Low Resolution Economics

Jacopo Bizzotto* Toomas Hinnosaar† Adrien Vigier‡

December 28, 2021§

Abstract

Economic agents seldom observe the precise actions of those whom they interact with. In this paper, we explore a broad class of settings comprising a leader and a follower, where the follower can observe the element of a partition containing the leader's action, but nothing beyond that. We characterize the set of outcomes resulting from any possible partition of the leader's action space, and illustrate our analysis with applications to oligopolies, public goods, and R&D.

JEL: C72, D43, D82

Keywords: Observability, sequential games, information, competition, public goods, R&D

1 Introduction

Individuals, firms, and countries often do not precisely observe the actions of those whom they strategically interact with. Reasons are numerous: data is typically recorded and stored in a coarse manner, legal and physical constraints can prevent access to detailed information, or observational costs may be large. How do such considerations affect economic outcomes?

To address the previous question, we examine a broad class of settings comprising a leader and a follower. Each player chooses an action in an interval. Actions could be strategic complements or substitutes, with positive or negative externalities. The leader's action space

*OsloMet, jacopo.bizzotto@oslomet.no.

†University of Nottingham and CEPR, toomas@hinnosaar.net.

‡University of Nottingham, adrien.vigier@nottingham.ac.uk.

§We are grateful to Matteo Escude for many helpful discussions on the topic of this paper. We also benefited from the comments and suggestions of Giacomo Calzolari, Albert Ma, Lucie Menager, and Daniel Monte.

is partitioned into observable subsets: before choosing an action, the follower learns which partition element contains the action chosen by the leader, but nothing beyond that. We focus on environments which admit a unique Nash outcome when the follower is completely uninformed about the leader's true action (i.e., when the aforementioned partition contains a single element); borrowing terminology from Industrial Organization, we refer to this outcome as the Cournot outcome.

To illustrate with a familiar example, consider the textbook model of duopoly, in which two firms sequentially choose quantities. The classic model of von Stackelberg (1934) supposes that the follower *perfectly* observes the leader's strategy; we are concerned with the consequences of relaxing this extreme assumption. For instance: whereas the number of plants that a firm builds is readily observable, the amount and type of machinery that it uses is far harder to ascertain; a company might be able to observe that a development center has been established, without knowing the precise amount of resources that were invested (e.g., the number and the quality of the engineers that its competitor has hired might be difficult to find out); the follower could have access to data concerning the operations of the leader in countries A and B, all the while lacking data regarding country C.

To see how coarse observability might affect outcomes, imagine that, in the example described above the follower only learns whether the leader chose a quantity that is below or above some cutoff. Suppose moreover that said cutoff is larger than the leader's Cournot quantity, but smaller than its Stackelberg quantity. We argue that, in this case, in no equilibrium outcome does the leader select the Stackelberg quantity. The logic is simple: if the follower were to best-respond to the Stackelberg quantity, the leader would want to deviate to a lower quantity (without being detected by the follower). Similar arguments show that the leader in fact produces at the lower bound of the upper interval (i.e., at the cutoff itself). Thus, partial information here erodes the first-mover advantage of the leader. As the cutoff gets larger, the quantity which the leader selects gradually increases, and the leader's first-mover advantage is restored. Furthermore, raising the cutoff still further (i.e., beyond the Stackelberg quantity) goes on increasing quantity produced: as long as the resulting profit is greater than the profit obtained under the Cournot outcome, the leader has no profitable deviation.

Returning to our general setting, we say in this paper that an action pair comprises a *plausible* outcome if it is an equilibrium outcome of the strategic environment induced by *some* partition of the leader's action space. We characterize the entire set of plausible outcomes, as well as the subset of outcomes that remain plausible under interval partitions (i.e., when each partition element takes the form of an interval).

The ability to identify the set of plausible outcomes in any given environment is important for at least two reasons. Firstly, an external analyst may have imperfect knowledge of the true information structure.[1] In this case, knowing the full set of plausible outcomes might help to evaluate the consequences of, say, introducing a new policy. Alternatively, the information structure could be an object of *design*. For instance, crowdfunding platforms such as Kickstarter have complete freedom regarding what to reveal about past contributions so as to encourage future contributions; similarly, the manager of a large company who oversees a production chain can selectively disclose information downstream of the production chain regarding activity upstream.

We first show that the set of plausible outcomes under *interval* partitions is exactly the set of outcomes at which the leader enjoys a payoff lying between the payoffs obtained when the follower is, respectively, fully informed and uninformed, which we refer to as the leader's Stackelberg and Cournot payoffs (for expository purposes, we use throughout the paper the same terminology as in the duopoly example). The corresponding set of plausible leader actions have the Cournot action as an extreme point; in this sense, said action might be a poor predictor for an analyst possessing no knowledge of the true information structure. Whether the Cournot action is an upper or a lower bound of all plausible actions under interval partitions depends on the strategic environment, and the sign of the externalities.[2] By contrast, the leader's Stackelberg action may be an interior point of the set of plausible actions under interval partitions.

When either payoff externalities are small or strategic interactions weak, any plausible outcome is also plausible outcome under interval partitions. The question of whether the full set of plausible outcomes is larger than the corresponding set under interval partitions boils down to finding out whether, in the environment considered, the leader can ever be held to less than the Cournot payoff: for this, the payoff externalities need to be sufficiently large and the strategic interactions have to be sufficiently strong. Intuitively, the leader's incentive to select an action belonging to one partition element rather than another is determined by the threat of a low payoff resulting from the follower's action in case the leader deviated. The larger the payoff externalities and the more sensitive the follower's optimal action to the partition element observed, the more potent the aforementioned threat.

[1] In the context of our model, the information structure is subsumed by the partition of the leader's action space into observable subsets.

[2] With strategic complements (respectively, substitutes) and positive (respectively, negative) externalities, the Cournot action is a lower bound. In all other cases, the Cournot action is an upper bound.

Our model allows for arbitrary partitions of the leader's action space. However, we show that any plausible outcome can be achieved through some partition such that at most one element of this partition is not an interval. This observation enables us to characterize the full set of plausible outcomes, and to provide an easy method by which this set may be computed. Furthermore, whenever the information structure is an object of design, the previous observation means that the designer, whoever she might be, can without loss restrict attention to relatively simple partitions.

The second part of the paper studies various applications in detail, starting with the duopoly example discussed earlier in this introduction. We examine the impact of the information structure on consumer and producer surplus. With increasing returns to scale, full information (leading to the Stackelberg outcome) maximizes producer surplus. We characterize the Pareto frontier, the class of information structures maximizing consumer surplus, and show that the maximum thus attained can represent a considerable improvement both on the Stackelberg and Cournot outcomes.

We then explore the implications of our analysis in the context of the provision of a public good. If individuals, teams, or countries, sequentially decide how much to invest in a public good, what is the optimal way to disclose information about the leader's contribution in order to maximize the sum of all contributions? We show that the optimal policy takes a simple form: tell the follower whether the leader's contribution is inside or outside an intermediate interval whose lower bound is greater than the Cournot contribution. The idea is as follows. Suppose the follower learns that the leader contributed outside of the interval. In that case, the follower believes that the leader contributed an amount above the interval and, consequently, chooses to contribute nothing. Knowing this, the leader picks a contribution inside the interval, thus contributing over and above the Cournot level. As the slope of the follower's reaction curve is below one, the resulting sum of contributions is larger than the sum obtained in the Cournot outcome. We also discuss the implications of our analysis in the context of R&D with technological spillovers.

The rest of the paper is organized as follows. The model is presented in Section 2. Section 3 introduces additional notation and gathers a number of preliminary remarks. Section 4 illustrates the basic workings of the model by way of a numerical example. Section 5 contains the general analysis and our main results; that section ends with a cookbook procedure to identify all plausible outcomes in any given application. We apply the results of our analysis in Section 6, where we revisit, respectively, textbook models of duopoly, public good, and R&D. Section 7 concludes.

Related Literature Questions related to the information structure of sequential games stretch back at least to Schelling (1960). Daughety (1990) and Varian (1994) compare situations in which the follower observes perfectly the action of the leader to situations in which the follower observes nothing (in contexts of oligopolies and public goods, respectively); in Bagwell (1995), on the other hand, the follower either perfectly observes the action of the leader, or observes the realization of a random variable that has full support over a fixed set of signals for every possible action that the leader might select. Makris and Renou (2018) neither places any restriction on the base game nor any restriction on the information structure, and provides the most widely applicable revelation principle we know of.[3] The design of information structures in dynamic environments is studied in Albano and Lizzeri (2001) in the context of quality certification, and in the work of Calzolari and Pavan (2006) exploring the problem of two principals contracting with an agent. This research was recently revived through the work of Zapechelnyuk (2020) and Saeedi and Shourideh (2020). Our notion of plausibility is directly inspired by Zapechelnyuk (2020). To the best of our knowledge, the present paper is the first to characterize, for a broad class of settings, the set of outcomes resulting from considering *all* possible partitions of the leader's action space.

