Theoretical Economics 13 (2018), 1151–1189 1555-7561/20181151

A complete characterization of equilibria in an intrinsic common agency screening game

David Martimort
Ecole des Hautes Etudes en Sciences Sociales, Paris School of Economics

Aggey Semenov
Department of Economics, University of Ottawa

Lars A. Stole
Booth School of Business, University of Chicago

We characterize the complete set of equilibrium allocations to an intrinsic common agency screening game as the set of solutions to *self-generating* optimization programs. We provide a complete characterization of equilibrium outcomes for regular environments by relying on techniques developed elsewhere for aggregate games and for the mechanism design delegation literature. The set of equilibria includes those with nondifferentiable payoffs and discontinuous choices, as well as equilibria that are smooth and continuous in types. We identify one equilibrium, the *maximal equilibrium*, which is the unique solution to a self-generating optimization program with the largest (or "maximal") domain, and the only equilibrium that is supported with biconjugate (i.e., least-concave) tariffs. The maximal equilibrium exhibits a n-fold distortion caused by each of the n principal's non-cooperative behavior in overharvesting the agent's information rent. Furthermore, in any equilibrium, over any interval of types in which there is full separation, the agent's equilibrium action corresponds to the allocation in the maximal equilibrium. Under reasonable conditions, the maximal equilibrium maximizes the agent's information rent within the class of equilibrium allocations. When the principals' most-preferred equilibrium allocation differs from the maximal equilibrium, we demonstrate that the agent's choice function exhibits an interval of bunching over the worst agent types, and elsewhere corresponds to the maximal allocation. The optimal region of bunching trades off the principals' desire to constrain inefficient n-fold marginalizations of the agent's rent against the inefficiency of pooling agent types.

Keywords. Intrinsic common agency, aggregate games, mechanism design for delegated decision-making, duality, equilibrium selection.

David Martimort: david.martimort@psemail.eu
Aggey Semenov: aggey.semenov@uottawa.ca
Lars A. Stole: lars.stole@chicagobooth.edu
For useful remarks, we thank seminar and conference participants at the Canadian Economic Theory 2017 Meeting at UBC, the Canadian Economic Association 2016 in Ottawa, and 2017 in Antigonish, and the 2017 SWET Meeting in Paris. We also thank Andrea Attar, Didier Laussel, four referees, and Johannes Horner for useful comments that lead us to considerably reshape the material of this paper. The usual disclaimer applies.

Available at http://econtheory.org. https://doi.org/10.3982/TE2266

JEL classification. D82, D86.

1. Introduction

We consider a canonical class of common agency games in which principals simultaneously offer contracts to a privately informed common agent whose action is publicly observable and contractible by all principals, and who must either accept all contract offers from the principals or choose not to participate. Common agency is thus *public* and *intrinsic*.[1] As a motivating example, suppose there are multiple government agencies (principals) that regulate a polluting public utility (the common agent) that has private information about the cost of production. If the firm decides to produce, it is under the joint control of all regulators. Regulators, however, may have conflicting objectives. For example, an environmental agency wishes on the margin to reduce output and negative externalities, while a public-utility commission instead prefers to increase output and consumer surplus. In this game regulators simultaneously offer menus of transfer–output pairs so as to influence the choice of the public utility. The public utility must either choose an output and abide by the consequences of each principal's menu or exit the market entirely.

One of the main theoretical difficulties of modeling non-cooperative scenarios is to characterize the multitude of equilibrium outcomes that can arise. In more familiar single-principal screening environments, the *revelation principle* defines the set of relevant communication strategies and describes feasible allocations by means of incentive compatibility constraints.[2] With multiple principals, however, the revelation principle is neither simple to apply nor particularly useful. Even though the *delegation principle* proposed in Martimort and Stole (2002)[3] does offer a simple and universal representation of the strategy spaces available to mechanism designers in common agency environments, this tool fails to give a complete representation of equilibrium allocations. As a result, the literature on common agency has primarily focused on specific equilibria in structured games rather than exploring the entire set of equilibrium possibilities. In particular, previous analyses have often restricted attention to differentiable equilibria both because of their tractability and because of the attractiveness of the simple economic insights that emerge. This restriction, however, is with loss of generality and the arbitrariness of such a selection raises concerns about the robustness of any implications deduced from the refined set. A more complete approach—the task of the present paper—is to characterize the entire set of equilibria, to make welfare comparisons across equilibria, and, where possible, to make broader statements that apply to all equilibria.

Insights from aggregate games

Our first step toward a full characterization of equilibria relies on the fundamental structure of intrinsic common agency games. As noted by Martimort and Stole (2012), these

[1] See Martimort (2006) for a review of these definitions and modeling choices in common agency games.
[2] Myerson (1982).
[3] Also sometimes referred to as *menu theorems* in the parlance of Peters (2001).

games are special cases of *aggregate games*. Because the agent only cares about the sum of the payments offered by the non-cooperating principals, incentive compatibility and participation constraints can only depend on the resulting *aggregate contract*. As a consequence, principal i's expected payoff depends only upon his own contract and the aggregate contract (i.e., the sum of contracts) offered by the other principals. Because each principal can always undo the aggregate contract offered by others using only his own tariff, a principal can implement any incentive-feasible allocation he would like. It follows that a necessary condition for any equilibrium is that all principals agree on inducing the same allocation. Such agreement has remarkable consequences. In particular, because each principal's virtual surplus function is maximized by the equilibrium allocation, it must also be that the sum of the principals' individual virtual surpluses is also maximized by this allocation. This aggregate virtual surplus function, however, is *not* the same as the virtual surplus function that would arise in a cooperative setting in which the principals jointly contract with the agent. Critically, the former corresponds to what a fictional principal would maximize if this principal valued rent extraction n times more than is the case. In equilibrium, everything happens as if the individual principals were to delegate their choice of allocation to a *surrogate* principal whose payoffs are distorted relative to the collective preference of the principals. This *aggregate concurrence principle*, as coined by Martimort and Stole (2012), is a key ingredient to characterize the set of equilibrium outcomes.[4]

Self-generating maximization programs

The aggregate concurrence principle provides necessary conditions for equilibria. We are, however, interested in the set of equilibria that may be decidedly smaller. To this end, we demonstrate that the solution set of carefully chosen self-generating maximization (hereafter SGM) programs corresponds to the equilibrium set of our common-agency game.[5] In our common-agency game, following the approach in Martimort and Stole (2012), we first establish that the solution set to an infinite-dimensional SGM program corresponds to the set of equilibrium allocations (Proposition 1).[6] Here, the

[4]The implication of aggregate concurrence was used by Bernheim and Whinston (1986) in a moral hazard setting, but it applies to a larger class of aggregate games and, in particular, to our present screening model.

[5]To be clear at the outset about our concept of an SGM program, consider a canonical maximization program with objective function ϕ defined over a domain X, with the additional feature that ϕ is parameterized by an arbitrary reference point, $\hat{x}$, which also lies in the choice domain X. We denote $\phi(x, \hat{x})$ to be the value of this objective evaluated at the choice x, given the reference point $\hat{x}$. The pair (ϕ, X) gives rise to a self-generating maximization program whose solution set is defined by the requirement that each element, x^*, satisfies

$$x^* \in \arg\max_{x \in X} \phi(x, x^*).$$

Hence, self-generating problems are optimization problems with a fixed point.

[6]Martimort and Stole (2012) used self-generating programs to prove equilibrium existence in intrinsic common-agency games under quite general conditions (general type spaces, action sets, and preferences) but they did not characterize equilibrium strategies and allocations. This paper goes beyond existence and

self-generating objective function is found by aggregating the principals' virtual programs, taking the aggregate contract as given. We demonstrate that the solutions to the SGM program in Proposition 1 are those incentive-compatible allocations of output and agent utility that maximize an objective function, which, in turn, depends upon an aggregate tariff that implements the given output–utility pair. Our main characterization result (Proposition 2) introduces an assumption on the bilinearity of the agent's preferences to reduce the SGM program to a remarkably simple, pointwise optimization program over the set of equilibrium actions. Characterizing equilibria with SGM programs embeds the fixed-point nature of equilibrium but it also imports the tractability and techniques found in solving simpler optimization problems in lower dimensions.

From an economic viewpoint, SGM problems look like the problem that n cooperating principals would face. There are two key differences, however. First, as already noted, in the SGM program, reductions in the agent's rent are weighted n times more than in the cooperative program. It is as if there is a surrogate principal that maximizes a payoff that is biased toward overharvesting the agent's information rent. This n-fold excess weighting captures the fact that, in the non-cooperative scenario, each principal attempts to extract the agent's information rent without consideration of the distortionary costs imposed on the other principals' payoffs. Second, unlike the cooperative program, nondifferentiabilities in the equilibrium aggregate tariff appear in the SGM program and can be self-enforcing in an equilibrium. As we will see, the lack of smoothness is the source of equilibrium multiplicity.

From a technical viewpoint, the fact that there is a single optimization problem (what we refer to as the "surrogate" problem) that summarizes equilibrium behavior (and not a collection of n different optimization problems, one for each principal) allows us to derive important properties of the surrogate principal's value function (e.g., absolute continuity, envelope condition). These properties, in turn, help us characterize equilibrium output with only minimal regularity conditions on the set of available contracts (i.e., upper semicontinuity). In particular, we do not impose differentiability of the tariffs at the outset. Indeed, making such an assumption a priori would prevent us from exhibiting nondifferentiable equilibria even though, as we will see below, there is a plethora of such equilibria, including some with attractive welfare properties.

Maximal equilibrium

Following our characterization of the equilibrium set, we focus on a single equilibrium—the *maximal equilibrium*—that we will see, forms a basis for all equilibria. We define and construct the maximal equilibrium as the allocation in which the agent's equilibrium choice set is at its largest (i.e., maximal) and the optimization program is unconstrained (Proposition 1). In regular environments characterized by a monotone hazard rate of the types distribution, the maximal equilibrium features an n-fold asymmetric

describes all equilibrium allocations in more structured environments than those analyzed in Martimort and Stole (2012). To get sharp predictions, we assume that the agent's preferences are bilinear in output and type. This allows us to import powerful tools from convex analysis and duality at minimal cost for exposition.

information distortion. This allocation is remarkable for at least two reasons. First, it has been the implicit focus of all applied research in public screening environments to date (Laffont and Tirole 1993, Chapter 17, Martimort and Semenov 2008, Martimort and Stole 2009a, 2009b, among others). Second, we establish in Proposition 4 that this equilibrium allocation is the unique equilibrium allocation that is supported with biconjugate (i.e., least-concave) tariffs: the maximal equilibrium is smooth, its choice allocation is continuous, and each principal offers a continuous, concave tariff.

Complete set of equilibria

Using the properties of the SGM program allows us to characterize specific features of *any* equilibrium. Because Proposition 2 establishes that an equilibrium is completely identified with the range of equilibrium choices made by the various types of agents, we can view the set of equilibrium allocations as a surrogate principal's optimization program given an equilibrium set of outputs. In this sense, the optimization program shares techniques that were recently developed in the mechanism design literature on delegated decision-making.[7] This literature has shown that the solutions to the optimal delegation problem are allocations that are either independent of the privately informed party's information or that instead correspond to the individuals' ideal point. This important insight carries over into our setting. At points where the equilibrium output is continuous and separating, the allocation can be identified with the *maximal* equilibrium allocation in Proposition 1. Elsewhere, equilibrium allocations entail discontinuities and bunching (Proposition 3). It is worth noting that the discontinuous equilibria can be constructed from the maximal allocation by introducing gaps in the range of equilibrium outputs. These allocations exhibit bunching, even in regular environments satisfying the monotone hazard rate condition, and discontinuities at points where the *surrogate surplus* nevertheless remains constant. Tariffs in these equilibria are not biconjugate. They entail large negative payments over the discontinuity gaps. These punishments prevent not only the agent from choosing outputs in the gap, but also the principals from deviating with contracts that would induce the agent to choose outputs in the gap. Imposing biconjugacy rules out the implicit coordination that principals might reach by specifying offers with infinitely negative payments. Biconjugacy makes more deviations attractive. It acts thus as a refinement of the equilibrium set.

Welfare comparisons

We conclude our analysis with a consideration of the welfare properties of the equilibrium set, viewed from both the agent's and the principals' perspectives. In Proposition 5, we show that the agent prefers the maximal equilibrium allocation to all neighboring, discontinuous equilibrium outcomes.

The preferred equilibrium of the principals (i.e., what they would choose collectively to maximize the sum of their payoffs) may differ from the maximal equilibrium. Because

[7] See Holmström (1984), Melumad and Shibano (1991), Martimort and Semenov (2006), Alonso and Matouschek (2008), and Amador and Bagwell (2013), among others.

equilibria must be solutions to the SGM program, we can determine the principals' optimal equilibrium as the solution to a delegation game in which the principals delegate to the surrogate (with distorted preferences) to select an equilibrium on their behalf. Indeed, everything happens as if non-cooperating principals were jointly delegating to their surrogate representative the decision to choose an output for each possible realization of the agent's type. Of course, the difference in objectives between the principals acting collectively and their fictional surrogate captures the loss due to non-cooperative behavior. Viewed as a delegation design problem, we draw on recent advances in the delegation literature (Amador and Bagwell 2013) to determine the optimal equilibrium. In Proposition 6, we demonstrate that the principals' preferred equilibrium generally differs from the maximal equilibrium, inducing a "floor" on outputs (by means of sufficiently large punishments for outputs below this floor), which prevents excessive rent extraction. Intuitively, by refusing to pay the agent for outputs that are too low, principals reach a minimal amount of coordination and attenuate their incentives to overharvest the agent's information rent. The optimal region of bunching trades off the principals' desire to constrain inefficient, n-fold marginalizations of the agent's rent against the inefficiency of pooling agent types.

Literature review

Existing characterization results for common agency models are quite fragmented and cover various contracting scenarios. Assuming symmetric information and *delegated common agency* with public contracts, Bernheim and Whinston (1986) and Laussel and Le Breton (1998, 2001) described payoffs for the so-called *truthful equilibria*, while Chiesa and Denicolò (2009) investigated the case of private contracts. The former authors focus on *truthful tariffs* with one justification being that they are coalition-proof as proved in Bernheim and Whinston (1986) to ensure that each principal's contribution exactly reflects his preferences over possible alternatives. Efficiency follows. The only remaining question is how the possibility that the agent may reject some offer redistributes surplus among players.

Under asymmetric information, the distributions of equilibrium payoffs can no longer be disentangled from the allocative distortions that arise at equilibrium. Martimort and Stole (2015) present necessary conditions that are satisfied by all equilibrium outcomes of a delegated public common agency game. Compared with the intrinsic counterpart, delegated public common agency games allow the agent to refuse any strict subset of the principals' offers if he wishes so. Martimort and Stole (2015) derive *maximal equilibria* in those contexts and observe that they differ from maximal equilibria in intrinsic games because these additional strategic possibilities require that tariffs must remain nonnegative. The necessary conditions in Martimort and Stole (2015) remain compact enough to describe both continuous and discontinuous equilibrium allocations just as in the intrinsic scenario that is our focus hereafter. Yet, these conditions differ from those presented below because there are fewer deviations available under delegated common agency. Moreover, these necessary conditions are not sufficient; sufficiency has to be checked directly in contrast with the analysis of intrinsic games developed hereafter where sufficiency is immediate. The difference comes from the fact that intrinsic games are *bijective* aggregate games in the parlance of Martimort and Stole (2012). Knowing the solution to the self-generating problem is enough to recover solutions to all principals' optimization problems. Delegated agency games do not satisfy bijectivity since the possibility of rejecting any offer implies that contracts are necessarily nonnegative. This makes it impossible for principals to undo all aggregate offers, while undoing aggregate tariffs is always feasible under intrinsic common agency.

