Regulating the Stakeholders' Firm*

Liam Lods† David Martimort‡

March 7, 2026

Abstract

This paper studies the optimal regulation of a firm subject to internal stakeholder influence. We embed a complete-information delegated common agency game, where stakeholders influence a manager via side-transfers, within a regulatory screening problem. Because the regulator designs the baseline schedule that anchors the influence game, she anticipates and free-rides on this inevitable internal lobbying to reduce transfers paid to the firm's manager. When the regulator has redistributive preferences favouring hostile stakeholders, her strategic depression of the manager's fallback option in the lobbying game generates countervailing incentives, resulting in optimal bunching for the most efficient types, which is distortive in expected terms. Finally, when multiple stakeholder groups compete, the regulator's ability to manipulate their marginal contributions breaks the standard neutrality result of menu auctions, demonstrating that the internal distribution of stakes alters optimal regulation.

1 Introduction

The traditional approach to Contract Theory, the Principal-Agent paradigm, views the Agent as a black box, a monolith maximising a single objective function taken as given. In the seminal regulation works of Baron and Myerson (1982) and Laffont and Tirole (1993), a Regulator designs an incentive scheme for this kind of monolithic firm. The reality of the firm is, however, far more complex, as emphasised by the *Theory of the Firm* (Laffont and Martimort 1997). According to this theory, a firm is a nexus of contracts where various stakeholders — shareholders, unions, local communities, activists — vie for influence over the firm's decisions, made by a manager (Bénabou and Tirole 2010; Magill, Quinzii, and Rochet 2015).

*We thank Michel Le Breton, François Salanié, Bruno Jullien, Juuso Välimäki, Johannes Hörner, Thomas Mariotti and Alexis Ghersengorin for helpful discussion and comments. This paper also benefited from remarks by participants to the TSE Microeconomic Theory Workshop and the TSE PhD General Workshop. This paper has been funded by the Agence Nationale de la Recherche under grants ANR-17-EURE-0010 (Investissements d'Avenir program) and LOBBIES-ANR-24-CE26-7833-04.

†Toulouse School of Economics, liam.lods@tse-fr.eu

‡Toulouse School of Economics, david.martimort@tse-fr.eu

This internal tug-of-war is not only a corporate governance issue, as it affects how the firm responds to regulation and how profits are redistributed among its components (Fleurbaey and Ponthière 2023). Consider Baron's (2001) description of *Private Politics*: rather than attempting to capture public institutions as in classic political economy models (Laffont and Tirole 1991; Grossman and Helpman 1994), activists bypass the state to target firms directly. An influenced manager will distort his decision to suit the activists' preferences over outcomes, and this influence will transfer some utility between the manager and other stakeholders. Because these groups write side-contracts with management to sway production decisions, their internal lobbying is highly relevant to a regulator, particularly when the latter is concerned about redistribution.

This paper asks how a Regulator should design policy when she knows the firm is subject to internal pressure. Should the Regulator treat the firm as a monolith, she would ignore the fact that the manager's incentives are being distorted by stakeholders. Given that the Regulator knows the exact structure of the firm in our model, she is able to perfectly counteract the stakeholders' influence on the outcome, and finds it optimal to do so in the absence of redistributive concerns. In the presence of such concerns, however, the Regulator actively free-rides on this inevitable influence. This is in sharp contrast with Tirole's (1986) *Collusion-Proofness Principle*, where side-transfers are assumed to be inefficient.

Our approach to regulation is to consider a model of procurement where the Regulator (the Principal, *she*) designs a non-linear transfer schedule to motivate a Manager (the Agent, *he*). One or two stakeholder groups try to influence the Manager's decision via a menu of side-transfers — when there are multiple groups, they compete *via* a menu auction à la Bernheim and Whinston (1986). Crucially, we model this influence as a game of *delegated common agency*: stakeholders make take-it-or-leave-it offers to the Manager, but they do not possess veto power over production. Should the Manager reject their influence, he simply produces according to the Regulator's baseline schedule. This assumption captures the position of weak stakeholders — NGOs, local communities, or minority shareholders, for instance — who can pressure management but lack the institutional authority to unilaterally halt the firm's operations. Finally, the Regulator acts as a Stackelberg leader regarding institutional design: she sets the tax schedule anticipating the equilibrium of the influence game played within the firm, knowing that it is played under complete information.

Our first set of results characterises the optimal regulation when there is a single stakeholder group. A central finding is that the Regulator always free-rides on the stakeholder's influence, for two reasons: first, because stakeholders exert influence through side-transfers, the Regulator can adjust her tax schedule to extract part of the generated surplus; second, these side-transfers are intrinsically less costly to her than taxes. Consequently, the expected produced quantity is always distorted in the direction of the stakeholders' preferences, as it allows the Regulator to substitute costly direct transfers with stakeholder contributions, preserving the Manager's incentives.

When stakeholders want to reduce production, but are preferred to the Manager by the Regulator — that is, the Regulator would like for stakeholders to get as much surplus from the interaction as possible — the latter faces a severe trade-off. Indeed, to let stakeholders capture this surplus, the Manager's outside option in the bargaining process should be depressed. This is achieved by decreasing the quantity produced by more efficient types, a force that is directly at odds with the standard monotonicity requirement. This is most salient when the Manager's outside option corresponds to a message that is never sent in equilibrium, a case occurring when efficiency is high. This induces *bunching at the top* of the distribution. In sharp contrast to standard ironing techniques (Myerson 1981; Toikka 2011), this bunching is *distortive in expected terms* with respect to the second-best, as it does not resolve an irregularity of the technology, but arises from the Regulator's deliberate sabotage of the Manager's bargaining position.

We then extend the model to allow for two stakeholder groups to simultaneously influence the Manager. A standard intuition in such influence games is that the outcome should be neutral with respect to the distribution of stakes; that is, the final production decision depends only on the aggregate pressure exerted by all groups combined. In our framework, however, this neutrality result breaks down. As the Regulator sits above the influence game and has redistributive preferences, the specific distribution of payoffs within the firm matters to her. We show that the Regulator manipulates the marginal contribution of each individual group through the tax schedule. Thus, separating a single stakeholder group into two smaller groups with the same aggregate stake alters the optimal quantity schedule, particularly when their preferences diverge. This non-neutrality highlights how treating stakeholders as monolithic may distort regulation, even when the firm is treated as a nexus of contracts.

1.1 Literature Review

Our framework connects the mechanism design approach to regulation with the political economy of common agency, and demonstrates how the internal structure of a firm generates a novel rationale for inflexible rules.

First, we contribute to the literature on private politics and regulation. Following Baron (2001), a body of literature studies how activists target firms directly rather than through public institutions. Closely related to our work, Egorov and Harstad (2017) study the interaction between activist boycotts, firm self-regulation and formal public regulation. While they take a dynamic, game-theoretic approach where regulation is modelled as a binary policy choice effectively ending the private influence game, we formalise regulation as a mechanism design problem under asymmetric information. In our setting, the Regulator does not pre-empt or ban private politics, but designs a non-linear tax schedule that anticipates and sets the stage for stakeholders' influence, using this influence to reduce the direct public transfers paid to the Manager. In this model, side-transfers from stakeholders to the Manager, traditionally viewed as socially wasteful collusion in standard hierarchies (Tirole 1986), are leveraged by the Regulator to reduce her cost of incentive provision. The design of regulation as a mechanism echoes seminal models of regulation of monolithic firms, such as Baron and Myerson (1982) or Laffont and Tirole (1993).

Second, our modelling of internal firm politics relies on the common agency framework (Bernheim and Whinston 1986; Laussel and Le Breton 2001). A defining feature of this literature, most famously applied to political lobbying by Grossman and Helpman (1994), is the neutrality of the outcome with respect to the distribution of stakes among lobbyists. In standard models, the influenced party sits at the top of the hierarchy and maximises aggregate surplus, while in our paper, the Regulator sits above the influence game and holds specific redistributive preferences. Her ability to design a baseline schedule upon which the menu auction relies alters the stakeholders' marginal contributions, thus breaking the standard neutrality result. This highlights the relevance of the allocation of stakes within a firm, echoing questions raised by Martimort and Semenov (2008) and Martimort and Stole (2024) regarding the optimal structuring of influence groups.

Finally, our model provides a novel organisational rationale for inflexible rules and bunching in regulatory schedules. In standard mechanism design, bunching is typically a remedy for exogenous irregularities, such as a failure of the hazard rate condition (Myerson 1981; Toikka 2011). This classic ironing procedure is designed to be neutral in expectation, merely flattening the schedule to preserve monotonicity without altering the expected virtual surplus over the pooled interval. Alternatively, as explored by Lewis and Sappington (1989) and generalised by Jullien (2000), optimal pooling may arise from countervailing incentives due to exogenous, type-dependent outside options. While our solution relies on similar mathematical mechanics, the economic origin of the distortion is fundamentally different. Indeed, bunching in our model is endogenous and *Principal*-driven, arising because the Regulator's redistributive desire to empower hostile stakeholders requires her to depress the Manager's fallback position in the influence game. This strategic manipulation of the bargaining outcome acts as an endogenous type-dependent reservation utility, inducing an inflexible rule for the most efficient types that is strictly distortive in expected terms.

Section 2 presents the model. Section 3 deals with the standard benchmark where Stakeholders are inactive. Section 4 characterises the solution to a situation with a single, unified Stakeholder group, while Section 5 looks at a situation with two Stakeholder groups, used to underline the considerations specifically arising from fragmentation of influence, in particular non-neutrality. Section 6 describes some applications, and Section 7 concludes. Proofs are in the Appendix.

2 Model

We consider a model inspired by Baron and Myerson (1982) with a Regulator (the principal, *she*) and a Manager (the agent, *he*). We add a third player to this standard hierarchy: a Stakeholder group (sometimes referred to as the stakeholders, *they*).

Technology and Information. The Regulator enjoys a surplus $S(q)$ when the Manager produces $q \in [0, \overline{q}]$ units of a good, where $\overline{q}$ represents maximum production capacity. We assume the underlying surplus function $S : \mathbb{R}_+ \to \mathbb{R}$ is strictly increasing and concave ($S' > 0 > S''$), with $S(0) = 0$, $\lim_{q\to 0} S'(q) = +\infty$ and $\lim_{q\to +\infty} S'(q) = 0$. This surplus can be interpreted as the consumers' utility on a secondary market.[1] We assume that the capacity constraint $\overline{q}$ is sufficiently large that it never binds for on-path optimal allocations, but serves to bound the domain of off-equilibrium threat allocations.

1. More generally, this function can represent the utility of unorganised citizens whose interests are safeguarded by the Regulator.

Similarly, the Stakeholder group derives an aggregate utility sq from production, with $s \in \mathbb{R}$.

When producing q units, the Manager incurs a cost θq. The marginal cost parameter θ is drawn from a commonly known cumulative distribution function F, assumed to be absolutely continuous with associated density function $f > 0$ on $[\underline{\theta}, \overline{\theta}]$. The realisation of θ is privately observed by the Manager and the Stakeholders, but not by the Regulator.[2] Finally, we assume that $s < \underline{\theta}$, which ensures that the stakeholders can never profitably fund any production without the Regulator's baseline schedule.

We make the standard regularity assumption:

Assumption 1. *The distribution F satisfies the Monotone Hazard Rate Property, that is:*

$$\frac{F(\theta)}{f(\theta)} \text{ is non-decreasing in } \theta.$$

Contracts. The contractual environment features two layers of influence:

- *The Regulator* offers a non-linear transfer schedule $T : [0, \overline{q}] \to \mathbb{R}$ based on the observable produced quantity q. The Manager's outside option if he rejects the regulation is not to produce ($q = 0$), yielding zero utility.
- *The Stakeholder group* moves after θ is realised. Possessing all bargaining power vis-à-vis the Manager, the stakeholders make a take-it-or-leave-it offer in the form of a side-transfer schedule $t(\cdot, \theta) : [0, \overline{q}] \to \mathbb{R}$. Because the Manager and Stakeholders share complete information, this schedule is offered knowing both θ and the Regulator's baseline schedule T.

The Regulator and Stakeholders cannot interact directly; they can only affect each other's payoff through the Manager's choice of q. In particular, the Regulator cannot directly elicit the stakeholders' superior information regarding θ.

Payoffs. Given a realised cost θ and the schedules T and t, the Manager's payoff for a chosen quantity $q \in [0, \overline{q}]$ is:

$$U(q, \theta \mid T, t) = T(q) + t(q, \theta) - \theta q.$$

The Stakeholder's utility is:

$$V(q, \theta \mid t) = sq - t(q, \theta).$$

The Regulator maximises a weighted social welfare function, placing Pareto weights $\alpha, \mu \in [0, 1)$ on the Manager and the Stakeholder group, respectively:

$$W(q, \theta \mid T, t) = S(q) - T(q) + \alpha U(q, \theta \mid T, t) + \mu V(q, \theta \mid t).$$

2. This assumption models the idea that stakeholders may have superior information that is learned after the regulation is drafted.

Substituting the players' utilities into the Regulator's objective function yields:

$$W(q,\theta \mid T,t) = S(q) - (1-\alpha)T(q) + (\mu s - \alpha\theta)q - (\mu - \alpha)t(q,\theta).$$

This explicitly highlights the Regulator's posture towards side-transfers: whenever $\alpha \neq \mu$, the Regulator is not indifferent about the internal lobbying game. Furthermore, as $\mu < 1$, side-transfers are inherently cheaper than direct ones from the point of view of the Regulator. Ultimately, the Regulator chooses the schedule T to maximise her expected payoff

$$\mathscr{W}(T) = \mathbb{E}_\theta \left[W(q^*(\theta),\theta \mid T,t^*)\right],$$

anticipating the equilibrium side-transfers t^* and production choice q^*.