The general approach of our paper mirrors that of Bergemann, Brooks and Morris (2015); the nature of our problem is, however, very different from theirs: in Bergemann *et al.* (2015) the information structure pertains to an exogenous state; by contrast, in our setting, the information structure pertains to an endogenous object (the leader's action), and thus affects decisions both downstream and *upstream*. Boleslavsky and Kim (2018) explore an interesting, somewhat hybrid model. More generally, our paper is connected in spirit to an important line of research taking a "base game" as given, and exploring what a designer may achieve by engineering various aspects of the actual game (including Nishihara (1997), Kamenica and Gentzkow (2011), Bergemann and Morris (2016), Salcedo (2017), Gallice and Monzón (2019), and Doval and Ely (2020)).

Our work is more broadly linked to a vast and influential literature studying information sharing in oligopolies, starting with Clarke (1983), Vives (1984), Gal-or (1986), and Shapiro (1986).[4] This literature sheds light on the way firms' access to information about their com-

[3]Our work also bears a connection with Forges (1986) and Myerson (1986). However, their mediator must elicit past actions. By contrast, in our model, the designer (if there is one) can observe the action which the leader takes. This distinction is key: for instance, in the kind of environments that we consider, Ui (2008) proved the existence of a single correlated equilibrium (which places probability one on the pure-strategy Nash equilibrium).

[4]See also Vives (1990), Sakai and Yamato (1989), Darrough (1993), Raith (1996), and Pae (2002), among

petitors' cost (alternatively, demand function) affects market outcomes. Much less is known about the effect of firms' information about their competitors' *strategy*.

2 Model

Base game. A *leader* and a *follower* respectively choose actions x and y in $\mathcal{X} := [\underline{b}, \overline{b}]$. We assume symmetric payoffs throughout the body of the paper; the asymmetric payoffs case is almost identical.[5] The payoffs are given by $u(x, y)$ for the leader and $u(y, x)$ for the follower, where u is a twice continuously differentiable function satisfying $u_{11} < 0$.[6] The latter assumption assures that $y \mapsto u(y, x)$ possesses a unique maximizer in $\mathcal{X}$, henceforth denoted by $R(x)$. We define

$$\phi(x) := R\big(R(x)\big).$$

Information structure. The information of the follower concerning the action of the leader is given by a partition Π of the action space $\mathcal{X}$, such that the follower observes the partition element $\pi(x)$ containing the leader's action x. Both the base game and the information structure are common knowledge. We refer to the partition comprising a single element as the uninformative partition; the partition comprising only singletons is referred to as the fully informative partition.

Regularity conditions. To maintain the tractability of the model, we focus on environments satisfying three regularity conditions: (RC1) $\phi(x^C) = x^C$ at a unique $x^C \in \mathcal{X}$; (RC2) u_{12} has fixed sign;[7] (RC3) u_2 has fixed sign. Condition (RC1) guarantees the existence of a unique Nash equilibrium in the game generated by the uninformative partition; in terms of the primitives of the model, condition (RC1) is for instance implied by $u_{11} < -|u_{12}|$.[8] Borrowing terminology from Industrial Organization, we refer to x^C (formally defined in (RC1)) as the Cournot action. Conditions (RC2) and (RC3) say that actions must either be strategic substitutes or strategic complements, and all externalities must either be positive or negative.

many others.

[5]We briefly discuss the asymmetric payoffs case in the online appendix.

[6]The differentiability assumption can be relaxed; all we need is for the payoffs to be strictly quasi-concave in own action.

[7]That is, either $u_{12}(x, y) > 0$ for all (x, y), or $u_{12}(x, y) < 0$ for all (x, y). The case in which $u_{12} = 0$ is uninteresting and we therefore ignore it.

[8]See Lemma 5 in the appendix. The condition given in the text is essentially the standard condition uncovered in Rosen (1965).

Plausible outcomes. Let $\beta(x) \in \mathcal{X}$ represent the follower's belief concerning the action that the leader took, when the leader's true action is equal to x. A pair (Π, β) comprising a partition Π and a belief system $\beta : \mathcal{X} \to \mathcal{X}$ is said to be *admissible* if, for all $x \in \mathcal{X}$:

(A1) : $\beta(x) = \beta(x'), \forall\, x' \in \pi(x)$;

(A2) : $\beta(x) \in \pi(x)$;

(A3) : $u\Big(\beta(x), R\big(\beta(x)\big)\Big) \geq u\Big(x, R\big(\beta(x)\big)\Big)$.

Condition (A1) captures our assumption that the follower cannot distinguish between two actions belonging to the same partition element. Condition (A2) requires the belief $\beta(x)$ induced by x to remain within the partition element that contains x. Condition (A3) may be understood as follows: consider $x' \in \pi(x)$, so that the follower cannot distinguish between actions x and x'; suppose moreover that, given the belief system β, choosing x' yields a strictly lower payoff to the leader than choosing x; then (A3) requires $\beta(x) \neq x'$, that is, upon observing the partition element $\pi(x)$, the follower cannot believe that the leader chose action x'.

A pair (Π, β) is said to *implement* x^* if

(B1) : (Π, β) is admissible;

(B2) : $\beta(x^*) = x^*$;

(B3) : $x^* \in \arg\max_{x \in \mathcal{X}}\ u\Big(x, R\big(\beta(x)\big)\Big)$.

Condition (B2) requires that the follower's belief be x^* when the leader picks x^*; condition (B3) ensures that the leader cannot profitably deviate from x^*. We can now formalize the central concept of our paper.

Definition 1. *An action x^* is* plausible *if a pair (Π, β) implements x^*. An outcome (x^*, y^*) is plausible if x^* is plausible and $y^* = R(x^*)$.*

As is readily checked, an outcome (x^*, y^*) is plausible if there exist a partition Π and a sequential equilibrium $\mathcal{E}$ of the extensive-form game which this information structure induces, such that the pair (Π, β) formed with the belief system β induced by $\mathcal{E}$ is admissible, and (x^*, y^*) is the outcome of $\mathcal{E}$.[9]

[9] Note that (x^*, y^*) is plausible if and only if there exists a partition Π such that (x^*, y^*) is a subgame perfect Nash equilibrium outcome of an extensive-form game in which the leader first picks a member of Π, and then an action x belonging to that partition element.

We say that an outcome is plausible under interval partitions if, in Definition 1, Π is an interval partition (that is, all elements of Π are intervals). Throughout, we use the acronym IP for interval partition.

3 Preliminaries

To streamline the exposition, we introduce in this section some additional notation. First, we define the leader's payoff associated with action x, conditional on the follower best-responding to x:

$$U(x) := u\big(x, R(x)\big).$$

We denote the upper contour set of x with respect to U by

$$\mathcal{Q}^U_{\geq}(x) := \big\{\tilde{x} :\ U(\tilde{x}) \geq U(x)\big\};$$

when chances of confusion are sufficiently small, we will often talk about the upper contour set of x, without explicit reference to U. The sets $\mathcal{Q}^U_{>}(x)$, $\mathcal{Q}^U_{\leq}(x)$, and $\mathcal{Q}^U_{<}(x)$ are similarly defined. Finally, let

$$\eta(\tilde{x}, x) := u\big(\tilde{x}, R(x)\big) - u\big(x, R(x)\big) = u\big(\tilde{x}, R(x)\big) - U(x).$$

Informally, $\eta(\tilde{x}, x)$ measures the leader's gain from deviating to $\tilde{x}$ assuming the follower's belief fixed at x. Note that condition (A3) can now be written as

$$\eta\big(x, \beta(x)\big) \leq 0. \tag{1}$$

We end this section with a simple lemma, which will play a key role in the process of identifying plausible outcomes.

Lemma 1. *A pair (Π, β) implements x^* if and only if it is admissible and*

$$x^* \in \arg\max_{x \in \beta(\mathcal{X})} U(x). \tag{2}$$

The proof is straightforward, and relegated to the appendix. Given an admissible pair (Π, β), we may think of $\beta(\mathcal{X})$ as representing the set of actions which the leader is able to "signal". Thus, we may interpret Lemma 1 as stating that, to be plausible, x^* must be optimal among the actions which the leader is able to signal under *some* admissible pair.

4 Example

In this section, we illustrate the basic workings of the model by way of a numerical example based on the textbook model of duopoly discussed in the introduction. Leader and follower are two identical firms, each with a cost of producing a quantity q given by $3q - (2/5)q^2$. The inverse demand function is given by $5 - Q$ (variable Q representing the total quantity produced). We set $\mathcal{X} = [0, 10/3]$.[10]

Letting $u(x, y)$ represent the profit of the leader,

$$u(x, y) = (2 - y)x - \tfrac{3}{5}x^2,$$

and so $u_{11} = -6/5 < -1 = u_{12}$. Regularity conditions (RC1)–(RC3) are therefore satisfied. Moreover, straightforward calculations yield $x^C = 10/11$.

