SCUOLA DI DOTTORATO
UNIVERSITÀ DEGLI STUDI DI MILANO-BICOCCA

Dipartimento di
Economia, Metodi Quantitativi e Strategie D'impresa

Dottorato di Ricerca in ECONOMIA E FINANZA PUBBLICA Ciclo XXX

Curriculum in ECONOMICS

ESSAYS ON INFORMATION ECONOMICS

KARAKOÇ PALMINTERI GÜLEN

Matricola: 798632

Supervisor: Prof. PICCOLO SALVATORE

Coordinator: Prof. MANERA MATTEO

Academic Year 2017/2018

To my parents who made me curious,
and to my husband who made me happy.

Acknowledgments

First and foremost, I am particularly indebted to my doctoral adviser Prof. Salvatore Piccolo without whom this thesis would have only been a dream. I am very grateful for his unwavering support and consistent guidance over these years, especially during tough times. I truly learned a great deal about economics from him and for that no amount of thanks is enough.

I would like to express my gratitude to Prof. Marco Pagnozzi and Prof. Giovanni Immordino, my co-authors, for the opportunity to work together. Thank you so much for taking the time to share your knowledge with me.

I am also grateful to Prof. Giacomo Calzolari, Prof. Nenad Kos and Prof. Giovanni Ursino for giving insightful suggestions and constructive criticism to improve this dissertation.

My sincere thanks also goes to Prof. Emanuele Tarantino for hosting me as a guest Ph.D. student at the University of Mannheim. I am grateful for the people of CDSE, faculty and PhD students, for their kindness, hospitality and especially for offering me a stimulating work environment.

Special thanks goes to my friends Aldo, Elena, Eva, Ginevra, Idil, Ivana, Okan, Özlem, Myrto and Nihan for their continuous encouragement, support and for all the energy they gave me in times of need.

I am deeply thankful to my parents, Şule and Yaman, for playing great role in my life and for all the sacrifices they made for me. They have been the best parents I could ever imagine. Many thanks to my brother Doğan and my sister-in-law Özlem for all the support no matter what, and to my parents-in-law, Mario and Letizia, for all the emotional support.

Last but not least, comes my husband Vincenzo. Thank you for being so loving, supportive, encouraging and patient husband. Your presence in my life made every moment enjoyable.

— Teşekkürler —

Co-Authorship Disclaimer

The first chapter of this dissertation is a joint work with Prof. Salvatore Piccolo, Department of Economics, University of Bergamo and CSEF, and with Prof. Marco Pagnozzi, Department of Economics, University of Naples Federico II and CSEF. We contributed equally to the writing and to the equilibrium analysis. A revised version of the working paper is presented at the Catholic University of Milan. The second chapter is a joint work Prof. Salvatore Piccolo, and with Prof. Giovanni Immordino, Department of Economics, University of Naples Federico II and CSEF. We contributed equally to the writing and to the equilibrium analysis. This chapter has been accepted for publication in *Revue Economique* and a revised version of the working paper presented at the University of Mannheim. The third chapter is based on my own work.

Abstract

This dissertation is composed of three chapters exploring the role of private information in various economic environments. In the first chapter, we study a dynamic vertical contracting environment in which a manufacturer deals with an exclusive retailer for two periods. In our model, a manufacturer designs a long-term contract with a retailer who is privately informed about demand and faces competition by an integrated entrant in the future. When demand is correlated across periods, information about past production affects firms' behavior after entry. We analyze the incentives of the incumbent players to share information with the entrant and show that the retailer benefits from transparency, but the manufacturer does not. Contrary to what intuition suggests, our results show that transparency with an integrated entrant harms consumers. When the entrant is not an integrated firm, whether transparency benefits consumers depends on the degree of demand persistency.

In the second chapter we study a simple mechanism design problem that describes the optimal behavior of a country targeted by a foreign terrorist group. The country is uncertain about the terrorists' strength and may decide to acquire such information from the community hosting the terrorists. We highlight a novel trade-off between target hardening — i.e., mitigating the incidence of an attack by strengthening internal controls and improving citizens' protection — and preemptive military measures aimed at eradicating the problem at its root — i.e., a strike in the terrorists' hosting country. We show that, conditional on being informed about the terrorists' strength, the country engages in a preemptive attack only when it faces a sufficiently serious threat and when the community norms favoring terrorists are weak. Yet, in contrast with the existing literature, we show that it is optimal for the country to acquire information only when these norms are strong enough and when its prior information about the terrorists' strength is sufficiently poor.

The last chapter of this dissertation, my current work-in-progress, analyzes the effect of uncertain biases on strategic information transmission. We consider a simple cheap talk model in which an uninformed decision maker seeks advice from one or two partially informed experts whose biases are unknown to the decision maker. Two types of bias has been considered: an expert may have either a moderate bias or an extreme bias. We characterize a semi-revealing equilibrium, in which an expert with moderate bias reports truthfully his private information and an expert with an extreme bias reports the same message regardless of his private information. Interest-

ingly, we find that fully-revealing equilibrium with one expert may be informationally superior to semi-revealing equilibrium with two experts. Specifically, uncertainty over biases allows experts to lie relatively more often as compared to fully-revealing case, whereby reducing the information content of the messages. However, with two experts, the decision maker has a higher chance to get truthful information from one of the experts which, in turn, may provide more information than the one-expert communication does. The net effect on the decision maker's ex-ante expected profit depends on the probability that the decision maker believes the expert to be moderate — i.e., whether the expert's report is informative or not. This result suggests that getting second opinion may not be always helpful for decision making.

Contents

List of Figures

Chapter 1

The Value of Transparency in Dynamic Contracting with Entry

1.1 Introduction

Incumbents facing the threat of future entry often engage in anticompetitive practices that protect their market power.[1] Limit pricing, excessive patenting, capacity building, exclusive dealings and other forms of vertical restraints are well known examples of barriers to entry.[2] Incumbents may also use information disclosure as a strategic tool to protect their dominant position. Although information sharing among firms has been extensively studied in static models of oligopoly (see, e.g., Vives, 2006, for a survey), little is known on firms' incentives to share information in dynamic environments, where incumbents may strategically disclose or hide information to potential entrants. Even less is known on the interplay between these incentives and vertical contracting.

Do incumbents want to share their private information with future competitors? What are the effects of this form of communication on consumers? What is the role of vertical contracting?

We analyze a dynamic vertical contracting environment in which a manufacturer deals with an

Joint work with S. Piccolo, University of Bergamo and CSEF, and with M. Pagnozzi, University of Naples Federico II and CSEF. We would like to thank Marie-Laure Allain, Giacomo Calzolari, Vincenzo Denicolò, Nenad Kos, David Martimort, Patrick Rey, Marcello Puca, Jérome Pouyet, Yossi Spiegel, and the audiences of the workshops on Competition and Bargaining in Vertical Chains (Dusserdolf, 2017) and on Advances in Industrial Organization (Bergamo, 2017) and PhD seminar participants at Catholic University of Milan for insightful comments and suggestions.

[1]Late or sequential entry is common in many industries (e.g., Geroski, 1995). In concentrated markets, for example, entry often occurs only when regulatory intervention creates scope for competition. In emerging markets, high-quality innovators often choose to delay production, in order to gather additional information about demand and consumers' needs (e.g., Dutta *et al.*, 1995). In industries where economies of scale restricts competition, entry is often subsidized by public policy aimed at reducing market concentration (e.g., Bernheim, 1984).

[2]See, e.g., Aghion and Bolton (1987), Gilbert and Harris (1984), Hart *et al.* (1990), Hoppe (2002), Milgrom and Roberts (1982), Ordover *et. al.* (1990), and Ziss (1996).

exclusive retailer for two periods. In the first period the retailer is a monopolist in the downstream market, while in the second period it faces competition by an integrated entrant.[3] Firms compete *à la* Cournot by selling a homogeneous product whose demand is uncertain and, in every period, is privately observed by the downstream players — i.e., by the retailer in both periods and by the entrant in the second period. The manufacturer designs a long-term contract to elicit the retailer's private information and, since demand is correlated over time, the retailer's second-period production depends on its report in the first period (that the manufacturer uses to update its beliefs about demand) — i.e., the optimal dynamic contract features memory (e.g., Baron and Besanko, 1984; Laffont and Tirole, 1996; Battaglini, 2005).

We show that firms' profits depend on the information about the first period possessed by the entrant. One may wonder why the entrant should care about this information, since it observes demand in the second period, when it chooses production. The reason is that, because the entrant does not observe demand in the first period, information about the retailer's first-period report (or directly about first-period production) affects the entrant's production and competition.[4] The long-term interaction between the incumbent players creates a contractual link between periods, and the role of information sharing hinges on this link.

Our main result is that the manufacturer and the retailer have diverging incentives to share information: the retailer would like to commit to inform the entrant about its first-period report, while the manufacturer has no incentive to disclose this information. The reason is that the manufacturer's profit and the retailer's information rent are affected by the entrant's production in opposite ways.

When the entrant is informed, it faces lower uncertainty about the incumbent's production and, hence, it produces a relatively less volatile output (as a function of second-period demand). This weakens the *competition effect* (highlighted in Martimort, 1996) and increases the retailer's rent in the second period compared to a situation without information sharing. By contrast, the manufacturer has no incentive to share information for two reasons. First, sharing information is detrimental to the manufacturer because it increases the retailer's rent. Second, *ceteris paribus*, the incumbent's profit is lower with information sharing because the entrant is more aggressive on average when it is informed: a business stealing effect.[5] Moreover, the manufacturer does not disclose information to the entrant even if it can sell it, since the entrant's willingness to pay for information is always lower than the manufacturer's reservation price.

[3]For example, the long-term relationship between the incumbent firms may arise because of the need to use a retailer with specific skills to customize a new product, requiring a fixed investment that is not worth paying for one period only. By contrast, an entrant may not need a specialized retailers once the product 'standard' has been developed by the incumbent. We also consider entry by a nonintegrated firm — see below.

[4]Notice that the entrant is only interested in information that signals the quantity produced by the retailer in the second period (like its report or the quantity produced in the first period), rather than information about past demand *per se*.

[5]This echoes the findings of the literature on information sharing in oligopoly — see, e.g., Gal-Or (1985), Li (1985), Shapiro (1986) and Vives (1984) — showing that with Cournot competition firms do not share information about demand because this increases correlation among their decisions, whereby reducing output and profits.

Sharing information reduces the incumbent's production because, other things being equal, it induces the entrant to increase production when the incumbent distorts it for rent extraction reasons. On balance, however, aggregate production is lower with information sharing because, holding constant the incumbent's production, the entrant's production decision is always efficient regardless of its information (since the entrant equalizes marginal revenue to marginal cost). In other words, although information sharing rebalances production between the incumbent and the entrant, it reduces market efficiency because it increases the retailer's rent via the competition effect. Therefore, consumer surplus and welfare are always lower with information sharing and, contrary to what is commonly believed, with an integrated entrant a welfare maximizing policy should reduce transparency and forbid incumbents to disclose information to entrants.

Our results hold even when the incumbents can disclose noisy information to the entrant, by only letting the entrant observe an imperfect signal (whose precision is chosen by the incumbents) of the retailer's first-period report or production.

We also extend our analysis by considering entry by a manufacturer that sells through its exclusive retailer, rather than by an integrated firm.[6] In this case, there are two competing hierarchies in the market in the second period. As in our main model, the new entrant does not observe demand in the first period, but demand in the second period is observed by all retailers, so that both manufacturers need to design contracts to elicit truthful information from them.

We show that, as with an integrated entrant, the incumbent retailer prefers to share information about the retailer's first-period report, while the incumbent manufacturer does not want to do so. With competing hierarchies, however, transparency may increase consumer surplus and welfare because information sharing allows the entrant manufacturer to elicit information from its retailer at a lower cost, which increases the entrant's output and tends to increase market efficiency. Since the incumbent reacts by reducing its own production, the effect of information sharing on aggregate production depends on the degree of demand persistency, which affects the entrant's information about first-period demand without information sharing. Hence, with competing hierarchies transparency has an ambiguous effect on welfare. Moreover, in this case the incumbent may have an incentive to sell information to the entrant.

Summing up, our analysis suggests that, in the presence of dynamic vertical contracting, disclosing information to entrants does not necessarily increase competition and consumer surplus. Hence, from a normative point of view, welfare may actually be reduced by mandatory disclosure rules forcing firms to adopt transparency standards that reveal information to potential entrants (see, e.g., Oh and Park, 2017).

The rest of the paper is organized as follows. After discussing the existing literature, Section 1.2 describes the baseline model and Section 1.3 analyzes the entrant's problem and discusses benchmarks without asymmetric information and without entry in the second period. Section 1.4

[6]For example, the entrant may be a foreign firm that needs a local retailer in order to enter the market and distribute its product.

provides the equilibrium analysis with and without information sharing. In Section 1.5, we describe the incumbents' incentives to share information and in Section 1.5.1 we analyze a market for information. Welfare is discussed in Section 3.6. We then consider various extensions: Section 1.7.1 considers competing hierarchies, Section 1.7.2 analyzes stochastic disclosure rules, Section 1.7.3 discusses ex post information sharing, and Section 1.7.4 extends the analysis to large uncertainty. The last section concludes. All proofs are in the Appendix.

Related Literature. We build on and contribute to three strands of literature. First, our paper relates to the literature on dynamic contracting with types correlated over time and full commitment by the principal. Baron and Besanko (1984) first characterized optimal contracts in a two-period environment. Laffont and Tirole (1996) applied dynamic contracting with adverse selection to the regulation of pollution rights and provided an interpretation of the optimal mechanisms in terms of markets with options. More recently, stemming from Battaglini (2005), the literature has evolved to multi-period models (both with discrete and continuous types) to investigate the memory and complexity of optimal dynamic contracts after long histories, convergence to efficiency, the effects of learning by doing, risk aversion and renegotiation, the limits of the 'first-order' approach, the impact of dynamics and enforcement risk on the contract incompleteness and stationarity (Arve and Martimort, 2016; Battaglini and Coate, 2008; Battaglini and Lamba, 2015; Garrett and Pavan, 2012; Gennaioli and Ponzetto, 2017; Eső and Szentes, 2017; Martimort *et al.*, 2017; Pavan *et al.*, 2014).[7] In our two-period model, most of these technical issues are not present: we chose to analyze a simple contracting environment to focus on the relationship between dynamics, transparency and product market competition.

Second, our analysis is related to the IO literature on information sharing in oligopoly. This literature shows that firms' incentives to share information about their common demand function (Novshek and Sonnenschein, 1982; Clarke, 1983; Vives, 1984; Gal-Or, 1985) or about their private costs of production (Fried, 1984; Gal-Or, 1986; Shapiro, 1986) depend on the nature of competition. Raith (1996) rationalizes the results of this vast literature in a unified framework. In contrast to our model, this literature typically assumes that firms are ex ante symmetric and play simultaneously. By introducing dynamic incentives and sequential entry, we introduce an endogenous asymmetry between firms that depends on the incumbent's contract and information sharing decision, which affect the entrant's behavior.[8] Hence, the novel and key feature of our environment is vertical contracting, which creates an endogenous relationship between information and competition. Without the contracting dimension, information sharing would play no role since firms' production in every period would only depend on current demand, which is observed by the entrant.

[7]See Bergemann and Pavan (2015) for a survey of the dynamic contracting literature.

[8]While in most of the existing literature on information sharing firms symmetrically exchange information and reciprocally learn about each other's characteristics, in our model the decision to share information is unilaterally taken by the incumbent, who has perfect information about the entrant.

Finally, we contribute to the literature on communication between vertical hierarchies, with endogenous information that principals have to obtain from privately informed retailers. Calzolari and Pavan (2006a) were the first to study this problem in a sequential contracting environment where principals may share the information obtained by contracting with a common agent (see also Calzolari and Pavan, 2006b, for a model with resale). They show how the information disclosed by one principal affects the contractual relationships between other players and analyze when a principal wants to offer full privacy to the agent.[9] When contracts are exclusive, Piccolo and Pagnozzi (2013) show that sharing information about costs affects contracting within competing organizations and induces agents' strategies to be correlated through the distortions imposed by principals to obtain information. In this environment, the incentives to share information depend on the nature of upstream externalities between principals and the correlation of agents' information.[10] In contrast to our model, both these papers focus on one-period relationships.

1.2 The Model

Players and Environment. Two incumbent players, a manufacturer M and its exclusive retailer R, contract for two periods. The manufacturer supplies a fundamental input to the retailer, which is used to produce a final good. There are constant returns to scale and the marginal cost of production is normalized to zero. In the first period $\tau = 1$, R is a monopolist in the downstream market. In the second period $\tau = 2$, an integrated firm E enters the market and firms compete by choosing quantities.[11] There are no entry costs.[12]

The inverse demand function in period $\tau = 1, 2$ is

$$P(\theta_\tau, Q_\tau) \triangleq \max\{0, \theta_\tau - Q_\tau\},$$

where Q_1 is R's production in the first period and $Q_2 \triangleq q_2 + q_E$ is aggregate production in the second period — i.e., the sum of R's second-period production q_2 and E's production q_E. The parameter $\theta_\tau \in \Theta \triangleq \left\{\underline{\theta}, \overline{\theta}\right\}$ is a measure of the magnitude of demand, with $\Delta\theta \triangleq \overline{\theta} - \underline{\theta} > 0$. The assumption of a linear demand function is standard in the literature on information sharing but it is not necessary for most of our results.[13]

[9] See also Bennardo *et al.* (2015) and Maier and Ottaviani (2009) for common agency models with moral hazard and communication.

[10] See also Piccolo *et al.* (2015) for a model with moral hazard and communication between competing hierarchies.

[11] In Section 1.7.1, we consider entry both in the upstream and in the downstream market — i.e., a non-integrated entrant.

[12] In Section 3.6, we discuss the implications of introducing (fixed) entry costs.

[13] These demand functions arise, for example, if in every period τ there is a representative consumer in the market with utility function $\theta_\tau Q_\tau - \frac{Q_\tau^2}{2} - p_\tau Q_\tau$, where p_τ is the market price.

Demand is correlated across periods. We assume that $\Pr\left[\theta_1 = \overline{\theta}\right] = \frac{1}{2}$, and let

$$\Pr\left[\theta_2 = \overline{\theta} | \theta_1 = \overline{\theta}\right] \triangleq \overline{\nu}, \qquad \Pr\left[\theta_2 = \underline{\theta} | \theta_1 = \underline{\theta}\right] \triangleq \underline{\nu},$$

where $(\overline{\nu}, \underline{\nu}) \in [0,1]^2$ and $\Delta\nu \triangleq \overline{\nu} - \underline{\nu}$. The parameters $\overline{\nu}$ and $\underline{\nu}$ can be interpreted as the degree of demand persistency: an increase in $\overline{\nu}$ (resp. $\underline{\nu}$) makes it is more likely that demand is high (resp. low) in the second period when it was high (low) in the first period. Notice that demand is positively correlated if

$$\Pr\left[\theta_2 = \overline{\theta} | \theta_1 = \overline{\theta}\right] > \Pr\left[\theta_2 = \overline{\theta} | \theta_1 = \underline{\theta}\right] \quad \Leftrightarrow \quad \overline{\nu} + \underline{\nu} > 1,$$

and it is negatively correlated otherwise.

In every period, R privately observes θ_τ while the manufacturer does not. The entrant observes θ_2, but not θ_1.

Contracts. We assume that M commits to a long-term contract with R. Following the literature — e.g., Baron and Besanko (1984) — we define a long term contract as a menu

$$\left\{q_1(m_1), t_1(m_1), q_2(m_2, m_1), t_2(m_2, m_1)\right\},$$

where, for every period τ, $m_\tau \in \Theta$ is R's report about θ_τ; while $q_\tau(\cdot)$ is the quantity produced by R and $t_\tau(\cdot)$ the transfer paid by R to M, both contingent on R's current and past reports.[14,15]

R is protected by limited liability in both periods, which avoids full surplus extraction in the second period (see, e.g., Laffont and Martimort, 2002)[16]. The contract is secret, so that E cannot directly observe it.

Communication. The incumbent players can disclose m_1 to E before market competition takes place in period 2. Following the IO literature — e.g., Vives (1984) and Raith (1996) among many others — in the baseline model we consider a '*all-or-nothing*' disclosure rule $d \in \{S, N\}$: either m_1 is fully disclosed to E ($d = S$) or it remains private information of R and M ($d = N$). In Section 1.7.2 we consider more general (stochastic) disclosure rules.

Assuming that firms can only share information about m_1 is without loss of generality in our framework, because of the contractual link between R's first-period report m_1 and its second-period production. Equivalently, this can be interpreted as incumbents disclosing the first-period

[14] In practice, this contract can be implemented by non-linear transfers $t_1(q_1)$ and $t_2(q_2, q_1)$, such that the second-period transfer depends on the first-period production.

[15] We do not consider more complex franchise contracts, like resale price maintenance (RPM), in order to avoid full extraction of R's surplus by M — see, e.g., Gal-Or (1991). For example, RPM contracts cannot be enforced when prices are too costly to verify. Even with RPM, however, information rents may still emerge if adverse selection is coupled with a moral hazard problem *à la* Laffont and Tirole (1986).

[16] Without limited liability, the retailer obtains no rent and the manufacturer implements the efficient outcome. Because of limited liability the retailer's payoff in any period must be non-negative for all m_1.

production to the entrant, which is a natural and realistic form of communication, since quantities are usually verifiable.[17] By contrast, the entrant is not interested in information about θ_1 *per se*, because the incumbent's production in the second period only depends on R's report on first-period demand. Of course, disclosing m_2 in the second period also has no effect since E directly observes θ_2.

As standard in the literature, we also assume that once an information sharing decision has been announced, it cannot be renegotiated after uncertainty about θ_1 and θ_2 realizes.[18] Commitment requires, for instance, the presence of a third party (such as a certification intermediary) that verifies communication.

Rather than assuming that a specific incumbent player chooses to share information, we first characterize the equilibrium under each disclosure policy, and then analyze the incumbents' private and joint incentives to share or sell information.

Timing and Profits. The timing of the game is as follows.

1. First period.
   - A disclosure policy $d \in \{S, N\}$ is announced.
   - R observes θ_1.
   - M offers a contract. If R accepts it, it reports m_1 to M.
   - Production occurs, and t_1 is paid.
2. Second period.
   - m_1 is disclosed if and only if $d = S$ and E updates its beliefs about θ_1.
   - R and E observe θ_2.
   - R reports m_2 to M.
   - Production occurs, and t_2 is paid.

All players are risk neutral and M and R discount future profit at a common discount factor $\delta \in (0, 1)$.[19] Hence, R's intertemporal payoff is

$$\sum_{\tau=1,2} \delta^{\tau-1} \left[P\left(\theta_\tau, Q_\tau\right) q_\tau - t_\tau\right],$$

[17]Given that, in practice, long-term contracts consist of menus that specify production in a period as a function of production in previous periods, disclosing information about quantity simply amounts to disclosing the contractual terms agreed between M and R.

[18]In Section 1.7.3 we consider secret renegotiations.

[19]This can be interpreted as a measure of the length of period 2 relative to period 1.

and M's intertemporal payoff is

$$\sum_{\tau=1,2} \delta^{\tau-1} t_\tau.$$

E's profit is $P(\theta_2, Q_2)\, q_E$.

Equilibrium. The solution concept is Perfect Bayesian Equilibrium (PBE). We focus on separating equilibria in which, for any disclosure policy: (i) M offers an incentive compatible contract; (ii) R accepts the contract and truthfully reports demand; (iii) quantity produced by firms in the second period are mutual best responses. We impose passive beliefs off equilibrium path, so that whenever R is offered an unexpected contract, it believes that E still follows its equilibrium strategy. This is a natural assumption since M's offer should not convey any information about E's behavior.

We make the following assumptions to simplify the analysis.

Assumption 1. Demand persistency is such that $\overline{\nu} \leqslant 4\underline{\nu} - 1$.

This assumption requires that the degree of demand persistency is neither too low when $\theta_1 = \underline{\theta}$ (i.e., $\underline{\nu} > \frac{1}{4}$) nor too high when $\theta_1 = \overline{\theta}$. These restrictions imply that R's equilibrium rent in the second period is always positive: a necessary condition for the game to feature a separating equilibrium in the second period.[20,21] Indeed, with pooling in the second period, information sharing becomes irrelevant since R's production in the second period does not depend on the first-period report.

Assumption 2. Demand uncertainty is small — i.e., $\Delta\theta \approx 0$.

This assumption is imposed to obtain closed form solutions when comparing the players' expected profit with and without information sharing. Essentially, the assumption allows us to compute (expected) payoffs by taking first-order Taylor approximations around $\Delta\theta = 0$.[22] We will show that information sharing has relevant welfare effects even in this limit case. In our dynamic framework, focusing on small uncertainty also implies that, in the first period, R only has an incentive to claim that demand is low when it is actually high, and not *vice versa*, yielding the standard 'no distortion at the top' result and that the incumbent never shuts down production. In Section 1.7.4 we consider the case of large uncertainty.

[20] See, e.g., Gal-Or (1999), Martimort (1996), and Kastl *et al.* (2011) for static models with similar assumptions.

[21] Notice that R's equilibrium rent in the second period is also positive if we impose positive correlation. For a similar approach, see Battaglini (2005).

[22] For a similar approach, see Laffont and Tirole (1988); Martimort (1999); and Martimort and Piccolo (2010).

1.3 Preliminaries

We first analyze E's behavior in the second period. Information sharing affects E's production since R's second-period production $q_2(\cdot)$ depends on its first-period report m_1.

Consider an equilibrium in which R truthfully reports demand in the first-period — i.e., such that $m_1 = \theta_1$ — and E expects R to produce $q_2(\theta_2, \theta_1)$ in the second period. With information sharing, E's problem is

$$\max_{q_E \geqslant 0} P(\theta_2, q_E + q_2(\theta_2, \theta_1))q_E,$$

whose solution yields a downward-sloping reaction function

$$q_E(\theta_2, q_2(\theta_2, \theta_1)) \triangleq \frac{\theta_2 - q_2(\theta_2, \theta_1)}{2}, \quad \forall (\theta_2, \theta_1) \in \Theta^2. \tag{1.1}$$

By contrast, with no information sharing, E must form a belief about θ_1 (which is equal to m_1 in equilibrium), given θ_2. Bayes' rule implies that E's posterior beliefs about θ_1 are

$$\Pr\left[\theta_1 = \overline{\theta} | \theta_2 = \overline{\theta}\right] = \frac{\overline{\nu}}{1 + \Delta\nu}, \quad \text{and} \quad \Pr\left[\theta_1 = \underline{\theta} | \theta_2 = \underline{\theta}\right] = \frac{\underline{\nu}}{1 - \Delta\nu}.$$

Hence, E's problem is

$$\max_{q_E \geqslant 0} \sum_{\theta_1} \Pr[\theta_1 | \theta_2] P(\theta_2, q_E + q_2(\theta_2, \theta_1))q_E,$$

whose solution yields a downward-sloping reaction function

$$q_E(\theta_2, \mathbb{E}[q_2(\cdot) | \theta_2]) \triangleq \frac{\theta_2 - \sum_{\theta_1} \Pr[\theta_1 | \theta_2] q_2(\theta_2, \theta_1)}{2}, \quad \forall \theta_2 \in \Theta. \tag{1.2}$$

The slope of this function depends on the degree of demand intertemporal correlation: the higher is the correlation, the more 'accurate' is E's inference on θ_1, given θ_2.

1.3.1 Benchmarks

Consider two useful benchmarks. First, suppose that, in every period, θ_τ is common knowledge. Then M fully extracts R's surplus and the optimal contract implements the monopoly outcome in the first period — i.e., $q^*(\theta_1) \triangleq \frac{\theta_1}{2}$ — and the symmetric Cournot outcome in the second period — i.e., both firms produce $q^C(\theta_2) \triangleq \frac{\theta_2}{3}$.

Second, suppose that there is no entry in the second period. Assume that in both periods only the incentive compatibility constraint of the high-demand type matters,[23] and let $U_1(\cdot)$ be R's equilibrium rent in the first period. Using a standard change of variables (e.g., Laffont and

[23]It can be checked that this is always the case under Assumption 2.

Martimort, 2002), M offers the contract that solves the following intertemporal problem:

$$\max_{q_1(\cdot),q_2(\cdot),U_1(\cdot)} \mathbb{E}\left[\sum_{\tau=1,2}\delta^{\tau-1}P\left(\theta_\tau,q_\tau\left(\cdot\right)\right)q_\tau\left(\cdot\right)\right]-\sum_{\theta_1}\Pr\left[\theta_1\right]\left[U_1\left(\theta_1\right)+\delta\Pr\left[\theta_2=\overline{\theta}|\theta_1\right]\Delta\theta q_2\left(\underline{\theta},\theta_1\right)\right],$$

subject to $U_1\left(\underline{\theta}\right)\geqslant 0$ and

$$U_1\left(\overline{\theta}\right)\geqslant \underbrace{U_1\left(\underline{\theta}\right)+\Delta\theta q_1\left(\underline{\theta}\right)}_{\text{Static rent}}+\underbrace{\delta\overline{\nu}\Delta\theta\left[q_2\left(\underline{\theta},\underline{\theta}\right)-q_2\left(\underline{\theta},\overline{\theta}\right)\right]}_{\text{Intertemporal rent}}. \tag{1.3}$$

It can be verified that both constraints bind and that, in the optimal dynamic contract, first-period quantities are

$$q_1^M(\overline{\theta})=q^*(\overline{\theta}),\quad \text{and}\quad q_1^M\left(\underline{\theta}\right)=q^*(\underline{\theta})-\frac{\Delta\theta}{2},$$

while second-period quantities are

$$q_2^M\left(\underline{\theta},\underline{\theta}\right)=q_1^M\left(\underline{\theta}\right)-\underbrace{\frac{1+\Delta\nu}{2\underline{\nu}}\Delta\theta}_{\text{Intertemporal distortion}},$$

and $q_2\left(\theta_2,\theta_1\right)=q^*\left(\theta_2\right)$ in all other states.

Hence, R always produces the monopoly quantity in a period in which demand is high — i.e., there is 'no distortion at the top'. By contrast, there is a standard (static) downward distortion of production in the first period when demand is low, while production in the second period is distorted only when demand is low in both periods. This intertemporal distortion arises because a higher quantity in state $(\underline{\theta},\underline{\theta})$ increases R's rent both in the second period and in the first period (since it makes it more attractive for R to report low demand in the first period, ceteris paribus).

The intertemporal distortion increases with $\overline{\nu}$ and decreases with $\underline{\nu}$. First, a high $\underline{\nu}$ implies a high probability of low demand in the second period following low demand in the first period, which reduces M's willingness to distort production. Second, other things being equal, a higher $\overline{\nu}$ increases R's intertemporal rent and induces M to increase quantity distortion to trade off efficiency and rent minimization.

1.4 Equilibrium Analysis

We now characterize the optimal contract offered by M without information sharing (Section 1.4.1) and with information sharing (Section 1.4.2).

1.4.1 No Information Sharing

With no information sharing, E's production depends on its expectation of the quantity produced by R, which depends on m_1 through the contract chosen by M (that E correctly expects in equilibrium). Let $q_E^N(\theta_2)$ be E's equilibrium production and denote by

$$\Delta q^N \triangleq q_E^N\left(\overline{\theta}\right) - q_E^N\left(\underline{\theta}\right)$$

the difference between E's production with high and low demand in the second period.

1.4.1.1 Retailer's Rent

Let $U_2(\cdot)$ be R's equilibrium rent in the second period. Following Martimort (1996), we first assume that R only has an incentive to under-report demand and then verify this conjecture ex post. Given a report m_1, R's relevant incentive and participation constraints in the second period are

$$U_2\left(\overline{\theta}, m_1\right) \geqslant U_2\left(\underline{\theta}, m_1\right) + \left(\Delta\theta - \Delta q^N\right) q_2\left(\underline{\theta}, m_1\right), \quad \forall m_1 \in \Theta,$$

$$U_2\left(\underline{\theta}, m_1\right) \geqslant 0, \quad \forall m_1 \in \Theta.$$

Since limited liability implies that $U_2\left(\underline{\theta}, m_1\right) = 0$ for every m_1, R's second period rent is

$$U_2\left(\overline{\theta}, m_1\right) \triangleq \underbrace{\Delta\theta q_2\left(\underline{\theta}, m_1\right)}_{\text{Information rent}} - \underbrace{\Delta q^N q_2\left(\underline{\theta}, m_1\right)}_{\text{Competition effect}}, \quad \forall m_1 \in \Theta. \tag{1.4}$$

This expression embeds two contrasting effects. First, R has an incentive to report low demand in the second period in order to pay a lower transfer. Other things being equal, this secures R a (standard) information rent which is increasing in the quantity produced when demand is low — see, e.g., Mussa and Rosen (1978) and Maskin and Riley (1985). Second, there is a *competition effect* (see, e.g., Gal-Or, 1999; Martimort, 1996; Martimort and Piccolo, 2010): when R under-reports demand in the second period, E produces more than M expects and the transfer offered to R does not take this effect into account. Hence, R's incentive to under-report demand is weaker than without entry. As a result, competition in the downstream market reduces R's information rent, and makes it less costly for M to elicit truthful information from R.

Consider now the first period. Let R's rent in the first period be

$$U_1\left(\theta_1, m_1\right) \triangleq P\left(\theta_1, q_1\left(m_1\right)\right) q_1\left(m_1\right) - t_1\left(m_1\right),$$

and $U_1(\theta_1) \triangleq U_1(\theta_1, m_1 = \theta_1)$, $\forall m_1 \in \Theta$. Taking into account its rent in the second period (1.4), R's intertemporal incentive constraint (ensuring that R truthfully reports demand in the first period) is

$$U_1(\theta_1) + \delta \Pr\left[\theta_2 = \overline{\theta}|\theta_1\right]\left(\Delta\theta - \Delta q^N\right) q_2(\underline{\theta}, \theta_1) \geqslant \\ U_1(\theta_1, m_1) + \delta \Pr\left[\theta_2 = \overline{\theta}|\theta_1\right]\left(\Delta\theta - \Delta q^N\right) q_2(\underline{\theta}, m_1), \quad \forall \theta_1, m_1 \in \Theta.$$

Assuming that the constraint only binds when demand is high,[24] the relevant first-period incentive compatibility constraint is

$$U_1\left(\overline{\theta}\right) \geqslant \underbrace{U_1(\underline{\theta}) + \Delta\theta q_1(\underline{\theta})}_{\text{Static Rent}} + \underbrace{\delta\overline{\nu}\left(\Delta\theta - \Delta q^N\right)\left[q_2(\underline{\theta}, \underline{\theta}) - q_2\left(\underline{\theta}, \overline{\theta}\right)\right]}_{\text{Intertemporal Rent}}, \tag{1.5}$$

while the relevant first-period participation constraint is $U_1(\underline{\theta}) \geqslant 0$.

R's incentive to under-report demand in the first period depends on two terms: the static rent of a single period relationship and the intertemporal rent that R obtains when demand is high in the second period, which happens with probability $\overline{\nu}$. The sign of this second term depends on how the first-period report affects production in the second period when demand is low. If $q_2(\underline{\theta}, \underline{\theta}) > q_2(\underline{\theta}, \overline{\theta})$, R's second-period rent is higher when it reports low rather than high demand in the first period and eliciting truthful information is more costly than in a static environment. By contrast, when $q_2(\underline{\theta}, \underline{\theta}) < q_2(\underline{\theta}, \overline{\theta})$, it is less costly for M to elicit truthful information. Of course, as the competition effect becomes stronger — i.e., as Δq^N increases — R's second-period rent decreases and the difference $q_2(\underline{\theta}, \underline{\theta}) - q_2(\underline{\theta}, \overline{\theta})$ has a weaker effect on R's first-period rent.

1.4.1.2 Optimal Long Term Contract

After a standard change of variables, M's intertemporal (relaxed) maximization problem is

$$\max_{q_1(\cdot), q_2(\cdot), U_1(\cdot)} \mathbb{E}\left[\sum_{\tau=1,2} \delta^{\tau-1} P\left(\theta_\tau, Q_\tau^N(\cdot)\right) q_\tau(\cdot)\right] + \\ -\sum_{\theta 1} \Pr[\theta_1]\left[U_1(\theta_1) + \delta \Pr\left[\theta_2 = \overline{\theta}|\theta_1\right]\left(\Delta\theta - \Delta q^N\right) q_2(\underline{\theta}, \theta_1)\right], \tag{1.6}$$

subject to (1.5) and the relevant participation constraint, where

$$Q_2^N(\theta_2, \theta_1) \triangleq q_2(\theta_2, \theta_1) + q_E^N(\theta_2).$$

Since both constraints bind at the optimum, it can be shown that first-period production is as in a market without entry, with no distortion at the top and downward distortion at the bottom — i.e., using the superscript N to denote the optimal quantities chosen by the manufacturer, $q_1^N\left(\overline{\theta}\right) = q^*\left(\overline{\theta}\right)$ and that $q_1^N(\underline{\theta}) = q_1^M(\underline{\theta})$.

[24] In the Appendix we check that under Assumption 2 this conjecture is verified in equilibrium.

Differentiating the objective function with respect to $q_2\left(\overline{\theta}, \theta_1\right)$ yields

$$P_{q_2}\left(\overline{\theta}, Q_2^N\left(\overline{\theta}, \theta_1\right)\right) q_2\left(\overline{\theta}, \theta_1\right) + P\left(\overline{\theta}, Q_2^N\left(\overline{\theta}, \theta_1\right)\right) = 0, \quad \forall \theta_1 \in \Theta, \tag{1.7}$$

where $P_{q_2}(\cdot)$ denotes the partial derivative with respect to q_2. Hence, when demand is high in the second period, R's production is not distorted (compared to the benchmark without incomplete information) regardless of the level of demand in the first period — i.e., $q_2^N(\overline{\theta}, \theta_1) = q^C(\overline{\theta})$ for every θ_1 — so that E's best response is $q_E^N(\overline{\theta}) = q^C(\overline{\theta})$.

Differentiating the objective function with respect to $q_2(\underline{\theta}, \overline{\theta})$ yields

$$P_{q_2}\left(\underline{\theta}, Q_2^N(\underline{\theta}, \overline{\theta})\right) q_2(\underline{\theta}, \overline{\theta}) + P\left(\underline{\theta}, Q_2^N(\underline{\theta}, \overline{\theta})\right) = 0. \tag{1.8}$$

Hence, if demand is high in the first period, the optimal dynamic contract rewards R in the second period even if demand is low in the second period — i.e., production is determined by the equalization of marginal revenues to marginal cost (which is normalized to zero).

Finally, differentiating the objective function with respect to $q_2\left(\underline{\theta}, \underline{\theta}\right)$ yields

$$P_{q_2}\left(\underline{\theta}, Q_2^N\left(\underline{\theta}, \underline{\theta}\right)\right) q_2\left(\underline{\theta}, \underline{\theta}\right) + P\left(\underline{\theta}, Q_2^N\left(\underline{\theta}, \underline{\theta}\right)\right) = \left(\Delta\theta - \Delta q^N\right) \frac{1 + \Delta\nu}{\underline{\nu}}. \tag{1.9}$$

As without entry, increasing the second-period output in state $(\underline{\theta}, \underline{\theta})$ has two effects: a higher $q_2\left(\underline{\theta}, \underline{\theta}\right)$ increases both R's second-period rent when demand is high in the second period and low in the first period, and R's intertemporal rent when demand is high in the first period. Both effects make it more profitable for R to under-report demand in the first period in order to enjoy higher rents in the future. Hence, a higher $\underline{\nu}$ reduces both the static and the intertemporal distortion, while a higher $\overline{\nu}$ increases the intertemporal distortion, which induces a higher distortion when demand is low in both periods.

Substituting $q_2^N(\overline{\theta}, \theta_1) = q_E^N(\overline{\theta}) = q^C(\overline{\theta})$ into (1.2), (1.8), (1.9) yields the following result.

Proposition 1.1. *Without information sharing,* $q_2^N\left(\underline{\theta}, \underline{\theta}\right) < q_2^N(\underline{\theta}, \overline{\theta}) < q^C\left(\underline{\theta}\right) < q_E^N\left(\underline{\theta}\right)$.

Hence, M distorts production downward (compared to a benchmark without incomplete information) when demand is low in both periods in order to optimally trade off efficiency and rent extraction. This distortion induces E to increase production when demand is low because E expects R to under-produce with positive probability. As a consequence, R faces a more aggressive competitor when demand is low in the second period, regardless of first-period demand, which induces it to reduce production.

1.4.2 Information Sharing

With information sharing, E's equilibrium production $q_E^S\left(\theta_2, m_1\right)$ depends both on demand in the second period and on M's report in the first period. This impacts R's second-period rent, and therefore it affects R's equilibrium production through the distortions chosen by M in order to trade off efficiency and (intertemporal) rent extraction. For any m_1, let

$$\Delta q^S(m_1) \triangleq q_E^S(\overline{\theta}, m_1) - q_E^S(\underline{\theta}, m_1)$$

be the difference between E's production with high and low demand in the second period.

1.4.2.1 Retailer's Rent

R's binding incentive compatibility constraint in the second period is[25]

$$U_2(\overline{\theta}, m_1) = \underbrace{\Delta\theta q_2(\underline{\theta}, m_1)}_{\text{Information rent}} - \underbrace{\Delta q^S(m_1)\, q_2(\underline{\theta}, m_1)}_{\text{Competition effect}}, \quad \forall m_1 \in \Theta.$$

In contrast to the case of no information sharing, the competition effect and, hence, R's intertemporal rent now depends on the effect of R's first-period report on E's production.