Timing. The game unfolds as follows:

1. The Regulator offers a baseline transfer schedule $T(q)$ to the Manager.
2. Nature draws the Manager's marginal cost θ, which is perfectly observed by the Manager and the Stakeholder group.
3. The Stakeholders offer a side-transfer schedule $t(q,\theta)$ to the Manager.
4. The Manager chooses whether to accept the baseline schedule as well as the side-transfers, and selects a production level q.
5. Production takes place and payments are made.

3 Standard Regulation Scenario

As a benchmark, consider a simple model à la Baron and Myerson (1982), in which the Stakeholder group is restricted to offering the null schedule $t \equiv 0$.[3] In this case, we can write the Manager's payoff as

$$U(q,\theta \mid T) = T(q) - \theta q.$$

His indirect utility for type θ facing the schedule T is given by

$$U(\theta \mid T) = \max_q \left\{T(q) - \theta q\right\}. \tag{3.1}$$

This means that the schedule T induces a quantity schedule $q : [\underline{\theta}, \overline{\theta}] \to [0, \overline{q}]$ resulting from the Manager's utility maximisation problem, characterised by

$$q(\theta \mid T) \in \arg\max_q \left\{T(q) - \theta q\right\}. \tag{3.2}$$

For notational brevity, we will hereafter write $U(\theta)$ and $q(\theta)$ when the underlying schedule T is unambiguous.

The stakeholders' utility then simplifies to

$$V(q(\theta)) = sq(\theta),$$

3. Equivalently, this can be interpreted as the stakeholders being inactive.

and the Regulator's expected payoff is

$$\mathscr{W}(T) = \int_{\underline{\theta}}^{\overline{\theta}} \left(S(q(\theta)) - T(q(\theta)) + \mu s q(\theta) + \alpha U(\theta)\right) dF(\theta),$$

or, substituting the Manager's indirect utility,

$$\mathscr{W}(T) = \int_{\underline{\theta}}^{\overline{\theta}} \left(S(q(\theta)) - (1-\alpha)T(q(\theta)) + (\mu s - \alpha\theta)q(\theta)\right) dF(\theta).$$

Full Information Setting. If θ is perfectly observed by the Regulator in this setting, she does not need to take expectations or leave information rents ($U^{fb}(\theta) \equiv 0$). The first-best quantity schedule $q^{fb}(\theta)$ is characterised by

$$S'(q^{fb}(\theta)) = \theta - \mu s, \quad \forall \theta \in [\underline{\theta}, \overline{\theta}]. \tag{3.3}$$

Note that the presence of stakeholders, even if they are inactive, amounts to a shift in the marginal cost of production of size $|\mu s|$.

Asymmetric Information Setting. If θ is the Manager's private information and stakeholders are not able to influence him, the model becomes a standard screening problem. Incentive compatibility comes from the fact that the Manager directly selects the quantity schedule to produce as defined in (3.2), or equivalently from the utility maximisation program (3.1). The indirect utility function (the informational rent) can thus be used to characterise incentive compatibility.

Lemma 1. *The Manager's indirect utility U is absolutely continuous, convex, and satisfies the integral representation*

$$U(\theta) = U(\overline{\theta}) + \int_{\theta}^{\overline{\theta}} q(\vartheta) d\vartheta. \tag{3.4}$$

Reciprocally, any allocation schedule (U, q) such that U is absolutely continuous, convex, and can be written as $U(\theta) = T(q(\theta)) - \theta q(\theta)$ with $-q(\theta) \in \partial U(\theta)$ can be implemented by a transfer schedule T.

The integral representation (3.4) can be rewritten using an envelope condition, as U being absolutely continuous ensures almost everywhere differentiability, with

$$\dot{U}(\theta) = -q(\theta) \quad \text{a.e.} \tag{3.5}$$

implying a *monotonicity condition on output*, namely,

$$q \text{ is non-increasing in } \theta. \tag{3.6}$$

Another constraint limiting the Regulator's choice of schedule is the *participation constraint*

$$\forall \theta \in [\underline{\theta}, \overline{\theta}], \quad U(\theta) \geq 0. \tag{3.7}$$

Given the integral representation (3.4) and the positivity of q, it is obvious that U is a decreasing function of θ. Hence, it is sufficient to restrict attention to the least efficient type, that is, to have

$$U(\overline{\theta}) \geq 0. \tag{3.8}$$

The Regulator wants to maximise her expected payoff over all implementable allocations. Substituting $T(q(\theta)) = U(\theta) + \theta q(\theta)$ into the objective, the Regulator's program is given by:

$$\begin{aligned}
\max_{q,U} \quad & \int_{\underline{\theta}}^{\overline{\theta}} \left(S(q(\theta)) - \theta q(\theta) - (1-\alpha)U(\theta) + \mu s q(\theta)\right) f(\theta) d\theta \\
\text{s.t.} \quad & \dot{U}(\theta) = -q(\theta) \text{ a.e.} \\
& U(\overline{\theta}) \geq 0 \\
& q \text{ is non-increasing in } \theta.
\end{aligned}$$

PROPOSITION 1. *Under Assumption 1, the solution to the Regulator's problem when the Stakeholders cannot offer transfers is characterised by*

$$S'(q^{sb}(\theta)) = \theta - \mu s + (1-\alpha)\frac{F(\theta)}{f(\theta)}, \quad \forall \theta \in [\underline{\theta}, \overline{\theta}]. \tag{3.9}$$

This is a standard second-best allocation under asymmetric information, defined by the virtual surplus, with the explicit difference that the marginal cost distribution is shifted by $-\mu s$. Note that this shift may be upward or downward, depending on the sign of s: a stakeholder "supporting" production effectively lowers the social cost of production, while an "opposing" stakeholder increases it. Because this additional term is also present in the first-best benchmark, the standard properties of the second-best solution remain intact:

- *No distortion at the top*: $q^{sb}(\underline{\theta}) = q^{fb}(\underline{\theta})$. As no type wants to mimic the most efficient type, there is no informational reason to distort its allocation.
- *No rent at the bottom*: $U^{sb}(\overline{\theta}) = 0$, which follows directly from the binding participation constraint. Because the least efficient type has no incentive to mimic more efficient types, there is no need to reward such a report.
- *downward distortion*: $\forall \theta \in (\underline{\theta}, \overline{\theta}]$, $q^{sb}(\theta) < q^{fb}(\theta)$ (as $\alpha < 1$), meaning that informational frictions generically reduce production to limit information rents.

These three properties will be central when comparing the optimal regulation with active stakeholders against the second-best benchmark. While the lack of rents at the bottom will survive, the other two properties may not hold.

4 Regulating a Structure

Allow now the Stakeholders to offer a non-linear side-transfer schedule to the Manager. While the Manager has no bargaining power in the influence subgame, he retains the right to reject the Stakeholders' offer and produce according to the Regulator's baseline schedule. The first step to solve this bargaining game is to define the payoff that is to be bargained over: the *indirect utility of the Stakeholder-Manager coalition*, given by:

$$U_C(\theta) = \max_q \{T(q) - (\theta - s)q\}. \tag{4.1}$$

This indirect utility is achieved by selecting a new quantity schedule defined by

$$q_C(\theta) \in \arg\max_q \{T(q) - (\theta - s)q\}. \tag{4.2}$$

Because the Stakeholders possess all bargaining power and the Manager has as an outside option $U(\theta)$, the side-transfers and resulting payoffs are characterised by:

LEMMA 2. *Given θ and T, the Manager chooses to produce $q_C(\theta)$, the efficient quantity for the coalition formed by the Stakeholders and the Manager. His indirect utility remains the same as when the Stakeholders are inactive,*

$$U(\theta) = \max_q \{T(q) - \theta q\},$$

while the Stakeholders are residual claimants, getting a utility of

$$V(\theta) = U_C(\theta) - U(\theta).$$

Furthermore, we have the equalities

$$U(\theta) = U_C(\theta + s) \text{ and } q(\theta) = q_C(\theta + s). \tag{4.3}$$

REMARK 1. *This lemma directly gives that intervention is optimal for Stakeholders in a screening environment. If the Stakeholders remain passive, their baseline utility evaluated at the Manager's optimal quantity is*

$$sq(\theta).$$

By intervening, they become the residual claimants of the rent, achieving, from equality (4.3)

$$V(\theta) = U(\theta - s) - U(\theta).$$

Because asymmetric information forces the Regulator to implement a non-increasing quantity schedule to ensure incentive compatibility, the Manager's informational rent $U(\theta)$ is necessarily convex, with almost everywhere derivative $\dot{U}(\theta) = -q(\theta)$. The standard subgradient characterisation of convex functions ensures that the tangent line always lies below the curve:

$$U(\theta - s) \geq U(\theta) + \dot{U}(\theta)(\theta - s - \theta).$$

Substituting the derivative yields

$$U(\theta - s) - U(\theta) \geq sq(\theta).$$

Noting that the left-hand side is simply $V(\theta)$, this shows that the presence of informational frictions induces optimality of active private politics for Stakeholders.

The Stakeholders' influence has a similar mechanical impact as in the benchmark, effectively shifting the Manager's marginal cost. However, because this shift now arises from direct side-transfers rather than passive internalisation, the Regulator must actively plan around it. Her objective reflects this when substituting the expression for V arising from Lemma 2:

$$\int_{\underline{\theta}}^{\overline{\theta}} \left[S(q_C(\theta)) - T(q_C(\theta)) + \alpha U(\theta) + \mu\left(U_C(\theta) - U(\theta)\right)\right] dF(\theta).$$

Although there are two indirect utility functions to keep track of, equation (4.3) relates them in a highly tractable way. Furthermore, by definition, we can write $T(q_C(\theta)) = U_C(\theta) + (\theta - s)q_C(\theta)$. Hence, the Regulator's objective can be decomposed as:

$$\int_{\underline{\theta}}^{\overline{\theta}} \left(S(q_C(\theta)) - (1-\mu)U_C(\theta) - (\theta - s)q_C(\theta)\right) f(\theta) d\theta - \int_{\underline{\theta}}^{\overline{\theta}} (\mu - \alpha) U_C(\theta + s) dF(\theta). \quad (4.4)$$

The first integral is a *shifted asymmetric information* term, fundamentally similar to standard screening results, up to the cost shift $\theta - s$ and the fact that the Stakeholders are now the residual claimants of rent. The second integral reflects the Regulator's *redistributive concerns* regarding bargaining between the Manager and the Stakeholders.

The tension between these two terms comes from our assumption regarding the *allocation of property rights between Manager and Stakeholders.*[4] Crucially, both the Manager's outside option in the bargaining process — $U(\theta)$ — and the coalitional welfare — $U_C(\theta)$ — are induced by the schedule T, evaluated at two different points due to equality (4.3). The Regulator must therefore navigate a rent-efficiency-redistribution trade-off where the redistributive term is lagged — or anticipated — by the stake s.

This shift has another technical consequence: the necessity of out-of-support evaluation. Indeed, one can rewrite the redistributive concern term in equation (4.4) as

$$\int_{\underline{\theta}+s}^{\overline{\theta}+s} (\mu - \alpha) U_C(\theta) dF(\theta - s).$$

The Regulator must account for the possibility that the Manager rejects the Stakeholders' offer, acting "alone" and thus forfeiting the cost shift $-s$. As it turns out, the optimal regulation given the bargaining described in Lemma 2 mainly depends on the sign of s. Two cases will be considered separately in the remainder of this section: *Support*, when $s > 0$, and *Opposition*, when $s < 0$.

Support. We first consider the case where $s > 0$. In this case, the Stakeholders want to increase production, and thus act as a cost-reducing mechanism. The rent profile should be described on $[\underline{\theta}, \overline{\theta} + s]$, since a Manager rejecting the Stakeholders' offer faces an increased cost.

4. Some works specifically focus on the optimal selection of property rights; see, for instance, Dworczak and Muir (2024). In our framework, what matters is that both parties walk away with some determined positive share of the surplus.

Proposition 2 (Support). *Under Assumption 1, and for $s \geq 0$ sufficiently close to 0,*[5] *the Regulator-optimal quantity schedule is characterised by*

$$\begin{cases} S'(q^s(\theta)) = \theta - s + (1-\mu)\dfrac{F(\theta)}{f(\theta)} + (\mu - \alpha)\dfrac{F(\theta - s)}{f(\theta)} & \theta \in [\underline{\theta}, \overline{\theta}], \\ q^s(\theta) = 0 & \theta \in (\overline{\theta}, \overline{\theta} + s]. \end{cases} \tag{4.5}$$

This characterisation of the optimal regulation directly follows, on $[\underline{\theta}, \overline{\theta}]$, from the decomposition (4.4): the first two terms reflect the usual rent-efficiency trade-off, while the third underlines redistributive concerns. To describe the features of this solution, we compare its properties to the second-best benchmark (3.9):

- Under Support, there is *upward distortion at the top*: the Stakeholders' offer is a transfer to the Manager, which is less costly to the Regulator than direct intervention, as $\mu < 1$.[6] As such, she free-rides on these contributions, creating perfect internalisation of the Stakeholders' preferences in the cost: $S'(q^s(\underline{\theta})) = \underline{\theta} - s$. Note, however, that this does not mean that μ is no longer a relevant parameter, since it measures the extent to which Stakeholders' transfers are costly to the Regulator as well as the cost of leaving rent.[7]

  More generally, note that on $[\underline{\theta}, \underline{\theta} + s]$, we have $F(\theta - s) = 0$, hence redistributive concerns are absent because this interval of virtual costs can only be attained by accepting the Stakeholders' offer. As a consequence, the solution on this interval coincides with a modified second-best benchmark where the effective type is $\theta - s$ and the weight placed on the rent is μ — the weight α placed on the Manager is irrelevant here.
- There is still *no rent at the bottom*: $U^s(\overline{\theta}) = 0$. Although the highest possible cost type a Manager could emulate by rejecting the Stakeholders' offer is $\overline{\theta} + s$, leaving any rent to such off-path reports yields only marginal redistributive benefits of order $\mu - \alpha < 1$. On the other hand, it would require leaving informational rents to all on-path actions, incurring a cost of order 1.[8] Thus, while there is now potentially a benefit to leaving rent at the bottom, it does not outweigh the cost, and Managers with high costs have extremely low outside options.
- *Distortion may be upward or downward*: of course, upward distortion at the top ensures that there cannot be a general downward distortion. However, this does not mean that distortion is upward at any point, as is shown in Figure 1.[9]

While pointwise distortion may go in either direction, we know that the average distortion associated with a stake $s > 0$ is upward:

5. Throughout, the "s close enough to 0" condition simply ensures the absence of distributional ironing à la Toikka (2011) for simplicity. One could easily include this type of bunching.