[0,19/6)	x^C
[19/6,10/3]	19/6

(Π', β')

[0,3/2)	x^C
[3/2,19/6)	3/2
[19/6,10/3]	19/6

(Π'', β'')

$[0, 1/5] \cup (7/11, 19/6)$	1/5
(1/5,7/11]	7/11
[19/6,10/3]	19/6

(Π''', β''')

Table 1

We next consider three admissible pairs and examine the outcomes which they implement. Table 1 fully describes said pairs: in each panel, the left column lists the elements of the relevant partition, while the right column lists the corresponding beliefs of the follower. The binary partition Π' reveals only whether $x \geq 19/6$. Partitions Π'' and Π''' refine Π': Π'' reveals whether $x > 3/2$, while Π''' reveals whether $x \in (1/5, 7/11]$. Figure 1 illustrates: in every panel, each color is associated with a given partition element, and each colored circle represents the corresponding belief of the follower. The colored curves indicate $u\Big(x, R\big(\beta(x)\big)\Big)$. Notice that each colored circle maximizes the colored curve on which it lies; all three pairs are therefore admissible.

[10] Quantities larger than 10/3 would lead to negative profits no matter what.

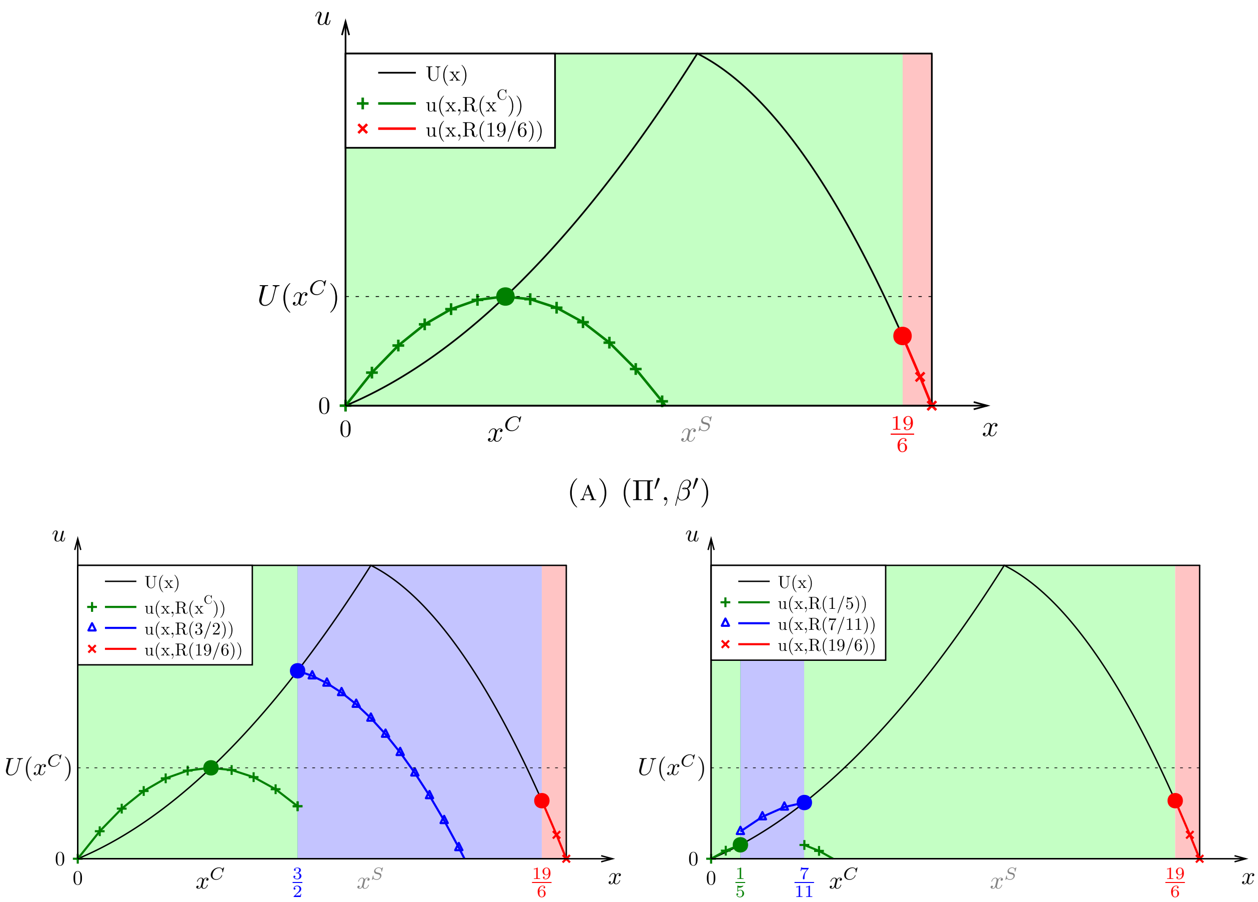

(A) (Π', β')

(B) (Π'', β'')

(C) (Π''', β''')

Figure 1: Duopoly Example

Using Lemma 1 enables us to read off from Figure 1 the action implemented by every pair: in each panel, pick the highest colored circle. The pair (Π', β') thus implements the Cournot action x^C, while the pair (Π'', β'') implements $3/2$, and the pair (Π''', β''') implements $19/6$. Noting that $U(3/2) > U(x^C) > U(19/6)$ enables us to make the following observations:

(i) (Π', β') and (Π'', β''), which both comprise an IP, implement actions belonging to the upper contour set of x^C;

(ii) (Π''', β'''), which does not comprise an IP, implements an action in the strict lower contour set of x^C.

We will see in the next section that observation (i) is an instance of a very general result: any action plausible under IPs must belong to the upper contour set of the Cournot action.

In fact, Theorem 1 in the next section will show that the converse holds as well: any action in the upper contour set of x^C is plausible under IPs. Observation (ii) is an instance of another general result: we show in Theorem 2 that the full set of plausible outcomes is typically larger than the corresponding set under IPs.

Finally, both Π'' and Π''' being finer partitions than Π', we see that increasing the amount of information available to the follower could make the leader better or worse off. This last observation refutes the intuitively appealing idea according to which increasing the amount of information available to the follower necessarily bolsters the leader's first-mover advantage. We will see that, in many cases, facing a more informed follower can be detrimental for the leader.

5 General Analysis

This section is divided into two subsections: Subsection 5.1 explores the set of plausible outcomes under IPs; Subsection 5.2 studies the full set of plausible outcomes.

5.1 Interval Partitions

We show first that, if Π is an interval partition, the admissibility criteria for a pair (Π, β) then take a very simple form: for each $x \in \mathcal{X}$, the belief $\beta(x)$ induced by x needs to be the action closest to x^C among all actions in the partition element containing x.

Lemma 2. *If Π is an interval partition, then (Π, β) comprises an admissible pair if and only if, for all $x \in \mathcal{X}$, conditions (A1) and (A2) hold, as well as $|\beta(x) - x^C| \leq |x - x^C|$.*

Proof: Let Π be an interval partition, and β a belief system such that (Π, β) satisfies (A1) and (A2) for all $x \in \mathcal{X}$. We prove below the "only if" part of the lemma, that is, we show that if, for all $x \in \mathcal{X}$, condition (A3) is satisfied, then condition $|\beta(x) - x^C| \leq |x - x^C|$ has to hold; the proof of the converse is similar and therefore omitted.

First, notice that, the best-response function R being continuous (by Berge's maximum theorem), $\phi(\cdot)$ is continuous as well. Hence:

$$\begin{cases} \phi(x) > x & \text{if } x < x^C, \\ \phi(x) < x & \text{if } x > x^C. \end{cases} \tag{3}$$

Second, $u(\cdot, R(x))$ being concave and maximized at $\phi(x)$, observe that $\eta(\cdot, x)$ is increasing over the interval $[\underline{b}, \phi(x)]$ and decreasing over $[\phi(x), \overline{b}]$. Combining the previous observations shows that:

(a) $x < x^C$ implies $\eta(x+\epsilon, x) > 0$ for all sufficiently small $\varepsilon > 0$;

(b) $x > x^C$ implies $\eta(x-\varepsilon, x) > 0$ for all sufficiently small $\varepsilon > 0$.

Now suppose that condition (A3) is satisfied for all $x \in \mathcal{X}$, that is, (1) holds for all $x \in \mathcal{X}$. Then property (a) gives $\beta(x) \neq x$ for every $x < x^C$ such that $\pi(x+\varepsilon) = \pi(x)$ for all sufficiently small $\varepsilon > 0$. Similarly, property (b) gives $\beta(x) \neq x$ for every $x > x^C$ such that $\pi(x-\varepsilon) = \pi(x)$ for all sufficiently small $\varepsilon > 0$. ■

We can now state our first main result:

Theorem 1. *The set of plausible actions under interval partitions coincides with the upper contour set of x^C with respect to U.*

Proof: Pick an arbitrary action x^* that is plausible under IPs. By Lemma 1, we can thus find an admissible pair (Π, β) where Π is an IP and such that (2) holds. On the other hand, Lemma 2 implies $\beta(x^C) = x^C$. Therefore, $U(x^*) \geq U(x^C)$.