R's intertemporal incentive constraint is

$$U_1(\theta_1) + \delta \Pr\left[\theta_2 = \overline{\theta}|\theta_1\right]\left[\Delta\theta - \Delta q^S(\theta_1)\right] q_2(\underline{\theta}, \theta_1) \geqslant \\ U_1(\theta_1, m_1) + \delta \Pr\left[\theta_2 = \overline{\theta}|\theta_1\right]\left[\Delta\theta - \Delta q^S(m_1)\right] q_2(\underline{\theta}, m_1), \quad \forall \theta_1, m_1 \in \Theta.$$

As before, we assume (and verify ex post) that R only has an incentive to misreport demand when demand is high. Hence, the relevant first-period incentive compatibility constraint is

$$U_1(\overline{\theta}) \geqslant \underbrace{U_1(\underline{\theta}) + \Delta\theta q_1(\underline{\theta})}_{\text{Static rent}} + \underbrace{\delta\overline{\nu}\left[(\Delta\theta - \Delta q^S(\underline{\theta}))q_2(\underline{\theta}, \underline{\theta}) - (\Delta\theta - \Delta q^S(\overline{\theta}))q_2(\underline{\theta}, \overline{\theta})\right]}_{\text{Intertemporal rent}}, \tag{1.10}$$

while the relevant first-period participation constraint is $U_1(\underline{\theta}) \geqslant 0$.

Other things being equal, R's intertemporal rent is increasing in $\Delta q^S(\overline{\theta})$ and decreasing in $\Delta q^S(\underline{\theta})$. The stronger is the competition effect when a high demand is reported in the first period, the higher is R's intertemporal rent because second-period rents are higher (in equilibrium). By contrast, the stronger is the competition effect when a low demand is reported in the first period, the lower is R's intertemporal rent, which *ceteris paribus* reduces R's incentive to mimic in the first period.

[25] For simplicity, we use the same notation for R's rent as in Section 1.4.1.1.

1.4.2.2 Optimal Long Term Contract

After a standard change of variables, M's intertemporal (relaxed) maximization problem is

$$\max_{q_1(\cdot),q_2(\cdot),U_1(\cdot)} \mathbb{E}\left[\sum_{\tau=1,2}\delta^{\tau-1}P\left(\theta_\tau, Q_\tau^S(\cdot)\right)q_\tau(\cdot)\right]+$$
$$-\sum_{\theta_1}\Pr[\theta_1]\left[U_1(\theta_1)+\delta\Pr\left[\theta_2=\overline{\theta}|\theta_1\right]\left(\Delta\theta-\Delta q^S(\theta_1)\right)q_2(\underline{\theta},\theta_1)\right], \tag{1.11}$$

subject to (1.10) and the relevant participation constraint, where

$$Q_2^S(\theta_2,\theta_1)\triangleq q_2(\theta_2,\theta_1)+q_E^S(\theta_2,\theta_1).$$

Since both constraints bind at the optimum, it is easy to show that first-period production is the same as without information sharing, and that in the second period there is no distortion at the top regardless of level of demand in the first period — i.e., using the superscript S to denote the optimal quantities with information sharing, $q_2^S(\overline{\theta},\theta_1)=q_E^S(\overline{\theta},\theta_1)=q^C(\overline{\theta})$ for every θ_1.

Differentiating with respect to $q_2\left(\underline{\theta},\overline{\theta}\right)$ yields

$$P_{q_2}\left(\underline{\theta},Q_2^S(\underline{\theta},\overline{\theta})\right)q_2(\underline{\theta},\overline{\theta})+P\left(\underline{\theta},Q_2^S(\underline{\theta},\overline{\theta})\right)=0. \tag{1.12}$$

With information sharing, since E's output depends on first-period demand, neither firm distorts production when demand is low in the first period and high in the second — i.e., $q_2^S(\underline{\theta},\overline{\theta})=q_E^S(\underline{\theta},\overline{\theta})=q^C(\underline{\theta})$. Differentiating with respect to $q_2(\underline{\theta},\underline{\theta})$ yields

$$P_{q_2}\left(\underline{\theta},Q_2^S(\underline{\theta},\underline{\theta})\right)q_2(\underline{\theta},\underline{\theta})+P\left(\underline{\theta},Q_2^S(\underline{\theta},\underline{\theta})\right)=(\Delta\theta-\Delta q^S(\underline{\theta}))\frac{1+\Delta\nu}{\underline{\nu}}. \tag{1.13}$$

Hence, the effects of the intertemporal distortions on production when demand is low in both periods are as in the case of no information sharing.

Substituting and solving jointly with E's first-order condition (1.1) yields the following result.

Proposition 1.2. *With information sharing,* $q_2^S(\underline{\theta},\underline{\theta})<q^C(\underline{\theta})<q_E^S(\underline{\theta},\underline{\theta})$.

As intuition suggests, E produces more than R when demand is repeatedly low since M distorts production downward to reduce R's intertemporal rent. By contrast, when demand is high in the first period and low in the second period, firms produce the same quantities because M does not distort production.

[26]Notice that while we consider information about demand, similar effects arises with information about costs. In fact, information about θ_1 allows the entrant to learn whether R's production will be distorted, which is analogous to knowing whether a competitor has high or low cost of production.

1.5 Incentives to Share Information

To analyze the effects of information sharing on the manufacturer's and the retailer's expected profits,[26] we start by comparing E's production with and without information sharing. Since production is never distorted when demand is high, we can focus on the quantity produced by E when demand is low.

Proposition 1.3. *When demand is low in the second period, E's production is higher (lower) with information sharing than without if demand is low (high) in the first period — i.e., $q_E^S(\underline{\theta},\overline{\theta}) < q_E^N(\underline{\theta}) < q_E^S(\underline{\theta},\underline{\theta})$. Moreover, E's average production is higher with information sharing than without — i.e., $\mathbb{E}_{\theta_1}\left[q_E^S(\underline{\theta},\theta_1)\right] > q_E^N(\underline{\theta})$.*

With information sharing, E knows the quantity that R produces in the second period. When demand is low in both periods, R's second-period production is distorted for rent extraction reasons and, since reaction functions are downward sloping, information sharing allows E to produce more. By contrast, without information sharing, instead, E is uncertain about R's production and has a lower incentive to expand its own production. When demand is low in the second period and high in the first period, E's production is lower with information sharing because R's second-period production is not distorted, and this induces E to produce less when he is informed.

Therefore, with information sharing R faces tougher (weaker) competition from E when demand is low (high) in the first period. Without information sharing, E responds to uncertainty about R's production by producing an intermediate quantity, between $q_E^S(\underline{\theta},\overline{\theta})$ and $q_E^S(\underline{\theta},\underline{\theta})$. In expectation, information sharing increases E's production since the first effect discussed above dominates, so that the entrant obtains a larger market share when it is informed.

To analyze players' incentives to share information, we now compare R's expected rents and M's expected profits with and without information sharing. In order to obtain closed-form solutions, we restrict to the case of small uncertainty (by Assumption 2).[27]

Proposition 1.4. *R wants to share information with E, while M does not.*

In the limit of small uncertainty only the competition effect shapes the impact of information sharing on the R's rent, while the effect on own quantities is second-order (see Section 1.7.4). Specifically, R would like to disclose m_1 because letting E know the quantity that R produces in the second period reduces the variability of E's production — i.e., the difference between $q^C(\overline{\theta})$ and E's production when demand is low. In fact, by Proposition 1.3, with information sharing E expands production when demand is low, compared to the case without information sharing, whereby reducing its (equilibrium) output variability — i.e.,

$$q_E^S(\underline{\theta},\underline{\theta}) > q_E^N(\underline{\theta}) \quad \Rightarrow \quad \Delta q^S(\underline{\theta}) < \Delta q^N.$$

[27]Hence, expected rents and profits are approximated by a first-order Taylor expansion for $\Delta\theta$ close to 0. We relax Assumption 2 in Section 1.7.4.

This weakens the competition effect (relative to the case without information sharing) and increases rents in the second period. Hence, R prefers to face an informed rather than an uninformed competitor in the second period.

For the manufacturer, by contrast, disclosing m_1 to E has two negative effects. First, since $q_E^N(\underline{\theta}) < \mathbb{E}_{\theta_1}\left[q_E^S(\underline{\theta}, \theta_1)\right]$, the entrant is (on average) more aggressive with information sharing. Hence, M can extract a lower surplus from R when E is informed about first-period demand: a *business stealing* effect. Second, holding revenues constant, sharing information is detrimental to M because it increases R's expected rent.

Therefore, in our environment transparency arises if the decision of whether to share information is taken by the downstream incumbent, but not if it is taken by the upstream incumbent that prefers to face an uninformed entrant. This highlights a conflict of interest between upstream and downstream firms in vertical relations facing entry: whether information is shared with entrants depends on which player owns privacy rights, and is accordingly entitled to disclose information within a vertical hierarchy.[28]

We now consider the effect of information sharing on the (expected) joint profit of M and R, to analyze whether their conflict of interest can be solved by ex ante contracting. Can M and R jointly agree to an information sharing decision with compensation for the damaged party, before demand realizes?[29]

Proposition 1.5. *The ex ante joint profit of M and R is lower with information sharing than without.*

Hence, with small uncertainty, the manufacturer and the retailer jointly gain by not sharing information with the entrant: M obtains privacy rights by offering an ex ante payment to R which compensates its loss for facing an uninformed entrant. When ex ante contracting is not possible, however, either because M is capital constrained or because privacy rights cannot be easily transferred, it is unlikely that M can prevent R from disclosing information to E.

Since the entrant can always (commit to) disregard the information received by the incumbent and implement the same outcome as without information sharing, we have the following result.

Proposition 1.6. *E obtains higher profit with information sharing than without.*

1.5.1 Market for Information

Since information about the first-period report by the retailer is valuable, the entrant is willing to pay for it. Do incumbent players have any incentive to sell information to the entrant, rather

[28]The fact that the retailer prefers to disclose information is in line with the literature finding that with Cournot competition firms want to exchange information about their stochastic costs (see, e.g., Shapiro, 1986).

[29]This is equivalent to analyzing whether M and R can agree, *behind the veil of ignorance*, to a system of ex ante transfers that harmonizes their interests, with R paying M to disclose m_1 to E, or vice versa. Of course, in order for this agreement to be feasible, players must not be capital constrained.

than simply share it at no cost? Coherently with our full commitment assumption, in order to address this question we assume that the incumbent can commit at the outset of the game to a price that the entrant has to pay in order to acquire information.

Of course, by Propositions 1.4 and 1.6, E and R have a joint interest to trade information, since they are both better off with information sharing. By contrast, since M's profits are lower with information sharing by Proposition 1.4, M has an incentive to sell information to E only if the highest price that E is willing to pay for information is higher than M's loss for facing an informed competitor — i.e.,

$$\underbrace{\Pi_E^S - \Pi_E^N}_{E\text{'s willingness to pay}} \geqslant \underbrace{\Pi_2^N - \Pi_2^S}_{M\text{'s reservation price}},$$

where, given a disclosure policy $d \in \{S; N\}$, Π_E^d denotes E's expected profit and Π_2^d denotes M's second-period expected profit.

Under Assumption 2, we have the following result.

Proposition 1.7. *It is never profitable for M to sell information to E.*

As discussed above, M's loss for information sharing is shaped by two effects: by sharing information M induces the entrant to be more aggressive and increases R's rent. Since the highest price that E is willing to pay for information can internalize the first effect, but not the second one, M has no incentive to sell information.

Finally, if M and R maximize joint profits, since R's rent is just a transfer between M and R, then the entrant can pay a price for information that fully compensates M's loss.

Proposition 1.8. *M and R have an incentive to jointly sell information to E.*

Trading information is jointly profitable for the incumbents and the entrant because it maximizes total profit in the industry: it allows firms to extract more surplus from consumers (as we are going to show), and it rebalances production from a less efficient firm (the incumbent who faces agency costs) to a more efficient one (the entrant who faces no agency costs and, on average, produces more when it is informed).

1.6 Welfare

In order to study the welfare effects of information sharing, since first-period production is the same with and without information sharing, we analyze how the incumbent's decision to disclose information impacts aggregate production in the second period.

Proposition 1.9. *Expected aggregate production is lower with information sharing than without.*

Information sharing reduces the incumbent's production and allows E to increase production when the incumbent distorts it. On balance, however, aggregate production is lower than without information sharing because, holding constant the incumbent's production, E's decision is always efficient regardless of its information (since it equalizes marginal revenue to marginal costs). In other words, information reduces market efficiency because it increases R's information rent via the competition effect, whereby reducing the incumbent's overall efficiency.

In the limit of small uncertainty, information sharing has an analogous effect on consumer surplus and total welfare.

Proposition 1.10. *Consumer surplus and total welfare are lower with information sharing than without.*

This suggests that information sharing between incumbents and new entrants should not necessarily be allowed, and that incumbents should not be forbidden to sell information about their past production or communication to future competitors. Hence, our analysis highlights a potential drawback of imposing transparency about past performance to incumbents.[30]

Notice that the result in Proposition 1.10 hinges on the absence of a fixed cost of entry. With a sufficiently high entry cost, not sharing information may foreclose entry, which always harms consumers. With a stochastic entry cost, however, the net effect of information sharing on consumer surplus depends on the relative likelihood of entry being blocked without information sharing.

1.7 Extensions

1.7.1 Competing Hierarchies

Suppose that the entrant is a vertical hierarchy rather than an integrated firm (see, e.g., Caillaud *et al.*, 1995; and Martimort, 1996): in period 2 a new manufacturer M_E enters the market and sells through its exclusive retailer R_E, who is privately informed about θ_2 but does not know θ_1. For example, this may happen when the entrant is a foreign firm that needs a local retailer in order to enter the market and distribute its product.

M_E offers to R_E a direct revelation mechanism

$$\{q_E(m_E, s), t_E(m_E, s)\},$$

which specifies a production level $q_E(\cdot)$ and a transfer $t_E(\cdot)$ contingent on R_E's report m_E about θ_2 and on the information $s \in \Theta \cup \{\varnothing\}$ revealed by the incumbent about R's first-period report

[30] Of course, transparency may be welfare beneficial in other contexts. For example, improving price and quality transparency unambiguously benefit consumers — e.g., Varian (1980), Schultz (2009) and Gu and Wenzel (2011). But while these models focus on firms' ability to inform consumers about product characteristics, in our environment communication is about past demand or performance.

(with $s = \varnothing$ denoting no information).[31] To focus on separating equilibria, we impose a condition equivalent to Assumption 1.

Assumption 3. The degree of demand persistency is such that $\overline{\nu} \leqslant \overline{\nu}^* \triangleq \min\left\{4\underline{\nu}^2 - 1, \frac{1}{2}\right\}$. Moreover, δ is not too large.

The assumption requires that the degree of demand persistency is neither too low when $\theta_1 = \underline{\theta}$ (i.e., $\underline{\nu} > \frac{1}{2}$) nor too high when $\theta_1 = \overline{\theta}$. This guarantees that, in the second period, retailers have an incentive to mis-report demand only if demand is high, and that retailers' information rent is positive (see, e.g., Kastl *et al.*, 2011). Moreover, the assumption on δ ensures that intertemporal rents are positive — i.e., that in the first period R has an incentive to mis-report only when demand is high (see Laffont and Martimort, 2002).[32]

M's maximization problem is the same as in our main model, regardless of whether information is shared or not. By contrast, the entrant solves a different maximization problem, since M_E is uninformed about θ_2 and has to induce R_E to truthfully report it. Hence, the entrant's production is also distorted for rent extraction reasons, and this distortion crucially depends on whether the incumbent shares information or not.

1.7.1.1 No Information Sharing

Without information sharing, when R truthfully reveals its private information, R_E's expected utility from truthfully reporting θ_2 is

$$U_E\left(\theta_2\right) \triangleq \sum_{\theta_1} \Pr\left[\theta_1|\theta_2\right] P\left(\theta_2, Q_2^N\left(\theta_2, \theta_1\right)\right) q_E\left(\theta_2\right) - t_E\left(\theta_2\right),$$

where, abusing notation, aggregate quantity is

$$Q_2^N\left(\theta_2, \theta_1\right) \triangleq q_E\left(\theta_2\right) + q_2^N\left(\theta_2, \theta_1\right).$$

Let

$$\Delta q_2^N \triangleq \mathbb{E}_{\theta_1}\left[q_E^N\left(\overline{\theta}, \theta_1\right)\right] - \mathbb{E}_{\theta_1}\left[q_E^N\left(\underline{\theta}, \theta_1\right)\right].$$

Conjecturing that only the incentive compatibility constraint in the high demand state binds, R_E's information rent is determined by

$$U_E(\overline{\theta}) \geqslant U_E\left(\underline{\theta}\right) + \left(\Delta\theta - \Delta q_2^N\right) q_E\left(\underline{\theta}\right), \tag{1.14}$$

which reflects a competing contracts effect (averaged over θ_1).

[31] Consistent with our main model, we assume passive beliefs off equilibrium path so that, whenever a retailer receives an unexpected offer, it believes that the other players follow equilibrium strategies.

[32] This is a sufficient condition that does not affect the main results of the analysis since only second-period outputs matter to determine the incentives to share/sell information, and the welfare effects of this choice.

By standard techniques, M_E's (relaxed) maximization problem is

$$\max_{q_E(\cdot)} \sum_{\theta_2} \Pr[\theta_2] \sum_{\theta_1} \Pr[\theta_1|\theta_2] P\left(\theta_2, Q_2^N(\theta_2,\theta_1)\right) q_E(\theta_2) - \Pr\left[\theta_2 = \overline{\theta}\right]\left(\Delta\theta - \Delta q_2^N\right) q_E(\underline{\theta}).$$

Differentiating with respect to $q_E(\overline{\theta})$ and using the first-order conditions (1.7) it follows that, in equilibrium, R_E's production is not distorted when demand is high in the second period. Since R's production is also efficient, in equilibrium $q_E^N(\overline{\theta}) = q_2^N(\overline{\theta},\theta_1) = q^C(\overline{\theta})$ for every θ_1. By contrast, differentiating with respect to $q_E(\underline{\theta})$ yields

$$\sum_{\theta_1} \Pr[\theta_1|\theta_2 = \underline{\theta}] \left[P_{q_E}\left(\underline{\theta}, Q_2^N(\underline{\theta},\theta_1)\right) q_E(\underline{\theta}) + P\left(\underline{\theta}, Q_2^N(\underline{\theta},\theta_1)\right)\right] =$$

$$= \underbrace{\Delta\theta - \left[q^C(\overline{\theta}) - \mathbb{E}_{\theta_1}\left[q_2^N(\underline{\theta},\theta_1)\right]\right]}_{\text{Distortion without information}}.$$

Hence, other things being equal, R_E's production is downward distorted when demand is low in the second period. Together with (1.8) and (1.9), this condition determines equilibrium production in the second period.

Proposition 1.11. *Without information sharing,* $q_E^N(\underline{\theta}) < q^C(\underline{\theta}) < q_2^N(\underline{\theta},\underline{\theta}) = q_2^N\left(\underline{\theta},\overline{\theta}\right)$.

Therefore, the incumbent overproduces (compared to the benchmark without incomplete information). Since M can exploit demand correlation to reduce R's rent while M_E cannot, without information sharing the entrant always produces less than the incumbent. This asymmetry provides the incumbent with a competitive edge that completely offsets the potential distortion stemming from asymmetric information and enables M to increase production when demand is low.

1.7.1.2 Information Sharing

With information sharing, when R truthfully reports its information, R_E's equilibrium utility is

$$U_E(\theta_2,\theta_1) \triangleq P\left(\theta_2, Q^S(\theta_2,\theta_1)\right) q_E(\theta_2,\theta_1) - t_E(\theta_2,\theta_1),$$

where, abusing notation, aggregate quantity is

$$Q_2^S(\theta_2,\theta_1) \triangleq q_E(\theta_2,\theta_1) + q_2^S(\theta_2,\theta_1).$$

As before, we assume (and verify ex post) that only the incentive compatibility constraint in the high demand state matters. Let

$$\Delta q_2^S(\theta_1) \triangleq q_2^S(\overline{\theta},\theta_1) - q_2^S(\underline{\theta},\theta_1), \quad \forall \theta_1 \in \Theta.$$

R_E's information rent is then determined by the following inequality

$$U_E(\overline{\theta}, \theta_1) \geqslant U_E(\underline{\theta}, \theta_1) + \left(\Delta\theta - \Delta q_2^S(\theta_1)\right) q_E(\underline{\theta}, \theta_1), \quad \forall \theta_1 \in \Theta. \tag{1.15}$$

With information sharing, the strength of the 'competing contracts effect' depends on demand in the first period. Hence, for every θ_1, M_E's (relaxed) maximization problem is

$$\max_{q_E(\cdot,\theta_1)} \sum_{\theta_2} \Pr[\theta_2|\theta_1] P\left(\theta_2, Q_2^S(\theta_2,\theta_1)\right) q_E(\theta_2,\theta_1) - \Pr\left[\theta_2 = \overline{\theta}|\theta_1\right] \left(\Delta\theta - \Delta q_2^S(\theta_1)\right) q_E(\underline{\theta}, \theta_1).$$

Differentiating with respect to $q_E\left(\overline{\theta}, \theta_1\right)$ it follows that, in equilibrium, R_E's production is not distorted when demand is high in the second period. Since R's second-period production is also efficient, in equilibrium $q_E^S(\overline{\theta}) = q_2^S(\overline{\theta}, \theta_1) = q^C(\overline{\theta})$ for every θ_1. By contrast, in the low demand state R_E's production is distorted for rent extraction reasons. Differentiating with respect to $q_E(\underline{\theta}, \theta_1)$ yields

$$P_{q_E}\left(\underline{\theta}, Q_2^S(\underline{\theta},\theta_1)\right) q_E(\underline{\theta},\theta_1) + P\left(\underline{\theta}, Q_2^S(\underline{\theta},\theta_1)\right) = \underbrace{\frac{\Pr\left[\theta_2 = \overline{\theta}|\theta_1\right]}{\Pr\left[\theta_2 = \underline{\theta}|\theta_1\right]}\left(\Delta\theta - \Delta q_2^S(\theta_1)\right)}_{\text{Distortion with information}}, \quad \forall \theta_1 \in \Theta.$$

Together with (1.12) and (1.13) this condition determines the equilibrium production in states $(\underline{\theta}, \underline{\theta})$ and $(\underline{\theta}, \overline{\theta})$. While M only distorts R's production when first-period demand is low, R_E's production is always distorted downward because M_E has an incentive to minimize R_E's static rent by reducing $q_E(\underline{\theta})$. The magnitude of this distortion depends on the likelihood ratio

$$\mathcal{L}(\theta_1) \triangleq \frac{\Pr\left[\theta_2 = \overline{\theta}|\theta_1\right]}{\Pr\left[\theta_2 = \underline{\theta}|\theta_1\right]} = \begin{cases} \frac{\overline{\nu}}{1-\overline{\nu}} \text{ if } \theta_1 = \overline{\theta} \\ \frac{1-\underline{\nu}}{\underline{\nu}} \text{ if } \theta_1 = \underline{\theta}. \end{cases} \tag{1.16}$$

Indeed, M_E distorts production more when the information about $m_1 = \theta_1$ received by the incumbent indicates that demand in the second period is relatively more likely to be high than low.

Proposition 1.12. *With information sharing, $q_2^S(\underline{\theta}, \overline{\theta}) > q_E^S(\underline{\theta}, \overline{\theta})$ and $q_2^S(\underline{\theta}, \underline{\theta}) < q_E^S(\underline{\theta}, \underline{\theta})$. Moreover, $q_2^S(\underline{\theta}, \overline{\theta}) > q^C(\underline{\theta})$, while $q_2^S(\underline{\theta}, \underline{\theta}) < q^C(\underline{\theta})$ if and only if $\overline{\nu} \geqslant \frac{(1-\underline{\nu})(1-2\underline{\nu})}{3\underline{\nu}-1}$.*

Hence, the incumbent produces more than the entrant in the second period if and only if demand in the first period is high. Moreover, when demand is relatively persistent — i.e., when $\overline{\nu}$ is sufficiently large — M_E obtains more precise on $\theta_2$from observing m_1, so that it has to pay a lower rent to R_E and can expand production. In this case, M under-produces if it faces an informed rival. By contrast, when the information conveyed by m_1 is less precise, M_E is more uncertain about demand, pays higher rents, and distorts production more in order to trade off efficiency and rent extraction. This, in turn, induces M to expand production (relatively to the complete information benchmark).

1.7.1.3 Value of Information

Consider the effects of information sharing on players' profits.

Proposition 1.13. *With information sharing, the entrant always produces more than without information sharing — i.e.,* $q_E^N(\underline{\theta}) < \min\left\{q_E^S\left(\underline{\theta},\overline{\theta}\right), q_E^S\left(\underline{\theta},\underline{\theta}\right)\right\}$.

Since knowing θ_1 allows M_E to elicit R_E's private information at a lower cost, information sharing induces the entrant to increase production, so that $\mathbb{E}_{\theta_1}\left[q_E^S\left(\underline{\theta},\theta_1\right)\right] > q_E^N(\underline{\theta})$. Hence, we have the following result (as in our baseline model).

Proposition 1.14. *R wants to share information, while M does not. Moreover, M_E obtains higher profit when it is informed.*

1.7.1.4 Welfare

To analyze the effects of information sharing on consumer surplus and total welfare, consider aggregate (expected) production.

Proposition 1.15. *There exist two thresholds $\underline{\nu}_0$ and $\overline{\nu}_0(\underline{\nu}) \leqslant \overline{\nu}^*$ such that aggregate production, consumer surplus and welfare are higher with information sharing if: (i) $\underline{\nu} \geqslant \underline{\nu}_0$ or (ii) $\underline{\nu} < \underline{\nu}_0$ and $\overline{\nu} \leqslant \overline{\nu}_0(\underline{\nu})$. Otherwise, aggregate production, consumer surplus and welfare are higher without information sharing.*

Figure 1.1 illustrates the region of parameters, identified in Proposition 1.15, where consumers prefer information sharing (for values of $\overline{\nu}$ and $\underline{\nu}$ consistent with Assumption 3). In contrast to the case in which the entrant is an integrated firm, with competing hierarchies information sharing may improve market efficiency. In fact, information sharing reduces R_E's information rent, which allows M_E to distort production less and tends to increase aggregate production. Since reaction functions are downward sloping, however, the increase of the entrant's production triggers a reduction of the incumbent's production. The net effect on aggregate production depends on the degree of demand persistency, which measures the precision of the information that the entrant obtains on θ_2 when it learns θ_1.

When demand is sufficiently persistent in state $\theta_1 = \underline{\theta}$ (i.e., $\underline{\nu}$ is large) and it is not too persistent in state $\theta_1 = \overline{\theta}$ (i.e., $\overline{\nu}$ is small), the likelihood ratio $\mathcal{L}(\cdot)$ is small. In this case, information sharing has a stronger impact on the entrant's production than on the incumbent's production, thus increasing aggregate production. By contrast, when the information obtained by M_E on θ_1 does not result in a sufficiently large increase in R_E's production, information sharing reduces aggregate production.

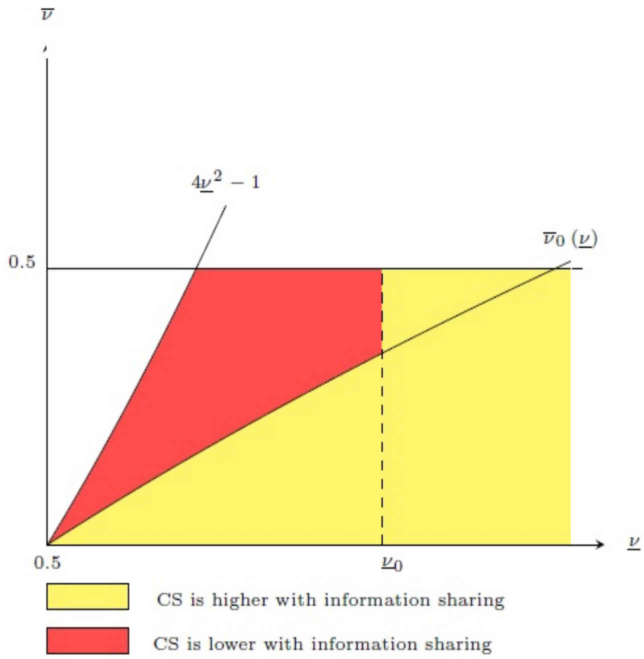

Figure 1.1: Effect of information sharing on consumer surplus.

The effect of information sharing on aggregate production (in the limit of small uncertainty) also determines its impact on consumer surplus and total welfare. Hence, in contrast to the case of an integrated entrant, with competing hierarchies transparency standards improve efficiency only if demand is sufficiently persistent in bad times and/or not too persistent in good times.

Finally, allowing M and R to contract ex ante reduces welfare since it induces them not to share information.

Proposition 1.16. *The ex ante joint profit of M and R is lower with information sharing than without.*

1.7.1.5 Market for Information

Assume now that incumbents can commit to a price at which to sell information to the entrant. Of course, R and M_E have a joint incentive to trade information, since they obtain a higher profit with information sharing. Moreover, as in our baseline model, the highest price that M_E is willing to pay for information cannot internalize M's loss due to the higher information rent for R.

Proposition 1.17. *It is never profitable for M to sell information to M_E.*

Suppose now that M and R maximize joint profits when selling information to M_E.

Proposition 1.18. *There exist two thresholds $\underline{\nu}_1$ and $\overline{\nu}_1(\underline{\nu}) \leqslant \overline{\nu}^*$ such that M and R have a joint incentive to sell information to M_E if and only if $\underline{\nu} \leqslant \underline{\nu}_1$ and $\overline{\nu} \geqslant \overline{\nu}_1(\underline{\nu})$.*

Figure 1.2 illustrates the region of parameters where M and R have a joint incentive to sell information (for values of $\overline{\nu}$ and $\underline{\nu}$ consistent with Assumption 3). When $\overline{\nu}$ is sufficiently large, firms trade information in order to gain market power vis-à-vis consumers, who are harmed by information sharing. By contrast, M and R do not sell information when (ceteris paribus) $\underline{\nu} \leqslant \underline{\nu}_1$ because an informed M_E distorts production more when $\underline{\nu}$ is small (see the expression of $\mathcal{L}(\underline{\theta})$ in (1.16)). In this case, it is less costly for the incumbents to face an informed competitor, which lowers the reservation price at which they are willing to sell information.[33]

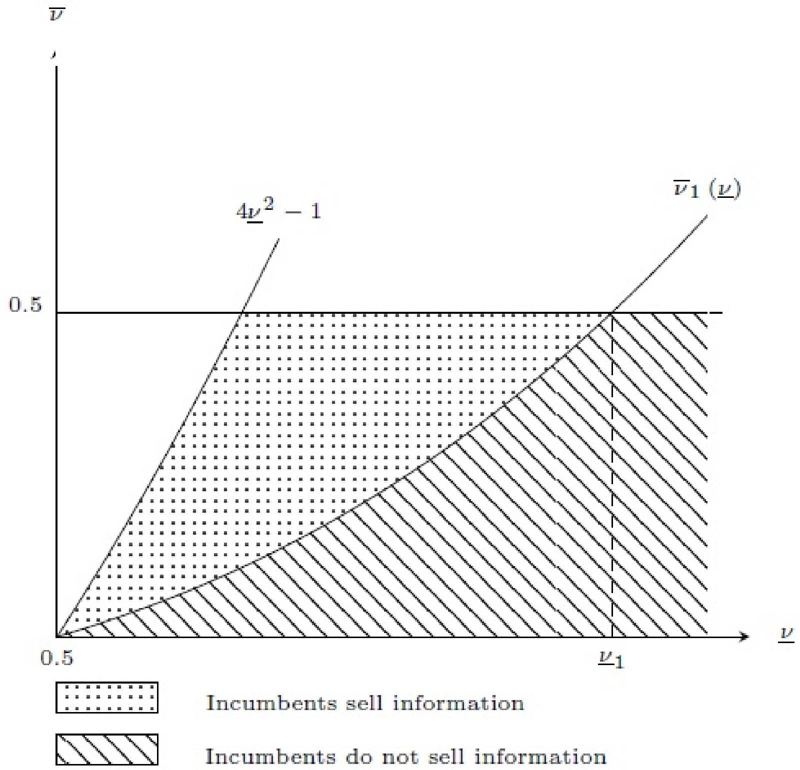

Figure 1.2: Incumbents' incentive to sell information.

Finally, the next result shows when the presence of a market where firms can trade information harms consumers.

Proposition 1.19. *When M and R have a joint incentive to sell information to M_E, information sharing harms consumers if $\underline{\nu} < \underline{\nu}_0$ and $\overline{\nu} > \overline{\nu}_0(\underline{\nu})$. When, M and R have no incentive to sell information to M_E, information sharing always benefits consumers.*

Figure 1.3 graphically summarizes the result of Proposition 1.19. When incumbents sell information to the entrant, the welfare effect depends on the degree of demand persistency. By

[33]Of course, M_E's willingness to pay for information also depends on the fact that information sharing reduces R_E's rent.

contrast, consumers are always worse off when incumbents do not sell information. Hence, a social planner should force incumbents to sell information when they are not willing to do so, despite the presence of a market for information.

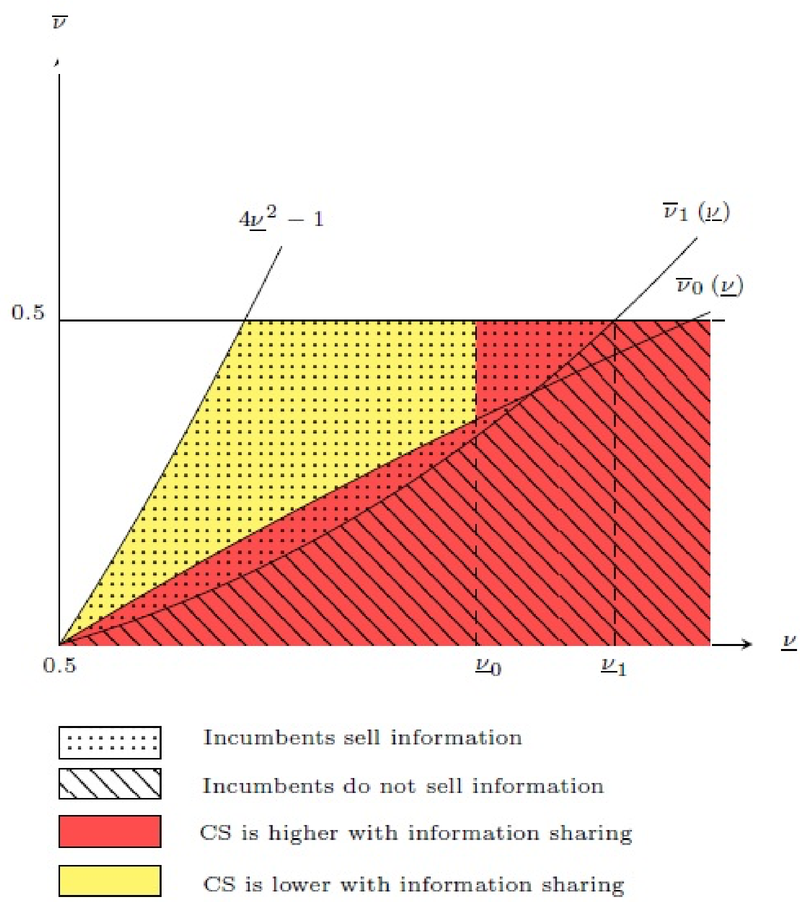

Figure 1.3: Does trading information harm consumers?

1.7.2 Stochastic Disclosure Rule

In our main model we have assumed an all-or-nothing disclosure rule. In this section, we consider a more sophisticated communication protocol that relies on a stochastic structure.[34] To simplify the analysis, we assume that $\overline{\nu} = \underline{\nu} = \nu$.

If it shares information, the incumbent commits to a disclosure rule that consists of a binary experiment with two signals, $\sigma \in \{\overline{\sigma}, \underline{\sigma}\}$, such that $\Pr\left[\sigma = \overline{\sigma} | m_1 = \overline{\theta}\right] \triangleq \alpha$ and $\Pr\left[\sigma = \underline{\sigma} | m_1 = \underline{\theta}\right] \triangleq \beta$ (see, e.g., Bergemann *et al.,* 2018, and Kastl *et al.*, 2018). The parameters α and β measure the informativeness, or accuracy, of the experiment. As a convention (and without loss of generality),

[34] As in our main model, we assume the entrant is an integrated firm.

we assume that $\alpha + \beta \geq 1$.[35] Consistent with the assumption of verifiable information in our main model, the outcome of the experiment is public — i.e., there are no further information frictions between the incumbents and the entrant. An experiment with $\alpha = \beta = 1$ is fully informative, which is equivalent to $d = S$ in our main model; while an experiment with $\alpha + \beta = 1$ is uninformative, which is equivalent to $d = N$ in our main model.

The long term contract is a menu

$$\{q_1(m_1), t_1(m_1), q_2(m_2, m_1, \sigma), t_2(m_2, m_1, \sigma)\},$$

that specifies a quantity and a transfer in the second period that are also contingent on the realized signal σ.

Since σ and θ_2 are independent conditionally on m_1, with information sharing sharing E observes two independent signals on m_1, that it uses to infer the quantity produced by the incumbent in the second period. Hence, the entrant's posterior is

$$\begin{aligned}
\Pr\left[m_1 = \overline{\theta}|\sigma, \theta_2\right] &= \frac{\Pr\left[\sigma, \theta_2|m_1 = \overline{\theta}\right]\Pr\left[m_1 = \overline{\theta}\right]}{\Pr\left[\sigma, \theta_2|m_1 = \overline{\theta}\right]\Pr\left[m_1 = \overline{\theta}\right] + \Pr\left[\sigma, \theta_2|m_1 = \underline{\theta}\right]\Pr\left[m_1 = \underline{\theta}\right]} \\
&= \frac{\Pr\left[\theta_2|\theta_1 = \overline{\theta}\right]\Pr\left[\sigma|\theta_1 = \overline{\theta}\right]}{\Pr\left[\theta_2|\theta_1 = \overline{\theta}\right]\Pr\left[\sigma|\theta_1 = \overline{\theta}\right] + \Pr\left[\theta_2|\theta_1 = \underline{\theta}\right]\Pr\left[\sigma|\theta_1 = \underline{\theta}\right]},
\end{aligned}$$

where we have used the fact that, in a truthful equilibrium where $m_1 = \theta_1$, $\Pr\left[m_1 = \overline{\theta}\right] = \Pr\left[\theta_1 = \overline{\theta}\right] = \frac{1}{2}$.

For any realization (σ, θ_2), the entrant's problem is

$$\max_{q_E \geqslant 0} \sum_{\theta_1} \Pr\left[\theta_1|\sigma, \theta_2\right] P(\theta_2, q_E + q_2(\theta_2, \theta_1, \sigma)) q_E,$$

whose first order condition yields

$$q_E(\theta_2, \sigma) \triangleq \frac{\theta_2 - \sum_{\theta_1} \Pr\left[\theta_1|\theta_2, \sigma\right] q_2(\theta_2, \theta_1, \sigma)}{2}, \quad \forall (\sigma, \theta_2).$$

Consider now the incumbent's problem. For any σ, let

$$\Delta q^S(\sigma) \triangleq q_E^S\left(\overline{\theta}, \sigma\right) - q_E^S\left(\underline{\theta}, \sigma\right)$$

be the difference between E's output with high and low demand in the second period. R's binding

[35] This is just a labelling of signals that ensures that upon observing signal $\overline{\sigma}$ (resp. $\underline{\sigma}$), the entrant assigns higher probability to $m_1 = \overline{\theta}$ (resp. $\underline{\theta}$).

incentive compatibility constraint in the second period (see the Appendix) is

$$U_2\left(\overline{\theta}, m_1, \sigma\right) = \Delta\theta q_2(\overline{\theta}, m_1, \sigma) - \underbrace{\left[q_E\left(\overline{\theta}, \sigma\right) - q_E\left(\underline{\theta}, \sigma\right)\right]}_{\Delta q(m_1, \sigma)} q_2\left(\underline{\theta}, m_1, \sigma\right),$$

where the competition effect now depends on the signal σ. Hence, the relevant first-period incentive compatibility constraint is

$$\begin{aligned} U_1\left(\overline{\theta}\right) = {} & \underbrace{\Delta\theta q_1\left(\underline{\theta}\right)}_{\text{Static rent}} + \\ & + \underbrace{\delta\nu\left[\textstyle\sum_\sigma \Pr\left[\sigma|\underline{\theta}\right]\left(\Delta\theta - \Delta q(\underline{\theta}, \sigma)\right) q_2\left(\underline{\theta}, \underline{\theta}, \sigma\right) - \sum_\sigma \Pr\left[\sigma|\overline{\theta}\right]\left(\Delta\theta - \Delta q(\overline{\theta}, \sigma)\right) q_2\left(\underline{\theta}, \overline{\theta}, \sigma\right)\right]}_{\text{Intertemporal rent}}, \end{aligned}$$

which is equivalent to (1.5) when $\alpha + \beta = 1$ and to (1.10) when $\alpha = \beta = 1$.