Among others, Stole (1991), Martimort (1992), Martimort and Stole (2009a), and Calzolari and Denicolò (2013) for private contracting and Laffont and Tirole (1993, Chapter 17), Laussel and Le Breton (1998), Martimort and Semenov (2008), Martimort and Stole (2009b) and Hoernig and Valletti (2011) for public contracting describe various differentiable equilibria that arise under asymmetric information in intrinsic common agency games with a continuum of types. None of these papers investigates the full set of equilibria as we do here. This step is possible by building on techniques similar to those in Martimort and Stole (2015) but now specialized to intrinsic common agency games. Laussel and Resende (2018) also tackle this problem in the specific context of competing manufacturers. Beside other technical differences, their approach in characterizing equilibrium allocations proceeds by deriving necessary conditions based on individual best responses that are stricter than ours. This leaves aside the issue of whether the allocations so found are indeed equilibria. Necessary and sufficient conditions are obtained altogether with our approach based on viewing equilibria as solutions to self-generating problems for bijective aggregate common agency games. Moreover, our approach allows us to directly identify equilibrium output profiles with implementable allocations of a simple mechanism design problem of delegated decision-making. This allows us to leverage valuable tools from this literature, first to describe all equilibrium allocations and second to find the best one from the principals viewpoint.

Organization

Section 2 presents the model. Section 3 describes the set of incentive-feasible allocations. We present there some of the duality tools that are used throughout the paper, defining in particular the notion of *biconjugacy*. We also briefly review the cooperative benchmark. Section 4 presents the *self-generating* optimization problems that represent equilibria. Section 5 characterizes those equilibria. Section 6 discusses equilibrium selection and welfare. Proofs are relegated to an Appendix.

2. An intrinsic common agency game

The focus of this paper is on common agency games with $n > 1$ principals (indexed by $i \in \{1, \dots, n\}$), each of whom contracts with a single common agent. We assume that common agency is *intrinsic* and the choice variable of the agent is *public* (i.e., commonly observable and contractible by all principals). For some of the interpretations below, it is useful to think of the agent as producing a good or service on behalf of the principals.

Preferences

All principals and the agent have quasi-linear preferences over output q and payments t_i that are defined, respectively, as

$$S_i(q) - t_i \quad \text{(principal } i) \ \forall i \in \{1, \dots, n\} \quad \text{and} \quad \sum_{i=1}^{n} t_i - \theta q \quad \text{(common agent)}.$$

The agent produces a good on behalf of the principals by selecting a q from an interval of feasible outputs, $\mathcal{Q} = [0, q_{\max}]$. We assume that the payoff functions S_i (for $i \in \{1, \dots, n\}$) are strictly concave and twice continuously differentiable over $\mathcal{Q}$, and that nonparticipation by the agent is equivalent to a choice of $q = 0$. Without loss of generality, we normalize principal payoffs so that $S_i(0) = 0$. Thus, $S_i(q)$ represents the net utility of principal i relative to the outside option of $q = 0$.[8] We denote the aggregate principals' payoffs by $S(q) = \sum_{i=1}^{n} S_i(q)$ and the aggregate payoff for all principals except i by $S_{-i}(q) = \sum_{j \neq i} S_j(q)$.

Contracts

Using the delegation principle, Martimort and Stole (2002) demonstrate that there is no loss of generality in studying pure-strategy common agency equilibria to require that principals' strategy spaces are restricted to tariffs from output to transfers. Let $\mathcal{T}$ be the set of all upper semicontinuous mappings, T_i, from $\mathcal{Q}$ into $\mathbb{R}$ (for $i \in \{1, \dots, n\}$). We denote an arbitrary array of contracts by $\mathbf{T} = (T_1, \dots, T_n) \in \mathcal{T}^n$.[9] An *aggregate contract* (or, in short, an *aggregate*) is defined as $T(q) = \sum_{i=1}^{n} T_i(q)$. We also use the notation $\mathbf{T}_{-i}$ and $T_{-i}(q) = \sum_{j \neq i} T_j(q)$ to denote, respectively, an array of contracts and the aggregate contract from all principals but i.

Timing and information

The timing is typical of principal–agent screening games, but now with n principals contracting instead of one. First, the agent privately learns his type (a cost parameter), θ,

[8]That the agent's utility function is bilinear in θ and q allows us to import many direct results from duality theory from convex analysis (for instance, our notion of biconjugacy below). These findings could be generalized to preferences for the agent of the sort $t_i + u(\theta, q)$ for some u function. The relevant generalization of convexity is u-convexity as discussed in Carlier (2001) and Basov (2005, Chapter 3). We can also easily generalize the agent's preferences to allow for the addition of a nonlinear function of q. Specifically, suppose that the agent's payoff is $\sum_{i=1}^{n} t_i + S_0(q) - \theta q$, where S_0 is a concave function normalized at $S_0(0) = 0$ that represents the agent's intrinsic benefit of production. Redefine payments from each principal and their respective payoff functions so that $\tilde{t}_i = t_i - \frac{S_0(q)}{n}$ and $\widetilde{S}_i(q) = S_i(q) + \frac{S_0(q)}{n}$ (for $i \in \{1, \dots, n\}$). One can verify that $\widetilde{S}_i(0) = 0$ and the expressions for the principals' and the agent's utility functions can be written, respectively, as $\tilde{S}_i(q) - \tilde{t}_i$ and $\sum_{i=1}^{n} \tilde{t}_i - \theta q$, which is the simpler form that we have adopted.

[9]We do not consider stochastic payment schedules because they have no value in our context with risk neutral players. Any stochastic payment schedule that would offer a lottery over payments for a given value of the agent's output could be replaced by the corresponding expected payment without changing payoffs and incentives. Also, we do not consider the possibility of writing contracts on contracts as in Szentes (2015). In some contexts (regulatory environments or competition in nonlinear pricing), such referencing of contracts is ruled out by institutional constraints.

that is distributed over the support $\Theta = [\theta_L, \theta_H]$ according to a continuous, commonly known distribution $F(\theta)$, with corresponding positive density $f(\theta)$. Let $\mathrm{E}_\theta[\cdot]$ denote the expectation operator with respect to the distribution of types.

Second, principals simultaneously offer the agent the tariffs, $T_i : \mathcal{Q} \to \mathbb{R}$, which are promises to pay $T_i(q)$ to the agent following the choice of $q \in \mathcal{Q} = [0, q_{\max}]$. Our assumption that common agency is *public* is captured by the fact that all principals contract on the same observed choice by the agent.

Third, the agent either accepts or rejects all of the principals' offers (i.e., common agency is *intrinsic*).[10] Refusing to participate results in zero transfers and a reservation payoff of zero to all players. Formally, we denote the agent's participation decision by the strategy δ, where $\delta = 1$ indicates acceptance and $\delta = 0$ indicates rejection. Thus, the agent's strategy is a pair of functions, $\{\delta, q\}$, mapping agent types and principal contract offers into $\{0, 1\} \times \mathcal{Q}$. If all contracts are accepted ($\delta = 1$), the agent then chooses $q \in \mathcal{Q}$ to maximize his utility and receives payments from each principal according to their contractual offers. Upper semicontinuity, together with the compactness of $\mathcal{Q}$, ensures that an optimal output exists. If, however, contracts are rejected ($\delta = 0$), then by default $q = 0$, all transfers are zero, and each player earns a normalized payoff of 0.

Equilibrium

Our focus in this paper is on equilibrium allocations that arise in a pure-strategy Perfect Bayesian equilibrium.

Definition 1. An equilibrium is an $n + 2$-tuple $\{\overline{T}_1, \dots, \overline{T}_n, \overline{q}_0, \overline{\delta}_0\}$ (with aggregate $\overline{T}(q) = \sum_{i=1}^n \overline{T}_i(q)$) such that the following statements hold:

(i) The functions $\overline{q}_0(\theta, \mathbf{T})$ and $\overline{\delta}_0(\theta, \mathbf{T})$ jointly maximize the agent's payoff:

$$\{\overline{q}_0(\theta, \mathbf{T}), \overline{\delta}_0(\theta, \mathbf{T})\} \in \underset{q \in \mathcal{Q}, \delta \in \{0,1\}}{\arg\max}\ \delta T(q) - \theta q \quad \forall \theta \in \Theta, \forall T \in \mathcal{T}.$$

(ii) The tariff $\overline{T}_i$ maximizes principal i's expected payoff given the other principals' contracts $\overline{\mathbf{T}}_{-i}$:

$$\overline{T}_i \in \underset{T_i \in \mathcal{T}}{\arg\max}\, \mathrm{E}_\theta\big[S_i\big(\overline{q}_0(\theta, T_i, \overline{\mathbf{T}}_{-i})\big) - \overline{\delta}_0(\theta, T_i, \overline{\mathbf{T}}_{-i}))T_i\big(\overline{q}_0(\theta, T_i, \overline{\mathbf{T}}_{-i})\big)\big] \quad \forall \theta \in \Theta.$$

For any equilibrium, $\{\overline{\mathbf{T}}, \overline{q}_0, \overline{\delta}_0\}$, we define the associated equilibrium allocation as the triplet $\overline{\delta}(\theta) = \overline{\delta}_0(\theta, \overline{\mathbf{T}})$, $\overline{q}(\theta) = \overline{q}_0(\theta, \overline{\mathbf{T}})$, and $\overline{U}(\theta) = \delta(\theta)\overline{T}(\overline{q}(\theta)) - \theta\overline{q}(\theta)$.

In what follows, it will be useful to refer to the set of type-allocation mappings that are implementable for some aggregate tariff, denoted $\mathcal{I}$, and to the subset of those type-allocation mappings that arise in some equilibrium, denoted $\mathcal{I}^{\text{eq}}$.

[10]Partial participation is not an option contrary to the scenario studied in Martimort and Stole (2015).

Definition 2. An allocation (U, q, δ), $U : \Theta \to \mathbb{R}$, $q : \Theta \to \mathcal{Q}$, and $\delta : \Theta \to \{0, 1\}$, is implementable if there is an aggregate tariff, $T : \mathcal{Q} \to \mathbb{R}$ such that

$$\big(q(\theta), \delta(\theta)\big) \in \underset{q\in\mathcal{Q},\delta\in\{0,1\}}{\arg\max} \ \delta T(q) - \theta q,$$

$$U(\theta) = \max_{q\in\mathcal{Q},\delta\in\{0,1\}} \delta T(q) - \theta q.$$

The set of all implementable allocations is denoted $\mathcal{I}$.

An allocation $(\overline{U}, \overline{q}, \overline{\delta})$ is an equilibrium allocation, or equilibrium implementable, if it is implementable by an aggregate tariff, $\overline{T}$, that arises at an equilibrium.

Proposition 1 below shows that, for any equilibrium allocation $(\overline{U}, \overline{q}, \overline{\delta})$, there exists another equilibrium allocation in which the agent always participates, $\delta(\theta) = 1$ for all θ. For this reason, we subsequently focus our attention on the pair $(\overline{U}, \overline{q})$ and suppress the type-allocation mapping for δ. Yet, our more general formulation remains useful for two reasons. First, it allows us to incorporate the agent's decision to participate as a requirement of implementability.[11] Second, it accounts for the possibility of equilibria where principals make nonserious offers. Indeed, there always exist uninteresting, trivial equilibria induced by a coordination failure in which two or more principals require sufficiently negative payments for each $q \in \mathcal{Q}$ so that it is not profitable for any principal to induce agent participation and $\delta = 0$ for such equilibrium allocations.

Full information allocation

The first-best allocation $(U^{\text{fb}}, q^{\text{fb}})$ is obtained when principals cooperate and know the agent's cost parameter. In this scenario, principals jointly request production at the first-best level, $q^{\text{fb}}(\theta)$, and set transfers that extract the agent's surplus. Assuming $S'(0) \geq \theta_H$ and $S'(q_{\max}) \leq \theta_L$ to avoid corner solutions,[12] we obtain

$$S'\big(q^{\text{fb}}(\theta)\big) = \theta \quad \text{and} \quad U^{\text{fb}}(\theta) = 0 \quad \forall \theta \in \Theta.$$ [13]

3. Implementability, duality and cooperative benchmark

The following lemma provides a standard characterization of the set of implementable allocations by means of familiar incentive and participation constraints.

[11] This is item (iii) in Lemma 1 below. This requirement was left implicit in the description of incentive-feasible allocations given in Martimort and Stole (2012). We find it useful to make this requirement explicit for completeness and clarity.

[12] In the sequel, we assume that the second condition (which prevents a corner at the upper bound of the feasible outputs) always holds. The first condition will be sometimes modified below to account for the fact that output remains positive under stringent conditions when there is asymmetric information. In this case, we will be explicit about such modification.

[13] This outcome is also one possible equilibrium of the intrinsic common agency game when it takes place under complete information. In sharp contrast to the analysis under asymmetric information that will follow, the principals' non-cooperative behavior need not entail any welfare loss. However, many other inefficient equilibria exist.

LEMMA 1. *An allocation* (U, q) *belongs to* $\mathcal{I}$ *if and only if*

(i) $U(\theta)$ *is absolutely continuous and differentiable almost everywhere with*

$$\dot{U}(\theta) = -q(\theta) \tag{1}$$

(ii) $U(\theta)$ *is convex (equivalently,* $q(\theta)$ *is nonincreasing)*

(iii) agent participation is optimal for all types:

$$U(\theta) \geq 0 \quad \forall \theta \in \Theta.$$

Tariffs

Once given an allocation with convex U satisfying (1), simple duality arguments allow us to recover the expression of a nonlinear tariff that implements this allocation. As a first step, we observe that any aggregate contract $T \in \mathcal{T}$ that implements an allocation (U, q) satisfies the inequality

$$U(\theta) = T\big(q(\theta)\big) - \theta q(\theta) \geq T(q) - \theta q \quad \forall q \in \mathcal{Q}, \forall \theta \in \Theta.$$

Equivalently,

$$T(q) \leq U(\theta) + \theta q \quad \forall q \in \mathcal{Q}, \forall \theta \in \Theta$$

with equality at $q = q(\theta)$. From this, we immediately obtain an upper bound $T^*(q)$ on all implementing contracts as

$$T^*(q) = \min_{\theta \in \Theta} U(\theta) + \theta q \quad \forall q \in \mathcal{Q}. \tag{2}$$

In fact, T^* is the least-concave upper semicontinuous tariff implementing (U, q) and thus

$$U(\theta) = \max_{q \in \mathcal{Q}} T^*(q) - \theta q \quad \forall \theta \in \Theta. \tag{3}$$

Using the language of convex analysis, the dual conditions (2) and (3) show that U and T^* are *conjugate* functions. Because T^* is a minimum of linear functions, it is itself concave.

Since the high-cost type's participation constraint is binding—a property that holds both in the common agency equilibria explored below and when principals cooperate—we have $U(\theta_H) = 0$. Hence, $T^*(0) = 0$ and the agent is always indifferent between accepting such offer T^* while producing zero output, and refusing to participate.

For further reference, observe also that (2) can be written by means of (3) in a more compact form that highlights the fact that U and T^* are conjugate functions:

$$T^*(q) = \min_{\theta \in \Theta} \Big\{ \max_{q' \in \mathcal{Q}} \{ T^*(q') - \theta q' \} + \theta q \Big\}.$$

Broadening the applicability of this *biconjugacy property*, we offer the following definition.

