MECS 570-1: Organizational Economics I

Michael Powell

Boundaries of the Firm, Part 4: Influence Costs; Organizational IO

Influence Costs (Updated Feb 7, 2017)

At the end of the discussion of the Transaction-Cost Economics approach to firm boundaries, I mentioned that there are two types of costs that can arise when unprogrammed adaptation is required: costs associated with inefficient ex post decision making (adaptation costs) and costs associated with rent-seeking behavior (haggling costs). The TCE view is that when these costs are high for a particular transaction between two independent firms, it may make sense to take the transaction in-house and vertically integrate. I then described a model of adaptation costs in which this comparative static arises. I will now describe Powell's (JLEO, 2015) model of rent-seeking behavior in which similar comparative statics arise.

This model brings together the TCE view of haggling costs between firms as the central costs of market exchange and the Milgrom and Roberts (AJS, 1988) view that influence costs—costs associated with activities aimed at persuading decision makers—represent the central costs of internal organization. Powell asserts that the types of decisions that managers in separate firms argue about typically have analogues to the types of decisions that managers in different divisions within the same firm argue about (e.g., prices versus transfer prices, trade credit versus capital allocation) and that there is no reason to think a priori that the ways in which they argue with each other differ across different governance structures. They may in fact argue in different ways, but this difference should be derived, not assumed.

The argument that this model puts forth is the following. Decisions are ex post non-contractible, so whoever has control will *exercise* control (this is in contrast to the Property Rights Theory in which ex post decisions arise as the outcome of ex post efficient bargaining). As a result, the party who does not have control will have the incentives to try to influence the decision(s) of the party with control.

Control can be allocated via asset ownership, and therefore you can take away someone's right to make a decision. However, there are **position-specific private benefits**, so you cannot take away the fact that they care about that decision. In principle, the firm could replace them with someone else, but that person would also care about that decision. Further, while you can take away the rights to make a decision, you cannot take away the ability of individuals to try to influence whoever has decision rights, at least not unless you are willing to incur additional costs. As a result, giving control to one party reduces that party's incentives to engage in influence activities, but it intensifies the now-disempowered party's incentives to do so.

As in the Property Rights Theory, decision-making power affects parties' incentives. Here, it affects their incentives to try to influence the other party. This decision-making power is therefore a scarce resource that should be allocated efficiently. In contrast to the Property Rights Theory, decisions are ex post non-contractible. Consequently, whoever has control will exercise their control and will make different decisions ex post. So allocating control also affects the quality of ex post decision making. There may be a tension between allocating control to improve ex post decision making and allocating control to reduce parties' incentives to engage in influence activities.

Yet control-rights allocations are not the only instrument firms have for curtailing influence activities—firms can also put in place rigid organizational practices that reduce parties' incentives to engage in influence activities, but these practices may have costs of their own. Powell's model considers the interaction between these two substitute instruments for curtailing influence activities, and he shows that unified control and rigid organizational practices may complement each other.

Description Two managers, L and R, are engaged in a business relationship, and two decisions, d_1 and d_2 have to be made. Managers' payoffs for a particular decision depends on an underlying state of the world, $s \in S$. s is unobserved; however, L and R can potentially commonly observe an informative but manipulable signal σ. Managers bargain ex ante over a **control structure**, $c \in \mathcal{C} = \{I_L, I_R, NI, RNI\}$ and an **organizational practice**, $p \in \mathcal{P} = \{O, C\}$. Under I_i, manager i controls both decisions; under NI, L controls d_1, and R controls d_2; and conversely under RNI. Under an **open-door organizational practice**, $p = O$, the signal σ is commonly observed by L and R, and under a **closed-door organizational practice**, $p = C$, it is not. A bundle $g = (c, p) \in \mathcal{G} \equiv \mathcal{C} \times \mathcal{P}$ is a **governance structure**. Assume that in the ex ante bargaining process, L makes an offer to R, which consists of a proposed governance structure g and a transfer $w \in \mathbb{R}$ to be paid to R. R can accept the offer or reject it in favor of outside option yielding utility 0.