The incumbent's maximization problem is

$$\begin{aligned} \max_{q_1(\cdot), q_2(\cdot), U_1(\cdot)} \mathbb{E}\left[\sum_{\tau=1,2} \delta^{\tau-1} P\left(\theta_\tau, Q_\tau\left(\cdot\right)\right) q_\tau\left(\cdot\right)\right] + \\ - \sum_{\theta_1} \Pr\left[\theta_1\right]\left[U_1\left(\theta_1\right) + \delta \Pr\left[\theta_2 = \overline{\theta}|\theta_1\right] \sum_\sigma \Pr\left[\sigma|\theta_1\right]\left[\Delta\theta - \Delta q\left(\theta_1, \sigma\right)\right] q_2\left(\underline{\theta}, \theta_1, \sigma\right)\right], \end{aligned}$$

where

$$Q_2\left(\theta_2, \theta_1, \sigma\right) \triangleq q_2\left(\theta_2, \theta_1, \sigma\right) + q_E\left(\theta_2, \sigma\right).$$

In the Appendix, we show that the first-order conditions of this problem are analogous to those in our main model. The difference is that the entrant's production now depends on the signal produced by the experiment.

Proposition 1.20. *The optimal experiment offered by M is uninformative — i.e., it features $\alpha + \beta = 1$. The optimal experiment offered by R is fully informative — i.e., it features $\alpha = \beta = 1$. The uninformative experiment maximizes consumer surplus and welfare.*

Hence, in the limit of small uncertainty, even with a more complex information structure, our main qualitative results obtained with the all-or-nothing disclosure rule, and their policy implications, hold.

1.7.3 Secret Renegotiation and Ex-post Disclosure

In our main model, we assumed that the incumbent players can commit ex-ante to an information disclosure rule. Even if commitment is a standard hypothesis in the existing literature on information sharing (see, e.g., Vives, 2006), one may wonder whether our results are robust to the possibility that the incumbent players (secretly) renege on their ex ante commitment to share or not information. In this section we show that, when at the beginning of period 2 — i.e., before learning θ_2 — the incumbent players can renege on the information sharing decision, but not on the terms of the optimal long term contract, only the equilibrium with information sharing survives.[36]

Proposition 1.21. *The equilibrium with information sharing characterized in Section 1.4.2 is robust to ex post renegotiation, while the equilibrium without information sharing is not.*

The reason why the equilibrium with information sharing is robust to ex-post renegotiation of the information sharing decision is straightforward. Consider an equilibrium in which the incumbent players commit to share information and M offers the long term contract characterized in Section 1.4.2. First, M has no incentive to renege on this commitment since players cannot modify the contractual terms and, hence, by refusing to share information M cannot increase the second-period transfer. But then R has no profitable deviation either, since the optimal long term contract is incentive compatible.

By contrast, the equilibrium without information sharing is not robust to ex-post renegotiation because R has a unilateral incentive to disclose information when demand in the first period is high. In fact, other things being equal, this reduces E's production relative to the no information sharing outcome characterized in Section 1.4.1, whereby increasing R's revenue.

This result suggests that when incumbent players can secretly renege on their ex-ante commitment not to share information, there is an even stronger incentive from a welfare point of view to ban communication with entrants.

1.7.4 Large Uncertainty

Our results hinge of the assumption of small uncertainty — i.e., $\Delta\theta$ small (Assumption 2) — that allowed us to analytically solve for players' expected profit and rents in Section 1.5. In this section we use numerical simulations to analyze the effects of large uncertainty.

To simplify the analysis, we assume that $\underline{\theta} = 1$ and $\overline{\nu} = \underline{\nu} = \nu$, so that Assumption 1 implies $\nu \geqslant \frac{1}{3}$. Moreover, we impose a no-shut down condition ensuring that $\Delta\theta$ is not so large that the incumbent shuts down production when demand is repeatedly low — i.e., $\Delta\theta \leqslant \Delta\theta_0(\nu) \triangleq \frac{3\nu-1}{4}$.[37] Hence, the two parameters of interest are ν and $\Delta\theta$ and we compare profits and rents with and without informations sharing when $\nu \geqslant \frac{1}{3}$ and $\Delta\theta \leqslant \Delta\theta_0(\nu)$ (see the Appendix for details).

[36] A similar result obtains when renegotiation occurs before θ_2 has realized.

[37] In the Appendix, we show that the incumbent never shuts down production in this case. In fact, for $\nu \geqslant \frac{1}{3}$,

$$q_2^S(\underline{\theta}, \underline{\theta}) = \frac{1}{3} - \frac{4}{3(3\nu - 1)}\Delta\theta < q_2^N(\underline{\theta}, \underline{\theta}) = \frac{1}{3} - \frac{3+\nu}{6\nu}\Delta\theta,$$

and $q_2^S(\underline{\theta}, \underline{\theta}) \geqslant 0$ if and only if $\Delta\theta \leqslant \Delta\theta_0(\nu)$.

Figure 1.4 shows that the retailer wants to share information if and only if uncertainty is sufficiently small — i.e.,

$$\Delta\theta \leqslant \Delta\theta_u(\nu) \triangleq \frac{2\nu(3\nu-1)}{14\nu+9\nu^2-3}.$$

Since $q_2^S(\underline{\theta},\underline{\theta}) < q_2^N(\underline{\theta},\underline{\theta})$, holding E's production constant R prefers not to share information because rents are increasing with quantity. However, E's production and, hence, the competition effect also depend on the entrant's information. When $\Delta\theta$ is small, R's rent is mainly shaped by the competition effect because the difference in the incumbent's quantities with an without information sharing only has a second order effect — i.e., $\left[q_2^S(\underline{\theta},\underline{\theta}) - q_2^N(\underline{\theta},\underline{\theta})\right] \to 0$ as $\Delta\theta \to 0$. By contrast, when $\Delta\theta$ grows large, the difference in the incumbent's quantities have a larger effect on R's rent and overcome the competition effect, so that R prefers not to share information. And the effect of the incumbent's quantities magnifies when demand is more persistent.

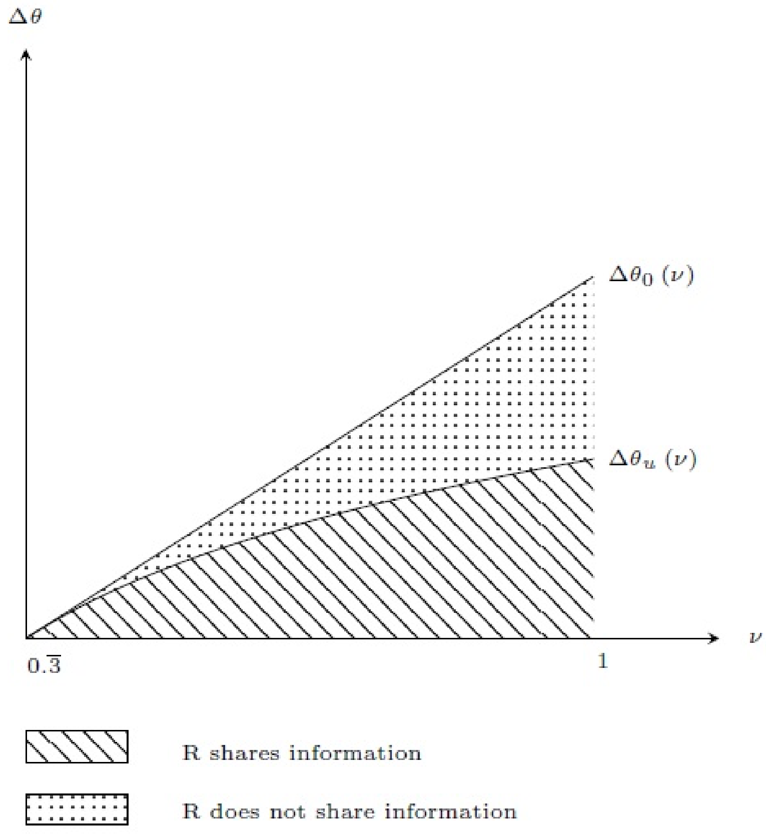

Figure 1.4: R's incentive to share information with large uncertainty

The manufacturer's incentive to share information is illustrated in Figure 1.5. M prefers not to share information if and only if $\Delta\theta$ is sufficiently small — i.e.,

$$\Delta\theta \leqslant \Delta\theta_\pi(\nu) \triangleq \frac{16\nu(3\nu-1)}{38\nu+63\nu^2-9}.$$

Of course, even with large uncertainty, sharing information allows E to be more aggressive, which harms M. As $\Delta\theta$ grows large, however, R's rent is higher without information sharing, as discussed above. Hence, for $\Delta\theta$ sufficiently large M has an incentive to share information.

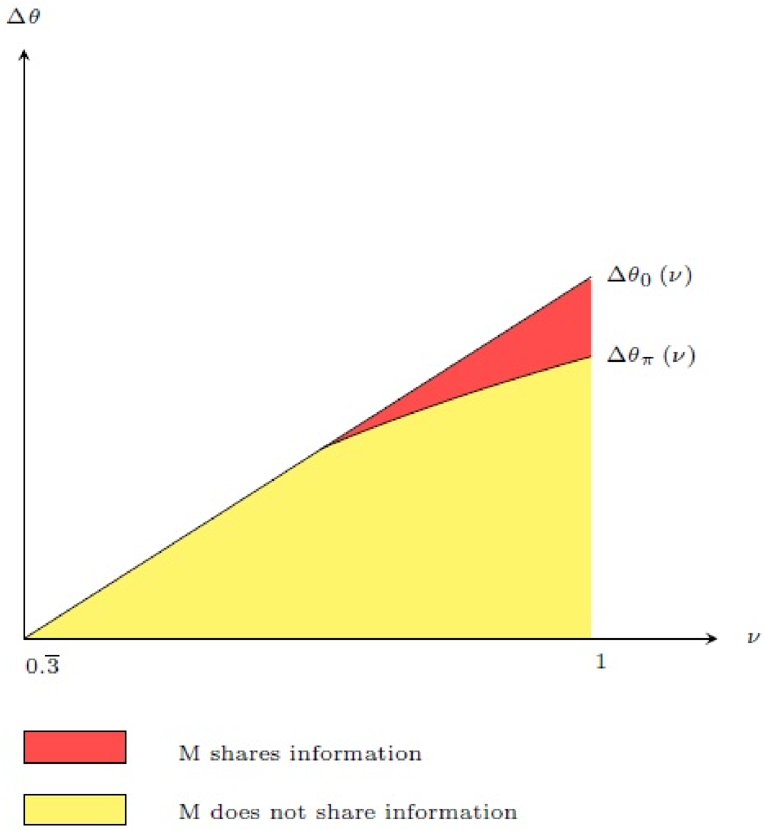

Figure 1.5: M's incentive to share information with large uncertainty

Finally, since $\Delta\theta_\pi(\nu) \geqslant \Delta\theta_u(\nu)$, M and R have a joint incentive not to share information for $\Delta\theta \in [\Delta\theta_u(\nu), \Delta\theta_\pi(\nu)]$ even if they do not contract ex ante, as shown in Figure 1.6.

1.8 Conclusions

It is commonly believed that forcing incumbents to be more transparent with entrants intensifies competition and increases consumer surplus, efficiency and total welfare. This presumption may be incorrect, however, when competition takes place between vertical hierarchies. Specifically, when incumbents contract over time with privately informed retailers or downstream units, forcing them to share information about past performances with an entrant may actually lower consumer surplus and total welfare. Interestingly, while downstream firms are willing to disclose their private information to entrants, upstream firms do not want to do so.

Although we developed our arguments in a manufacturer-retailer framework, the scope of our analysis is broader. Our insights apply to any environment involving entry by a competing orga-

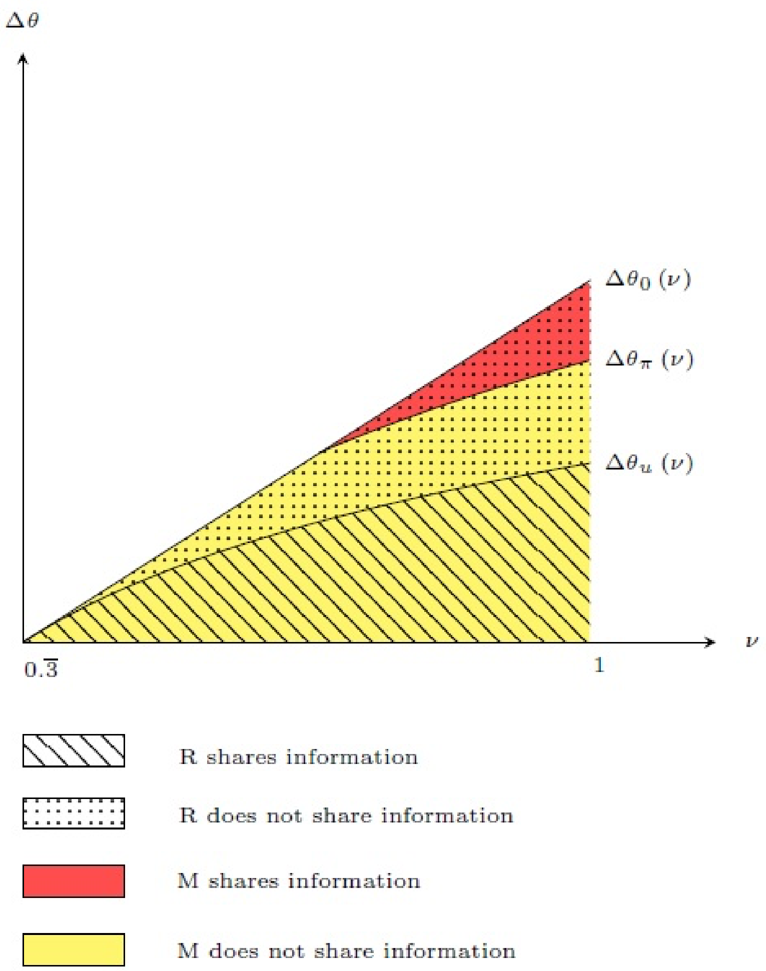

Figure 1.6: Joint incentives to share information with large uncertainty

nization with horizontal externalities, where principals deal with exclusive and privately informed agents, like procurement contracting, executive compensations, patent licensing, and insurance or credit relationships.

1.9 Appendix

Posterior Probabilities. Since in a separating equilibrium M's report is truthful, — i.e., $m_1 = \theta_1$ — E's beliefs are computed through the Bayes rule:

$$\Pr\left[\theta_1 = \overline{\theta} | \theta_2 = \overline{\theta}\right] \triangleq \frac{\Pr\left[\theta_2 = \overline{\theta} | \theta_1 = \overline{\theta}\right] \Pr\left[\theta_1 = \overline{\theta}\right]}{\sum_{\theta_1} \Pr\left[\theta_2 = \overline{\theta} | \theta_1\right] \Pr\left[\theta_1\right]} = \frac{\overline{\nu}}{1 + \Delta\nu},$$

$$\Pr\left[\theta_1 = \overline{\theta} | \theta_2 = \underline{\theta}\right] \triangleq \frac{\Pr\left[\theta_2 = \underline{\theta} | \theta_1 = \overline{\theta}\right] \Pr\left[\theta_1 = \overline{\theta}\right]}{\sum_{\theta_1} \Pr\left[\theta_2 = \underline{\theta} | \theta_1\right] \Pr\left[\theta_1\right]} = \frac{1 - \overline{\nu}}{1 - \Delta\nu},$$

with $\Pr\left[\theta_1 = \underline{\theta} | \theta_2\right] = 1 - \Pr\left[\theta_1 = \overline{\theta} | \theta_2\right]$ for every θ_2.

Proof of Proposition 1.1. Both constraints (1.5) and $U_1(\underline{\theta}) \geqslant 0$ bind at the optimum. Maximizing (1.6) with respect to $q_1\left(\overline{\theta}\right)$ and $q_1(\underline{\theta})$ yields

$$q_1^N\left(\overline{\theta}\right) = \frac{\overline{\theta}}{2} > q_1^N(\underline{\theta}) = \frac{\underline{\theta} - \Delta\theta}{2}.$$

Maximizing (1.6) with respect to $q_2\left(\overline{\theta}, \theta_1\right)$, $q_2\left(\underline{\theta}, \overline{\theta}\right)$ and $q_2(\underline{\theta}, \underline{\theta})$ yields

$$\overline{\theta} - 2q_2\left(\overline{\theta}, \theta_1\right) - q_E^N\left(\overline{\theta}\right) = 0, \quad \forall \theta_1 \in \Theta, \tag{A1}$$

$$\underline{\theta} - 2q_2\left(\underline{\theta}, \overline{\theta}\right) - q_E^N(\underline{\theta}) = 0, \tag{A2}$$

$$\underline{\theta} - 2q_2(\underline{\theta}, \underline{\theta}) - q_E^N(\underline{\theta}) - \frac{1 + \Delta\nu}{\underline{\nu}}\left(\Delta\theta - \Delta q^N\right) = 0. \tag{A3}$$

Using E's reaction function (1.2), we obtain $q_2^N\left(\overline{\theta}, \theta_1\right) = q_E^N\left(\overline{\theta}\right)$ and

$$q_E^N(\underline{\theta}) = q^C(\underline{\theta}) + \frac{1 + \Delta\nu}{3(1 - 2\Delta\nu)}\Delta\theta,$$

$$q_2^N(\underline{\theta}, \overline{\theta}) = q^C(\underline{\theta}) - \frac{1 + \Delta\nu}{6(1 - 2\Delta\nu)}\Delta\theta,$$

$$q_2^N(\underline{\theta}, \underline{\theta}) = q^C(\underline{\theta}) - \frac{(1 + \Delta\nu)(3(1 - \Delta\nu) + \underline{\nu})}{6\underline{\nu}(1 - 2\Delta\nu)}\Delta\theta.$$

It is then immediate to verify that

$$q_2^N\left(\underline{\theta}, \overline{\theta}\right) < q^C(\underline{\theta}) < q_E^N(\underline{\theta}),$$

and, by Assumption 1,

$$q_2^N(\underline{\theta},\underline{\theta}) - q_2^N\left(\underline{\theta},\overline{\theta}\right) = -\frac{(1+\Delta\nu)(1-\Delta\nu)}{2\underline{\nu}(1-2\Delta\nu)}\Delta\theta < 0.$$

R's second-period rent is strictly positive since, by Assumption 1,

$$\Delta\theta - \Delta q^N = \frac{1-\Delta\nu}{1-2\Delta\nu}\Delta\theta > 0.$$

R's first-period rent is

$$U_1^N\left(\overline{\theta}\right) = \Delta\theta q_1^N(\underline{\theta}) + \delta\overline{\nu}\left(\Delta\theta - \Delta q^N\right)\left[q_2^N(\underline{\theta},\underline{\theta}) - q_2^N\left(\underline{\theta},\overline{\theta}\right)\right].$$

By first-order Taylor approximation around $\Delta\theta = 0$,

$$U_1^N\left(\overline{\theta}\right) \approx \lim_{\Delta\theta\to 0} U_1^N\left(\overline{\theta}\right) + \Delta\theta \lim_{\Delta\theta\to 0}\frac{\partial U_1^N\left(\overline{\theta}\right)}{\partial\Delta\theta},$$

where $\lim_{\Delta\theta\to 0} U_1^N\left(\overline{\theta}\right) = 0$. Letting $\underline{\theta} = \overline{\theta} - \Delta\theta$,

$$\lim_{\Delta\theta\to 0}\frac{\partial U_1^N\left(\overline{\theta}\right)}{\partial\Delta\theta} = q^M\left(\overline{\theta}\right).$$

Therefore, for $\Delta\theta$ small (Assumption 2) R's first-period rent is also strictly positive. ■

Proof of Proposition 1.2. Both constraints (1.10) and $U_1(\underline{\theta}) \geqslant 0$ are binding at the optimum. Maximizing (1.11) with respect to $q_1(\cdot)$, it is straightforward to show that first-period production is the same as without information sharing. Maximizing (1.11) with respect to $q_2\left(\overline{\theta},\theta_1\right)$, $q_2\left(\underline{\theta},\overline{\theta}\right)$ and $q_2(\underline{\theta},\underline{\theta})$ yields

$$\bar{\theta} - 2q_2\left(\overline{\theta},\theta_1\right) - q_E^S\left(\overline{\theta},\theta_1\right) = 0, \quad \forall\theta_1\in\Theta, \tag{A4}$$

$$\underline{\theta} - 2q_2\left(\underline{\theta},\overline{\theta}\right) - q_E^S\left(\underline{\theta},\overline{\theta}\right) = 0, \tag{A5}$$

$$\underline{\theta} - 2q_2(\underline{\theta},\underline{\theta}) - q_E^S(\underline{\theta},\underline{\theta}) - \frac{1+\Delta\nu}{\underline{\nu}} \times \left[\Delta\theta - \Delta q^S(\underline{\theta})\right] = 0. \tag{A6}$$

Using E's reaction function (1.1), we obtain

$$q_2^S(\underline{\theta},\underline{\theta}) = q^C(\underline{\theta}) - \frac{4(1+\Delta\nu)}{3(4\underline{\nu}-\overline{\nu}-1)}\Delta\theta,$$

$$q_E^S(\underline{\theta},\underline{\theta}) = q^C(\underline{\theta}) + \frac{2(1+\Delta\nu)}{3(4\underline{\nu}-\overline{\nu}-1)}\Delta\theta.$$

Moreover,

$$q_2^S\left(\overline{\theta},\theta_1\right) = q_E^S\left(\overline{\theta},\theta_1\right) = q^C\left(\overline{\theta}\right), \quad \forall\theta_1\in\Theta,$$

and $q_2^S\left(\underline{\theta},\overline{\theta}\right)=q_E^S\left(\underline{\theta},\overline{\theta}\right)=q^C\left(\underline{\theta}\right)$. By Assumption 1,

$$\begin{aligned} q_2^S\left(\underline{\theta},\underline{\theta}\right)-q_2^S\left(\underline{\theta},\overline{\theta}\right) &= -\frac{4\left(1+\Delta\nu\right)}{3\left(4\underline{\nu}-\overline{\nu}-1\right)}\Delta\theta<0, \\ q_E^S\left(\underline{\theta},\underline{\theta}\right)-q_E^S\left(\underline{\theta},\overline{\theta}\right) &= \frac{2\left(1+\Delta\nu\right)}{3\left(4\underline{\nu}-\overline{\nu}-1\right)}\Delta\theta>0. \end{aligned}$$

Hence, E produces more than R when demand is repeatedly low.

R's second-period rent is strictly positive since, by Assumption 1,

$$\Delta\theta-\Delta q^S(\overline{\theta})=\frac{2}{3}\Delta\theta>0,\quad \Delta\theta-\Delta q^S\left(\underline{\theta}\right)=\frac{2\underline{\nu}}{4\underline{\nu}-\overline{\nu}-1}\Delta\theta>0.$$

R's first-period rent is

$$U_1^S(\overline{\theta})=\Delta\theta q_1^S\left(\underline{\theta}\right)+\delta\overline{\nu}\left\{\left[\Delta\theta-\Delta q^S(\underline{\theta})\right]q_2^S\left(\underline{\theta},\underline{\theta}\right)-\left[\Delta\theta-\Delta q^S(\overline{\theta})\right]q_2^S\left(\underline{\theta},\overline{\theta}\right)\right\}.$$

By a first-order Taylor approximation around $\Delta\theta=0$,

$$U_1^S(\overline{\theta})\approx\lim_{\Delta\theta\to0}U_1^S(\overline{\theta})+\Delta\theta\lim_{\Delta\theta\to0}\frac{\partial U_1^S(\overline{\theta})}{\partial\Delta\theta},$$

where $\lim_{\Delta\theta\to0}U_1^S(\overline{\theta})=0$. Letting $\underline{\theta}=\overline{\theta}-\Delta\theta$,

$$\lim_{\Delta\theta\to0}\frac{\partial U_1^S\left(\overline{\theta}\right)}{\partial\Delta\theta}=q^M\left(\overline{\theta}\right)+\delta\overline{\nu}\left[\frac{\partial q_E^S(\underline{\theta},\underline{\theta})}{\partial\Delta\theta}-\frac{\partial q_E^S\left(\underline{\theta},\overline{\theta}\right)}{\partial\Delta\theta}\right]q^C(\overline{\theta}),$$

Hence, for $\Delta\theta$ small,

$$U_1^S\left(\overline{\theta}\right)\approx q^M(\overline{\theta})\Delta\theta+\delta\overline{\nu}\left[\frac{\partial q_E^S(\underline{\theta},\underline{\theta})}{\partial\Delta\theta}-\frac{\partial q_E^S\left(\underline{\theta},\overline{\theta}\right)}{\partial\Delta\theta}\right]q^C(\overline{\theta})\Delta\theta.$$

Since $\frac{\partial}{\partial\Delta\theta}q_E^S\left(\underline{\theta},\underline{\theta}\right)>0$ and $\frac{\partial}{\partial\Delta\theta}q_E^S\left(\underline{\theta},\overline{\theta}\right)<0$ by Assumption 1, for $\Delta\theta$ small R's first-period rent is strictly positive. ■

Proof of Proposition 1.3. We compare E's equilibrium quantities with and without information sharing. When demand is low in both periods, by Assumption 1,

$$q_E^S\left(\underline{\theta},\underline{\theta}\right)-q_E^N\left(\underline{\theta}\right)=\frac{\left(1+\Delta\nu\right)\left(1-\overline{\nu}\right)}{\left(4\underline{\nu}-\overline{\nu}-1\right)\left(1-2\Delta\nu\right)}\Delta\theta>0.$$

When demand is low only in the second period, by Assumption 1,

$$q_E^S(\underline{\theta},\overline{\theta})-q_E^N\left(\underline{\theta}\right)=-\frac{1+\Delta\nu}{3\left(1-2\Delta\nu\right)}\Delta\theta<0.$$

For the last part of the proposition,

$$\mathbb{E}\left[q_E^S\left(\underline{\theta},\theta_1\right)|\underline{\theta}\right]-q_E^N\left(\underline{\theta}\right)=\Pr\left[\theta_1=\overline{\theta}|\theta_2=\underline{\theta}\right]q_E^S\left(\underline{\theta},\overline{\theta}\right)+\Pr\left[\theta_1=\underline{\theta}|\theta_2=\underline{\theta}\right]q_E^S\left(\underline{\theta},\underline{\theta}\right)-q_E^N\left(\underline{\theta}\right)$$
$$=\frac{\left(1-\overline{\nu}\right)\left(1+\Delta\nu\right)^2}{3\left(4\underline{\nu}-\overline{\nu}-1\right)\left(1-\Delta\nu\right)\left(1-2\Delta\nu\right)}\Delta\theta,$$

which is strictly positive by Assumption 1. ■

Proof of Proposition 1.4. First, we compare R's ex ante rent with and without information sharing. R's rent without information sharing is

$$\mathcal{V}^N\triangleq\frac{\delta}{2}\left(1-\underline{\nu}\right)\left(\Delta\theta-\Delta q^N\right)q_2^N\left(\underline{\theta},\underline{\theta}\right)$$
$$+\frac{1}{2}\{\underbrace{\Delta\theta q^M\left(\underline{\theta}\right)+\delta\overline{\nu}\left(\Delta\theta-\Delta q^N\right)\left[q_2^N\left(\underline{\theta},\underline{\theta}\right)-q_2^N\left(\underline{\theta},\overline{\theta}\right)\right]}_{\triangleq U_1^N\left(\overline{\theta}\right)}+\delta\overline{\nu}\left(\Delta\theta-\Delta q^N\right)q_2^N\left(\underline{\theta},\overline{\theta}\right)\}.$$

By a first-order Taylor approximation around $\Delta\theta=0$,

$$\mathcal{V}^N\approx\lim_{\Delta\theta\to 0}\mathcal{V}^N+\Delta\theta\lim_{\Delta\theta\to 0}\frac{\partial\mathcal{V}^N}{\partial\Delta\theta},$$

where $\lim_{\Delta\theta\to 0}\mathcal{V}^N=0$ and, letting $\underline{\theta}=\overline{\theta}-\Delta\theta$,

$$\lim_{\Delta\theta\to 0}\frac{\partial\mathcal{V}^N}{\partial\Delta\theta}=\frac{1}{2}q^M(\overline{\theta})+\frac{\delta}{2}q^C(\overline{\theta})\left[\overline{\nu}\left(1+\frac{\partial q_E^N(\underline{\theta})}{\partial\Delta\theta}\right)+\left(1-\underline{\nu}\right)\left(1+\frac{\partial q_E^N(\underline{\theta})}{\partial\Delta\theta}\right)\right].$$

Hence,

$$\mathcal{V}^N\approx\frac{1}{2}q^M(\overline{\theta})\Delta\theta+\frac{\delta}{2}q^C(\overline{\theta})\left[\overline{\nu}\left(1+\frac{\partial q_E^N(\underline{\theta})}{\partial\Delta\theta}\right)+\left(1-\underline{\nu}\right)\left(1+\frac{\partial q_E^N(\underline{\theta})}{\partial\Delta\theta}\right)\right]\Delta\theta. \tag{A7}$$

Similarly, R's rent with information sharing is

$$\mathcal{V}^S\triangleq\frac{\delta}{2}\left(1-\underline{\nu}\right)\left[\Delta\theta-\Delta q^S\left(\underline{\theta}\right)\right]q_2^S\left(\underline{\theta},\underline{\theta}\right)$$
$$+\frac{1}{2}\{\underbrace{\Delta\theta q_1^M\left(\underline{\theta}\right)+\delta\overline{\nu}\left[(\Delta\theta-\Delta q^S(\underline{\theta}))q_2^S\left(\underline{\theta},\underline{\theta}\right)-(\Delta\theta-\Delta q^S(\overline{\theta}))q_2^S\left(\underline{\theta},\overline{\theta}\right)\right]}_{\triangleq U_1^S\left(\overline{\theta}\right)}$$
$$+\delta\overline{\nu}(\Delta\theta-\Delta q^S(\overline{\theta}))q_2^S\left(\underline{\theta},\overline{\theta}\right)]\}.$$

As before, $\lim_{\Delta\theta\to 0}\mathcal{V}^S=0$ and

$$\lim_{\Delta\theta\to 0}\frac{\partial\mathcal{V}^S}{\partial\Delta\theta}=\frac{1}{2}q^M(\overline{\theta})+\frac{\delta}{2}q^C(\overline{\theta})\left[\overline{\nu}\left(1+\frac{\partial q_E^S(\underline{\theta},\overline{\theta})}{\partial\Delta\theta}\right)+\left(1-\underline{\nu}\right)\left(1+\frac{\partial q_E^S(\underline{\theta},\underline{\theta})}{\partial\Delta\theta}\right)\right].$$

Hence,

$$\mathcal{V}^S\approx\frac{1}{2}q^M(\overline{\theta})\Delta\theta+\frac{\delta}{2}q^C(\overline{\theta})\left[\overline{\nu}\left(1+\frac{\partial q_E^S(\underline{\theta},\overline{\theta})}{\partial\Delta\theta}\right)+\left(1-\underline{\nu}\right)\left(1+\frac{\partial q_E^S(\underline{\theta},\underline{\theta})}{\partial\Delta\theta}\right)\right]\Delta\theta. \tag{A8}$$

Comparing (A7) and (A8),

$$\mathcal{V}^N - \mathcal{V}^S \approx \frac{\delta}{2} q^C(\overline{\theta}) \left[\overline{\nu} \frac{\partial}{\partial \Delta\theta} \left(q_E^N(\underline{\theta}) - q_E^S(\underline{\theta}, \underline{\theta}) \right) + (1-\underline{\nu}) \frac{\partial}{\partial \Delta\theta} \left(q_E^N(\underline{\theta}) - q_E^S(\underline{\theta}, \underline{\theta}) \right) \right] \Delta\theta \tag{A9}$$
$$= -\frac{\delta q^C(\overline{\theta})(1-\overline{\nu})(1+\Delta\nu)^2}{2(1-2\Delta\nu)(4\underline{\nu}-\overline{\nu}-1)} \Delta\theta < 0,$$

where we have used equilibrium quantities from Propositions 1.1 and 1.2 and Assumption 1.

Second, we compare M's expected profit with and without information sharing. By a first-order Taylor approximation around $\Delta\theta = 0$, M's expected profit without information sharing is

$$\Pi^N \approx \lim_{\Delta\theta \to 0} \Pi^N + \Delta\theta \lim_{\Delta\theta \to 0} \frac{\partial \Pi^N}{\partial \Delta\theta},$$

where

$$\lim_{\Delta\theta \to 0} \Pi^N = q^M(\overline{\theta})^2 + \delta q^C \left(\overline{\theta}\right)^2,$$

and, using $\underline{\theta} = \overline{\theta} - \Delta\theta$ and the Envelope Theorem,

$$\lim_{\Delta\theta \to 0} \frac{\partial \Pi^N}{\partial \Delta\theta} = -q^M(\overline{\theta}) - \delta \left[1 + \tfrac{\partial q_E^N(\underline{\theta})}{\partial \Delta\theta} \right] q^C(\overline{\theta}).$$

Hence,

$$\Pi^N \approx q^M(\overline{\theta})^2 + \delta q^C(\overline{\theta})^2 - \left[q^M\left(\overline{\theta}\right) + q^C(\overline{\theta})\delta \left(1 + \tfrac{\partial q_E^N(\underline{\theta})}{\partial \Delta\theta} \right) \right] \Delta\theta. \tag{A10}$$

With information sharing, since

$$\lim_{\Delta\theta \to 0} \Pi^S = q^M(\overline{\theta})^2 + \delta q^C(\overline{\theta})^2,$$

and

$$\lim_{\Delta\theta \to 0} \frac{\partial \Pi^S}{\partial \Delta\theta} = -q^M(\overline{\theta}) - q^C(\overline{\theta})\delta \left[\overline{\nu} \left(1 + \tfrac{\partial q_E^S(\underline{\theta},\underline{\theta})}{\partial \Delta\theta} \right) + (1-\overline{\nu}) \left(1 + \tfrac{\partial q_E^S\left(\underline{\theta},\overline{\theta}\right)}{\partial \Delta\theta} \right) + \left(1 + \tfrac{\partial q_E^S(\underline{\theta},\underline{\theta})}{\partial \Delta\theta} \right) \right],$$

M's expected profit is

$$\Pi^S \approx q^M(\overline{\theta})^2 + \delta q^C(\overline{\theta})^2 - q^M(\overline{\theta})\Delta\theta +$$
$$- q^C(\overline{\theta})\delta \left[\overline{\nu} \left(1 + \tfrac{\partial q_E^S(\underline{\theta},\underline{\theta})}{\partial \Delta\theta} \right) + (1-\overline{\nu}) \left(1 + \tfrac{\partial q_E^S\left(\underline{\theta},\overline{\theta}\right)}{\partial \Delta\theta} \right) + \left(1 + \tfrac{\partial q_E^S(\underline{\theta},\underline{\theta})}{\partial \Delta\theta} \right) \right] \Delta\theta.$$

Comparing this with (A10),

$$\Pi^N - \Pi^S \approx q^C(\overline{\theta})\delta \left[\tfrac{\partial q_E^S(\underline{\theta},\underline{\theta})}{\partial \Delta\theta} - \tfrac{\partial q_E^N(\underline{\theta})}{\partial \Delta\theta} - \frac{1-\overline{\nu}}{2} \frac{\partial}{\partial \Delta\theta} \left(q_E^S(\underline{\theta}, \underline{\theta}) - q_E^S\left(\underline{\theta}, \overline{\theta}\right) \right) \right] \Delta\theta \tag{A11}$$
$$= \frac{2\delta q^C(\overline{\theta})(1-\overline{\nu})(1+\Delta\nu)^2}{3(4\underline{\nu}-\overline{\nu}-1)(1-2\Delta\nu)} \Delta\theta > 0,$$

where we have used equilibrium quantities from Propositions 1.1 and 1.2 and Assumption 1. ■

Proof of Proposition 1.5. For any $d \in \{S, N\}$, the ex ante joint profit of M and the R is $\Pi^d + \mathcal{V}^d$. For $\Delta\theta$ small, using Taylor approximations and the results of Proposition 1.4,

$$\left(\Pi^N + \mathcal{V}^N\right) - \left(\Pi^S + \mathcal{V}^S\right) \approx \frac{\delta q^C\left(\overline{\theta}\right)(1-\overline{\nu})(1+\Delta\nu)^2}{6(1-2\Delta\nu)(4\underline{\nu}-\overline{\nu}-1)}\Delta\theta.$$

This is strictly positive by Assumption 1. ■

Proof of Proposition 1.6. Let Π_E^d, $d \in \{S, N\}$, be the entrant's ex ante profit. Let $\underline{\theta} = \bar{\theta} - \Delta\theta$. By Taylor approximations around $\Delta\theta = 0$ we have

$$\begin{aligned}\Pi_E^N &\approx \lim_{\Delta\theta\to 0} \Pi_E^N + \lim_{\Delta\theta\to 0} \frac{\partial \Pi_E^N}{\partial \Delta\theta}\Delta\theta \\ &= q^C\left(\overline{\theta}\right)^2 - \frac{1}{2}\left[(1-\overline{\nu})\left(1+\frac{\partial q_2^N\left(\underline{\theta},\overline{\theta}\right)}{\partial\Delta\theta}\right) + \underline{\nu}\left(1+\frac{\partial q_2^N\left(\underline{\theta},\underline{\theta}\right)}{\partial\Delta\theta}\right)\right] q^C\left(\overline{\theta}\right)\Delta\theta,\end{aligned}$$

and

$$\begin{aligned}\Pi_E^S &\approx \lim_{\Delta\theta\to 0} \Pi_E^S + \lim_{\Delta\theta\to 0} \frac{\partial \Pi_E^S}{\partial \Delta\theta}\Delta\theta \\ &= q^C\left(\overline{\theta}\right)^2 - \frac{1}{2}\left[(1-\overline{\nu})\left(1+\frac{\partial q_2^S\left(\underline{\theta},\overline{\theta}\right)}{\partial\Delta\theta}\right) + \underline{\nu}\left(1+\frac{\partial q_2^S\left(\underline{\theta},\underline{\theta}\right)}{\partial\Delta\theta}\right)\right] q^C\left(\overline{\theta}\right)\Delta\theta.\end{aligned}$$

Hence,

$$\begin{aligned}\Pi_E^N - \Pi_E^S &\approx \left[(1-\overline{\nu})\frac{\partial}{\partial\Delta\theta}\left(q_2^S\left(\underline{\theta},\overline{\theta}\right) - q_2^N\left(\underline{\theta},\overline{\theta}\right)\right) + \underline{\nu}\frac{\partial}{\partial\Delta\theta}\left(q_2^S\left(\underline{\theta},\underline{\theta}\right) - q_2^N\left(\underline{\theta},\underline{\theta}\right)\right)\right]\frac{q^C\left(\bar{\theta}\right)\Delta\theta}{2} \\ &= -\frac{q^C\left(\overline{\theta}\right)(1+\Delta\nu)^2(1-\overline{\nu})}{3(1-2\Delta\nu)(4\underline{\nu}-\overline{\nu}-1)}\Delta\theta,\end{aligned}$$

which is negative under Assumption 1. ■

Proof of Proposition 1.7. Let ρ be the price for information about θ_1, that E pays to M. E is willing to buy information if and only if

$$\rho \leqslant \Pi_E^S - \Pi_E^N.$$

M is willing to sell information if and only if

$$\rho \geqslant \Pi_2^N - \Pi_2^S.$$

Hence, M and E are willing to trade at price $\rho > 0$ if and only if

$$\Pi_2^N - \Pi_2^S \leqslant \Pi_E^S - \Pi_E^N.$$

Under Assumption 2 this condition simplifies to

$$\frac{2q^C(\overline{\theta})\left(1-\overline{\nu}\right)\left(1+\Delta\nu\right)^2}{3\left(1-2\Delta\nu\right)\left(4\underline{\nu}-\overline{\nu}-1\right)}\Delta\theta \leqslant \frac{q^C\left(\overline{\theta}\right)\left(1-\overline{\nu}\right)\left(1+\Delta\nu\right)^2}{3\left(1-2\Delta\nu\right)\left(4\underline{\nu}-\overline{\nu}-1\right)}\Delta\theta,$$

which is never satisfied. ■

Proof of Proposition 1.8. M and R are willing to trade information if and only if

$$\mathcal{J}_2^N - \mathcal{J}_2^S \leqslant \Pi_E^S - \Pi_E^N,$$

where J_2^d is the joint profit of M and R in the second period. Using the results of Propositions 1.5 and 1.6, under Assumption 2 this condition simplifies to

$$\frac{q^C\left(\overline{\theta}\right)\left(1-\overline{\nu}\right)\left(1+\Delta\nu\right)^2}{6\left(1-2\Delta\nu\right)\left(4\underline{\nu}-\overline{\nu}-1\right)}\Delta\theta \leqslant \frac{q^C\left(\overline{\theta}\right)\left(1+\Delta\nu\right)^2\left(1-\overline{\nu}\right)}{3\left(1-2\Delta\nu\right)\left(4\underline{\nu}-\overline{\nu}-1\right)}\Delta\theta,$$

which is always satisfied. ■

Proof of Proposition 1.9. For any disclosure policy $d \in \{S, N\}$, expected aggregate production in the second period is

$$\mathcal{Q}^d \triangleq \sum_{\theta_1\in\Theta} \Pr\left[\theta_1\right] \sum_{\theta_2\in\Theta} \Pr\left[\theta_2|\theta_1\right] Q_2^d\left(\theta_2,\theta_1\right).$$

For $\Delta\theta$ small

$$\mathcal{Q}^d \approx \lim_{\Delta\theta\to 0} \mathcal{Q}^d + \Delta\theta \lim_{\Delta\theta\to 0} \frac{\partial}{\partial\Delta\theta}\mathcal{Q}^d,$$

with $\lim_{\Delta\theta\to 0}\mathcal{Q}^S = \lim_{\Delta\theta\to 0}\mathcal{Q}^N = 2q^C\left(\overline{\theta}\right)$, and

$$\lim_{\Delta\theta\to 0}\frac{\partial\mathcal{Q}^N}{\partial\Delta\theta} = \frac{1}{2}[\left(1-\overline{\nu}\right)\underbrace{\frac{\partial}{\partial\Delta\theta}\left(q_2^N\left(\underline{\theta},\overline{\theta}\right)+q_E^N\left(\underline{\theta}\right)\right)}_{\frac{\partial}{\partial\Delta\theta}Q^N(\underline{\theta},\overline{\theta})} + \underline{\nu}\underbrace{\frac{\partial}{\partial\Delta\theta}\left(q_2^N\left(\underline{\theta},\underline{\theta}\right)+q_E^N\left(\underline{\theta}\right)\right)}_{\frac{\partial}{\partial\Delta\theta}Q^N(\underline{\theta},\underline{\theta})}],$$

$$\lim_{\Delta\theta\to 0}\frac{\partial\mathcal{Q}^S}{\partial\Delta\theta} = \frac{1}{2}[\left(1-\overline{\nu}\right)\underbrace{\frac{\partial}{\partial\Delta\theta}\left(q_2^S\left(\underline{\theta},\overline{\theta}\right)+q_E^S\left(\underline{\theta},\overline{\theta}\right)\right)}_{\frac{\partial}{\partial\Delta\theta}Q^S(\underline{\theta},\overline{\theta})} + \underline{\nu}\underbrace{\frac{\partial}{\partial\Delta\theta}\left(q_2^S\left(\underline{\theta},\underline{\theta}\right)+q_E^S\left(\underline{\theta},\underline{\theta}\right)\right)}_{\frac{\partial}{\partial\Delta\theta}Q^S(\underline{\theta},\underline{\theta})}].$$

Hence,

$$\mathcal{Q}^S - \mathcal{Q}^N \approx \left[\left(1-\overline{\nu}\right)\frac{\partial}{\partial\Delta\theta}\left(Q^S\left(\underline{\theta},\overline{\theta}\right)-Q^N\left(\underline{\theta},\overline{\theta}\right)\right) + \underline{\nu}\frac{\partial}{\partial\Delta\theta}\left(Q^S\left(\underline{\theta},\underline{\theta}\right)-Q^N\left(\underline{\theta},\underline{\theta}\right)\right)\right]\frac{\Delta\theta}{2} \quad \text{(A12)}$$

$$= -\frac{\left(1-\overline{\nu}\right)\left(1+\Delta\nu\right)^2}{6\left(4\underline{\nu}-\overline{\nu}-1\right)\left(1-2\Delta\nu\right)}\Delta\theta < 0,$$

where we have used equilibrium quantities from Propositions 1.1 and 1.2 and Assumption 1. ■

Proof of Proposition 1.10. Without loss of generality, we focus on the second period, since production in the first period is the same with and without information sharing. For any $d \in$

$\{S,N\}$, since the inverse demand is linear, expected consumer surplus is

$$\mathcal{CS}^d = \sum_{\theta_1\in\Theta} \Pr[\theta_1] \sum_{\theta_2\in\Theta} \Pr[\theta_2|\theta_1] \frac{Q_2^d(\theta_2,\theta_1)^2}{2}.$$

For $\Delta\theta$ small

$$\mathcal{CS}^d \approx \lim_{\Delta\theta\to 0} \mathcal{CS}^d + \Delta\theta \lim_{\Delta\theta\to 0} \frac{\partial \mathcal{CS}^d}{\partial \Delta\theta},$$

with $\lim_{\Delta\theta\to 0} \mathcal{CS}^d = 2q^C\left(\overline{\theta}\right)^2$ and

$$\lim_{\Delta\theta\to 0} \frac{\partial \mathcal{CS}^d}{\partial \Delta\theta} = q^C\left(\overline{\theta}\right)\left[(1-\overline{\nu})\frac{\partial Q^d\left(\underline{\theta},\overline{\theta}\right)}{\partial\Delta\theta} + \underline{\nu}\frac{\partial Q^d\left(\underline{\theta},\underline{\theta}\right)}{\partial\Delta\theta}\right].$$

Hence,

$$\mathcal{CS}^N - \mathcal{CS}^S \approx \frac{q^C\left(\overline{\theta}\right)(1-\overline{\nu})(1+\Delta\nu)^2}{3\left(4\underline{\nu}-\overline{\nu}-1\right)(1-2\Delta\nu)}\Delta\theta,$$

which is strictly positive by Assumption 1.

