Essays in Industrial Organisation

Marta Troya Martinez

University of Oxford

May 2012

A thesis submitted in fulfilment of the requirements for the degree of

D.Phil. in Economics.

Abstract

The first chapter considers a common agency model where two competition authorities share information about a firm under investigation. It shows that information-sharing can sometimes be welfare detrimental unless the authorities coordinate their enforcement policies as well as share information. The reason behind is that the authorities may have different leniency levels and the firm may decide to provide less precise information to one in an attempt to appeal the other. Furthermore it shows that the authorities may want to distort their policies in order to prevent the firm from obscuring the information it provides.

The second chapter studies the seller's incentives to provide misleading advice about complex goods such as consumer electronics, banking or phone services. It shows how the incentives to give biased and imprecise advice are affected by the possibility of ex-post litigation, when a court or consumer protection authority investigates how biased the advice is and penalise accordingly. It finds that a more biased advice will also be less precise, thus, a stricter punishment for deceiving consumers also increases precision.

The third chapter analyses the impact of trade credit on a relational contract between two vertically related firms. The firms operate in an environment with unobservable shocks, like a developing country or a black market, which create moral hazard in the repayment decision. It shows that the quantity sold in the market will be distorted downwards in order to curb the constrained firm's incentives to steal the credit and derives the optimal repayment scheme.

Acknowledgements

I would like to begin by thanking my supervisor, David Myatt, for patiently listening and rightfully challenging the ideas contained in this thesis. I am also indebted to Paul Klemperer and John Vickers for their supervision and advice. I feel privileged to have benefitted from the sharp intellectual insights and invaluable opinions from them three. At Oxford, I have also benefitted from helpful comments and discussions with Meg Meyer, Jim Malcomson and Peter Eso. Finally, I am grateful to my earlier teachers Matthew Ellman and Doh-Shin Jeon for encouraging me to pursue postgraduate study and research.

I am grateful to the many seminar and conferences audiences where I have presented the different chapters of my thesis. In particular, I strongly appreciate helpful comments from Amelia Fletcher, Filip Lachowski, Lakshmi Iyer, Vicente Cunat, Natalia Fabra, Rosa Ferrer, Maria Goltsman, Martin Hellwig, Marc Moller and Bent Nielsen. I have been lucky to meet extraordinary people in Oxford who have become my dear friends. Thanks to Martina, Babak, Nicola, Paolo, Mauro, Alan, Maria Ana, Oswaldo and Andres for sharing the ups and downs of the dissertation.

The financial support of the Bank of Spain and Ramon Areces Foundation is gratefully acknowledged. I would also like to express my gratitude to a number of grants from St. Antony's College as well as the Department of Economics at Oxford.

Foremost, I am ever thankful to my family for their love: my parents, who taught me the precious qualities of hardwork and perseverance; my brother, David; my proud grandmothers and late grandfathers. The most special thank you is to Misha. Without his encouragement and unconditional support I would have not been able to do it, thank you.

This thesis is dedicated to M&M.

Declaration

I declare that this thesis represents my own work, and that (with the exception of Chapter 2, which is based on my M.Phil. thesis) none of it has already been accepted. or is currently being submitted, for any degree or diploma or certificate or other qualification in this University or elsewhere.

Signature   Date

Word count (366 words per page x 141 pages): 51606 words

Preface

Chapter 2 has been presented at the Annual Conference of the Royal Economic Society (Guildford, 2009), the Annual Meeting of the European Economic Association (Barcelona, 2009), the Industrial Organisation Workshop (Lecce, 2009), the Annual Competition and Regulation European Summer School and Conference (Chania, 2009), the International Industrial Organization Conference (Vancouver, 2010), the Economist Group Meeting of the Office of Fair Trading (London, 2009), the Gorman Workshop (Oxford, 2010) and the Center for Economic Policy Research Meeting in Applied Industrial Organization (Toulouse, 2010). This version of the paper was circulated as a Discussion Paper #496 by the Department of Economics of the University of Oxford.

Chapter 3 has been presented at the workshop "Advances in Industrial Organisation" (Vienna, 2010), the International Industrial Organization Conference (Vancouver, 2010), the Annual Competition and Regulation European Summer School and Conference (Chania, 2010), Jornadas de Economia Industrial (Madrid, 2010), Royal Economic Society Meetings (Cambridge, 2012) and the seminars at the Max Planck Institute for Research on Collective Goods (Bonn, 2010) and IIES (Stockholm, 2010). This paper has been previously circulated under the title "Bias, Noise and Litigation".

Chapter 4 has been presented at the Gorman Workshop (Oxford, 2011), the Center for the Study of African Economies Conference (Oxford, 2011), the International Industrial Organization Conferences (Boston, 2011), Jornadas de Economia Industrial (Valencia, 2011) and the Royal Economic Society Meetings (Cambridge, 2012).

Contents

[1] The third possibility, namely, that only quality of the match is observed (or can be used in court) does not help here: since all consumers buy for all signal realizations, the ex post distribution of quality is the same as the prior one. This would be different if consumers buy only above a certain threshold as in Section 3.4.

Chapter 1

Introduction

This thesis is comprised of three chapters in the fields of theoretical industrial organization and consumer policy. Each chapter stands alone. Chapter 3 was coauthored with Mikhail Drugov (Department of Economics, University of Carlos III, Madrid). Chapter 2 and 4 are my own work.

1.1 Information-Sharing Between Competition Authorities

This chapter explores the consequences of an information-sharing agreement between competition authorities using a signal jamming model. These agreements have recently proliferated due to the increasing number of multi-jurisdictional cases (like the one introduced in the EC Regulation 1/2003). I consider a situation where two authorities are reviewing the same merger. Each authority clears the merger if the information sent by the firm suggests that the expected welfare in its country will be enhanced and the firm can secretly manipulate the precision with which it transmits this information (i.e., choose the variance of the signal). The authorities differ in their leniency towards the merger and I focus on the cases where the authorities disagree about the decision. Under no information-sharing, the firm chooses an extreme level of precision: very high (low) for the most (least) lenient authority; interestingly the least lenient authority may benefit from committing to a more lenient policy to obtain more precise information. Under information-sharing, the firm has to send the same information and thus chooses the same precision for both authorities. The choice of precision depends on the level of cooperation in the decision-making between the countries. If the authorities exert their veto power, the firm gives very imprecise information. If the authorities cooperate in the decision-making, the choice of precision may be non-monotonic and intermediate levels of precision are chosen, making information-sharing potentially desirable.

These findings can be applied to a variety of situations like a politician talking to different audiences.

1.2 Vague Lies: How to Advise Consumers When They Complain

Chapter 3 studies the seller's incentives to provide misleading advice about complex goods such as consumer electronics, banking or phone services. This is a relevant policy issue because these products systematically top the list of consumer's complaints worldwide (see, for example, the yearly lists elaborated by the Office of Fair Trading in the UK or the Federal Trade Commission in the US). We look at how the incentives to give biased and imprecise advice are affected by the possibility of ex-post litigation, when a court or consumer protection authority might infer how biased the advice is and penalize accordingly. Hence, unlike the existing research, we endogenize the cost of bias and imprecision by explicitly modelling how they are affected by the court's decision. In particular, biasing the advice upwards increases the revenues but also the cost of litigation by making it easier for the authority to convict the seller. Making the advice less precise does not affect the revenues in equilibrium but interferes with the authority's inference and affects the expected fine in a non-monotonic way. In particular, making the advice less precise makes it harder to convict the seller but increases the expected fine when the seller is found guilty. We find that, in the equilibrium, biasing the advice and making it noisier are complements; in particular, a higher buyers' heterogeneity, a stricter standard of proof employed by the authority and a larger share of credulous consumers make the advice more biased and less precise.

1.3 Vertical Relational Contracts and Trade Credit

Chapter 4 provides a first analysis on the impact of supplier trade credit on the inter-firm relational contracts. Both trade credit and relational contracts (which are enforced thanks to the gains from future trade) are widely used by vertically related firms in developing countries where there is no rule of law and an underdeveloped financial market. I also consider another stylized fact of these countries: the existence of adverse shocks that make the firms unable to honour their agreements, such as a warehouse being robbed. These adverse shocks may not be observed (e.g. non-prosecuted criminal offences) by the other contracting party which creates moral hazard, as the downstream firms may not return the trade credit and pretend it was a shock. This chapter derives and explores the form of these relational contracts. As compared to the standard relational incentive contracts literature, trade credit imposes a constraint on the available punishments. As a result, the firms do not maximize the total profits subject to the self-enforceability constraint, which translates, for instance, in the existence of double marginalization (i.e., price above marginal cost) even despite the use of two-part tariffs.