Next, pick $x^* \in \mathcal{Q}^U_{\geq}(x^C) \setminus \{x^C\}$ (note that x^C is trivially plausible under IPs). To fix ideas, suppose $x^* > x^C$ (the argument is the same in the other case). Consider the interval partition Π comprising two partition elements: $\mathcal{X}_1 = [\underline{b}, x^*)$ and $\mathcal{X}_2 = [x^*, \overline{b}]$. Let β be the belief system given by $\beta(x) = x^C$ for all $x \in \mathcal{X}_1$, and $\beta(x) = x^*$ for all $x \in \mathcal{X}_2$. As $x^* > x^C$, applying Lemma 2 shows that (Π, β) constitutes an admissible pair. Since $U(x^*) \geq U(x^C)$, we see in view of Lemma 1 that $\big(x^*, R(x^*)\big)$ is plausible under IPs. ■

The Cournot outcome is by definition plausible under the uninformative partition. Therefore, one way of interpreting Theorem 1 is to say that, restricting attention to IPs: starting from the uninformative partition and increasing the amount of information available to the follower makes the leader (weakly) better off. Our next proposition generalizes the previous observation. Below, extending Definition 1, we say that x' is plausible under refinements (respectively, interval refinements) of Π if there exists a pair (Π', β') which implements x', where Π' is a partition (respectively, interval partition) finer than Π.

Proposition 1. *Consider an interval partition Π, and a belief system β such that (Π, β) implements action x^*. The set of plausible actions under interval refinements of Π coincides with the upper contour set of x^*.*

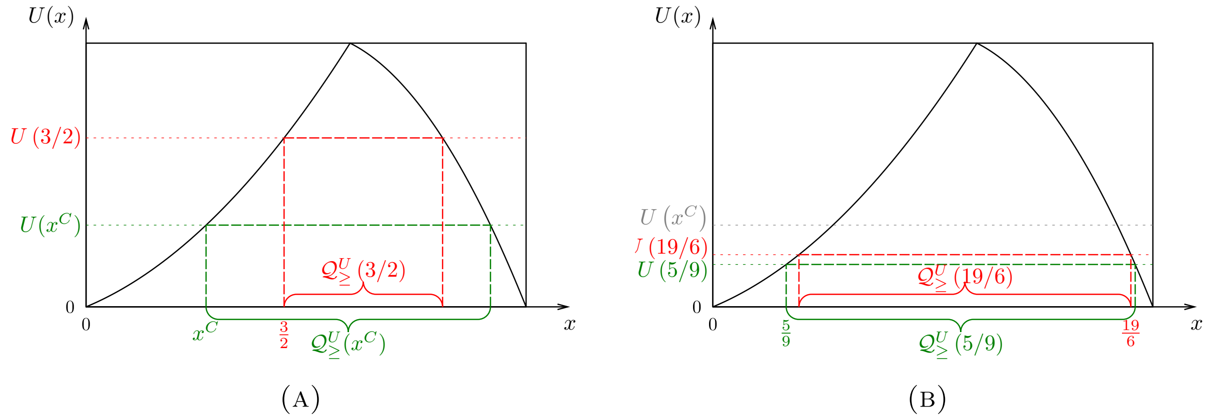

FIGURE 2: UPPER CONTOUR SETS

Proof: Let Π' be an interval partition finer than Π, and β' a belief system such that (Π', β') implements the action x', say. Then combining Lemmata 1 and 2 shows that $x' \in \mathcal{Q}^U_{\geq}(x^*)$. The proof that every action in the upper contour set of x^* is plausible under interval refinements of Π is similar to the second part of the proof of Theorem 1, and therefore omitted. ■

Figure 2, panel A, refers to the example from Section 4 and illustrates two upper contour sets: by Proposition 1, the upper contour set of x^C (respectively, 3/2) coincides with the set of plausible actions under interval refinements of Π' (respectively, Π''). Notice that, in this example, the Cournot action is the greatest lower bound of the plausible actions under IPs. The latter observation is an instance of a general result: a corollary of Theorem 1 is that the Cournot action is either the greatest lower bound of the plausible actions under IPs, or this set's least upper bound.[11] We note that the Cournot outcome is, in this sense, a poor predictor, for an external analyst possessing no knowledge of the true information structure.

5.2 Beyond Interval Partitions

Theorem 1 characterizes the set of plausible outcomes under IPs. We next explore the full set of plausible outcomes. For an action in $\mathcal{Q}^U_{<}(x^C)$ to be plausible, an admissible structure must exist that prevents the leader from signalling x^C; exploiting this observation leads to the following lemma:

[11]The Cournot action is the greatest lower bound if, as in our example, $u_2 \cdot u_{12} > 0$, and is the least upper bound otherwise.

Lemma 3. *Let*

$$\mathcal{S} := \big\{x:\ \eta(x^C, x) \le 0\big\} \cap \big\{x:\ u\big(x^C, R(x)\big) < U(x^C)\big\}.$$

Whenever $\mathcal{S} = \varnothing$, every plausible outcome is plausible under interval partitions.

Proof: We prove the contrapositive, namely, we show that if some plausible outcome is not at the same time plausible under IPs, then $\mathcal{S} \neq \varnothing$. Suppose such an outcome exists, write it $\big(x^*, R(x^*)\big)$, and pick an admissible pair (Π, β) such that (2) holds (we can do this, by Lemma 1). Then

$$\beta(x^C) \in\ \big\{x:\ \eta(x^C, x) \le 0\big\} \cap \mathcal{Q}^U_{\le}(x^*).$$

On the other hand, since x^* is not plausible under IPs, Theorem 1 yields $x^* \in \mathcal{Q}^U_{<}(x^C)$. Combining the previous remarks, we obtain $\beta(x^C) \in\ \big\{x:\ \eta(x^C, x) \le 0\big\} \cap \mathcal{Q}^U_{<}(x^C)$, which, in turn, yields

$$u\Big(x^C, R\big(\beta(x^C)\big)\Big) \le u\Big(\beta(x^C), R\big(\beta(x^C)\big)\Big) = U\big(\beta(x^C)\big) < U(x^C).$$

Finally, coupling the two highlighted expressions above shows that $\beta(x^C)$ belongs to $\mathcal{S}$. ■

We proceed to characterize the set of plausible outcomes when $\mathcal{S}$ is non-empty. We start by introducing some additional notation. It is readily seen that, for every $x \in \mathcal{S}$:

$$\eta\big(\gamma(x), x\big) = 0 \ \text{ for exactly one action } \gamma(x) \in \mathcal{X} \setminus \{x\};$$

we refer to $\gamma(x)$ as the *conjugate action* of x. When the follower's belief equals x, the leader is indifferent between choosing x or the conjugate of x. Whenever $\mathcal{S} \neq \varnothing$, let

$$\underline{U} := \inf_{\tilde{\gamma} \in \gamma(\mathcal{S})} U(\tilde{\gamma}).$$

To streamline the exposition, we assume that the previous infinimum is attained. We are now ready to state our second main result.

Theorem 2. *Assume that x^C is interior. If $\mathcal{S} = \varnothing$, the set of plausible actions coincides with the upper contour set of x^C with respect to U; otherwise, the set of plausible actions coincides with the upper level set of $\underline{U}$, i.e., $\{x : U(x) \ge \underline{U}\}$.*

The first part of the theorem is obtained by coupling Theorem 1 with Lemma 3; being longer and slightly more technical than previous proofs, the proof of the second part of the theorem is relegated to the appendix.

We show in the appendix that, whenever $\mathcal{S}$ is non-empty, $\underline{U}$ is generically strictly smaller than the Cournot payoff (i.e., $U(x^C)$). As the Cournot outcome is by definition plausible under the uninformative partition, we conclude from Theorem 2 that increasing the amount of information available to the follower might end up making the leader worse off. The next proposition records the counterpart of Proposition 1 for general partitions of the leader's action space.

Proposition 2. *Consider a partition* Π*, and a belief system* β *such that* (Π, β) *implements action* x^**. The set of plausible actions under refinements of* Π *is contained within the upper level set of* $\underline{U}$*, and contains the strict upper contour set of* x^**.*

Proof: The first part of the proposition follows from Theorem 2. Next, pick $x' \in \mathcal{Q}^U_>(x^*)$. Observe that, by Lemma 1, $x' \notin \beta(\mathcal{X})$. Let $\mathcal{X}_1$ denote the element of Π containing x', and Π' the partition obtained from Π when dividing up $\mathcal{X}_1$ into $\{x'\}$ and $\mathcal{X}_1 \setminus \{x'\}$. Finally, let β' be the belief system such that $\beta'(x) = \beta(x)$ for all $x \in \mathcal{X} \setminus \{x'\}$, and $\beta'(x') = x'$. One now easily verifies that (Π', β') implements x'. ■

Figure 2, panel B, returns to the example from Section 4. Straightforward calculations (in Online Appendix A) show that $\mathcal{S} = [0, x^C)$ and $\underline{U} = U(5/9)$. The figure illustrates two upper contour sets: by Proposition 2, the set of plausible actions under refinements of Π''' is a superset of the upper contour set of 19/6, and a subset of the upper contour set of 5/9.

In various applications, the way the leader's action space is divided up into observable subsets might be an object of design. The complexity of the partitions needed to generate specific outcomes then becomes central. We will say that a partition Π is a *quasi interval partition* if all but at most one partition elements of Π are intervals. We then say that an outcome is plausible under quasi interval partitions if, in Definition 1, Π is a quasi interval partition. We can now state this section's last result; the proof is in the appendix.