Total (expected) welfare in the second period — i.e., the sum of M's expected profit, R's expected rent, E's expected profit and the expected consumer surplus — is

$$\mathcal{TW}^d \triangleq \sum_{\theta_1\in\Theta} \Pr[\theta_1] \sum_{\theta_2\in\Theta} \Pr[\theta_2|\theta_1] \left[\theta_2 Q_2^d(\theta_2,\theta_1) - \frac{1}{2}Q_2^d(\theta_2,\theta_1)^2\right].$$

For $\Delta\theta$ small, using a first-order Taylor approximation around $\Delta\theta = 0$,

$$\mathcal{TW}^d \approx \lim_{\Delta\theta\to 0} \mathcal{TW}^d + \Delta\theta \lim_{\Delta\theta\to 0} \frac{\partial \mathcal{TW}^d}{\partial \Delta\theta},$$

with $\lim_{\Delta\theta\to 0} \mathcal{TW}^d = 4q^C\left(\overline{\theta}\right)^2$ and

$$\lim_{\Delta\theta\to 0} \frac{\partial \mathcal{TW}^d}{\partial \Delta\theta} = \sum_{\theta_1} \Pr[\theta_1] \sum_{\theta_2} \Pr[\theta_2 = \underline{\theta}|\theta_1] \left[q^C\left(\overline{\theta}\right)\frac{\partial Q_2^d(\underline{\theta},\theta_1)}{\partial\Delta\theta} - 2q^C\left(\overline{\theta}\right)\right].$$

Hence,

$$\begin{aligned}\mathcal{TW}^N - \mathcal{TW}^S &\approx q^C\left(\overline{\theta}\right)\left[(1-\overline{\nu})\frac{\partial}{\partial\Delta\theta}\left(Q_E^N\left(\underline{\theta},\overline{\theta}\right) - Q_E^S\left(\underline{\theta},\overline{\theta}\right)\right) + \underline{\nu}\frac{\partial}{\partial\Delta\theta}\left(Q_E^N\left(\underline{\theta},\underline{\theta}\right) - Q_E^S\left(\underline{\theta},\underline{\theta}\right)\right)\right]\frac{\Delta\theta}{2}, \\ &= \frac{1}{2}\left[\mathcal{CS}^N - \mathcal{CS}^S\right] > 0.\end{aligned}$$

■

Proof of Proposition 1.11. M's maximization problem does not depend on the information sharing decision. Without information sharing, M_E's (relaxed) maximization program is

$$\max_{q_E(\cdot),U_E(\cdot)} \sum_{\theta_2\in\Theta} \Pr[\theta_2] \sum_{\theta_1\in\Theta} \Pr[\theta_1|\theta_2]\, P\left(\theta_2, q_E(\theta_2) + q_2^N(\theta_2,\theta_1)\right) q_E(\theta_2) +$$
$$-\Pr\left[\theta_2=\overline{\theta}\right]\left[\Delta\theta - \sum_{\theta_1\in\Theta} \Pr\left[\theta_1|\theta_2=\overline{\theta}\right] q_2^N\left(\overline{\theta},\theta_1\right) + \sum_{\theta_1\in\Theta} \Pr\left[\theta_1|\theta_2=\underline{\theta}\right] q_2^N\left(\underline{\theta},\theta_1\right)\right] q_E\left(\underline{\theta}\right).$$

Maximizing with respect to $q_E\left(\overline{\theta}\right)$ and $q_E\left(\underline{\theta}\right)$ yields

$$\underbrace{\frac{\overline{\nu}}{1+\Delta\nu}}_{\Pr\left[\theta_1=\overline{\theta}|\theta_2=\overline{\theta}\right]} \times \left[\overline{\theta} - 2q_E(\overline{\theta}) - q_2^N(\overline{\theta},\overline{\theta})\right] + \underbrace{\frac{1-\nu}{1+\Delta\nu}}_{\Pr\left[\theta_1=\underline{\theta}|\theta_2=\overline{\theta}\right]} \times \left[\overline{\theta} - 2q_E(\overline{\theta}) - q_2^N(\overline{\theta},\underline{\theta})\right] = 0,$$

$$\left[\underline{\theta} - 2q_E\left(\underline{\theta}\right)\right] - \left[\Delta\theta - q^C(\overline{\theta})\right] - 2\left[\underbrace{\frac{1-\overline{\nu}}{1-\Delta\nu}}_{\Pr\left[\theta_1=\overline{\theta}|\theta_2=\underline{\theta}\right]} \times q_2^N(\underline{\theta},\overline{\theta}) + \underbrace{\frac{\nu}{1-\Delta\nu}}_{\Pr\left[\theta_1=\underline{\theta}|\theta_2=\underline{\theta}\right]} \times q_2^N\left(\underline{\theta},\underline{\theta}\right)\right] = 0.$$

Using (A1)-(A3), it follows that $q_E^N\left(\overline{\theta}\right) = q_2^N\left(\overline{\theta},\theta_1\right)$, $q_E^N\left(\underline{\theta}\right) = q^C\left(\underline{\theta}\right) - \frac{2}{3}\Delta\theta$ and

$$q_2^N\left(\underline{\theta},\overline{\theta}\right) = q_2^N\left(\underline{\theta},\underline{\theta}\right) = q^C\left(\underline{\theta}\right) + \tfrac{1}{3}\Delta\theta.$$

By direct comparison of these quantities,

$$q_E^N\left(\underline{\theta}\right) < q^C\left(\overline{\theta}\right) < q_2^N\left(\underline{\theta},\overline{\theta}\right) = q_2^N\left(\underline{\theta},\underline{\theta}\right).$$

R obtains a non-negative rent in the second period since $\Delta\theta - \Delta q^N = 0$. Similarly, R_E rent is

$$\left[\Delta\theta - \sum_{\theta_1} \Pr\left[\theta_1|\theta_2=\overline{\theta}\right] q_2^N\left(\overline{\theta},\theta_1\right) + \sum_{\theta_1} \Pr\left[\theta_1|\theta_2=\underline{\theta}\right] q_2^N\left(\underline{\theta},\theta_1\right)\right] q_E^N\left(\underline{\theta}\right) = \frac{2}{3}\Delta\theta q_E^N\left(\underline{\theta}\right),$$

which is strictly positive. Finally, R's rent in the first period is

$$U_1^N\left(\overline{\theta}\right) = \Delta\theta q_1^N\left(\underline{\theta}\right) + \delta\overline{\nu}\left[\Delta\theta - \Delta q^N\right]\left[q_2^N\left(\underline{\theta},\underline{\theta}\right) - q_2^N\left(\underline{\theta},\overline{\theta}\right)\right],$$

which is positive for δ not too large. ■

Proof of Proposition 1.12. M's maximization problem is the same as in the baseline model. With information sharing, M_E's (relaxed) maximization problem is

$$\max_{q_E(\cdot),U_E(\cdot)} \sum_{\theta_2\in\Theta} \Pr[\theta_2|\theta_1]\, P\left(\theta_2, q_E(\theta_2,\theta_1) + q_2^S(\theta_2,\theta_1)\right) q_E(\theta_2,\theta_1) +$$
$$-\Pr\left[\theta_2=\overline{\theta}|\theta_1\right]\left[\Delta\theta - \left(q_2^S\left(\overline{\theta},\theta_1\right) - q_2^S\left(\underline{\theta},\theta_1\right)\right)\right] q_E\left(\underline{\theta},\theta_1\right).$$

Maximizing with respect to $q_E\left(\overline{\theta},\overline{\theta}\right)$ and $q_E\left(\underline{\theta},\overline{\theta}\right)$ yields

$$\overline{\theta}-2q_E\left(\overline{\theta},\overline{\theta}\right)-q_2^S\left(\overline{\theta},\overline{\theta}\right)=0,$$
$$\left(1-\overline{\nu}\right)\left(\underline{\theta}-2q_E\left(\underline{\theta},\overline{\theta}\right)-q_2^S\left(\underline{\theta},\overline{\theta}\right)\right)-\overline{\nu}\left(\Delta\theta-\Delta q_2^S\left(\overline{\theta}\right)\right)=0.$$

Using (A4) and (A5), it follows that $q_E^S\left(\overline{\theta},\overline{\theta}\right)=q_2^S\left(\overline{\theta},\overline{\theta}\right)=q^C\left(\overline{\theta}\right)$,

$$q_E^S\left(\underline{\theta},\overline{\theta}\right)=q^C\left(\underline{\theta}\right)-\frac{4\overline{\nu}}{3\left(3-4\overline{\nu}\right)}\Delta\theta,$$
$$q_2^S\left(\underline{\theta},\overline{\theta}\right)=q^C\left(\underline{\theta}\right)+\frac{2\overline{\nu}}{3\left(3-4\overline{\nu}\right)}\Delta\theta.$$

Maximizing with respect to $q_E\left(\overline{\theta},\underline{\theta}\right)$ and $q_E\left(\underline{\theta},\underline{\theta}\right)$ yields

$$\overline{\theta}-2q_E\left(\overline{\theta},\underline{\theta}\right)-q_2^S\left(\overline{\theta},\underline{\theta}\right)=0,$$
$$\underline{\nu}\left(\underline{\theta}-2q_E\left(\underline{\theta},\underline{\theta}\right)-q_2^S\left(\underline{\theta},\underline{\theta}\right)\right)-\left(1-\underline{\nu}\right)\left(\Delta\theta-\Delta q_2^S\left(\underline{\theta}\right)\right)=0.$$

Using (A4) and (A6), it follows that $q_E^S\left(\overline{\theta},\underline{\theta}\right)=q_2^S\left(\overline{\theta},\underline{\theta}\right)=q^C\left(\overline{\theta}\right)$,

$$q_E^S\left(\underline{\theta},\underline{\theta}\right)=q^C\left(\underline{\theta}\right)+\frac{2\left(2\underline{\nu}^2-3\underline{\nu}+\overline{\nu}+1\right)}{3\left(4\underline{\nu}^2-\overline{\nu}-1\right)}\Delta\theta,$$
$$q_2^S\left(\underline{\theta},\underline{\theta}\right)=q^C\left(\underline{\theta}\right)-\frac{2\left(3\underline{\nu}-2\underline{\nu}^2-\overline{\nu}-1+3\overline{\nu}\underline{\nu}\right)}{3\left(4\underline{\nu}^2-\overline{\nu}-1\right)}\Delta\theta.$$

Direct comparison of these outputs together with Assumption 3 yields

$$q_2^S\left(\underline{\theta},\underline{\theta}\right)-q_E^S\left(\underline{\theta},\underline{\theta}\right)=-\frac{2\overline{\nu}\underline{\nu}}{4\underline{\nu}^2-\overline{\nu}-1}\Delta\theta<0,$$
$$q_2^S\left(\underline{\theta},\overline{\theta}\right)-q_E^S\left(\underline{\theta},\overline{\theta}\right)=\frac{2\overline{\nu}}{3-4\overline{\nu}}\Delta\theta>0,$$
$$q^C\left(\underline{\theta}\right)-q_2^S\left(\underline{\theta},\overline{\theta}\right)=-\frac{2\overline{\nu}}{3\left(3-4\overline{\nu}\right)}\Delta\theta<0,$$

and

$$q^C\left(\underline{\theta}\right)-q_2^S\left(\underline{\theta},\underline{\theta}\right)=\frac{2\left(3\underline{\nu}-2\underline{\nu}^2-\overline{\nu}-1+3\overline{\nu}\underline{\nu}\right)}{3\left(4\underline{\nu}^2-\overline{\nu}-1\right)}\Delta\theta>0\quad\Leftrightarrow\quad\overline{\nu}\geqslant\frac{\left(1-\underline{\nu}\right)\left(1-2\underline{\nu}\right)}{3\underline{\nu}-1}.$$

In order to show that retailers obtain strictly positive rents in the second period, notice that

$$U_2^S(\overline{\theta},\overline{\theta}) \triangleq \left[\Delta\theta - \Delta q_E^S(\overline{\theta})\right] q_2^S(\underline{\theta},\overline{\theta}) = \frac{2-4\overline{\nu}}{3-4\overline{\nu}} q_2^S(\underline{\theta},\overline{\theta})\Delta\theta,$$

$$U_2^S(\overline{\theta},\underline{\theta}) \triangleq \left[\Delta\theta - \Delta q_E^S(\underline{\theta})\right] q_2^S(\underline{\theta},\underline{\theta}) = \frac{2\underline{\nu}(2\underline{\nu}-1)}{4\underline{\nu}^2-\overline{\nu}-1} q_2^S(\underline{\theta},\underline{\theta})\Delta\theta,$$

$$U_E^S(\overline{\theta},\overline{\theta}) \triangleq \left[\Delta\theta - \Delta q_2^S(\overline{\theta})\right] q_E^S(\underline{\theta},\overline{\theta}) = \frac{2(1-\overline{\nu})}{3-4\overline{\nu}} q_E^S(\underline{\theta},\overline{\theta})\Delta\theta,$$

$$U_E^S(\overline{\theta},\overline{\theta}) \triangleq \left[\Delta\theta - \Delta q_2^S(\underline{\theta})\right] q_E^S(\underline{\theta},\underline{\theta}) = \frac{2\underline{\nu}(2\underline{\nu}-\overline{\nu}-1)}{4\underline{\nu}^2-\overline{\nu}-1} q_E^S(\underline{\theta},\underline{\theta})\Delta\theta,$$

which are all strictly positive under Assumption 3.

Finally, R's rent in the first period is

$$U_1^S(\overline{\theta}) = \Delta\theta q_1^S(\underline{\theta}) + \delta\overline{\nu}\left[(\Delta\theta - \Delta q^S(\underline{\theta}))q_2^S(\underline{\theta},\underline{\theta}) - (\Delta\theta - \Delta q^S(\overline{\theta}))q_2^S(\underline{\theta},\overline{\theta})\right],$$

which is positive for δ not too large. ■

Proof of Proposition 1.13. When demand is low only in the second period,under Assumption 3,

$$q_E^N(\underline{\theta}) - q_E^S(\underline{\theta},\overline{\theta}) = -\frac{2-4\overline{\nu}}{3-4\overline{\nu}}\Delta\theta < 0$$

$$q_E^N(\underline{\theta}) - q_E^S(\underline{\theta},\underline{\theta}) = -\frac{2\underline{\nu}(2\underline{\nu}-1)}{4\underline{\nu}^2-\overline{\nu}-1}\Delta\theta < 0.$$

■

Proof of Proposition 1.14. Using (A11) and the equilibrium quantities in Propositions 1.11 and 1.12,

$$\Pi^N - \Pi^S \approx q^C(\overline{\theta})\,\delta\left[\frac{-2\overline{\nu}^3+\overline{\nu}^2(4\underline{\nu}+1)+\overline{\nu}(-14\underline{\nu}^2+\underline{\nu}+2)+10\underline{\nu}^2-3\underline{\nu}-1}{(3-4\overline{\nu})(4\underline{\nu}^2-\overline{\nu}-1)}\right]\Delta\theta. \tag{A13}$$

Since the denominator is positive by Assumption 3, the sign of (A13) depends on the sign of

$$\xi(\overline{\nu},\underline{\nu}) \triangleq -2\overline{\nu}^3+\overline{\nu}^2(4\underline{\nu}+1)+\overline{\nu}\left(-14\underline{\nu}^2+\underline{\nu}+2\right)+10\underline{\nu}^2-3\underline{\nu}-1,$$

with

$$\frac{\partial\xi(\overline{\nu},\underline{\nu})}{\partial\overline{\nu}} = -6\overline{\nu}^2+\overline{\nu}(8\underline{\nu}+2)-14\underline{\nu}^2+\underline{\nu}+2.$$

It can be shown that $\frac{\partial\xi(\overline{\nu},\underline{\nu})}{\partial\overline{\nu}} < 0$ in our relevant region of parameters. Since $\xi(4\underline{\nu}^2-1,\underline{\nu}) = 4\underline{\nu}^3(2\underline{\nu}-1)(7-16\underline{\nu}^2) > 0$ and $\xi(0.5,\underline{\nu}) = \frac{3}{2}\underline{\nu}(2\underline{\nu}-1) > 0$, (A13) is positive and, hence, M does not want to share information.

Using (A9) and the equilibrium quantities in Propositions 1.11 and 1.12,

$$\mathcal{V}^N - \mathcal{V}^S \approx -\frac{q^C\left(\overline{\theta}\right)\delta\underline{\nu}\left(2\underline{\nu}-1\right)\left(1+\Delta\nu\right)}{4\underline{\nu}^2-\overline{\nu}-1}\Delta\theta, \tag{A14}$$

which is negative under Assumption 3. Hence, R's ex ante rent is higher with information sharing.

We now compare M_E's expected profit with and without information sharing. Notice that

$$\lim_{\Delta\theta\to 0}\Pi_E^N = \lim_{\Delta\theta\to 0}\Pi_E^S = q^C\left(\overline{\theta}\right)^2,$$

and

$$\lim_{\Delta\theta\to 0}\frac{\partial\Pi_E^N}{\partial\Delta\theta} = -q^C\left(\overline{\theta}\right)\left[\Pr\left[\theta_1=\overline{\theta}|\theta_2=\underline{\theta}\right]\left(1+\frac{\partial q_2^N\left(\underline{\theta},\overline{\theta}\right)}{\partial\Delta\theta}\right)+\Pr\left[\theta_1=\underline{\theta}|\theta_2=\underline{\theta}\right]\left(1+\frac{\partial q_2^N\left(\underline{\theta},\underline{\theta}\right)}{\partial\Delta\theta}\right)\right],$$

$$\lim_{\Delta\theta\to 0}\frac{\partial\Pi_E^S}{\partial\Delta\theta} = -q^C\left(\overline{\theta}\right)\left[1+\frac{1}{2}\frac{\partial}{\partial\Delta\theta}\left(q_2^S\left(\underline{\theta},\overline{\theta}\right)+q_2^S\left(\underline{\theta},\underline{\theta}\right)\right)\right].$$

Hence, using a Taylor approximation around $\Delta\theta=0$ and the equilibrium quantities from Proposition 1.11 and 1.12,

$$\begin{aligned}\Pi_E^S - \Pi_E^N &\approx q^C\left(\overline{\theta}\right)\left[\frac{1-\overline{\nu}}{1-\Delta\nu}\frac{\partial q_2^N\left(\underline{\theta},\overline{\theta}\right)}{\partial\Delta\theta}+\frac{\underline{\nu}}{1-\Delta\nu}\frac{\partial q_2^N\left(\underline{\theta},\underline{\theta}\right)}{\partial\Delta\theta}-\frac{1}{2}\frac{\partial}{\partial\Delta\theta}\left(q_2^S\left(\underline{\theta},\overline{\theta}\right)+q_2^S\left(\underline{\theta},\underline{\theta}\right)\right)\right]\Delta\theta \\ &= q^C\left(\overline{\theta}\right)\frac{\overline{\nu}^2\left(3-4\underline{\nu}\right)-\overline{\nu}\left(4\underline{\nu}^2+\underline{\nu}-1\right)+2\underline{\nu}^2+3\underline{\nu}-2}{\left(3-4\overline{\nu}\right)\left(4\underline{\nu}^2-\overline{\nu}-1\right)}\Delta\theta. \end{aligned} \tag{A15}$$

Since the denominator is positive by Assumption 3, the sign of (A15) depends on the sign of

$$\mu\left(\overline{\nu},\underline{\nu}\right) \triangleq \overline{\nu}^2\left(3-4\underline{\nu}\right)-\overline{\nu}\left(4\underline{\nu}^2+\underline{\nu}-1\right)+2\underline{\nu}^2+3\underline{\nu}-2,$$

where it can be shown that, in the relevant region of parameters,

$$\frac{\partial\mu\left(\overline{\nu},\underline{\nu}\right)}{\partial\overline{\nu}} = \overline{\nu}\left(6-8\underline{\nu}\right)-4\underline{\nu}^2-\underline{\nu}+1<0.$$

Hence, since

$$\mu\left(4\underline{\nu}^2-1,\underline{\nu}\right) = 2\underline{\nu}^2\left(2\underline{\nu}-1\right)\left(7-16\underline{\nu}^2\right)>0$$

and

$$\mu\left(0.5,\underline{\nu}\right) = \frac{3}{4}\left(2\underline{\nu}-1\right)>0,$$

M_E's expected profit is higher with information sharing. ■

Proof of Proposition 1.15. Using (A12) and the equilibrium quantities in Propositions 1.11 and 1.12,

$$\mathcal{Q}^N-\mathcal{Q}^S \approx \frac{2\overline{\nu}^3-\overline{\nu}^2\left(16\underline{\nu}^2-4\underline{\nu}+1\right)+\overline{\nu}\left(16\underline{\nu}^3+2\underline{\nu}^2+\underline{\nu}-2\right)-12\underline{\nu}^3+8\underline{\nu}^2-3\underline{\nu}+1}{2\left(3-4\overline{\nu}\right)\left(4\underline{\nu}^2-\overline{\nu}-1\right)}\Delta\theta. \tag{A16}$$

The sign of (A16) depends on the sign of the numerator

$$\chi\left(\overline{\nu},\underline{\nu}\right)\triangleq 2\overline{\nu}^3-\overline{\nu}^2\left(16\underline{\nu}^2-4\underline{\nu}+1\right)+\overline{\nu}\left(16\underline{\nu}^3+2\underline{\nu}^2+\underline{\nu}-2\right)-12\underline{\nu}^3+8\underline{\nu}^2-3\underline{\nu}+1.$$

Following the proof of Proposition 1.12, first let $\underline{\nu}\leqslant 0.6$ so that $\overline{\nu}^*=4\underline{\nu}^2-1$. In this case,

$$\chi\left(0,\underline{\nu}\right)=-\left(2\underline{\nu}-1\right)\left(6\underline{\nu}^2-\underline{\nu}+1\right)<0,$$
$$\chi\left(4\underline{\nu}^2-1,\underline{\nu}\right)=2\underline{\nu}^2\left(7-16\underline{\nu}^2\right)\left(2\underline{\nu}-1\right)^2>0.$$

Moreover,

$$\frac{\partial\chi\left(\overline{\nu},\underline{\nu}\right)}{\partial\overline{\nu}}=6\overline{\nu}^2-\overline{\nu}\left(32\underline{\nu}^2-8\underline{\nu}+2\right)+16\underline{\nu}^3+2\underline{\nu}^2+\underline{\nu}-2.$$

Setting this equation equal to 0 and solving for $\overline{\nu}$ yields the critical points

$$\overline{\nu}_{\min}\triangleq\tfrac{8}{3}\underline{\nu}^2+\tfrac{1}{6}\sqrt{256\underline{\nu}^4-224\underline{\nu}^3+36\underline{\nu}^2-14\underline{\nu}+13}-\tfrac{2}{3}\underline{\nu}+\tfrac{1}{6}>0,$$
$$\overline{\nu}_{\max}\triangleq\tfrac{8}{3}\underline{\nu}^2-\tfrac{1}{6}\sqrt{256\underline{\nu}^4-224\underline{\nu}^3+36\underline{\nu}^2-14\underline{\nu}+13}-\tfrac{2}{3}\underline{\nu}+\tfrac{1}{6}>0.$$

Since

$$\lim_{\overline{\nu}\to\overline{\nu}_{\min}}\frac{\partial^2\chi\left(\overline{\nu},\underline{\nu}\right)}{\partial\overline{\nu}^2}=2\sqrt{256\underline{\nu}^4-224\underline{\nu}^3+36\underline{\nu}^2-14\underline{\nu}+13}>0,$$
$$\lim_{\overline{\nu}\to\overline{\nu}_{\max}}\frac{\partial^2\chi\left(\overline{\nu},\underline{\nu}\right)}{\partial\overline{\nu}^2}=-2\sqrt{256\underline{\nu}^4-224\underline{\nu}^3+36\underline{\nu}^2-14\underline{\nu}+13}<0,$$

$\chi\left(\overline{\nu},\underline{\nu}\right)$ has a relative minimum at $\overline{\nu}=\overline{\nu}_{\min}$ and relative maximum at $\overline{\nu}=\overline{\nu}_{\max}$. Finally, for $\underline{\nu}\leqslant 0.6$ the critical points are outside the interval of interest — i.e.,

$$\overline{\nu}_{\min}-\left(4\underline{\nu}^2-1\right)=\tfrac{1}{6}\sqrt{256\underline{\nu}^4-224\underline{\nu}^3+36\underline{\nu}^2-14\underline{\nu}+13}-\tfrac{2}{3}\underline{\nu}-\tfrac{4}{3}\underline{\nu}^2+\tfrac{7}{6}>0,$$
$$\overline{\nu}_{\max}-\left(4\underline{\nu}^2-1\right)=-\tfrac{1}{6}\sqrt{256\underline{\nu}^4-224\underline{\nu}^3+36\underline{\nu}^2-14\underline{\nu}+13}-\tfrac{2}{3}\underline{\nu}-\tfrac{4}{3}\underline{\nu}^2+\tfrac{7}{6}>0.$$

Hence, by the mean-value theorem there exists a unique $\overline{\nu}_0$ such that $\chi\left(\overline{\nu},\underline{\nu}\right)<0$ (so that aggregate production is higher with information sharing) if and only if $\overline{\nu}\leqslant\overline{\nu}_0\left(\underline{\nu}\right)$.

Second, consider the case where $\underline{\nu}>0.6$ so that $\overline{\nu}\leqslant\frac{1}{2}$. Notice that

$$\chi\left(0.5,\underline{\nu}\right)=\tfrac{1}{2}\underline{\nu}\left(2\underline{\nu}-1\right)\left(3-4\underline{\nu}\right)<0\quad\Leftrightarrow\quad\underline{\nu}>\underline{\nu}_0\triangleq 0.75.$$

Let $\underline{\nu}\geqslant\underline{\nu}_0$. The function $\chi\left(\overline{\nu},\underline{\nu}\right)$ has two critical points $\overline{\nu}_{\min}$ and $\overline{\nu}_{\max}$ and is always negative because

$$\overline{\nu}_{\min}-\tfrac{1}{2}=\tfrac{1}{6}\sqrt{256\underline{\nu}^4-224\underline{\nu}^3+36\underline{\nu}^2-14\underline{\nu}+13}-\tfrac{2}{3}\underline{\nu}+\tfrac{8}{3}\underline{\nu}^2-\tfrac{1}{3}>0,$$
$$\overline{\nu}_{\max}-\tfrac{1}{2}=-\tfrac{1}{6}\sqrt{256\underline{\nu}^4-224\underline{\nu}^3+36\underline{\nu}^2-14\underline{\nu}+13}-\tfrac{2}{3}\underline{\nu}-\tfrac{4}{3}\underline{\nu}^2+\tfrac{2}{3}>0.$$

Hence, in this region of parameters, consumers are better off with information sharing. Next, let $\underline{\nu}<\underline{\nu}_0$. Since $\chi\left(0,\underline{\nu}\right)<0$ and $\chi\left(0.5,\underline{\nu}\right)>0$, when $\overline{\nu}<\frac{1}{2}$ the function $\chi\left(\overline{\nu},\underline{\nu}\right)$ crosses the

$\overline{\nu}$-axis at least once. This point is unique because the relative maximum and minimum are outside the interval of interest — i.e., $\overline{\nu}_{\min} > \frac{1}{2}$ and $\overline{\nu}_{\max} > \frac{1}{2}$. Hence, if $\underline{\nu} < \underline{\nu}_0$ there exist a unique threshold $\overline{\nu}_0$ such that aggregate production is higher with information sharing if $\overline{\nu} \leqslant \overline{\nu}_0(\underline{\nu})$. Using numerical approximations, Figure 1.1 illustrates the region of parameters where consumers benefit from information sharing.

Finally, using Taylor approximations (see the proof of Proposition 1.10), there are linear relationships between total welfare and aggregate production,

$$\mathcal{TW}^N - \mathcal{TW}^S = q^C\left(\overline{\theta}\right)\left[\mathcal{Q}^N - \mathcal{Q}^S\right],$$

and between consumer surplus and aggregate production,

$$\mathcal{CS}^N - \mathcal{CS}^S = 2q^C\left(\overline{\theta}\right)\left[\mathcal{Q}^N - \mathcal{Q}^S\right].$$

Hence, whenever information sharing increases aggregate production, it also increases consumer surplus and total welfare. ■

Proof of Proposition 1.16. For any $d \in \{S, N\}$, the ex ante joint profit of M and R is $\mathcal{J}^d = \Pi^d + \mathcal{V}^d$. For $\Delta\theta$ small, using Taylor approximations and the results of Proposition 1.14, we have

$$\left[\Pi^N + \mathcal{V}^N\right] - \left[\Pi^S + \mathcal{V}^S\right] \approx \frac{\delta q^C\left(\overline{\theta}\right)\left(-2\overline{\nu}^3 + \overline{\nu}^2\left(8\underline{\nu}^2 + 1\right) - \overline{\nu}\left(8\underline{\nu}^3 + 8\underline{\nu}^2 - 2\right) + 6\underline{\nu}^3 + \underline{\nu}^2 - 1\right)}{\left(3 - 4\overline{\nu}\right)\left(4\underline{\nu}^2 - \overline{\nu} - 1\right)}\Delta\theta.$$

The sign of this expression depends on the numerator

$$\Psi\left(\overline{\nu}, \underline{\nu}\right) \triangleq -2\overline{\nu}^3 + \overline{\nu}^2\left(8\underline{\nu}^2 + 1\right) - \overline{\nu}\left(8\underline{\nu}^3 + 8\underline{\nu}^2 - 2\right) + 6\underline{\nu}^3 + \underline{\nu}^2 - 1,$$

with

$$\frac{\partial\Psi\left(\overline{\nu}, \underline{\nu}\right)}{\partial\overline{\nu}} = -6\overline{\nu}^2 + \overline{\nu}\left(16\underline{\nu}^2 + 2\right) - 8\underline{\nu}^3 - 8\underline{\nu}^2 + 2.$$

It can be shown that $\frac{\partial\Psi(\overline{\nu},\underline{\nu})}{\partial\overline{\nu}} < 0$ in the relevant region of parameters. Hence, since

$$\Psi\left(4\underline{\nu}^2 - 1, \underline{\nu}\right) = \underline{\nu}^2\left(2\underline{\nu} - 1\right)\left(7 - 16\underline{\nu}^2\right) > 0$$

and

$$\Psi\left(0.5, \underline{\nu}\right) = \underline{\nu}^2\left(2\underline{\nu} - 1\right) > 0,$$

the ex ante joint profit of M and R is higher without information sharing. ■

Proof of Proposition 1.17. Let ρ be the price for information about θ_1, that M_E pays to M. M_E is willing to buy information if and only if

$$\rho \leqslant \Pi_E^S - \Pi_E^N.$$

M is willing to sell information if and only if

$$\rho \geqslant \Pi_2^N - \Pi_2^S.$$

Hence, M and M_E are willing to trade at price $\rho > 0$ if and only if

$$\Pi_2^N - \Pi_2^S \leqslant \Pi_E^S - \Pi_E^N.$$

Under Assumption 3, this inequality simplifies to

$$\frac{q^C\left(\overline{\theta}\right)\left(2\overline{\nu}^3 - \overline{\nu}^2\left(8\underline{\nu} - 2\right) - \overline{\nu}\left(-10\underline{\nu}^2 + 2\underline{\nu} + 1\right) - 8\underline{\nu}^2 + 6\underline{\nu} - 1\right)}{\left(3 - 4\overline{\nu}\right)\left(4\underline{\nu}^2 - \overline{\nu} - 1\right)}\Delta\theta \geq 0. \qquad \text{(A17)}$$

The sign of this expression depends on the numerator

$$\kappa\left(\overline{\nu}, \underline{\nu}\right) \triangleq 2\overline{\nu}^3 - \overline{\nu}^2\left(8\underline{\nu} - 2\right) - \overline{\nu}\left(-10\underline{\nu}^2 + 2\underline{\nu} + 1\right) - 8\underline{\nu}^2 + 6\underline{\nu} - 1,$$

with

$$\frac{\partial\kappa\left(\overline{\nu}, \underline{\nu}\right)}{\partial\overline{\nu}} = 6\overline{\nu}^2 - \overline{\nu}\left(16\underline{\nu} - 4\right) + 10\underline{\nu}^2 - 2\underline{\nu} - 1.$$

It can be shown that in the relevant region of parameters $\frac{\partial\kappa(\overline{\nu},\underline{\nu})}{\partial\overline{\nu}} > 0$. Hence, since

$$\kappa\left(4\underline{\nu}^2 - 1, \underline{\nu}\right) = -2\underline{\nu}^2\left(7 - 16\underline{\nu}^2\right)\left(2\underline{\nu} - 1\right)^2 < 0$$

and

$$\kappa\left(0.5, \underline{\nu}\right) = -\frac{3}{4}\left(2\underline{\nu} - 1\right)^2 < 0,$$

inequality (A17) is never satisfied. ■

Proof of Proposition 1.18. M_E, M and R have a joint incentive to trade information if and only if

$$\mathcal{J}_2^N - \mathcal{J}_2^S \leqslant \Pi_E^S - \Pi_E^N.$$

Under Assumption 3, using the results of Proposition 1.14, this inequality simplifies to

$$\frac{q^C\left(\overline{\theta}\right)\left[2\overline{\nu}^3 - \overline{\nu}^2\left(8\underline{\nu}^2 + 4\underline{\nu} - 2\right) - \overline{\nu}\left(-8\underline{\nu}^3 - 4\underline{\nu}^2 + \underline{\nu} + 1\right) - 6\underline{\nu}^3 + \underline{\nu}^2 + 3\underline{\nu} - 1\right]}{\left(3 - 4\overline{\nu}\right)\left(4\underline{\nu}^2 - \overline{\nu} - 1\right)}\Delta\theta \geq 0. \qquad \text{(A18)}$$

The sign of the this expression depends on the sign of the numerator

$$\varpi\left(\overline{\nu}, \underline{\nu}\right) \triangleq 2\overline{\nu}^3 - \overline{\nu}^2\left(8\underline{\nu}^2 + 4\underline{\nu} - 2\right) - \overline{\nu}\left(-8\underline{\nu}^3 - 4\underline{\nu}^2 + \underline{\nu} + 1\right) - 6\underline{\nu}^3 + \underline{\nu}^2 + 3\underline{\nu} - 1.$$

First, let $\nu \leqslant 0.6$, so that $\overline{\nu}^* = 4\underline{\nu}^2 - 1$. In this case,

$$\varpi\left(0, \underline{\nu}\right) = -\left(2\underline{\nu} - 1\right)\left(\underline{\nu} + 3\underline{\nu}^2 - 1\right) < 0,$$
$$\varpi\left(4\underline{\nu}^2 - 1, \underline{\nu}\right) = \underline{\nu}^2\left(2\underline{\nu} - 1\right)\left(7 - 16\underline{\nu}^2\right) > 0,$$

and

$$\frac{\partial \varpi\left(\overline{\nu},\underline{\nu}\right)}{\partial \overline{\nu}} = 6\overline{\nu}^2 - \overline{\nu}\left(16\underline{\nu}^2 + 8\underline{\nu} - 4\right) + 8\underline{\nu}^3 + 4\underline{\nu}^2 - \underline{\nu} - 1 = 0$$

$$\Leftrightarrow \begin{cases} \overline{\nu}_{\min} \triangleq \frac{2}{3}\underline{\nu} + \frac{1}{6}\sqrt{2}\sqrt{32\underline{\nu}^4 + 8\underline{\nu}^3 - 20\underline{\nu}^2 - 5\underline{\nu} + 5} + \frac{4}{3}\underline{\nu}^2 - \frac{1}{3} > 0, \\ \overline{\nu}_{\max} \triangleq \frac{2}{3}\underline{\nu} - \frac{1}{6}\sqrt{2}\sqrt{32\underline{\nu}^4 + 8\underline{\nu}^3 - 20\underline{\nu}^2 - 5\underline{\nu} + 5} + \frac{4}{3}\underline{\nu}^2 - \frac{1}{3} > 0. \end{cases}$$

Since

$$\lim_{\overline{\nu}\to\overline{\nu}_{\min}} \frac{\partial^2 \varpi\left(\overline{\nu},\underline{\nu}\right)}{\partial \overline{\nu}^2} = 2\sqrt{2}\sqrt{-5\underline{\nu} - 20\underline{\nu}^2 + 8\underline{\nu}^3 + 32\underline{\nu}^4 + 5} > 0,$$
$$\lim_{\overline{\nu}\to\overline{\nu}_{\max}} \frac{\partial^2 \varpi\left(\overline{\nu},\underline{\nu}\right)}{\partial \overline{\nu}^2} = -2\sqrt{2}\sqrt{-5\underline{\nu} - 20\underline{\nu}^2 + 8\underline{\nu}^3 + 32\underline{\nu}^4 + 5} < 0,$$

$\varpi\left(\overline{\nu},\underline{\nu}\right)$ has a relative minimum at $\overline{\nu} = \overline{\nu}_{\min}$ and relative maximum at $\overline{\nu} = \overline{\nu}_{\max}$. The relative minimum is outside the interval of interest — i.e.,

$$\overline{\nu}_{\min} - \left(4\underline{\nu}^2 - 1\right) = -\frac{2}{3}\left(-\underline{\nu} - \frac{1}{4}\sqrt{2}\sqrt{-5\underline{\nu} - 20\underline{\nu}^2 + 8\underline{\nu}^3 + 32\underline{\nu}^4 + 5} + 4\underline{\nu}^2 - 1\right) > 0.$$

Hence, by the mean-value theorem, there exists a unique $\overline{\nu}_1$ such that $\varpi\left(\overline{\nu},\underline{\nu}\right) > 0$ if and only if $\overline{\nu}_1 \leqslant \overline{\nu}$.