These findings can be applied to other less studied markets where contracts cannot be written and traders are subject to unobservable shocks like in black markets such as drug trade.

Chapter 2

Information-Sharing Between Competition Authorities

2.1 Introduction

> "Most officials believe that the issue of confidentiality is the chief limitation of enforcement cooperation agreements and hence it is submitted that the majority of effort should be concentrated on overcoming this particular obstruction to effective cooperation between antitrust agencies."
>
> Marsden and Whelan (2005), p.24.

With globalization, competition has had an increasingly international dimension. A clear example is the existence of international cartels, such as the vitamins cartel which took place between January 1990 and February 1999, as well as the increasing number of mergers that involve more than one jurisdiction.[1] Both examples suggest the need to enhance cooperation between competition agencies.[2]

[1] UNCTAD (2000) shows that the share of cross-border mergers increased to 78% of the world FDI in the late 1990's.

[2] For instance, the U.S. agencies have about 120 mergers' notifications to foreign governments in a two-year period. In some 64 of them there is additional contact with the foreign agency where publicly-available information is exchanged. In about 40 cases the agencies engage in some level of cooperation, but confidential business information is only exchanged if a waiver is granted, which happens in some 16 cases. See OECD (2003).

There has been a proliferation of multilateral platforms where various policy issues are discussed such as the ICN (International Competition Network) and the OECD (Organization for Economic Cooperation and Development). These policy forums produce useful instruments[3]; nonetheless, their usefulness is constrained by their non-binding aspect. Similarly, many bilateral agreements have emerged. As pointed out by the quotation above, one of the main limitations of these agreements is the impossibility of exchanging confidential information between competition authorities. One of the most prominent examples of this type of agreement is the one between the E.U. and the U.S.[4]

The issue of exchange of confidential information has been raised in many occasions. In 2002, Advanced Micro Devices (AMD) requested court documents, gathered during a U.S. antitrust case against Intel several years before, to be transferred to the European Commission as support for a complaint against Intel. AMD believed that many of the issues in the U.S. case were similar to the questions under investigation by the E.C. This request was made under 28 U.S.C. § 1782, which allows a district court to order the production of documents "for use in a proceeding in a foreign or international tribunal" upon request by "any interested person". This request was first denied by the district court and then reversed by the Ninth Circuit Court of Appeals.[5]

A minority of competition policy agreements expressly provide for the exchange of confidential information. Within this group we find the bilateral agreement between the U.S. and Australia[6], the trilateral agreement between Iceland, Norway and Denmark[7] and, more recently, the agreement between competition authorities

---

[3] Such as those introduced by the OECD to deal with exchange of information in hard-core cartels: http://www.oecd.org/dataoecd/1/33/35590548.pdf.

[4] Agreement Regarding the Application of their Competition Laws, 23 Sept. 1991, 4 CMLR 823, 30 ILM 1487 and the E.U.-U.S. Positive Comity Agreement, 4th June 1998, (1998) OJ L173/28, (1999) 4 CMLR 502.

[5] In the Supreme court opinion two justices concurred with the judgment, one dissented, and one took no part in it. Intel Corp. v. Advanced Micro Devices, Inc., No. 02-572, available at: www.supremecourtus.gov/opinions/03pdf/02-572.pdf

[6] Agreement on Mutual Antitrust Enforcement Assistance, available at www.apeccp.org.tw/doc/USA/Cooperation/usaus7.htm

[7] www.globalcompetitionforum.org/regions/europe/Denmark/Agreemen1.pdf

of the European Member States[8].

The goal of this paper is to explore the information-sharing agreement's impact on the incentives of the firm to provide precise information when undertaking a multinational merger.

We start by analyzing the case where the firm ("he") deals with a single authority ("she") or, equivalently, when the authorities do not share information. The welfare implications of the merger correspond to an underlying real-valued state variable. In order to have the merger assessed, the multinational firm has to provide information about this variable to the competition authority. Neither the firm nor the authority have private information[9] concerning this variable. For instance, the firm may be giving information about the merger's market effect or the merger's effect on third parties, for which the firm may not have more information than the competition authority.[10] The sole mechanism available to transmit this information is through a noisy signal[11]. The authority only observes a random realization of the signal. The signal is unbiased, that is, the firm cannot misrepresent the merger's welfare implications on average[12]. However, the firm can choose to secretly manipulate the precision with which he transmits this information to the authority at no cost. We model this by allowing the firm to choose the variance[13] of the signal, for instance, the reports can be drafted in a very imprecise way and this choice is secret because the authority does not actually know how precise is the information that the firm has. Also, the firm can overload the authority with too many documents, knowing that the authority has a limited amount of time to review them. Indeed, a larger quantity of information does not necessarily imply more precision as some of it may

[8] E.C. Regulation 1/2003, available at http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri =OJ:L:2003:001:0001:0025:EN:PDF.

[9] In Section 2.5, we discuss how the results change if we allow the firm to have private information about this variable.

[10] We rule out the case where the firm's profit may be negatively correlated with the total welfare to acknowledge the possibility of efficiency effects following the merger.

[11] The signal has a normal additive structure.

[12] Adding bias to the signal involes lying. This is very costly for a firm that is in the middle of an investigation as there are high chances to be discovered in which case it is exposed to a high penalty.

[13] Noise, variance and lack of precision are interchangeable.

be irrelevant.

The policy decision is binary (to clear the merger or not) and depends only on the realization of the signal. The authority adopts a cut-off rule whereby, if the realization of the signal is above some threshold, she clears the merger and she blocks it otherwise.

We find that the optimal variance chosen by the firm does not depend on how good or bad the average merger is (i.e. how far the average merger is above or below the policy threshold) but rather on whether it is good or bad (i.e. above or below the threshold). In particular, if the average merger is bad, the firm chooses the risky strategy of the largest variance to have more chances to be thought good. By contrast, if the average merger is good, the firm chooses the lowest variance to increase the likelihood of obtaining a high realization of the signal. Furthermore, this strategy does not change depending on whether the authority can commit to it ex-ante. We also establish the optimal response (i.e. policy threshold) of the authority to the firm's behavior. If the authority cannot commit to a policy ex-ante, the optimal threshold when the average merger is welfare detrimental (enhancing) is stricter (more lenient) than the full information threshold. When there is uncertainty about the undesirability of the merger, the authority's ability to commit (for instance, by issuing detailed guidelines) makes her set a more lenient threshold in order to induce the firm to provide more precise information. Otherwise, the authority sets the ex-post optimal threshold of the no-commitment case[14].

Then, we analyze an international merger, where a multinational wants to undertake a merger that needs the approval of two countries at the same time. We consider the case where one authority is more lenient than the other (i.e. their policy thresholds differ), either because their tolerance to mergers differ or because the same underlying state variable has different welfare implications in the two countries. Information-sharing means that the authorities receive the same realization of

[14] In comparison with no sharing information but each authority having veto power over the decision.

the signal[15] sent by the multinational, so the multinational is left to make a unique choice of variance for both countries. We explore the impact of the information-sharing regime on the incentives of the firm to provide precise information[16].

When the authorities share information, we find that the agreement only has an impact on the firm's behavior when the average merger is "conflicting", that is, it is good for one country but bad for the other. This impact will depend on the particular rule that the authorities use to deal with the case of disagreement (i.e. when the realization of the signal lies between their thresholds). We consider different possible rules to deal with disagreement.

If the authorities have veto power (i.e. each country can unilaterally block the merger), then information-sharing induces the firm to send a very imprecise signal, making the more lenient country strictly worse off.

Another possibility is for cooperation to go beyond the information-sharing stage, by having some informal bargaining/persuasion process (not explicitly modelled in this paper) taking place between the competition authorities. We model the outcome of the bargaining process as the merger being approved with some probability in case of disagreement.

First, we take this probability as being fixed and exogenous, which may be interpreted as the authorities' relative bargaining power. The choice of the variance in this case may be non-monotonic in the expected merger, and intermediate values of variance may be chosen in equilibrium. When the average merger is relatively bad for the less lenient country, the firm chooses either low variance or high variance. In particular, if the countries have equal bargaining power and the distance between authorities' policies is very large (or the lowest variance is low enough), the firm first chooses high variance in order to maximize the chances to be above the more lenient threshold (where the merger is cleared for sure) and, as the average merger improves,

[15] We do not consider the case where each authority receives a different realization from the same signal to remove the obvious benefits that come from having more information.

[16] We do not consider how the policy thresholds change following the agreement. The policy thresholds are the same for the national and the international mergers (which are a minority), so the authorities may not want to be strategic in their decision as this may have a large impact on their national welfare.

he switches to low variance in order to maximize the chances of being cleared at least with some probability. As the average merger gets less welfare detrimental for the stricter country, it becomes safer to play a riskier strategy by increasing the variance to an intermediate level. Furthermore, this variance decreases as the average merger approaches to the policy threshold of the stricter country. This is because the chances of having the merger cleared are already high and, by reducing the variance, the probability of having it blocked by both authorities is decreased.