6. If $\mu = 1$, it is easy to see that $q^s \equiv q^{fb}$ on $[\underline{\theta}, \underline{\theta} + s]$, even if $\alpha < 1$, which directly follows from the fact that the Stakeholder group is the residual claimant of the rent.

7. Recall that the rent-efficiency trade-off does not consider the Manager, as the Stakeholders are the residual claimants of rent.

8. Recall that actions from types in $(\overline{\theta}, \overline{\theta} + s]$ are not expected by the Regulator.

9. The general condition to have downward distortion at the bottom is $f(\overline{\theta})s < 1 + \frac{\mu - \alpha}{1-\mu}F(\overline{\theta} - s)$, where $f(\overline{\theta})$ is obtained by left-continuity of f. The right-hand side is always strictly positive, meaning that for s small enough this inequality holds. If it does not, we have $q^s \geq q^{fb}$ over the whole interval.

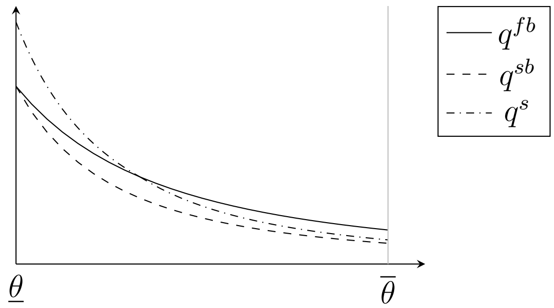

Figure 1: Quantity schedules under complete information, standard asymmetric information, and support.

$\mu = 2/3,\ \alpha = 1/2,\ s = 1/3,\ f(\theta) = \mathbb{1}\{1 \le \theta \le 2\}$

Lemma 3. *The expected distortion increases with s; that is,*

$$\frac{d}{ds}\mathbb{E}\left[q^s(\theta)\right] > \frac{d}{ds}\mathbb{E}\left[q^{sb}(\theta)\right] > 0.$$

Note that when $s = 0$, we have $q^s \equiv q^{sb}$. As such, we obtain an important result:

Corollary 1. *When $s > 0$, ex-ante quantities are higher when the Stakeholder group is able to sway the Manager; that is, $\mathbb{E}\left[q^s(\theta)\right] > \mathbb{E}\left[q^{sb}(\theta)\right]$.*

Opposition. We now consider the case where the Stakeholders want to reduce production, $s < 0$. The setup is similar to the Support regime: the rent profile must now be described on $[\underline{\theta} + s, \overline{\theta}]$, as a Manager who rejects the Stakeholders' offer faces a lowered cost. There is, however, a key difference: in Proposition 2, the "off-path region" constrained the rent profile across the entire support. Under Opposition, this is no longer true: given that $s < 0$, there are no implied informational rents associated with off-path actions taken in $[\underline{\theta} + s, \underline{\theta})$, meaning the primary constraint that may bind the Regulator is the *monotonicity constraint.*

Proposition 3 (Opposition). *Under Assumption 1, for $s \le 0$ sufficiently close to zero, the Regulator-optimal quantity schedule depends on the relative values of μ and α:*

- *If $\alpha \ge \mu$, it is characterised by*

$$\begin{cases} q^o(\theta) = \overline{q} & \theta \in [\underline{\theta} + s, \underline{\theta}), \\ S'(q^o(\theta)) = \theta - s + (1-\mu)\dfrac{F(\theta)}{f(\theta)} + (\mu - \alpha)\dfrac{F(\theta - s)}{f(\theta)} & \theta \in [\underline{\theta}, \overline{\theta}]. \end{cases} \tag{4.6}$$

- *If $\alpha < \mu$, the solution displays bunching at the top,*

$$\begin{cases} q^o(\theta) = b & \theta \in [\underline{\theta}, \vartheta(b)), \\ S'(q^o(\theta)) = \theta - s + (1-\mu)\dfrac{F(\theta)}{f(\theta)} + (\mu - \alpha)\dfrac{F(\theta - s)}{f(\theta)} & \theta \in [\vartheta(b), \overline{\theta}], \end{cases} \tag{4.7}$$

where $\vartheta(b) \in [\underline{\theta}, \overline{\theta}]$ *and* b *are jointly pinned down by*

$$\begin{cases} S'(b) = \vartheta(b) - s + (1-\mu)\dfrac{F(\vartheta(b))}{f(\vartheta(b))} + (\mu - \alpha)\dfrac{F(\vartheta(b) - s)}{f(\vartheta(b))}, \\ \displaystyle\int_{\underline{\theta}+s}^{\vartheta(b)} \big[(S'(b) - (\theta - s))f(\theta) - (1-\mu)F(\theta) - (\mu-\alpha)F(\theta - s)\big]\, d\theta = 0. \end{cases} \tag{4.8}$$

When the Regulator values the Manager more than the Stakeholders ($\alpha \geq \mu$), the outside option provided is maximal in the off-path region, as seen in (4.6): there are no associated costs, and this does not violate the monotonicity constraint imposed by Lemma 1 on q. If the Regulator values the Stakeholders more than the Manager ($\alpha < \mu$), however, there is a severe tension between the monotonicity constraint and her desire to minimise the Manager's outside option, which induces *bunching at the top.*

The characterisation of this bunching in (4.8) relies on standard techniques; see, for instance, Guesnerie and Laffont (1984). It is pinned down by continuity at the threshold $\vartheta(b)$ and a *no distortion on average* condition over the pooled interval. Its economic interpretation, however, is novel. First, unlike in Myerson (1981), this bunching does not arise from exogenous distributional irregularities — such as a failure of the monotone hazard rate property. Second, while this pooling resembles Lewis and Sappington's (1989) inflexible rules, their countervailing incentives stem from the Agent's exogenous, type-dependent outside options. In our framework, the countervailing incentives are endogenous and *Principal-driven*: it is not the Agent's inherent conflict that necessitates non-responsiveness, but the Regulator's deliberate and strategic depression of the Manager's fallback position to redistribute surplus to the Stakeholders.

Once again, we summarise our findings using our core properties, with an example in Figure 2:

- Under Opposition, there is *distortion at the top*, which can take the extreme form of *bunching at the top.* The argument from the Support regime still holds: the Stakeholders' preferences are more heavily internalised than in the second-best benchmark. The difference here is that redistributive concerns remain active at the top. This means that we do not revert to a result akin to the first-best with a shifted cost. Even in the absence of bunching, we have $S'(q^o(\underline{\theta})) = \underline{\theta} - s + (\mu - \alpha)\frac{F(\underline{\theta}-s)}{f(\underline{\theta})}$, which is generically different from $\underline{\theta} - \mu s$.
- There is still *no rent at the bottom.* This is even more direct here than in the Support regime, as the highest possible effective cost type the Regulator may encounter is simply $\overline{\theta}$.
- As in the Support regime, the general direction of pointwise distortions is ambiguous. While we guarantee downward distortion at the bottom, $q^{sb}(\overline{\theta}) > q^o(\overline{\theta})$,[10] what happens at the top is more complex. If bunching occurs, downward distortion at the top is guaranteed, but it may not hold otherwise.[11]

10. This directly follows from $S'(q^o(\overline{\theta})) = \overline{\theta} - s + 1 - \alpha > \overline{\theta} - \mu s + 1 - \alpha = S'(q^{sb}(\overline{\theta}))$.

11. The condition for downward distortion at the top without bunching is $sf(\underline{\theta}) < \frac{\mu-\alpha}{1-\mu}F(\underline{\theta} - s)$. Bunching at the top occurs if $\mu > \alpha$, in which case the right-hand side is positive and the inequality trivially holds for any $s < 0$.

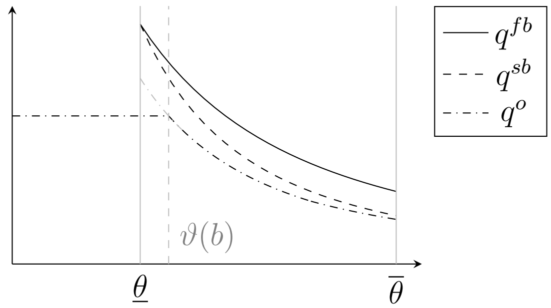

Figure 2: Quantity schedules under complete information, standard asymmetric information, and opposition.

$\mu = 2/3,\ \alpha = 1/2,\ s = -1/3,\ f(\theta) = \mathbb{1}\{1 \leq \theta \leq 2\}$

Similarly to Lemma 3, we can show that the ability to sway the Manager works in favour of the Stakeholders *ex-ante*:

Lemma 4. *The expected distortion increases with s; that is,*

$$\frac{d}{ds}\mathbb{E}\left[q^o(\theta)\right] > \frac{d}{ds}\mathbb{E}\left[q^{sb}(\theta)\right] > 0.$$

Once again, $q^o \equiv q^{sb}$ when $s = 0$. Hence, we have:

Corollary 2. *When $s < 0$, ex-ante quantities are lower when the Stakeholder group is able to sway the Manager; that is, $\mathbb{E}\left[q^o(\theta)\right] < \mathbb{E}\left[q^{sb}(\theta)\right]$.*

5 Multiple Stakeholder Groups

In the previous section, we used a straightforward bargaining solution, disregarding potential friction or collusion between the Stakeholders themselves. It is, however, unrealistic to assume that the Stakeholders form a monolithic group — especially given our departure from the monolithic view of the firm. To provide insights into more complex structures, allowing for multiple Stakeholder groups to simultaneously try to influence the Manager appears essential. Given the linear structure of the stakes, focusing on two groups is sufficient to describe the possible outcomes. We therefore consider a firm where the Stakeholders are divided into two groups, each associated with a stake s_i and a Pareto weight μ_i, $i \in \{1, 2\}$, representing the value to the Regulator of each group's utility.

When multiple groups coexist, it is no longer sufficient to describe the bargaining process as a single take-it-or-leave-it offer, since the bargaining power of the groups must be taken into account. Such models of competition to influence an agent — or common agency — usually take the form of a menu auction *à la* Bernheim and Whinston (1986), where the word "menu" simply reflects the fact that the offers from each Stakeholder group are functions of the Manager's action. In our model, this menu auction takes place given θ and T; that is, in step 3 of the timing described in Section 2. We first describe this new bargaining format.

Menu Auction as Within-Firm Bargaining. In a menu auction, both Stakeholder groups act as principals who simultaneously and non-cooperatively offer menus t_1 and t_2. The Manager is once again free to reject each offered schedule separately, but does not act beyond this acceptance decision.[12] Given that the Stakeholder groups may agree with each other, the equilibrium concept is *Coalition-Proof Nash Equilibrium* in the static game where actions are the transfer schedules.[13]

We then rely on tools from cooperative game theory built by Laussel and Le Breton (2001) to describe this type of game as a transferable utility cooperative game. In such a game, the relevant quantities are the maximum welfare of all possible coalitions. Given that we consider two players — plus the Manager, who is necessarily included in any active coalition as the agent selecting the produced quantity — there are four possible coalitional utility functions:[14]

$$\begin{cases} U_{12}(\theta) = \max_q \left\{ T(q) - (\theta - s_1 - s_2)q \right\}, \\ U_1(\theta) = \max_q \left\{ T(q) - (\theta - s_1)q \right\}, \\ U_2(\theta) = \max_q \left\{ T(q) - (\theta - s_2)q \right\}, \\ U(\theta) = \max_q \left\{ T(q) - \theta q \right\}, \end{cases} \tag{5.1}$$

and associated optimal quantities produced:

$$\begin{cases} q_{12}(\theta) \in \arg\max_q \left\{ T(q) - (\theta - s_1 - s_2)q \right\}, \\ q_1(\theta) \in \arg\max_q \left\{ T(q) - (\theta - s_1)q \right\}, \\ q_2(\theta) \in \arg\max_q \left\{ T(q) - (\theta - s_2)q \right\}, \\ q(\theta) \in \arg\max_q \left\{ T(q) - \theta q \right\}. \end{cases}$$

As was the case when there was a single Stakeholder group, the coalitional marginal cost corresponds to a shifted version of the Manager's type:

$$\begin{aligned} U(\theta) = U_1(\theta + s_1) = U_2(\theta + s_2) = U_{12}(\theta + s_1 + s_2), \\ q(\theta) = q_1(\theta + s_1) = q_2(\theta + s_2) = q_{12}(\theta + s_1 + s_2). \end{aligned} \tag{5.2}$$

From Laussel and Le Breton (2001), we have a simple dichotomy when two groups are competing: either the transferable utility game is convex, or it is concave, depending on the sign of $s_1 s_2$ (that is, depending on whether the Stakeholder groups agree with each other or not).

Lemma 5. *The transferable utility cooperative game described by the coalitional welfares in* (5.1) *is:*

12. The Manager's passive role in this subgame directly mirrors his role in the grand mechanism.

13. If we do not rely on Coalition-Proofness, groups who agree may well fail to influence the Manager due to coordination issues. For an extended discussion on this equilibrium selection, see Bernheim and Whinston (1986).

14. To streamline notation and align with standard cooperative game theory, we denote the indirect utility of a coalition by its participating Stakeholder groups, dropping the subscript C used in the single-stakeholder case.