Proposition 3. *Every plausible outcome is plausible under quasi interval partitions.*

In the rest of the paper, we use the acronym QIP for quasi interval partition.

5.3 A Computational Cookbook Procedure

The previous section's analysis gives us the following cookbook procedure for computing all plausible outcomes in any given application:

Step 1. Compute the best-response function R, as well as x^C solving $\phi(x^C) = x^C$.

Step 2. Compute the set $\mathcal{S}$ as follows:

1. if $u_{12} \cdot u_2 > 0$ then $\mathcal{S} = \big\{x < x^C : \ \eta(x^C, x) \leq 0\big\}$;
2. if $u_{12} \cdot u_2 < 0$ then $\mathcal{S} = \big\{x > x^C : \ \eta(x^C, x) \leq 0\big\}$.

Step 3. Compute $U(x) = u\big(x, R(x)\big)$.

Step 4. Either $\mathcal{S}$ is empty, or it is not:

1. If $\mathcal{S} = \varnothing$, the set of plausible actions is then the upper contour set of x^C with respect to U.
2. If $\mathcal{S} \neq \varnothing$, first compute for each $x \in \mathcal{S}$ the conjugate action $\gamma(x) \neq x$ such that $\eta\big(\gamma(x), x\big) = 0$, then solve
$$\underline{U} = \inf_{\tilde{\gamma} \in \gamma(\mathcal{S})} U(\tilde{\gamma});$$
the set of plausible actions is the upper level set of $\underline{U}$.

Online Appendix A applies this procedure to the example from Section 4.

6 Applications

We pursue in this section our exploration of the duopoly example and apply our analysis to models of $R\&D$ and public good provision.

6.1 Quantity Competition

We consider here a general version of the duopoly example examined in Section 4. A firm producing quantity q faces cost $3q - \frac{c}{2}q^2$ and sells at price $4 - (1-d)Q - dq$, so that:

$$u(x, y) = x - (1-d)xy - \Big(1 - \frac{c}{2}\Big)x^2.$$

Parameter $d \in [0,1)$ measures the degree of product differentiation; $c \in (0, 1+d)$ determines the returns to scale.[12]

We first show that the set of plausible actions expands as d decreases and as c increases. To do so, we use the following lemma:

Lemma 4. *Assume (i) x^C is interior, (ii) $u_{22} = 0$, and (iii) $R'(x) = k$ for all x such that $R(x) \in \big(\underline{b}, \overline{b}\big)$. Then $\mathcal{S} \neq \varnothing$ if and only if $|k| > 1/\sqrt{2}$.*

Proof: Note that $\eta(x^C, x^C) = 0$ and $\frac{\partial \eta(x^C, x)}{\partial x}\Big|_{x=x^C} = -u_1(x^C, R(x^C)) = 0$. Moreover,

$$\frac{\partial^2 \eta(x^C, x)}{\partial x^2}\bigg|_{x=x^C} < 0 \iff \underbrace{-\frac{u_{12}(x^C, R(x^C))}{u_{11}(x^C, R(x^C))}}_{=R'(x^C)} R'(x^C) = [R'(x^C)]^2 > \frac{1}{2} \iff |k| > \frac{1}{\sqrt{2}}.$$

Hence, whenever the latter condition holds, we can find $\varepsilon > 0$ such that $\eta(x^C, x) < 0$ for all $x \in (-\varepsilon, \varepsilon)$. The latter remark in turn implies $\mathcal{S} \neq \varnothing$.

If $|k| \leq \frac{1}{\sqrt{2}}$ then for all x such that $R(x) \in (\underline{b}, \overline{b})$,

$$\frac{\partial^2 \eta(x^C, x)}{\partial x^2} = \underbrace{-u_{11}(x, R(x))}_{>0} \underbrace{(1 - 2k^2)}_{\geq 0} + k^2[u_{22}(x^C, R(x)) - u_{22}(x, R(x))].$$

As $u_{22} = 0$, the previous expression is non-negative. Moreover, we have $R'(x) = 0$ for all x such that $R(x) \in \{\underline{b}, \overline{b}\}$, and so

$$\frac{\partial^2 \eta(x^C, x)}{\partial x^2} = -u_{11}(x, R(x)) > 0.$$

We conclude that $\eta(x^C, x)$ is (weakly) convex in x for all x, and attains its minimum value of 0 at $x = x^C$. This shows that $\mathcal{S} = \varnothing$. ■

One readily verifies that conditions (i)-(iii) in Lemma 4 are all satisfied in the example of this subsection. It ensues that the full set of plausible actions is strictly larger than the set of plausible actions under IPs whenever d is sufficiently small or c sufficiently large.

Figure 3, panel A (respectively, panel B), depicts the sets of plausible actions given $c = 4/5$ (resp., $d = 0$) for all possible values of the parameter d (resp., c). The actions which are

[12] In Section 4, we fixed $d = 0$ and $c = 4/5$.

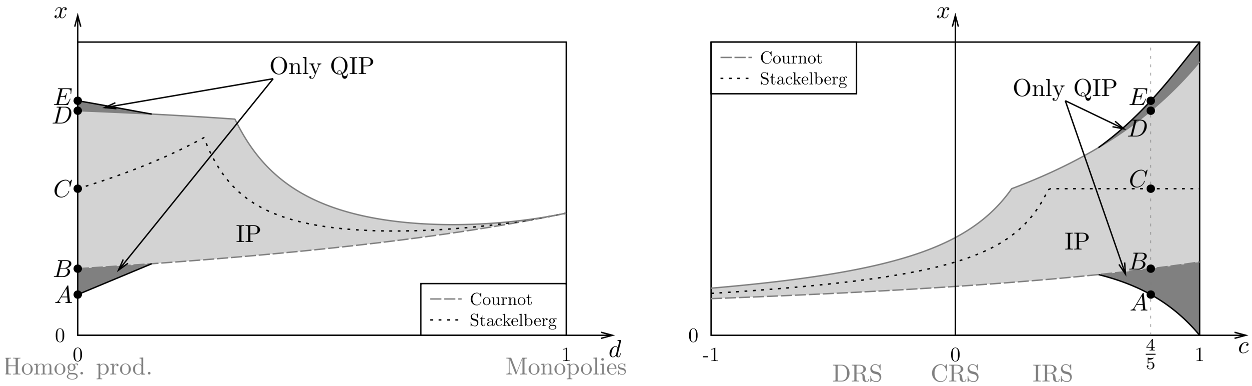

(A) Differentiation d (for $c = 4/5$) (B) Returns-to-scale c (for $d = 0$)

Figure 3: Plausible Leader Quantities

plausible under IPs are represented in light gray; the dark gray regions represent the actions which are plausible only under QIPs.

Consumer surplus (CS) and producer surplus (PS) can be computed from the action x^* which a pair (Π, β) implements:

$$CS(x^*) = \int_0^{x^*+R(x^*)} \big[P(z) - P\big(x^* + R(x^*)\big)\big]\, dz,$$
$$PS(x^*) = u\big(x^*, R(x^*)\big) + u\big(R(x^*), x^*\big).$$

Figure 4 shows the combinations of consumer surplus and producer surplus attainable for $d = 0$ and $c = 4/5$ as the information structure changes. Points A to E in the figure correspond to their counterparts in Figure 3. The section of the frontier between B and D corresponds to the outcomes that are plausible under IPs. The Pareto frontier coincides with the section joining C to E, and correspond to the outcomes in which only the leader produces a positive quantity. Total welfare, $CS + PS$, reaches a maximum at E.

6.2 Technological Spillovers

In this application, leader and follower are firms first making cost-reducing investments, then competing in quantities.[13] The firms invest sequentially and choose their quantities simultaneously. We are interested in the role of information with regard to the investment of the leader. Here, actions represent investments. The leader pays $x^2/2$ to ensure a per-unit production

[13]This example is from Shy (1996), section 9.3.1.

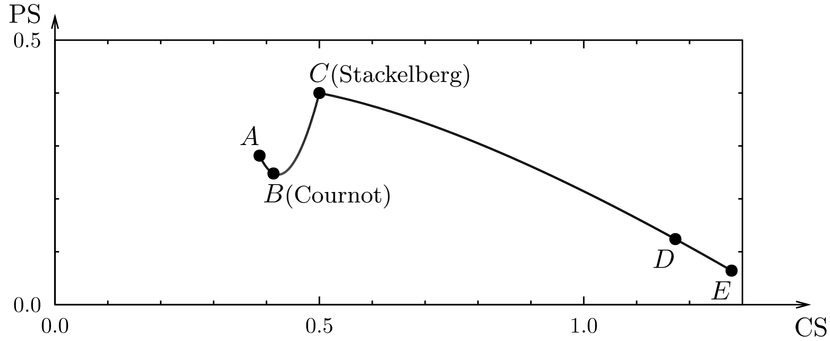

FIGURE 4: THE SURPLUS FRONTIER FOR $d = 0$ AND $c = 4/5$

cost equal to

$$c(x, y) = 50 - x - \beta y,$$

where $\beta \in \big((3 - \sqrt{7})/2, 1\big)$ measures the extent of the technological spillover. The inverse demand function is $p(Q) = 100 - Q$. Following investments $\{x, y\}$, the leader earns gross profits equal to:

$$\pi(x, y) = \frac{(50 + (2 - \beta)x + (2\beta - 1)y)^2}{9}.$$

Letting $u(x, y)$ represent the net profits:

$$u(x, y) = \pi(x, y) - \frac{x^2}{2}.$$

The investment of the competitor has here a direct positive externality due to the technological spillover, and an indirect negative externality, as lower costs make for tougher competition. For $\beta > 1/2$, the positive externality dominates. In this case, a larger investment by the competitor leads to more production, which, in turn, justifies greater investment: positive externality comes with strategic complementarity. For $\beta < 1/2$ instead, the negative externality dominates and investments are strategic substitutes.