Given a governance structure, each manager chooses a level of **influence activities**, λ_i, at private cost $k(\lambda)$, which is convex, symmetric around zero, and satisfies $k(0) = k'(0) = 0$. Influence activities are chosen prior to the observation of the public signal without any private knowledge of the state of the world, and they affect the conditional distribution of σ_p given the state of the world s. The managers cannot bargain over a signal-contingent decision rule ex ante, and they cannot bargain ex post over the decisions to be taken or over the allocation of control.

Timing The timing of the model is as follows:

1. L makes an offer of a governance structure $g \in \mathcal{G}$ and a transfer $w \in \mathbb{R}$ to R. g and w are publicly observed. R chooses whether to accept $(d = 1)$ or reject $(d = 0)$ this offer in favor of outside option yielding utility 0. $d \in D = \{0, 1\}$ is commonly observed.

2. L and R simultaneously choose influence activities $\lambda_L, \lambda_R \in \mathbb{R}$ at cost $k(\lambda)$; λ_i is privately observed by i.

3. L and R publicly observe signal σ_p.

4. Whoever controls decision ℓ chooses $d_\ell \in \mathbb{R}$.

5. Payoffs are realized.

Functional-Form Assumptions The signal under $p = O$ is linear in the state of the world, the influence activities, and noise: $\sigma_O = s + \lambda_L + \lambda_R + \varepsilon$. All random variables are independent and normally distributed with mean zero: $s \sim N(0, h^{-1})$ and $\varepsilon \sim N(0, h_\varepsilon^{-1})$. The signal under $p = C$ is uninformative, or $\sigma_C = \emptyset$. For the purposes of Bayesian updating, the signal-to-noise ratio of the signal is $\varphi_p = h_\varepsilon / (h + h_\varepsilon)$ under $p = O$ and, abusing notation, can be thought of as $\varphi_p = 0$ under $p = C$. Influence costs are quadratic, $k(\lambda_i) = \lambda_i^2/2$, and each manager's payoffs gross of influence costs are

$$U_i(s, d) = \sum_{\ell=1}^{2} \left[-\frac{\alpha_i}{2} (d_\ell - s - \beta_i)^2 \right], \alpha_i > 0, \beta_i \in \mathbb{R}.$$

Both managers prefer each decision to be tailored to the state of the world, but given the state of the world, manager i prefers that $d_1 = d_2 = s + \beta_i$, so there is disagreement between the two managers. Define $\Delta \equiv \beta_L - \beta_R > 0$ to be the **level of disagreement,** and assume that $\alpha_L \geq \alpha_R$: manager L cares more about the losses from not having her ideal decision implemented. Further, assume that **managers operate at similar scales**: $\alpha_R \leq \alpha_L \leq \sqrt{3}\alpha_R$.

Although there are four possible control-rights allocations, only two will ever be optimal: unifying control with manager L or dividing control by giving decision 1 to L and decision 2 to R. Refer to unified control as **integration,** and denote it by $c = I$, and refer to divided control as **non-integration**, and denote it by $c = NI$. Consequently, there are effectively four governance structures to consider:

$$\mathcal{G} = \{(I, O), (I, C), (NI, O), (NI, C)\}.$$

Solution Concept A governance structure $g = (c, p)$ induces an extensive-form game between L and R, denoted by $\Gamma(g)$. A **Perfect-Bayesian Equilibrium** of $\Gamma(g)$ is a belief profile μ^*, an offer $(g^*, \theta^*), w^*$ of a governance structure and a transfer, a pair of influence-activity strategies $\lambda_L^* : \mathcal{G} \times \mathbb{R} \times D \times \Delta(s) \to \mathbb{R}$ and $\lambda_R^* : \mathcal{G} \times \mathbb{R} \times D \times \Delta(s) \to \mathbb{R}$, and a pair of decision rules $d_\ell^* : \mathcal{G} \times \mathbb{R} \times D \times \mathbb{R} \times \Sigma \times \Delta(s) \to \mathbb{R}$ such that the influence-activity strategies and the decision rules are sequentially optimal for each player given his/her beliefs, and μ^* is derived from the equilibrium strategy using Bayes's rule whenever possible.