We also consider the case where the probability of clearing the merger depends on the particular realization of the signal. This captures the idea that the less lenient country is more willing to clear the merger if the realization falls near its threshold as compared to when it falls very far from it. In particular, we consider the case where this probability is an odd increasing function of the signal realization. From the point of view of the firm, it is as if there were a unique but uncertain threshold. Because of the rotational symmetry of the function, the firm considers the expected threshold and behaves as in the case of one authority with respect to this threshold.

To sum up, if authorities exert their veto power, then sharing information is a bad idea because the firm will send very imprecise information. Further cooperation modifies the firm's payoff structure in such a way that the firm makes a greater use of intermediate and lower levels of noise (as compared to the veto power case) and, therefore, it can be potentially good if the lenient country is not always the same one.

2.1.1 Related literature

This is a signal-jamming model, like the ones used in the career concerns literature[17], where the firm jams the signal, not by manipulating its mean but by changing its variance, and hence its information content.

The choice of variability as a strategic variable has been considered in a large variety of setups. For instance, Anderson and Cabral (2007) analyze a model of

[17] See Holmstrom (1982, 1999) and Dewatripont et al. (1999).

R&D races where the two contestants, taking as given the level of R&D expenditure, need to choose a level of risk. Tsetlin et al. (2004) consider the choice of variability of the performance distribution in a multi-round contest. In the compensation literature, Gaba and Kalra (1999) introduce the level of dispersion of the probability distribution of sales as a choice variable besides the level of effort. The general conclusion of all these papers is that the players that are at a disadvantage tend to optimally choose more risky strategies than those who are in a favorable position. We also find this "gambling for resurrection" behavior in our national merger framework. However, this behavior may disappear in the multinational merger case if authorities decide to cooperate both in the decision making as well as at the information-sharing stage.

Johnson and Myatt (2006) also consider the incentives of a firm to provide its potential customers with more or less precise information. They show how the seller's supply of information affects the shape of the distribution of buyers' expected valuations and hence generates rotations of the demand curve. These rotations generate a convexity in the monopolist's profits, which explains the optimality of the monopolist's extreme choices of variance in Lewis and Sappington (1994). Contrary to Johnson and Myatt (2006), in our model there is another strategic player: the competition authority. In particular, our firm will choose a level of precision in order to maximize the chances that the realization of the signal falls above the policy threshold level chosen by the authority. On the other hand, the authority will choose the threshold in order to maximize the (ex-post or ex-ante) expected welfare. In the Johnson and Myatt paper, the monopolist is choosing both the precision and the threshold (i.e. the price) in order to maximize the ex-post expected profits.

As far as we are aware, another distinctive feature of our model with respect to the previous work is that the noise (and hence, the distribution of the signal) is not observed by the receiver. In our framework, observability of the noise renders the problem trivial because the competition authority could simply condition the policy threshold on the noise and block any merger with an imprecise report.

Finally, this paper considers the consequences of information-sharing between receivers/principals and hence it is related to the literature that compares private with public communication. However, comparisons with this literature are difficult because, so far, its focus has been on frameworks of mechanism design (see, for instance, Calzolari and Pavan (2006), Maier and Ottaviani (2009)), cheap talk (see Farrell and Gibbons(1989)), and signalling (see Gertner, Gibbons and Scharfstein (1988) and Spiegel and Spulber (1997)).

The paper proceeds as follows. Section 2.2 introduces the model. Section 2.3 studies the national merger. This section first presents the results for the no-commitment case and then highlights the differences resulting from the authority's ability to commit. Section 2.4 introduces and analyses the multinational merger setup. Section 2.5 discusses the consequences of the firm having private information. Finally, Section 2.6 concludes. All the proofs can be found in the Appendix.

2.2 The model

We consider a multinational firm (M) proposing a merger that must be cleared under the competition law of the country. The welfare consequences of the merger can be summarized in the real-valued state variable θ, with support on $[-\infty, +\infty]$. The mergers with $\theta > 0$ are welfare enhancing while the ones with $\theta < 0$ are welfare detrimental. Furthermore, the higher the θ, the more desirable it is for the country to have the merger cleared by the competition regulator (R). If R were to observe the true θ (full information framework), she would like to clear the merger with probability one whenever $\theta \geq 0$ and block it otherwise. The multinational always has a positive net benefit normalized to 1 from the merger[18] and, therefore, he would like to have the merger cleared for any θ. The firm does not have private information about θ[19]. For instance, the firm may be giving information about the merger's market effect or the merger's effect on third parties, for which the firm

[18] Otherwise, he would not have proposed it in the first place.

[19] This assumption is standard in the career concerns literature. We discuss in Section 2.5 the consequences of relaxing it.

may not have more information than the competition authority. The parameter θ is distributed according to a Normal distribution, $f(\theta)$, with mean μ and variance η^2 and this distribution is common knowledge.

The actions available to the players are the following. M sends a message containing the information about θ to R but he does so through a noisy signal[20]. In particular, the signal has the following form:

$$S = \theta + \varepsilon$$

where ε is a random variable distributed according to a Normal distribution with mean zero and variance V. M can secretly choose V, which reflects the precision with which the message is sent (the larger the V, the less informative is the realization of the signal about θ). For instance, the reports can be drafted in a very imprecise way and this choice is secret because R does not know how much information M has. Also, M can overload R with too many documents (some of them irrelevant), knowing that R has a limited amount of time to review them. For simplicity, we assume that M can only choose V within this interval $[V_L, V_H]$, where $0 < V_L < V_H$.[21] The firm cannot, however, exaggerate the welfare benefits of the merger by shifting upwards the mean of ε (i.e. lying is not possible and therefore, the signal should be on average equal to θ). This is because the chances of being discovered and punish accordingly[22] are high given that the firm is in the middle of an investigation[23].

[20] The signal is noisy because the production technology (i.e. the way in which the firm compiles and transmits the information about the merger) is noisy.

[21] There are many justifications for $V_L > 0$. For instance, nobody knows θ exactly until the merger takes place or M is not able to compile and transmit the information perfectly (e.g., it is costly).

[22] See, for instance, the fines and convictions resulting from providing falso or misleading information to the Office of Fair Trading or the Competition Commission in the United Kingdom in the Enterprise Act 2002 at Sections 110, 111 and 117 (available at http://www.legislation.gov.uk/ukpga/2002/40/part/3/chapter/5/crossheading/investigation-powers).

[23] Another reason for ruling out the possiblity of biasing the signal upwards is because, without penalization, the incentives for increasing the bias are clear, completely monotonic, and independent from the noise dimension. Moreover, the firm would like to bias the report in the same direction for both authorities. We want to focus on the actions that differ depending on the authority that is in charge. For a model that considers both the manipulation of the mean and the precision of the signal see Drugov and Troya-Martinez (2012). The authors analyse how the incentives of a seller to provide biased and/or imprecise advice to consumers are affected by the

Therefore, the distribution of the signal S, $g(s,V)$, is Normal with mean μ and variance $\eta^2 + V$ and $G(s,V)$ is the cumulative distribution.

R observes the realization of the signal s, updates her beliefs about θ, and chooses the appropriate probability of clearance $p(s)$. We consider the case where $p(s)$ is a cut-off rule[24], where $\widehat{s}$ is optimally determined by R:[25]

$$p(s) = \begin{cases} 0 \text{ if } s < \widehat{s} \\ 1 \text{ if } s \geq \widehat{s} \end{cases}$$

The timing of the game depends on whether R has the ability to commit to a particular threshold before the information is transmitted by M. Under no commitment, the timing is as follows. First, M chooses the variance with which he sends the signal. A realization of the signal, s, is observed by R. Given this realization, R updates her beliefs about θ and decides which policy threshold $\widehat{s}$ to use. Finally, the payoffs are realized. The solution of this model is a Nash equilibrium where the conjectures of each of the players are correct in equilibrium. Under commitment, first R chooses a policy threshold $\widehat{s}$ and M then chooses V. A realization of the signal, s, is observed by R. Given this realization, R clears the merger whenever s is above $\widehat{s}$. Finally, the payoffs are realized. We solve this model backwards.

possibility of facing ex-post litigation, where a court will infer how much lying has taken place and penalise accordingly.