Second, let $\underline{\nu} > 0.6$, so that $\overline{\nu} \leqslant \frac{1}{2}$. Notice that

$$\varpi(0.5,\underline{\nu}) = \tfrac{1}{4}\left(2\underline{\nu} - 1\right)\left(3 - 4\underline{\nu}^2\right) < 0 \quad \Leftrightarrow \quad \underline{\nu} > \underline{\nu}_1 \triangleq 0.87.$$

When $\underline{\nu} > \underline{\nu}_1$, $\varpi\left(\overline{\nu},\underline{\nu}\right)$ has two critical points $\overline{\nu}_{\min}$ and $\overline{\nu}_{\max}$ and is always negative because

$$\overline{\nu}_{\min} - \tfrac{1}{2} = \tfrac{2}{3}\left(\underline{\nu} + \tfrac{1}{4}\sqrt{2}\sqrt{-5\underline{\nu} - 20\underline{\nu}^2 + 8\underline{\nu}^3 + 32\underline{\nu}^4 + 5} + 2\underline{\nu}^2 - \tfrac{5}{4}\right) > 0,$$
$$\overline{\nu}_{\max} - \tfrac{1}{2} = \tfrac{2}{3}\left(\underline{\nu} - \tfrac{1}{4}\sqrt{2}\sqrt{-5\underline{\nu} - 20\underline{\nu}^2 + 8\underline{\nu}^3 + 32\underline{\nu}^4 + 5} + 2\underline{\nu}^2 - \tfrac{5}{4}\right) > 0.$$

Therefore, condition (A18) is not satisfied and players do not have incentive to sell information.

When $\underline{\nu} \leqslant \underline{\nu}_1$, $\varpi\left(0,\underline{\nu}\right) < 0$ and $\varpi\left(0.5,\underline{\nu}\right) > 0$. Hence, the function $\varpi\left(\overline{\nu},\underline{\nu}\right)$ crosses the $\overline{\nu}$ axis at least once. This point is unique because the relative minimum is outside the interval of interest — i.e., $\overline{\nu}_{\min} > \frac{1}{2}$.

Summing up, there exist a unique $\overline{\nu}_1\left(\underline{\nu}\right)$ such that M_E, M and R have a joint incentive to trade information if $\underline{\nu} \leqslant \underline{\nu}_1$ and $\overline{\nu} \geqslant \overline{\nu}_1\left(\underline{\nu}\right)$. The region of parameters where M and R sell information to M_E is illustrated in Figure 1.2 by numerical approximations. ■

Proof of Proposition 1.19. Using numerical approximations of the implicit functions defined by $\varpi\left(\overline{\nu},\underline{\nu}\right) = 0$ and $\chi\left(\overline{\nu},\underline{\nu}\right) = 0$, Figure 1.3 shows that $\overline{\nu}_1\left(\underline{\nu}\right) \leqslant \overline{\nu}_0\left(\underline{\nu}\right)$ for $\underline{\nu} \leqslant \underline{\nu}_1$, which proves the result. (The coding for the numerical approximations is available upon request.) ■

Proof of Proposition 1.20. Differentiating M's objective function, it is easy to show that $q_2\left(\overline{\theta},\theta_1,\sigma\right) = q_E\left(\overline{\theta},\sigma\right) = q^C\left(\overline{\theta}\right)$ for every θ_1 and σ. While, for every $\sigma \in \{\overline{\sigma},\underline{\sigma}\}$ we have

$$P_{q_2}\left(\underline{\theta}, Q_2^S\left(\underline{\theta},\overline{\theta},\sigma\right)\right) q_2\left(\underline{\theta},\overline{\theta},\sigma\right) + P\left(\underline{\theta}, Q_2^S\left(\underline{\theta},\overline{\theta},\sigma\right)\right) = 0, \tag{A19}$$

and

$$P_{q_2}\left(\underline{\theta}, Q_2^S\left(\underline{\theta}, \underline{\theta}, \sigma\right)\right) q_2\left(\underline{\theta}, \underline{\theta}, \sigma\right) + P\left(\underline{\theta}, Q_2^S\left(\underline{\theta}, \underline{\theta}, \sigma\right)\right) = \frac{\Delta\theta - \Delta q^S\left(\underline{\theta}, \sigma\right)}{\nu}. \tag{A20}$$

Solving (A19) and (A20) together with E's first-order conditions,

$$\begin{aligned} q_2\left(\underline{\theta}, \overline{\theta}, \overline{\sigma}\right) &= q^C\left(\underline{\theta}\right) + \frac{1-\beta}{3\left(1-\beta-3\left(\alpha\left(1-\nu\right)+\nu\left(1-\beta\right)\right)\right)}\Delta\theta \\ q_2\left(\underline{\theta}, \overline{\theta}, \underline{\sigma}\right) &= q^C\left(\underline{\theta}\right) - \frac{\beta}{3\left(3-\beta-3\left(\alpha\left(1-\nu\right)+\nu\left(1-\beta\right)\right)\right)}\Delta\theta \\ q_2\left(\underline{\theta}, \underline{\theta}, \underline{\sigma}\right) &= q^C\left(\underline{\theta}\right) - \frac{3\left(1-\nu\right)\left(1-\alpha\right)+4\beta\nu}{3\nu\left(3-\beta-3\left(\alpha\left(1-\nu\right)+\nu\left(1-\beta\right)\right)\right)}\Delta\theta \\ q_2\left(\underline{\theta}, \underline{\theta}, \overline{\sigma}\right) &= q^C\left(\underline{\theta}\right) + \frac{3\alpha\left(1-\nu\right)+4\nu\left(1-\beta\right)}{3\nu\left(1-\beta-3\left(\alpha\left(1-\nu\right)+\nu\left(1-\beta\right)\right)\right)}\Delta\theta \end{aligned}$$

and

$$\begin{aligned} q_E\left(\underline{\theta}, \underline{\sigma}\right) &= q^C\left(\underline{\theta}\right) + \frac{2\beta}{3\left(3-\beta-3\left(\alpha\left(1-\nu\right)+\nu\left(1-\beta\right)\right)\right)}\Delta\theta \\ q_E\left(\underline{\theta}, \overline{\sigma}\right) &= q^C\left(\underline{\theta}\right) - \frac{2\left(1-\beta\right)}{3\left(1-\beta-3\left(\alpha\left(1-\nu\right)+\nu\left(1-\beta\right)\right)\right)}\Delta\theta. \end{aligned}$$

Substituting these quantities into M's expected profit, maximizing with respect to α and β, respectively, and assuming that $\Delta\theta \to 0$, in an interior solution we have

$$\left(\alpha+\beta-1\right)\left(3\left(1-\nu\right)\left(2\beta-1\right)\alpha+\left(1-\beta\right)\left(3\left(1-\nu\right)-2\beta\left(1-3\nu\right)\right)\right)=0, \tag{A21}$$

and

$$\left(\alpha+\beta-1\right)\left(\left(6\beta\nu-9\nu-2\beta+7\right)\alpha-\left(1-3\nu\right)\left(1-\beta\right)-6\left(1-\nu\right)\alpha^2\right)=0. \tag{A22}$$

Solving with respect to α and β, the system of equations (A21)-(A22) features two critical points $(\alpha = 0, \beta = 1)$ and $(\alpha = 1, \beta = 0)$. Let M's expected profit in the second period be

$$\Pi\left(\alpha,\beta\right) \triangleq \frac{1-\nu}{2}\left[\alpha q_2\left(\underline{\theta}, \overline{\theta}, \overline{\sigma}\right)^2 + \left(1-\alpha\right) q_2\left(\underline{\theta}, \overline{\theta}, \underline{\sigma}\right)^2\right] + \\ + \frac{\nu}{2}\left[\beta q\left(\underline{\theta}, \underline{\theta}, \underline{\sigma}\right)^2 + \left(1-\beta\right) q\left(\underline{\theta}, \underline{\theta}, \overline{\sigma}\right)^2\right]$$

Notice that these two solutions are payoff-equivalent since they both imply an uninformative experiment — i.e.,

$$\Pi\left(\alpha=0,\beta=1\right)=\Pi\left(\alpha=1,\beta=0\right).$$

Moreover, it can be shown that

$$\Pi\left(\alpha=1,\beta=1\right)-\Pi\left(\alpha=1,\beta=0\right) \approx -\frac{2\left(1-\nu\right)}{9\left(3\nu-1\right)}\overline{\theta}\Delta\theta < 0.$$

Hence, M prefers not to disclose information since $\nu > \frac{1}{3}$ by Assumption 1.

Consider now R's incentive to share information. R's second-period expected rent is

$$\mathcal{U}(\alpha,\beta) \triangleq \frac{1-\beta}{2}\left[\Delta\theta-\left(q^{C}(\underline{\theta})-q(\underline{\theta},\overline{\sigma})\right)\right] q(\underline{\theta},\underline{\theta},\overline{\sigma})+ \\ +\frac{\beta}{2}\left[\Delta\theta-\left(q^{C}(\underline{\theta})-q(\underline{\theta},\underline{\sigma})\right)\right] q_2(\underline{\theta},\underline{\theta},\underline{\sigma}).$$

Maximizing with respect to α and β, respectively, and assuming $\Delta\theta \to 0$, in an interior solution, we have

$$(\alpha+\beta-1)(3(1-\nu)(1-2\beta)-(1-\beta)(3(1-\nu)-2\beta(1-3\nu)))=0, \tag{A23}$$

and

$$(\alpha+\beta-1)\left((3\nu-1)(1-\beta)-6(1-\nu)\alpha^2+(7-9\nu-2\beta(1-3\nu))\alpha\right)=0. \tag{A24}$$

The system of equations (A23)-(A24) features two payoff-equivalent solutions $(\alpha=0,\beta=1)$ and $(\alpha=1,\beta=0)$. Notice, however, that

$$\mathcal{U}(\alpha=1,\beta=1)-\mathcal{U}(\alpha=0,\beta=1)=\frac{1-\nu}{3(3\nu-1)}\overline{\theta}\Delta\theta>0.$$

Hence, $\mathcal{U}(\alpha,\beta)$ has a global maximum at $\alpha=\beta=1$, so that R would like to share information perfectly.

Finally, consider the effect of information sharing in consumer surplus. As before, it can be easily shown that for $\Delta\theta$ small, the effect on consumer and total welfare is equivalent to the effect on aggregate quantity — i.e.,

$$\mathcal{Q}(\alpha,\beta) \triangleq \frac{1-\nu}{2}\left(\alpha\left(q_2\left(\underline{\theta},\overline{\theta},\overline{\sigma}\right)+q(\underline{\theta},\overline{\sigma})\right)+(1-\alpha)\left(q_2\left(\underline{\theta},\overline{\theta},\underline{\sigma}\right)+q(\underline{\theta},\underline{\sigma})\right)\right)+ \\ +\frac{\nu}{2}\left(\beta\left(q_2(\underline{\theta},\underline{\theta},\underline{\sigma})+q(\underline{\theta},\underline{\sigma})\right)+(1-\beta)\left(q_2(\underline{\theta},\underline{\theta},\overline{\sigma})+q(\underline{\theta},\overline{\sigma})\right)\right).$$

Maximizing with respect to α and β, respectively, in an interior solution for $\Delta\theta \to 0$ we have

$$(\alpha+\beta-1)(3(1-\nu)(1-2\beta)\alpha-(1-\beta)(3(1-\nu)-2\beta(1-3\nu)))=0, \tag{A25}$$

and

$$(\alpha+\beta-1)\left((3\nu-1)(1-\beta)+(6\beta\nu-9\nu-2\beta+7)\alpha-6(1-\nu)\alpha^2\right)=0. \tag{A26}$$

The system of equations (A25)-(A26) features two payoff-equivalent solutions $(\alpha=0,\beta=1)$ and $(\alpha=1,\beta=0)$. Moreover,

$$\mathcal{Q}(\alpha=1,\beta=1)-\mathcal{Q}(\alpha=1,\beta=0)\approx-\frac{1}{6}\frac{1-\nu}{3\nu-1}\Delta\theta<0,$$

so that consumer surplus is maximized by an uninformative experiment. ■

Proof of Proposition 1.21. Consider first the outcome without information sharing characterized in Section 1.4.1. In order to show that it is not robust to ex-post renegotiation, consider R and suppose that: (*i*) $\theta_1=\overline{\theta}$, and (*ii*) $m_1=\overline{\theta}$. Then, R has an incentive to disclose m_1 to E if and only if

$$\underbrace{\overline{\nu}\left[\Delta\theta - \left(q^C\left(\overline{\theta}\right) - q_E^N\left(\underline{\theta}\right)\right)\right] q_2^N\left(\underline{\theta},\overline{\theta}\right)}_{\text{Second-period equilibrium rent}} <$$

$$\underbrace{\sum_{\theta_2} \Pr\left[\theta_2|\overline{\theta}\right]\left[P\left(\theta_2, q_2^N\left(\theta_2,\overline{\theta}\right) + q_E^R\left(\theta_2,\overline{\theta}\right)\right) q_2^N\left(\theta_2,\overline{\theta}\right) - t_2^N\left(\theta_2,\overline{\theta}\right)\right]}_{\text{Deviation profit}},$$

where

$$q_E^R\left(\theta_2,\overline{\theta}\right) = \arg\max_{q_E}\left\{P\left(\theta_2, q_E + q_2^N\left(\theta_2,\overline{\theta}\right)\right) q_E\right\}.$$

Notice that, by definition

$$t_2^N\left(\overline{\theta},\overline{\theta}\right) \equiv P\left(\overline{\theta}, q_2^N\left(\overline{\theta},\overline{\theta}\right) + q_E^R\left(\overline{\theta},\overline{\theta}\right)\right) q_2^N\left(\overline{\theta},\overline{\theta}\right) - \left[\Delta\theta - \left(q^C\left(\overline{\theta}\right) - q_E^N\left(\underline{\theta}\right)\right)\right] q_2^N\left(\underline{\theta},\overline{\theta}\right),$$

and

$$t_2^N\left(\underline{\theta},\overline{\theta}\right) \equiv P\left(\underline{\theta}, q_2^N\left(\underline{\theta},\overline{\theta}\right) + q_E^R\left(\underline{\theta},\overline{\theta}\right)\right) q_2^N\left(\underline{\theta},\overline{\theta}\right).$$

R's incentive to disclose m_1 then rewrites as

$$0 < \left(1-\overline{\nu}\right) q_2^N\left(\underline{\theta},\overline{\theta}\right)\left(q_E^N\left(\underline{\theta}\right) - q_E^R\left(\underline{\theta},\overline{\theta}\right)\right). \tag{A27}$$

It can be shown that

$$q_E^R\left(\underline{\theta},\overline{\theta}\right) \equiv \frac{\underline{\theta} - q_2^N\left(\underline{\theta},\overline{\theta}\right)}{2} = q^C\left(\underline{\theta}\right) + \frac{\left(1+\Delta\nu\right)\left(3\left(1-\Delta\nu\right)+\underline{\nu}\right)}{12\underline{\nu}\left(1-2\Delta\nu\right)}\Delta\theta,$$

so that

$$q_E^N\left(\underline{\theta}\right) - q_E^R\left(\underline{\theta},\overline{\theta}\right) = \frac{1+\Delta\nu}{4\left(1-2\Delta\nu\right)},$$

which is positive under Assumption 1. Hence, (A27) holds.

The equilibrium with information sharing is robust to ex-post renegotiation because M cannot improve its profit from concealing information to E since, by assumption, the second period transfer cannot be reneged on. Hence, M does not deviate. This implies that, by the intertemporal incentive compatibility constraint, R cannot deviate either. ■

Large uncertainty. We derive the functions plotted in Figures 1.4 and 1.5. Under the parametric restrictions imposed in Section 1.7.4, the first-order conditions without information sharing imply

$$q_E^N\left(\overline{\theta}\right) = \frac{1}{3} + \frac{1}{3}\Delta\theta, \qquad q_2^N\left(\overline{\theta},\theta_1\right) = \frac{1}{3} + \frac{1}{3}\Delta\theta, \qquad \forall\theta_1 \in \Theta,$$

$$q_E^N\left(\underline{\theta}\right) = \frac{1}{3} + \frac{1}{3}\Delta\theta, \qquad q_2^N\left(\underline{\theta},\overline{\theta}\right) = \frac{1}{3} - \frac{1}{6}\Delta\theta, \qquad q_2^N\left(\underline{\theta},\underline{\theta}\right) = \frac{1}{3} - \frac{3+\nu}{6\nu}\Delta\theta.$$

Similarly, the first-order conditions with information sharing imply

$$q_E^S\left(\overline{\theta},\theta_1\right) = \frac{1}{3} + \frac{1}{3}\Delta\theta, \qquad q_2^S\left(\overline{\theta},\theta_1\right) = \frac{1}{3} + \frac{1}{3}\Delta\theta, \qquad \forall\theta_1 \in \Theta,$$

$$q_E^S\left(\underline{\theta},\overline{\theta}\right)=q_2^S\left(\underline{\theta},\overline{\theta}\right)=\frac{1}{3},\qquad q_E^S\left(\underline{\theta},\underline{\theta}\right)=\frac{1}{3}+\frac{2}{3\left(3\nu-1\right)}\Delta\theta,\qquad q_2^S\left(\underline{\theta},\underline{\theta}\right)=\frac{1}{3}-\frac{4}{3\left(3\nu-1\right)}\Delta\theta.$$

Notice that $q_2^S\left(\underline{\theta},\underline{\theta}\right)<q_2^N\left(\underline{\theta},\underline{\theta}\right)$ for $\nu\geqslant\frac{1}{3}$. Hence, we need to impose $\Delta\theta\leqslant\Delta\theta_0\left(\nu\right)\triangleq\frac{3\nu-1}{4}$ to guarantee that the incumbent does not shut down production when demand is repeatedly low.

Consider now R's expected rent. With no information sharing, R's second-period rent is strictly positive since

$$U_2^N\left(\overline{\theta},\theta_1\right)=\Delta\theta-\Delta q^N=\Delta\theta>0.$$

R's first-period rent with no information sharing is

$$U_1^N\left(\overline{\theta}\right)=\frac{\Delta\theta}{2}-\Delta\theta^2,$$

which is strictly positive for $\Delta\theta\leqslant\Delta\theta_0\left(\nu\right)$, where $\Delta\theta_0\left(\nu\right)>0$.

With information sharing, R's second-period rent is strictly positive since, by Assumption 1,

$$U_2^S\left(\overline{\theta},\overline{\theta}\right)=\Delta\theta-\Delta q^S\left(\overline{\theta}\right)=\frac{2}{3}\Delta\theta>0,$$
$$U_2^S\left(\overline{\theta},\underline{\theta}\right)=\Delta\theta-\Delta q^S\left(\underline{\theta}\right)=\frac{2\nu}{3\nu-1}\Delta\theta>0,$$

R's first-period rent is

$$U_1^S(\overline{\theta})=\frac{\left(93\nu^2-58\nu+9\right)\Delta\theta-\left(129\nu^2-54\nu+9\right)\Delta\theta^2}{18\left(3\nu-1\right)^2}.$$

The sign of this expression depends on the numerator

$$\sigma\left(\nu,\Delta\theta\right)\triangleq-\Delta\theta^2\left(129\nu^2-54\nu+9\right)+\Delta\theta\left(93\nu^2-58\nu+9\right),$$

with

$$\frac{\partial\sigma\left(\nu,\Delta\theta\right)}{\partial\Delta\theta}=-\Delta\theta\left(258\nu^2-108\nu+18\right)+93\nu^2-58\nu+9>0$$

in the relevant region of parameters. Hence, since $\sigma\left(\nu,0\right)=0$, and

$$\sigma\left(\nu,\frac{3\nu-1}{4}\right)=\frac{1}{16}\left(3\nu-1\right)^3\left(45-43\nu\right)>0,$$

R's first-period information rent with information sharing is positive.

We now compare R's ex ante rent with and without information sharing. For any $d\in\{S,N\}$ R's ex ante rent is

$$\mathcal{V}^d=\sum_{\theta_1}\Pr\left[\theta_1\right]\left[U_1^d\left(\theta_1\right)+\sum_{\theta_2}\Pr\left[\theta_2=\overline{\theta}|\theta_1\right]U_2^d\left(\overline{\theta},\theta_1\right)\right].\qquad\text{(A28)}$$

Using the first and second period information rents just derived, it can be shown that

$$\mathcal{V}^N - \mathcal{V}^S = \frac{(1-\nu)\left(9\nu^2 + 14\nu - 3\right)\Delta\theta + (1-\nu)\left(2\nu - 6\nu^2\right)}{12\nu\left(3\nu - 1\right)^2}\Delta\theta.$$

Setting the numerator equal to 0 and solving for $\Delta\theta$ yields

$$\Delta\theta_u\left(\nu\right) \triangleq \frac{2\nu\left(3\nu-1\right)}{14\nu + 9\nu^2 - 3},$$

which is positive in the relevant region of parameters. Figure 1.4 plots the threshold $\Delta\theta_u\left(\nu\right)$ such that R's ex ante rent is the same with and without information sharing — i.e., $\mathcal{V}^N = \mathcal{V}^S$.

Second, we compare M's expected profit with and without information sharing. Without loss of generality, we focus on the second period, since production in the first period is the same with and without information sharing. For any disclosure policy $d \in \{S, N\}$, M's expected profit is

$$\Pi^d = \sum_{\theta_1} \Pr\left[\theta_1\right] \sum_{\theta_2} \Pr\left[\theta_2|\theta_1\right] q_2^d\left(\theta_2, \theta_1\right)^2.$$

Hence,

$$\Pi^N - \Pi^S = \frac{\left((1-\nu)\left(9 - 63\nu^2 - 38\nu\right)\Delta\theta + \left(48\nu^2 - 16\nu\right)(1-\nu)\right)}{72\nu\left(3\nu-1\right)^2}\Delta\theta.$$

Setting the numerator equal to 0 and solving for $\Delta\theta$ yields

$$\Delta\theta_\pi\left(\nu\right) \triangleq \frac{16\nu\left(3\nu-1\right)}{38\nu + 63\nu^2 - 9},$$

which is positive in the relevant region of parameters. Figure 1.5 plots the threshold $\Delta\theta_\pi\left(\nu\right)$ such that M's expected profit is the same with and without information sharing — i.e., $\Pi^N = \Pi^S$.

Finally, showing that $\Delta\theta_\pi\left(\nu\right) \geqslant \Delta\theta_u\left(\nu\right)$ is immediate. ■

References

[1] Arve, M., and D. Martimort (2016). "Dynamic procurement under uncertainty: optimal design and implications for incomplete contracts." *American Economic Review*, 106, 3238-3274.

[2] Baron, D.P., and D. Besanko (1984). "Regulation and information in a continuing relationship." *Information Economics and Policy,* 1, 267-302.

[3] Battaglini, M. (2005). "Long term contracting with markovian consumers." *American Economic Review*, 95, 637-658.

[4] Battaglini, M., and S. Coate (2008). "A dynamic theory of public spending, taxation, and debt." *American Economic Review*, 98, 201-236.

[5] Battaglini, M., and L. Lamba (2015). "Optimal dynamic contracting: the first-order approach and beyond." Mimeo, Cornell University and Penn State University.

[6] Bennardo, A., M. Pagano and S. Piccolo (2015). "Multiple-bank lending, creditor rights and information sharing." *Review of Finance*, 19, 519-570.

[7] Bergemann, D., and A. Pavan (2015). "Introduction to symposium on dynamic contracts and mechanism design." *Journal of Economic Theory,* 159, 679-701.

[8] Bergemann D., A. Bonatti, and A. Smolin (2017). "The Design and Price of Information." Mimeo, MIT Sloan.

[9] Bernheim, B.D., (1984). "Strategic deterrence of sequential entry into an industry." *RAND Journal of Economics,* 15, 1-11.

[10] Caillaud, B., B. Jullien and P. Picard (1995). "Competing vertical structures: precommitment and renegotiation." *Econometrica,* 63, 621-646.

[11] Calzolari, G., and A. Pavan (2006a). "On the optimality of privacy in sequential contracting." *Journal of Economic Theory,* 130, 168-204.

[12] Calzolari, G., and A. Pavan (2006b). "Monopoly with resale." *RAND Journal of Economics*, 37, 362-375.

[13] Clarke, R. (1983). "Collusion and the incentives for information sharing." *Bell Journal of Economics,* 14, 383-394.

[14] Dutta, K.D., S. Lach and A. Rustichini (1995). "Better late than early: vertical differentiation in the adoption of a new technology." *Journal of Economics and Management Strategy,* 4, 563-589.

[15] Eső, P., and B. Szentes (2017). "Dynamic contracting: an irrelevance result." *Theoretical Economics,* 12, 109-39.

[16] Fried, D. (1984). "Incentives for information production and disclosure in a duopolistic environment." *Quarterly Journal of Economics,* 99, 367-381.

[17] Gal-Or, E. (1985). "Information sharing in oligopoly." *Econometrica,* 53, 329-343.

[18] Gal-Or, E. (1986). "Information transmission–Cournot and Bertrand equilibria." *Review of Economic Studies,* 53, 85-92.

[19] Gal-Or, E., (1991). "Vertical restraints with incomplete information". *The Journal of Industrial Economics*, 39, 503–516.

[20] Gal-Or, E. (1999). "Vertical integration or separation of the sales function as implied by competitive forces." *International Journal of Industrial Organization,* 17, 641-662.

[21] Garrett, D., and A. Pavan (2012). "Managerial turnover in a changing world." *Journal of Political Economy,* 120, 879-925.

[22] Gennaioli, N., and G. Ponzetto (2017). "Optimally Vague Contracts and the Law." Mimeo, Bocconi University.

[23] Geroski, P.A. (1995). "What do we know about entry?" *International Journal of Industrial Organization,* 13, 421-440.

[24] Gu, Y., and T. Wenzel (2011). "Transparency, price-dependent demand and product variety." *Economics Letters,* 110, 216-219.

[25] Kastl, J., D. Martimort and S. Piccolo (2011). "When should manufacturers want fair trade? New insights from asymmetric information." *Journal of Economics and Management Strategy,* 3, 649-677.

[26] Kastl, J., M. Pagnozzi and S. Piccolo (2018). "Selling information to competitive firms." *RAND Journal of Economics*, 49, 254-282.

[27] Laffont, J.-J., and D. Martimort (2002). The Theory of Incentives: The Principal-Agent Model. Princeton University Press, Princeton.

[28] Laffont, J.-J., and J. Tirole (1988). The dynamics of incentive contracts. *Econometrica,* 56, 1153-1175.

[29] Laffont, J.-J., and J. Tirole (1996). Pollution permits and compliance strategies. *Journal of Public Economics,* 62, 85-125.

[30] Li, L. (1985). "Cournot oligopoly with information sharing." *RAND Journal of Economics,* 16, 521-536.

[31] Maier, N., and M. Ottaviani (2009). "Information sharing in common agency: when is transparency good?" *Journal of the European Economic Association,* 7(1), 162-187.

[32] Martimort, D. (1996). "Exclusive dealing, common agency, and multiprincipals incentive theory." *RAND Journal of Economics,* 27, 1-31.

[33] Martimort, D. (1999). "Renegotiation design with multiple regulators." *Journal of Economic Theory,* 88, 261-293.

[34] Martimort, D., A. Semenov and L. Stole (2017). "A theory of contracts with limited enforcement". *Review of Economic Studies,* 84, 816-852.

[35] Maskin, E.S., and J.G. Riley (1985). "Auction theory with private values." *American Economic Review*, 75, 150-155.

[36] Mussa, M., and S. Rosen (1978). "Monopoly and product quality." *Journal of Economic Theory,* 18, 301-317.

[37] Novshek, W., and H. Sonnenschein (1982). "Fulfilled expectations Cournot duopoly with information acquisition and release." *Bell Journal of Economics*, 13, 214-218.

[38] Oh, F.D., and J. Park (2017). "Credit ratings and corporate disclosure behaviour: evidence from regulation fair disclosure in Korea" *Applied Economics*, 49, 3481-3494.

[39] Pavan, A., I. Segal and J. Toikka (2014). "Dynamic mechanism design: a Myersonian approach." *Econometrica,* 63, 1-27.

[40] Piccolo, S., and M. Pagnozzi (2013). "Information sharing between vertical hierarchies." *Games and Economic Behavior,* 79, 201-222.

[41] Piccolo, S., E. Tarantino and G. Ursino (2015). "The value of transparency in multidivisional firms." *International Journal of Industrial Organization,* 41, 9-18.

[42] Raith, M. (1996). "A general model of information sharing in oligopoly." *Journal of Economic Theory,* 71, 260-288.

[43] Schultz, C. (2009). "Transparency and product variety." *Economic Letters,* 102, 165-168.

[44] Shapiro, C. (1986). "Exchange of cost information in oligopoly." *Review of Economic Studies,* 53, 433-446.

[45] Varian, H. (1980). "A model of sales." *American Economic Review,* 70, 651-659.

[46] Vives, X. (1984). "Duopoly information equilibrium: Cournot and Bertrand." *Journal of Economic Theory,* 34, 71-94.

[47] Vives, X. (2006). Information sharing among firms. The New Palgrave Dictionary of Economics.

Chapter 2

Terrorism, Counterterrorism and Optimal Striking Rules

2.1 Introduction

After September 11, the war on terror has became a primary objective of the western world. Since then, a wide campaign against religious terrorism (primarily Al-Qaeda) has been undertaken, involving open and covert military operations, new security legislation, efforts to block the financing of terrorism, etc. Yet, despite remarkable efforts, the problem seems far from being solved. After the death of Osama bin Laden in 2011, the rise of the ISIS (Islamic State of Iraq and Syria) and the recent dramatic attacks to the heart of Europe have increased again the alert, and thrown serious doubts on the way the risks associated with the new threat have been assessed. After more than 15 years, the major obstacle to the implementation of timely and effective measures is still the 'hidden face' of terror. When will they strike? Where? And, how aggressively?

Answering these questions is difficult, especially in the absence of reliable information on terrorists' activities, network, strength, equipment, etc. This is why, in many cases, it could be useful to collaborate with the local communities hosting the terrorists. Noncombatants, indeed, are likely to own insider knowledge that would hardly be acquired by intelligence agencies — see, e.g., Kalyvas (2006). According to Berman *et al.* (2011), noncombatants are responsive and active actors. Popkin (1979) argues that they make rational decisions regarding the direction and degree of their cooperation, while Galula (1964) and Petersen (2001) show that their propensity to do so varies at the individual level and shifts across space and time.

Joint work with S. Piccolo, University of Bergamo and CSEF, and with G. Immordino, University of Naples Federico II and CSEF. We are indebted to Giacomo Calzolari, Nenad Kos and David Martimort for their insightful comments and suggestions. Comments received by an anonymous referee are gratefully acknowledged. The audience attending the Mannheim Graduate Students Seminar and the conference in honor of David Martimort (Paris, 2017) provided useful suggestions to improve the paper.

Hence, the interaction between target countries, terrorists, and the communities whose cooperation they compete for can be best understood by accounting for their preferences and incentives (Nagl, 2002; Sepp, 2005; Petraeus, 2006; Fridovich and Krawchuk, 2007; Cassidy, 2008; McMaster, 2008). To investigate these incentives we study a stylized mechanism design problem that describes the optimal behavior of a country targeted by a foreign terrorist group. The country is uncertain about the terrorists' strength (measured by the magnitude of the damage produced by an attack) and may decide to acquire such information from the community hosting the terrorists. As a prize for collaboration, the target country provides the community with basic public goods, infrastructures, technological knowledge, etc.

We highlight a novel trade-off between target hardening, which requires the country to invest public funds to mitigate the incidence of the attack — e.g., by strengthening internal controls and improving citizens' protection — and preemptive attacks aimed at eradicating the problem at its root. Specifically, we show that, conditional on being informed about the terrorists' strength, the country engages in a preemptive attack only when the threat that it faces is sufficiently serious — i.e., the potential damage to civilians and public infrastructures is relatively large — and when the community norms of noncombatants favoring terrorists are weak — i.e., when the terrorists are not very much rooted in the country hosting them. The key point hinges on the mechanism design approach we adopt: in order to elicit truthful information about the strength of the terrorist group, the target country has to grant an information rent to the noncombatants, who would otherwise have an incentive to overstate the strength of the terrorists in order to amplify the risk of retaliation (which occurs when the strike fails) and induce the country to provide larger provisions of public goods to guarantee their collaboration.

Hence, in order to reduce these rents, the country is forced to engage in preemptive attacks relatively more often than in the first-best (i.e., the scenario in which the terrorists' strength is common knowledge) when facing strong terrorist groups. Indeed, *ceteris paribus*, being aggressive to strong terrorist groups mitigates (or even nullifies) the incentive of the noncombatants to overstate the risk of retaliation: a *rent saving* effect. Yet, a preemptive strike is also costly because a military intervention would encroach on the community's internal norms, whose members might either be resilient to foreign interference with their territorial sovereignty, or have ties with the terrorists: a *norm breaching* effect. The optimal policy trades off these two effects. Clearly, although the county has an incentive to strike strong terrorist groups, it does less so as the community features stronger norms.

Building on this insight we then turn to characterize the country's information acquisition decision. The benchmark against which we compare the country's expected payoff from acquiring information is that in which there is no interaction with noncombatants, and the policy is chosen *behind the veil of ignorance* — i.e., without knowing the terrorists' strength. We show that acquiring information is optimal only when: (i) the community features strong enough internal norms or non-negligible ties with the terrorist group; (ii) the target country's prior information about the terrorists' strength is sufficiently poor. Intuitively, the country's decision between acquiring information and remaining uninformed is shaped by the following trade-off. On the one hand, by not acquiring information the country bears the risk of making a decision that is inefficient ex-post. On the other hand, when the country decides to remain uninformed, whatever decision is taken, it saves on the cost of acquiring information, which as argued above depends on the strength of the community norms and the rent that the community has to be granted in order to truthfully report its private information.

Surprisingly, when the community of noncombatants features sufficiently weak internal norms, the country prefers not to acquire information. The intuition is straightforward. In this case, conditional on acquiring information, the country is most likely to strike. Yet, when it does so, its ex post payoff is lower than the payoff it obtains when a preemptive strike is conducted behind the veil of ignorance: no public good has to be provided to the community since no information is transmitted. By contrast, as the community norms become stronger, the country is relatively less likely to strike upon acquiring information. Indeed, the value of acquiring information lies in the ability of the informed country to tailor its target hardening choice to the terrorists' damage because this eliminates the risk of making a wrong decision (as opposed to the case of no information acquisition). As a result, the larger is the set of contingencies (states of nature) in which the country decides not to strike, the higher is the value of information. Since it is relatively more appealing for the country to invest in target hardening as the community norms become stronger, the value of information increases too.

Hence, the model delivers an inverse relationship between the incentive to acquire information about the strength of terrorist groups and the norms featured by the communities hosting these terrorists. This result contrasts with the equilibrium characterization offered by Berman *et al.* (2011) who find that norms favoring rebel control reduce incentives to disclose information in a model where information is not elicited through incentive compatible mechanisms but it is spontaneously disclosed by the community in a 'cheap talk' fashion. Their model features an equilibrium in which communities disclose information if and only if norms are sufficiently weak. By contrast, in our environment countries have an incentive to acquire this information if and only if norms are sufficiently strong.

Finally, notice that although our model is very stylized, we believe that it captures some salient aspects of reality and provides a few novel insights that might help designing future empirical work.

First, the distinction between defensive and offensive strategies and their optimal use contingent on the severity of the terrorists potential damage, is a natural aspect of the problem in real life. The way these instruments are used in real life seems to vary a lot over time and across countries. In a recent report by the NATO, it is argued that countering strategies need to develop a wider range of responses commensurate with each threat type, blending military, law enforcement and other civilian agencies where appropriate.[1] In spite of this, it is reported that while some

[1] http://www.coedat.nato.int/publication/researches/04-FutureTrends.pdf

nations have been engaged in countering insurgencies in Iraq and Afghanistan in recent years, the overwhelming response by other civil and military leaders since then is to avoid such conflicts, and refocus mostly on conventional target hardening measures. Although this evidence might be driven by decreasing returns to scale (both on target hardening and offence level) our model offers an alternative explanation for it, that does not necessarily rely on decreasing returns.[2]

Second, allowing the target country to decide not to acquire information — i.e., the extensive margin of the information acquisition problem — complements the cheap talk approach featured in earlier contributions, and seems more compelling when there is a substantial imbalance between the bargaining power of the country and the community of noncombatants.

Third, our simple set-up is flexible enough to accommodate additional aspects of the problem that will be discussed in Section 2.4 and that we hope to address in future research.

The paper is organized as follows. In the next section (Section 2.2) we relate our paper to the existing literature. In Section 2.3 we set up the model. In Section 2.3.1, we analyze the first-best benchmark. The analysis under asymmetric information is developed in Section 2.3.2, where we characterize the optimal behavior of the uninformed country and the informed one. In Section 2.3.3, we identify the incentive of the target country to get informed and the relevant comparative statics. Section 2.4 discusses possible extensions of the model, traces a research agenda for future work and concludes.

2.2 Related Literature

There exists a fairly developed literature in economics studying terrorism and counterterrorism, to which our paper relates. The closest contribution is Berman *et al.* (2011), who provide a model that captures a three-way interaction between insurgent organizations, government forces and a local community. As already said, an important difference between our model and Berman *et al.* (2011) is that we take a mechanism design approach, while they analyze a game in which the community decides how much information to disclose in a 'cheap-talk' fashion. Hence, in this sense the two models are complementary. Nevertheless, while we make a formal difference between target hardening and military measures, they model the country's counterterrorism activity as a unidimensional decision variable. In contrast with our results, in their model norms favoring rebel control reduce the community incentive to disclose information.

Many historians and political scientists who study war have recognized that states know that war will entail costs, and even if they expect offsetting benefits they still have an incentive to avoid the costs. The central question, then, is what prevents states from reaching an ex ante agreement that avoids the costs will be paid ex post if they go to war? In his important article

[2] Besides its benchmarking role, which we discuss in Remark 1 below, substitutability may for example emerge when either defense or offense display high (unmodeled) fixed costs that might make hard for a capital constrained country to invest in both instruments.

Fearon (1995) recognizes three answers to this puzzle. First, private information about relative capabilities and incentives to misrepresent such information. Second, commitment problems, i.e. situations in which mutually preferable bargains are unattainable because states might have an incentive to renege. Third, issue indivisibilities. This article focuses on private information and does not consider other mechanisms.

Some relevant contribution to the asymmetries of information strand of the literature are the highly acclaimed study on the causes of war by Blainey (1988) and the two papers by Morrow (1992 and 1994) on the possibility of solving crises by the parties trading concessions on different issues "linkage" and the formalization of cooperation in the face of problems of distribution and information. Closer to this paper is the work by Wang and Bier (2011) who consider a model of incomplete information in which the defender is uncertain about the attacker's weights on which target to attack. They consider a dynamic game with incomplete information in which the defender first chooses how to allocate his defensive resources, and then an attacker chooses which target to attack according to his preferences. The defender's uncertainty about attacker's preferences is modeled by a subjective distribution representing both defender uncertainty about the attacker weights on the various known attributes and also defender ignorance about any unobserved attributes. In our model there is only one target, and the uncertainty is about the potential harm produced by the attack. Yet, while they allow only to defend the target, we also consider the possibility of a preemptive attack (on uncertainty about preferences see also Lapan and Sandler, 1993; Brown et. al., 2006; Bier et al., 2007; Farrow, 2008).

Like us, Powell (2006) considers a model where a target country can engage in preemptive attacks. He adopts a bargaining approach and focus on commitment problems to argue that even if violence occurs in equilibrium, it is inefficient in some cases. In our model, this is not the case when the country acquires information, while it can well be the case with an uninformed country. The main trade-off between acquiring information and acting behind the veil of ignorance hinges precisely on this tension (on preemptive strikes, see also Powell, 1991 and Fearon, 1995). Powell (2006) also proposes a distinction about which types of war can be interestingly accounted for by asymmetries of information versus difficulty to commit. An informational approach may explain the early phases of some conflicts but does not provide a convincing description of long conflicts. The reason being that informational asymmetries are unlikely to persist for a long conflict.