This model is a signal-jamming game, like the career concerns model earlier in the class. Further, the assumptions we have made will ensure that players want to choose relatively simple strategies. That is, they will choose public influence-activity strategies $\lambda_L^* : \Delta(s) \to \mathbb{R}$ and $\lambda_R^* : \Delta(s) \to \mathbb{R}$ and decision rules $d_\ell^* : \mathcal{G} \times \Sigma \times \mathbb{R} \times \Delta(s) \to \mathbb{R}$.

The Program Take a governance structure g as given. Suppose manager i has control of decision ℓ under governance structure g. Let $\lambda^{g*} = (\lambda_L^{g*}, \lambda_R^{g*})$ denote the equilibrium level of influence activities. Manager i chooses d_ℓ to minimize her expected loss given her beliefs about the vector of influence activities, which I denote by $\hat{\lambda}(i)$. She therefore chooses d_ℓ^* to solve

$$\max_{d_\ell} E_s \left[-\frac{\alpha_i}{2} (d_\ell - s - \beta_i)^2 \middle| \sigma_p, \hat{\lambda}(i) \right].$$

She will therefore choose

$$d_\ell^{g*}\left(\sigma_p; \hat{\lambda}(i)\right) = E_s \left[s \middle| \sigma_p, \hat{\lambda}(i) \right] + \beta_i.$$

The decision that manager i makes differs from the decision manager $j \neq i$ would make if she had the decision right for two reasons. First, $\beta_i \neq \beta_j$, so for a given set of beliefs, they prefer different decisions. Second, out of equilibrium, they may differ in their beliefs about λ. Manager i knows λ_i but only has a conjecture about λ_j. These differences in out-of-equilibrium beliefs are precisely the channel through which managers might hope to change decisions through influence activities.

Since random variables are normally distributed, we can make use of the normal updating formula to obtain an expression for $E_s\left[s|\,\sigma_P,\hat{\lambda}\left(i\right)\right]$. In particular, it will be a convex combination of the prior mean, 0, and the modified signal $\hat{s}\left(i\right)=\sigma_p-\hat{\lambda}_L\left(i\right)-\hat{\lambda}_R\left(i\right)$, which of course must satisfy $\hat{\lambda}_i\left(i\right)=\lambda_i$. The weight that i's preferred decision strategy attaches to the signal is given by the φ_p, so

$$d_\ell^{g*}\left(\sigma_p;\hat{\lambda}\left(i\right)\right)=\varphi_p\cdot\hat{s}\left(i\right)+\beta_i.$$

Given decision rules $d_\ell^{g*}\left(\sigma_p;\lambda^{g*}\right)$, we can now set up the program that the managers solve when deciding on the level of influence activities to engage in. Manager j chooses λ_j to solve

$$\max_{\lambda_j} E_{s,\varepsilon}\left[\sum_{\ell=1}^{2}-\frac{\alpha_j}{2}\left(d_\ell^{g*}\left(\sigma_p;\lambda^{g*}\right)-s-\beta_j\right)^2\right]-k\left(\lambda_j\right).$$

Taking first-order conditions, we get:

$$\begin{aligned}\left|k'\left(\lambda_j^{g*}\right)\right| &= \left|E_{s,\varepsilon}\left[\sum_{\ell=1}^{2}-\alpha_j\underbrace{\left(d_\ell^{g*}\left(\sigma_p;\lambda^{g*}\right)-s-\beta_j\right)}_{=0\text{ if } j \text{ controls } g;\ =\Delta\text{ otherwise}}\underbrace{\frac{\partial d_\ell^{g*}}{\partial\sigma}}_{\varphi_p}\underbrace{\frac{\partial\sigma}{\partial\lambda_j}}_{=1}\right]\right| \\ &= N_{-j}^c\alpha_j\Delta\varphi_p,\end{aligned}$$

where N_{-j}^c is the number of decisions manager j does not control under control structure c.