[24]The monotone likelihood ratio property (MLRP) holds in the setup. This is a sufficient condition for a cut-off rule to be optimal if the authority cannot commit to a threshold ex-ante (see Milgrom (1981)). It is not clear whether there is any loss of generality in the case where the authority can commit to a policy. Obviously, if the behaviour of the firm does not change under the more general rule $\widetilde{p}(s)$, then a cut-off rule dominates $\widetilde{p}(s)$, as it is ex-post optimal. Section 2.4.2 shows that if $\widetilde{p}(s)$ is non-decreasing, continous and odd function, and the thresholds are not too far from each other, the behaviour of the firm does not change. Example 2.1 in the Appendix provides another example where $\widetilde{p}(s) = Z(s)$ and $Z(s)$ is a cummulative Normal distribution.

For comparison purposes, we restrict ourselves to a cut-off rule also in the commitment case.

[25]R does not infer the noise used based on the unique realisation of the signal. This assumption is not restrictive because, in the no-commitment case, R has point beliefs about V which will not be affected by letting R make inferences from s. In the commitment case, it would be difficult for the authority to justify blocking the merger of an extreme positive realisation, given that all realisations can happen with positive probability and MLRP holds.

2.3 Benchmark: national merger

2.3.1 No-commitment regime

The problem of the firm

The firm always benefits from the merger and, therefore, chooses the level of precision that maximizes the probability of having the merger cleared. M makes a conjecture about the policy threshold used by R, $\widehat{s}^c$, where the superscript stands for conjecture. Using Bayes rule, M maximizes the probability of having the merger cleared:

$$P(V) = \int_{\widehat{s}^c}^{+\infty} g\left(s, V\right) ds$$

The objective function is decreasing in V whenever the average merger is a good merger (i.e. $\widehat{s}^c < \mu$) and increasing in V when the average merger is a bad merger (i.e. $\widehat{s}^c > \mu$). Therefore, if $\widehat{s}^c < \mu$, M will choose the minimum available variance V_L. Conversely, if $\widehat{s}^c > \mu$, M will choose the maximum available variance V_H.

Lemma 2.1 *Under no-commitment, the optimal variance is:*

$$V^*(\mu, \widehat{s}^c) = \begin{cases} V_H \text{ } if \text{ } \mu < \widehat{s}^c \\ V_L \text{ } if \text{ } \mu > \widehat{s}^c \end{cases}$$

If $\widehat{s}^c = \mu$, *M is indifferent between any variance in* $[V_L, V_H]$.

Proof. See Appendix. ■

For simplicity, we will assume that if $\widehat{s}^c$ is exactly μ the firm will choose V_L.

The intuition supporting this optimal strategy is that, by increasing the variance, an expected bad merger obtains more chances to be thought good. By contrast, decreasing the variance cuts down an expected good merger's chance of being considered bad. Therefore, the optimal V does not depend on how far the particular μ is from $\widehat{s}^c$, only on whether μ lies below or above $\widehat{s}^c$. This is due to the "bang-bang" payoff structure created by the cut-off rule. This "gambling for resurrection" result is in line with the findings of the literature that considers the choice of variability as a strategic variable[26].

The problem of the competition authority

Given the realization s and the conjecture that R forms about M's choice V^c, she needs to decide whether to clear the merger or not. The ex-post expected welfare of the merger is:

$$W(s, \widehat{s}^c) = \int_{-\infty}^{+\infty} \theta f\left(\theta | \, s, V^c(\mu, \widehat{s}^c)\right) d\theta$$

where $f\left(\theta | \, s, V\right)$ is the posterior distribution which is Normal with mean $\frac{s\eta^2+\mu V}{\eta^2+V}$ and variance $\frac{\eta^2 V}{\eta^2+V}$.[27]

If $W(s, \widehat{s}^c)$ is positive, R wants to clear the merger. Conversely, if $W(s, \widehat{s}^c)$ is negative, R blocks the merger. The authority determines the threshold $\widehat{s}$ so that she is indifferent between blocking or clearing the merger, that is, so that $W(s, \widehat{s}^c)$ is zero:

$$W(s, \widehat{s}^c) \mid_{s=\widehat{s}(\widehat{s}^c)} = \frac{\widehat{s}(\widehat{s}^c)\eta^2 + \mu V^c(\mu, \widehat{s}(\widehat{s}^c))}{\eta^2 + V^c(\mu, \widehat{s}(\widehat{s}^c))} = 0 \qquad (2.1)$$

Denote the solution of equation (2.1) as $\widehat{s}(\widehat{s}^c)$. The equilibrium threshold $\widehat{s}^*$ is, then, the fixed point of this solution:

$$\widehat{s}(\widehat{s}^*) = \widehat{s}^*$$

Proposition 2.1 *Under no commitment, the equilibrium threshold $\widehat{s}^*$ is:*

$$\widehat{s}^* = \begin{cases} \frac{-\mu V_H}{\eta^2} & \text{if } \mu < 0 \\ \frac{-\mu V_L}{\eta^2} & \text{if } \mu \geq 0 \end{cases}$$

Proof. It is straightforward to solve equation (2.1). ■

Figure 2.1 depicts the optimal threshold $\widehat{s}^*$ as a function of the average merger

[26] See Section 2.1.1.

[27] See Example 7.2.16 of Casella and Berger (2002) for the exact derivation of the Normal Bayes estimators.

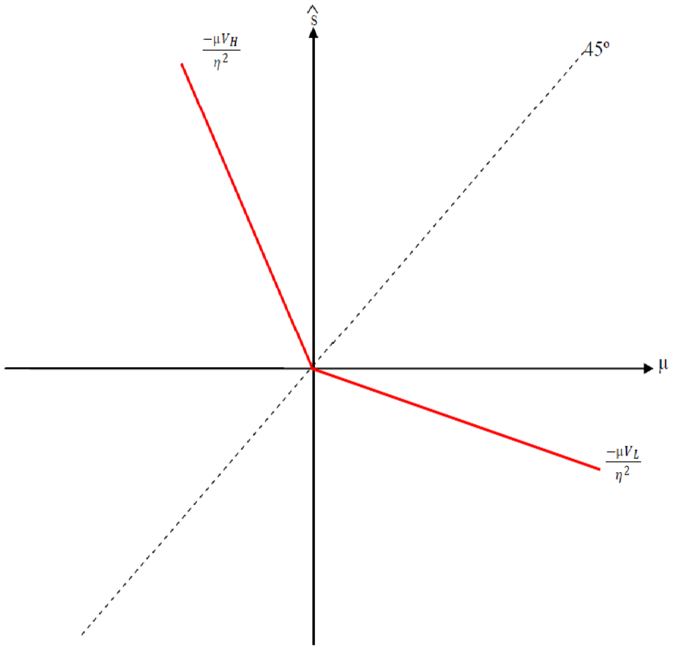

Figure 2.1: Optimal threshold under no commitment

μ.[28] If the bulk of mergers are bad mergers ($\mu < 0$), then the authority increases the standard of proof of a good merger by setting the policy threshold above the neutral merger 0. Conversely, when the average merger is welfare enhancing ($\mu > 0$), the authority sets a very lenient standard of proof, tolerating even negative signals.

Finally, note that as V_L tends to 0, the report becomes extremely informative and $\widehat{s}^*$ increases to zero (i.e. R becomes stricter with the on average good mergers as the probability of "unlucky realizations", that is, $s < \theta$, decreases) when $\mu \geq 0$. Similarly, as V_H tends to $+\infty$, the report submitted by M becomes uninformative and $\widehat{s}^*$ tends to $+\infty$, so all the mergers that may raise competition issue on average will be blocked when $\mu < 0$.

[28] By Lemma 2.1, above the 45^o line, $\widehat{s}^* > \mu$, so the firm chooses V_H and below the 45^o line, $\widehat{s}^* \leq \mu$, so M chooses V_L. By Proposition 2.1, the optimal threshold, $\widehat{s}^*$, has the opposite sign to the average merger and increases in absolute value with the average merger and $V^*(\mu, \widehat{s}^*)$.

2.3.2 Commitment regime

The problem of the firm

Suppose that R can commit in advance to a policy threshold $\widehat{s}$, for instance, by issuing very detailed merger guidelines. Then, M chooses V taking into account the actual $\widehat{s}$ instead of the conjecture $\widehat{s}^c$. The problem is identical to the one solved in Section 2.3.1 and therefore it is omitted here.

Lemma 2.2 *Under commitment, the optimal variance is:*

$$V^*(\mu, \widehat{s}) = \begin{cases} V_H \text{ if } \mu < \widehat{s} \\ V_L \text{ if } \mu > \widehat{s} \end{cases}.$$

If $\widehat{s} = \mu$, *M is indifferent between any variance in* $[V_L, V_H]$.

Proof. It is omitted as it is the same as Lemma 2.1. ∎

Again, let us break indifference by assuming that M will choose V_L.