In Figure 5 the different shades of gray have the same interpretation as in Figure 3. Note that the magnitude of the spillovers has a non-monotonic effect on the size of the set of actions plausible under IP. In this example, both $|R'(x)|$ and $|u_2(x, y)|$ are decreasing functions of β for $\beta < 1/2$ and increasing functions for $\beta > 1/2$. Spillovers thus reduce the externality and the size of the shifts in the best-response whenever $\beta < 1/2$, and have the opposite effect for $\beta > 1/2$. For $\beta = 1/2$, the two externalities cancel out, the competitor's choice is irrelevant,

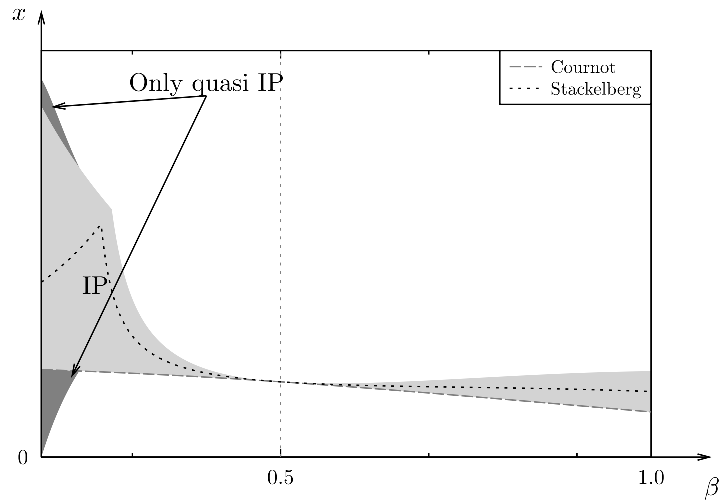

Figure 5: Plausible Leader Investments

and so is the information structure. Finally, in this example conditions (i)-(iii) in Lemma 4 hold, thus $\mathcal{S} \neq \varnothing$ whenever $|\beta - 1/2|$ is sufficiently large. In particular, there exists a threshold $\overline{\beta} < 1/2$ such that the set of plausible actions is strictly larger than the set of actions plausible under IPs for $\beta < \overline{\beta}$.[14]

6.3 Public Good Provision

In this subsection, leader and follower are endowed with a unit each of a good that can be consumed privately or contributed towards a public good. Letting actions represent contributions, $\mathcal{X} = [0, 1]$ and:

$$u(x, y) = \log(1 - x) + a \cdot \log(x + by),$$

where $a > 0$ captures the relative preference for the public good, and $b \in (0, 1]$ measures the degree of non-excludability of the public good.

The different shades of gray in Figure 6 have the same interpretation as in Figure 3. Panels A and B show, respectively, that the set of actions plausible under IPs expands both with b and a. In this example, $|u_2(x, y)| = ab/(x + by)$ and $|R'(x)| = b/(1 + a)$. Non excludability thus induces larger externalities and larger shifts in the best-response: the size of the light gray region is accordingly an increasing function of parameter b. A stronger relative preference for the public good obviously increases the magnitude of the externality, but discourages large

[14] To see this it is sufficient to note that $|R'| < 1/\sqrt{2}$ for $\beta = 1$.

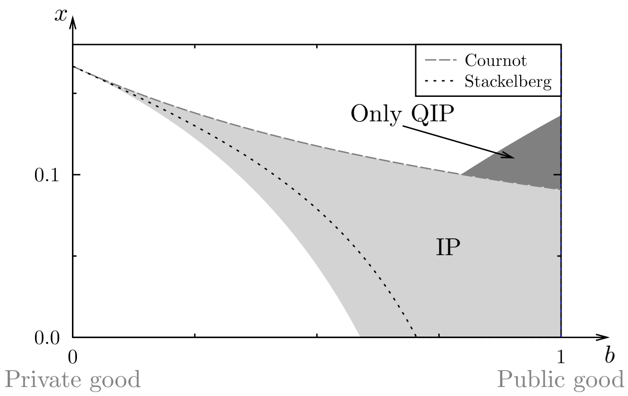

(A) Non-Excludability b (for $a = 1/5$)

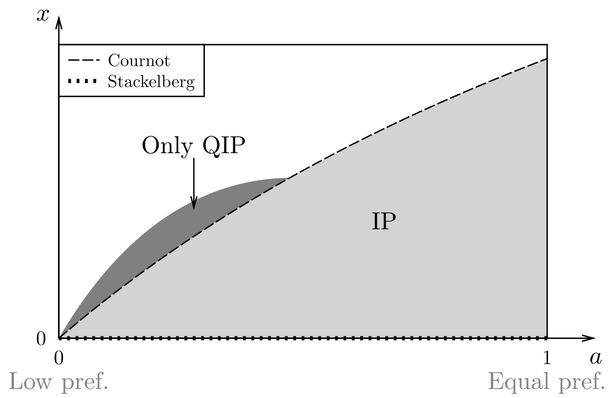

(B) Preference a (for $b = 1$)

Figure 6: Plausible Leader Contributions

shift in the best-response; for $b = 1$, the first effect dominates.[15]

6.3.1 The Designer Problem

We explore here the problem of a designer who selects a pair (Π, β) with the objective of maximizing the overall contribution to the public good. To fix ideas, leader and follower could be two divisions of the same company, each deciding autonomously how many resources to contribute towards a joint project. The designer could be a manager choosing the information structure with an eye to maximizing the joint contributions.

As the slope of the best-response function is larger than -1, the overall contribution is an increasing function of the leader's contribution. It is easy to verify that the Cournot action is the upper bound of the set of actions plausible under IPs. If all plausible actions are plausible under IPs, then the designer cannot improve on the uninformative partition.

Consider now parameters for which the set of plausible actions is strictly larger than the set of actions plausible under IPs. One can show that a contribution-maximizing partition exists that only reveals whether the contribution of the leader belongs to some interval of contributions larger than the Cournot one, or not. Rather than a general proof, we provide an example. Fix $b = 1$ and $a = 1/3$: Figure 7, Panel (A) presents an admissible pair (Π, β). Partition Π reveals whether $x \in [x^*, 1/3)$, where x^* is implicitly defined by $\eta(x^*, 1/3) = 0$, or not. For each partition element, the circle in Figure 7, Panel (B) maximizes the colored curve.

[15] This example violates condition (iii) in Remark 1. One can verify that the dark gray region is not empty whenever $\eta(x^c, a/b) < 0 \Leftrightarrow \left(1 - \frac{a(1+a)}{b(1+b)}\right)\left(\frac{(1+a)}{b} + 1\right)^a < 1$.

$[0, x^*) \cup [1/3, 1]$	$1/3$
$[x^*, 1/3)$	x^*

(A)

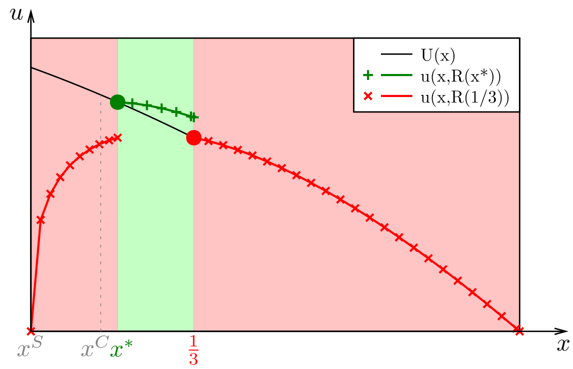

(B)

FIGURE 7: HIGHEST PLAUSIBLE x^*, FOR $a = 1/3$ AND $b = 1$

As the circle in correspondence of x^* is the highest, the leader picks action x^*. In words, if the follower learns that the leader contributed outside of the interval, he believes that the leader contributed a and, consequently, chooses to contribute nothing. Anticipating this, the leader picks a contribution inside the interval, thus contributing over and above the Cournot level. In fact, one can check that this partition maximizes the overall contribution.

7 Conclusions

We study the effect of varying the information structure in a broad class of leader-follower settings. Our model assumes that the follower only observes the element of a partition containing the action chosen by the leader. We characterize the set of outcomes resulting from any possible partition of the leader's action space. We show that, in some cases, any plausible outcome is plausible under interval partitions. In other cases, more complex partitions must be considered. We provide a cookbook procedure for computing all plausible outcomes in any given application. Several textbook models fall within the class of settings we analyze.