Finally, at $t=1$, L will make an offer g,w to

$$\max_{g,w} E_{s,\varepsilon}\left[\sum_{\ell=1}^{2}-\frac{\alpha_L}{2}\left(d_\ell^{g*}\left(\sigma_p;\lambda^{g*}\right)-s-\beta_L\right)^2\right]-k\left(\lambda_L^{g*}\right)-w$$

subject to R's participation constraint:

$$E_{s,\varepsilon}\left[\sum_{\ell=1}^{2}-\frac{\alpha_R}{2}\left(d_\ell^{g*}\left(\sigma_p;\lambda^{g*}\right)-s-\beta_R\right)^2\right]-k\left(\lambda_R^{g*}\right)+w\geq 0.$$

w will be chosen so that the participation constraint holds with equality, so that L's problem becomes:

$$\underbrace{\max_{g} E_{s,\varepsilon}\left[\sum_{i\in\{L,R\}}\sum_{\ell=1}^{2}-\frac{\alpha_i}{2}\left(d_\ell^{g*}\left(\sigma_p;\lambda^{g*}\right)-s-\beta_i\right)^2\right]-\sum_{i\in\{L,R\}}k\left(\lambda_i^{g*}\right).}_{W(g)}$$

The **Coasian Program** is then

$$\max_{g\in\mathcal{G}} W\left(g\right).$$

Solution Managers' payoffs are quadratic. The first term can therefore be written as the sum of the mean-squared errors of d_1^{g*} and d_2^{g*} as estimators of the **ex post surplus-maximizing decision**, which is

$$s+\frac{\alpha_L}{\alpha_L+\alpha_R}\beta_L+\frac{\alpha_R}{\alpha_L+\alpha_R}\beta_R$$

for each decision. As a result, the first term can be written as the sum of a bias term and a variance term (recall that for two random variables X and Y, $Var\left(Y\right)=E\left[Var\left(Y|\,X\right)\right]+Var\left(E\left[Y|\,X\right]\right)$):

$$W\left(c,p\right)=-\left(ADAP\left(p\right)+ALIGN\left(c\right)+INFL\left(c,p\right)\right),$$

where after several lines of algebra, the expressions for these terms are:

$$\begin{aligned}
ADAP\left(p\right) &= \left(\alpha_L+\alpha_R\right)\frac{1-\varphi_p}{h}\\
ALIGN\left(c\right) &= \frac{\alpha_R}{2}\Delta^2+\frac{\alpha_L}{2}\Delta^2\mathbf{1}_{c=NI}+\frac{\alpha_R}{2}\Delta^2\mathbf{1}_{c=I}\\
INFL\left(c,p\right) &= \left(\frac{1}{2}\left(\alpha_R\Delta\varphi_p\right)^2+\frac{1}{2}\left(\alpha_L\Delta\varphi_p\right)^2\right)\mathbf{1}_{c=NI}+\frac{1}{2}\left(2\alpha_R\Delta\varphi_p\right)^2\mathbf{1}_{c=I}.
\end{aligned}$$

$ADAP\left(p\right)$ represents the costs associated with basing decision making on a noisy signal. $ADAP\left(p\right)$ is higher for $p=C$, because under $p=C$, even the noisy signal is unavailable.

$ALIGN\,(c)$ represents the costs associated with the fact that ex post, decisions will always be made in a way that are not ideal for someone. Whether they are ideal for manager L or R depends on the control structure c. Finally, $INFL\,(c,p)$ are the influence costs, $k\,(\lambda_L^{g*}) + k\,(\lambda_R^{g*})$. When $p = C$, these costs will be 0, since there is no signal to manipulate. When $p = O$, these costs will depend on the control structure.