The problem of the competition authority

The competition authority chooses $\widehat{s}$ to maximize the ex-ante expected welfare, taking into account the behavior of the firm:

$$\begin{aligned} EW(\widehat{s}, V^*(\mu, \widehat{s})) &= \int_{\widehat{s}}^{+\infty} \left(\frac{s\eta^2 + \mu V^*(\mu, \widehat{s})}{\eta^2 + V^*(\mu, \widehat{s})} \right) g\left(s, V^*\left(\mu, \widehat{s}\right)\right) ds \qquad (2.2) \\ \text{subject to} &: \quad V^*(\mu, \widehat{s}) = \begin{cases} V_H \text{ if } \mu < \widehat{s} \\ V_L \text{ if } \mu \geq \widehat{s} \end{cases} \end{aligned}$$

where the term in brackets is the expected welfare given a realization s, $W(s, \widehat{s})$.[29]

[29] The objective function is equivalent to $\int_{-\infty}^{+\infty} \theta f\left(\theta | \, s, V^*(\mu, \widehat{s})\right) d\theta$, where $f\left(\theta | \, s, V^*(\mu, \widehat{s})\right)$ is the posterior distribution of θ.

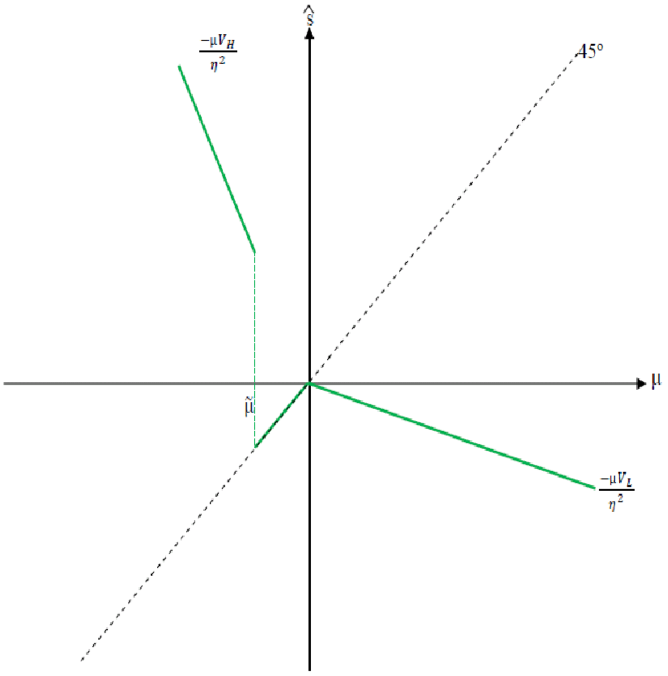

Figure 2.2: Optimal threshold under commitment

Proposition 2.2 *There exist $\widetilde{\mu}$ such that, under commitment, the optimal threshold $\widehat{s}^{**}$ is:*

$$\widehat{s}^{**} = \begin{cases} \frac{-\mu V_H}{\eta^2} & \text{if } \mu < \widetilde{\mu} \\ \mu & \text{if } \mu \in [\widetilde{\mu}, 0] \\ \frac{-\mu V_L}{\eta^2} & \text{if } \mu > 0 \end{cases}$$

$\widehat{s}^{**}$ is non-monotonic in the expected welfare and $\widetilde{\mu}$ strictly increases (decreases) with V_L (V_H).

Proof. See Appendix. ■

As in Figure 2.1, we depict $\widehat{s}^{**}$ as a function of μ in Figure 2.2.

Note that the authority values the level of precision of the signal as this allows her to make fewer mistakes when deciding whether to clear the merger or not.

When the average merger is welfare enhancing ($\mu > 0$), by Lemma 2.2, M chooses V_L. Since, this is as informative as the report can be, R sets the ex-post optimal threshold found in Section 2.3.1.

By the same Lemma, when the average merger is welfare detrimental, M chooses V_H. This introduces a new trade-off for the competition authority because she may decide to distort the ex-post optimal policy upwards in order to extract more precise information from the firm. Indeed, we find that, when there is substantial uncertainty about the undesirability of the average merger (μ is negative and close to zero, i.e., $\mu \in [\widetilde{\mu}, 0]$), then R gains from committing to a more lenient threshold (not only more lenient than the ex-post optimal threshold but also than the full information threshold) in order to induce M to reduce his equilibrium variance from V_H to V_L. On the other hand, if the merger is on average clearly welfare detrimental ($\mu < \widetilde{\mu}$), the authority will again set the ex-post optimal threshold found in Section 2.3.1.

Therefore in some cases, the authority, by gaining commitment power, decides to distort her policy by decreasing her standards (in an ex-post non-optimal way) in order to fix the information problem (i.e. to improve the informativeness of the signal).

Corollary 2.1 *Policy comparison:*

$\widehat{s}^* = \widehat{s}^{**}$ $\forall \mu \in (-\infty, \widetilde{\mu}) \cup (0, +\infty)$ *and* $\widehat{s}^* > \widehat{s}^{**} \forall \mu \in [\widetilde{\mu}, 0]$.

Finally, note that if V_L increases, the range of mergers for which the authority is more lenient than the ex-post optimal, $[\widetilde{\mu}, 0]$, shrinks as there is less gain from obtaining a less precise signal. Similarly, this range will expand if V_H increases.

2.4 Multinational merger

In this section we consider the situation where a multinational wants to undertake the same merger in two different countries (or jurisdictions) and these countries differ in terms of their policies. Without loss of generality, we assume that Country 1 will block the merger if the signal is below $\widehat{s}_1$ and clear it otherwise, while Country 2 will do the same using the threshold $\widehat{s}_2$, where $\widehat{s}_1 < \widehat{s}_2$.

We assume that the authorities can commit to the thresholds. The thresholds are then set ex-ante based on a generic (as compared to merger specific) anticipated distribution of mergers.[30] The difference in thresholds can be interpreted as Country 1 being in general more lenient in its merger policy than Country 2. For instance, if the welfare function of Country i is $a_i + b_i\theta$, where $i = 1, 2$, then $a_1 > a_2$ and $b_2 > b_1 > 0$[31], where a_i could be interpreted as how lenient is the authority and b_i how sensitive is the authority to the welfare's changes.[32] Another possible interpretation is that mergers are expected to have different competitive effects in the two countries. For instance, this would be the case if Country 1 has a highly trained competition authority with extensive resources that has been very successful in ensuring that markets are competitive. In this case, a_i and b_i do not reflect the taste of the authority but rather the general level of competitiveness and openness of the different markets in the Country.

In practice, some mergers may raise competition issues at a national (or sub-national) level that can be solved by a local action (for instance, to impose a remedy that forces the firm to undertake a national divestiture if the merger is cleared somewhere else) without the need to reach an agreement with other competition authorities. In this paper we focus on the type of mergers that do need the agreement of all the jurisdictions in order for the merger to happen. An example of such a merger was the one proposed between two U.S.-based companies, General Electric and Honeywell, with the E.U. prohibiting the merger,[33] and the U.S. Department of Justice approving it[34].[35]

---

[30] If the authorities could not commit to a policy, then they would set the thresholds on a case-by-case basis taking into account the specific distribution of the particular merger. If $\widehat{s}_1 < \mu < \widehat{s}_2$, then Country 1 would like to set $\widehat{s}_1 = -\infty$ while Country 2 $\widehat{s}_2 = +\infty$, to counteract the other country's interference.

[31] To be precise, we need $\frac{a_1}{b_1} > \frac{a_2}{b_2}$.

[32] It is easy to check that the resulting equilibrium threshold under no-commitment is:

$$\widehat{s}_i = \begin{cases} -\mu\frac{V_H}{\eta^2} - \frac{a_i}{b_i}\frac{\eta^2+V_H}{\eta^2} & \text{if } \mu < -\frac{a_i}{b_i} \\ -\mu\frac{V_L}{\eta^2} - \frac{a_i}{b_i}\frac{\eta^2+V_L}{\eta^2} & \text{if } \mu \geq -\frac{a_i}{b_i} \end{cases}$$

Under commitment, if the ratio $\frac{a_1}{b_1}$ is very close to $\frac{a_2}{b_2}$ there may be a region of μ where both authorities will coincide in setting $\widehat{s}_i = \mu$. To focus on the interesting cases, we assume that the welfare functions are different enough so that this does not occur and $\widehat{s}_1 < \widehat{s}_2$ for all μ.

[33] *GE/Honeywell*, Case No COMP/M.2220, European Commission Decision available at http://www.ec.europa.eu/competition/mergers/cases/decisions/m2220_en.pdf

[34] Subject to GE divesting Honeywell's helicopter engine business and licensing a new competitor to maintain and repair certain Honeywell engines.