As for target hardening, Hastings and Chan (2013) develop a simple cheap-talk model that highlights the relationship between target hardening and the value that a terrorist group derives from attacking it (but see also Arce and Sandler, 2003; Berman and Laitin, 2005; Enders and Sandler, 1993 and 2004). They compare how the expected value of attacking a hardened target varies depending on whether terrorists just maximize the physical damage inflicted to the target, or if they also attach a symbolic value to the attack (even when it fails). They find that there are some benefits in hardening a target since it decreases the probability of an attack and, almost by definition, raises the loss of the terrorist group. Differently, they also stress that the marginal benefits of hardening the target decreases because of the significant rise in the symbolic value of the target to terrorist group. We do not model this interaction and assume that terrorists always attack in order to produce the most harmful damage to the target. However, we focus on the information acquisition aspect that is neglected by Hastings and Chan (2013).

2.3 The Model

In order to capture in the simplest possible way the basic trade-off between target hardening and offensive measures, as well as the pros and cons of acquiring salient information about terrorists from non-combatants, the model that we will develop is highly stylized. A country is targeted by a foreign terrorist group whose attack produces an harm (damage) $\mathcal{H}(\theta, d) \leqslant 0$, with $\mathcal{H}_\theta(\theta, d) < 0$ and $\mathcal{H}_d(\theta, d) \geqslant 0$. The parameter θ is a random variable, distributed on the compact support $\Theta \triangleq [\underline{\theta}, \overline{\theta}] \subseteq \Re^+$ with cdf $F(\theta)$ and pdf $f(\theta)$, which measures (other things being equal) the severity (incidence) of the harm. As a convention, we assume that the higher θ, the more harmful the attack, which explains why $\mathcal{H}(\cdot)$ is decreasing in θ. The interpretation of θ is that it reflects the military strength (violence) of the terrorist group. In order to shield against the attack, the country can invest public funds in counterterrorism activities $d \geqslant 0$ (target hardening) that mitigate the incidence of the attack, which explains why $\mathcal{H}(\cdot)$ is increasing in d — e.g., introduction of metal detectors, mandatory passenger screening, fortification and protection of government buildings, etc. For simplicity, we assume that the monetary cost of defense is linear, and it simply equals d.

In addition to counterterrorism activities, the country can also engage in a preemptive attack striking the terrorists in the hosting country. We denote by $s \in \{0, 1\}$ the striking decision: $s = 1$ if a strike occurs, $s = 0$ otherwise. Again, we assume that the monetary cost of conducing a military campaign abroad is linear and equal to λ. We posit that λ is larger than 1 to capture the idea that — differently from the cost of defense — the actual cost of the military campaign weights also non-pecuniary aspects such as political dissent stemming from the possible loss of lives during the campaign, and so on.

The country has no *a priori* information about the strength of the terrorist group. Yet, this information can be acquired by dealing with the community of noncombatants hosting the terrorists. These people own private information about θ and are willing to disclose it as long as the target country provides public goods g as a prize for collaboration.

Following the mechanism design approach developed by Myerson (1981) it is convenient to model the choice of a strike as a mixed strategy and denote by $\alpha \triangleq \Pr[s = 1] \in [0, 1]$ the probability

[3] Indeed, this 'convexification' allows us to deal with a smooth maximization problem for the Government that is relatively easier to solve than the problem where $\alpha \in \{0, 1\}$. However, as shown in Proposition 2.3, in equilibrium the Government will always play a pure strategy.

of a strike.[3] The community is risk neutral and has a utility function

$$u\left(g,\alpha,\theta\right)\triangleq g-\alpha x-\theta\left(1-\alpha\right).$$

Following Berman *et al.* (2011), the parameter x can be interpreted as a measure of the weight that the community assigns to its territorial sovereignty or, in other words, as the strength of the norms against foreign control of their territory. Alternatively, x could simply reflect a measure of how rooted the terrorist group is within the community: the higher x, the more rooted in the community the terrorists are. This cost is paid when the strike occurs — i.e., with probability α. By contrast, when the country decides not to strike, the terrorist group exerts retaliation over the community who bears a loss equal, for simplicity, to the strength θ of the terrorist group. Indeed, obvious reputation concerns may induce terrorists to discipline the community when it collaborates with the country.[4]

As we will explain below, acquiring information allows the country to always take the best (interim) decision. The country's payoff is

$$\mathcal{V}\left(\alpha,d,g,\theta\right)\triangleq\left(1-\alpha\right)\mathcal{H}\left(\theta,d\right)-d-\lambda\alpha-g,$$

Hence, the country must first decide whether to acquire information about θ or not, and then it chooses the intensity (probability) of the strike and how much to invest in counterterrorism.

Without loss of generality we assume that, if the country decides to acquire information, it offers by the Revelation Principle (Laffont and Martimort, 2002) a direct mechanism

$$\mathcal{M}\triangleq\left\{g(m),\alpha(m)\right\}_{m\in\Theta},$$

which specifies an amount of public goods $g\left(\cdot\right)$ and a probability of strike $\alpha\left(\cdot\right)$ both contingent on the community report m about the strength of the terrorist group θ. Of course, given $g\left(\cdot\right)$ and $\alpha\left(\cdot\right)$, the country will optimally set the counterterrorism activity $d\left(\cdot\right)$ so as to maximize its (expected) payoff. As we will explain below, since d has no direct impact on the community's payoff, there is no loss of generality in restricting attention to mechanisms where the country does not commit to a defense level *vis-à-vis* the community.[5]

For tractability, and without loss of insights, throughout we assume that

- **(A1)** The harm is quadratic
$$\mathcal{H}\left(\theta,d\right)\triangleq-\frac{1}{2}\left[\theta-d\right]^{2}.$$

A quadratic loss function is typically used in the literature to obtain (tractable) closed form

[4] Kalyvas (2006) argues that rebels and terrorist groups typically engage in violence against noncombatants to discourage their collaboration with target countries.

[5] Considering such an extended mechanism would also seem unrealistic since the country's actual choice of d is hardly verifiable by the community.

solutions. In our context, this specification has a two important implications. First, it allows to obtain results that depend only on the first two moments of the distribution of θ. Second, a quadratic harm also guarantees concavity of the country's objective function and hence uniqueness of the optimal policy. Notice that this quadratic function also captures the idea that Governments want to properly match the target hardening effort to the actual threat imposed by the terrorist — e.g., because providing too much defense might impose unnecessary costs on society. Hence, despite the simple structure, a quadratic harm is coherent with a matching game.[6]

- **(A2)** $F(\theta)$ exhibits an increasing (inverse) hazard rate — i.e., $h(\theta) \triangleq F(\theta)/f(\theta)$ is increasing in θ. Moreover, it features mean μ and variance σ^2, with $\Delta\theta \triangleq \overline{\theta} - \underline{\theta} \geqslant 0$.

  Assuming an increasing (inverse) hazard rate is standard in most screening models.

- **(A3)** The harm is not too strong in expectation — i.e., $\mu < \frac{1}{2} + \lambda$. Moreover, $\sigma^2 \geqslant \sqrt{2\lambda}$ and

$$\underline{\theta} \geqslant \max\left\{\sqrt{2\lambda}, \sqrt{2(x+\lambda)+1} - 1\right\}.$$

Altogether, these parametric restrictions simply guarantee that there exist a non-empty region of the model parameters such that the country strikes regardless of whether it is informed or not. The complementary region of parameters in which there is no strike, or the country prefers neither to strike nor to defend, is uninteresting for our purposes.

The solution concept is Subgame Perfect Nash Equilibrium (SPNE).

Remark 1. As it will be clear soon, the fact that $\mathcal{H}(\cdot)$ only depends on d and that the strike always succeeds — i.e., it eradicates the problem at its root with certainty — *de facto* imply that the utility of the country is such that defense and attack are substitutes rather than complements. But, one could imagine that a preemptive strike increases the terrorists' cost of attacking the target, or that striking displays decreasing returns.[7] In both cases, which are equally plausible, many results of the paper would change since it could be optimal for the country to employ both measures at the same time (especially when the harm is very high). Yet, substitutability seems to provide a useful benchmark to understand the basic trade off between attack and defense, and hence it seems the very first step to make in order to understand the problem. We hope to address the complementarity issue in future research.

2.3.1 The First-Best Benchmark

As a benchmark, consider first the case in which the country knows θ. In that case, there is no need to waste funds in providing public goods to the community since the strength of the terrorist

[6]We thank an anonymous referee for suggesting this interpretation.

[7]E.g., the cost of striking is convex in the intensity of the strike or the probability of defeating the terrorists is concave with respect to the intensity of the strike.

group is common knowledge. Hence, for every θ the optimal d and α solve

$$\max_{\alpha\in[0,1],d\geqslant 0} \mathcal{V}(\alpha,d,\theta) \triangleq \max_{\alpha\in[0,1],d\geqslant 0} (1-\alpha)\,\mathcal{H}(\theta,d) - \alpha\lambda - d,$$

where, abusing slightly notation, we have defined $\mathcal{V}(\alpha,d,\theta) \triangleq \mathcal{V}(\alpha,d,g=0,\theta)$.

Differentiating with respect to α we have

$$\frac{\partial\mathcal{V}(\alpha,d,\theta)}{\partial\alpha} = \underbrace{-\mathcal{H}(\theta,d)}_{\text{Harm avoidance}} - \lambda \lesseqgtr 0. \tag{2.1}$$

Clearly, a higher probability of striking lowers the harm since, by assumption, the attack eradicates the terrorists' threat (harm avoidance). Hence, it is optimal to strike if and only if the benefit that the country obtains from the avoidance of the harm is larger than the cost λ of the military campaign.

Differentiating with respect to d we have

$$\frac{\partial\mathcal{V}(\alpha,d,\theta)}{\partial d} = \underbrace{(1-\alpha)\,\mathcal{H}_d(\theta,d)}_{\text{Target hardening}} - 1 \leqslant 0. \tag{2.2}$$

The optimal defense level trades off two effects: the cost of investing public funds in counterterrorism activities against the reduced incidence of the harm due to these activities (target hardening). The $\leqslant$ sign comes from the fact that $\mathcal{H}(\theta,d)$ is strictly concave in d. Then, the solution is either 0 — i.e., (2.2) holds with a strict inequality — or interior — i.e., (2.2) holds with equality.

Combining (2.1) and (2.2), we can establish the following result.

Proposition 2.1. *When the country knows θ (complete information), the optimal policy is such that $\alpha^{FB}(\theta) = 1$ and $d^{FB}(\theta) = 0$ if and only if*

$$\theta \geqslant \theta^{FB} \triangleq \frac{1}{2} + \lambda.$$

Otherwise, for every $\theta \leqslant \theta^{FB}$, it is $\alpha^{FB}(\theta) = 0$ and $d^{FB}(\theta) = \theta - 1 > 0$.

The solution of the first-best problem is clearly 'bang-bang' since the country's payoff is linear in α. As intuition suggests, when there is no uncertainty about the terrorists' strength, it is optimal to strike only when the harm is sufficiently large, namely when the terrorists are strong enough. Clearly, other things being equal, as the harm becomes more severe — i.e., as θ grows large — the country has more incentive to strike, while it has less incentive to do so as the cost of the military campaign rises — i.e., as λ grows large. Of course, when the strike does not take place, the optimal level of public funds invested in target hardening is increasing with harm.

2.3.2 Asymmetric Information

Suppose now that the country has no information about θ. As explained before, there are two viable options. The country can either base its strategy on the prior information it owns about θ, or it can acquire information from the community and make a decision that is ex post optimal. Yet, in order to ensure participation of the community to the deal and elicit a truthful report, the country has to give up an information rent to the community, which might either not accept the contract or misreport θ. In what follows we characterize the optimal counterterrorism policy for each information acquisition choice and then compare the two outcomes.

Uninformed country. If the country decides to be uninformed, its optimization problem is similar to that solved in the complete information benchmark with the difference that the harm has to be taken in expected value. Hence, the optimal policy under no information acquisition solves the following maximization problem:

$$\max_{\alpha\in[0,1],d\geqslant 0} \mathcal{V}(\alpha,d) \triangleq \max_{\alpha\in[0,1],d\geqslant 0} \int_\theta \mathcal{V}(\alpha,d,\theta)\,dF(\theta)\,.$$

Differentiating with respect to α we have again a trade off between the harm avoidance effect (in expectation though) and the cost of striking — i.e.,

$$\frac{\partial\mathcal{V}(\alpha,d)}{\partial\alpha} = -\int_\theta \mathcal{H}(\theta,d)\,dF(\theta) - \lambda \lesseqgtr 0.$$

Differentiating with respect to d we have a trade off between the target hardening effect (again in expectation) and the cost of defense — i.e.,

$$\frac{\partial\mathcal{V}(\alpha,d)}{\partial d} = (1-\alpha)\int_\theta \mathcal{H}_d(\theta,d)\,dF(\theta) - 1 \leqslant 0.$$

Hence, it is not difficult to guess that the solution of the uninformed country's maximization problem has a structure that is similar to that of the first-best. The difference being that, since the terrorists' strength is unknown, the costs and benefits associated with the use of each policy instrument must be taken in expected terms. Because the country is risk averse — i.e., the harm $\mathcal{H}(\cdot)$ is quadratic — this means that the variance of θ now plays an important role in the analysis.

Proposition 2.2. *When the country is uninformed about θ, the optimal counterterrorism policy is such that $\alpha^N = 1$ and $d^N = 0$ if and only if*

$$\sigma^2 \geqslant \sigma_0^2 \triangleq 1 - 2(\mu-\lambda)\,.$$

and $\alpha^N = 0$ and $d^N = \mu - 1$ otherwise.

The behavior of an uninformed country is rather simple, and is shaped by the following trade-off. On the one hand, investing in defense only entails some risk for the country because defense is costly and it does not completely neutralize the harm. To see why, notice that the expected harm evaluated at d^N is

$$\begin{aligned}\int_\theta \mathcal{H}\left(\theta, d^N\right) dF(\theta) &\triangleq -\int_\theta \frac{[\theta-(\mu-1)]^2}{2} dF(\theta) \\ &= -\int_\theta \frac{[\theta-\mu+1]^2}{2} dF(\theta) \\ &= -\frac{\sigma^2+1}{2},\end{aligned}$$

which is decreasing in σ^2 since the country is risk averse (i.e., the harm is quadratic in θ). On the other hand, a strike completely neutralizes the harm, but the military campaign is costly. Hence, the higher σ^2 relative to λ, the more appealing the strike.

Finally, we can compute the expected payoff of the uninformed country:

$$\mathcal{V}^N \triangleq \begin{cases} \frac{1}{2}-\mu-\frac{\sigma^2}{2} & \Leftrightarrow \sigma^2 \leqslant \sigma_0^2 \\ -\lambda & \Leftrightarrow \sigma^2 \geqslant \sigma_0^2 \end{cases},$$

which is weakly decreasing in λ, μ and σ^2.

As explained before, the uninformed country strikes, and completely eradicates the threat, if and only if there is enough uncertainty about the terrorists' strength ($\sigma^2 \geqslant \sigma_0^2$) — e.g., because the group of terrorists is new and very little is known about them. In this case, the country's expected payoff is simply equal to the expected cost of the military campaign. Differently, when the uncertainty about the strength of the terrorist group is low ($\sigma^2 \leqslant \sigma_0^2$), the country prefers to take the risk of being attacked and invests only in defense to reduce the incidence of the attack. Hence, the expected payoff falls with the average strength of the terrorists μ and with its variance, which measures the risk to which the country is exposed when there is no strike.

Informed country. Suppose now that the country acquires information from the community before deciding whether to strike or not. The information released by the community is truthful if and only if the mechanism $\mathcal{M}$ is incentive compatible. Let

$$u(\theta, m) \triangleq g(m) - x\alpha(m) - \theta(1-\alpha(m)),$$

denote the community's expected payoff when it reports m and the true state of nature is θ. And, abusing slightly notation, denote by

$$u(\theta) \triangleq u(\theta, m=\theta)$$

the community's information rent when it truthfully reveals the terrorists' strength. An incentive compatible policy requires the following standard implementability conditions

$$\left.\frac{\partial u(\theta, m)}{\partial m}\right|_{m=\theta} = 0 \quad \Longrightarrow \quad u(\theta) = u(\overline{\theta}) + \int_{\theta}^{\overline{\theta}} (1 - \alpha(z))\, dz, \tag{2.3}$$

and

$$\left.\frac{\partial^2 u(\theta, m)}{\partial m^2}\right|_{m=\theta} \leqslant 0 \quad \Longrightarrow \quad \dot{\alpha}(\theta) \geqslant 0. \tag{2.4}$$

Condition (2.3) simply reflects the local incentive compatibility constraint: it guarantees that the community has no incentive to misreport locally the strength of the terrorist group.[8] Essentially, this condition delivers the information rent that the country has to give up to the community in order to elicit truthful information. Notice that the community's rent is decreasing with the strength of the terrorist group — i.e., $\dot{u}(\theta) \leqslant 0$ — which means that communities hosting weaker terrorist groups are able to extract higher rents when collaborating with target countries. The intuition is that stronger terrorists retaliate more harshly. Hence, the community has an incentive to overstate the risk of retaliation in order to be offered a better deal (i.e., more public good provision). In addition, it is useful to observe that the rent enjoyed by the community is decreasing in the probability of a strike. Indeed, the larger is $\alpha(\cdot)$, the more protected the community feels against the risk of retaliation, and the less costly it is for the country to elicit truthful information.

Condition (2.4), instead, reflects the so called 'monotonicity' constraint which states that any incentive feasible mechanism must be such that the stronger is the terrorist group, the more likely it is that a strike will occur. This requirement is equivalent to the standard Spence-Mirleess single-crossing condition.

The country's maximization problem is

$$\max_{\alpha(\cdot)\in[0,1], g(\cdot)\geqslant 0, d(\cdot)\geqslant 0} \int_{\theta} \mathcal{V}(\alpha(\theta), d(\theta), g(\theta), \theta)\, dF(\theta),$$

subject to (2.3), (2.4) and

$$u(\theta) \geqslant 0 \quad \forall \theta \in \Theta. \tag{2.5}$$

Neglecting the monotonicity condition (2.4) and using a standard change of variable — i.e., optimizing with respect to $u(\cdot)$ instead of $g(\cdot)$ — simple integration by parts of the rent in (2.3) (see, Laffont and Martimort, 2002) allows to rewrite the country's maximization problem as

$$\max_{\alpha(\cdot)\in[0,1], d(\cdot)\geqslant 0} \int_{\theta} \{\mathcal{V}(\alpha(\theta), d(\theta), \theta) - \alpha(\theta) x - (1 - \alpha(\theta)) [\theta + h(\theta)]\}\, dF(\theta). \tag{2.6}$$

[8] It can be easily shown that incentive compatibility is satisfied also globally if it holds locally (see, e.g., Laffont and Martimort, 2002).

Notice that, as usual in these models, we have optimally set $u(\overline{\theta}) = 0$ because the community hosting the most violent terrorists can only under-report its retaliation loss, but this is clearly not optimal. As a result, the $\overline{\theta}$-type is left with no rent.

Again, the objective function of the country is linear in $\alpha(\cdot)$ and its structure reflects the standard trade off between efficiency and rent-extraction. Essentially, in order to elicit truthful information, the country must give up an information rent to the community, which will in turn affect the optimal striking rule. Yet, relative to the first-best benchmark, dealing with the community is costly and the structure of this cost depends on the probability of a strike. Differentiating with respect to $\alpha(\cdot)$ we have

$$\underbrace{-\mathcal{H}(\theta, d(\theta)) - \lambda}_{\text{First-best rule}} - \underbrace{x}_{\text{Norm breaking}} + \underbrace{\theta + h(\theta)}_{\text{Rent saving}} \lesseqgtr 0. \tag{2.7}$$

The new trade off that the country faces when it acquires information is essentially between the cost of breaching the community's norms and the information rent that it has to pay when deciding to invest in target hardening only. First, since the community is averse to hosting a foreign army, the level of public good that has to be provided by the country when it strikes must compensate the community for the loss due to the breach of its norms. Second, when the country does not strike, retaliation by the terrorist group occurs. Hence, the public good that has to be provided in this case exceeds the participation level θ, because it also takes into account the rent that is necessary to pay in order to induce the community to report truthfully the terrorists' strength. Other things being equal, the higher θ the more willing the country is to strike because a relatively higher probability of striking in state θ mitigates the incentive to mimic of the community with type below θ, as reflected by the increasing inverse hazard rate $h(\theta)$ that measures the mass of types below θ.

Differentiating with respect to $d(\cdot)$ it is easy to show that the same condition as in the first-best obtains. This echoes the dichotomy result in Laffont and Tirole (1993, Chapter 3). Essentially, since counterterrorism does not affect the community's payoff, only the striking decision is used as a screening device, while the defense level is set at its first-best rule. Notice that this same outcome would occur if the community committed to an extended mechanism that includes also the defense level as a screening instrument, which explains why it is without loss of generality to exclude d from the mechanism $\mathcal{M}$.

We can thus establish the following result:

Proposition 2.3. *Define*

$$0 < \underline{x} \triangleq 2\underline{\theta} - \frac{1}{2} - \lambda < \overline{x} \triangleq \underline{x} + 2\Delta\theta + h\left(\overline{\theta}\right).$$

When the country acquires information, the optimal policy is such that:

- *If* $x \leqslant \underline{x}$ *then* $\alpha^I(\theta) = 1$, $d^I(\theta) = 0$ *and* $g^I(\theta) = x$ *for every* $\theta \in \Theta$.
- *If* $x \in (\underline{x}, \overline{x})$ *then:* $\alpha^I(\theta) = 1$, $d^I(\theta) = 0$ *and* $g^I(\theta) = x$ *when* $\theta \geqslant \theta^*$. *Otherwise,* $\alpha^I(\theta) = 0$, $d^I(\theta) = \theta - 1$ *and*

$$g^I(\theta) = \theta + \int_{\theta}^{\theta^*} \left(1 - \alpha^I(\theta)\right) dz = \theta^*. \tag{2.8}$$

The threshold $\theta^* \in \left(\underline{\theta}, \overline{\theta}\right)$ *is the unique solution of*

$$2\theta + h(\theta) = \frac{1}{2} + x + \lambda,$$

and it is increasing in x *and* λ.

- *If* $x \geqslant \overline{x}$, *then* $\alpha^I(\theta) = 0$ *and* $d^I(\theta) = \theta - 1$ *for every* $\theta \in \Theta$.

The intuition behind this result is rather simple. As in the first-best, linearity of the country's payoff implies a bang-bang solution also when getting informed is costly. Yet, with asymmetric information, the optimal strategy depends not only on the terrorists' actual strength, but also on the extent of the community norms. More precisely, depending on the magnitude of x either a pooling or a separating outcome may occur. Notice that, in both cases, the monotonicity condition (2.4) is satisfied.

First, when the community's internal norms are sufficiently weak ($x \leqslant \underline{x}$) — e.g., because its members are not too resilient to host a foreign army or because they have weak ties with the terrorists — the optimal policy requires a pooling outcome in which the country strikes no matter how strong the terrorist group is. As intuition suggests, the threshold $\underline{x}$ is decreasing in λ and increasing in $\underline{\theta}$. Intuitively, as λ grows large it becomes more costly to strike and therefore a pooling outcome such that the country always strikes regardless of θ becomes less appealing. Differently, as the harm becomes more severe — i.e., as the support of θ shifts to the right (higher values of $\underline{\theta}$) — the information rent enjoyed by the community when a strike does not occur increases, and this makes the country (other things being equal) less willing to take the risk of investing only in target hardening.

Second, when the community norms are not so weak to induce the country to strike no matter what, but neither so large to make a strike not worth at all — i.e., $x \in (\underline{x}, \overline{x})$ — the optimal policy is such that the country intervenes militarily if and only if the terrorists are sufficiently strong. That is, a strike occurs when θ exceeds the threshold θ^* that is determined endogenously to trade off the cost that the country has to pay in order to compensate the community members for violating their norms and the information rents, induced by the risk of retaliation, that it would have to give up when the information received by the community is used to fine tune defense only. As intuition suggests, the threshold θ^* is increasing in x and λ since is relatively less profitable to strike if the community features stronger internal norms and/or if the military cost of the campaign is higher.

Third, when the community norms are very strong — i.e., $x \geqslant \overline{x}$ — striking becomes excessively costly, and the optimal policy is such that the country prefers a strategy based on target hardening only. The threshold $\overline{x}$ is increasing with the width of the terrorists' types $\Delta\theta$ since a wider support of types is associated with larger rents, and it is increasing with the (inverse) hazard rate function which is also a direct measure of rents as explained above.

Interestingly, the level of public goods provided by the country is increasing in the norms' severity x and (weakly) increasing in the cost of the strike λ. That is, the more the terrorists are rooted in the local communities that host them, the larger is the public goods provision requested in order to reveal truthful information. Similarly, the investment in public goods is higher when the target country faces a higher military cost to implement a successful strike because, in this case, it is more likely that it will invest only in target hardening leaving the community exposed to retaliation by the terrorists.

Finally, by inspecting (2.7) we can compare the first-best with the behavior of an uninformed country which decides to acquire information. Letting θ_0 be the value of the terrorists' strength at which the solution of the second- and first-best problems are equivalent, that is, the solution of

$$x = \theta + h(\theta),$$

it can be immediately shown that

Corollary 2.4. *For any $\theta > \theta_0$ (resp. $<$) in the second-best striking becomes more (resp. less) appealing than in the first-best.*

Hence, when the country acquires information, weak terrorist groups are attacked relatively less often than in the first-best, while the country behaves more aggressively *vis-à-vis* stronger and more violent groups. Obviously, although these distortions will not affect the defense level on the intensive margin, they affect the extensive margin — i.e., the extent to which each instrument is used.

Summing up, we can compute the country's expected payoff from acquiring information. The result stated below will turn quite useful when we will provide the interpretation for the comparison between the country's expected payoff with and without information acquisition.

Proposition 2.5. *When the country acquires information, its expected payoff is*

$$\mathcal{V}^I \triangleq \underbrace{-\lambda}_{\text{Striking cost}} + \underbrace{\int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta + F(\theta^*)\, h(\theta^*) - x}_{\text{Rent saving}},$$

which is increasing in θ^*.

Hence, the country's expected payoff from acquiring information can be written as the cost of striking, plus the expected rent-saving effect that a strike brings about. That is, the community's expected information rent

$$\int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta + F(\theta^*)\, h(\theta^*),$$

minus the cost x of breaching its internal norms. Of course, the larger is the set of states in which it is optimal to invest in target hardening only, the stronger is the rent saving effect, whereby increasing the benefit that the country obtains when it acquires information.

2.3.3 Optimal Information Acquisition Rule

Now we turn to analyze the information acquisition decision of the country. The comparison between V^I and V^N yields (see the Appendix):

$$\mathcal{V}^I - \mathcal{V}^N \triangleq \begin{cases} \int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta + F(\theta^*)\, h(\theta^*) - x & \Leftrightarrow \sigma^2 \geqslant \sigma_0^2 \\ \int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta + F(\theta^*)\, h(\theta^*) - (x+\lambda) - \frac{1}{2} + \mu + \frac{\sigma^2}{2} & \Leftrightarrow \sigma^2 \leqslant \sigma_0^2 \end{cases}.$$

Hence, we can establish the following result:

Proposition 2.6. *The optimal information acquisition rule has the following features:*

- *For $\sigma^2 \geqslant \sigma_0^2$ acquiring information is optimal — i.e., $\mathcal{V}^I \geqslant \mathcal{V}^N$ — only if $\overline{\theta} \leqslant \sigma_0^2/2$, and if $x \geqslant x^*$, with $x^* \in (\underline{x}, \overline{x})$ being the unique solution in x of*

$$\int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta + F(\theta^*)\, h(\theta^*) - x = 0.$$

  Otherwise, it is never optimal to acquire information.

- *For $\sigma^2 \leqslant \sigma_0^2$ acquiring information is optimal — i.e., $\mathcal{V}^I \geqslant \mathcal{V}^N$ — only if $\overline{\theta} \leqslant \sigma_0^2/2$ and $x \geqslant x^*$, and if*

$$\sigma^2 \geqslant \sigma_0^2 - 2\left[\int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta + F(\theta^*)\, h(\theta^*) - x\right].$$

  Otherwise, it is never optimal to acquire information.

The country's decision between acquiring information and remaining uninformed is shaped by the following trade-off. On the one hand, by not acquiring information the country bears the risk of making a decision that is inefficient ex-post. On the other hand, when the country decides to remain uninformed, whatever decision is taken, it saves on the cost of acquiring information, which as argued above depends on the strength of the community norms and the rent that the community has to be granted in order to truthfully report its private information.

Surprisingly, when the community's norms are sufficiently weak the country prefers not to acquire information. To see why recall that θ^* is increasing in x. Hence, for low values of x the country strikes relatively more often and, when it does, its ex post payoff is $-(x+\lambda)$. Therefore, conditional on striking, the country is always worse off when it acquires information — i.e., $-\lambda > -(\lambda + x)$ — which is the case also in expectation when x is low enough to induce θ^* being close to $\underline{\theta}$. This suggests that acquiring information should be optimal if the rent saving effect discussed above is strong enough — i.e., if $x \geqslant x^*$. Yet, x sufficiently large is not the only condition that is needed in order for the country to be willing to acquire information: $\overline{\theta}$ must also be not too large. The reason is simple: strong norms (high x) imply that the set of contingencies in which the informed country invests in target hardening only is large, and in these contingencies a rent is paid to the community. Now, if $\overline{\theta}$ is large, these rents are large too (see, e.g., equation 2.3), hence acquiring information becomes costly. In the region of parameters where the uninformed country does not strike, the condition for the country to acquire information becomes even stronger: in addition to a sufficiently high x and a not too large $\overline{\theta}$ it must also be σ^2 not too small. The reason is simple: if there is little uncertainty (σ^2 small) about the terrorists' strength, then acquiring information becomes too costly relative to being uninformed because in the latter case the risk to which the country is exposed when it invests in target hardening only is small.

Turning to the comparative statics, we can show the following result.

Proposition 2.7. *x^* is decreasing in λ.*

Intuitively, as λ grows large, there is more incentive to acquire information (other things being equal) because by not doing so, the country is more exposed to the risk of making the wrong choice.

Summing up, the model delivers an inverse relationship between the incentive to acquire information about the strength of terrorist groups and the norms featured by the communities hosting these terrorists. This result contrasts with the equilibrium characterization obtained by Berman *et al.* (2011) who find that norms favoring rebel control reduce incentives to disclose information. Their model features an equilibrium in which communities disclose information if and only if norms are sufficiently weak. By contrast, in our model countries have an incentive to acquire this information if and only if norms are sufficiently strong. As explained before, what makes the difference is the rent saving effect and the fact that in our model the target country can use two alternative instruments in order to shield against the terrorists' violence. In addition: (*i*) conditional on striking, the country would like to be uninformed because by so doing it would save the cost it has to pay in order to compensate the non-combatants for breaking their norms; (*ii*) as x becomes smaller, θ^* becomes smaller too, so that when the country acquires information it is relatively more likely to strike, which is however an ex-post inefficient strategy. Hence, the country has a weaker incentive to acquire information when the community features weaker norms.

2.4 Concluding Remarks

In this paper we have developed a mechanism design approach to study the relationship between terrorism, counterterrorism and information acquisition. By so doing, we have highlighted a novel tension between target hardening and preemptive attacks. Specifically, we showed that a country targeted by a group of terrorists engages in preemptive attacks, which eradicate the threat at its root, only when it faces sufficiently strong terrorists and when the community of noncombatants from which information about terrorists is elicited features weak norms, or if these people have poor connections with the terrorists. Yet, in contrast with the existing literature, in our model it is optimal for the target country to acquire information about the strength of the terrorists only when noncombatants feature strong enough internal norms and when there is enough uncertainty about the terrorists' strength. This suggests that using public goods provision as a way to attract communities of noncombatants in order to elicit from them information about terrorists is not always ex ante efficient.

Although the model is highly stylized, we believe it captures some basic aspects of reality and that, in addition, it provides a solid basis for a broader research agenda. First, the distinction between defensive and offensive strategies and the fact that each of these instruments is used depending on the severity of the terrorists' potential damage, seem to be natural aspects of the real life problem. Second, allowing the target country to gather information about the terrorists' strength — i.e., the extensive margin on the information acquisition decision — seems in some cases more realistic than the cheap-talk approach featured in earlier contributions, especially when the community hosting terrorists is relatively smaller than the target country. Finally, the model also provides a flexible framework to address — in future research — related issues that have not been addressed here. Namely, how results would change with complementarty between the two instruments, the decision making problem of a country targeted by more than one terrorist group, the introduction of violence as one of the terrorists' decision variables, the design of a coalition against terror and the related common agency problem, the effects of dynamics, etc.

2.5 Appendix

Proof of Proposition 2.1. Differentiating the objective function $\mathcal{V}(\cdot)$ with respect to α we have

$$\frac{1}{2}[\theta - d]^2 - \lambda. \tag{B1}$$

Differentiating with respect to d we have

$$(1-\alpha)[\theta - d] - 1. \tag{B2}$$

First, notice that if $\alpha = 1$, then (B2) is negative, so that $d = 0$. Hence, $\alpha = 1$ is admissible if and only if

$$\frac{\theta^2}{2} - \lambda \geqslant 0 \quad \forall \theta, \tag{B3}$$

which is always true as long as $\underline{\theta} \geqslant \sqrt{2\lambda}$ as imposed in assumption **A3.**

Next, we show that if $d > 0$, then $\alpha = 0$. Suppose, by contradiction, that this is not the case. For given $\alpha < 1$, the objective function is strictly concave in d. Hence, if $d > 0$ it must be equal to

$$\widehat{d}(\theta) \triangleq \theta - \frac{1}{1-\alpha}.$$

Evaluating the objective function at $\widehat{d}(\theta)$ we have

$$\mathcal{V}(\alpha, \widehat{d}(\theta), \theta) = \frac{1}{2(1-\alpha)} - \theta - \lambda\alpha. \tag{B4}$$

Moreover, notice that if $d > 0$ then $\alpha > 0$ if and only if (B1) is positive. Therefore,

$$\alpha \geqslant 1 - \frac{1}{\sqrt{2\lambda}}. \tag{B5}$$

Maximizing (B4) with respect to α subject to (B5) we have

$$\frac{1}{2(1-\alpha)^2} - \lambda + \eta = 0, \tag{B6}$$

where $\eta \geqslant 0$ is the Lagrangian multiplier associated to (B5). If $\eta > 0$ then (B5) binds and (B6) implies $\eta = 0$: a contradiction. Moreover, if $\eta = 0$ and $\alpha > 0$, the first order condition (B6) yields another contradiction since $\alpha = 1 - \frac{1}{\sqrt{2\lambda}}$. Hence, $\alpha = 0$ whenever $d > 0$.

Finally, to show the result we need to sign the difference

$$\mathcal{V}(\alpha = 1, d = 0, \theta) - \mathcal{V}(\alpha = 0, d = \theta - 1, \theta) = \theta - \lambda - \frac{1}{2},$$

which yields immediately the result. ■

Proof of Proposition 2.2. Differentiating the objective function $\mathcal{V}(\alpha, d)$ with respect to α we have

$$\frac{1}{2}\int_{\theta}(\theta-d)^2\,dF(\theta)-\lambda, \tag{B7}$$

while differentiating with respect to d we have

$$(1-\alpha)\int_{\theta}(\theta-d)\,dF(\theta)-1. \tag{B8}$$

First, notice that if $\alpha=1$, then (B8) is negative, so that $d=0$. Hence, $\alpha=1$ is admissible if and only if

$$\frac{\sigma^2+\mu^2}{2}-\lambda\geqslant 0, \tag{B9}$$

which is true as long as $\sigma^2\geqslant 2\lambda$ as imposed in assumption **A3.**

In what follows, we show that if $d>0$, then $\alpha=0$. Suppose, by contradiction, that this is not the case. Then for given $\alpha<1$, the objective function is strictly concave in d. Hence, if $d>0$ it must be equal to

$$\widehat{d}\triangleq\mu-\frac{1}{1-\alpha}.$$

Substituting for $\widehat{d}$ into the objective function, we have

$$\mathcal{V}(\alpha,\widehat{d})=\frac{1}{2(1-\alpha)}-\lambda\alpha-\mu. \tag{B10}$$

Next, notice that if $d>0$, then $\alpha>0$ if and only if (A4) is positive. Hence, if the following holds

$$\alpha\geqslant 1-\frac{1}{\sqrt{2\lambda}}. \tag{B11}$$

Now, maximizing (B10) with respect to α subject to (B11), one gets

$$\frac{1}{2(1-\alpha)^2}-\lambda+\eta=0, \tag{B12}$$

where $\eta\geqslant 0$ is the Lagrangian multiplier associated to (B11). If $\eta>0$ then (B11) binds and (B12) implies that $\eta=0$: a contradiction. Moreover, if $\eta=0$ and $\alpha>0$, then the condition (B12) implies that $\alpha=1-\frac{1}{\sqrt{2\lambda}}$. Hence, $\alpha=0$ whenever $d>0$.

Finally, to show the result we just need to sign the following difference

$$\mathcal{V}(\alpha=1,d=0)-\mathcal{V}(\alpha=0,d=\mu-1)=\mu+\frac{\sigma^2}{2}-\lambda-\frac{1}{2},$$

which yields immediately the result. In fact,

$$\mathcal{V}(\alpha=1,d=0)\geqslant\mathcal{V}(\alpha=0,d=\mu-1)\quad\Leftrightarrow\quad\sigma^2\geqslant\sigma_0^2\triangleq 1-2(\mu-\lambda),$$

where it is immediate to verify that $\sigma_0^2>0$ under Assumption **A3**. ■

Proof of Proposition 2.3. Optimizing pointwisely the objective function in (2.6) with respect to $\alpha(\cdot)$, we obtain

$$\frac{[\theta - d(\theta)]^2}{2} - (x+\lambda) + \theta + h(\theta), \tag{B13}$$

while optimizing with respect to $d(\cdot)$, we have

$$(1-\alpha(\theta))(\theta - d(\theta)) - 1. \tag{B14}$$

Again, if $\alpha(\theta) = 1$, then (B14) is negative, so that $d(\theta) = 0$. Hence, $\alpha(\theta) = 1$ is admissible if and only if

$$\frac{\theta^2}{2} - (x+\lambda) + \theta + h(\theta) \geqslant 0, \ \theta \in [\underline{\theta}, \overline{\theta}], \tag{B15}$$

which is true as long as $\underline{\theta} \geqslant \sqrt{2(x+\lambda)+1} - 1$ as implied by assumption **A3**.

In what follows, we show that if $d(\theta) > 0$, then $\alpha(\theta) = 0$. Suppose, by contradiction, that this is not the case. Notice that for given $\alpha(\theta) < 1$, the objective function is strictly concave in $d(\theta)$. Hence, if $d(\theta) > 0$, strict concavity of the objective function implies that it must be equal to

$$\widehat{d}(\theta) \triangleq \theta - \frac{1}{1-\alpha(\theta)}.$$

Substituting for $\widehat{d}(\theta)$ into the objective function, we have

$$\int_\theta \left\{ \mathcal{V}(\alpha(\theta), \widehat{d}(\theta), \theta) - \alpha(\theta) x - (1-\alpha(\theta)) [\theta + h(\theta)] \right\} dF(\theta). \tag{B16}$$

Next, we notice that if $d(\theta) > 0$, then $\alpha(\theta) > 0$ if and only if (B13) is positive — i.e., if

$$\alpha(\theta) \leqslant 1 - \frac{1}{\sqrt{2(x+\lambda-(\theta+h(\theta)))}}, \quad \forall \theta \in \Theta. \tag{B17}$$

Maximizing pointwisely (B16) with respect to $\alpha(\theta)$ subject to (B17), we have

$$\frac{1}{2(1-\alpha(\theta))^2} - \lambda - x + \theta + h(\theta) - \eta(\theta) = 0. \tag{B18}$$

where $\eta(\theta) \geqslant 0$ is the Lagrangian multiplier associated to (B17). If $\eta(\theta) > 0$ then (B17) is binding and (B18) yields $\eta(\theta) = 0$: a contradiction. Moreover, if $\eta(\theta) = 0$ for some θ and $\alpha(\theta) > 0$, then the condition (B12) implies that

$$\alpha(\theta) = 1 - \frac{1}{\sqrt{2(x+\lambda-(\theta+h(\theta)))}},$$

which is again a contradiction. Hence, $\alpha(\theta) = 0$ whenever $d(\theta) > 0$.