There will be two trade-offs of interest.

Influence-cost–alignment-cost trade-off First, let us ignore $ADAP\,(p)$ and look separately at $ALIGN\,(c)$ and $INFL\,(c,p)$. To do so, let us begin with $INFL\,(c,p)$. When $p = C$, these costs are clearly 0. When $p = O$, they are:

$$\begin{aligned} INFL\,(I,O) &= \frac{1}{2}\,(2\alpha_R\Delta\varphi_O)^2 \\ INFL\,(NI,O) &= \frac{1}{2}\,(\alpha_L\Delta\varphi_O)^2 + \frac{1}{2}\,(\alpha_R\Delta\varphi_O)^2\,. \end{aligned}$$

Divided control minimizes influence costs, as long as managers operate at similar scale:

$$INFL\,(I,O) - INFL\,(NI,O) = \frac{1}{2}\left(3\,(\alpha_R)^2 - (\alpha_L)^2\right)(\Delta\varphi_O)^2 > 0.$$

Next, let us look at $ALIGN\,(c)$. When $c = I$, manager L gets her ideal decisions on average, but manager R does not:

$$ALIGN\,(I) = \alpha_R\Delta^2.$$

When $g = NI$, each manager gets her ideal decision correct on average for one decision but not for the other decision:

$$ALIGN\,(NI) = \frac{\alpha_L + \alpha_R}{2}\Delta^2.$$

When $\alpha_L = \alpha_R$, so that $ALIGN\,(I) = ALIGN\,(NI)$, we have that $INFL\,(I,O) - INFL\,(NI,O) > 0$, so that influence costs are minimized under non-integration. When

$p = C$, so that there are no influence costs, and $\alpha_L > \alpha_R$, $ALIGN\,(I) < ALIGN\,(NI)$, so that alignment costs are minimized under integration. Unified control reduces ex post alignment costs and divided control reduces influence costs, and there is a trade-off between the two.

Influence-cost–adaptation-cost trade-off Next, let us ignore $ALIGN\,(c)$ and look separately at $ADAP\,(p)$ and $INFL\,(c, p)$. Since

$$ADAP\,(p) = (\alpha_L + \alpha_R)\,\frac{1 - \varphi_p}{h},$$

adaptation costs are higher under closed-door practices, $p = C$. But when $p = C$, influence costs are reduced to 0. Closed-door practices therefore eliminate influence costs but reduce the quality of decision making, so there is a trade-off here as well. Finally, it is worth noting that when $p = C$, influence activities are eliminated, so the parties might as well unify control, since doing so reduces ex post alignment costs. That is, closed-door policies and integration are complementary.

The following figure illustrates optimal governance structures for different model parameters. The figure has three boundaries, each of which correspond to different results from the literature on influence activities and organizational design. The vertical boundary between (I, O) and (NI, O) is the "Meyer, Milgrom, and Roberts boundary": a firm rife with politics should perhaps disintegrate.