[35] See Muris (2001) for more detail.

Since what drives the variance decision is how μ compares with the threshold, we focus on the case where the information-sharing agreement makes a difference in the firm's choice. This occurs when there is a conflicting merger in the sense that an average merger is good for Country 1 but bad for Country 2, that is, $\widehat{s}_1 < \mu < \widehat{s}_2$.[36]

When the competition authorities do not share information about the firm, the problem of the multinational is separable and we are back to the national merger case. The multinational needs to convince each authority individually, regardless of the way in which authorities reach an agreement ex-post. In other words, he chooses a level of precision for each country so as to maximize the probability of having the merger cleared in each country. By Lemma 2.1, the firm will choose V_L in Country 1 and V_H in Country 2.

When the authorities share information, they receive the same realization of the signal (i.e. the same report). We do not consider the case where each authority receives a different realization from the same signal to remove the benefits that simply come from having more information. The problem of the multinational is no longer separable and he needs to choose a unique level of precision so as to maximize the probability of having the merger cleared. This probability will depend on the process by which a final decision on the international merger is reached and we consider several possibilities below.

In what follows, we analyze the impact that information-sharing has on the behavior of the firm, keeping the competition authorities policies fixed, that is, we consider the firm's reaction to $(\widehat{s}_1, \widehat{s}_2)$. We have not considered how $\widehat{s}_1$ and $\widehat{s}_2$ adjust to take into account the international dimension of the problem for two reasons: first, information-sharing agreements do not usually contemplate changes in policies and, second, the policies are the same for national and international mergers (which are only a minority of the mergers reviewed), which makes them less sensitive to the agreement.

[36]Note that if $\mu < \widehat{s}_1$ (or $\widehat{s}_2 < \mu$), the firm chooses V_H (or V_L) in both countries and the information-sharing agreement makes no difference.

2.4.1 Information-sharing with veto power

Consider first the case where the authorities only cooperate in exchanging information but not in their decision process, i.e. the authorities use their veto power. The only disagreement that can arise is that Country 1 wants to clear the merger, while Country 2 does not want this. Country 2 using its veto power translates into the merger being cleared only if both competition authorities agree that the merger should be cleared (i.e. if the signal is above $\widehat{s}_2$) and as a result the firm chooses V_H.

2.4.2 Information-sharing with cooperation in the decision-making

We turn now to the case where there is a bargaining or persuasion process taking place between the authorities, for example, due to the repeated interaction between them. We model this decision as taking place in two stages: first, the competition authorities decide unilaterally whether they should clear the merger and, then, if the decisions differ, they discuss their arguments until they reach an agreement.[37] From the point of view of the firm, the outcome of this bargaining is a conflicting merger being cleared with some probability. We consider two ways of how this probability is determined.

Constant probability

We first analyze the case where, if there is disagreement, the merger is cleared with probability α (even though Country 2 does not want this) and blocked with probability $1-\alpha$. The parameter α can be interpreted as a measure of the bargaining power of the authority in Country 1 versus the one in Country 2. M maximizes the probability of having the merger cleared:

[37] For instance, in the WorldCom / Sprint merger review, the co-operation between the Europan Commission and the US Department of Justice involved such an extensive sharing of information (thanks to a confidentiality waiver granted by the parties) that allowed both case teams to discuss in-depth the merits of the case and to reach consistent assessments of the competitive impact of the transaction on the area of joint concern. For this, and many more examples where information-sharing has lead to a more active cooperation that have resulted in decisions to clear mergers, see OECD (2001).

$$P^{\alpha}(V) = \alpha \int_{\widehat{s}_1}^{\widehat{s}_2} g\left(s, V\right) ds + \int_{\widehat{s}_2}^{+\infty} g\left(s, V\right) ds \tag{2.3}$$

for $\alpha \in (0,1)$.

Since μ lies in between the thresholds, the first part of his objective function is decreasing in V, while the second part is increasing in V. In other words, when the firm chooses a high variance, this decreases the mass of the probability distribution of the signal in the interval $[\widehat{s}_1, \widehat{s}_2]$, but at the same time increases the mass in the tails, so the chances that the realization lies in the interval $[\widehat{s}_2, +\infty)$ increase. Sharing information, makes the firm face a new trade-off between choosing V so as to maximize the probability in the interval $[\widehat{s}_1, \widehat{s}_2]$ or in the interval $[\widehat{s}_2, +\infty)$.

Denote the middle point of the interval $[\widehat{s}_1, \widehat{s}_2]$ by $s_M = \frac{\widehat{s}_1 + \widehat{s}_2}{2}$. The behavior of the multinational for $\mu \in [\widehat{s}_1, \widehat{s}_2]$ is summarized in the next Proposition.

Proposition 2.3 *For $\mu \in [\widehat{s}_1, \widehat{s}_2]$, the choice of variance may be non-monotonic in μ. In particular, for $\mu \in [\widehat{s}_1, s_M)$, M chooses either V_L or V_H. For $\mu \in (s_M, \widehat{s}_2]$, M may choose some intermediate variance $V(\mu,\alpha) = \frac{(\widehat{s}_2-\mu)^2-(\mu-\widehat{s}_1)^2}{2\ln\left(\frac{1-\alpha}{\alpha}\frac{\widehat{s}_2-\mu}{\mu-\widehat{s}_1}\right)} - \eta^2$ (provided $V(\mu,\alpha) \in (V_L, V_H)$).*

Proof. See Appendix. ■

This result relies on the fact that the curvature of the firm's objective function depends on μ. In particular, the objective function is quasi-convex in V for $\mu \in (\widehat{s}_1, s_M)$ and quasi-concave in V for $\mu \in (s_M, \widehat{s}_2)$.

Thus, for $\mu \in (\widehat{s}_1, s_M)$, M chooses either V_H for all μ or it starts with V_H and then switches to V_L if the following condition holds:

$$\underbrace{1 - G\left(\widehat{s}_2; V_H\right) - \left(1 - G\left(\widehat{s}_2; V_L\right)\right)}_{\text{Relative cost of choosing } V_L} <$$

$$\underbrace{\alpha\left[G\left(\widehat{s}_2; V_L\right) - G\left(\widehat{s}_1; V_L\right) - \left(G\left(\widehat{s}_2; V_H\right)\right) - G\left(\widehat{s}_1; V_H\right)\right]}_{\text{Expected relative benefit of choosing } V_L}$$

This condition states that, by choosing V_L, as opposed to V_H, M loses chances of being cleared with probability 1 but may increase the chances of being cleared with probability α if μ is close enough to s_M and depending on how far away from μ is the crossing point between $g(s; V_L)$ and $g(s; V_H)$. The switch happens more often when α is larger.

The following two corollaries state the precise variance that the firm chooses, when the authorities have equal bargaining power, that is, for $\alpha = \frac{1}{2}$. Denote $c(V_L, V_H)$ as the distance between μ and the crossing point of $g(s; V_L)$ and $g(s; V_H)$.

Corollary 2.2 *For $\mu \in (\widehat{s}_1, s_M)$ and $\alpha = \frac{1}{2}$, a sufficient condition for M to choose V_H for any $\mu \in (\widehat{s}_1, s_M)$ is:*

$$\widehat{s}_2 - \mu < c(V_L, V_H)$$

where $c(V_L, V_H)$ increases with V_H and V_L.

There exists a threshold $\overline{\mu}$ such that a sufficient condition for M to first choose V_H until $\mu = \overline{\mu}$ and then switch to V_L for $\mu \in (\overline{\mu}, s_M)$ is:

$$c(V_L, V_H) < \mu - \widehat{s}_1$$

These conditions are not necessary.

Proof. See Appendix. ■

Figure 2.3 depicts this result.

Note from Corollary 2.2 that for a given pair of policy thresholds $\{\widehat{s}_1, \widehat{s}_2\}$, the larger V_H, the larger $c(V_L, V_H)$ and hence, the more likely it is that we are in the case where the firm chooses V_H for all $\mu \in (\widehat{s}_1, s_M)$. In the same way, the smaller V_L or the larger the conflict between authorities $(\widehat{s}_2 - \widehat{s}_1)$, the more likely it is that we are in the regime where the firm chooses first V_H and then V_L. Therefore, if the authorities have equal bargaining power and their conflict is large (or it is possible to send very precise reports), the choice of the variance for the conflicting mergers is non-monotonic in the average merger. This non-monotonicity is the result of the trade-off between maximizing the probability at the interval $[\widehat{s}_1, \widehat{s}_2]$ or at the interval $[\widehat{s}_2, +\infty)$.