(i) concave if interests are (strictly) conflicting; that is,

$$s_1 s_2 < 0 \implies U_{12}(\theta) + U(\theta) \leq U_1(\theta) + U_2(\theta).$$

(ii) convex if interests are congruent; that is,

$$s_1 s_2 \geq 0 \implies U_{12}(\theta) + U(\theta) \geq U_1(\theta) + U_2(\theta).$$

This dichotomy with respect to the sign of the stakes follows directly from the convexity of the informational rent function U and equation (5.2), meaning that incentive compatibility places a structural restriction on the political game played within the firm.[15]

The two cases are quite distinct; hence, we study them in turn.

5.1 Conflicting Interests

Interests are said to be conflicting when $s_1 s_2 < 0$, that is, when one group wants to increase production while the other wants to decrease it. There is a clear-cut conflict about the optimal decision. Assume without loss of generality that $s_1 > 0 > s_2$ so that the group supporting production is 1 and the one opposing it is 2. Lemma 5 tells us that the transferable utility cooperative game described in (5.1) is concave. In this case, equilibrium payoffs are uniquely defined, using Laussel and Le Breton's (2001) Proposition 3.3, as

$$V_i(\theta) = U_{12}(\theta) - U_{-i}(\theta), \; i \in \{1, 2\}$$

for each Stakeholder group, and

$$V_m(\theta) = U_{12}(\theta) - V_1(\theta) - V_2(\theta) = U_1(\theta) + U_2(\theta) - U_{12}(\theta) \geq U(\theta)$$

for the Manager. By the concavity of the game established in Lemma 5, pitting the groups against each other now allows the Manager to get some rent above his outside option $U(\theta)$, while each Stakeholder group walks away with their marginal contribution to coalitional welfare. Rewriting in terms of U_{12} and plugging in the Regulator's objective yields:

$$\int_{\underline{\theta}}^{\overline{\theta}} \left[S(q_{12}(\theta)) - T(q_{12}(\theta)) + (\mu_1 + \mu_2 - \alpha)U_{12}(\theta)\right] f(\theta) d\theta \\ + \sum_{i=1}^{2} \int_{\underline{\theta}+s_i}^{\overline{\theta}+s_i} (\alpha - \mu_i) U_{12}(\theta) f(\theta - s_i) d\theta. \quad (5.3)$$

Now, note that Lemma 1 still holds, as equality (5.2) can be used to write U_{12} in terms of U, which is still the same informational rent function as in the second-best benchmark. To ensure a well-behaved problem where informational rent extraction remains costly to the Regulator, we assume $1 + \alpha - \mu_1 - \mu_2 > 0$. Then, standard techniques can once again be applied to obtain:

15. This can be related to Laffont and Martimort (1999), where the convexity of U implies that the separation of regulatory rights among different bodies that may be subject to influence is optimal.

Proposition 4 (Conflicting Interests). *Under Assumption 1, for $s_1 \geq 0 \geq s_2$ close enough to zero, the Regulator-optimal quantity schedule depends on the relative values of μ_2 and α:*

- *If $\alpha \geq \mu_2$, it is characterised by*

$$\begin{cases} q^{ci}(\theta) = \overline{q} & \theta \in [\underline{\theta} + s_2, \underline{\theta}), \\ S'(q^{ci}(\theta)) = \theta - s_1 - s_2 + (1 + \alpha - \mu_1 - \mu_2)\dfrac{F(\theta)}{f(\theta)} & \\ \qquad\qquad + \displaystyle\sum_{i=1}^{2}(\mu_i - \alpha)\frac{F(\theta - s_i)}{f(\theta)} & \theta \in [\underline{\theta}, \overline{\theta}] \\ q^{ci}(\theta) = 0 & \theta \in (\overline{\theta}, \overline{\theta} + s_1]. \end{cases} \tag{5.4}$$

- *If $\alpha < \mu_2$ instead, it is given by*

$$\begin{cases} q^{ci}(\theta) = b & \theta \in [\underline{\theta} + s_2, \vartheta(b)], \\ S'(q^{ci}(\theta)) = \theta - s_1 - s_2 + (1 + \alpha - \mu_1 - \mu_2)\dfrac{F(\theta)}{f(\theta)} & \\ \qquad\qquad + \displaystyle\sum_{i=1}^{2}(\mu_i - \alpha)\frac{F(\theta - s_i)}{f(\theta)} & \theta \in (\vartheta(b), \overline{\theta}], \\ q^{ci}(\theta) = 0 & \theta \in (\overline{\theta}, \overline{\theta} + s_1], \end{cases} \tag{5.5}$$

where b and $\vartheta(b) \in [\underline{\theta}, \overline{\theta}]$ are pinned down by

$$\begin{cases} S'(b) = \vartheta(b) - s_1 - s_2 + (1 + \alpha - \mu_1 - \mu_2)\dfrac{F(\vartheta(b))}{f(\vartheta(b))} \\ \qquad\qquad + \displaystyle\sum_{i=1}^{2}(\mu_i - \alpha)\frac{F(\vartheta(b) - s_i)}{f(\vartheta(b))}, \\ \displaystyle\int_{\underline{\theta}+s_2}^{\vartheta(b)} \Big[\big(S'(b) - (\theta - s_1 - s_2)\big) f(\theta) \\ \qquad -(1 + \alpha - \mu_1 - \mu_2)F(\theta) - \displaystyle\sum_{i=1}^{2}(\mu_i - \alpha)F(\theta - s_i)\Big] d\theta = 0. \end{cases} \tag{5.6}$$

Conflicting Interests combines the off-path regions and arguments from the Support and Opposition regimes: above $\overline{\theta}$, increasing the Manager's outside option is too expensive to be justified; below $\underline{\theta}$, the Regulator has to balance redistributive concerns and the monotonicity of q, which may lead to bunching. On-path and in the absence of bunching, there are now four terms: the first one is the shifted cost $\theta - s_1 - s_2$, showing perfect internalisation of *both* groups. The last two are redistributive terms: there is additive separability of the redistributive concerns across Stakeholder groups, as there is no direct bargaining between them.

The remaining term, $(1 + \alpha - \mu_1 - \mu_2)\frac{F(\theta)}{f(\theta)}$, governs the aggregate allocation of rent and reveals a shift in the model's mechanics. In standard settings, the virtual coefficient is $1 - \alpha$, reflecting the fact that the Manager is the sole claimant of rent.

When Stakeholders are gathered in a single group, this coefficient is $1-\mu$, reflecting that the Stakeholder group extracts all surplus from the Manager. However, under conflicting interests, the competing Stakeholders bid away the surplus: increasing the coalitional rent U_{12} at a point actually *increases* the marginal contribution of both Stakeholders simultaneously, strictly *reducing* the Manager's realised utility V_m. There is thus no single residual claimant absorbing the rent. Instead, the Regulator faces a transfer cost of 1, plus a marginal *penalty* to the Manager α, offset only by the joint gains to the Stakeholders, $\mu_1 + \mu_2$.

Figure 3 plots an example for the specific case where $\mu_1 = \mu_2$ and $s_1 = -s_2$, so that one could expect the Stakeholder groups to "cancel each other out". Indeed, fixing T, and as underlined by Bernheim and Whinston (1986) and Grossman and Helpman (1994), a menu auction showcases *neutrality with respect to the distribution of stakes.* This is trivial to show here as well when $s_1 = -s_2$:

$$q_{12}(\theta) \in \arg\max_q \{T(q) - (\theta - s_1 - s_2)q\} = \arg\max_q \{T(q) - \theta q\} \ni q(\theta).$$

This neutrality, however, is only an intermediary result, as T is not an exogenously given constant as in the studies cited above, but comes from the Regulator's optimisation. Anticipating competition, she may — and generically does — distort the regulation, in turn affecting the framework in which the influence game takes place. Figure 3 shows that this may be the case even when the Regulator does not favour a specific group, as long as she at least cares about redistribution between the Stakeholders and the Manager.[16] To better understand, note that Grossman and Helpman's (1994) notion of neutrality in a menu auction regards the action taken and not the distribution of payoffs, which is precisely what matters to the Regulator here. The core difference is then that Grossman and Helpman's (1994) model takes the influenced agent to be "at the top" of a hierarchy, whereas in our model, the Manager is subject to a screening contract designed by an uninformed external Regulator.

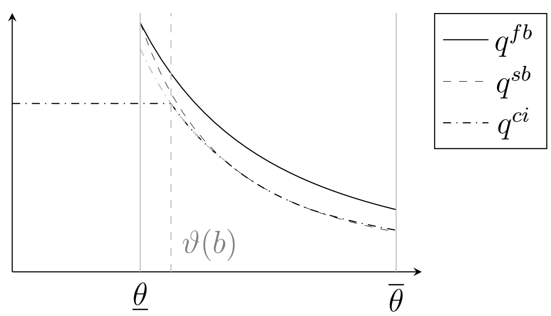

Figure 3: Quantity schedules under complete information, standard asymmetric information, and conflicting interests.

$\mu_1 = \mu_2 = 2/3,\ \alpha = 1/2,\ s_1 = -s_2 = 1/3,\ f(\theta) = \mathbb{1}\{1 \leq \theta \leq 2\}$

In short, this extra step generates a deep non-neutrality result: because the Regulator is constrained by asymmetric information, the initial distribution of stakes and rights

16. This is, of course, the general idea of this paper: were the Regulator not considering redistribution within the firm, we would be back to the standard second-best benchmark. The Regulator would nonetheless still free-ride on the Stakeholders' contributions.

within the firm fundamentally alters the optimal regulatory mechanism.[17]

5.2 Congruent Interests

When $s_1 s_2 > 0$, interests are said to be congruent, as the two Stakeholder groups agree on the direction of the swaying. As stakes are taken to be linear, they also agree on the extent to which influence should take place. While conflicting interests created competition to influence, the game under congruent interests is a *coordination game* with two layers: first, on the aggregate amount of influence to exert, then on the distribution of the costs associated with that aggregate level of influence. Coalition-Proofness deals with the first layer: because a coalition formed by both groups should not be able to improve their joint payoff, equilibria featuring over- or under-influencing are ruled out. This translates into a *no-rent property* (Laussel and Le Breton 2001) arising from the convexity of the game as shown in Lemma 5, stating that the Manager should get his outside option; that is,

$$V_m(\theta) = U(\theta).$$

The second layer, agreeing on "who pays what", creates a continuum of possible equilibria corresponding to the *core* of the cooperative game, characterised by

$$\begin{cases} V_1(\theta) \geq U_1(\theta) - U(\theta), \\ V_2(\theta) \geq U_2(\theta) - U(\theta), \\ V_1(\theta) + V_2(\theta) = U_{12}(\theta) - U(\theta). \end{cases} \tag{5.7}$$

The system (5.7) admits solutions precisely due to the convexity of the game.[18] Note that $U_i(\theta) - U(\theta)$ is the outside option of each group, which must leave $U(\theta)$ to the Manager to be able to influence him alone. On the other hand, the upper bound on group i's payoff is obtained by leaving group $-i$ with their outside option, making group i the residual claimant of the remaining surplus.

To span this continuum of equilibria, we only need a single one-dimensional parameter, as there are only two players bargaining over a fixed total payoff — by the no-rent property.

Definition 1. *Let* $\beta \in [0, 1]$ *be the* bargaining power of Stakeholder 1*, used to derive the payoff redistribution through the expressions*

$$\begin{cases} V_1(\theta) = \beta \left[U_{12}(\theta) - U_2(\theta)\right] + (1 - \beta) \left[U_1(\theta) - U(\theta)\right], \\ V_2(\theta) = (1 - \beta) \left[U_{12}(\theta) - U_1(\theta)\right] + \beta \left[U_2(\theta) - U(\theta)\right]. \end{cases} \tag{5.8}$$

Setting $\beta = 0$ yields group 2's preferred equilibrium, where group 1 is held strictly to its outside option, while $\beta = 1$ represents group 1's preferred equilibrium, where it extracts all residual surplus. Another interesting point to consider is $\beta = \frac{1}{2}$, which

17. The issue of allocating property rights prior to an interaction has been considered in Dworczak and Muir (2024). Here, such an allocation could be used by the firm to obtain a more advantageous regulation, by creating artificial conflict, for instance.

18. Were the game concave instead as when interests are conflicting, we would have $V_1(\theta) + V_2(\theta) \geq U_1(\theta) - U(\theta) + U_2(\theta) - U(\theta) \geq U_{12}(\theta) - U(\theta)$, hence we would get at most a single solution.

corresponds to the *Shapley Value* (Shapley 1953) of this transferable utility cooperative game, a point often considered in cooperative game theory.[19]

For now, let $\beta \in [0,1]$. The Regulator's objective writes, using the envelope condition, as

$$\int_{\underline{\theta}}^{\overline{\theta}} \big[S(q_{12}(\theta)) - (\theta - s_1 - s_2)q_{12}(\theta) \\ \overbrace{- (1 - \beta\mu_1 - (1-\beta)\mu_2)\, U_{12}(\theta) - (\beta\mu_2 + (1-\beta)\mu_1 - \alpha)\, U(\theta)}^{\text{Aggregate Stakeholder impact}} \\ + \underbrace{\beta(\mu_2 - \mu_1)U_2(\theta) + (1-\beta)(\mu_1 - \mu_2)U_1(\theta)}_{\text{Redistribution between groups}}\big]\, dF(\theta).$$

While the redistribution term across groups is as expected, it is interesting to note that the aggregate impact does not amount to what a unified Stakeholder group would achieve, unless $\beta = \frac{1}{2}$ or $\mu_1 = \mu_2$. This wedge arises because the payoff to be shared among the groups is $U_{12}(\theta) - U(\theta)$. Notice that in the coefficient for $U(\theta)$, the bargaining power of group i is multiplied by the Pareto weight of group $-i$. Economically, this means that the group with more bargaining power captures the coalitional surplus $U_{12}(\theta)$, but also forces the *weaker* group to shoulder more of the burden of the Manager's informational rent. Thus, as long as the Regulator weights the groups asymmetrically — $\mu_1 \neq \mu_2$ — the mere presence of two distinct but congruent Stakeholder groups alters the aggregate impact beyond redistributive concerns themselves, because the Regulator must account for the fact that the weaker Stakeholder group is paying for the Manager's rent.