Definition 3. An aggregate contract T is *biconjugate* if and only if $T(0)=0$ and

$$T(q)=\min_{\theta\in\Theta}\Bigl\{\max_{q'\in\mathcal{Q}}\bigl\{T\bigl(q'\bigr)-\theta q'\bigr\}+\theta q\Bigr\}\quad\forall q\in\mathcal{Q}.$$

An allocation $(\overline{U},\overline{q})\in\mathcal{I}^{\mathrm{eq}}$ is a *biconjugate* equilibrium if and only if it is equilibrium implemented by an aggregate biconjugate tariff $\overline{T}$.

Remarks

Biconjugate functions are concave functions such that $T^*(0)=0$.[14] Observe that a biconjugate contract, T^*, takes finite values since $\mathcal{Q}$ is bounded. The tariff T^* is also concave over the convex hull of $q(\Theta)\cup\{0\}$ and is linear for intervals of q that lie outside $q(\Theta)$. As a comparison, consider now the following tariff T_0, taking values over the extended real line:

$$T_0(q)=\begin{cases}T^*(q) & \text{if } q\in q(\Theta)\cup\{0\},\\ -\infty & \text{otherwise.}\end{cases}\tag{4}$$

It is straightforward to check that T_0 also implements the allocation (U,q). But while T_0 inherits the concavity of T^* over all connected subsets of $q(\Theta)$, T_0 is not itself biconjugate. The tariff T_0 differs from T^* in the sense that, had the agent trembled in choosing outputs, choices outside of the equilibrium range $q(\Theta)$ would be severely punished. To illustrate, had $q(\Theta)$ taken only a finite number of values, T_0 would be a familiar *forcing contract.* Such forcing contracts are inconsistent with biconjugacy.

Cooperative outcome

Suppose that principals cooperate in designing contracts. Under asymmetric information, the optimal cooperative allocation $(\overline{U}^c,\overline{q}^c)$ is a solution to

$$(\mathcal{P}^c)\qquad\max_{(U,q)\in\mathcal{I}}E_\theta\bigl[S\bigl(q(\theta)\bigr)-\theta q(\theta)-U(\theta)\bigr].$$

The solution to this monopolistic screening problem is well known. For tractability, we assume the distribution of types satisfies the standard monotone hazard rate property.[15] Formally, we assume

$$H(\theta)=\frac{F(\theta)}{f(\theta)}$$

[14]The reader may wonder why we refer to these functions using the property of biconjugacy rather than the more evocative notion of minimally concave functions through the origin. In an earlier version of this paper, we explored common-agency games with discrete type spaces. Equilibrium contracts in this setting are equivalent to menus with finite output–tariff pairs and so the appropriate notion of concavity over the domain $\mathcal{Q}$ is unclear without more details. Biconjugacy provides the exact notion of concavity that is required for analogous results in the discrete type setting.

[15]See Bagnoli and Bergstrom (2005).

is nondecreasing and differentiable for all $\theta \in \Theta$. Equipped with this regularity condition, we state the well known characterization of the cooperative solution. The cooperative output satisfies

$$\begin{cases} S'(q^c(\theta)) = \theta + H(\theta) & \text{if } S'(0) \geq \theta + H(\theta), \\ q^c(\theta) = 0 & \text{otherwise.} \end{cases}$$

The agent's corresponding rent profile for the cooperative setting is

$$U^c(\theta) = \int_\theta^{\theta_H} q^c(\tilde{\theta})\, d\tilde{\theta}.$$

The monotone hazard condition ensures that q^c is everywhere nonincreasing (and hence U^c is convex) as required by Lemma 1. The expression of the least-concave nonlinear tariff T^c that implements this cooperative allocation is easily recovered from (2).

4. Equilibria as solutions to self-generating problems

Martimort and Stole (2012) demonstrate that intrinsic common agency games are aggregate games whose equilibria can be identified with the solution set to *self-generating* maximization problems. Specializing the necessary and sufficient conditions in their Theorem 2′ to our present setting, we obtain the following characterization of the *entire* set of equilibrium allocations as solutions of such problems.[16]

Proposition 1. *The pair* $(\overline{U}, \overline{q})$ *is an equilibrium allocation if and only if there exists an aggregate tariff* $\overline{T}$ *satisfying* $\overline{T}(0) = 0$, *which implements* $(\overline{U}, \overline{q})$ *and is such that* $(\overline{U}, \overline{q})$ *solves the self-generating maximization problem*

$$(\overline{\mathcal{P}}) \qquad \max_{(U,q)\in\mathcal{I}} E_\theta\big[S(q(\theta)) - \theta q(\theta) - nU(\theta) + (n-1)\big(\overline{T}(q(\theta)) - \theta q(\theta)\big)\big].$$

The maximization problem $(\overline{\mathcal{P}})$ bears some strong similarity with the cooperative mechanism design problem $(\mathcal{P}^c)$. The difference comes from the fact that $(\overline{\mathcal{P}})$ is now *self-generating*: its solution is implemented by an aggregate $\overline{T}$ that also appears in the maximand, embedding the fixed-point nature of equilibrium. Importantly, the fact that intrinsic common agency games are aggregate games allows a significant reduction of the difficulties faced when characterizing such fixed points. Instead of having n optimality conditions determining individual best responses for each principal, only one optimization problem remains after aggregation. This simplification allows us to import powerful techniques from optimization. In particular, this maximization problem defines a value function (the value of the maximand for each possible realization of θ) that is more regular than what we a priori imposed on aggregate tariffs. While aggregate tariffs are only restricted to be upper semicontinuous, the value function is absolutely continuous and admits a derivative almost everywhere; this is a critical step in the proof

[16]It is worth noting that this proposition does not rely on the monotone hazard rate assumption.

of Proposition 2 below where we are able to get an even more precise description of equilibrium allocations.

In contrast to $(\mathcal{P}^c)$, the objective function in $(\overline{\mathcal{P}})$ now features n-times the extraction of the agent's information rent. Each principal individually introduces distortions to extract this rent, ignoring the costs of such distortions on rivals. The n-fold term captures the resulting *tragedy of the commons* that arises with each principal overharvesting the agent's information rent. Because this strategic effect is embedded in the self-generating maximization (SGM) program, everything thus happens as if a *surrogate principal* was now in charge of maximizing the principals' collective payoff with the proviso that the agent's information rent is now weighted negatively by n. The final term in the maximand consists of $n-1$ times the agent's payoff at the induced allocation and is the source of multiple equilibria. This term has no consequence if either there is only one principal or the aggregate tariff is differentiable.[17] If, however, multiple principals offer nondifferentiable (possibly discontinuous) tariffs, then it is possible the third term generates an allocation that is consistent with these nonsmooth tariffs. For example, if multiple principals set sufficiently negative payments for some set of outputs, then it is an equilibrium for each principal to do so and the solution to the optimization program will not implement such outputs.

Necessity

The necessity part of Proposition 1 can be obtained by summing the individual optimization problems of all principals. An equilibrium allocation, since it maximizes each principal's problem, also maximizes his/her sum. This summation introduces the n-fold distortion. In any equilibrium with nonzero output, the agent's information rent will thus be overweighted by a factor of n (instead of a coefficient of 1 that would arise had principals cooperated). It is this noncooperative information-rent externality that the principals would like to mitigate in their equilibrium selection; an issue we come back to in Section 6.2 below.

Sufficiency

Establishing the sufficiency argument is more subtle. Sufficiency bears on the fact that, under intrinsic agency, the objectives of every principal are aligned in equilibrium with that of the surrogate principal who maximizes $(\overline{\mathcal{P}})$. In other words, nothing is lost by aggregating individual objectives. From a more technical point of view, sufficiency is obtained by reconstructing each principal's individual maximization problem from $(\overline{\mathcal{P}})$ itself. Doing so requires us to propose expressions of individual equilibrium tariffs that are derived from the aggregate, that solve the self-generating problem $(\overline{\mathcal{P}})$, and that are individual best responses to the tariffs other principals are offering. To this end, consider the construction of tariffs

$$\overline{T}_j(q) = S_j(q) - \frac{1}{n}\big(S(q) - \overline{T}(q)\big) \quad \forall j \in \{1, \dots, n\}. \tag{5}$$

[17] If the tariff is differentiable, incentive compatibility implies this term has zero marginal contribution in equilibrium.

Summing over j yields an aggregate of $\overline{T}$. Summing instead over all principals but i gives

$$\overline{T}_{-i}(q) = S_{-i}(q) - \frac{n-1}{n}\big(S(q) - \overline{T}(q)\big).$$

By undoing the aggregate offer $\overline{T}_{-i}$ of his rivals, principal i can always offer any aggregate T he likes, thereby inducing any implementable allocation (U, q). This construction gives principal i an expected payoff of

$$\begin{aligned}
&\mathrm{E}_\theta\big[S_i\big(q(\theta)\big) - T\big(q(\theta)\big) + \overline{T}_{-i}\big(q(\theta)\big)\big] \\
&\quad \equiv \mathrm{E}_\theta\left[S\big(q(\theta)\big) - T\big(q(\theta)\big) - \frac{n-1}{n}\big(S\big(q(\theta)\big) - \overline{T}\big(q(\theta)\big)\big)\right],
\end{aligned}$$

where the right-hand side equality follows from our previous equation for $\overline{T}_{-i}$. Expressing payments in terms of the agent's rent, we may simplify this payoff as

$$\frac{1}{n}\mathrm{E}_\theta\big[S\big(q(\theta)\big) - \theta q(\theta) - nU(\theta) + (n-1)\big(\overline{T}\big(q(\theta)\big) - \theta q(\theta)\big)\big]. \tag{6}$$

The objective function (6) exactly replicates that of the surrogate principal up to a factor $\frac{1}{n}$. Therefore, principal i's incentives to induce a particular implementable allocation (U, q) are identical to those of this representative. As a result, all principals get the same payoffs with the above construction:

$$S_i(q) - \overline{T}_i(q) = \frac{1}{n}\big(S(q) - \overline{T}(q)\big) \quad \forall i \in \{1, \dots, n\}.$$

5. Characterization of the equilibrium set

This section characterizes the complete equilibrium set. This set is large and diverse. In addition to the differentiable equilibria that have been the focus of the existing literature, there is an infinity of equilibria with discontinuous outputs or bunching of types. We provide a characterization theorem for the entire set, illuminating economic features common to all equilibria and characterizing features that are unique to particular equilibrium selections.

The solution set to the self-generating program $(\overline{\mathcal{P}})$ is difficult to characterize. Leveraging our assumption of bilinear preferences and our regularity assumption on H, however, we can simplify the SGM program, reducing it to problem in pointwise optimization. This affords us a much sharper characterization of the equilibrium allocations. The main characterization result of this paper follows:

Proposition 2. *An allocation $(\overline{U}, \overline{q})$ belongs to $\mathcal{I}^{\mathrm{eq}}$ if and only if*

$$\overline{q}(\theta) \in \underset{q \in \overline{q}(\Theta)}{\arg\max}\, S(q) - \big(\theta + nH(\theta)\big)q \quad \forall \theta \in \Theta, \tag{7}$$

$$\overline{U}(\theta) = \int_\theta^{\theta_H} \overline{q}(\tilde{\theta})\, d\tilde{\theta} \quad \forall \theta \in \Theta. \tag{8}$$

Surrogate principal's incentive constraints

Condition (7) represents a greatly simplified, pointwise SGM program that embeds the strategic interactions of the principals into a simple optimization program. Note that the domain of this program is restricted to a self-generating set of equilibrium outputs. In contrast, the self-generation in Proposition 1 was over the more complicated object of an aggregate tariff. It is as if a surrogate representative of the principals is optimizing on their behalf but with an overemphasis on rent extraction (n weight rather than unity). At any type realization θ, this surrogate principal, whose decisions reflect the non-cooperative behavior of the principals, should prefer to choose the equilibrium output $\overline{q}(\theta)$ rather than any other output that would have been chosen had any other type realized. This explains why in the maximand of (7), the domain of maximization is over $\overline{q}(\Theta)$.

To evaluate those best choices, the surrogate principal a priori considers the maximand of the self-generating problem $(\overline{\mathcal{P}})$. The first remarkable finding is that the surrogate principal's incentive constraints (7) are now written ex post instead of ex ante as in the maximand $(\overline{\mathcal{P}})$. This transformation requires to replacing the cost parameter θ by a new expression that entails an n-fold information distortion due to the principals' non-cooperative behavior, namely $\theta + nH(\theta)$. The monotone hazard rate condition that this *modified virtual cost parameter* remains nondecreasing and thus $\overline{q}$ is itself nonincreasing. Henceforth, any solution to a self-generating problem $(\overline{\mathcal{P}})$ is obtained as the solution to the relaxed problem $(\overline{\mathcal{P}}^r)$, where the convexity requirement for $\overline{U}$ can be omitted.

The second remarkable simplification incorporated into (7) is that the extra term $(n-1)(\overline{T}(q) - \theta q)$ that is found in the maximand of $(\overline{\mathcal{P}})$ has now disappeared. Intuitively, $\overline{q}(\theta)$ is also a maximizer for this last term since it has to be the agent's equilibrium choice. Although no assumption on differentiability of the equilibrium aggregate tariff $\overline{T}(q)$ is ever made, everything happens as if an envelope condition could be used to simplify the writing of the surrogate principal's incentive constraints.

The third notable fact is that, although the optimization domain in $(\overline{\mathcal{P}})$ is $\mathcal{Q}$, only outputs in $\overline{q}(\Theta)$ are used to write (7). This is so because $\overline{q}(\Theta)$ certainly differs from $\mathcal{Q}$ when $\overline{T}$ specifies sufficiently large, negative payments over $\mathcal{Q}/\overline{q}(\Theta)$.

Equilibrium allocations

The characterization of equilibrium allocations by means of the surrogate principal's incentive constraints (7) bears strong similarities to the characterization of implementable allocations found in the literature on mechanism design for delegated decision-making problems as in Holmström (1984), Melumad and Shibano (1991), Martimort and Semenov (2006), Alonso and Matouschek (2008), and Amador and Bagwell (2013). This literature demonstrates how an uninformed party can delegate decision-making to a privately informed party in circumstances of conflicting preferences, asymmetric information, and when no incentive payments are available to align objectives. In our context, the conflict of interest comes from the fact that, although principals would like to cooperate, they are unable to do so when each of them can deviate to a bilateral agreement with the agent. The non-cooperative outcome is captured by the optimizing behavior of a surrogate principal. Yet, while cooperating principals maximize a virtual surplus worth

$$S(q) - \big(\theta + H(\theta)\big)q, \tag{9}$$

the surrogate principal cares about a surrogate surplus that entails the *modified virtual cost parameter*

$$S(q) - \big(\theta + nH(\theta)\big)q. \tag{10}$$

This surrogate surplus accounts for the fact that non-cooperating principals extract n times the agent's rent while cooperating principals only care about extracting that rent once. This difference in their concerns for rent extraction is the source of conflict between cooperating principals and their surrogate representative.

Although there is no asymmetric information per se between the cooperating principals and their surrogate, the latter implements at equilibrium an allocation that is a *pointwise* optimum of the surrogate surplus. This maximization thus induces a set of incentive constraints that are reminiscent of those found in the aforementioned delegation design literature. Borrowing techniques that were developed there provides a clear characterization of equilibrium outputs. Everything happens thus as if the surrogate principals was informed on the agent's cost himself although he replaces this cost parameter by its non-cooperative virtual version.

