In particular, by combining features of the incomplete-contract theory of firms' organizational designs and boundaries, together with the rational-expectations theory of the price mechanism, we have developed a model that incorporates two, reciprocal considerations. First, firms operate in the context of the market (specifically, the informativeness of the price mechanism affects parties' optimal governance structures). And second, the buyers in the market for an intermediate good are firms (specifically, parties' governance structures affect how they behave in this market and hence the informativeness of the price mechanism).

In the primary interpretation of our model in terms of organizational design we provide a formal explanation for why similar (possibly *ex ante* identical) firms choose different structures and strategies. Our analysis also demonstrates that viewing an individual firm, or transaction, as the unit of analysis can be misleading. Because of the interaction between

firm-level governance choices and the industry-wide informativeness of the price mechanism, equilibrium governance choices are shaped by industry-wide factors.

We also showed that our model can be reinterpreted to address firms' boundaries. Again, considering the endogenous informativeness of prices implies that both property-rights theory and transaction-cost economics abstract from potentially important issues by focusing on the transaction as the unit of analysis.

To develop and analyze our model, we imposed several strong assumptions that might be relaxed in future work. For example, to eliminate a market for machines, we assumed that machines are dyad-specific. Also, we have ignored the possibility of strategic information transmission before or during the price-formation process. We hope to explore these and other possibilities in future work.

References

- Acemoglu, Daron, Philippe Aghion, Claire Lelarge, John Van Reenan, and Fabrizio Zilibotti, “Technology, Information and the Decentralization of the Firm,” *Quarterly Journal of Economics*, 122 (2007), 1759–1799.
- Aghion, Philippe and Jean Tirole, “Formal and Real Authority in Organizations,” *Journal of Political Economy*, 105 (1997), 1–29.
- Alonso, Ricardo, Wouter Dessein, and Niko Matouschek, “When does Coordination require Centralization?,” *American Economic Review*, 98 (2008), 145–179.
- Angrist, Joshua D. and Jorn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricists Companion* (Princeton University Press, 2009).
- Antras, Pol, “Firms, Contracts and Trade Structure,” *Quarterly Journal of Economics*, 118 (2003), 1375–1418.
- Avenel, Eric, “Strategic Vertical Integration Without Foreclosure,” *Journal of Industrial Economics*, 56 (2008), 247–262.
- Baker, George P. and Thomas N. Hubbard, “Make versus Buy in Trucking: Asset Ownership, Job Design, and Information,” *American Economic Review*, 93 (2003), 551–572.
- Bloom, Nicholas, Raffaella Sadun, and John Van Reenen, “The Organization of Firms across Countries,” *NBER working paper no.15129*, (2009).
- Bolton, Patrick and Mathias Dewatripont, Authority in Organizations“ (2012).”
- Bresnahan, Timothy, Erik Brynjolfsson, and Lorin M Hitt, “Information Technology, Workplace Organization and the Demand for Skilled Labor: Firm-level Evidence,” *Quarter*, 117 (2002), 339–376.

- Christensen, Clayton M., *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail* (Harvard Business Press, 1997).
- Forbes, Silke J. and Mara Lederman, "Adaptation and Vertical Integration in the Airline Industry," *American Economic Review*, 99 (2009), 1831–1849.
- and —, "The Performance Implications of Vertical Integration: Evidence from the Airline Industry," *RAND Journal of Economics*, 41 (2010), 765–790.
- Gibbons, Robert, Niko Matouschek, and John Roberts, Decisions in Organizations“ (2012).”
- Gibbons, Robert S., Richard T. Holden, and Michael L. Powell, "Integration and Information: Markets and Hierarchies Revisited," *NBER working paper no.15779*, (2010).
- Grossman, Gene M. and Elhanan Helpman, "Integration Versus Outsourcing in Industry Equilibrium," *Quarterly Journal of Economics*, 117 (2002), 85–120.
- Grossman, Sanford J., "An Introduction to the Theory of Rational Expectations Under Asymmetric Information," *Review of Economic Studies*, 48 (1981), 541–559.
- and Joseph E. Stiglitz, "Information and Competitive Price Systems," *American Economic Review*, 66 (1976), 246–253.
- and —, "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70 (1980), 393–408.
- and Oliver D. Hart, "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94 (1986), 691–719.
- Hart, Oliver and John Moore, "Property Right and the Nature of the Firm," *Journal of Political Economy*, 98 (1990), 1119–1158.
- Holmstrom, Bengt and Jean Tirole, "Transfer Pricing and Organizational Form," *Journal of Law, Economics and Organization*, 7 (1991), 201–228.