To streamline the analysis, we can isolate the structural impact of the internal bargaining game. We define the *net marginal effect of the shifted outside options* on the Regulator's virtual surplus as the function:

$$\phi(\theta) \equiv (\alpha - [(1-\beta)\mu_1 + \beta\mu_2])\, F(\theta - s_1 - s_2) \\ - \beta(\mu_1 - \mu_2)F(\theta - s_1) + (1-\beta)(\mu_1 - \mu_2)F(\theta - s_2). \quad (5.9)$$

As we will see, this function captures the complexity introduced by fracturing a unified group, driving the quantity distortions both on- and off-path.

As in the case with a single Stakeholder group, the analysis naturally splits depending on the direction of the stakes. Assuming without loss of generality that $s_1 \geq s_2$, we will separately study the subcases $s_1 \geq s_2 \geq 0$ and $0 \geq s_1 \geq s_2$. Mirroring the regimes established in Section 3, we will call the former *Widespread Support* and the latter *Widespread Opposition.*

WIDESPREAD SUPPORT. As the name implies, Widespread Support is closely related to the Support regime, and the logic of Proposition 2 echoes through the following result.

19. Recall that there are four possible coalitions, with payoffs described in (5.1). It is then straightforward to compute the Shapley value in this game with two players.

Proposition 5 (Widespread Support). *Under Assumption 1, for $s_1 \geq s_2 \geq 0$ close enough to zero, the Regulator-optimal quantity schedule is characterised by:*

$$\begin{cases} S'(q^{ws}(\theta)) = \theta - s_1 - s_2 + (1 - \beta\mu_1 - (1-\beta)\mu_2)\dfrac{F(\theta)}{f(\theta)} - \dfrac{\phi(\theta)}{f(\theta)} & \theta \in [\underline{\theta}, \overline{\theta}], \\ q^{ws}(\theta) = 0 & \theta \in (\overline{\theta}, \overline{\theta} + s_1 + s_2]. \end{cases} \tag{5.10}$$

An interesting property to notice is that having $\mu_1 = \mu_2$ brings us exactly back to the Support case. Indeed, given Coalition-Proofness, the Manager's utility is fixed, and the Regulator will distort the transfer schedule T only if redistribution across groups *directly* matters to her. As such, the multiplicity of Stakeholder groups who happen to all agree to increase production is not relevant if the Regulator is not willing to distinguish them by favouring some groups rather than others:

Corollary 3. *All Stakeholders supporting production on whom the Regulator puts the same Pareto weight can be aggregated in a single group, whose stake is the sum of the members' stakes.*

This corollary can be related to neutrality, as it states that the Regulator's ability to distinguish Stakeholders who agree with each other only matters when she explicitly puts different weight on them.

We now focus on the case $\mu_1 \neq \mu_2$. In this situation, there is a *Regulator-preferred continuation equilibrium*, described in the following lemma:

Lemma 6 (Regulator-preferred continuation equilibrium). *The Regulator-preferred continuation equilibrium when $\mu_1 \neq \mu_2$ is characterised by $\beta^* = \mathbb{1}_{\mu_1 > \mu_2}$.*

This lemma states that the Regulator prefers the internal bargaining game to assign all bargaining power to the Stakeholder group on whom she places the highest Pareto weight. The economic intuition for this selection follows directly from the aggregate impact term in the envelope condition: by granting all bargaining power to the favoured group — suppose $\mu_1 > \mu_2$ so this favoured group is group 1, as in Figure 4 — group 2 is pushed down to their outside option and thus shoulders the entire burden of the Manager's informational rent. Because the Regulator places a lower weight on group 2, leaving rent to the Manager becomes less costly from a social welfare perspective.

However, resolving the global comparison between quantities under Widespread Support (q^{ws}) and a unified group (q^s) yields an important insight: the impact of fracturing a monolithic Stakeholder group on actual production is generically ambiguous. Even assuming the internal political game resolves at the Regulator-preferred equilibrium — minimising the social weight applied to the rent — the overall virtual costs also depend on the shape of the type distribution across the discrete shifts induced by the stakes. Because the fractional stakes impose intermediate evaluation points, the convexity or concavity of the distribution over the intervals $[\theta - s_1 - s_2, \theta]$ can either inflate or deflate the virtual cost. Thus, splitting a unified group does not systematically raise or lower production; rather, the political game played within the firm may either push q^{ws} above or below q^s, depending on the specific informational environment.

In the specific uniform distribution example illustrated in Figure 4, evaluating the schedule at the Regulator-preferred continuation equilibrium results in production that is locally increased or unchanged compared to the unified group benchmark, q^s.[20,21] Note, however, that there is no such distortion at the top: $S'(q^{ws}(\underline{\theta})) = \theta - s_1 - s_2$, because a Manager selecting $q^{ws}(\underline{\theta})$ can only act optimally if he is accepting *both* Stakeholder offers. Once again, this relates to the case of Support and Figure 1.

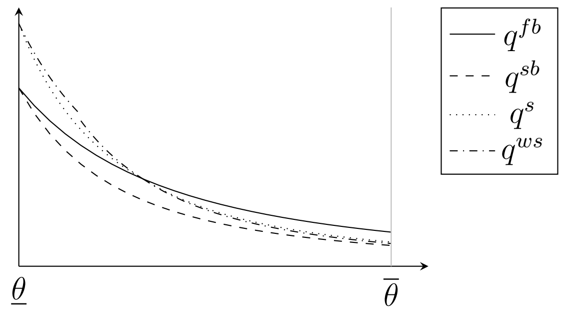

Figure 4: Quantity schedules under complete information, standard asymmetric information, support and widespread support.

$\mu_1 = 14/15,\ \mu_2 = 2/5,\ \alpha = 1/2,\ s_1 = s_2 = 1/6,\ f(\theta) = \mathbb{1}\{1 \leq \theta \leq 2\}$ and $\beta = 1$

Widespread Opposition. Recall our global assumption that $s_1 \geq s_2$. In the context of Widespread Opposition, where both Stakeholder groups oppose production, this implies $0 \geq s_1 \geq s_2$, meaning that group 2 is the most strongly opposed and holds the largest absolute stake.

Once again, the insights from Proposition 3 provide a baseline for understanding this regime. However, the conditions dictating the off-path quantity schedule — and the resulting bunching regions — are significantly more complex. While Propositions 3 and 4 relied on a simple comparison between α (the Pareto weight of the Manager) and μ or μ_2 (the Pareto weight of the opposing Stakeholder group) the coordination game under congruent interests implies that the marginal value of the Manager's outside option to the Regulator is no longer constant over the off-path interval $[\underline{\theta} + s_1 + s_2, \underline{\theta})$.

This variation was irrelevant under Widespread Support for the same reason it was ignored in standard Support: the cost of leaving rent to the Manager is always strictly positive, meaning that the Regulator optimally sets the off-path quantity to zero without violating monotonicity. Under Widespread Opposition, however, the Regulator may actually *want* to increase the Manager's outside option to induce a different redistribution of payoffs within the firm. Therefore, it is crucial to understand exactly *when* the monotonicity constraint binds.

In a relaxed solution, the decision to set the off-path quantity to zero is determined

20. To construct an example where comparison with the case of a unified group is straightforward, we set $s_1 = s_2$ and $\frac{\mu_1+\mu_2}{2} = 2/3$. In general, the unified group should be subject to a Pareto weight ν such that $\mu_1 s_1 + \mu_2 s_2 = \nu(s_1 + s_2)$ to be perfectly comparable.

21. In Figure 4, the distribution F is linear over the interval $[\underline{\theta}, \overline{\theta}]$, but constant over $[\underline{\theta} - s_1 - s_2, \underline{\theta})$. The kink at $\underline{\theta}$ generates local convexity, increasing q^{ws} above q^s at the bottom of the cost distribution.

by the function ϕ defined in (5.9):

$$\forall \theta < \underline{\theta},\ q^r(\theta) = 0 \iff \phi(\theta) < 0. \tag{5.11}$$

When $\mu_1 = \mu_2 \equiv \mu$, the inequality cleanly turns into the condition from Proposition 3: $(\alpha - \mu)F(\theta - s_1 - s_2) < 0$. The complexity under asymmetric weights arises because the function ϕ involves the cumulative distribution function evaluated at intermediate shifts, thus is not necessarily cleanly increasing.

To understand why this is problematic, note that the inequality effectively pits the probability of the Manager rejecting *both* Stakeholders' offers ($F(\theta - s_1 - s_2)$) against the probability of rejecting only one ($F(\theta - s_1)$ and $F(\theta - s_2)$). As the cumulative distribution function may be non-linear, these terms could grow at different rates as θ progresses through the off-path region. Depending on the shape of the distribution, the sign of this inequality can flip back and forth. The relaxed solution may then oscillate between 0 and $\overline{q}$, forcing a Myersonian ironing procedure to smooth out these non-monotonicities into regions of partial bunching rather than a single block.

This translates into the following result:

Proposition 6 (Widespread Opposition). *Under Assumption 1, for $0 \geq s_1 \geq s_2$ close enough to zero, the Regulator-optimal quantity schedule is given by*

$$\begin{cases} q^{wo}(\theta) = \overline{q} & \theta \in [\underline{\theta} + s_1 + s_2, \theta^*), \\ q^{wo}(\theta) = b & \theta \in [\theta^*, \vartheta(b)], \\ S'(q^{wo}(\theta)) = \theta - s_1 - s_2 & \\ \qquad\qquad + (1 - \beta\mu_1 - (1-\beta)\mu_2)\dfrac{F(\theta)}{f(\theta)} - \dfrac{\phi(\theta)}{f(\theta)} & \theta \in (\vartheta(b), \overline{\theta}], \end{cases} \tag{5.12}$$

where the bunching parameters $\theta^* \in [\underline{\theta} + s_1 + s_2, \underline{\theta}]$, $b \in (0, \overline{q})$ and $\vartheta(b) \in [\underline{\theta}, \overline{\theta}]$ are jointly determined by the system:

$$\begin{cases} \theta^* \in \arg \max_{\theta \in [\underline{\theta} + s_1 + s_2, \underline{\theta}]} \displaystyle\int_{\underline{\theta} + s_1 + s_2}^{\theta} \phi(t)dt, \\ S'(b) = \vartheta(b) - s_1 - s_2 + (1 - \beta\mu_1 - (1-\beta)\mu_2)\dfrac{F(\vartheta(b))}{f(\vartheta(b))} - \dfrac{\phi(\vartheta(b))}{f(\vartheta(b))}, \\ \displaystyle\int_{\theta^*}^{\vartheta(b)} \left[(S'(b) - (\theta - s_1 - s_2))f(\theta) - (1 - \beta\mu_1 - (1-\beta)\mu_2)F(\theta) + \phi(\theta) \right] d\theta = 0. \end{cases} \tag{5.13}$$

The structure of Proposition 6 highlights a complex ironing procedure over the off-path tail. Because the objective function is linear with respect to q on the interval $[\underline{\theta} + s_1 + s_2, \underline{\theta})$, the relaxed solution is bang-bang, driven entirely by the sign of the coefficient ϕ.

Bunching is triggered when the relaxed solution oscillates, violating the monotonicity constraint. A sufficient condition for partial bunching to occur (*i.e.*, $\theta^* < \underline{\theta}$) is that the marginal objective evaluated at the boundary of the support is strictly negative: $\phi(\underline{\theta}) < 0$. If this holds, the cumulative objective function $\theta \mapsto \int_{\underline{\theta}+s_1+s_2}^{\theta} \phi(t)dt$ is strictly decreasing at the boundary, ensuring the global maximum θ^* must lie in the interior

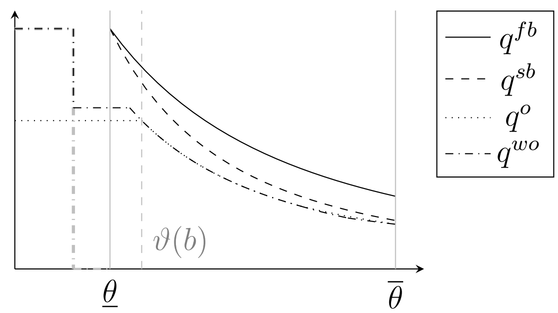

Figure 5: Quantity schedules under complete information, standard asymmetric information, opposition, and widespread opposition.
$\mu_1 = 14/15,\ \mu_2 = 2/5,\ \alpha = 1/2,\ s_1 = s_2 = -1/6,\ f(\theta) = \mathbb{1}\{1 \leq \theta \leq 2\}$ and $\beta = 1$

of the off-path tail. While the condition $\phi(\underline{\theta}) < 0$ may be highly nonlinear, it can be simplified via a first-order Taylor approximation for small stakes, yielding a highly intuitive local sufficient condition for bunching:

$$|s_1|\,(\alpha - \mu_1) + |s_2|\,(\alpha - \mu_2) < 0. \tag{5.14}$$

For infinitesimally small stakes, this linearised condition extends the intuition from the single-group Opposition regime, where bunching occurs if the opposing Stakeholder is valued more than the Manager ($\mu > \alpha$): the Regulator's propensity to bunch the schedule depends on her relative valuation of each group against the Manager ($\alpha - \mu_i$) weighted by their relative share of the aggregate stake ($|s_i| \,/\, |s_1 + s_2|$). The larger a group's proportion of the aggregate off-path threat, the more they matter in the Regulator's choice of off-path production.

However, it is crucial to emphasize that this linear collapse arises from a local approximation. For non-marginal stakes, even the sufficient condition $\phi(\underline{\theta}) < 0$ — which is not necessary — is driven by the non-linearities of the cumulative distribution function across the evaluation points.[22] The marginal objective ϕ pits the probability of a joint rejection against the probability of an individual rejection by each group. These terms grow at different rates depending on the shape of the type distribution, thus the fragmentation of the Stakeholder group has far-reaching consequences on the implemented schedule.