Maximal equilibrium

Following a path taken by Martimort and Stole (2015) in their analysis of delegated common agency games, one may select among all equilibria described in Proposition 2 by considering maximization in (7) over the full domain $\mathcal{Q}$. We denote this output allocation by q^m, i.e., the unconstrained maximum of the strictly concave objective (10):

$$q^m(\theta) \in \arg\max_{q\in\mathcal{Q}} S(q) - \big(\theta + nH(\theta)\big)q \quad \forall \theta \in \Theta.$$

The resulting range of optimal outputs in this unconstrained optimization program is also a self-generating solution to the program in (7), and so we term it the *maximal* solution.[18] Thanks to the monotone hazard rate assumption, this maximal output allocation, which is characterized by

$$\begin{cases} S'\big(q^m(\theta)\big) = \theta + nH(\theta) & \text{if } S'(0) \geq \theta + nH(\theta), \\ q^m(\theta) = 0 & \text{otherwise,} \end{cases} \tag{11}$$

[18]While we use the same notion of *maximal* solutions in Martimort and Stole (2015) as in the present paper, note that the virtual surpluses in the SGM program of the delegated agency setting of Martimort and Stole (2015) and in the present paper are different. In particular, the space of contracts in the intrinsic common agency game is larger than in the delegated scenario since the latter only includes nonnegative tariffs.

is again nonincreasing. Compared with a cooperative outcome, the maximal equilibrium allocation features a distortion that is now proportional to n times the hazard rate. This captures the fact that, at equilibrium, each principal adds his own distortion for rent extraction reasons.

Define now the rent allocation U^m as

$$U^m(\theta) = \int_\theta^{\theta_H} q^m(\tilde{\theta})\, d\tilde{\theta}.$$

Given our construction of the surrogate's unconstrained optimal allocation, (U^m, q^m), it is an immediate consequence of Proposition 2 that this allocation is an equilibrium outcome in the common agency game.

Corollary 1. *The maximal allocation (U^m, q^m) is an equilibrium allocation.*

Observe that any equilibrium allocation must satisfy requirement (7), and therefore on any interval where it is continuous and separating, it must equal the maximal allocation.

This maximal equilibrium has been the focus of the earlier common agency literature. For instance, Martimort and Stole (2012) used this particular selection to prove existence of an equilibrium to intrinsic common agency games under broad conditions, though they made no attempt to characterize the properties of the equilibrium set, in sharp contrast to the present paper. In more applied work, Laffont and Tirole (1993, Chapter 17) modeled privatization as a common agency game between shareholders and regulators controlling the firm's manager. Their conclusion that joint control leads to low-powered incentives relies on the selection of the smooth maximal equilibrium or an extreme discontinuous forcing equilibrium.

More detailed characterization

The next proposition provides a complete and detailed characterization of all equilibrium output profiles.

Proposition 3 (General characterization of equilibrium output profiles). *An output allocation $\overline{q} : \Theta \to \mathcal{Q}$ is an equilibrium outcome if and only if it satisfies the following properties:*

(i) The allocation $\overline{q}$ is nonincreasing, differentiable a.e. with at most a countable number of downward-jump discontinuities.

(ii) At any point of differentiability,

$$\dot{\overline{q}}(\theta)\big(S'\big(\overline{q}(\theta)\big) - \theta - nH(\theta)\big) = 0. \tag{12}$$

(iii) At any discontinuity for an interior type, $\theta_0 \in (\theta_L, \theta_H)$, bunching arises on both sides of θ_0 satisfying

$$\overline{q}(\theta) = q^m(\theta_1) \quad \forall \theta \in [\theta_1, \theta_0) \quad \textit{and} \quad \overline{q}(\theta) = q^m(\theta_2) \quad \forall \theta \in (\theta_0, \theta_2]$$

for some θ_1 and θ_2 such that $\theta_1 < \theta_0 < \theta_2$.[19] *The surrogate surplus is continuous around θ_0:*

$$S\big(q^m(\theta_1)\big) - \big(\theta_0 + nH(\theta_0)\big)q^m(\theta_1) = S\big(q^m(\theta_2)\big) - \big(\theta_0 + nH(\theta_0)\big)q^m(\theta_2). \tag{13}$$

Equilibrium outputs

From (12), any equilibrium output profile is either flat over some range, in which case it is unresponsive to the agent's private information, or when it is decreasing and continuous in θ, it corresponds to the maximal equilibrium output. To illustrate with a scenario of some relevance for what follows, an equilibrium output profile can be obtained simply by putting a floor on the maximal equilibrium output. In that case, principals are unable to implement outputs that are too low.

Discontinuities in the equilibrium output $\overline{q}$ have also a quite specific structure. First, the fact that $\overline{q}$ is nonincreasing (from incentive compatibility) implies that such discontinuities are necessarily countable in number. Second, the fact that the surrogate principal's surplus is maximized at an optimal choice implies that such discontinuities must nonetheless preserve the continuity of the surrogate surplus:

$$\max_{q \in q(\Theta)} S(q) - \big(\theta + nH(\theta)\big)q.$$

Outputs on both sides of such discontinuities satisfy the simple condition (13), which expresses the fact that the surrogate principal should be indifferent between moving output on either side of the gap.

Equilibrium ranges

Introducing a discontinuity of $\overline{q}$ at a given type θ_0 amounts to withdrawing an interval $(q^m(\theta_2), q^m(\theta_1))$ from the range of q^m to obtain the range of $\overline{q}$. This may be done by imposing sufficiently negative payments for every $q \in (q^m(\theta_2), q^m(\theta_1))$. Clearly, each principal is willing to offer such payments if other principals are expected to do so. In this manner, we may generate arbitrary equilibrium allocations by introducing gaps in the maximal allocation. The maximal equilibrium contains the ranges of all discontinuous equilibria that are constructed by introducing gaps.[20]

Another immediate consequence of Proposition 3 is that for any subset $\overline{Q}$ of $q^m(\Theta)$ that is the union of a countable number of intervals, there is a unique equilibrium allocation $\overline{q}$ whose range is $\overline{Q}$ itself. A specific case arises when principals offer forcing contracts at a finite number of outputs. Even though they do not provide any in-depth analysis of those equilibria, Laffont and Tirole (1993, Chapter 17) already devoted an appendix to discuss an interesting subclass of equilibria that are implemented by means

[19]The output $\overline{q}$ can be made either right-continuous ($\overline{q}(\theta_0) = q^m(\theta_1)$) or left-continuous ($\overline{q}(\theta_0) = q^m(\theta_2)$) with, of course, no consequences on payoffs for both the agent and the principals.

[20]There are discontinuous equilibria that feature degenerate pooling on a single point and cannot be generated by introducing gaps into the maximal equilibrium. The simplest such equilibrium is one in which $q(\theta) = 0$ for all types, $T_i(0) = 0$ and $T_i(q) = -\infty$ for all $q \neq 0$.

of forcing contracts. Their analysis is unfortunately incomplete. Forcing contracts induce allocations with exhibit bunching almost everywhere, but they only represent a special case of the more complete analysis of Propositions 2 and 3 above. Moreover, and in sharp contrast to ours, their analysis is silent on the possible welfare comparison of those nondifferentiable equilibria to the smooth maximal allocation.

Tariffs

An aggregate payment $\overline{T}$ can easily be reconstructed from any equilibrium allocation $(\overline{U}, \overline{q})$, where $\overline{q}$ satisfies (12) outside discontinuities and (13) at any discontinuity point. First, using the rent profile $\overline{U}$ obtained from (8), the duality argument in (2) gives us a nonlinear price schedule T^* that implements $(\overline{U}, \overline{q})$. Second, principals are prevented from deviating to a contract that would induce the agent to choose outputs within a discontinuity gap by imposing that the aggregate tariff $\overline{T}$ entails infinitely negative payments for $q \notin \overline{q}(\Theta) \cup \{0\}$ as requested from (4). Finally, conditions (5) allow one to reconstruct from this aggregate tariff the equilibrium tariffs offered by each principal.

Comparison with the cooperative outcome

It is interesting to ascertain the validity of our findings in Propositions 2 and 3 in the limiting case where $n = 1$ (recall that our maintained assumption has been $n > 1$). Of course, the maximal equilibrium allocation corresponds to the cooperative solution for this case. When $n = 1$, the necessary conditions in Propositions 2 and 3 continue to hold, but the sufficient conditions are no longer valid. In fact, when $n > 1$, some of the equilibrium allocations are obtained by the very fact that a given principal may not be free to choose any output because other principals have stipulated infinitely negative payments at this output; a threat that is only available when $n > 1$.

A simple numerical example

As an illustration of the various allocations that were characterized above, let us consider the following uniform-quadratic example. Suppose that $n = 2$, $S_1(q) = S_2(q) = 4q - \frac{1}{4}q^2$, and θ is distributed uniformly on $[1, 5]$ so that $H(\theta) = \theta - 1$. It is straightforward to derive the allocations for the first-best outcome, the cooperative optimum, and the maximal noncooperative equilibrium as

$$q^{\text{fb}}(\theta) = 8 - \theta, \qquad q^c(\theta) = \max\{9 - 2\theta, 0\}, \qquad q^m(\theta) = \max\{10 - 3\theta, 0\},$$

which are illustrated in Figure 1. The range of the maximal equilibrium is $q^m(\Theta) = [0, 7]$. The (aggregate) tariff for the cooperative solution and the maximal equilibrium are, respectively, given by

$$T^c(q) = \frac{9}{2}q - \frac{q^2}{4} \quad \text{and} \quad T^m(q) = \frac{10}{3}q - \frac{q^2}{6}.$$

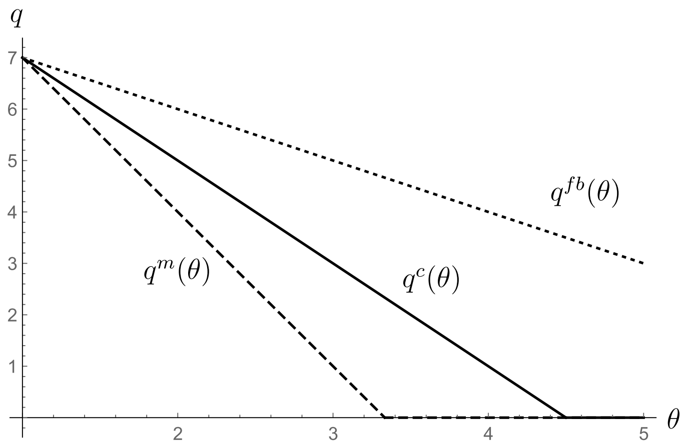

FIGURE 1. Numerical example: Output allocations for the first-best benchmark, the cooperative benchmark, and the maximal equilibrium.

Now consider taking off some outputs in that range so as to construct an equilibrium whose range is $\overline{q}(\Theta) = q^m(\Theta)/(2, 5)$. It is straightforward to check that $\overline{q}$ is discontinuous at $\theta_0 = \frac{13}{6}$, which leaves the agent indifferent between moving on either side of the discontinuity. The corresponding equilibrium output is thus

$$\overline{q}(\theta) = \begin{cases} q^m(\theta) & \text{if } \theta \in \left[1, \frac{5}{3}\right] \cup \left[\frac{8}{3}, 5\right], \\ 5 & \text{if } \theta \in \left[\frac{5}{3}, \frac{13}{6}\right], \\ 2 & \text{if } \theta \in \left(\frac{13}{6}, \frac{8}{3}\right]. \end{cases} \tag{14}$$

In this uniform-quadratic setting, we can also verify that an implementing aggregate tariff satisfies

$$\overline{T}(q) = \begin{cases} T^m(q) & \text{if } q \in q^m(\Theta)[0, 2] \cup [5, 7], \\ -\infty & \text{otherwise.} \end{cases}$$

6. Equilibrium selection

We now present two approaches to possibly select within the large set of equilibria found above. First, under reasonable assumptions, the maximal equilibrium is also the best equilibrium from the agent's viewpoint. In sharp contrast, under mild conditions the best equilibrium from the principals' viewpoint is never the maximal equilibrium.

6.1 *Maximal equilibrium, biconjugacy, and agent optimality*

The maximal equilibrium is implemented by a biconjugate tariff $T^m(q) = \min_{\theta\in\Theta} U^m(\theta) + \theta q$. To the contrary, other equilibria that feature discontinuity gaps within the range of $q^m(\Theta)$ cannot be implemented in an equilibrium with the corresponding biconjugate tariff $T^*(q) = \min_{\theta\in\Theta} \overline{U}(\theta) + \theta q$. These equilibria are instead implemented by an aggregate tariff $\overline{T}$ such that $\overline{T} < T^*$ over some range.

To see why, consider an equilibrium allocation $\overline{q}$ whose range $\overline{q}(\Theta)$ is obtained by withdrawing an interval $(q^m(\theta_2), q^m(\theta_1))$ from $q^m(\Theta)$. Let us again denote by θ_0 the type who is indifferent between $q^m(\theta_2)$ and $q^m(\theta_1)$. Suppose now that the corresponding aggregate equilibrium tariff $\overline{T}$, which entails infinitely negative payments over this discontinuity gap, is replaced with the least-concave tariff T^* that also implements $\overline{q}$. By construction, $\overline{T}$ induces a self-generating program whose solution is $(\overline{U}, \overline{q})$. However, replacing $\overline{T}$ by T^* in that self-generating program, although it changes nothing from the agent's viewpoint, may modify its solution. Since T^* entails finite payments, an output $q \in (q^m(\theta_2), q^m(\theta_1))$ may now become more attractive. More precisely, T^* is linear over $(q^m(\theta_2), q^m(\theta_1))$ and the agent with type θ_0 is actually indifferent between all options in $(q^m(\theta_2), q^m(\theta_1))$. In this case, a principal may now find it attractive to induce this type to choose an output within this discontinuity gap so that $(\overline{U}, \overline{q})$ is no longer a solution of a new self-generating program obtained with T^* replacing $\overline{T}$.

Imposing biconjugacy on the aggregate tariff thus certainly refines the equilibrium set. Indeed, it limits the possibility that principals have to collectively prevent a deviation by offering an aggregate tariff with infinitely negative payments over a discontinuity gap.

Proposition 4. *The maximal equilibrium (U^m, q^m) is the only equilibrium sustained by a biconjugate aggregate tariff.*

Geometry and welfare properties

Consider an equilibrium $(\overline{U}, \overline{q})$ whose range $\overline{q}(\Theta)$ is obtained by withdrawing an interval $(q^m(\theta_1), q^m(\theta_2))$ from $q^m(\Theta)$ and let us again denote by θ_0 the point of discontinuity for $\overline{q}$. We now want to compare the rent profiles $\overline{U}$ and U^m. (See also Figure 2.)

A first observation is that those profiles are of course identical on $[\theta_2, \theta_H]$ since $\overline{U}(\theta) = \int_\theta^{\theta_H} \overline{q}(\tilde{\theta})\, d\tilde{\theta}$, $U^m(\theta) = \int_\theta^{\theta_H} q^m(\tilde{\theta})\, d\tilde{\theta}$ and outputs are the same on that interval; $\overline{q}(\tilde{\theta}) = q^m(\theta)$ on $[\theta_2, \theta_H]$. Second, from Proposition 3, the equilibrium output $\overline{q}$ lies below (resp. above) q^m over $[\theta_0, \theta_2]$ (resp. $[\theta_1, \theta_0]$). The overall impact of those distortions on how $\overline{U}$ compares with U^m is thus ambiguous. The next example nicely illustrates how this ambiguity can be solved.