- Joskow, Paul L., “Vertical Integration and Long-term Contracts: The Case of Coal-burning Electric Generating Plants,” *Journal of Law Economics and Organization*, 1 (1985), 33–80.
- Klein, Benjamin, Robert G. Crawford, and Armen A. Alchian, “Vertical Integration, Appropriable Rents, and the Competitive Contracting Process,” *Journal of Law and Economics*, 21 (1978), 297–326.
- Legros, Patrick and Andrew F. Newman, “Competing for Ownership,” *Journal of the European Economic Association*, 6 (2008), 1279–1308.
- and —, “A Price Theory of Vertical and Lateral Integration,” *Boston University working paper*, (2009).
- McLaren, John, “"Globalization" and Vertical Structure,” *American Economic Review*, 90 (2000), 1239–1254.
- Milgrom, Paul and John Roberts, “An Economic Approach to Influence Activities in Organizations,” *American Journal of Sociology*, 94 (1988), S154–S179.
- Ordover, Janusz A., Garth Saloner, and Stephen C. Salop, “Equilibrium Vertical Foreclosure,” *American Economic Review*, 80 (1990), 127–142.
- Rantakari, Heikki, “Governing Adapatation,” *Review of Economic Studies*, 75 (2008), 1257–1285.
- Syverson, Chad, “Product Substitutability and Productivity Dispersion,” *Review of Economics and Statistics*, 86 (2004), 534–550.
- Williamson, Oliver, *Markets and Hierarchies: Analysis and Antitrust Implications* (New York: Free Press, 1975).

9 Figures

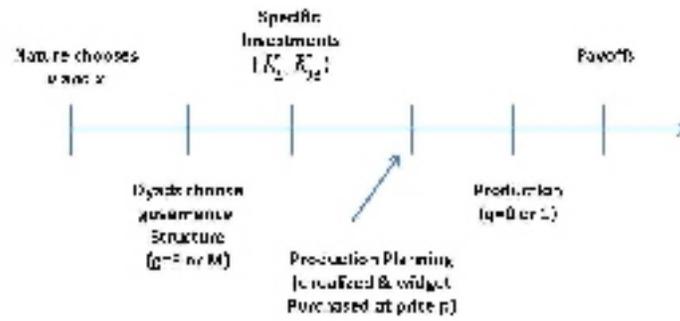


Figure 1: Timeline

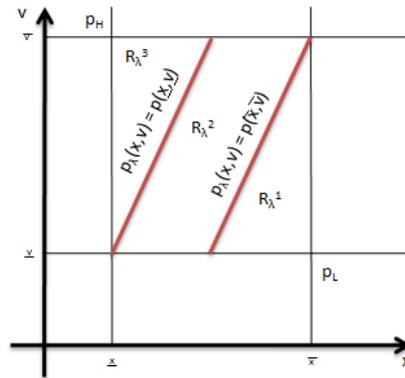


Figure 2: Regions of Piecewise-Linear Pricing Function

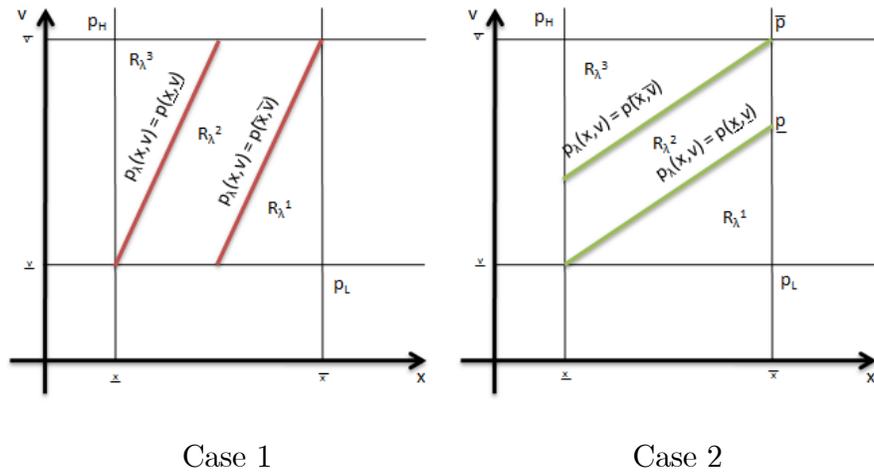


Figure 3

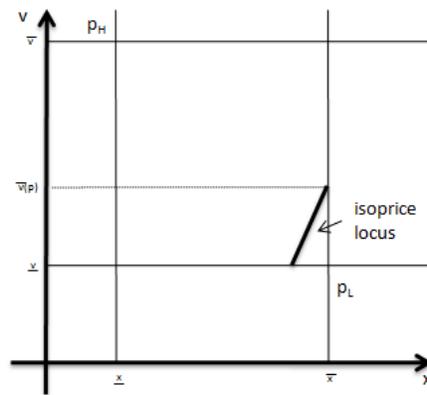


Figure 4