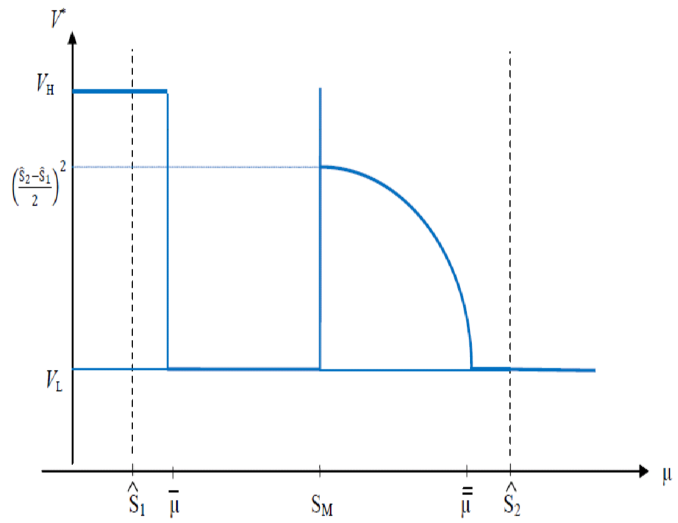

Figure 2.3: Optimal choice of variance with constant probability

The intuition behind the choice of variance is the following. For the average merger μ between $\widehat{s}_1$ and s_M, the objective function is quasi-convex in V. Therefore, the firm chooses an extreme variance. First consider the average merger at the left endpoint, $\mu = \widehat{s}_1$. For this μ, the probability that the realization of the signal is above $\widehat{s}_1$ is $\frac{1}{2}$ for all the values of V. However, the larger the variance, the higher the probability of obtaining a realization above $\widehat{s}_2$ where the merger is cleared with probability 1 rather than $\frac{1}{2}$. Therefore, he will choose V_H. By continuity, the same intuition holds for the average mergers that are in between $\widehat{s}_1$ and $\overline{\mu}$, where $\overline{\mu}$ is defined in equation (2.13). If the minimum variance is low enough (or the conflict between authorities is high enough), then it is optimal for the firm to switch to V_L after $\overline{\mu}$ in order to maximize the chances of being cleared at least by Country 1. It is at this point that the new trade-off becomes effective.

The firm with an average merger just in the middle, $\mu = s_M$, is indifferent between any variance. The reason for this is that the probability of having the merger cleared is $\frac{1}{2}$ for all possible values of V due to the symmetry of the probability function at this point.

As μ exceeds s_M, the objective function of the firm becomes quasi-concave in V. Because of the proximity of μ with the interval $[\widehat{s}_2, +\infty)$, it becomes safer to play a riskier strategy by increasing V (without reaching the maximum level). Finally, as μ keeps increasing towards $\widehat{s}_2$, the value of this optimal intermediate variance decreases.[38] This is because the chances of having the realization in the interval $[\widehat{s}_1, +\infty)$ are already high and, by reducing the variance, the probability of having the realization in the tail $(-\infty, \widehat{s}_1)$ is decreased. The variance decreases until it hits the minimum level at $\overline{\overline{\mu}}$.[39]

Finally, when the average merger is at the right endpoint, $\mu = \widehat{s}_2$, the probability of having a realization of the signal below $\widehat{s}_2$ is $\frac{1}{2}$; however the higher the variance, the higher the probability that the realization of the signal is below $\widehat{s}_1$ where the merger is blocked, rather than cleared with probability $\frac{1}{2}$. As a result, the firm chooses V_L.

The information-sharing agreement has modified the payoff structure, which is no longer "bang-bang" as in Section 2.3. As a result, the firm makes a greater use of intermediate levels of variability in a non-monotonic way.

Increasing probability

Now we consider the case where the bargaining power of Country 1 (Country 2) increases (decreases) monotonically with the particular realization s. This reflects the fact that the closer s is to $\widehat{s}_2$, the less reluctant (more convincing) will Country 2 (Country 1) be about clearing the merger as compared to a realization s very close to $\widehat{s}_1$. We consider the case of a continuous, increasing and odd (with respect

[38] Note that:

$$\frac{dV\left(\mu, \frac{1}{2}\right)}{d\mu} = \frac{-\left(\widehat{s}_2 - \widehat{s}_1\right)\left[2\ln x - x + \frac{1}{x}\right]}{2\left(\ln x\right)^2}$$

where $x = \frac{\widehat{s}_2 - \mu}{\mu - \widehat{s}_1}$. Since we focus on $\mu \in (s_M, \widehat{s}_2]$, then $0 < x < 1$ and $\frac{dV\left(\mu, \frac{1}{2}\right)}{d\mu} < 0$.

[39] Since M would like to choose a variance at $\widehat{s}_2$ equal to $-\eta^2$, there exists $\overline{\overline{\mu}}$ such that:

$$V\left(\overline{\overline{\mu}}, \frac{1}{2}\right) = \frac{(\widehat{s}_2 - \overline{\overline{\mu}})^2 - (\overline{\overline{\mu}} - \widehat{s}_1)^2}{2\ln\left(\frac{\widehat{s}_2 - \overline{\overline{\mu}}}{\overline{\overline{\mu}} - \widehat{s}_1}\right)} - \eta^2 = V_L$$

$\left(s_M, \frac{1}{2}\right))$ function $\rho(s)$. The expected probability of clearance then becomes:

$$P^c(V) = \int_{\widehat{s}_1}^{\widehat{s}_2} \rho(s)g(s,V)ds + \int_{\widehat{s}_2}^{+\infty} g(s,V)ds$$

Contrary to the previous case, now increasing the variance does not unambiguously decrease the probability of being cleared if the realization s falls in the interval $(\widehat{s}_1, \widehat{s}_2)$ because these realizations have attached different probabilities of clearance.

Proposition 2.4 *The optimal variance chosen by M is:*

$$V^*(\mu, s_M) = \begin{cases} V_H \text{ if } \mu < s_M \\ V_L \text{ if } \mu > s_M \end{cases}$$

If $\mu = s_M$, then M is indifferent between any variance in $[V_L, V_H]$.

Proof. See Appendix. ■

See Figure 2.4. When μ is below s_M, increasing the variance increases the probability of falling above $\widehat{s}_2$ and decreases the probability of having realizations in the first half of $[\widehat{s}_1, \widehat{s}_2]$, where the probability of clearance is very low anyway. This makes increasing the variance less damaging, so the firm chooses V_H. Instead, when μ is above s_M, increasing the variance decreases the probability of having realizations in the second half of $[\widehat{s}_1, \widehat{s}_2]$, where the probability of clearance is large and increasing, so the firm chooses V_L as a result. From the point of view of the firm, it is as if there were a unique threshold but there is uncertainty about where exactly this lies. Because of the symmetry of $\rho(s)$ around $\left(s_M, \frac{1}{2}\right)$, the firm takes the expectation and behaves as in the national merger case with respect to this unique threshold. That is, the firm uses high variance whenever he thinks the merger is not going to pass through $(\mu < s_M)$ and low variance otherwise.

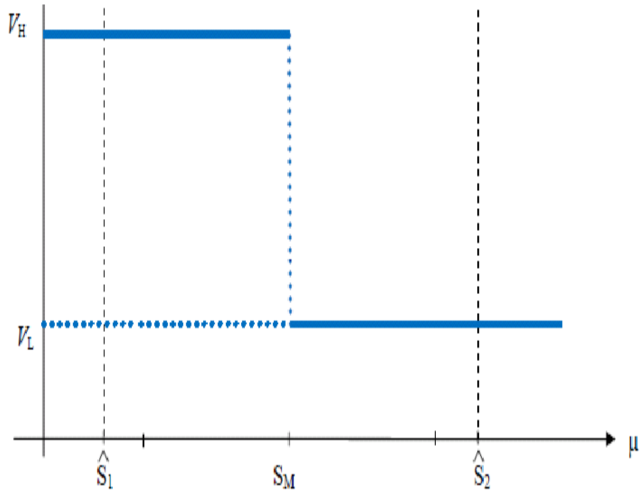

Figure 2.4: Optimal choice of variance with increasing probability

2.4.3 Information-sharing with cooperation in the policy-making

So far, we have assumed that each authority is fixing its policy threshold unilaterally. In this section, we consider the situation where the authorities fix a unique threshold, $\widehat{s}^J$, for both countries, for instance, in the European case, this would correspond to the European Commission setting a common policy. This policy will be chosen so that their joint ex-post welfare is maximized:

$$W(s,\widehat{s}^J)\mid_{s=s(\widehat{s}^c)}=\gamma\left[a_1+b_1\left(\frac{s\eta^2+\mu V^*(\mu,s)}{\eta^2+V^*(\mu,s)}\right)\right]+(1-\gamma)\left[a_2+b_2\left(\frac{s\eta^2+\mu V^*(\mu,s)}{\eta^2+V^*(\mu,s)}\right)\right]$$

where γ and $1-\gamma$ are the weights given to the welfare of Country 1 and 2, respectively. It is easy to check that, given the firm's behavior in Lemma 2.2, the optimal threshold is:

$$\widehat{s}^J=\begin{cases}-\mu\frac{V_H}{\eta^2}-\frac{\gamma a_1+(1-\gamma)a_2}{\gamma b_1+(1-\gamma)b_2}\frac{\eta^2+V_H}{\eta^2} & \text{if } \mu<-\frac{\gamma a_1+(1-\gamma)a_2}{\gamma b_1+(1-\gamma)b_2}\\ -\mu\frac{V_L}{\eta^2}-\frac{\gamma a_1+(1-\gamma)a_2}{\gamma b_1+(1-\gamma)b_2}\frac{\eta^2+V_L}{\eta^2} & \text{if } \mu\geq-\frac{\gamma a_1+(1-\gamma)a_2}{\gamma b_1+(1-\gamma)b_2}\end{cases}$$

2.4.4 Discussion

It stands from the previous analysis that if the competition authorities have veto power over the decision, then sharing information is never beneficial as Country 1 is strictly worse off due to the imprecise submitted information (while Country 2 is indifferent).