References

Albano, G. L. and Lizzeri, A. (2001) Strategic Certification and Provision of Quality, *International Economic Review*, **42**, 267–283.

Bagwell, K. (1995) Commitment and Observability in Games, *Games and Economic Behavior*, **8**, 271–280.

Bergemann, D., Brooks, B. and Morris, S. (2015) The Limits of Price Discrimination, *American Economic Review*, **105**, 921–957.

Bergemann, D. and Morris, S. (2016) Bayes Correlated Equilibrium and the Comparison of Information Structures in Games, *Theoretical Economics*, **11**, 487–522.

Boleslavsky, R. and Kim, K. (2018) Bayesian Persuasion and Moral Hazard, available at SSRN: 2913669.

Calzolari, G. and Pavan, A. (2006) On the Optimality of Privacy in Sequential Contracting, *Journal of Economic Theory*, **130**, 168–204.

Clarke, R. N. (1983) Collusion and the Incentives for Information Sharing, *Bell Journal of Economics*, **14**, 383–394.

Darrough, M. N. (1993) Disclosure Policy and Competition: Cournot vs. Bertrand, *Accounting Review*, **68**, 534–561.

Daughety, A. F. (1990) Beneficial Concentration, *American Economic Review*, **80**, 1231–1237.

Doval, L. and Ely, J. C. (2020) Sequential Information Design, *Econometrica*, **88**, 2575–2608.

Forges, F. (1986) An Approach to Communication Equilibria, *Econometrica*, **54**, 1375–1385.

Gal-or, E. (1986) Information Transmission—Cournot and Bertrand Equilibria, *Review of Economic Studies*, **53**, 85–92.

Gallice, A. and Monzón, I. (2019) Co-operation in Social Dilemmas Through Position Uncertainty, *Economic Journal*, **129**, 2137–2154.

Kamenica, E. and Gentzkow, M. (2011) Bayesian Persuasion, *American Economic Review*, **101**, 2590–2615.

Makris, M. and Renou, L. (2018) Information Design in Multi-stage Games, arXiv preprint arXiv: 2102.13482.

Myerson, R. B. (1986) Multistage Games with Communication, *Econometrica*, **54**, 323–358.

Nishihara, K. (1997) A Resolution of N-Person Prisoners' Dilemma, *Economic Theory*, **10**, 531–540.

Pae, S. (2002) Optimal Disclosure Policy in Oligopoly Markets, *Journal of Accounting Research*, **40**, 901–932.

Raith, M. (1996) A General Model of Information Sharing in Oligopoly, *Journal of Economic Theory*, **71**, 260–288.

Rosen, J. B. (1965) Existence and Uniqueness of Equilibrium Points for Concave N-Person Games, *Econometrica*, **33**, 520–534.

Saeedi, M. and Shourideh, A. (2020) Optimal Rating Design, arXiv preprint arXiv: 2008.09529.

Sakai, Y. and Yamato, T. (1989) Oligopoly, Information and Welfare, *Journal of Economics*, **49**, 3–24.

Salcedo, B. (2017) Implementation Without Commitment in Moral Hazard Environments, manuscript.

Schelling, T. C. (1960) *The Strategy of Conflict*, Harvard University Press, Cambridge, Mass.

Shapiro, C. (1986) Exchange of Cost Information in Oligopoly, *Review of Economic Studies*, **53**, 433–446.

Shy, O. (1996) *Industrial Organization: Theory and Applications*, MIT Press, Cambridge, Mass.

Ui, T. (2008) Correlated Equilibrium and Concave Games, *International Journal of Game Theory*, **37**, 1–13.

Varian, H. R. (1994) Sequential Contributions to Public Goods, *Journal of Public Economics*, **53**, 165–186.

Vives, X. (1984) Duopoly Information Equilibrium: Cournot and Bertrand, *Journal of Economic Theory*, **34**, 71–94.

Vives, X. (1990) Trade Association Disclosure Rules, Incentives to Share Information, and Welfare, *RAND Journal of Economics*, **21**, 409–430.

von Stackelberg, H. (1934) *Marktform und Gleichgewicht*, J. Springer.

Zapechelnyuk, A. (2020) Optimal Quality Certification, *American Economic Review: Insights*, **2**, 161–176.

Appendix

The following elementary remark will be useful: given an admissible pair (Π, β), the image of $\mathcal{X}$ under β consists of the fixed points of β. We state it here for future reference.

Remark 1. *If* (Π, β) *is admissible, then* $\beta(\mathcal{X}) = \{x \in \mathcal{X} : \beta(x) = x\}$.

Lemma 5 (Rosen (1965))**.** *If* $u_{11} < -|u_{12}|$ *then* $\phi(x^C) = x^C$ *at a unique* $x^C \in \mathcal{X}$.

Proof of Lemma 5: Let

$$G_u := \begin{bmatrix} u_{11} & u_{12} \\ u_{12} & u_{11} \end{bmatrix}.$$

Applying Theorem 2 in Rosen (1965), a sufficient condition for the existence of a unique x^C such that $\phi(x^C) = x^C$ is for $G_u + G_u^T$ to be negative definite. Yet, the matrix G_u being symmetric, the sum $G_u + G_u^T$ is negative definite if and only if G_u is negative definite. The eigenvalues of G_u are $u_{11} + u_{12}$ and $u_{11} - u_{12}$. Hence assuming $u_{11} < -|u_{12}|$ ensures that $\rho(x^C) = x^C$ at a unique $x^C \in \mathcal{X}$. ■

Proof of Lemma 1: The "only if" part of the lemma follows from Remark 1. We here prove the "if" part of the lemma. Consider an admissible pair (Π, β), and an action x^* such that (2) holds. Condition (B2) being implied by Remark 1, it only remains to show that condition (B3) holds as well.

Consider first $\tilde{x} \in \pi(x^*)$. Then $\beta(\tilde{x}) = \beta(x^*) = x^*$. By virtue of condition (A3), we thus obtain

$$u\Big(x^*, R\big(\beta(x^*)\big)\Big) = u\Big(\beta(x^*), R\big(\beta(x^*)\big)\Big) \geq u\Big(\tilde{x}, R\big(\beta(x^*)\big)\Big) = u\Big(\tilde{x}, R\big(\beta(\tilde{x})\big)\Big).$$

Next, pick $\tilde{x} \in \mathcal{X} \setminus \pi(x^*)$. The same logic as above shows that

$$u\Big(\beta(\tilde{x}), R\big(\beta(\tilde{x})\big)\Big) \geq u\Big(\tilde{x}, R\big(\beta(\tilde{x})\big)\Big),$$

so (2) yields

$$u\big(x^*, R(x^*)\big) \geq u\Big(\beta(\tilde{x}), R\big(\beta(\tilde{x})\big)\Big) \geq u\Big(\tilde{x}, R\big(\beta(\tilde{x})\big)\Big).$$

Yet, $x^* = \beta(x^*)$ (by Remark 1), hence $u\Big(x^*, R\big(\beta(x^*)\big)\Big) \geq u\Big(\tilde{x}, R\big(\beta(\tilde{x})\big)\Big)$. ■

Proof of Theorem 2: We will show the proof for the case of strategic complements with positive externalities ($u_{12} > 0$ and $u_2 > 0$); the proof is similar for all the other cases. Pick

$\hat{x} \in \mathcal{S}$ such that $U\big(\gamma(\hat{x})\big) = \underline{U}$. To shorten notation, let $\hat{\gamma} := \gamma(\hat{x})$. Part one of the proof will show that all actions in $\mathcal{Q}^U_{\geq}(\hat{\gamma})$ are plausible; part two will prove that all plausible actions belong to $\mathcal{Q}^U_{\geq}(\hat{\gamma})$.

The first part of the proof starts with two intermediary claims.

Claim 1: $\mathcal{S} \subseteq \{x : \; x < \gamma(x) \leq x^C\}$. Let x belong to $\mathcal{S}$. Reason by contradiction, and suppose that $x \geq x^C$. Note that, since $u_{12} > 0$, the best-response function R is a non-decreasing function. As $u_2 > 0$, we thus obtain $u\big(x^C, R(x)\big) \geq u\big(x^C, R(x^C)\big) = U(x^C)$, contradicting $x \in \mathcal{S}$. This shows that $x < x^C$. In turn, using (3) yields $\phi(x) > x$. Yet, $\eta(\cdot, x)$ is increasing over $[\underline{b}, \phi(x)]$, equal to 0 at x, and decreasing over $[\phi(x), \bar{b}]$. As (i) $\eta\big(\gamma(x), x\big) = 0$ (by definition of $\gamma(x)$), (ii) $\eta(x^C, x) \leq 0$ (since $x \in \mathcal{S}$), and (iii) $x^C > x$ (by the first step above), we then obtain $x < \gamma(x) \leq x^C$.