Numerical example (continued). Take the equilibrium output $\overline{q}$ defined in (14). We can check that

$$\int_{\frac{5}{3}}^{\frac{8}{3}} \overline{q}(\tilde{\theta})\, d\tilde{\theta} = \left(\frac{13}{6} - \frac{5}{3}\right) \times 5 + \left(\frac{8}{3} - \frac{13}{6}\right) \times 2 \equiv \frac{7}{2} = \int_{\frac{5}{3}}^{\frac{8}{3}} q^m(\tilde{\theta})\, d\tilde{\theta}.$$

That equality is remarkable. It means that the rent profiles $\overline{U}$ and U^m are the same not only on $[\frac{8}{3}, 5]$, but also on $[1, \frac{5}{3}]$. Over the interval $(\frac{5}{3}, \frac{8}{3})$, instead $\overline{U} < U^m$. The argument could be generalized in a straightforward manner to any possible countable number

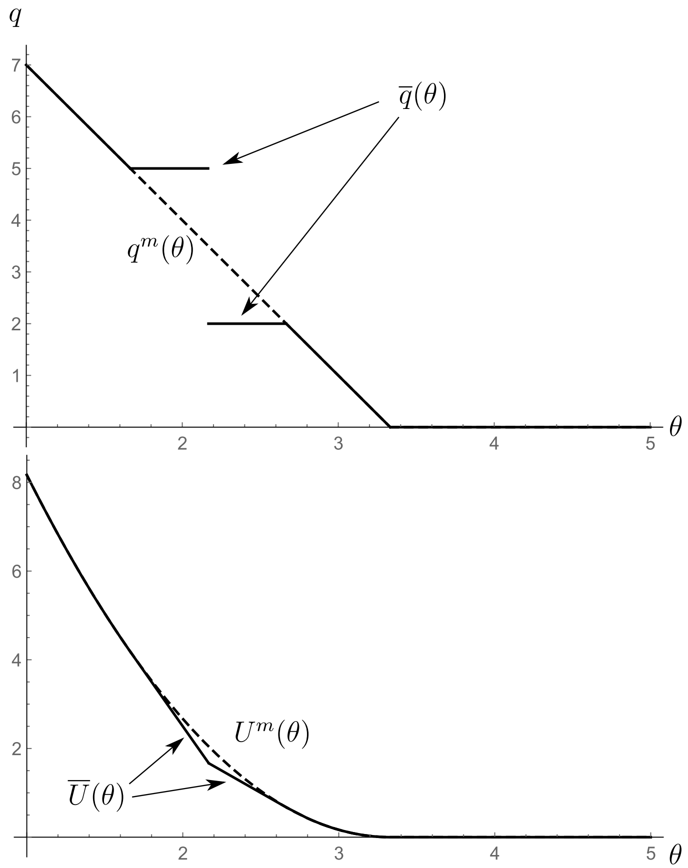

FIGURE 2. Numerical example: Output and informational rent for the maximal equilibrium (dashed) and the discontinuous equilibrium (solid).

of discontinuities and to any quadratic surplus function if the types distribution is uniform. In other words, the maximal equilibrium yields an upper bound on all equilibrium rent profiles.

The geometry of these results is interesting in itself. The graph of the rent profile $\overline{U}$ is obtained from that of U^m by replacing the strictly convex part of U^m over $[\frac{5}{3}, \frac{8}{3}]$ by the maximum of the two tangents at the points $\frac{5}{3}$ and $\frac{8}{3}$, namely $\max\{U^m(\frac{5}{3}) - 5(\theta - \frac{5}{3}); U^m(\frac{8}{3}) - 2(\theta - \frac{8}{3})\}$. At the discontinuity point θ_0, the two tangents cross, capturing the fact that θ_0 is indifferent between choosing $q^m(\frac{5}{3}) = 5$ and $q^m(\frac{8}{3}) = 2$. ◊

The fact that the aggregate surplus function is quadratic and the distribution of types is uniform plays a key role in the example above. Provided that the addition of the discontinuity gap remains small enough–the surplus function is locally quadratic and the distribution is uniform—we might thus hope to get similar striking results. To this end, consider introducing a discontinuity gap around a given point $\theta_0 \in (\theta_L, \theta_H)$ and the corresponding equilibrium $(\overline{U}, \overline{q})$. This gap again takes the form $(q^m(\theta_1(\varepsilon)), q^m(\theta_2(\varepsilon)))$, where $\theta_2(\varepsilon) = \theta_0 + \varepsilon$. We are interested in the case where ε is small enough and index accordingly the equilibrium rent and output $(\overline{U}_\varepsilon, \overline{q}_\varepsilon)$.[21]

Proposition 5. *Suppose that H is twice differentiable with $\ddot{H} \geq 0$ and that S is thrice differentiable with $S''' \leq 0$. Then, for ε sufficiently small,*

$$\overline{U}_\varepsilon(\theta) \leq U^m(\theta) \quad \forall \theta \in \Theta.$$

Moreover, if preferences are quadratic and types are uniformly distributed, the agent prefers the maximal equilibrium globally.

Under mild conditions, introducing a small discontinuity gap in the equilibrium range of outputs cannot improve the agent's rent. Using familiar duality arguments, the corresponding (least-concave) aggregate tariffs can be ranked as

$$T^*_\varepsilon(q) = \min_{\theta \in \Theta} \overline{U}_\varepsilon(\theta) + \theta q \leq T^*(q) = \min_{\theta \in \Theta} \overline{U}(\theta) + \theta q.$$

The tariff in the maximal equilibrium is thus the upper envelope of all least-concave aggregate tariffs that implement other equilibria. Going further, Proposition 5 states that the uniform-quadratic assumption is sufficient to extend the ε-discontinuity result globally.

6.2 *Principals' ex ante optimality*

Another possibility is to select among equilibria in terms of the expected net surplus they give to the principals. Indeed, if principals could meet ex ante and negotiate over the equilibrium to be played, a reasonable prediction would be that they would agree to play the equilibrium that maximizes their ex ante collective payoff. Therefore, we now investigate what this ex ante best equilibrium allocation is for the principals. In this respect, Proposition 3 shows that for any set $\overline{\mathcal{Q}}$ such that $\overline{\mathcal{Q}} \subset q^m(\Theta)$ and $\overline{\mathcal{Q}}$ is a union of countable intervals, there is a unique equilibrium allocation $\overline{q}$ such that $\overline{q}(\Theta) = \overline{\mathcal{Q}}$. This allocation fully defines the corresponding aggregate transfer as we have seen. Thus, the equilibrium selection problem can be reduced to the principals optimally choosing a delegation set $\overline{\mathcal{Q}}$ to offer to the surrogate representative, who then chooses an allocation that solves a self-generating maximization program $(\overline{\mathcal{P}})$, where the equilibrium tariff $\overline{T}$ has domain $\overline{\mathcal{Q}}$. Restated in this form, we may apply a recent result from Amador and Bagwell (2013)[22] to conclude that the optimal delegation set is a connected interval that puts a floor on outputs. To do so, it is sufficient to make the following assumption on the types distribution.

Assumption 1. *For almost all $\theta \in \Theta$,*

$$\big(n\dot{H}(\theta) + 1\big)\big(\dot{H}(\theta) + 1\big) \geq (n-1)H(\theta)\ddot{H}(\theta).$$

[21]Although Proposition 5 deals with the case of a single discontinuity gap, it is straightforward to again generalize these findings to account for more discontinuities.

[22]See Martimort and Semenov (2006) and Alonso and Matouschek (2008) for earlier, slightly stronger conditions along those lines.

In tandem with our assumption of a monotone hazard rate, Assumption 1 requires that H is not too convex. It is satisfied by several well known distributions with positive supports, including uniform, exponential, Pareto, inverse Weibull, and chi-square.[23]

Proposition 6. *Suppose that Assumption 1 holds. The principals' best equilibrium from an ex ante point of view is characterized by an output interval,* $\overline{\mathcal{Q}} = [q^m(\hat{\theta}), q^m(\theta_L)]$, *with* $\hat{\theta} > \theta_L$ *and*

$$\mathrm{E}_\theta\big[S'\big(q^m(\hat{\theta})\big) - \theta - H(\theta) | \theta \geq \hat{\theta}\big] = 0. \tag{15}$$

When $q^m(\theta_H) > 0$, *we have* $\hat{\theta} < \theta_H$ *and the maximal equilibrium is never ex ante optimal.*

Proposition 6 tells us that the range of the maximal equilibrium, $q^m(\Theta)$, generates too much informational-rent distortion from an ex ante point of view for the principals. They would prefer to put a floor on the set of equilibrium outputs, thereby reducing the sensitivity of the output allocation on the agent's underlying type, so as to mitigate the problem of excessive rent extraction. Such a floor can be implemented by an aggregate tariff that sets $\overline{T}(q) = -\infty$ for any $q \in (0, q^m(\hat{\theta}))$.

Numerical example (continued). Using the definition of q^m given in (11), condition (15) can be written in terms of the cutoff $\hat{\theta}$ only as

$$E_\theta(2\theta - 1 | \theta \geq \hat{\theta}) = \min\{8, 3\hat{\theta} - 2\}.$$

Tedious computations show that the cutoff is $\hat{\theta} = 3$. The output profile at the *ex ante* best non-cooperative equilibrium is a truncation of the maximal equilibrium profile:

$$\overline{q}(\theta) = \max\{1, 10 - 3\theta\}.$$

At that best equilibrium, the principals shut down payments for q less than $\hat{q} = q^m(3) = 1$. Bunching arises over the upper tail of the distribution $[3, 5]$:

$$\overline{T}(q) = \begin{cases} 0 & \text{if } q = 0, \\ \dfrac{11}{6} + \dfrac{10}{3}q - \dfrac{q^2}{6} & \text{if } q \in [1, 7], \\ -\infty & \text{otherwise.} \end{cases}$$

Observe also that $\hat{q} > q^m(5) = 0$, so that by restricting the equilibrium set of outputs, principals are able to implement outputs that are sometimes above the cooperative solution.

[23]This property is verified for the chi-square distributions using numerical methods for degrees of freedom $d \neq 2$.

Appendix: Proofs

The proof of Lemma 1 is standard and thus omitted. See Rochet (1987) or Milgrom and Segal (2002).

Proof of Proposition 1. The proof follows similar steps to those in Martimort and Stole (2012, Theorem 2′), though for completeness, we explicitly treat here the agent's participation decision δ, which was left implicit in this paper.

Necessity. Given the aggregate tariff $\overline{T}_{-i}$ offered by competing principals, principal i's net gain with the agent of type θ when $(q, \delta) \in \mathcal{Q} \times \{0, 1\}$ is chosen is given by

$$S_i(q) - \delta\overline{T}(q) \equiv S_i(q) - \theta q + \delta\overline{T}_{-i}(q) - \big(\delta\overline{T}(q) - \theta q\big).$$

For $(\overline{U}, \overline{q}, \overline{\delta})$ to be an equilibrium allocation, it must be that principal i desires to implement this allocation, which must thus solve

$$(\overline{U}, \overline{q}, \overline{\delta}) \in \underset{(U,q,\delta)\in\mathcal{I}}{\arg\max}\, E_\theta\big[S_i\big(q(\theta)\big) - \theta q(\theta) + \delta(\theta)\overline{T}_{-i}\big(q(\theta)\big) - U(\theta)\big].$$

Note that every principal i faces the same domain of maximization, $\mathcal{I}$. The difference between the programs of any two principals, i and j, is entirely embedded in the differences in the aggregates $\overline{T}_{-i}$ and $\overline{T}_{-j}$. Following Martimort and Stole's (2012) analysis of general aggregate games, an equilibrium allocation must necessarily maximize the sum of the principals' programs. Thus, $(\overline{U}, \overline{q}, \overline{\delta})$ must solve

$$(\overline{U}, \overline{q}, \overline{\delta}) \in \underset{(U,q,\delta)\in\mathcal{I}}{\arg\max}\, E_\theta\big[S\big(q(\theta)\big) - \theta q(\theta) - nU(\theta) + (n-1)\big(\delta(\theta)\overline{T}\big(q(\theta)\big) - \theta q(\theta)\big)\big]. \quad (16)$$

The solution to this problem always has at least one type, say $\hat{\theta}$, such that

$$\overline{U}(\hat{\theta}) = 0. \quad (17)$$

Indeed, if it were not the case, then the whole rent profile $\overline{U}$ could be reduced uniformly by $\epsilon > 0$ without modifying output, and this modification would improve the value of the program. This means that, at least one principal would have an incentive to deviate by uniformly reducing his payment to the agent by ϵ.

From a remark in the text, any aggregate tariff that implements an equilibrium $(\overline{U}, \overline{q})$ solution to the SGM problem above must satisfy

$$T(q) \le T^*(q) = \min_{\theta\in\Theta} \overline{U}(\theta) + \theta q$$

with equality at all $q \in \overline{q}(\Theta)$. In particular, this condition should hold for the equilibrium aggregate tariff $\overline{T}$. From (17), it follows that $T^*(0) = \min_{\theta\in\Theta} \overline{U}(\theta) = 0$ and, thus,

$$\overline{T}(0) \le 0. \quad (18)$$

Consider now the new tariff $\tilde{T}$ obtained from $\overline{T}$ as

$$\tilde{T}(q) = \begin{cases} 0 & \text{if } q = 0, \\ \overline{T}(q) & \text{otherwise.} \end{cases}$$

Because of (18), $\tilde{T}$ is itself upper semicontinuous. Moreover, we have

$$\overline{T}(q) \leq \tilde{T}(q) \leq T^*(q) \quad \forall q \in \mathcal{Q}.$$

Because both $\overline{T}$ and T^* implement $(\overline{U}, \overline{q}, \overline{\delta})$, we deduce from those inequalities that $\tilde{T}$ also does so. Under $\tilde{T}$, every agent type chooses to participate, $\tilde{\delta}(\theta) = 1$, because he always has the option to choose $q = 0$ and thereby get his reservation payoff that is normalized at zero. Because the objective function in (16) has the same expected value at $(\overline{U}, \overline{q}, \overline{\delta})$ using $\overline{T}$ as it does at $(\overline{U}, \overline{q}, \tilde{\delta})$ using $\tilde{T}$, we conclude that

$$(\overline{U}, \overline{q}) \in \underset{(U,q)\in\mathcal{I}}{\arg\max}\, E_\theta\big[S\big(q(\theta)\big) - \theta q(\theta) - nU(\theta) + (n-1)\big(\tilde{T}\big(q(\theta)\big) - \theta q(\theta)\big)\big].$$

This new problem is again self-generating and $(\overline{U}, \overline{q})$ is equilibrium implemented by $\tilde{T}$.

Sufficiency. Consider a solution $(\overline{U}, \overline{q})$ to $(\overline{\mathcal{P}})$ that is implemented by the aggregate tariff $\overline{T}$. Note that because $(\overline{U}, \overline{q})$ is implemented by $\overline{T}$ with $\overline{T}(0) = 0$, we are considering the case where the agent always participates, $\delta = 1$. Let us now construct individual tariffs $\overline{T}_i$ satisfying

$$\overline{T}_i(q) = S_i(q) - \frac{1}{n}\big(S(q) - \overline{T}(q)\big) \quad \forall i \in \{1, \ldots, n\}.$$

By construction,

$$\sum_{i=1}^{n} \overline{T}_i(q) = \overline{T}(q).$$

We show that this contract profile $(\overline{T}_1, \ldots, \overline{T}_n)$ is an equilibrium. Suppose indeed that all principals j for $j \neq i$ offer $\overline{T}_j$. At a best response, principal i induces an allocation (U, q, δ) that solves

$$(\mathcal{P}_i) \qquad \max_{(U,q,\delta)\in\mathcal{I}} E_\theta\big[S_i\big(q(\theta)\big) - \theta q(\theta) - U(\theta) + \overline{T}_{-i}\big(q(\theta)\big)\big].$$

Inserting the expressions of $\overline{T}_j$ (for $j \neq i$) using our construction above, the allocation that principal i would like to induce should solve

$$\max_{(U,q,\delta)\in\mathcal{I}} E_\theta\big[S\big(q(\theta)\big) - \theta q(\theta) - nU(\theta) + (n-1)\big(\overline{T}\big(q(\theta)\big) - \theta q\big)\big].$$

But this is the same maximization program in (16), and hence principal i's choice $\overline{T}_i$ is a best response against $\overline{T}_{-i}$. □

Proof of Proposition 2. *Necessity.* Because of (1), any implementable rent profile is nonincreasing. It follows that, necessarily, for any solution to $(\overline{\mathcal{P}})$ where the agent's rent is minimized, we must have

$$\overline{U}(\theta_H) = 0. \tag{19}$$

From (1) and (19), we thus get (8). Inserting this expression of the rent into the maximand of $(\overline{\mathcal{P}})$ and integrating by parts shows that any solution to the relaxed problem

$(\overline{\mathcal{P}}^r)$ obtained when the convexity requirement on U has been omitted should also solve pointwise the problems

$$\overline{q}(\theta) \in \arg\max_{q \in \mathcal{Q}} S(q) - \big(\theta + nH(\theta)\big)q + (n-1)\big(\overline{T}(q) - \theta q\big) \quad \text{a.e.}, \tag{20}$$

where $\overline{T}$ implements $(\overline{U}, \overline{q})$.