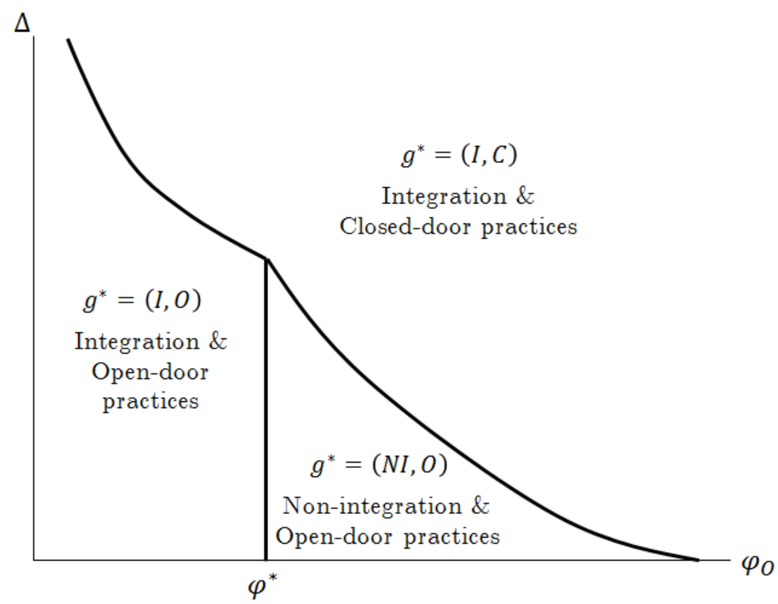

The diagonal boundary between (I, O) and (I, C) is the "Milgrom and Roberts boundary": rigid decision-making rules should sometimes be adopted within firms. These two boundaries highlight the idea that non-integration and rigid organizational practices are substitute instruments for curtailing influence costs: sometimes a firm prefers to curtail influence activities with the former and sometimes with the latter. Finally, the boundary between (NI, O) and (I, C) is the "Williamson boundary." If interactions across firm boundaries, which are characterized by divided control and open lines of communication, invite high levels of influence activities, then it may be optimal instead to unify control *and* adopt closed-door practices.

At the end of the day, any theory of the firm has to contend with two polar questions. First, why are all transactions not carried out in the market? Second, why are all transactions not carried out within a single large firm? TCE identifies "haggling costs" as an answer to the first question and "bureaucratic costs of hierarchy" as an answer to the second. Taking a parallel approach focused on the costs of internal organization, Milgrom and Roberts identify "influence costs" as an answer to the second question and "bargaining costs" between firms as an answer to the first. The model presented above blurs the distinction between TCE's "haggling costs" and Milgrom and Roberts's "influence costs" by arguing that the types of decisions over which parties disagree across firm boundaries typically have within-firm analogs, and the methods parties employ to influence decision makers within firms are not exogenously different than the methods they employ between firms.

This perspective implies, however, that unifying control *increases* influence costs, in direct contrast to Williamson's claim that "fiat [under integration] is frequently a more efficient way to settle minor conflicts": modifying firm boundaries without adjusting practices does not solve the problem of haggling. However, adopting rigid organizational practices in addition to unifying control provides a solution. Fiat (unified control) appears effective at eliminating haggling, precisely because it is coupled with bureaucracy. This influence-cost approach to firm boundaries therefore suggests that bureaucracy is not a cost of integration. Rather, it is an endogenous response to the actual cost of integration, which is high levels of influence activities.

Finally, we can connect the implications of this model to the empirical implications of the TCE approach. As with most theories of the firm, directly testing the model's underlying causal mechanisms is inherently difficult, because many of the model's dependent variables, such as the levels of influence activities and the optimality of ex post decision making, are unlikely to be observed by an econometrician. As a result, the model focuses on predictions regarding how potentially observable independent variables, such as environmental uncertainty and the level of ex post disagreement, relate to optimal choices of potentially observable dependent variables, such as the integration decision or organizational practices.

In particular, the model suggests that if interactions across firm boundaries involve high levels of influence costs, it may be optimal to unify control and adopt closed-door practices. This may be the case when the level of ex post disagreement (Δ) is high and when the level of ex ante uncertainty (h) is low. The model therefore predicts a positive relationship between integration and measures of ex post disagreement and a negative relationship between integration and measures of ex ante uncertainty.

The former prediction is consistent with the TCE hypothesis and is consistent with the findings of many empirical papers, which we will soon discuss. The second prediction contrasts with the TCE hypothesis that greater environmental uncertainty leads to more contractual incompleteness and more scope for ex post haggling, and therefore makes integration a relatively more appealing option. This result is in line with the failure of empirical TCE papers to find consistent evidence in favor of TCE's prediction that integration and uncertainty are positively related.