Finally, to show the result we need to sign the following difference

$$\mathcal{V}(\alpha(\theta) = 1, d(\theta) = 0, \theta) - \mathcal{V}(\alpha(\theta) = 0, d(\theta) = \theta - 1, \theta) = -(x+\lambda) - \left[\frac{1}{2} - (2\theta + h(\theta))\right],$$

which yields $\alpha(\theta) = 1$ and $d(\theta) = 0$ if and only if

$$2\theta + h(\theta) \geqslant \frac{1}{2} + x + \lambda. \tag{B19}$$

Notice that as admits a solution $\theta^* \in \left(\underline{\theta}, \overline{\theta}\right)$ if and only if

$$2\underline{\theta} < \frac{1}{2} + x + \lambda \quad \Rightarrow \quad x \geqslant \underline{x} \triangleq 2\underline{\theta} - \frac{1}{2} - \lambda,$$

which is always positive by Assumption **A3**; and

$$2\overline{\theta} + h\left(\overline{\theta}\right) \geqslant \frac{1}{2} + x + \lambda \quad \Rightarrow \quad x \geqslant \overline{x} \triangleq \underline{x} + 2\Delta\theta + h\left(\overline{\theta}\right) > 0.$$

Hence, the optimal policy is such that: (i) $\alpha(\theta) = 1$ and $d(\theta) = 0$ for every $\theta \in \Theta$ when $x \leqslant \underline{x}$; (ii) $\alpha(\theta) = 1$ and $d(\theta) = 0$ if and only if $\theta \geqslant \theta^* \in int\Theta$, while $\alpha(\theta) = 0$ and $d(\theta) = \theta - 1$ otherwise, for $x \in (\underline{x}, \overline{x})$; (iii) $\alpha(\theta) = 0$ and $d(\theta) = \theta - 1$ for every $\theta \in \Theta$ when $x \geqslant \overline{x}$. ■

Proof of Corollary 2.4. The proof of this result is straightforward and will be omitted.

Proof of Proposition 2.5. To begin with notice that, with information acquisition, the country's expected payoff writes as

$$\mathcal{V}^I \triangleq -\int_{\underline{\theta}}^{\theta^*} \left\{\theta - \frac{1}{2} + \theta^*\right\} dF(\theta) - \int_{\theta^*}^{\overline{\theta}} [\lambda + x]\, dF(\theta).$$

Simple integration by parts imply

$$\int_{\underline{\theta}}^{\theta^*} \theta dF(\theta) = \theta^* F(\theta^*) - \int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta. \tag{B20}$$

Hence, using (B20) together with the definition of θ^*, simple algebraic manipulations imply

$$\mathcal{V}^I = \int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta + F(\theta^*)\, h(\theta^*) - (\lambda + x).$$

Differentiating with respect to θ^* we have

$$\frac{\partial \mathcal{V}^I}{\partial \theta^*} = 2F(\theta^*) + F(\theta^*)\, \dot{h}(\theta^*),$$

which immediately proves the result since $\dot{h}(\theta^*) > 0$ by assumption **A2**. ■

Proof of Proposition 2.6. Suppose first that $\sigma^2 \geqslant \sigma_0^2$, so that the uninformed country always chooses to strike. Hence,

$$\mathcal{V}^I - \mathcal{V}^N = \int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta + F(\theta^*)\, h(\theta^*) - x.$$

Notice that $\mathcal{V}^I - \mathcal{V}^N = -\underline{x}$ when $x \leqslant \underline{x}$ so that $\theta^* = \underline{\theta}$. Instead,

$$\left.\mathcal{V}^I\right|_{\theta^*=\overline{\theta}} - \mathcal{V}^N = \int_{\underline{\theta}}^{\overline{\theta}} F(\theta)\, d\theta + h\left(\overline{\theta}\right) - \overline{x}.$$

for $x \geqslant \overline{x}$ since $\theta^* = \overline{\theta}$. Integrating by parts, this yields

$$\left.\mathcal{V}^I\right|_{\theta^*=\overline{\theta}} - \mathcal{V}^N = \overline{\theta} F\left(\overline{\theta}\right) - \int_{\underline{\theta}}^{\overline{\theta}} \theta dF(\theta) + h\left(\overline{\theta}\right) - \overline{x} = -\overline{\theta} - \mu + \frac{1}{2} + \lambda,$$

which is positive if and only if $2\overline{\theta} \leqslant \sigma_0^2$. In addition, notice that $\left.\mathcal{V}^I\right|_{\theta^*=\overline{\theta}} - \mathcal{V}^N$ is increasing in θ^*. Therefore, if $2\overline{\theta} \leqslant \sigma_0^2$, then there exists a unique x^* that solves

$$\int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta + F(\theta^*)\, h(\theta^*) - x = 0,$$

and is such that $\mathcal{V}^I \geqslant \mathcal{V}^N$ as long as $x \geqslant x^*$. Clearly, if $2\overline{\theta} > \sigma_0^2$ then $\left.\mathcal{V}^I\right|_{\theta^*=\overline{\theta}} < \mathcal{V}^N$ regardless of x.

Next, suppose that $\sigma^2 < \sigma_0^2$. In this case,

$$\mathcal{V}^I - \mathcal{V}^N = \int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta + F(\theta^*)\, h(\theta^*) - (x + \lambda) - \frac{1}{2} + \mu + \frac{\sigma^2}{2},$$

rearranging we have

$$\mathcal{V}^I \geqslant \mathcal{V}^N \quad \Leftrightarrow \quad \sigma^2 \geqslant \sigma_0^2 - 2\left[\int_{\underline{\theta}}^{\theta^*} F(\theta)\, d\theta + F(\theta^*)\, h(\theta^*) - x\right]. \tag{B21}$$

Hence, for $\sigma^2 < \sigma_0^2$, we have $\mathcal{V}^I \geqslant \mathcal{V}^N$ only if $2\overline{\theta} < \sigma_0^2$ and $x \geqslant x^*$, and if (B21) holds; $\mathcal{V}^I < \mathcal{V}^N$ otherwise. ■

References

[1] ARCE, D. G.,SANDLER, T. 2003, *Terrorism and Game Theory,* Simulation and Gaming, **34**(3), 319-337.

[2] BERMAN, E., LAITIN, D. 2005, *Hard Targets: Theory and Evidence on Suicide Attacks*, La Jolla, CA, University of California, San Diego.

[3] BERMAN E., SHAPIRO J.N., FELTER J.H. 2011, *Can Hearts and Minds Be Bought? The Economics of Counterinsurgency in Iraq*, Journal of Political Economy, **119**(4), 766-819.

[4] BIER V.M., OLIVEROS S., SAMUELSON L. 2007, *Choosing what tobprotect: Strategic defensive allocation against an unknown attacker*, J. Public Econom. Theory, **9**(4) 563–587.

[5] BLAINEY, G. 1988, *Causes of War*, Simon and Schuster.

[6] BROWN G., MATTHEW C., SALMERON J., WOOD K. 2006, *Defending critical infrastructure*, Interfaces, **36**(6) 530–544.

[7] CASSIDY, R. 2008, *Counterinsurgency and the Global War on Terror*, Westport, CT: Praeger.

[8] ENDERS, W., SANDLER, T., 1993, *The effectiveness of anti-terrorism policies: Vector-autoregression-intervention analysis*, American Political Science Review, **87**(4), 829-844.

[9] ENDERS W., SANDLER T. 1995, *Terrorism: Theory and Applications*, Handbook of Defense Economics, **1**, 213-249.

[10] ENDERS, W., SANDLER, T. 2004, *What do we know about the substitution effect in transnational terrorism?,* In Silke, A. (Ed.) Research on terrorism: Trends, achievenments and failures. London, Frank Cass.

[11] FARROW, S. 2008, *The economics of homeland security expenditures: Foundational expected cost-effectiveness approaches*, Contemporary Econom. Policy, **25**(1), 14–26.

[12] FEARON, J.D. 1995, *Rationalist Explanations for War*, International Organization, **49**(3), 379-414.

[13] Fridovich, D.P., Krawchuk F. 2007, *Winning in the Pacific: The Special Operations Forces Indirect Approach,* Joint Force Quarterly, **44**, 24–27.

[14] Galula, D. 1964, *Counterinsurgency Warfare: Theory and Practice*, New York, Praeger.

[15] Hastings J.V., Chan R.J. 2013, *Target Hardening and Terrorist Signaling: The Case of Aviation Security*, Terrorism and Political Violence, **25**(5), 777-797.

[16] Kalyvas S. 2006, *The Logic of Violence in Civil War*, Cambridge University Press.

[17] Laffont J.J., Martimort D. 2002, *The Theory of Incentives*, Princeton University Press.

[18] Laffont J.J., Tirole J. 1993, *A Theory of Incentives in Procurement and Regulation*, MIT Press.

[19] Lapan, H.E., Sandler T. 1993, *Terrorism and signaling*, Eur. J. Political Econom., **9**(3) 383–397.

[20] McMaster H.R. 2008, *On War: Lessons to Be Learned*, Survival, **50**, 19-30.

[21] Myerson R.B., 1981, *Optimal auction design,* Mathematics of Operations Research, **6.1**, 58-73.

[22] Morrow, J. D. 1992, *Signaling Difficulties with Linkage in Crisis Bargaining*, International Studies Quarterly, **36**(2), 153-172.

[23] Morrow, J. D. 1994, Modeling the forms of international cooperation: distribution versus information, International Organization, **48**(3), 387-423.

[24] Nagl, J. 2002, *Counterinsurgency Lessons from Malaya and Vietnam: Learning to Eat Soup with a Knife*, Westport, CT: Praeger.

[25] Petersen, R. 2001, *Resistance and Rebellion: Lessons from Eastern Europe*, Studies in Rationality and Social Change, Cambridge, Cambridge Univ. Press.

[26] Petraeus, D. 2006, *Learning Counterinsurgency: Observations from Soldiering in Iraq*, Military Rev. (January–February), 45–55.

[27] Popkin S.L. 1979, *The Rational Peasant: The Political Economy of Rural Society in Vietnam*, Berkeley, University of California Press.

[28] Powell R. 1991, *Absolute and Relative Gains in International Relations Theory*, American Political Science Review, **85**, 1303-1320.

[29] Powell R. 2006, *War as a Commitment Problem*, International Organization, **60**, 169-203.

[30] Sandler T. 2015, *Terrorism and Counterterrorism: an Overview*, Oxford Economic Papers, **67**(1), 1-20.

[31] Sepp K. 2005, *Best Practices in Counterinsurgency*, Military Review, 8-12.

[32] Schneider F., Brück T., Meierrieks D. 2015, *The Economics of Counterterrorism: A Survey*, Journal of Economic Surveys, **29**(1), 131-157.

[33] Wang C., Bier V.M. 2011, *Target-hardening Decisions Based on Uncertain Multiattribute Terrorist Utility*, Decision Analysis, **8**(4), 286-302.

Chapter 3

Cheap Talk with Multiple Experts and Uncertain Biases

3.1 Introduction

Conventional wisdom suggests that getting a second opinion is helpful for decision making and is common in many real-life situations. In healthcare markets, for instance, patients often seek a second advice to find the right diagnosis. Universities often ask more than one recommendation letter before making tenure decisions, and customers often talk to several salesmen to find the product that better fits their needs. All these examples suggest that decision makers may wish to consult more than one expert in order to make sound decisions. However, experts often have different preferences vis- à- vis the decision maker and this makes communication difficult.[1] In particular, when the talk is cheap and unverifiable, biased experts may have incentives to strategically alter their advice in order to push the decision maker towards a certain direction.

Many existing models explain why, and under which conditions, an uninformed decision maker benefits from consulting multiple experts before making a decision (See e.g., Sobel, 2013, for a survey). However, most of these models assume that the experts' biases are known, whereas little is known about the communication when the bias of the expert is private information. Do experts have incentives to share their private information with the decision maker? What is the effect of this information asymmetry on the decision maker's behavior? Is it better to consult two experts or just one?

We address these issues by analyzing a simple cheap talk model adopted from Austen-Smith

I am particularly grateful to Giacomo Calzolari, Nenad Kos, Salvatore Piccolo and Giovanni Ursino for their constructive criticism and useful suggestions. I would like to thank to Koray Aktas, Simone Boccaletti, Diogo Britto, Ferdinando Colombo, Ginevra Gallassi and Aldo Pignataro for insightful comments and suggestions. All remaining errors are my own.

[1]For instance, a salesman may promote a specific product in order to get a higher commission.

(1993). We consider an environment in which an uninformed decision maker seeks advice from either one or two partially informed experts before taking a payoff relevant action. Each expert receives a private binary signal about the state of the world and provides information to the decision maker through simultaneous cheap talk. The decision maker and the experts have different preferences (or biases) over actions. The decision maker's bias is common knowledge across players, while each expert is privately informed about his bias. Two different types of bias are considered: an expert is either moderately biased (hereafter moderate expert), whose bias is small, or extremely biased (hereafter extreme expert), whose bias is large in absolute terms. More precisely, an expert's bias measures how distant his preferences are relative to those of the decision maker. In other words, a moderate expert is less biased than an extreme expert.

Building on this insight, we focus on two informative equilibria in which the decision maker can learn some information from the experts' messages. First, we consider a fully-revealing equilibrium in which experts of either type truthfully reveal their privately observed signals about the state of the world and the decision maker believes them. Second, we consider a semi-revealing equilibrium in which a moderate expert is willing to send informative messages to the decision maker depending on his privately observed signal, while an extreme expert reports the same message independent of his private information, so no information can be inferred from his message.

We first examine the effect of uncertain biases on the decision maker's action and on the experts' truth-telling incentives. We show that in a fully-revealing equilibrium the conditions for the existence of such equilibrium are not different from those one would obtain if the biases were known. By contrast, in a semi-revealing equilibrium the fact that the decision maker is uncertain about the experts' biases affects the incentives to disclose information. In particular, the interval that supports truth-telling as an equilibrium is small compared to that of the fully-revealing equilibrium. The reason is that in a semi-revealing equilibrium the decision maker knows that with some probability each expert reports a message which does not necessarily reflect the privately observed signal and the decision maker updates his/her beliefs accordingly. As a consequence, this makes incentives to lie stronger and the truth-telling condition tighter.

After characterizing the conditions for the existence of fully-revealing and semi-revealing equilibria, we go on to make welfare comparisons. In order to make a welfare comparison, we use the ex-ante expected utility of the decision maker as a welfare measure. Interestingly, we find that the fully-revealing equilibrium with one expert may be informationally superior to the semi-revealing equilibrium with two experts. Specifically, uncertainty over biases allows experts to lie relatively more often as compared to fully-revealing case, whereby reducing the information content of the messages. With two experts, however, the decision maker has a higher chance to get truthful information from one of the experts which, in turn, may provide more information than the one-expert

[2] In Section 3.4, the informational properties of these equilibria with experts are compared to single expert. Specifically, we show that both in a fully-revealing and semi-revealing equilibria, the ex-ante expected utility of the decision maker with two experts is higher than that with a single expert.

communication does.[2] The net effect on the decision maker's ex-ante expected profit depends on the probability that the decision maker believes the expert to be moderate — i.e., whether the expert's report is informative or not.

The rest of the paper is organized as follows. After discussing the related literature, Section 3.2 describes the baseline model. Section 3.3.1 characterizes the conditions under which a fully-revealing equilibrium exists. In Section 3.3.2, we characterize the conditions under which a semi-revealing equilibrium exists. Section 3.4 discusses the welfare. The last section concludes. All proofs are in the Appendix.

Related Literature. We build on and contribute to two strands of literature. First, this paper relates to the literature on cheap talk with multiple experts. Gilligan and Krehbiel (1989) first characterized cheap talk model with two perfectly informed experts in a one-dimensional environment. Krishna and Morgan (2001) considers a cheap talk model with two perfectly informed experts to show that when the decision maker consults two experts are who are biased in the same direction, the most informative equilibrium is obtained by consulting the less biased expert alone. Gick (2006) studies a cheap talk model in which an uninformed decision maker seeks advice from two perfectly informed experts. He shows that having a second expert, even if he is more biased than the first one, improves the information structure when the communication is simultaneous.[3] The analysis in this paper is related to that in Austen-Smith (1993), who considers a uniform state space, and assumes that the experts are partially informed about the underlying state, as this paper does. However, we allow the decision maker to be uncertain about the experts' biases. Specifically, Austen-Smith (1993) shows that simultaneously consulting two experts leads to higher welfare than consulting only one expert, while in this paper we find that there exist circumstances under which two-expert communication is not necessarily superior to one-expert communication.

Second, this paper is related to cheap talk literature with uncertain individual preferences. There is a growing literature that considers experts' reputational/career concerns as a source of uncertainty. For instance, Sobel (1985), Bénabou and Laroque (1992), Morris (2001), Gentzkow and Shapiro (2006), Ottaviani and Sørensen (2006) consider uncertainty about expert types and focus on the reputational incentives which this paper does not address. Few papers focus on the informativeness of the communication with uncertain biases. In particular, Morgan and Stocken (2008) and Dimitrakas and Sarafidis (2005) show that revelation of the expert's bias weakens the communication when the magnitude of the bias is uncertain. Interestingly, Li (2004) and Li and Madarász (2008) characterize cheap talk equilibria with uncertain (and exogenous) biases in one expert mechanism.[4] Both these papers consider uniform state space and allow two values of the bias as this paper does. However, they assume the expert can perfectly observe the state. They show that the revelation of the bias always weakens the communication when there is uncertainty

[3]See also Li (2004) for a cheap talk model with multiple experts with sequential communication.

[4]Hence, they do not provide welfare comparison between one and two-experts mechanisms.

on the direction of the bias. In this paper, however, we show that transparency of biases enlarges the truth-telling interval, and, hence improves the incentives to truthfully communicate with the decision maker.

3.2 The Model

Players and Environment. Consider a decision maker (female), D, who seeks advice from two (male) experts, A_1 and A_2. The decision maker takes an action $y \in \mathbb{R}$ that affects the payoffs of all players. The state of the world, θ, is a random variable and uniformly distributed on $[0,1]$, with density $f(\theta) = 1$. The decision maker has no further information about θ, while each expert privately observes a binary signal about the state. The parameter $s_i \in \mathcal{S} \triangleq \{0,1\}$ denotes the signal observed by A_i such that each signal is equally likely $\Pr[s_i] = \frac{1}{2}, s_i \in \mathcal{S}, i = 1,2$.

Following Austen-Smith (1993), we assume that the signals are conditionally independent across experts given the underlying state θ. Each signal s_i has the following conditional probability

$$\Pr[s_i|\theta] = \theta^{s_i}(1-\theta)^{1-s_i} \qquad s_i \in \mathcal{S}. \tag{3.1}$$

Conditional on the state θ, therefore, the joint probability distribution of the signals is such that

$$\Pr[s_i, s_j|\theta] = \theta^{s_i+s_j}(1-\theta)^{2-s_i-s_j} \qquad s_i, s_j \in \mathcal{S}. \tag{3.2}$$

Based upon the realized signal, each expert simultaneously reports a message to the decision maker. Let m_i be A_i's message, and, for simplicity, we consider a binary message space such that $m_i \in \mathcal{M} \triangleq \{0,1\}$[5].

Based upon the received messages, the decision maker takes an action $y(m_i, m_j)$ that affects the payoffs of all players.

All players have quadratic loss utility functions. Specifically, D's utility is

$$\mathcal{U}_D(y,\theta,b_D) \triangleq -(y-\theta-b_D)^2,$$

and A_i's utility is

$$\mathcal{U}_i(y,\theta,b_i) \triangleq -(y-\theta-b_i)^2, \quad i=1,2.$$

The quadratic loss utility function is commonly used in the cheap talk literature (e.g., Crawford and Sobel, 1982; Austen-Smith, 1993; Farrell and Gibbons, 1996; Morgan and Stocken, 2008;

[5]The use of binary messages is without loss of generality because the state of the world is uniformly distributed on the unit interval and the signal space is assumed to be binary. Hence, the decision maker's uncertainty is just relative to these binary signals about the state so that a binary message space available has enough elements to transmit any information available for the experts. See, e.g., Kawamura (2011) for a formal proof.

among many others) since it allows to obtain (tractable) closed form solutions. The quadratic loss utility function has an important implication because it guarantees the concavity of D's objective function and hence uniqueness of the optimal action. Hence, given quadratic loss specification, in state θ, the decision maker's most preferred action is $\theta + b_D$ and A_i's most preferred action is $\theta + b_i$.

The parameter $b_D \geqslant 0$ represents the decision maker's bias and is common knowledge across players. The parameter $b_i \in \mathcal{B} \triangleq \{b_M, b_E\}$, $i = 1, 2$, represents A_i's bias and measures how distant his preferences are relative to those of the decision maker. Crucially, A_i's bias is his own private information and is drawn from the following distribution

$$\Pr[b_i = b_M] \triangleq \nu \triangleq 1 - \Pr[b_i = b_E] \qquad i = 1, 2.$$

Hence, A_i knows his own bias, while D and A_j have only a prior about that.[6] Specifically, if $b_i = b_M$, the bias is moderate and A_i is said to be a moderate expert, while if $b_i = b_E$, the bias is extreme and A_i is said to be an extreme expert where a moderate expert is assumed to be less biased than an extreme expert — i.e., $|b_M - b_D| < |b_E - b_D|$. Finally, all players are expected utility maximizers.

Timing. The timing is as follows.

- Nature randomly chooses θ according to a uniform distribution on $[0, 1]$.
- Each expert observes s_i.
- Each expert simultaneously sends m_i to the decision maker.
- Based upon the received messages, D takes an action $y \in \mathbb{R}$.

Equilibrium. The solution concept is Perfect Bayesian Equilibrium (PBE). For simplicity, we consider only pure strategies for the experts (See e.g., Austen-Smith (1993), Li (2004) among many others).[7]

As it is common in cheap talk models, multiple equilibria exist. In particular, a babbling equilibrium always exists, in which the messages do not depend on the experts' private information about the underlying state. Indeed, given such strategy, it is optimal for the decision maker to ignore the messages, but then babbling is actually a best response for the experts. However, we focus on two informative equilibria: (i) fully-revealing equilibrium in which experts of either type truthfully report their signals about the underlying state and the decision maker believes them; (ii) semi-revealing equilibrium in which a moderate expert truthfully reports his private signal while an extreme expert reports the same message regardless of his private information about the state.

[6] For a similar approach, see Morris (2001), Morgan and Stocken (2003) and Li (2004).

[7] Notice that, all messages fall on the equilibrium path. Hence, no off-equilibrium path beliefs are required.

Without loss of generality, in the analysis that follows we assume that the extreme expert is rightward biased — i.e., $b_D < b_E$. Assuming a rightward biased extreme expert is with no loss of generality because experts' payoffs are symmetric and the message space is binary.[8] As it will be clear soon, in a semi-revealing equilibrium, a rightward biased extreme expert always reports, with a slight abuse of notation, $m_E = 1$ independent of his signal — i.e., such that he wants as high action as possible relative to the decision maker. In fact, when he observes a signal equal to 1, he wants to report $m_E = 1$ instead of 0 because, by doing so, he is able to shift the decision maker's action rightward. Moreover, we do not impose any restrictions on the direction of the moderate bias. This is due to the fact that a moderate expert, in equilibrium, is willing to send both messages (both 0 and 1) depending on his privately observed signal. Therefore, he wants as high (resp. low) action as possible if $b_D < b_M$ (resp. $b_D > b_M$).

3.3 Equilibrium Analysis

We now characterize the decision maker's optimal action after receiving any messages then analyze experts incentives to communicate in fully-revealing and semi-revealing equilibria with one and two experts.[9]

3.3.1 Fully-Revealing Equilibrium

To gain intuition about the central result of the paper, we first analyze a simple case in which the experts simultaneously and truthfully report their private signals — i.e., such that $m_i = s_i$ and $m_j = s_j$ in equilibrium — and the decision maker believes them. Since the experts' messages reflect the true realizations of the signals, D's best response to such strategy is

$$\begin{aligned} y^F(s_i, s_j) &= \arg\max_{y \in \mathbb{R}} \int_\theta -(y - \theta - b_D)^2 f(\theta | s_i, s_j)\, d\theta, \\ &= b_D + \mathrm{E}[\theta | s_i, s_j] \qquad \forall (s_i, s_j) \in \mathcal{S}^2, \end{aligned} \tag{3.3}$$

where, abusing slightly notation, we define $y^F_{s_i,s_j} \triangleq y^F(s_i, s_j)$ and the superscript F denotes the optimal action taken by the decision maker after being truthfully informed about the signals. The expression in (3.3) simply reflects that when D receives truthful messages from the experts, her optimal action is just the conditional expectation of the state shifted by her own bias b_D.

The following lemma characterizes the decision maker's optimal action after being truthfully informed by one or two experts.

[8]Hence, the equilibrium in which the extreme expert is leftward biased expert is just the mirror image of the equilibrium with the rightward biased extreme expert.

[9]It is worth pointing out that the model with one expert is identical to the model with two experts. A detailed equilibrium analysis with one expert can be found in Appendix.

Lemma 3.1. *In a fully-revealing equilibrium, when D consults only one expert, her optimal actions are*

$$y_0^F = b_D + \frac{1}{3}, \qquad y_1^F = b_D + \frac{2}{3},$$

while when D simultaneously consults two experts, her optimal actions are

$$y_{0,0}^F = b_D + \frac{1}{4}, \qquad y_{0,1}^F = y_{1,0}^F = b_D + \frac{1}{2}, \qquad y_{1,1}^F = b_D + \frac{3}{4}.$$

Hence, in a fully-revealing equilibrium (both with one and two experts) uncertainty about the expert's types has no consequence on the optimal actions because D believes that experts truthfully report their private signals regardless of their type. Moreover, the optimal actions are such that $y_{0,0}^F < y_{0,1}^F < y_{1,1}^F$. The reason is simple: when D receives two different signals, she takes an action based on her prior beliefs about the state. Instead, when D receives two identical signals from the experts, she has a more precise idea regarding the state because both experts report their signals truthfully. As a result, this shifts the decision maker's action rightward when she receives $(s_i, s_j) = (1, 1)$, and shifts it leftward when she receives $(s_i, s_j) = (0, 0)$ from the experts. A similar logic applies when D consults one expert.

Consider now the experts' incentives to reveal the observed signals. Without loss of generality, we focus on the truth-telling incentives of A_i because the experts are ex-ante symmetric. Notice that from A_i's perspective A_j truthfully reports his signal — i.e., in equilibrium $m_j = s_j$. Hence, there exists a fully-revealing equilibrium if there is an incentive for A_i to report truthfully $m_i = s_i$ instead of false message $m_i = 1 - s_i$ along the equilibrium path. This condition is

$$\begin{aligned} &\sum_{s_j \in \mathcal{S}} \int_\theta -\left(y_{s_i,s_j}^F - \theta - b_i\right)^2 f\left(s_j, \theta | s_i\right) d\theta \geqslant \\ &\geqslant \sum_{s_j \in \mathcal{S}} \int_\theta -\left(y_{1-s_i,s_j}^F - \theta - b_i\right)^2 f\left(s_j, \theta | s_i\right) d\theta, \qquad b_i \in B. \end{aligned} \tag{3.4}$$

Let

$$\Delta y^F\left(s_i, s_j\right) \triangleq y_{1-s_i,s_j}^F - y_{s_i,s_j}^F,$$

be the difference between D's action after receiving false and correct signal from A_i given that A_j reports his signal truthfully in equilibrium. Taking into account D's optimal action after hearing the truthful messages (3.3), integrating and rearranging terms the above constraint simplifies to

$$\left(b_i - b_D\right) \sum_{s_j \in \mathcal{S}} \Pr\left[s_j | s_i\right] \underbrace{\Delta y^F\left(s_i, s_j\right)}_{\text{Overshooting Effect}} \leqslant \sum_{s_j \in \mathcal{S}} \Pr\left[s_j | s_i\right] \frac{\Delta y^F\left(s_i, s_j\right)^2}{2}. \tag{3.5}$$

Condition (3.5) reflects that A_i's incentive to report his private signal is shaped by D's reaction to receiving false information from A_i — i.e., the *overshooting effect* (highlighted in Morgan and Stocken, 2008): a deviation from a truthful message may shift the decision maker's action too far from the expert's ideal action. More specifically, an expert with rightward bias (resp. leftward bias) may prefer a higher (resp. lower) action than the decision maker, but the displacement in decision maker's action caused by an undetectable lie might be too large relative to the case of truth-telling, which is not desirable for either the expert or the decision maker. As we shall explain below, this makes truth-telling an optimal strategy for an expert who has preferences close to the those of the decision maker. Other things being equal, the sign of the overshooting effect depends on A_i's privately observed signal. More precisely, if $\Delta y^F(s_i, s_j) > 0$, the overshooting effect is positive and an expert with leftward bias $b_i < b_D$ has no incentive to misreport because sending false message shifts the decision maker's optimal action rightward. Similarly, if $\Delta y^F(s_i, s_j) < 0$, the overshooting effect is negative and an expert with rightward bias has no incentive lie because, in this case, reporting a false signal to the decision maker cannot be incentive compatible.

The following proposition characterizes a fully-revealing equilibrium with one and two experts.

Proposition 3.2. *(i) When D consults one expert, a fully-revealing equilibrium exists if and only if*

$$|b_1 - b_D| \leqslant \frac{1}{6}, \qquad b_1 \in \mathcal{B}.$$

(ii) When D simultaneously consults two experts, a fully-revealing equilibrium exists if and only if

$$|b_i - b_D| \leqslant \frac{1}{8}, \qquad b_i \in \mathcal{B}, i = 1, 2.$$

There are two key aspects to note about Proposition 3.2. First, the maximal distance in preferences (both with one expert and two experts) compatible with full information revelation does not depend on the parameter ν because D believes that an expert of either type truthfully reports his private signal. Accordingly, the impact of each message on D's optimal action is very high. This, in turn, makes truth-telling an optimal strategy for an expert who has preferences close to the those of the decision maker because he cannot do better than report his true signal due to the *overshooting effect*. Since the information asymmetry has no impact on the equilibrium, the conditions for its existence are not different from those that would obtain if biases were known.

Second, when D consults one expert, the magnitude of the overshooting effect is $\left|\Delta y^F(s_1)\right| \triangleq \frac{1}{3}$, $s_1 \in \mathcal{S}$, while when D consults two experts it is $\left|\Delta y^F(s_i, s_j)\right| \triangleq \frac{1}{4}$, $(s_i, s_j) \in \mathcal{S}^2$. This tells us, when D consults one expert, the displacement in decision maker's action caused by an undetectable lie is large compared to the case with two experts. This, in turn, increases A_i's incentives to misreport. Hence, A_i's preferences should be even more close to those of decision maker's (as compared to one expert) in order to reveal his private information. As a result, having multiple experts makes the truth-telling conditions tighter relative to the case where D consults only one expert.

3.3.2 Semi-Revealing Equilibrium

Consider now a semi-revealing equilibrium, in which the moderate expert truthfully reports his signal, while the extreme expert (rightward biased) reports $m_E = 1$ independent of his private signal. The structure of D's maximization problem is similar to that solved in a fully-revealing equilibrium with the difference that she must form beliefs about the signals (s_i, s_j) given the message pair (m_i, m_j) because the messages may not necessarily reflect the privately observed signals. The Bayes rule then implies the following posterior

$$\Pr[s_i, s_j | m_i, m_j] \triangleq \frac{\Pr[m_i, m_j | s_i, s_j] \Pr[s_i, s_j]}{\sum_{(s_i, s_j) \in \mathcal{S}^2} \Pr[m_i, m_j | s_i, s_j] \Pr[s_i, s_j]},$$

where $\Pr[s_i, s_j]$ is the joint probability of the signals.

To understand the updating process, notice that when D receives $(m_i, m_j) = (0,0)$ from the experts, she will be sure that these messages come from two moderate experts who tell the truth. As a consequence, the messages convey full information about the signals — i.e.,

$$\Pr[s_i = 0, s_j = 0 | 0, 0] = 1 \quad \text{and} \quad \Pr[s_i = 1, s_j = 1 | 0, 0] = 0. \tag{3.6}$$

When, instead, the decision maker receives $(m_i, m_j) = (1,1)$ from the experts, she is uncertain about the types/biases of the experts. As a consequence, she must update beliefs discounting the possibility of receiving uninformative message(s). In this case, by Bayes' rule D's posterior beliefs are

$$\Pr[s_i = 0, s_j = 0 | 1, 1] = \frac{(1-\nu)^2}{\nu^2 - 3\nu + 3}, \qquad \Pr[s_i = 1, s_j = 1 | 1, 1] = \frac{1}{\nu^2 - 3\nu + 3}. \tag{3.7}$$

Notice that

$$\frac{d \Pr[s_i = 0, s_j = 0 | 1, 1]}{d\nu} < 0 \qquad \text{and} \qquad \frac{d \Pr[s_i = 1, s_j = 1 | 1, 1]}{d\nu} > 0.$$

Hence, when D receives $(m_i, m_j) = (1,1)$, an increase of the probability of being moderate makes her more confident that the signals are $(s_i, s_j) = (1,1)$ and *vice versa*. Similar reasoning applies (See the Appendix) for the mixed messages and signals.

Hence, D's problem is

$$y^S(m_i, m_j) = \arg\max_{y \in \mathbb{R}} \int_\theta -(y - \theta - b_D)^2 f(\theta | m_i, m_j)\, d\theta,$$

whose solution yields,

$$y^S(m_i, m_j) = b_D + \underbrace{\sum_{(s_i,s_j)\in\mathcal{S}^2} \Pr[s_i, s_j | m_i, m_j]\,\mathrm{E}[\theta | s_i, s_j]}_{\triangleq \mathrm{E}_\nu[\theta | m_i, m_j]}, \tag{3.8}$$

where, abusing notation, we define $y^S_{m_i,m_j} \triangleq y^S(m_i, m_j)$ and the superscript S denotes the optimal actions taken by the decision maker in a semi-revealing equilibrium.

The following lemma describes the decision maker's optimal actions in a semi-revealing equilibrium.

Lemma 3.3. *In a semi-revealing equilibrium, when D consults only one expert, her optimal actions are*

$$y^S_0 = b_D + \frac{1}{3}, \qquad y^S_1 = b_D + \frac{3-\nu}{3(2-\nu)},$$

while when D consults two experts, her optimal actions are

$$y^S_{0,0} = b_D + \frac{1}{4}, \qquad y^S_{0,1} = y^S_{1,0} = b_D + \frac{2-\nu}{2(3-2\nu)}, \qquad y^S_{1,1} = b_D + \frac{\nu^2 - 4\nu + 6}{4(\nu^2 - 3\nu + 3)}.$$

Hence, even with uncertain biases, we have $y^S_{0,0} < y^S_{0,1} < y^S_{1,1}$. Clearly, when D receives $(m_i, m_j) = (0, 0)$ from the experts, she will be sure that the messages are sent by two moderate experts who tell the truth because the extreme expert is rightward biased and has no incentive to report 0. In this case, the decision maker's optimal action in a semi-revealing equilibrium coincides with her optimal action in a fully-revealing equilibrium — i.e., $y^S_{0,0} = y^F_{0,0}$ — as expected. By contrast, when D receives any other messages that contains at least one message equal to 1, she discounts the possibility of receiving false information and hence the experts' messages have a lower impact on the action taken by the decision maker. This implies that D's optimal action in a semi-revealing equilibrium is lower than the one in a fully-revealing equilibrium — i.e., $y^F_{1,0} > y^S_{1,0}$ and $y^F_{1,1} > y^S_{1,1}$ for all $\nu \in (0, 1)$. Notice also that the higher are the chances of being moderate, the more 'accurate' the inference that D can make on the messages given the signals. Hence, the optimal actions converge to those found in a fully-revealing equilibrium as ν tends to 1.

A similar reasoning applies when D consults single expert. More precisely, when D receives $m_1 = 0$ from the expert, she will be sure that the expert is moderate and is reporting truthfully. Hence, the decision maker will assign probability 1 to $b_1 = b_M$. In this case, not surprisingly, D's optimal action in a semi-revealing equilibrium with one expert coincides with her optimal action in a fully-revealing equilibrium — i.e., $y^F_0 = y^S_0$. By contrast, when D receives message $m_1 = 1$, she discounts the possibility that the expert is extreme (in which case the message reveals no information), and hence the expert's message has a lower impact on the final decision than the one in a fully-revealing equilibrium. An important point here is to note that the optimal actions are $y^S_1 < y^S_{1,1}$ and $y^S_{1,0} < y^S_1$ for all $\nu \in (0, 1)$. The first inequality follows from the fact that, when D consults two experts she has a higher chance to get truthful information from one of the experts.

The second inequality follows from observing that, when D receives any messages that contains at least one message equal to 0, she can infer with certainty that the message is sent by a moderate expert. Hence, the decision maker's optimal action is lower when she receives $(m_i, m_j) = (1, 0)$ than receiving only one message $m_1 = 1$.

Consider now the experts' incentives to reveal their private signals. As before, we focus on the truth-telling incentives of A_i since experts are ex-ante symmetric. Suppose now that A_i is moderate — i.e., such that $b_i = b_M$. Given that A_j's bias is his private information, from A_i's perspective A_j is either moderate with probability ν or extreme with probability $1 - \nu$. Hence, A_i's expected utility when reporting $m_i = s_i$ is higher than his expected utility when reporting a false message $m_i = 1 - s_i$ if

$$\sum_{b_j \in \mathcal{B}} \Pr[b_j] \sum_{s_j \in \mathcal{S}} \int_\theta -\left(y^S_{m_i, m_j} - \theta - b_M\right)^2 f(s_j, \theta | s_i)\, d\theta \ \geqslant$$
$$\geqslant \sum_{b_j \in \mathcal{B}} \Pr[b_j] \sum_{s_j \in \mathcal{S}} \int_\theta -\left(y^S_{1-m_i, m_j} - \theta - b_M\right)^2 f(s_j, \theta | s_i)\, d\theta. \tag{3.9}$$

For any m_j, let

$$\Delta y^S(m_i, m_j) \triangleq y^S_{1-m_i, m_j} - y^S_{m_i, m_j},$$

be the difference between D's action after receiving false and correct messages from A_i. Taking into account D's optimal action after hearing the signals (3.8), integrating and rearranging terms, (3.9) simplifies to

$$(b_M - b_D) \sum_{b_j \in \mathcal{B}} \Pr[b_j] \sum_{s_j \in \mathcal{S}} \Pr(s_j | s_i) \underbrace{\Delta y^S(m_i, m_j)}_{\text{Overshooting Effect}} \leqslant$$
$$\leqslant \sum_{b_j \in \mathcal{B}} \Pr[b_j] \sum_{s_j \in \mathcal{S}} \Pr(s_j | s_i) \left\{ \frac{\Delta y^S(m_i, m_j)^2}{2} + \Delta y^S(m_i, m_j) \underbrace{\left(\mathrm{E}_\nu[\theta | m_i, m_j] - \mathrm{E}[\theta | s_i, s_j]\right)}_{\text{goes to 0 as } \nu \longrightarrow 1} \right\}. \tag{3.10}$$

In contrast to the case of fully-revealing, the overshooting effect now depends on the parameter ν. Since A_j's bias is his private information and the other players have only a prior about that, A_j's report plays an important role on A_i's incentive to truthfully report his private information. Specifically, the overshooting effect is stronger when A_j reports $m_j = 1$ than when he reports

[10] In fact, for $\nu \in (0, 1)$, we have

$$\frac{1}{4(3 - 2\nu)} \triangleq \left|\Delta y^S(m_i, 0)\right| < \left|\Delta y^S(m_i, 1)\right| \triangleq \frac{\nu^2 - 6\nu + 6}{4(3 - 2\nu)(\nu^2 - 3\nu + 3)}.$$

$m_j = 0$ — i.e., $\left|\Delta y^S(m_i, 0)\right| < \left|\Delta y^S(m_i, 1)\right|$[10]. This is due to the fact that when D receives $m_j = 1$, she anticipates the risk that it is an uninformative message and discounts accordingly A_j's message. This, in turn, puts more weight on A_i's message so that a lie from A_i has a stronger impact on the decision maker's action.

To complete the characterization of the semi-revealing equilibrium, suppose that A_i is extreme — i.e., such that $b_i = b_E$. We need to check that, A_i has no incentive to report $m_i = 0$ when his private signal is $s_i = 0$. Hence, A_i's expected utility from reporting $m_i = 1$ is higher than his utility from reporting truthfully $m_i = 0$ if

$$\sum_{b_j \in \mathcal{B}} \Pr[b_j] \sum_{s_j \in \{0,1\}} \int_\theta -\left(y^S_{0,m_j} - \theta - b_E\right)^2 f(s_j, \theta | s_i = 0)\, d\theta <$$
$$< \sum_{b_j \in \mathcal{B}} \Pr[b_j] \sum_{s_j \in \{0,1\}} \int_\theta -\left(y^S_{1,m_j} - \theta - b_E\right)^2 f(s_j, \theta | s_i = 0)\, d\theta, \qquad (3.11)$$

which, integrating and rearranging terms, simplifies to

$$(b_E - b_D) \sum_{b_j \in \mathcal{B}} \Pr[b_j] \sum_{s_j \in \{0,1\}} \Pr(s_j|0)\, \Delta y^S(0, m_j) >$$
$$> \sum_{b_j \in \mathcal{B}} \Pr[b_j] \sum_{s_j \in \{0,1\}} \Pr(s_j|0) \left\{ \frac{\Delta y^S(1, m_j)^2}{2} + \Delta y^S(1, m_j)\left(\mathrm{E}_\nu[\theta|0, m_j] - \mathrm{E}[\theta|0, s_j]\right) \right\}. \qquad (3.12)$$

Clearly, when $s_i = 1$, a rightward biased extreme expert has an incentive to report $m_E = 1$. The following proposition characterizes a semi-revealing equilibrium with one and two experts.