Because the monotone hazard rate property holds (i.e., $\dot{H}(\theta) \geq 0$), $\theta + nH(\theta)$ is increasing in θ. Therefore, it immediately follows from standard revealed preferences arguments that $\overline{q}(\theta)$ that solves (20) is necessarily nondecreasing. Thus, the solution to the relaxed problem $(\overline{\mathcal{P}}^r)$ also solves $(\overline{\mathcal{P}})$ with the addition of the convexity requirement for $\overline{U}$.

Define now the value function for the program (20) as

$$\overline{V}(\theta) \equiv \max_{q \in \mathcal{Q}} S(q) - \big(\theta + nH(\theta)\big)q + (n-1)\big(\overline{T}(q) - \theta q\big). \tag{21}$$

Because the maximand on the right-hand side of (21) is absolutely continuous in θ, upper semicontinuous in q, and $\mathcal{Q}$ is compact, $\overline{V}(\theta)$ is itself absolutely continuous (Milgrom and Segal 2002). Moreover, given that $(\overline{U}, \overline{q})$ is an incentive-compatible allocation that solves this program, we have

$$\overline{V}(\theta) = S\big(\overline{q}(\theta)\big) - \big(\theta + nH(\theta)\big)\overline{q}(\theta) + (n-1)\overline{U}(\theta).$$

Because $\overline{V}$ is absolutely continuous, it is almost everywhere differentiable and it admits the integral representation

$$\overline{V}(\theta) - \overline{V}(\theta') = -\int_{\theta'}^{\theta} \big[\big(1 + n\dot{H}(\tilde{\theta})\big)\overline{q}(\tilde{\theta}) + (n-1)\overline{q}(\tilde{\theta})\big]\, d\tilde{\theta} \quad \forall(\theta, \theta').$$

Because $\overline{U}$ is also absolutely continuous, we have

$$\overline{U}(\theta) - \overline{U}(\theta') = -\int_{\theta'}^{\theta} \overline{q}(\tilde{\theta})\, d\tilde{\theta} \quad \forall(\theta, \theta').$$

Note that

$$\begin{aligned} & S\big(\overline{q}(\theta)\big) - \big(\theta + nH(\theta)\big)\overline{q}(\theta) - \big[S\big(\overline{q}(\theta')\big) - \big(\theta' + nH(\theta')\big)\overline{q}(\theta')\big] \\ &\quad = \overline{V}(\theta) - \overline{V}(\theta') - (n-1)\big[\overline{U}(\theta) - \overline{U}(\theta')\big]. \end{aligned}$$

Thus $S(\overline{q}(\theta)) - (\theta + nH(\theta))\overline{q}(\theta)$ is also absolutely continuous and admits the integral representation

$$\begin{aligned} & S\big(\overline{q}(\theta)\big) - \big(\theta + nH(\theta)\big)\overline{q}(\theta) - \big[S\big(\overline{q}(\theta')\big) - \big(\theta' + nH(\theta')\big)\overline{q}(\theta')\big] \\ &\quad = -\int_{\theta'}^{\theta} \big(1 + n\dot{H}(\tilde{\theta})\big)\overline{q}(\tilde{\theta})\, d\tilde{\theta}. \end{aligned}$$

Observe that

$$\big(H(\theta) - H(\theta')\big)\overline{q}(\theta') = \int_{\theta'}^{\theta} \dot{H}(\tilde{\theta})\overline{q}(\theta')\, d\tilde{\theta}.$$

Using this latter condition, the monotone hazard rate property $\dot{H}(\theta) \geq 0$, and the fact that $\overline{q}$ is nonincreasing, we obtain

$$S(\overline{q}(\theta)) - (\theta + nH(\theta))\overline{q}(\theta) - [S(\overline{q}(\theta')) - (\theta + nH(\theta))\overline{q}(\theta')]$$
$$= \int_{\theta'}^{\theta} (1 + n\dot{H}(\tilde{\theta}))(\overline{q}(\theta') - \overline{q}(\tilde{\theta}))\, d\tilde{\theta} \geq 0.$$

Because any $q' \in \overline{q}(\Theta)$ can be identified with some $\theta' \in \Theta$ so that $q' = \overline{q}(\theta')$, the inequality implies that $\overline{q}(\theta)$ satisfies (7) pointwise in θ.

Sufficiency. Consider any allocation $(\overline{U}, \overline{q})$ that satisfies (7), and (8). Simple revealed preferences arguments from (7) together with the monotone hazard rate property $\dot{H}(\theta) \geq 0$ imply that $\overline{q}$ is nonincreasing. Then (8) implies that $\overline{U}$ is convex. We now prove that this allocation is equilibrium implementable. First, we construct an aggregate transfer by duality as in (2) and obtain $\overline{T}$ from (4). By construction $\overline{q}(\theta)$ is a maximizer of $\overline{T}(q) - \theta q$ over $\overline{q}(\Theta)$. Second, individual contracts are then recovered by using (5). Third, we need to check that the optimality conditions (20) for the surrogate principal's optimization problem are satisfied. This last step is an immediate consequence of the fact that $\overline{q}(\theta)$ is a maximizer for both $S(q) - (\theta + nH(\theta))q$ and $\overline{T}(q) - \theta q$ over $\overline{q}(\Theta)$ so that it also maximizes a convex combination of both objectives. □

Proof of Corollary 1. Consider $T^m(q) = \min_{q \in \mathcal{Q}} U^m(\theta) + \theta q$. We now define a SGM problem associated to that tariff and check that (U^m, q^m) is a solution. Indeed, once one has taken care of the expression of the rent and integrating by parts, this SGM problem can be rewritten as in (20) in terms of output only as

$$q^m(\theta) \in \arg\max_{q \in \mathcal{Q}} S(q) - (\theta + nH(\theta))q + (n-1)(T^m(q) - \theta q).$$

That q^m is indeed a solution then follows from the fact that $q^m(\theta) \in \arg\max_{q \in \mathcal{Q}} S(q) - (\theta + nH(\theta))q$ and $q^m(\theta) \in \arg\max_{q \in \mathcal{Q}} (n-1)(T^m(q) - \theta q)$. □

Proof of Proposition 3. *Necessity.* That $\overline{q}$ should be nonincreasing follows from the first step in the proof of Proposition 2. Thus, $\overline{q}$ is almost everywhere differentiable (countable number of possible discontinuities). At any point θ where $\overline{q}$ is differentiable, the first-order necessary condition for optimality of the incentive problem

$$\theta \in \arg\max_{\hat{\theta} \in \Theta} S(\overline{q}(\hat{\theta})) - (\theta + nH(\theta))\overline{q}(\hat{\theta}) \quad \forall \theta \in \Theta$$

gives us (12).

Consider now the value function $\overline{V}$ as defined in (21). The function $\overline{V}$ is continuous and, thus, at a point of discontinuity θ_0 for $\overline{q}$, we have

$$\lim_{\theta \to \theta_0^-} \overline{V}(\theta) = \lim_{\theta \to \theta_0^+} \overline{V}(\theta). \tag{22}$$

Because $\overline{q}$ is nonincreasing, it is almost everywhere differentiable and such points of discontinuity are necessarily isolated. On the right- and the left-hand neighborhoods of

θ_0, (12) thus applies and either $\dot{\overline{q}}(\theta) = 0$ or $\overline{q}(\theta) = q^m(\theta)$. Moreover, at a point at which $\overline{q}$ is continuous but not differentiable, it must be that either the right or the left derivative is zero. We are going to prove that bunching arises both on a right- and a left-hand neighborhood of θ_0. We proceed by contradiction.

To this end, suppose first that bunching arises on the left neighborhood only and, thus, call it $\overline{q}(\theta_0^-) = \lim_{\theta\to\theta_0^-} \overline{q}(\theta)$ with $\overline{q}(\theta_0^-) > q^m(\theta_0)$ because $\overline{q}$ cannot be nondecreasing at such a discontinuity. Observe that $\overline{q}(\theta_0^-) = q^m(\theta_1)$ for some type $\theta_1 < \theta_0$ such that $\theta_1 = \max\{\theta \text{ s.t. } q^m(\theta) \geq \overline{q}(\theta_0^-)\}$ and that $q^m(\theta_1) = \overline{q}(\theta)$ for all $\theta \in [\theta_1, \theta_0)$.

Because the agent's rent $\overline{U}$ is also continuous at θ_0, we have

$$\overline{U}(\theta_0) = \lim_{\theta\to\theta_0^-} \overline{T}\big(\overline{q}(\theta)\big) - \theta\overline{q}(\theta) = \overline{T}\big(\overline{q}(\theta_0^-)\big) - \theta\overline{q}(\theta_0^-)$$

and

$$\overline{U}(\theta_0) = \lim_{\theta\to\theta_0^+} \overline{T}\big(\overline{q}(\theta)\big) - \theta\overline{q}(\theta) = \overline{T}\big(q^m(\theta_0)\big) - \theta q^m(\theta_0).$$

Therefore, we get

$$\overline{T}\big(\overline{q}(\theta_0^-)\big) - \theta\overline{q}(\theta_0^-) = \overline{T}\big(q^m(\theta_0)\big) - \theta q^m(\theta_0).$$

Inserting this equality into (22) and simplifying yields

$$\lim_{\theta\to\theta_0^-} S\big(\overline{q}(\theta)\big) - \big(\theta + nH(\theta)\big)\overline{q}(\theta) = \lim_{\theta\to\theta_0^+} S\big(q^m(\theta)\big) - \big(\theta + nH(\theta)\big)q^m(\theta).$$

Expressing those right- and left-hand side limits gives us

$$S\big(\overline{q}(\theta_0^-)\big) - \big(\theta_0 + nH(\theta_0)\big)\overline{q}(\theta_0^-) = S\big(q^m(\theta_0)\big) - \big(\theta_0 + nH(\theta_0)\big)q^m(\theta_0). \tag{23}$$

Because S is strictly concave, $q^m(\theta_0)$ is the unique maximizer of $S(q) - (\theta_0 + nH(\theta_0))q$ and (23) necessarily implies that $\overline{q}(\theta_0^-) = q^m(\theta_0)$. A contradiction with our starting premise that $\overline{q}(\theta_0^-) > q^m(\theta_0)$ at a discontinuity.

Similarly, we could also rule out the case where bunching only arises on the right-hand neighborhood of θ_0 at a value $\overline{q}(\theta_0^+) = \lim_{\theta\to\theta_0^+} \overline{q}(\theta)$.

Taking stock of these findings, we necessarily have $\overline{q}(\theta_0^-) > q^m(\theta_0) > \overline{q}(\theta_0^+)$ at a discontinuity point θ_0. This implies in passing that θ_0 must be such that $q^m(\theta_0) > 0$ and discontinuities do not lie on the lower boundary of the output space. Moreover, bunching arises on both sides of θ_0, which means $\overline{q}(\theta) = \overline{q}(\theta_0^-)$ (resp. $\overline{q}(\theta) = \overline{q}(\theta_0^+)$) for θ on this left-hand (resp. right-hand) neighborhood. Because q^m is strictly decreasing, there thus exist $\theta_1 < \theta_0 < \theta_2$ such that $\overline{q}(\theta_0^-) = q^m(\theta_1)$ and $\overline{q}(\theta_0^+) = q^m(\theta_2)$. In fact $\overline{q}(\theta) = q^m(\theta_1)$ for all $\theta \in [\theta_1, \theta_0)$. Suppose not. Then $\overline{q}$ would have a downward discontinuity at some $\theta_0' \in (\theta_1, \theta_0)$. The same argument as above shows that at any such putative discontinuity, we should have $\overline{q}(\theta_0'^-) > q^m(\theta_0') > \overline{q}(\theta_0'^+)$ and $\overline{q}(\theta_0'^+) \geq q^m(\theta_1)$. Since q^m is decreasing, this is a contradiction with the definition of θ_0'.

Because the agent's rent $\overline{U}$ is continuous at θ_0, we now have

$$\overline{U}(\theta_0) = \lim_{\theta\to\theta_0^-} \overline{T}\big(\overline{q}(\theta)\big) - \theta\overline{q}(\theta) = \overline{T}\big(q^m(\theta_1)\big) - \theta q^m(\theta_1)$$

and

$$\overline{U}(\theta_0) = \lim_{\theta\to\theta_0^+} \overline{T}\big(\overline{q}(\theta)\big) - \theta\overline{q}(\theta) = \overline{T}\big(q^m(\theta_2)\big) - \theta q^m(\theta_2).$$

It follows that

$$\overline{T}\big(q^m(\theta_1)\big) - \theta q^m(\theta_1) = \overline{T}\big(q^m(\theta_2)\big) - \theta q^m(\theta_2).$$

Inserting this equality into (22) and simplifying now yields

$$\lim_{\theta\to\theta_0^-} S\big(\overline{q}(\theta)\big) - \big(\theta + nH(\theta)\big)\overline{q}(\theta) = \lim_{\theta\to\theta_0^+} S\big(\overline{q}(\theta)\big) - \big(\theta + nH(\theta)\big)\overline{q}(\theta).$$

Expressing those right- and left-hand side limits gives us

$$S\big(q^m(\theta_1)\big) - \big(\theta_0 + nH(\theta_0)\big)q^m(\theta_1) = S\big(q^m(\theta_2)\big) - \big(\theta_0 + nH(\theta_0)\big)q^m(\theta_2),$$

which is (13).

Sufficiency. Sufficiency directly follows from the sufficiency part of the proof of Proposition 2. From the equilibrium output schedule $\overline{q}$, we reconstruct $\overline{U}(\theta) = \int_\theta^{\theta_H} \overline{q}(\tilde{\theta})\, d\tilde{\theta}$, the aggregate tariff $\overline{T}(q) = \min_{\theta\in\Theta} \overline{U}(\theta) + \theta q$, and individual tariffs so that $S_i(q) - \overline{T}_i(q) = S(q) - \overline{T}(q)$. □

Proof of Proposition 4. Proposition 1 shows that the maximal equilibrium (U^m, q^m) is an equilibrium. Since we have both $T^m(q) = \min_{\theta\in\Theta} U^m(\theta) + \theta q$ and $U^m(\theta) = \max_{q\in\mathcal{Q}} T^m(q) - \theta q$, (U^m, q^m) is in fact a biconjugate equilibrium. Consider now another equilibrium $(\overline{U}, \overline{q})$ whose range $\overline{q}(\Theta)$ presents one discontinuity gap $(q^m(\theta_2), q^m(\theta_1))$ for some (θ_1, θ_2). (The case of several discontinuity gaps follows similar steps and is omitted.) From our previous analysis, the corresponding discontinuity of $\overline{q}$ is at some $\theta_0 \in (\theta_1, \theta_2)$ with $\overline{q}(\theta_0^+) = q^m(\theta_2)$ and $\overline{q}(\theta_0^-) = q^m(\theta_1)$. Clearly, $\overline{T} \neq T^*$, where T^* is the least-concave tariff that implements $\overline{q}$ and that is defined as $T^*(q) = \min_{\theta\in\Theta} \overline{U}(\theta) + \theta q$.