Claim 2: U is increasing over the interval $[\underline{b}, x^C]$. Consider $x < x^C$ and $\varepsilon > 0$ sufficiently small that $u(x+\varepsilon, R(x)) > u(x, R(x))$ (such an ε exists, by remark (a) in the proof of Lemma 2). Then, R being non-decreasing and $u_2 > 0$, we obtain

$$U(x+\varepsilon) = u\big(x+\varepsilon, R(x+\varepsilon)\big) \geq u\big(x+\varepsilon, R(x)\big) > u\big(x, R(x)\big) = U(x).$$

We are now ready to show that all actions in $\mathcal{Q}^U_{\geq}(\hat{\gamma})$ are plausible. Consider $x^* \in \mathcal{Q}^U_{\geq}(\hat{\gamma}) \setminus \mathcal{Q}^U_{\geq}(x^C)$ (we already know that all actions in $\mathcal{Q}^U_{\geq}(x^C)$ are plausible, by Theorem 1). Let $\mathcal{X}_1 := \{\hat{x}\} \cup \mathcal{Q}^U_{>}(x^*)$, and Π the partition of $\mathcal{X}$ such that $\pi(x) = \mathcal{X}_1$ for all $x \in \mathcal{X}_1$, and $\pi(x) = \{x\}$ for all x $\in \mathcal{X} \setminus \mathcal{X}_1$. Lastly, let β be the belief system given by

$$\beta(x) = \begin{cases} \hat{x} & \text{if } x \in \mathcal{X}_1, \\ x & \text{if } x \in \mathcal{X} \setminus \mathcal{X}_1. \end{cases}$$

The pair (Π, β) plainly satisfies conditions (A1) and (A2). Hence, by Lemma 1, to prove that x^* is plausible, it is enough to show that: (a) $\eta(\tilde{x}, \hat{x}) \leq 0$ for all $\tilde{x} \in \mathcal{X}_1$, (b) $x^* \in \arg\max_{x \in \{\hat{x}\} \cup (\mathcal{X} \setminus \mathcal{X}_1)} U(x)$. As $x^* \in \mathcal{Q}^U_{\geq}(\hat{\gamma})$, observe that any $\tilde{x} \in \mathcal{Q}^U_{>}(x^*)$ belongs to $\mathcal{Q}^U_{\geq}(\hat{\gamma})$. Thus, as $\gamma(\hat{x}) \leq x^C$ (Claim 1), every $\tilde{x} \in \mathcal{Q}^U_{>}(x^*)$ must satisfy $\tilde{x} \geq \gamma(\hat{x})$ (Claim 2). Exploiting the properties of $\eta(\cdot, \hat{x})$ in a manner analogous to what we did in the proof of Claim 1 then establishes property (a). We next show that (b) holds as well. Coupling Claims 1 and 2 yields $U\big(\gamma(\hat{x})\big) > U(\hat{x})$. Since $x^* \in \mathcal{Q}^U_{\geq}(\hat{\gamma})$, we thus obtain $U(x^*) > U(\hat{x})$. The latter remark evidently implies property (b), since $\mathcal{X} \setminus \mathcal{X}_1 \subseteq \mathcal{Q}^U_{\leq}(x^*)$.

The second part of the proof again starts with two intermediary claims.

Claim 3: If (Π, β) *implements* x^**, then* $\beta(x) < x$ *for all* $x \in \mathcal{Q}^U_{>}(x^*)$. Let (Π, β), x and x^* satisfy the conditions of the claim. Reason by contradiction, and suppose that $\beta(x) \geq x$. Then, R being non-decreasing, $R\big(\beta(x)\big) \geq R(x)$, giving in turn (since $u_2 > 0$),

$$u\Big(x, R\big(\beta(x)\big)\Big) \geq u\big(x, R(x)\big) > u\big(x^*, R(x^*)\big) = u\Big(x^*, R\big(\beta(x^*)\big)\Big);$$

the latter expression contradicts condition (B3).

Claim 4: If (Π, β) *is admissible and* $\beta(x) < \min\{x^C, x\}$*, then* $\gamma\big(\beta(x)\big) \in \big(\beta(x), x\big]$. Let (Π, β) and x satisfy the conditions of the claim. Then, by (3), $\phi\big(\beta(x)\big) > \beta(x)$. Moreover, $\eta\big(\cdot, \beta(x)\big)$ is increasing over the interval $\big[\underline{b}, \phi\big(\beta(x)\big)\big]$, equal to 0 at $\beta(x)$, and decreasing over $\big[\phi\big(\beta(x)\big), \overline{b}\big]$. As (i) $\eta\Big(\gamma\big(\beta(x)\big), \beta(x)\Big) = 0$ (by definition of $\gamma\big(\beta(x)\big)$), (ii) $\eta\big(x, \beta(x)\big) \leq 0$ (since (Π, β) is admissible), and (iii) $x > \beta(x)$ (by assumption), we then obtain $\beta(x) < \gamma\big(\beta(x)\big) \leq x$.

We are now ready to prove that all plausible actions belong to $\mathcal{Q}^U_{\geq}(\hat{\gamma})$. Reason by contradiction, and suppose that some plausible action x^* belongs to $\mathcal{Q}^U_{<}(\hat{\gamma})$. The combination of Claim 2, $\hat{\gamma} > \underline{b}$, and continuity of U, shows in this case that we can find $x^\dagger < \hat{\gamma}$ satisfying

$$x^\dagger \in \mathcal{Q}^U_{>}(x^*) \cap \mathcal{Q}^U_{<}(\hat{\gamma}).$$

Now let (Π, β) implement x^*. Then, Claim 3 shows that $\beta(x^\dagger) < x^\dagger$. In turn, since $x^\dagger < \hat{\gamma} \leq x^C$, using Claim 4 yields

$$\beta(x^\dagger) < \gamma\big(\beta(x^\dagger)\big) \leq x^\dagger < \hat{\gamma} \leq x^C.$$

Noting firstly that we can write $\mathcal{S} = \big\{x : \ \eta(x^C, x) \leq 0, \ R(x) < R(x^C)\big\}$, and secondly that R is strictly increasing over $[\underline{b}, x^C]$ (we make use here of the assumption that x^C is interior), we thus obtain[16]

$$\beta(x^\dagger) \in \mathcal{S}.$$

Yet, at the same time, Claim 2 coupled with $\gamma\big(\beta(x^\dagger)\big) < \hat{\gamma} \le x^C$ give

$$U\Big(\gamma\big(\beta(x^\dagger)\big)\Big) < U(\hat{\gamma}).$$

The combination of the last two highlighted expressions contradicts the definition of $\hat{\gamma}$. ■

Proof of Proposition 3: The proposition follows from the proof of Theorem 2. ■

[16]The argument is the same as in Claim 4. By (3), $\phi\big(\beta(x^\dagger)\big) > \beta(x^\dagger)$. Moreover, $\eta\big(\cdot, \beta(x^\dagger)\big)$ is increasing over $\big[\underline{b}, \phi\big(\beta(x^\dagger)\big)\big]$, equal to 0 at $\beta(x^\dagger)$, and decreasing over $\big[\phi\big(\beta(x^\dagger)\big), \overline{b}\big]$. As (i) $\eta\Big(\gamma\big(\beta(x^\dagger)\big), \beta(x^\dagger)\Big) = 0$ (by definition of $\gamma\big(\beta(x^\dagger)\big)$), (ii) $\eta\big(x^\dagger, \beta(x^\dagger)\big) \le 0$ (since (Π, β) is admissible), and (iii) $x^C \ge x^\dagger > \beta(x^\dagger)$ (by a previously highlighted expression), we then obtain $\eta\big(x^C, \beta(x^\dagger)\big) \le 0$.

Online Appendix A

In this appendix, we apply the cookbook procedure to the example from Section 4.

Step 1: The best-response to x is:

$$R(x) = \begin{cases} \frac{5(2-x)}{6} & \text{if } x \in [0,2], \\ 0 & \text{if } x \in (2, \frac{10}{3}] \end{cases},$$

and $\phi(x^C) = x^C \Leftrightarrow x^C = \frac{10}{11}$.

Step 2: As $u_{12} \cdot u_2 = x \geq 0$, then $\mathcal{S} = \left\{ x < x^C : \ \eta(x^C, x) \leq 0 \right\}$. Straightforward calculations show $\mathcal{S} = [0, x^C)$.

Step 3: Define

$$U(x) = u(x, R(x)) = \begin{cases} (1 + \frac{7x}{10})\frac{x}{3} & \text{if } x \in [0,2] \\ (2 - \frac{3x}{5})x & \text{if } x \in (2, \frac{10}{3}]. \end{cases}$$

Step 4: For any $x \in \mathcal{S}$

1. $\eta(\gamma(x), x) = 0 \Leftrightarrow \gamma(x) = \frac{10+7x}{18}$, hence $\gamma'(x) > 0$ and $\gamma(x) \in [0, x^C)$;
2. $U'(x) > 0$.

These observations imply $\frac{dU(\gamma(x))}{dx} > 0$, which, in turn, yields $\underline{U} = U(\gamma(0)) = U\left(\frac{5}{9}\right)$. Solving $U\left(\frac{5}{9}\right) = U(x)$ for $x > x^C$ yields $x = \frac{1}{3}\left(5 + \sqrt{\frac{95}{6}}\right)$; the upper level set of $\underline{U}$ is then equal to $\left[\frac{5}{9}, \frac{1}{3}\left(5 + \sqrt{\frac{95}{6}}\right)\right]$.