As illustrated in Figure 5, when the boundary condition is met, the non-monotonicity of the virtual cost forces a Myersonian ironing procedure: the optimal schedule bridges the high off-path threat $\overline{q}$ to the on-path unbunched schedule via a bunched plateau at level b. As in the case of Opposition, the bunching is non-neutral despite obeying a no-distortion-on-average condition, for the latter takes into account off-path, never-realised production.

22. Note that even a uniform distribution is subject to such non-linearities, due to the kink of the cumulative distribution function at $\underline{\theta}$.

6 Applications

While the model has been developed using the stylised terminology of a Regulator, a Manager, and Stakeholders, the mathematical framework applies broadly to any principal-agent relationship embedded in a wider political ecosystem. Our analysis yields concrete implications across several domains of applied economics. In this section, we discuss how our results provide a formal mechanism-design foundation for phenomena observed in environmental regulation, decentralised governance, and corporate contracting.

6.1 Environmental Regulation and Private Politics

A rich empirical and theoretical literature highlights that formal environmental regulation does not operate in a vacuum. Firms are frequently subject to forms of informal regulation (Pargal and Wheeler 1996; Feddersen and Gilligan 2001), where local communities and NGOs exert direct pressure on polluting facilities, a phenomenon often conceptualised as "private politics" (Baron 2001; Baron and Diermeier 2007). As Lyon and Maxwell (2004) emphasise, such activism can serve as a substitute for, or a complement to, formal public policy.

Our framework outlines how a formal environmental protection agency (the Regulator) should adjust its baseline emission standards when a firm (the Manager) is targeted by such activist campaigns. If a unified NGO is pressuring the firm to reduce emissions, the Regulator free-rides on this pressure. The agency can then set a more ambitious formal standard, knowing the NGO's threat implicitly acts as a subsidy on the firm's green investments. In this respect, the presence of an NGO engaging in private (rather than public) politics helps environmental policy.

However, environmental activism is rarely monolithic: for instance, a firm might face simultaneous pressure from a local community group and a national environmental NGO. As the Widespread Support regime reveals, even if these two groups have perfectly congruent goals (here, reducing emissions), the Regulator cannot treat them as a single unified force. As Proposition 5 shows, the fragmentation of the activists alters the virtual costs of the mechanism, with a generically ambiguous impact on overall emission reductions.

6.2 Decentralised Governance and Political Capture

The model also translates to the political economy of federalism. Consider a central government (the Regulator) designing a policy mandate for a local mayoral office (the Manager). The local government is inherently better informed about the local cost of public good provision. However, decentralising power makes the local official vulnerable to lobbying by local interest groups, leading to varying degrees of political capture (Bardhan and Mookherjee 2000).

Our results on Conflicting Interests speak directly to the political dynamics of local governance when multiple factions fail to coordinate (Dixit, Grossman, and Helpman 1997). Suppose, for instance, that the central government mandates an infrastructure project. Local business owners (group 1) lobby for expanding the project, while local taxpayer watchdogs (group 2) lobby to scale it down.

Proposition 4 shows that this local political conflict changes the central government's allocation problem. Because the competing local factions bid away the mayor's surplus, the mayor ceases to be the sole residual claimant of the informational rent. Instead, the central government faces a strictly higher virtual cost of implementation. Thus, when local politics are highly polarised and conflicting, the central government cannot simply rely on local factions to neutrally "cancel each other out". Instead, the central government must actively distort the policy mandate to offset the efficiency losses generated by the local political gridlock.

6.3 Corporate Governance and Shareholder Activism

Finally, the framework naturally extends to corporate governance. The Regulator is the Board of Directors, the Manager is the CEO, and the Stakeholders are activist hedge funds (Brav, Jiang, Partnoy, and Thomas 2008) or ESG (Environmental, Social, and Governance) funds trying to influence the CEO's strategic direction.[23]

As Hart and Zingales (2017) note, modern shareholders hold highly heterogeneous preferences. However, even when a subset of shareholders shares a unified broad objective — such as multiple ESG funds opposing a CEO's environmentally damaging expansion — this block is rarely coordinated. Our model shows that the Board's optimal CEO compensation package depends on the relative stakes of these fragmented funds. If the ESG funds are fragmented, the Board knows that the CEO might appease one fund while ignoring another. The Board must therefore design a compensation contract that anticipates the internal bids for influence. The mathematical ironing from our propositions translates into real-world contract rigidities: the Board may choose to offer an inflexible, "flat" compensation package for highly efficient CEOs, because the fragmentation of the shareholder base creates severe non-monotonicities in the Board's virtual costs of leaving informational rent. Rather than smoothly adjusting the CEO's targets, the Board is forced to deploy bunched contracts to balance its own redistributive preferences against the competing activist pressures.

7 Conclusion

This paper incorporates insights from the Theory of the Firm into a standard regulatory screening model, demonstrating that the existence of within-firm bargaining is highly relevant to mechanism design. Because Stakeholders can influence the Manager, they are internalised by the Regulator above their natural Pareto weight; they "hit the Regulator in the wallet" by directly altering the cost of leaving informational informational rent. This introduces redistributive concerns into the drafted regulation, as the Regulator free-rides on the Stakeholders' contributions. When valued Stakeholders oppose production, this dynamic leads to *bunching at the top.* This regulatory rigidity is driven entirely by the political economy of the firm, rather than standard distributional anomalies.

23. Shareholder activism has become increasingly important over recent decades (Stowell and Stowell 2024).

Expanding the framework to multiple Stakeholder groups yields similar insights, but underlines a generic *non-neutrality* of fragmented influence. This non-neutrality is most severe when Stakeholders possess conflicting interests, as their competition to sway the Manager bids away his surplus, affecting the Regulator's virtual costs. However, even when Stakeholders agree on the direction of influence, fragmentation alters the optimal contract. While equally valued Stakeholders seamlessly aggregate, asymmetric weights force the Regulator to navigate intricate between-group redistributions. As we established, this fragmentation creates non-monotonicities, rendering the net effect on production ambiguous, and — in the case of Widespread Opposition — potentially requiring a complex Myersonian ironing procedure and yielding two distinct bunching regions at the top.

These findings open several avenues for future research. Having closed a model with linear stakes, a natural immediate step involves comprehensive comparative statics to assess how the redistribution of stakes impacts optimal production and the distribution of payoffs. Thinking longer term, the framework could be extended to accommodate a generalised, non-linear stake function $s(q, \theta)$ under suitable regularity assumptions, which would capture more realistic bargaining dynamics. Finally, relaxing complete information assumptions — modelling uninformed Stakeholders who do not perfectly observe θ, or a Regulator who is not perfectly informed about the stakes s — could introduce new frictions reflecting the realities of modern organisational ecosystems.

References

Bardhan, Pranab, and Dilip Mookherjee. 2000. "Capture and Governance at Local and National Levels." *The American Economic Review* 90 (2): 135–139.

Baron, David P. 2001. "Private Politics, Corporate Social Responsibility, and Integrated Strategy." *Journal of Economics & Management Strategy* 10 (1): 7–45.

Baron, David P., and Daniel Diermeier. 2007. "Strategic Activism and Nonmarket Strategy." *Journal of Economics & Management Strategy* 16 (3): 599–634.

Baron, David P., and Roger B. Myerson. 1982. "Regulating a Monopolist with Unknown Costs." *Econometrica* 50 (4): 911–930.

Bénabou, Roland, and Jean Tirole. 2010. "Individual and Corporate Social Responsibility." *Economica* 77 (305): 1–19.

Bernheim, B. Douglas, and Michael D. Whinston. 1986. "Menu Auctions, Resource Allocation, and Economic Influence." *The Quarterly Journal of Economics* 101 (1): 1–31.

Brav, Alon, Wei Jiang, Frank Partnoy, and Randall Thomas. 2008. "Hedge Fund Activism, Corporate Governance, and Firm Performance." *The Journal of Finance* 63 (4): 1729–1775.

Dixit, Avinash, Gene M. Grossman, and Elhanan Helpman. 1997. "Common Agency and Coordination: General Theory and Application to Government Policy Making." *Journal of Political Economy* 105 (4): 752–769.

Dworczak, Piotr, and Ellen Muir. 2024. *A mechanism-design approach to property rights.* SSRN Scholarly Paper. Rochester, NY.

Egorov, Georgy, and Bård Harstad. 2017. "Private Politics and Public Regulation." *The Review of Economic Studies* 84 (4 (301)): 1652–1682.

Feddersen, Timothy J., and Thomas W. Gilligan. 2001. "Saints and Markets: Activists and the Supply of Credence Goods." *Journal of Economics & Management Strategy* 10 (1): 149–171.

Fleurbaey, Marc, and Grégory Ponthière. 2023. "The Stakeholder Corporation and Social Welfare." *Journal of Political Economy* 131 (9): 2556–2594.

Grossman, Gene M., and Elhanan Helpman. 1994. "Protection for Sale." *The American Economic Review* 84 (4): 833–850.

Guesnerie, Roger, and Jean-Jacques Laffont. 1984. "A complete solution to a class of principal-agent problems with an application to the control of a self-managed firm." *Journal of Public Economics* 25 (3): 329–369.

Hart, Oliver, and Luigi Zingales. 2017. "Companies Should Maximize Shareholder Welfare Not Market Value." *Journal of Law, Finance and Accounting* 2 (2): 247–275.

Jullien, Bruno. 2000. "Participation Constraints in Adverse Selection Models." *Journal of Economic Theory* 93 (1): 1–47.

Laffont, Jean-Jacques, and David Martimort. 1997. "The Firm as a Multicontract Organization." *Journal of Economics & Management Strategy* 6 (2): 201–234.

———. 1999. "Separation of Regulators against Collusive Behavior." *The RAND Journal of Economics* 30 (2): 232–262.

Laffont, Jean-Jacques, and Jean Tirole. 1991. "The Politics of Government Decision-Making: A Theory of Regulatory Capture." *The Quarterly Journal of Economics* 106 (4): 1089–1127.

———. 1993. *A Theory of Incentives in Procurement and Regulation.* MIT Press. ISBN: 978-0-262-12174-3.

Laussel, Didier, and Michel Le Breton. 2001. "Conflict and Cooperation: The Structure of Equilibrium Payoffs in Common Agency." *Journal of Economic Theory* 100 (1): 93–128.

Lewis, Tracy R., and David E. M. Sappington. 1989. "Inflexible Rules in Incentive Problems." *American Economic Review* 79 (1): 69.

Lyon, Thomas P., and John W. Maxwell. 2004. *Corporate Environmentalism and Public Policy.* Cambridge: Cambridge University Press. ISBN: 978-0-521-81947-3,

Magill, Michael, Martine Quinzii, and Jean-Charles Rochet. 2015. "A Theory of the Stakeholder Corporation." *Econometrica* 83 (5): 1685–1725.

Martimort, David, and Aggey Semenov. 2008. "Ideological uncertainty and lobbying competition." *Journal of Public Economics* 92 (3): 456–481.

Martimort, David, and Lars Stole. 2024. *Menu Auctions under Asymmetric Information.*

Milgrom, Paul, and Ilya Segal. 2002. "Envelope Theorems for Arbitrary Choice Sets." *Econometrica* 70 (2): 583–601.

Myerson, Roger B. 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6 (1): 58–73.

Pargal, Sheoli, and David Wheeler. 1996. "Informal Regulation of Industrial Pollution in Developing Countries: Evidence from Indonesia." *Journal of Political Economy* 104 (6): 1314–1327.

Shapley, L. S. 1953. "17. A Value for n-Person Games." In *Contributions to the Theory of Games, Volume II,* 307–318. Princeton University Press. ISBN: 978-1-4008-8197-0,

Stowell, David P., and Paul Stowell. 2024. "Chapter 13 - Shareholder Activism and Impact on Corporations." In *Investment Banks, Hedge Funds, and Private Equity (Fourth Edition),* edited by David P. Stowell and Paul Stowell, 309–326. Academic Press. ISBN: 978-0-323-88451-8,

Tirole, Jean. 1986. "Hierarchies and Bureaucracies: On the Role of Collusion in Organizations." *Journal of Law, Economics, & Organization* 2 (2): 181–214.

Toikka, Juuso. 2011. "Ironing without control." *Journal of Economic Theory* 146 (6): 2510–2526.

A Proofs

Proof of Lemma 1. Necessity. Convexity of U directly follows from (3.1), as it is a maximum of linear functions of θ. The fact that U is absolutely continuous, almost everywhere differentiable and such that $\dot{U}(\theta) = -q(\theta)$ at every point of differentiability follows from the Envelope Theorem (Milgrom and Segal 2002), which in turn implies that the Integral Representation (3.4) holds.