We want to prove that replacing the implementing tariff $\overline{T}$ as proposed in (4) by T^* changes the solution to the SGM problem $(\overline{\mathcal{P}})$ and, more precisely, that $(\overline{U}, \overline{q})$ is no longer a solution to the new SGM problem $(\overline{\mathcal{P}}^*)$ so constructed. First, observe that T^* is concave and linear over $(q^m(\theta_2), q^m(\theta_1))$ with slope $\theta_0 \in (\theta_1, \theta_2)$, i.e., $T^{*\prime}(q) = \theta_0$ for $q \in (q^m(\theta_2), q^m(\theta_1))$. Second, we notice that the surrogate surplus at θ_0 is itself strictly concave. It writes as

$$S(q) - \big(\theta_0 + nH(\theta_0)\big)q + (n-1)\big(\overline{T}(q) - \theta_0 q\big).$$

This expression is maximized at $q^m(\theta_0) \in (q^m(\theta_H), q^m(\theta_L))$ and, thus, we should have $\overline{q}(\theta_0) = q^m(\theta_0)$ if $(\overline{U}, \overline{q})$ was the solution to the self-generating problem $(\overline{\mathcal{P}}^*)$. However, at the nonmaximal equilibrium $(\overline{U}, \overline{q})$, θ_0 should be indifferent between the two boundaries of the discontinuity hold, namely $q^m(\theta_1)$ and $q^m(\theta_2)$, and cannot take any option within the discontinuity gap $(q^m(\theta_1), q^m(\theta_2))$. A contradiction. Hence, $(\overline{U}, \overline{q})$ is no longer a solution to $(\overline{\mathcal{P}}^*)$. □

Proof of Proposition 5. First, we notice that $\overline{U}_\varepsilon$ and U^m are identical on $[\theta_2(\varepsilon), \theta_H]$ since $\overline{U}_\varepsilon(\theta) = \int_\theta^{\theta_H} \overline{q}_\varepsilon(\tilde{\theta})\, d\tilde{\theta}$, $U^m(\theta) = \int_\theta^{\theta_H} q^m(\tilde{\theta})\, d\tilde{\theta}$ and outputs are the same on that interval, namely $\overline{q}_\varepsilon(\tilde{\theta}) = q^m(\theta)$ on $[\theta_2(\varepsilon), \theta_H]$. Second, observe that over the discontinuity gap $[\theta_1(\varepsilon), \theta_2(\varepsilon)]$, the equilibrium rent $\overline{U}_\varepsilon$ satisfies

$$\overline{U}_\varepsilon\big(\theta_1(\varepsilon)\big) - \overline{U}_\varepsilon\big(\theta_2(\varepsilon)\big) = q^m\big(\theta_2(\varepsilon)\big)\big(\theta_2(\varepsilon) - \theta_0\big) + q^m\big(\theta_1(\varepsilon)\big)\big(\theta_0 - \theta_1(\varepsilon)\big). \tag{24}$$

Third, for the discontinuity gap so considered, condition (13) now writes as

$$S\big(q^m\big(\theta_1(\varepsilon)\big)\big) - \big(\theta_0 + nH(\theta_0)\big)q^m\big(\theta_1(\varepsilon)\big) = S\big(q^m\big(\theta_2(\varepsilon)\big)\big) - \big(\theta_0 + nH(\theta_0)\big)q^m\big(\theta_2(\varepsilon)\big). \tag{25}$$

Observe that a discontinuity can only arise at a point θ_0 such $q^m(\theta_0) > 0$. Thus q^m given in (9) is positive and everywhere differentiable (and twice times so when S is three time so) in a neighborhood of θ_0 with $\dot{q}^m(\theta) < 0$ there. We thus have

$$\big(S'\big(q^m(\theta)\big) - \theta\big)\dot{q}^m(\theta) = nH(\theta)\dot{q}^m(\theta) \quad \forall \theta \in \Theta.$$

We also have

$$\big(S'\big(q^m(\theta)\big) - \theta\big)\dot{q}^m(\theta) = \frac{d}{d\theta}\big(S\big(q^m(\theta)\big) - \theta q^m(\theta)\big) + q^m(\theta) \quad \forall \theta \in \Theta.$$

Hence,

$$nH(\theta)\dot{q}^m(\theta) = \frac{d}{d\theta}\big(S\big(q^m(\theta)\big) - \theta q^m(\theta)\big) + q^m(\theta) \quad \forall \theta \in \Theta.$$

Integrating over the interval $[\theta_1(\varepsilon), \theta_2(\varepsilon)]$, we obtain

$$n\int_{\theta_1(\varepsilon)}^{\theta_2(\varepsilon)} H(\theta)\dot{q}^m(\theta)\, d\theta = S(q^m\big(\theta_2(\varepsilon)\big) - \theta_2 q^m\big(\theta_2(\varepsilon)\big) - (S(q^m\big(\theta_1(\varepsilon)\big) - \theta_1 q^m\big(\theta_1(\varepsilon)\big)$$
$$+ \int_{\theta_1(\varepsilon)}^{\theta_2(\varepsilon)} q^m(\theta)\, d\theta.$$

Simplifying using (24) and (25), we finally obtain

$$n\int_{\theta_1(\varepsilon)}^{\theta_2(\varepsilon)} H(\theta)\dot{q}^m(\theta)\, d\theta = nH(\theta_0)\big(q^m\big(\theta_2(\varepsilon)\big) - q^m\big(\theta_1(\varepsilon)\big)\big) - \big(\overline{U}\big(\theta_1(\varepsilon)\big) - \overline{U}\big(\theta_2(\varepsilon)\big)\big)$$
$$+ \int_{\theta_1(\varepsilon)}^{\theta_2(\varepsilon)} q^m(\theta)\, d\theta.$$

Let us now define

$$\Delta(\varepsilon) = \overline{U}_\varepsilon\big(\theta_1(\varepsilon)\big) - \overline{U}_\varepsilon\big(\theta_2(\varepsilon)\big) - \big(U^m\big(\theta_1(\varepsilon)\big) - U^m\big(\theta_2(\varepsilon)\big)\big).$$

Since $\overline{U}_\varepsilon$ and U^m are identical on $[\theta_2(\varepsilon), \theta_H]$, the sign of $\Delta(\varepsilon)$ also measures whether $\overline{U}_\varepsilon(\theta_1(\varepsilon))$ might lie above (when $\Delta(\varepsilon) \geq 0$) or below (when $\Delta(\varepsilon) \leq 0$) $U^m(\theta_1(\varepsilon))$.

From the previous steps, we get

$$\Delta(\varepsilon) = n\big(\psi\big(\theta_2(\varepsilon)\big) - \psi\big(\theta_1(\varepsilon)\big)\big),$$

where

$$\psi(\theta)=H(\theta_0)q^m(\theta)-\int_{\theta_0}^{\theta}H(\theta)\dot{q}^m(\theta)\,d\theta.$$

Observe that

$$\dot{\psi}(\theta_0)=0,\quad \text{and}\quad \ddot{\psi}(\theta_0)=-\dot{H}(\theta_0)\dot{q}^m(\theta_0)>0.$$

We now write $\theta_1(\varepsilon)=\theta_0-\varphi(\varepsilon)$ for some function φ when $\theta_2(\varepsilon)=\theta_0+\varepsilon$. Observe that (25) implies that $\varphi(\varepsilon)$ goes to zero with ε. We also define

$$\Gamma(\varepsilon)=S\big(q^m(\theta_0+\varepsilon)\big)-\big(\theta_0+nH(\theta_0)\big)q^m(\theta_0+\varepsilon).$$

Tedious computations show that

$$\Gamma'(0)=0,\qquad \Gamma''(0)=\dot{q}^m(\theta_0)\big(1+n\dot{H}(\theta_0)\big)$$

and

$$\Gamma'''(0)=\dot{q}^m(\theta_0)\big(2S''\big(q^m(\theta_0)\big)\ddot{q}^m(\theta_0)+n\ddot{H}(\theta_0)\big).$$

Third-order Taylor expansions immediately give us

$$\Gamma(\varepsilon)=\Gamma(0)+\frac{\Gamma''(0)}{2}\varepsilon^2+\frac{\Gamma'''(0)}{6}\varepsilon^3+o\big(\varepsilon^3\big)$$

and

$$\Gamma\big(-\varphi(\varepsilon)\big)=\Gamma(0)+\frac{\Gamma''(0)}{2}\varphi^2(\varepsilon)-\frac{\Gamma'''(0)}{6}\varphi^3(\varepsilon)+o\big(\varepsilon^3\big).$$

Observe that (25) can then be written as $\Gamma(\varepsilon)=\Gamma(-\varphi(\varepsilon))$ or, using the above Taylor expansions,

$$\frac{\Gamma''(0)}{2}\varepsilon^2+\frac{\Gamma'''(0)}{6}\varepsilon^3=\frac{\Gamma''(0)}{2}\varphi^2(\varepsilon)-\frac{\Gamma'''(0)}{6}\varphi^3(\varepsilon)+o\big(\varepsilon^3\big).$$

Up to terms of order of magnitude greater than 2, we thus obtain

$$\varphi(\varepsilon)=\varepsilon+\frac{n\ddot{H}(\theta_0)+2S''\big(q^m(\theta_0)\big)\ddot{q}^m(\theta_0)}{1+n\dot{H}(\theta_0)}\frac{\varepsilon^2}{6}+o\big(\varepsilon^2\big).$$

Turning now to a Taylor expansion of $\Delta(\varepsilon)$ up to terms of magnitude higher than 2, we get

$$\Delta(\varepsilon)=n\frac{\ddot{\psi}(\theta_0)}{2}\big(\varepsilon-\varphi(\varepsilon)\big)2\varepsilon+o\big(\varepsilon^2\big).$$

Observe that $\ddot{H}(\theta)\geq 0$ and $S'''\leq 0$ implies $\ddot{q}^m\leq 0$ since $S'''(q^m(\theta))(\dot{q}^m(\theta))^2+S''(q^m(\theta))\ddot{q}^m(\theta)=n\ddot{H}(\theta)$. Thus $\varphi(\varepsilon)>\varepsilon$. Hence, $\Delta(\varepsilon)$ is negative for ε positive and small enough. Finally, $\overline{U}_\varepsilon(\theta_1(\varepsilon)) \le U^m(\theta_1(\varepsilon))$ and thus

$$\overline{U}_\varepsilon(\theta) = \overline{U}_\varepsilon\big(\theta_1(\varepsilon)\big) + \int_\theta^{\theta_1(\varepsilon)} q^m(\theta)\,d\theta \le U(\theta) = U^m\big(\theta_1(\varepsilon)\big) + \int_\theta^{\theta_1(\varepsilon)} q^m(\theta)\,d\theta,$$

which concludes the proof of the first part of the proposition.

To establish global preferences under uniform-quadratic preferences, we first show that introducing a discontinuity in the maximal equilibrium allocation weakly lowers the agent's utility. Suppose that this discontinuity is introduced at θ_0 and the associated intervals of bunching from Proposition 3 are $\overline{q}(\theta) = q^m(\theta_1)$ on $[\theta_1, \theta_0)$ and $\overline{q}(\theta) = q^m(\theta_2)$ on $(\theta_0, \theta_2]$. Following a similar argument to that provided for the numerical example in the text, it is sufficient to show that

$$\int_{\theta_1}^{\theta_0} q^m(\theta_1)\,d\theta + \int_{\theta_0}^{\theta_2} q^m(\theta_2)\,d\theta = \int_{\theta_1}^{\theta_2} q^m(\theta)\,d\theta$$

to ensure that the rent profiles $\overline{U}$ and U^m are the same not only on $[\underline{\theta}, \theta_1]$ and $[\theta_2, \overline{\theta}]$. Over the interval (θ_1, θ_2), instead $\overline{U} < U^m$, since the rent profile $\overline{U}$ is obtained from that of U^m by replacing the strictly convex part of U^m over $[\theta_1, \theta_2]$ by the maximum of the two tangents at the points θ_1 and θ_2, namely

$$\overline{U}(\theta) = \max\big\{U^m(\theta_1) - q^m(\theta_1)(\theta - \theta_1); U^m(\theta_2) - q^m(\theta_2)(\theta - \theta_2)\big\}, \quad \theta \in [\theta_1, \theta_2].$$

At the discontinuity point θ_0, the two tangents cross, capturing the fact that θ_0 is indifferent between choosing $q^m(\theta_1)$ and $q^m(\theta_2)$.

In the uniform-quadratic case, for $q^m(\theta) > 0$, $q^m(\theta)$ is linear. The above condition is therefore satisfied if and only if

$$\theta_0 = \frac{\theta_1 + \theta_2}{2}.$$

An equilibrium requirement of the new discontinuous allocation is that the surrogate principal is indifferent between inducing $q^m(\theta_1)$ and $q^m(\theta_2)$ at θ_0 (Proposition 3). This condition is equivalent to

$$\theta_0 + nH(\theta_0) = \frac{S\big(q^m(\theta_1)\big) - S\big(q^m(\theta_2)\big)}{q^m(\theta_1) - q^m(\theta_2)}.$$

The first-order condition defining q^m implies that the left-hand side equals $S'(q^m(\theta_0))$, and thus

$$S'\big(q^m(\theta_0)\big) = \frac{S\big(q^m(\theta_1)\big) - S\big(q^m(\theta_2)\big)}{q^m(\theta_1) - q^m(\theta_2)}.$$

Because S is quadratic, this requirement reduces to

$$q^m(\theta_0) = \frac{q^m(\theta_1) + q^m(\theta_2)}{2}.$$

Because q^m is linear in θ, this condition is equivalent to $\theta_0 = \frac{\theta_1+\theta_2}{2}$. Thus, we have established that the new equilibrium allocation is unchanged over $\Theta \setminus (\theta_1, \theta_2)$, but $\overline{U}(\theta) < U^m(\theta)$ for all $\theta \in (\theta_1, \theta_2)$. The agent is worse off.

Repeating the above argument, we may introduce a second discontinuity over any region of full separation and again establish that the new twice-discontinuous equilibrium allocation is worse for the agent than the once-discontinuous allocation, which is worse than the maximal allocation. Proceeding inductively, we can establish that any arbitrary equilibrium allocation with k discontinuities is worse for the agent than the original maximal allocation. □

Proof of Proposition 6. For the sake of clarity and because we shall soon import their results, we (mainly) use the notations of Amador and Bagwell (2013). We thus transform θ into the type, $\gamma = -(\theta + nH(\theta))$, and let $\phi(\gamma)$ denote the inverse mapping. The corresponding density for γ is given by $\tilde{f}(\gamma) = f(\phi(\gamma))|\phi'(\gamma)|$. Let $\tilde{F}(\gamma)$ denote the distribution of γ. This distribution has support $[\gamma_H, \gamma_L]$ with $\theta_H = \phi(\gamma_H) > \theta_L = \phi(\gamma_L)$, where ϕ is decreasing. Accordingly, we define the surrogate principal's preferences as

$$S(q) + \gamma q.$$

Observe that this expression is maximized at $q^m(\phi(\gamma))$ such that $S'(q^m(\phi(\gamma))) + \gamma = 0$ and that $q^m(\phi(\gamma))$ is everywhere positive when $q^m(\theta_H) > 0$.