Proposition 3.4. *(i) When D consults only one expert, there exist two thresholds $\alpha_1(\nu)$ and $\beta_1(\nu)$, with $0 < \alpha_1(\nu) < \beta_1(\nu)$, such that a semi-revealing equilibrium in which the extreme expert reports $m_E = 1$ exists if and only if*

$$-\beta_1(\nu) \leqslant b_M - b_D \leqslant \alpha_1(\nu) \qquad \textit{and} \qquad b_E - b_D > \alpha_1(\nu).$$

(ii) When D simultaneously consults two experts, there exist two thresholds $\alpha_2(\nu)$ and $\beta_2(\nu)$, with $\alpha_2(\nu) < \alpha_1(\nu)$ and $\beta_2(\nu) < \beta_1(\nu)$, such that a semi-revealing equilibrium in which the extreme expert reports $m_E = 1$ exists if and only if

$$-\beta_2(\nu) \leqslant b_M - b_D \leqslant \alpha_2(\nu) \qquad \textit{and} \qquad b_E - b_D > \alpha_2(\nu).$$

Moreover, $\alpha_1(\nu)$ and $\alpha_2(\nu)$ are increasing in ν, while $\beta_1(\nu)$ and $\beta_2(\nu)$ are decreasing in ν.

Figure 3.1 illustrates the region of parameters, identified in Propositions 3.2 and 3.4, where fully-revealing and semi-revealing equilibria exist with one and two experts.[11] Uncertainty about the experts' biases has two effects on information transmission. First, the interval that supports truth-telling as an equilibrium shrinks as the probability of being moderate tends to zero. The intuition is straightforward: in a semi-revealing equilibrium D knows that with some probability, each expert reports a message which does not necessarily reflect the privately observed signal. Hence, D updates her beliefs discounting the possibility of receiving uninformative messages. This implies that each message has a lower impact on D's action in equilibrium. This, in turn, makes the incentives to lie stronger when an expert observes a signal that would shift the decision maker's action in an undesired direction if reported truthfully. Moreover, observe that the thresholds $\alpha_2(\nu)$ is increasing in ν and $\beta_2(\nu)$ is decreasing in ν. That is, when the probability of being moderate increases, the truth-telling interval in a semi-revealing equilibrium enlarges, and eventually, it coincides with the truth-telling interval in a fully-revealing equilibrium.[12]

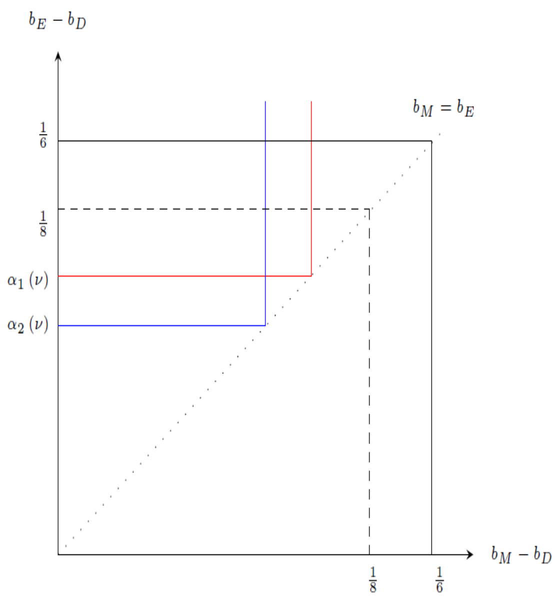

For $b_M > b_D$ and $b_E > b_D$ in the region bounded by black lines there exists a fully-revealing equilibrium with one expert; in the region bounded by dashed lines, there exists a fully-revealing equilibrium with two experts; in the region bounded by red lines, there exists a semi-revealing equilibrium with one expert; the region bounded by blue lines, semi-revealing equilibrium with two experts is an equilibrium.

Figure 3.1: Truth-telling thresholds for fully-revealing and semi-revealing equilibria.

Second, the conditions for truth-telling are tighter when the decision maker consults two experts rather than just one. To understand why, recall that an expert, say A_i, has only a prior

[11] For the sake of clarity, we focus on the situation where both types of experts are biased in the same direction relative to the decision maker — i.e., $b_D < b_M < b_E$.

[12] In the Appendix we derive a closed from solution for the thresholds $\alpha_i(\nu)$ and $\beta_i(\nu), i = 1, 2$, as a function of the bias parameter ν. In fact, our model is based on the quadratic-uniform setting and this permits to obtain closed form solutions for the threshold equilibria. Hence, closed from solutions deliver additional comparative statics to those mentioned above.

about the type of the other expert. When A_i is consulted alone, D's optimal action is conditioned only on his report and this makes him relatively sure of the consequence of the message that he sends to D. When there are two experts, instead, A_i is unsure about the weight of his message because D's optimal action depends on A_j's report too. As a result, the presence of another expert with unknown bias makes incentives to lie stronger relative to the communication with one expert.

3.4 Welfare Comparison

In order to study the welfare effects, we now turn to compare the decision maker's ex-ante expected utility among the types of equilibria defined in Proposition 3.2 and 3.4. First, we compare D's ex-ante expected utility with one and two experts within each equilibria. We have the following result.

Proposition 3.5. *Both in fully-revealing and semi-revealing equilibria, consulting two experts is informationally superior to consulting just one.*

Not surprisingly, in a fully-revealing equilibrium, two experts provide more information to the decision maker than a single expert. In other words, although two-expert communication reduces the size of the interval which supports truth-telling as an equilibrium, it induces D to take an action as a combination of two truthful messages. By doing so, D can have a more precise idea about the underlying state, allowing her to take a more precise action. Therefore, the fully-revealing equilibrium with two experts is informationally superior to the fully-revealing equilibrium with one expert.

The same conclusion holds even when the experts report noisy information to the decision maker. Although the magnitude of the overshooting effect is attenuated in a semi-revealing equilibrium relative to a fully-revealing equilibrium, the improvement in information transmission in two-expert communication is due to the fact that D has a higher chance to get truthful information from one of the experts.

Figure 3.2 plots the welfare maximizing equilibrium within each interval defined in Propositions 3.2 and 3.4. For the sake of clarity, we focus on the situation where both types of experts are biased in the same direction relative to the decision maker — i.e., $b_D < b_M < b_E$.

We now compare D's expected utility in a fully-revealing equilibrium with one expert and a semi-revealing equilibrium with two experts. Hence, we can establish the following result:

Proposition 3.6. *There exists a threshold $\widetilde{\nu}$ such that fully-revealing with one expert is informationally superior to semi-revealing equilibrium with two experts if $\nu \leqslant \widetilde{\nu}$.*

Surprisingly, when the probability of being moderate is sufficiently low, the decision maker prefers to consult a single expert. To see why, let us first consider the region of parameters where

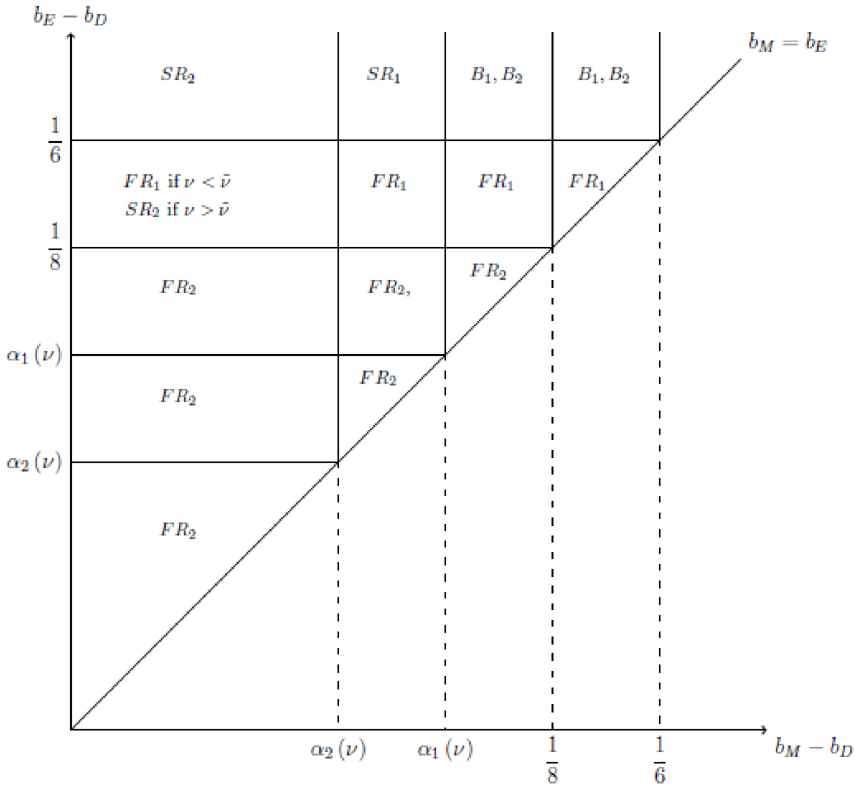

Figure 3.2: Welfare maximizing equilibria.

the two equilibria exist — i.e., a fully-revealing revealing equilibrium with one expert and a semi-revealing equilibrium with two experts (See Figure 3.2). It is immediate to see that these two equilibria obtain when

- a moderate type has preferences close enough to those of the decision maker to induce him to report truthfully his signal regardless of the strategy of the other expert — i.e., $b_M - b_D < \alpha_2(\nu)$, and
- an extreme expert has distant enough preferences that induces him not to report truthfully if the other expert does so *but* close enough that, if consulted alone, he would report truthfully his signal — i.e., such that $\frac{1}{8} < b_E - b_D < \frac{1}{6}$.

Now, in the aforementioned region of parameters, if ν is low the truth-telling interval in a semi-revealing equilibrium is small too because the threshold $\alpha_2(\nu)$ is increasing in ν. Hence, for low values of ν, the moderate expert has lower incentives to report truthfully his signal, and, hence, the information content of his message decreases. In this case, consulting two experts with uncertain biases increases the likelihood of receiving false information from the experts, which, in turn, lowers the ex-ante expected profit of the decision maker. Therefore, when ν is sufficiently low, the decision maker prefers to consult a single expert who report truthfully his signal. By contrast, when $\nu > \widetilde{\nu}$, the decision maker prefers to consult two experts with uncertain biases rather than a single expert because the experts distort information less when ν is high. In this case, D has a higher chance to get truthful information from the experts, who provide more information than a single expert.

Taken together, our result suggests that it may be optimal to consult a single expert rather than two experts whenever the biases of the experts are not too similar, the extreme expert is not too extreme and the probability of being moderate is sufficiently low.

3.5 Concluding Remarks and Further Research

It is commonly believed that seeking advice from multiple sources improves the information transmission between the uninformed party and the informed parties. This presumption may be incorrect, especially when there is uncertainty about the experts' biases. Specifically, we showed that the decision maker may prefer to consult a single expert rather than two experts if the probability that the decision maker believes the expert to be extreme (resp. moderate) is sufficiently high (resp. low). The reason is that consulting two experts with uncertain biases increases the likelihood of receiving distorted information from the extreme experts, and hence, a semi-revealing with two experts comes close to a babbling equilibrium. This suggests that talking to multiple experts in order to elicit information from them about the true state is not always ex-ante efficient.

Some extensions of our analysis would deserve further attention. First, one could generalize the analysis with one decision maker and three or more experts. Our conjecture is that, in a fully-revealing equilibrium, experts' incentives to misreport about their privately observed signals decrease as the number of experts increases. As a result, as the number of experts becomes larger, each agent may have less influence on the final decision. On the contrary, in a semi-revealing equilibrium with three or more experts, uncertainty about the experts' biases may create additional and possibly non-monotonic effect on the truth-telling interval. In order to see the net effect of having three or more experts on the decision maker's expected profit, however, it is necessary to study the overshooting effect in a generalized setup with many experts. Second, we focused on the situation in which experts reveal information only about their private signals about the state but it could be interesting to introduce the possibility that experts with different biases convey information about their biases too. One way to do this is to extend our model to multidimensional cheap talk setup (highlighted in Battaglini, 2002) by considering the uncertainty parameter with two components and the final decision with a two-dimensional vector. In order to drag the decision maker towards a certain direction in each dimension, the experts may have even stronger incentives to disclose their private information with the decision maker relative to the unidimensional case. This extension is left for future research.

3.6 Appendix

Proof of Lemma 3.1. *(i) One Expert.* Suppose that D consults one expert who truthfully reports his signal. Since the utility function is concave in y, the (expected) utility maximizing action of the decision maker after receiving $m_1 = s_1$ can be defined as follows

$$\begin{aligned} y^F_{s_1} &= \arg\max_{y\in\mathbb{R}} \int_\theta -(y-\theta-b_D)^2 f(\theta|s_1)\,d\theta, \\ &= b_D + \underbrace{\int_\theta \theta f(\theta|s_1)\,d\theta}_{\triangleq \mathrm{E}[\theta|s_1]}, \qquad \forall s_1 \in \mathcal{S}, \end{aligned} \tag{C1}$$

where the conditional density of θ given the signal s_1 is

$$f(\theta|s_1) = \frac{\Pr[s_1|\theta]\, f(\theta)}{\int_\theta \Pr[s_1|\theta]\, f(\theta)\,d\theta}.$$

Using the conditional probability distribution of the signal from (3.1), we obtain

$$f(\theta|s_1=0) = 2(1-\theta), \qquad\qquad f(\theta|s_1=1) = 2\theta. \tag{C2}$$

Substituting (C2) into (C1), it is immediate to verify that

$$y^F_0 = b_D + \underbrace{\frac{1}{3}}_{\mathrm{E}[\theta|s_1=0]}, \qquad\qquad y^F_1 = b_D + \underbrace{\frac{2}{3}}_{\mathrm{E}[\theta|s_1=0]}, \tag{C3}$$

as claimed.

(ii) Two Experts. From (3.3) we know that D's optimal action after receiving $m_i = s_i$ and $m_j = s_j$ is

$$y^F_{s_i,s_j} = b_D + \underbrace{\int_\theta \theta f(\theta|s_i,s_j)\,d\theta}_{\triangleq \mathrm{E}[\theta|s_i,s_j]}, \tag{C4}$$

where the conditional density of θ given the signals s_i and s_j is

$$f(\theta|s_i,s_j) = \frac{\Pr[s_i,s_j|\theta]\, f(\theta)}{\int_\theta \Pr[s_i,s_j|\theta]\, f(\theta)\,d\theta},$$

Using the conditional probability distribution of the signals from (3.2), we obtain

$$f(\theta|s_i=0,s_j=0) = 3(1-\theta)^2, \qquad f(\theta|s_i=1,s_j=1) = 3\theta^2, \tag{C5}$$

$$f(\theta|s_i=0,s_j=1) = f(\theta|s_i=1,s_j=0) = 6\theta(1-\theta). \tag{C6}$$

Substituting (C5) and (C6) into (C4) yields the decision maker's optimal actions

$$y^F_{0,0} = b_D + \underbrace{\frac{1}{4}}_{\mathrm{E}[\theta|0,0]}, \qquad y^F_{0,1} = y^F_{1,0} = b_D + \underbrace{\frac{1}{2}}_{\mathrm{E}[\theta|0,1]}, \qquad y^F_{1,1} = b_D + \underbrace{\frac{3}{4}}_{\mathrm{E}[\theta|1,1]}, \tag{C7}$$

as claimed.■

Proof of Proposition 3.2. *(i) One Expert.* Consider A_1's incentive to report truthfully his private signal. A_1's expected utility from reporting $m_1 = s_1$ is higher than his expected utility from reporting $m_1 = 1 - s_1$ if and only if

$$\int_\theta -\left(y^F_{s_1} - \theta - b_1\right)^2 f\left(s_1|\theta\right) d\theta \geqslant \int_\theta -\left(y^F_{1-s_1} - \theta - b_1\right)^2 f\left(s_1|\theta\right) d\theta, \qquad \forall s_1 \in \mathcal{S}, b_1 \in \mathcal{B},$$

which substituting $f\left(s_1|\theta\right) = f\left(\theta|s_1\right)\Pr\left[s_1\right]$ by Bayes' rule and integrating yields

$$-\left(y^F_{s_1} - \mathrm{E}\left[\theta|s_1\right] - b_1\right)^2 \Pr\left[s_1\right] \geqslant -\left(y^F_{1-s_1} - \mathrm{E}\left[\theta|s_1\right] - b_1\right)^2 \Pr\left[s_1\right].$$

Using D's best response from (C1) and rearranging terms, we obtain

$$\left(b_D - b_1\right)^2 \Pr\left[s_1\right] \leqslant (b_D + \underbrace{\mathrm{E}\left[\theta|1-s_1\right] - \mathrm{E}\left[\theta|s_1\right]}_{\triangleq \Delta y^F(s_1)} - b_1)^2 \Pr\left[s_1\right], \tag{C8}$$

Expanding squares and rearranging terms, (C8) further simplifies to

$$\left(b_1 - b_D\right)\Delta y^F\left(s_1\right) \leqslant \frac{\Delta y^F\left(s_1\right)^2}{2}, \tag{C9}$$

where we have used the fact that $\Pr\left[s_1\right] = \frac{1}{2}, \forall s_1 \in \mathcal{S}$. Solving (C9) jointly with D's optimal actions from Lemma 3.1, it is immediate to verify that when $s_1 = 0$ truth-telling by A_1 requires

$$b_1 - b_D \leqslant \frac{1}{6},$$

while, when he observes $s_1 = 1$, truth-telling condition is

$$b_1 - b_D \geqslant -\frac{1}{6},$$

where $b_1 \in \mathcal{B}$. The result follows immediately.

(ii) Two experts. Without loss of generality, we focus on A_i's incentive to report truthfully his signal, because experts are ex-ante symmetric. A_i's expected utility from reporting $m_i = s_i$ is higher than his expected utility from reporting a false message $m_i = 1 - s_i$ if and only if

$$\sum_{s_j \in \mathcal{S}} \int_\theta -\left(y^F_{s_i,s_j} - \theta - b_i\right)^2 f\left(s_j,\theta|s_i\right) d\theta \geqslant \sum_{s_j \in \mathcal{S}} \int_\theta -\left(y^F_{1-s_i,s_j} - \theta - b_i\right)^2 f\left(s_j,\theta|s_i\right) d\theta, \tag{C10}$$

which, substituting $f\left(s_j,\theta|s_i\right) = f\left(\theta|s_i,s_j\right)\Pr\left[s_j|s_i\right]$ by Bayes' rule and following the same steps as we did above, simplifies to

$$(b_i - b_D) \sum_{s_j \in \mathcal{S}} \Pr[s_j|s_i] \, \Delta y^F(s_i, s_j) \leqslant \sum_{s_j \in \mathcal{S}} \Pr[s_j|s_i] \, \frac{\Delta y^F(s_i, s_j)^2}{2}. \tag{C11}$$

In order to compute $\Pr[s_j|s_i]$, notice first that conditional probability distribution of the signals can be written as follows

$$\Pr(s_i, s_j|\theta) = \frac{f(s_i, s_j, \theta)}{f(\theta)}. \tag{C12}$$

Then, using (C12) together with the fact that $f(\theta) = 1$, we obtain

$$\Pr[s_j|s_i] = \int_\theta f(s_j, \theta|s_i) \, d\theta = \int_\theta \frac{f(s_i, s_j, \theta)}{\Pr(s_i)} d\theta = \Pr[s_i] \int_\theta \Pr(s_i, s_j|\theta) \, d\theta. \tag{C13}$$

Using (3.2) together with $\Pr[s_i] = \frac{1}{2}$, $s_i \in \mathcal{S}$, it can be easily verified that

$$\Pr[s_j = 0|s_i = 0] = \Pr[s_j = 1|s_i = 1] = \frac{2}{3}, \tag{C14}$$

$$\Pr[s_j = 1|s_i = 0] = \Pr[s_j = 0|s_i = 1] = \frac{1}{3}. \tag{C15}$$

Finally, substituting (C14), (C15) into (C11) and using D's optimal actions from Lemma **3.1**, when $s_i = 0$, truth-telling by A_i requires

$$b_i - b_D \leqslant \frac{1}{8},$$

while, when $s_i = 1$, truth-telling by A_i requires

$$b_i - b_D \geqslant -\frac{1}{8}.$$

where $b_i \in \mathcal{B}, i = 1, 2$. The result follows immediately. ■

Proof of Lemma 3.3. *(i) One Expert.* In a semi-revealing equilibrium, D's maximization problem after receiving $m_1 \in \mathcal{M}$ is

$$\begin{aligned} y^S_{m_1} &= \arg\max_{y \in \mathbb{R}} \int_\theta -(y - \theta - b_D)^2 f(\theta|m_1) \, d\theta, \\ &= b_D + \underbrace{\sum_{s_1 \in \mathcal{S}} \Pr[s_1|m_1] \, \mathrm{E}[\theta|s_1]}_{\triangleq \mathrm{E}_\nu[\theta|m_1]}, \qquad \forall m_1 \in \mathcal{M}. \end{aligned} \tag{C16}$$

Bayes rule implies that D's posterior beliefs about s_1 can be written as follows

$$\Pr[s_1|m_1] = \frac{\Pr[m_1|s_1] \Pr[s_1]}{\sum_{s_1 \in \mathcal{S}} \Pr[m_1|s_1] \Pr[s_1]}. \tag{C17}$$

When D receives $m_1 = 1$, her posteriors beliefs are

$$\Pr[s_1 = 1|m_1 = 1] = \frac{1}{2-\nu}, \qquad \Pr[s_1 = 0|m_1 = 1] = \frac{1-\nu}{2-\nu}, \tag{C18}$$

while when she receives $m_1 = 0$, her posterior beliefs are

$$\Pr[s_1 = 1|m_1 = 0] = 0, \qquad \Pr[s_1 = 0|m_1 = 0] = 1. \tag{C19}$$

Substituting the posterior beliefs (C18) and (C19) into (C16), and using the conditional expectations $\mathrm{E}[\theta|s_i, s_j]$ from the proof of Lemma 3.2, we have

$$y_0^S = b_D + \underbrace{\frac{1}{3}}_{\mathrm{E}_\nu[\theta|m_1=0]}, \qquad y_0^S = b_D + \underbrace{\frac{3-\nu}{3(2-\nu)}}_{\mathrm{E}_\nu[\theta|m_1=1]},$$

as claimed.

(ii) Two Experts. From (3.8) we know that D's optimal action after receiving m_i and m_j is

$$y_{m_i,m_j}^S = b_D + \underbrace{\sum_{(s_i,s_j)\in\mathcal{S}^2} \mathrm{E}[\theta|s_i, s_j] \Pr[s_i, s_j|m_i, m_j]}_{\mathrm{E}_\nu[\theta|m_i,m_j]}, \qquad \forall (m_i, m_j) \in \mathcal{M}^2. \tag{C20}$$

Bayes' rule implies that D's posterior beliefs about the signals can be written as follows

$$\Pr[s_i, s_j|m_i, m_j] = \frac{\Pr[m_i, m_j|s_i, s_j] \Pr[s_i, s_j]}{\sum_{(s_i,s_j)\in\mathcal{S}^2} \Pr[m_i, m_j|s_i, s_j] \Pr[s_i, s_j]}. \tag{C21}$$

Given that the extreme expert's babbling strategy is to report $m_E = 1$, when D receives $(m_i, m_j) = (1, 1)$, her posterior beliefs about (s_i, s_j) are

$$\Pr[s_i = 1, s_j = 1|1, 1] = \frac{1}{\nu^2 - 3\nu + 3}, \qquad \Pr[s_i = 0, s_j = 0|1, 1] = \frac{(1-\nu)^2}{\nu^2 - 3\nu + 3}, \tag{C22}$$

and

$$\Pr[s_i = 0, s_j = 1|1, 1] = \Pr[s_i = 1, s_j = 0|1, 1] = \frac{1-\nu}{2(\nu^2 - 3\nu + 3)}. \tag{C23}$$

By the same token, when D receives $(m_i, m_j) = (0, 1)$, the posteriors are

$$\Pr[s_i = 0, s_j = 0|0, 1] = \frac{2(1-\nu)}{3-2\nu}, \qquad \Pr[s_i = 0, s_j = 1|0, 1] = \frac{1}{3-2\nu}, \tag{C24}$$

and zero, otherwise. Since the message space is binary, symmetric argument applies to the case where decision maker receives $(m_i, m_j) = (1, 0)$. Finally, when D receives $(m_i, m_j) = (0, 1)$ the posteriors are

$$\Pr[s_i = 0, s_j = 0|0, 0] = 1, \tag{C25}$$

and zero, otherwise. Next, we need to compute the joint probability of the signals. Notice that, Bayes rule implies that $\Pr[s_i, s_j]$ can be written as follows:

$$\Pr[s_i, s_j] = \Pr[s_j|s_i]\Pr[s_i]. \tag{C26}$$

Then substituting $\Pr[s_j|s_i]$ from equations (C14) and (C15) into (C26), it follows that

$$\Pr[s_i = 1, s_j = 1] = \Pr[s_i = 0, s_j = 0] = \frac{1}{3}, \tag{C27}$$

$$\Pr[s_i = 0, s_j = 1] = \Pr[s_i = 1, s_j = 0] = \frac{1}{6}. \tag{C28}$$

Finally, substituting $\mathrm{E}[\theta|s_1, s_2]$ from the proof of Lemma 3.1 and the joint probability of the signals (C27) and (C28) into (C20), it is immediate to verify that

$$y^S_{0,0} = b_D + \underbrace{\frac{1}{4}}_{\mathrm{E}_\nu[\theta|0,0]} \qquad y^S_{0,1} = y^S_{1,0} = b_D + \underbrace{\frac{2-\nu}{2(3-2\nu)}}_{\mathrm{E}_\nu[\theta|0,1]} \qquad y^S_{1,1} = b_D + \underbrace{\frac{\nu^2 - 4\nu + 6}{4\nu^2 - 12\nu + 12}}_{\mathrm{E}_\nu[\theta|1,1]},$$

as desired. ■

Proof of Proposition 3.4. *(i) One Expert.* Suppose that A_1 is moderate — i.e., such that $b_1 = b_M$. Then A_1 has an incentive to report truthfully if and only if

$$\int_\theta -\left(y^S_{m_1} - \theta - b_M\right)^2 f(s_1|\theta)\, d\theta \geqslant \int_\theta -\left(y^S_{1-m_1} - \theta - b_M\right)^2 f(s_1|\theta)\, d\theta,$$

which following the same steps as we did in the proof of Proposition 3.2 and rearranging terms, simplifies to

$$\Pr[s_1](b_M - b_D)\Delta y^S(m_1) \leqslant$$
$$\leqslant \Pr[s_1]\left\{\frac{\Delta y^S(m_1)^2}{2} + \Delta y^S(m_1)\left(\mathrm{E}_\nu[\theta|m_1] - \mathrm{E}[\theta|s_1]\right)\right\}, \tag{C29}$$

where $\Delta y^S(m_1) \triangleq y^S_{1-m_1} - y^S_{m_1}$. Now substituting the optimal actions from Lemma 3.1 and Lemma 3.3 into (C29), it follows that whenever $s_1 = 0$, truth-telling by the moderate expert requires

$$b_M - b_D \leqslant \alpha_1(\nu) \triangleq \frac{1}{2(6-3\nu)}.$$

Similarly, when $s_1 = 1$ is observed, truth-telling by the moderate expert requires

$$b_M - b_D \geqslant -\beta_1(\nu),$$

where

$$\beta_1(\nu) \triangleq \frac{3-2\nu}{2(6-3\nu)}.$$

Moreover, $\alpha_1(\nu)$ is increasing in ν and $\beta_1(\nu)$ is decreasing in ν — i.e.,

$$\frac{d}{d\nu}[\alpha_1(\nu)] = \frac{1}{6(2-\nu)^2} > 0, \qquad \frac{d}{d\nu}[\beta_1(\nu)] = -\frac{1}{6(2-\nu)^2} < 0,$$

as expected. To complete the proof, we need to check that the extreme expert has no incentive to report $m_1 = 0$ when his signal is $s_1 = 0$. Adopting the same logic used above, this required condition is

$$(b_E - b_D)\,\Delta y^S(0) > \frac{\Delta y^S(0)^2}{2} + \Delta y^S(0)\left(\mathrm{E}_\nu[\theta|0] - \mathrm{E}[\theta|0]\right). \tag{C30}$$

Substituting the optimal actions from Lemma 3.1 and Lemma 3.3 into (C30), whenever $s_1 = 0$, babbling condition required by the extreme expert is

$$b_E - b_D > \alpha_1(\nu).$$

Finally, when $s_1 = 1$, a rightward biased extreme expert has a strict incentive to report $m_1 = 1$.

(ii) Two Experts. Without loss of generality, we focus on A_i's incentives to disclose his private information, since experts have symmetric payoffs. Consider first that A_i is a moderate such that $b_i = b_M$. Given that A_j's type is his private information, A_i's incentive compatibility constraints writes as

$$\sum_{b_j\in\mathcal{B}} \Pr[b_j] \sum_{s_j\in\mathcal{S}} \int_\theta -\left(y^S_{m_i,m_j} - \theta - b_M\right)^2 f(s_j,\theta|s_i)\,d\theta \geqslant$$
$$\geqslant \sum_{b_j\in\mathcal{B}} \Pr[b_j] \sum_{s_j\in\mathcal{S}} \int_\theta -\left(y^S_{1-m_i,m_j} - \theta - b_M\right)^2 f(s_j,\theta|s_i)\,d\theta. \tag{C31}$$

Following the same steps as we did in the proof of Proposition 3.2, the above constraint can be rewritten as follows

$$(b_M - b_D)\sum_{b_j\in\mathcal{B}} \Pr[b_j] \sum_{s_j\in\{0,1\}} \Pr(s_j|s_i)\,\Delta y^S(m_i,m_j) \leqslant$$
$$\leqslant \sum_{b_j\in\mathcal{B}} \Pr[b_j] \sum_{s_j\in\{0,1\}} \Pr(s_j|s_i)\left\{\frac{\Delta y^S(m_i,m_j)^2}{2} + \Delta y^S(m_i,m_j)\left(\mathrm{E}_\nu[\theta|m_i,m_j] - \mathrm{E}[\theta|s_i,s_j]\right)\right\}. \tag{C32}$$

Using the optimal actions from Lemma 3.1 and Lemma 3.3 and $\Pr[s_j|s_i]$ from equations (C14) and (C15) , when $s_i = 0$, truth-telling by the moderate expert requires

$$b_M - b_D \leqslant \alpha_2(\nu) \triangleq \frac{5\nu^4 - 34\nu^3 + 84\nu^2 - 90\nu + 36}{8(3-2\nu)(\nu^2 - 3\nu + 3)(3\nu^2 - 8\nu + 6)}.$$

By the same token, when $s_i = 1$, truth-telling by the moderate expert requires

$$b_M - b_D \geqslant -\beta_2(\nu),$$

where

$$\beta_2(\nu) \triangleq \frac{5\nu^4 - 33\nu^3 + 80\nu^2 - 87\nu + 36}{8(3-2\nu)(\nu^2 - 3\nu + 3)(2-\nu)}.$$

Moreover, it can be shown that

$$\frac{d}{d\nu}[\alpha_2(\nu)] = \frac{30\nu^8 - 408\nu^7 + 2329\nu^6 - 7374\nu^5 + 14\,262\nu^4 - 17\,316\nu^3 + 12\,906\nu^2 - 5400\nu + 972}{8(3-2\nu)^2(\nu^2 - 3\nu + 3)^2(3\nu^2 - 8\nu + 6)^2} \geqslant 0,$$

and

$$\frac{d}{d\nu}\left[\beta_2\left(\nu\right)\right]=\frac{\nu^6+10\nu^5-112\nu^4+384\nu^3-627\nu^2+504\nu-162}{8\left(2-\nu\right)^2\left(3-2\nu\right)^2\left(\nu^2-3\nu+3\right)^2}<0.$$

To complete the proof, we need to check that extreme expert has no incentive to report $m_i = 0$ when his signal is $s_i = 0$. Adopting the same logic used above, A_i's expected utility from reporting $m_E = 1$ is higher than his expected utility when reporting truthfully $m_i = 0$ if

$$\left(b_M-b_D\right)\sum_{b_j\in\mathcal{B}}\Pr\left[b_j\right]\sum_{s_j\in\{0,1\}}\Pr\left(s_j|0\right)\Delta y^S\left(0,m_j\right)>$$
$$>\sum_{b_j\in\mathcal{B}}\Pr\left[b_j\right]\sum_{s_j\in\{0,1\}}\Pr\left(s_j|0\right)\frac{\Delta y^S\left(0,m_j\right)^2}{2}+\Delta y^S\left(0,m_j\right)\left(\mathrm{E}_\nu\left[\theta|0,m_j\right]-\mathrm{E}\left[\theta|0,s_j\right]\right).$$

Using the optimal actions from Lemma 3.1 and Lemma 3.3, when $s_1 = 1$, babbling by the rightward biased extreme expert requires

$$b_E-b_D>\alpha_2\left(\nu\right),$$

while when $s_1 = 1$, the rightward biased extreme expert has an incentive to report truthfully his signal. ■

Proofs of Propositions 3.5 and 3.6. We first compare D's expected utility from consulting one and two experts within each equilibrium. Let $\mathrm{E}U_i^F, i = 1, 2$, be the decision maker's ex-ante expected utility in a fully-revealing equilibrium. More precisely, in a fully-revealing equilibrium, D's expected profit from consulting one expert is

$$\mathrm{E}U_1^F\triangleq\int_\theta\sum_{s_1\in\mathcal{S}}-\left(y_{s_1}^F-\theta-b_D\right)\Pr\left[s_1|\theta\right]f\left(\theta\right)d\theta, \tag{C33}$$

which using (3.1) and using the results of Lemma 3.1 yields

$$\mathrm{E}U_1^F=-\frac{1}{18}. \tag{C34}$$

Similarly, in a fully-revealing equilibrium, D's expected profit from consulting two experts is

$$\mathrm{E}U_2^F\triangleq\int_\theta\sum_{(s_i,s_j)\in\mathcal{S}^2}-\left(y_{s_i,s_j}^F-\theta-b_D\right)\Pr\left[s_i,s_j|\theta\right]f\left(\theta\right)d\theta. \tag{C35}$$

Using (3.2) and using the optimal actions from Lemma 3.1, we have

$$\mathrm{E}U_2^F=-\frac{1}{24}. \tag{C36}$$

Comparing this with (C34),

$$\mathrm{E}U_2^F-\mathrm{E}U_1^F=\frac{1}{72}>0.$$

Therefore, in a fully-revealing equilibrium, D's ex-ante expected utility is higher with two experts. Now let $\mathrm{E}U_i^S, i = 1, 2$ be the decision maker's ex-ante expected utility in a semi-revealing equilibrium. More precisely, in a semi-revealing equilibrium D's ex-ante expected profit from consulting one expert is

$$\mathrm{E}U_1^S \triangleq \int_\theta \sum_{m_1 \in \mathcal{M}} -\left(y_{m_1}^S - \theta - b_D\right) \Pr\left[m_1|\theta\right] f\left(\theta\right) d\theta, \tag{C37}$$

where

$$\Pr\left[m_1|\theta\right] = \sum_{s_1 \in \mathcal{S}} \Pr\left[m_1|s_1\right] \Pr\left[s_1|\theta\right]. \tag{C38}$$

Substituting the conditional probability distribution of s_1 from (3.1) and the corresponding prior beliefs into (C38), we have

$$\Pr\left[m_1 = 1|\theta\right] = \theta + (1-\nu)(1-\theta) \qquad \text{and} \qquad \Pr\left[m_1 = 1|\theta\right] = \nu(1-\theta).$$

Hence,

$$\mathrm{E}U_1^S = -\frac{3-2\nu}{18(2-\nu)}, \tag{C39}$$

where we have used the optimal actions from Lemma 3.3. Similarly, in a semi-revealing equilibrium D's ex-ante expected profit from consulting two experts is

$$\mathrm{E}U_2^S \triangleq \int_\theta \sum_{(m_i,m_j) \in \mathcal{M}^2} -\left(y_{m_i,m_j}^S - \theta - b_D\right) \Pr\left[m_i, m_j|\theta\right] f\left(\theta\right) d\theta,$$

where

$$\Pr\left[m_i, m_j|\theta\right] = \sum_{(s_i,s_j) \in \mathcal{S}^2} \Pr\left[m_i, m_j|s_i, s_j\right] \Pr\left[s_i, s_j|\theta\right]. \tag{C40}$$

Substituting the conditional probability distribution of (s_i, s_j) from (3.2) and the corresponding prior probabilities into (C40), we have

$$\Pr\left[m_i = 1, m_j = 1|\theta\right] = \left(1 - \nu(1-\theta)\right)^2, \qquad \Pr\left[m_i = 0, m_j = 0|\theta\right] = \nu^2(1-\theta)^2$$
$$\Pr\left[m_i = 0, m_j = 0|\theta\right] = \Pr\left[m_i = 1, m_j = 0|\theta\right] = \nu(1-\theta)\left(1-\nu(1-\theta)\right),$$

Then using the optimal actions from Lemma 3.3, we obtain

$$\mathrm{E}U_2^S = -\frac{36(1-\nu)^2 + 13\nu^2(1-\nu) + 2\nu^2}{48(3-2\nu)(\nu^2 - 3\nu + 3)}. \tag{C41}$$

Comparing (C39) and (C41),

$$\mathrm{E}U_2^S - \mathrm{E}U_1^S = \frac{(3-\nu)\left(\nu^3 + 6\nu(2-\nu)(1-\nu)\right)}{144(2-\nu)(3-2\nu)(\nu^2 - 3\nu + 3)},$$

which is positive. Therefore, in a semi–revealing equilibrium, D's ex-ante expected utility is higher with two experts. Finally, we compare D's ex-ante expected utility in a semi-revealing equilibrium with two experts with her ex-ante expected utility when she consults one expert who reports truthfully his signal. Direct comparison of (C34) and (C41) yields

$$\mathrm{E}U_2^S - \mathrm{E}U_1^F = \frac{1}{144} \frac{96\nu - 81\nu^2 + 23\nu^3 - 36}{(3-2\nu)(\nu^2 - 3\nu + 3)}. \tag{C42}$$

Since the denominator is positive, the sign of (C42) depends on the sign of

$$\mu(\nu) \triangleq 96\nu - 81\nu^2 + 23\nu^3 - 36.$$

Notice that

$$\mu(0) = -36 < 0,$$
$$\mu(1) = 2 > 0.$$

Moreover,

$$\frac{d\mu(\nu)}{d\nu} = 3\left(23\nu^2 - 54\nu + 32\right) > 0.$$

Hence, by mean value theorem there exists a unique $\widetilde{\nu} \triangleq 0.74$ such that $\mu(\nu) < 0$ (so that the decision maker's ex-ante expected utility is higher with one accurate expert) if and only if $\nu < \widetilde{\nu}$. ■

References

[1] Austen-Smith, D. (1993). "Interested Experts and Policy Advice: Multiple Referrals under Open Rule." *Games and Economics Behaviour*, 5, 3-43.

[2] Battaglini, M. (2002). "Multiple Referrals and Multidimensional Cheap Talk." *Econometrica*, 70, 13791401.

[3] Bénabou, R., and G. Laroque (1992). "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility." *Quarterly Journal of Economics*, 107, 921-958.

[4] Crawford, V.P., and J. Sobel (1982). "Strategic Information Transmission." *Econometrica*, 50, 1431-1451.

[5] Dimitrakas, V., and Y. Sarafidis (2005). "Advice from an Expert with Unknown Motives." Mimeo, INSEAD.

[6] Galeotti, A., C. Chiglino and F. Squintani (2013). "Strategic Information Transmission in Networks." *Journal of Economic Theory*, 148, 1751-1769.

[7] Gick, W. (2006). "Two Experts are Better than One: Multi-Sender Cheap Talk under Simultaneous Disclosure." Mimeo, Dartmouth University.

[8] Gentzkow, M. and J.M. Shapiro (2006). "Media Bias and Reputation." *Journal of Political Economy*, 114, 280-316.

[9] Gilligan, T.W., and K. Krehbiel (1989). "Asymmetric Information and Legislative Rules with a Heterogeneous Committee." *American Journal of Political Science*, 33, 459-490.

[10] Kawamura, K. (2011). "A Model of Public Consultation: Why is Binary Communication so Common." *The Economic Journal*, 121, 819-842.

[11] Krishna, V., and J. Morgan (2001). "A model of Expertise." *Quarterly Journal of Economics*, 116, 747-775.

[12] Morgan, J., and P.C. Stocken (2003). "An Analysis of Stock Recommendations." *RAND Journal of Economics*, 34, 183-203.

[13] Morgan, J., and P.C. Stocken (2008). "Information Aggregation in Polls." *American Economic Review*, 98, 864-896.

[14] Morris, S. (2001). "Political Correctness." *Journal of Political Economy*, 109, 231-265.

[15] Li, M. (2004). "To Disclose or Not to Disclose: Cheap Talk with Uncertain Biases." Mimeo, Concordia University.

[16] Li, M. (2008). "Two(talking) heads are not better than one." *Economics Bulletin*, 3, 1-8.

[17] Li, M. and K. Madarász (2008). "When mandatory disclosure hurts: Expert advice and conflicting interests." *Journal of Economic Theory*, 139, 47-74.

[18] Ottaviani, M., and P.N. Sørensen (2006). "Reputational Cheap Talk." *RAND Journal of Economics*, 37, 155-175.

[19] Sobel, J. (1985). "A Theory of Credibility." *Review of Economic Studies*, 52, 557-573.

[20] Sobel, J. (2013). "Giving and Receiving Advice." Advances in Economics and Econometrics: Tenth World Congress, Cambridge University Press, Chapter 10.