Sufficiency. Let $(\theta, \hat{\theta}) \in \Theta^2$, then the integral representation yields

$$U(\theta) = U(\hat{\theta}) + \int_{\theta}^{\hat{\theta}} q(\vartheta) d\vartheta.$$

Convexity of U ensures existence of the sub-differential, and as $-q(\hat{\theta}) \in \partial U(\hat{\theta})$, we obtain

$$U(\theta) \geq U(\hat{\theta}) - q(\hat{\theta})(\theta - \hat{\theta}).$$

By defining the transfer schedule as $T(q(\theta)) \equiv U(\theta) + \theta q(\theta)$, this inequality implies $U(\theta) \geq T(q(\hat{\theta})) - \theta q(\hat{\theta})$, thus recovering (3.1). ☐

Proof of Proposition 1. From the integral representation (3.4), we obtain the equality by integration by parts

$$\int_{\underline{\theta}}^{\overline{\theta}} U(\theta) dF(\theta) = U(\overline{\theta}) + \int_{\underline{\theta}}^{\overline{\theta}} q(\theta) F(\theta) d\theta. \tag{A.1}$$

The Regulator's objective then writes

$$\int_{\underline{\theta}}^{\overline{\theta}} \left[S(q(\theta)) - \left(\theta - \mu s + (1-\alpha)\frac{F(\theta)}{f(\theta)} \right) q(\theta) \right] f(\theta) d\theta - U(\overline{\theta}).$$

One can observe that it is optimal to set $U(\overline{\theta}) = 0$, and pointwise maximisation yields (3.9) when ignoring the monotonicity constraint (3.6). Under Assumption 1, this condition is satisfied. ☐

Proof of Lemma 2. Fix θ and T. The problem of the Stakeholder is

$$\max_{t} \; sq - t(q;\theta) \tag{A.2}$$

$$\text{s.t. } q \in \arg\max_{\chi} \{T(\chi) + t(\chi;\theta) - \theta\chi\} \tag{A.3}$$

$$T(q) + t(q;\theta) - \theta q \geq U(\theta) \tag{A.4}$$

Let

$$U_t(\theta) \equiv \max_{q} \{T(q) + t(q;\theta) - \theta q\},$$

and

$$q_t(\theta) \in \arg\max_q \{T(q) + t(q;\theta) - \theta q\}.$$

Constraint (A.3) then amounts to

$$t(q_t(\theta);\theta) = U_t(\theta) + \theta q_t(\theta) - T(q_t(\theta)),$$

and (A.2) can be written as a function of U_t and q_t,

$$s q_t(\theta) + T(q_t(\theta)) - \theta q_t(\theta) - U_t(\theta).$$

It is then trivial to observe that (A.4) is binding at the optimum, hence the program (A.2)-(A.3)-(A.4) can be written in short as

$$\max_{q_t} s q_t(\theta) + T(q_t(\theta)) - \theta q_t(\theta) - U(\theta).$$

It follows that $q_t(\theta) = q_C(\theta)$, hence the Stakeholder's indirect utility is

$$V(\theta) = U_C(\theta) - U(\theta).$$

Equalities (4.3) trivially follow from the definitions (3.1), (4.1), (3.2) and (4.2). □

Proof of Proposition 2. Recall the formulation of the Regulator's payoff (4.4):

$$\int_{\underline{\theta}}^{\overline{\theta}} [S(q_C(\theta)) - (\theta - s)q_C(\theta)] f(\theta)d\theta - (1-\mu)\int_{\underline{\theta}}^{\overline{\theta}} U_C(\theta)dF(\theta) - (\mu - \alpha)\int_{\underline{\theta}}^{\overline{\theta}} U_C(\theta + s)dF(\theta).$$

From (A.1) and (4.3), this can be written as

$$\int_{\underline{\theta}}^{\overline{\theta}} [S(q_C(\theta)) - (\theta - s)q_C(\theta)] f(\theta)d\theta - (1-\mu)\left[\int_{\underline{\theta}}^{\overline{\theta}} q_C(\theta)F(\theta)d\theta + U_C(\overline{\theta})\right]$$
$$- (\mu - \alpha)\left[\int_{\underline{\theta}+s}^{\overline{\theta}+s} q_C(\theta)F(\theta - s)d\theta + U_C(\overline{\theta} + s)\right]. \quad \text{(A.5)}$$

Given that $s \geq 0$, we have $\theta + s \geq \theta$, meaning that we can use

$$U_C(\overline{\theta}) = U_C(\overline{\theta} + s) + \int_{\overline{\theta}}^{\overline{\theta}+s} q_C(\theta)d\theta$$

and set the participation constraint at $\overline{\theta} + s$ binding ($U_C(\overline{\theta} + s) = 0$) to show that (A.5) is equal to

$$\int_{\underline{\theta}}^{\overline{\theta}+s} [(S(q_C(\theta)) - (\theta - s)q_C(\theta)) f(\theta) - [(1-\mu)F(\theta) + (\mu - \alpha)F(\theta - s)] q_C(\theta)] d\theta. \quad \text{(A.6)}$$

For $\theta > \overline{\theta}$, since $f(\theta) = 0$ and $F(\theta) = 1$, this function simplifies to

$$(-(1-\mu) - (\mu - \alpha)F(\theta - s)) q_C(\theta).$$

Given that $(\alpha, \mu) \in [0,1)^2$, this is maximised for $q_C(\theta) = 0$. Indeed,

$$-(1-\mu) - (\mu-\alpha)F(\theta - s) < 0 \iff (\alpha - \mu)F(\theta - s) < 1 - \mu$$
$$\iff \begin{cases} \alpha \leq \mu \\ \text{or} \\ F(\theta - s) < \underbrace{\dfrac{1-\mu}{\alpha - \mu}}_{>1} \text{ and } \alpha > \mu \end{cases}.$$

Over $[\underline{\theta}, \overline{\theta}]$, pointwise optimisation yields (4.5), which has a solution as long as $s < \underline{\theta}$. Finally, for s small enough, and under Assumption 1, the implied q satisfies the monotonicity constraint.[24] □

Proof of Lemma 3. First note that we are obviously taking on-path expectations, meaning that the behaviour of q outside of the support $[\underline{\theta}, \overline{\theta}]$ is irrelevant. Then, we can write from equations (3.9) and (4.5):

$$\mathbb{E}\left[S'(q^{sb}(\theta))\right] = \int_{\underline{\theta}}^{\overline{\theta}} \left[(\theta - \mu s)f(\theta) + (1-\alpha)F(\theta)\right] d\theta,$$

$$\mathbb{E}\left[S'(q^{s}(\theta))\right] = \int_{\underline{\theta}}^{\overline{\theta}} \left[(\theta - s)f(\theta) + (1-\mu)F(\theta) + (\mu-\alpha)F(\theta - s)\right] d\theta.$$

Taking derivatives with respect to s yields

$$\frac{d}{ds}\mathbb{E}\left[S'(q^{sb}(\theta))\right] = -\mu < 0$$

and

$$\begin{aligned}\frac{d}{ds}\mathbb{E}\left[S'(q^{s}(\theta))\right] &= \int_{\underline{\theta}}^{\overline{\theta}} -f(\theta) - (\mu-\alpha)f(\theta - s)d\theta \\ &= -1 - (\mu - \alpha)\int_{\underline{\theta}}^{\overline{\theta}} f(\theta - s)d\theta \\ &= -1 - (\mu-\alpha)\left[F(\overline{\theta} - s) - F(\underline{\theta} - s)\right].\end{aligned}$$

Because $s > 0$, we have $\underline{\theta} - s < \underline{\theta}$, implying $F(\underline{\theta} - s) = 0$. Thus, this evaluates to:

$$\frac{d}{ds}\mathbb{E}\left[S'(q^{s}(\theta))\right] = -1 - (\mu-\alpha)F(\overline{\theta} - s).$$

Then, we have

$$\frac{d}{ds}\mathbb{E}\left[S'(q^{s}(\theta))\right] < \frac{d}{ds}\mathbb{E}\left[S'(q^{sb}(\theta))\right] \iff (\alpha - \mu)F(\overline{\theta} - s) < 1 - \mu.$$

If $\alpha \leq \mu$, this last inequality is trivially true as $\mu < 1$. If $\alpha > \mu$, it is implied by $\alpha < 1$ and $F \leq 1$.

Now, since $S'' < 0$ by assumption and S does not depend on s, this gives the result.[25] □

Proof of Proposition 3. When $s \leq 0$, we use

$$U_C(\overline{\theta} + s) = U_C(\overline{\theta}) + \int_{\overline{\theta}+s}^{\overline{\theta}} q_C(\theta)d\theta$$

and set $U_C(\overline{\theta}) = 0$[26] to rewrite (A.5) as

$$\int_{\underline{\theta}+s}^{\overline{\theta}} \left[\left(S(q_C(\theta)) - (\theta - s)q_C(\theta)\right)f(\theta) - \left[(1-\mu)F(\theta) + (\mu - \alpha)F(\theta - s)\right]q_C(\theta)\right] d\theta. \qquad \text{(A.7)}$$

24. If this is not the case, standard ironing à la Toikka (2011) yields the solution.
25. Recall that even if bunching à la Toikka (2011) happens, it entails no distortion in the average quantity.
26. Binding participation constraint for the worst type.

First order conditions, ignoring the monotonicity constraint, now yields the marginal objective:

$$\begin{cases} (\alpha-\mu)F(\theta-s) & \theta\in[\underline{\theta}+s,\underline{\theta}), \\ \left[S'(q_C(\theta))-(\theta-s)\right]f(\theta)-(1-\mu)F(\theta)-(\mu-\alpha)F(\theta-s) & \theta\in[\underline{\theta},\overline{\theta}]. \end{cases}$$

Setting the second line to zero yields the interior solution for $\theta\in[\underline{\theta},\overline{\theta}]$.

For s close enough to 0, the implied q_C satisfies the monotonicity condition on $[\underline{\theta},\overline{\theta}]$. If furthermore $\alpha\geq\mu$, the marginal objective on the lower interval is non-negative, and the resulting solution (4.6) can be extended as non-increasing in θ on the entire interval $[\underline{\theta}+s,\overline{\theta}]$.

However, if $\alpha<\mu$, the marginal objective on $[\underline{\theta}+s,\underline{\theta})$ is strictly negative. The relaxed solution dictates $q_C(\theta)=0$ on this lower interval, violating monotonicity as $q_C(\underline{\theta})>0$ under the Inada conditions on S. To restore monotonicity, bunching must be introduced in the form of

$$q^B(\theta;b)\equiv\begin{cases} b & \theta\in[\underline{\theta}+s,\underline{\theta}) \\ \min\{b;q_C(\theta)\} & \theta\in[\underline{\theta},\overline{\theta}] \end{cases}$$

with $b\geq q_C(\overline{\theta})$. Let the intersection threshold be defined as

$$\vartheta(b)=\sup\left\{\theta\in[\underline{\theta},\overline{\theta}]\mid q^B(\theta;b)=b\right\}\vee\underline{\theta}=\inf\left\{\theta\in[\underline{\theta},\overline{\theta}]\mid q^B(\theta;b)=q_C(\theta)\right\}\wedge\overline{\theta}.$$

By definition, we have $q_C(\vartheta(b))=b$, hence the first condition in (4.8) must hold. The optimal plateau b maximises the Regulator's objective function:

$$\int_{\underline{\theta}+s}^{\vartheta(b)}\left[S(b)-b(\theta-s)\right]f(\theta)-b\left[(1-\mu)F(\theta)+(\mu-\alpha)F(\theta-s)\right]d\theta$$
$$+\int_{\vartheta(b)}^{\overline{\theta}}\left[S(q_C(\theta))-q_C(\theta)(\theta-s)\right]f(\theta)-q_C(\theta)\left[(1-\mu)F(\theta)+(\mu-\alpha)F(\theta-s)\right]d\theta.$$

Differentiating this objective with respect to b and setting it to zero yields the second condition in (4.8), completing the proof. □

Proof of Lemma 4. Focus on the relaxed solution, ignoring bunching at the top. In that case, (4.6) gives:

$$\mathbb{E}\left[S'(q^o(\theta))\right]=\int_{\underline{\theta}}^{\overline{\theta}}\left[(\theta-s)f(\theta)+(1-\mu)F(\theta)+(\mu-\alpha)F(\theta-s)\right]d\theta,$$

while the expected second best quantity writes:

$$\mathbb{E}\left[S'(q^{sb}(\theta))\right]=\int_{\underline{\theta}}^{\overline{\theta}}\left[(\theta-\mu s)f(\theta)+(1-\alpha)F(\theta)\right]d\theta.$$

Differentiating with respect to s gives:

$$\frac{d}{ds}\mathbb{E}\left[S'(q^{sb}(\theta))\right]=-\mu<0$$

and

$$\begin{aligned}\frac{d}{ds}\mathbb{E}\left[S'(q^o(\theta))\right]&=\int_{\underline{\theta}}^{\overline{\theta}}\left[-f(\theta)-(\mu-\alpha)f(\theta-s)\right]d\theta\\&=-1-(\mu-\alpha)\int_{\underline{\theta}}^{\overline{\theta}}f(\theta-s)d\theta\\&=-1-(\mu-\alpha)\left[1-F(\underline{\theta}-s)\right].\end{aligned}$$

Then, we obtain:

$$\frac{d}{ds}\mathbb{E}\left[S'(q^o(\theta))\right] < \frac{d}{ds}\mathbb{E}\left[S'(q^{sb}(\theta))\right] \iff (\alpha-\mu)(1-F(\underline{\theta}-s)) < 1-\mu.$$

This latter inequality is always true: if $\alpha \leq \mu$ this follows from $\mu < 1$, if $\alpha > \mu$ it follows from $\alpha < 1$ and $F \geq 0$.

Since $S'' < 0$ by assumption and S does not directly depend on s, a steeper drop in expected marginal surplus implies a larger increase in expected quantity for the relaxed solution. Finally, recall from the proof of Proposition 3 that if bunching at the top occurs, the ironed schedule on the relevant support is defined as $q^B(\theta; b) = \min\{b, q_C(\theta)\}$. Because the bunched quantity is pointwise weakly lower than the relaxed quantity across the entire distribution, the introduction of our bunching can only lower the expected quantity further. □

Proof of Lemma 5. Assume throughout that $s_1 > 0$.[27]

(i) Suppose $s_1 s_2 < 0$. Given our assumption $s_1 > 0$, this amounts to $s_2 < 0$. We can then order the points as

$$\theta + s_2 < \min\{\theta, \theta + s_1 + s_2\} \leq \max\{\theta, \theta + s_1 + s_2\} < \theta + s_1.$$

From this ordering, there exists $\lambda \in (0,1)$ such that:

$$\theta = \lambda(\theta + s_2) + (1-\lambda)(\theta + s_1).$$

Furthermore, as $\theta = (\theta + s_1) + (\theta + s_2) - (\theta + s_1 + s_2)$, the remaining interior point shares the inverted convex weights:

$$\theta + s_1 + s_2 = (1-\lambda)(\theta + s_2) + \lambda(\theta + s_1).$$

By convexity of U_{12}, this implies that:

$$U_{12}(\theta) + U_{12}(\theta + s_1 + s_2) \leq U_{12}(\theta + s_1) + U_{12}(\theta + s_2),$$

establishing strong subadditivity.