Following the notation in Amador and Bagwell (2013), the ex ante cooperating principals' preferences can now be rewritten as

$$w(\gamma, q) = S(q) + \gamma q + (n-1)H\big(\phi(\gamma)\big)q.$$

Because the surrogate principal's best action, $q^m(\phi(\gamma))$, is strictly positive for all γ when $q^m(\theta_H) > 0$, the standard regularity conditions stated in Assumption 1 of Amador and Bagwell (2013) are all satisfied in our setting.

Consider an interval $[\gamma_2, \gamma_1] \subset [\gamma_H, \gamma_L]$, with $\gamma_2 < \gamma_1$. Following Amador and Bagwell (2013), we now consider the conditions

(c1*) $\tilde{F}(\gamma) - w_q(\gamma, q^m(\phi(\gamma)))\tilde{f}(\gamma)$ is nondecreasing for all $\gamma \in [\gamma_H, \gamma_L]$[24]

(c2) if $\gamma_1 < \gamma_L$, $\gamma - \gamma_1 \geq \int_{\gamma}^{\gamma_L} w_q(\tilde{\gamma}, q^m(\phi(\gamma_1)))\frac{\tilde{f}(\tilde{\gamma})}{1-\tilde{F}(\gamma)}\,d\tilde{\gamma}$, $\forall \gamma \in [\gamma_1, \gamma_L]$ with equality at γ_1

(c2′) if $\gamma_1 = \gamma_L$, $w_q(\gamma_L, q^m(\phi(\gamma_L))) \geq 0$

(c3) if $\gamma_2 > \gamma_H$, $\gamma - \gamma_2 \leq \int_{\gamma_H}^{\gamma} w_q(\tilde{\gamma}, q^m(\phi(\gamma_2)))\frac{\tilde{f}(\tilde{\gamma})}{\tilde{F}(\gamma)}\,d\tilde{\gamma}$, $\forall \gamma \in [\gamma_H, \gamma_2]$ with equality at γ_2

(c3′) if $\gamma_2 = \gamma_H$, $w_q(\gamma_H, q^m(\phi(\gamma_H))) \leq 0$.

Equipped with these conditions, Proposition 1 in Amador and Bagwell (2013), which determines whether leaving discretion to the surrogate principal on a particular delegation set is optimal from the point of view of the ex ante cooperating principals, can be restated in our context as follows.

[24]Condition (c1*) is a strengthening of condition (c1) in Amador and Bagwell (2013) since it imposes monotonicity on all the domain $[\gamma_H, \gamma_L]$ and not only for $\gamma \in [\gamma_2, \gamma_1]$.

Proposition 7. *If conditions (c1*), (c2), (c2′), (c3), and (c3′) are satisfied, then the surrogate principal's preferred output* $q^m(\phi(\gamma))$ *is optimal over the interval* $[\gamma_2, \gamma_1]$, *where* $\gamma_1 > \gamma_2$.

We wish to prove that the optimal delegation set in our setting is of the form $[\gamma_2, \gamma_L]$ (or, alternatively, that the bunching area expressed in terms of θ is of the form $[\theta_2, \theta_H]$). Consequently, we need to verify conditions (c1*), (c2), (c2′), (c3), and (c3′) at the conjectured solution ($\gamma_2 > \gamma_H$, $\gamma_1 = \gamma_L > \gamma_2$).

First, if $\gamma_1 = \gamma_L$, condition (c2′) becomes $w_q(\gamma_L, q^m(\phi(\gamma_L))) \geq 0$, which holds as an equality since the cooperating principals and the surrogate representative have perfectly aligned preferences for $\gamma_L = \phi(\theta_L)$ (indeed $S'(q^m(\theta_L)) - \theta_L = 0$).

Second, if we conjecture that $\gamma_2 > \gamma_H$, it must be that condition (c3) holds, i.e.,

$$\gamma - \gamma_2 \leq \int_{\gamma_H}^{\gamma} w_q(\tilde{\gamma}, q^m(\phi(\gamma_2))\frac{\tilde{f}(\tilde{\gamma})}{\tilde{F}(\gamma)}\, d\tilde{\gamma}, \quad \forall \gamma \in [\gamma_H, \gamma_2] \text{ with equality at } \gamma_2.$$

Observe that

$$w_q\big(\tilde{\gamma}, q^m\big(\phi(\gamma_2)\big)\big) = S'\big(q^m\big(\phi(\gamma_2)\big)\big) + \tilde{\gamma} + (n-1)H\big(\phi(\tilde{\gamma})\big).$$

Since $S'(q^m(\phi(\gamma_2))) + \gamma_2 = 0$, we thus have

$$w_q\big(\tilde{\gamma}, q^m\big(\phi(\gamma_2)\big)\big) = \tilde{\gamma} - \gamma_2 + (n-1)H\big(\phi(\tilde{\gamma})\big).$$

Condition (c3) can thus be rewritten as

$$\gamma - \gamma_2 \leq \int_{\gamma_H}^{\gamma} \big(\tilde{\gamma} - \gamma_2 + (n-1)H\big(\phi(\tilde{\gamma})\big)\big)\frac{\tilde{f}(\tilde{\gamma})}{\tilde{F}(\gamma)}\, d\tilde{\gamma}, \quad \forall \gamma \in [\gamma_H, \gamma_2] \text{ with equality at } \gamma_2.$$

Take $\gamma_2 = \phi(\hat{\theta})$, where $\hat{\theta}$ is defined in (15). Then observe that (15) also writes as

$$\int_{\hat{\theta}}^{\theta_H} \big(\hat{\theta} + nH(\hat{\theta}) - \theta - H(\theta)\big)f(\theta)\, d\theta = 0.$$

Changing variables, this condition becomes

$$\int_{\gamma_H}^{\gamma_2} \big(\tilde{\gamma} - \gamma_2 + (n-1)H\big(\phi(\tilde{\gamma})\big)\big)\frac{\tilde{f}(\tilde{\gamma})}{\tilde{F}(\gamma)}\, d\tilde{\gamma} = 0.$$

Therefore, condition (c3) holds as an equality when $\gamma = \gamma_2$, as requested. Observe now that condition (c3) also holds for $\gamma \in [\gamma_H, \gamma_2]$ provided that (c1*) is true.

To see why, notice that

$$-\tilde{F}(\gamma)(\gamma - \gamma_2) + \int_{\gamma_H}^{\gamma} \big(\tilde{\gamma} - \gamma_2 + (n-1)H\big(\phi(\tilde{\gamma})\big)\big)\tilde{f}(\tilde{\gamma})\, d\tilde{\gamma}$$

has a derivative with respect to γ worth

$$-\tilde{F}(\gamma) + (n-1)H\big(\phi(\gamma)\big)\tilde{f}(\gamma) = -\tilde{F}(\gamma) + w_q\big(\gamma, q^m\big(\phi(\gamma)\big)\big)\tilde{f}(\gamma).$$

Thus that condition (c1*) holds amounts to checking that

$$\tilde{F}(\gamma) - (n-1)H\big(\phi(\gamma)\big)\tilde{f}(\gamma)$$

is nondecreasing. Differentiating, we shall prove

$$\tilde{f}(\gamma) - (n-1)\big(\dot{H}\big(\phi(\gamma)\big)\tilde{f}(\gamma)\phi'(\gamma) + H\big(\phi(\gamma)\big)\tilde{f}'(\gamma)\big) \geq 0.$$

Using $\tilde{f} = -f\phi'$, we have $\tilde{f}'(\gamma) = -f'(\theta)\phi'^2 - f(\theta)\phi''(\gamma)$. Because

$$\phi'(\gamma) = \frac{-1}{1 + n\dot{H}\big(\phi(\gamma)\big)},$$

we have $\phi''(\gamma) = \phi'^3 n\ddot{H}(\phi(\gamma))$. Substituting these relationships into our inequality and simplifying, we reduce our required condition to Assumption 1. Hence, condition (c1*) is satisfied.

We now prove that condition (15) is the defining first-order condition for finding an optimal floor delegation set. Consider thus a floor equilibrium with an allocation

$$\overline{q}(\theta) = \max\{q^m(\theta), q^m(\theta^*)\} \quad \text{for some } \theta^* \in \Theta$$

together with the rent minimization requirement (8). The maximal equilibrium is simply obtained by choosing $\theta^* = \theta_H$ as a special case. The sum of the principal's profits evaluated at such equilibria can be expressed in terms of the floor θ^* only as

$$\begin{aligned} V(\theta^*) = &\int_{\theta_L}^{\theta^*} \big(S\big(q^m(\theta)\big) - \big(\theta + H(\theta)\big)q^m(\theta)\big)f(\theta)\,d\theta \\ &+ \int_{\theta^*}^{\theta_H} \big(S\big(q^m(\theta^*)\big) - \big(\theta + H(\theta)\big)q^m(\theta^*)\big)f(\theta)\,d\theta. \end{aligned}$$

Differentiating with respect to θ^* yields

$$\dot{V}(\theta^*) = \dot{q}^m(\theta^*)\int_{\theta^*}^{\theta_H} \big(S'\big(q^m(\theta^*)\big) - \theta - H(\theta)\big)f(\theta)\,d\theta.$$

From this, we immediately get that

$$\dot{V}(\theta_L) = \dot{q}^m(\theta_L)\int_{\theta_L}^{\theta_H} \big(\theta_L - \theta - H(\theta)\big)f(\theta)\,d\theta > 0, \tag{26}$$

where the last inequality follows from $\theta_L - \theta - H(\theta) < 0$ and $\dot{q}^m(\theta_L) < 0$. Second, we also obtain:

$$\dot{V}(\theta_H) = 0 \quad \text{with } \ddot{V}(\theta_H) = -(n-1)\dot{q}^m(\theta_H) > 0$$

when $\dot{H}(\theta) \geq 0$ and $q^m(\theta_H) > 0$ so that $q^m(\theta)$ is strictly decreasing in that neighborhood of θ_H. Hence, although θ_H is a local extremum of V, it corresponds to a minimum. It follows that the maximal equilibrium is never optimal. From (26), we derive the existence of a maximum $\hat{\theta}$ that is necessarily interior. Rewriting the condition $\dot{V}(\hat{\theta}) = 0$ gives us (15). □

References

Alonso, Ricardo and Niko Matouschek (2008), "Optimal delegation." *Review of Economic Studies*, 75, 259–293. [1155, 1166, 1174]

Amador, Manuel and Kyle Bagwell (2013), "The theory of optimal delegation with an application to tariff caps." *Econometrica*, 81, 1541–1599. [1155, 1156, 1166, 1174, 1185]

Bagnoli, Mark and Ted Bergstrom (2005), "Log-concave probability and its applications." *Economic Theory*, 26, 445–469. [1162]

Basov, Suren (2005), *Multidimensional Screening*, volume 22 of Studies in Economic Theory. Springer-Verlag, Berlin Heidelberg. [1158]

Bernheim, B. Douglas and Michael Whinston (1986), "Common agency." *Econometrica*, 54, 923–942. [1153, 1156]

Calzolari, Giacomo and Vincenzo Denicolò (2013), "Competition with exclusive contracts and market-share discounts." *American Economic Review*, 103, 2384–2411. [1157]

Carlier, Guillaume (2001), "A general existence result for the principal–agent problem with adverse selection." *Journal of Mathematical Economics*, 35, 129–150. [1158]

Chiesa, Gabriella and Vincenzo Denicolò (2009), "Trading with a common agent under complete information: A characterization of Nash equilibria." *Journal of Economic Theory*, 144, 296–311. [1156]

Hoernig, Steffen and Tomaso Valletti (2011), "When two-part tariffs are not enough: Mixing with nonlinear pricing." *B.E. Journal of Theoretical Economics*, 11, 1–20. [1157]

Holmström, Bengt (1984), "On the theory of delegation." In *Bayesian Models in Economic Theory* (Marcel Boyer and Richard Kihlstrom, eds.). North-Holland Publishing Co, Amsterdam. [1155, 1166]

Laffont, Jean-Jacques and Jean Tirole (1993), *A Theory of Incentives in Procurement and Regulation*. MIT Press. [1155, 1157, 1168, 1169]

Laussel, Didier and Michel Le Breton (1998), "Efficient private production of public goods under common agency." *Games and Economic Behavior*, 25, 194–218. [1156, 1157]

Laussel, Didier and Michel Le Breton (2001), "Conflict and cooperation: The structure of equilibrium payoffs in common agency." *Journal of Economic Theory*, 100, 93–128. [1156]

Laussel, Didier and Joana Resende (2018), "Complementary monopolies with asymmetric information. A private common agency model." Unpublished paper, SSRN 3121668. [1157]

Martimort, David (1992), "Multi-Principaux avec Anti-Selection." *Annales d'Economie et de Statistiques*, 28, 1–37. [1157]

Martimort, David (2006), "Multi-contracting mechanism design." In *Advances in Economic Theory Proceedings of the World Congress of the Econometric Society* (Richard

Blundell, Whitney K. Newey, and Torsten Persson, eds.), 57–101, Cambridge University Press, New York, New York. [1152]

Martimort, David and Aggey Semenov (2006), "Continuity in mechanism design without transfers." *Economic Letters*, 93, 182–189. [1155, 1166, 1174]

Martimort, David and Aggey Semenov (2008), "Ideological uncertainty and lobbying competition." *Journal of Public Economics*, 92, 456–481. [1155, 1157]

Martimort, David and Lars A. Stole (2002), "The revelation and delegation principles in common agency games." *Econometrica*, 70, 1659–1673. [1152, 1158]

Martimort, David and Lars A. Stole (2009a), "Market participation under delegated and intrinsic common agency games." *RAND Journal of Economics*, 40, 78–102. [1155, 1157]

Martimort, David and Lars A. Stole (2009b), "Selecting equilibria in common agency games." *Journal of Economic Theory*, 144, 604–634. [1155, 1157]

Martimort, David and Lars A. Stole (2012), "Representing equilibrium aggregates in aggregate games with applications to common agency." *Games and Economic Behavior*, 76, 753–772. [1152, 1153, 1154, 1157, 1160, 1163, 1168, 1176]

Martimort, David and Lars A. Stole (2015), "Menu auctions and influence games with private information." Unpublished paper, SSRN 2569703. [1156, 1157, 1159, 1167]

Melumad, Nahum and Toshiyuki Shibano (1991), "Communication in settings with no transfers." *RAND Journal of Economics*, 22, 173–198. [1155, 1166]

Milgrom, Paul and Ilya Segal (2002), "Envelope theorems for arbitrary choice sets." *Econometrica*, 70, 583–601. [1176, 1178]

Myerson, Roger (1982), "Optimal coordination mechanisms in generalized principal–agent problems." *Journal of Mathematical Economics*, 10, 67–81. [1152]

Peters, Micahel (2001), "Common agency and the revelation principle." *Econometrica*, 69, 1349–1372. [1152]

Rochet, Jean-Charles (1987), "A necessary and sufficient condition for rationalizability in a quasi-linear context." *Journal of Mathematical Economics*, 16, 191–200. [1176]

Stole, Lars A. (1991), "Mechanism design under common agency." Unpublished paper, University of Chicago, Booth School of Business. [1157]

Szentes, Balázs (2015), "Contractible contracts in common agency problems." *Review of Economic Studies*, 82, 391–422. [1158]

Co-editor Johannes Hörner handled this manuscript.

Manuscript received 24 August, 2015; final version accepted 22 January, 2018; available online 13 March, 2018.