(ii) Suppose instead $s_1 s_2 \geq 0$. Given our assumption $s_1 > 0$, this implies $s_2 \geq 0$. The ordering of the points is now:

$$\theta \leq \min\{\theta + s_1, \theta + s_2\} \leq \max\{\theta + s_1, \theta + s_2\} \leq \theta + s_1 + s_2.$$

Using the exact same sum-of-arguments property ($(\theta + s_1) + (\theta + s_2) = \theta + (\theta + s_1 + s_2)$), the interior points can be expressed as symmetric convex combinations of the extremes:

$$\exists \lambda \in [0,1], \quad \begin{cases} \theta + s_1 = \lambda\theta + (1-\lambda)(\theta + s_1 + s_2), \\ \theta + s_2 = (1-\lambda)\theta + \lambda(\theta + s_1 + s_2). \end{cases}$$

Applying convexity of U_{12} again, this yields

$$U_{12}(\theta + s_1) + U_{12}(\theta + s_2) \leq U_{12}(\theta) + U_{12}(\theta + s_1 + s_2),$$

establishing superadditivity; that is, convexity of the TU game.

□

27. The case $s_1 < 0$ can be proved in the same way. In the case $s_1 = 0$, both the subadditivity and superadditivity conditions hold trivially for any s_2.

Proof of Proposition 4. Lemma 1 still holds, as $U_{12}(\theta) = U(\theta - s_1 - s_2)$, allowing us to write

$$U_{12}(\theta) = U_{12}(\overline{\theta} + s_1) + \int_{\theta}^{\overline{\theta}+s_1} q_{12}(\vartheta)d\vartheta.$$

Integration by parts then yields

$$\int_{\underline{\theta}}^{\overline{\theta}+s_1} U_{12}(\theta)dF(\theta) = U_{12}(\overline{\theta} + s_1) + \int_{\underline{\theta}}^{\overline{\theta}+s_1} q_{12}(\theta)F(\theta)d\theta,$$

hence (5.3) can be written as[28]

$$\int_{\underline{\theta}}^{\overline{\theta}} \left[S(q_{12}(\theta)) - (\theta - s_1 - s_2)q_{12}(\theta)\right] f(\theta)d\theta - (1 + \alpha - \mu_1 - \mu_2) \int_{\underline{\theta}}^{\overline{\theta}+s_1} q_{12}(\theta)F(\theta)d\theta$$
$$+ \int_{\underline{\theta}+s_2}^{\overline{\theta}+s_1} q_{12}(\theta)\left[(\alpha - \mu_1)F(\theta - s_1) + (\alpha - \mu_2)F(\theta - s_2)\right] d\theta - (1 - \alpha)U_{12}(\overline{\theta} + s_1).$$

Because $\alpha < 1$, the participation constraint at $\overline{\theta} + s_1$ is binding at the optimum, allowing the last term to be dropped. Pointwise differentiating with respect to q_{12}, ignoring the monotonicity constraints, yields the marginal objective:

$$\begin{cases} (\alpha - \mu_2)F(\theta - s_2) & \theta \in [\underline{\theta} + s_2, \underline{\theta}), \\ \left[S'(q_{12}(\theta)) - (\theta - s_1 - s_2)\right] f(\theta) - (1 + \alpha - \mu_1 - \mu_2)F(\theta) & \\ \qquad\qquad - \sum_{i=1}^{2}(\mu_i - \alpha)F(\theta - s_i) & \theta \in [\underline{\theta}, \overline{\theta}], \\ -(1 - \mu_1) - (\mu_1 - \alpha)F(\theta - s_1) & \theta \in (\overline{\theta}, \overline{\theta} + s_1]. \end{cases} \tag{A.8}$$

Setting the middle line to zero yields the interior solution. As was the case with Proposition 2, the marginal objective on $(\overline{\theta}, \overline{\theta} + s_1]$ is strictly negative, yielding $q_{12}(\theta) = 0$. Similarly, the first condition has to be dealt with as in Proposition 3:

- If $\alpha \geq \mu_2$, then the marginal objective is weakly positive, driving $q_{12}(\theta)$ to its maximum feasible level $\overline{q}$, and the monotonicity constraint is slack.
- If instead $\alpha < \mu_2$, then the relaxed solution entails $q_{12}(\theta) = 0$, meaning that the monotonicity constraint is violated for the most efficient types. Let then

  $$q^B(\theta; b) \equiv \begin{cases} b & \theta \in [\underline{\theta} + s_2, \underline{\theta}), \\ \min\{b; q_{12}(\theta)\} & \theta \in [\underline{\theta}, \overline{\theta} + s_1], \end{cases}$$

  with $b \geq q_{12}(\overline{\theta})$, and

  $$\vartheta(b) = \sup\left\{\theta \in [\underline{\theta}, \overline{\theta}] \mid q^B(\theta; b) = b\right\} \vee \underline{\theta} = \inf\left\{\theta \in [\underline{\theta}, \overline{\theta}] \mid q^B(\theta; b) = q_{12}(\theta)\right\} \wedge \overline{\theta}.$$

  Once again, $q_{12}(\vartheta(b)) = b$ by definition, implying the first condition in (5.6). The Regulator's expected payoff over the bunched schedule writes

  $$\int_{\underline{\theta}+s_2}^{\vartheta(b)} \left[S(b) - b(\theta - s_1 - s_2)\right] f(\theta)$$
  $$- b\left[(1 + \alpha - \mu_1 - \mu_2)F(\theta) + (\mu_1 - \alpha)F(\theta - s_1) + (\mu_2 - \alpha)F(\theta - s_2)\right] d\theta$$
  $$+ \int_{\vartheta(b)}^{\overline{\theta}+s_1} \left[S(q_{12}(\theta)) - q_{12}(\theta)(\theta - s_1 - s_2)\right] f(\theta)$$
  $$- q_{12}(\theta)\left[(1 + \alpha - \mu_1 - \mu_2)F(\theta) + (\mu_1 - \alpha)F(\theta - s_1) + (\mu_2 - \alpha)F(\theta - s_2)\right] d\theta.$$

28. Using $T(q_{12}(\theta)) = U_{12}(\theta) + (\theta - s_1 - s_2)q_{12}(\theta)$, from $q_{12}(\theta) = q(\theta - s_1 - s_2)$.

Differentiating this objective with respect to b and setting it to zero yields the second condition in (5.6), completing the proof.

$\square$

Proof of Proposition 5. Once again, Lemma 1 holds, giving

$$U_{12}(\theta) = U(\overline{\theta} + s_1 + s_2) + \int_{\theta}^{\overline{\theta}+s_1+s_2} q_{12}(\vartheta)d\vartheta.$$

Integration by parts and binding participation constraint at the worst type give the payoff function (with ϕ defined as in (5.9)):

$$\int_{\underline{\theta}}^{\overline{\theta}+s_1+s_2} \Big\{ \left[S(q_{12}(\theta)) - q_{12}(\theta)(\theta - s_1 - s_2)\right] f(\theta) \\ - q_{12}(\theta)(1 - \beta\mu_1 - (1-\beta)\mu_2)F(\theta) + q_{12}(\theta)\phi(\theta)\Big\} d\theta.$$

On $[\underline{\theta}, \overline{\theta}]$ for s_1 and s_2 close enough to zero so that $q_{12}(\theta)$ satisfies the monotonicity condition, we recover (5.10). For $\theta \in (\overline{\theta}, \overline{\theta} + s_1 + s_2]$, the marginal objective evaluates to

$$-(1 - \beta\mu_1 - (1-\beta)\mu_2) + \phi(\theta).$$

By expanding $\phi(\theta)$, this expression can be rearranged into:

$$\alpha F(\theta - s_1 - s_2) - 1 + \mu_1 \left[(1-\beta)\left(F(\theta - s_2) - F(\theta - s_1 - s_2)\right) + \beta(1 - F(\theta - s_1))\right] \\ + \mu_2 \left[\beta(F(\theta - s_1) - F(\theta - s_1 - s_2)) + (1-\beta)\left(1 - F(\theta - s_2)\right)\right].$$

Note that μ_1 and μ_2 are multiplied by convex combinations of positive numbers. This quantity is thus maximised — with respect to μ_1 and μ_2 — when $\mu_1 = \mu_2 = 1$, in which case it is negative because $\alpha < 1$ and $F(\theta - s_1 - s_2) \leq 1$. Hence, we recover

$$q_{12}(\theta) = 0 \text{ on } (\overline{\theta}, \overline{\theta} + s_1 + s_2].$$

$\square$

Proof of Lemma 6. Recall that the Regulator's objective writes

$$\int_{\underline{\theta}}^{\overline{\theta}} \big[S(q_{12}(\theta)) - (\theta - s_1 - s_2)q_{12}(\theta) \\ - (1 - \beta\mu_1 - (1-\beta)\mu_2)U_{12}(\theta) - (\beta\mu_2 + (1-\beta)\mu_1 - \alpha)U(\theta) \\ + \beta(\mu_2 - \mu_1)U_2(\theta) + (1-\beta)(\mu_1 - \mu_2)U_1(\theta)\big] dF(\theta).$$

Taking its derivative with respect to β yields

$$\int_{\underline{\theta}}^{\overline{\theta}} (\mu_1 - \mu_2)\left[U_{12}(\theta) + U(\theta) - U_1(\theta) - U_2(\theta)\right] dF(\theta).$$

Given that we are in the case of congruent interests, the continuation game is convex as stated in Lemma 5, implying that

$$\forall \theta \in [\underline{\theta}, \overline{\theta}],\ U_{12}(\theta) + U(\theta) - U_1(\theta) - U_2(\theta) \geq 0.$$

As such, the derivative of the Regulator's objective with respect to β has the sign of $\mu_1 - \mu_2$. Combined with the restriction $\beta \in [0, 1]$, this corner solution yields[29]

$$\beta^* = \mathbb{1}_{\mu_1 > \mu_2}.$$

$\square$

29. If $\mu_1 = \mu_2$, the Regulator is indifferent between all values of β.

Proof of Proposition 6. Recall from Lemma 1 and integration by parts that for $s_1, s_2 \leq 0$, the Regulator's objective can be written as (with ϕ defined as in (5.9)):

$$\int_{\underline{\theta}+s_1+s_2}^{\overline{\theta}} \Big\{ \left[S(q_{12}(\theta)) - q_{12}(\theta)(\theta - s_1 - s_2)\right] f(\theta) \\ - q_{12}(\theta) \left(1 - \beta\mu_1 - (1-\beta)\mu_2\right) F(\theta) + q_{12}(\theta)\phi(\theta) \Big\} d\theta.$$

Pointwise differentiation with respect to q_{12}, ignoring monotonicity constraints, yields the marginal objective:

$$\left[S'(q_{12}(\theta)) - (\theta - s_1 - s_2)\right] f(\theta) - (1 - \beta\mu_1 - (1-\beta)\mu_2)F(\theta) + \phi(\theta).$$

On the support $[\underline{\theta}, \overline{\theta}]$, setting this marginal objective to zero yields the interior schedule in (5.12). However, on the lower tail $[\underline{\theta}+s_1+s_2, \underline{\theta})$, both the density and the distribution function are nil, thus the marginal objective collapses to the linear coefficient $\phi(\theta)$.

If $\phi(\theta) \geq 0$ everywhere on this tail, the relaxed schedule entails $q_{12}(\theta) = \overline{q}$ and the monotonicity constraint is slack. If instead $\phi(\theta) < 0$ somewhere on the tail, the relaxed solution drops to zero, violating monotonicity with the positive interior schedule on $[\underline{\theta}, \overline{\theta}]$.

In that case, bunching must be introduced to restore monotonicity. Let $\Phi(\theta) \equiv \int_{\underline{\theta}+s_1+s_2}^{\theta} \phi(\vartheta)d\vartheta$. Because the objective is linear on the tail, standard ironing arguments (see, *e.g.*, Myerson 1981) state that the optimal monotone schedule, given a boundary plateau b, takes the maximum feasible value $\overline{q}$ up to a global maximiser of the cumulative marginal objective, $\theta^* \in \arg\max_{\theta \in [\underline{\theta}+s_1+s_2, \underline{\theta}]} \Phi(\theta)$, before dropping to b. The fully ironed schedule is thus:

$$q^B(\theta; b) \equiv \begin{cases} \overline{q} & \text{if } \theta \in [\underline{\theta} + s_1 + s_2, \theta^*), \\ b & \text{if } \theta \in [\theta^*, \vartheta(b)], \\ \hat{q}_{12}(\theta) & \text{if } \theta \in (\vartheta(b), \overline{\theta}], \end{cases}$$

where $\hat{q}_{12}$ is the interior solution on $[\underline{\theta}, \overline{\theta}]$ and $\vartheta(b) \equiv \sup\left\{\theta \in [\underline{\theta}, \overline{\theta}] \mid \hat{q}_{12}(\theta) \geq b\right\}$. By definition, $\hat{q}_{12}(\vartheta(b)) = b$, ensuring continuity.

The optimal plateau b maximises the Regulator's expected payoff over the piecewise bunched schedule. Differentiating the Regulator's objective over $q^B(\theta; b)$ with respect to b and setting it to zero yields the integral condition:

$$\int_{\theta^*}^{\underline{\theta}} \phi(\theta)d\theta + \int_{\underline{\theta}}^{\vartheta(b)} \frac{\partial \psi}{\partial q}(\theta; b)d\theta = 0,$$

where $\psi(\theta; q)$ is the integrand of the objective function on the main support. Rearranging this first-order condition yields the third condition in (5.13), completing the proof. □