If the authorities also cooperate in the decision taking stage, then whether sharing information is desirable ultimately depends on the average merger that the authorities are facing and on the particular rule used to reach a decision. To illustrate this point, imagine that the merger's welfare consequences in Country 1 are $3 + \theta$ while in Country 2 are θ. Imagine that the firm can choose a variance on the range $[2, 10]$ and that $\eta^2 = 4$.

If the ex-ante average merger is $\mu = -0.5$, then the authorities would set the national thresholds $\widehat{s}_1 = -2$ and $\widehat{s}_2 = 1.25$, so in the absence of information-sharing, the firm chooses variance 2 for Country 1 and variance 10 for Country 2. Imagine now that the authorities share information and the persuasion process is such that the probability of clearance in case of disagreement is constant and equal to 0.5. Then, given that $s_M = -1.5$, the firm chooses an intermediate variance equal to 3.22 and the ex-ante expected welfare is 1.822 and 0.058 for Country 1 and 2, respectively. Instead, if the probability of clearance is increasing and linear $\left(\rho(s) = \frac{s-s_1}{s_2-s_1}\right)$, the firm chooses a variance of 2 and the expected welfare of Country 1 and 2 are 2.088 and 0.193, respectively. Therefore, the authorities are better-off if their bargaining process is such that the probability of clearing a conflicting merger is increasing and linear. Obviously, the lowest joint expected welfare is obtained when the countries have veto power (1.404) and the largest when the authorities coordinate on the policy threshold (2.540, with $\gamma = 0.5$) where $\widehat{s}^J = -2$ and the firm uses variance 2.

Let us consider instead an ex-ante average merger $\mu = -2$. The thresholds set by the authorities are $\widehat{s}_1 = -3.5$ and $\widehat{s}_2 = 5$. Since $s_M = 0.75$, in the constant probability case the firm chooses minimum variance (condition (2.11) is satisfied) and expected welfare of Country 1 and 2 are 0.641 and -0.456, respectively. If instead the probability of clearance is linear, then the firm choose variance 10 and ex-ante expected welfare are 0.567 and -0.251. Hence, the authorities are jointly better off with the constant probability case. Again, the lowest expected welfare is under the veto regime (0.117) and the highest under the coordination of policies (0.445) where $\widehat{s}^J = -0.25$ and the firm chooses variance 10.

2.5 Extension: the case of private information

In this Section, we explore the robustness of our results in the assumption that the firm is not more informed than the competition authority about the merger's welfare effects.[40]

If we assume that θ is the private information of the firm then, qualitatively, the firm's behavior does not change. The optimal variance chosen by the firm will depend on the merger's type θ in the same way that it was depending before on the average merger μ. In particular, in the case of the national merger, a firm with a bad merger will choose the risky strategy of high variance, while a firm with a good merger will choose low variance to increase the likelihood of obtaining a high realization of the signal. In the same way, in the case of the multinational merger, the information-sharing agreement will have no impact on the types of merger that are either good or bad for both countries at the same time. However, for those firms whose merger is good for one country but bad for the other, the choice of the variance may be non-monotonic in their type in the same way as was found in Proposition 2.3.

However, with this new assumption, there is a significant change in the authority's behavior because the monotone likelihood ratio property (MLRP) does not generally hold, due to the way in which the firm optimally responds to the threshold rule. In particular, a very large s is more likely to come from a (bad) merger below the threshold because the distribution of the signal sent by such a merger has fatter tails. If the MLRP does not hold, then the cut-off rule may not be optimal. However, if we nonetheless restrict the authority to using the cut-off rule, then we can show that the equilibrium threshold, when the average merger is welfare detrimental, is stricter than the full information threshold and, contrary to what we found

[40] The proofs for this section are available upon request.

in Section 2.3, it does not change with the authority's ability to commit.

The intuition for this result is as follows. There are two types of effects following an increase in the commitment threshold, $\widehat{s}^{**}$: the direct and the strategic effect. The direct effect is the result from the trade-off between the benefit of decreasing the clearance probability of a bad merger against the cost of decreasing the clearance probability of a good merger. The direct benefit can be interpreted as a type II error of clearing a merger that should be blocked. By increasing $\widehat{s}^{**}$ we make this error less likely. Similarly, the direct cost can be interpreted as a type I error of blocking a merger when in fact it should be allowed and by increasing $\widehat{s}^{**}$ we make this error more likely.

The strategic (or indirect) effect is the result of the change in the firm's strategy resulting from moving merger types from above to below the threshold (i.e. when the firm switches from V_L to V_H, a good type θ has less chances of being cleared). Given that the authority values precision, as this allows her to take more informed decisions, the fact that in this new framework she chooses the same threshold regardless of whether or not she can commit to a policy may be puzzling. It seems natural to think that the mechanism highlighted in the Corollary 2.1 would still apply (that is, if under no commitment, the authority sets a threshold $\widehat{s}^{*}$, she would gain from committing to a lower threshold $\widehat{s}^{**}$ as this would induce the types in the interval $[\widehat{s}^{**}, \widehat{s}^{*}]$ to reduce their equilibrium variance from V_H to V_L). However, when the authority is considering a marginal decrease in the threshold, she only takes into account the strategic effects of the marginal type $\widehat{s}^{*}$ which consists in decreasing the variance from V_H to V_L. However, for $\widehat{s}^{*}$ (the type at the margin) the probability that the realization of the signal is above $\widehat{s}^{*}$ is $\frac{1}{2}$ in both cases. Therefore, since the marginal change in the firm's behavior generated by the ability to commit cancels out, the criterion used to set such a threshold is the same under both regimes. In particular, by lowering the threshold, the authority only trades off the increase in the type II error and the decrease in the type I error (i.e. the direct effect).

2.6 Conclusions

The goal of this paper has been to study the impact of an information-sharing agreement on the incentives of the firm to provide precise information to the authorities. We find that the agreement has no impact on the firm with an average merger that is welfare enhancing for both countries at the same time. This would explain why in some cases, where the firm is sure that a merger does not raise competition issues, he voluntarily grants a confidentiality waiver to the authorities. However, the agreement affects the behavior of the firm whose average merger is welfare enhancing for one country but welfare detrimental for the other. The impact in terms of information provision will depend on how authorities reach an agreement in the case of disagreement.

If there is no further cooperation and authorities exert their veto power, then the firm sends very imprecise information to both authorities. If the authorities cooperate further, and if they have equal bargaining power, then the choice of the variance may be non-monotonic in the average merger. Furthermore, despite the choice of variance being costless, intermediate levels of precision are chosen in equilibrium. Finally, if the probability is increasing on the particular report then the firm behaves as in the national merger case where the new cut-off rule is the expectation of the countries' rules.

Given their policies, do the agencies benefit from the agreement? The authorities value the level of precision because this allows them to make more accurate decisions; therefore, sharing information but retaining their veto power makes the authorities strictly worse off. If there is further cooperation, the firm makes more use of intermediate and lower variances. As a result, only the less lenient country benefits from the agreement because, before, she was receiving information with the minimum level of precision. The more lenient country will not benefit from the agreement because, without the agreement, she was receiving the information with the maximum level of precision. Information-sharing and cooperation could potentially be beneficial if the lenient country is not always the same one. This is not the case here where the authorities set their thresholds using the information of the general population of mergers. It could be the case if, in practice, authorities set their policies at an industry level (i.e. only taking into account the population of mergers in a given industry) and each country has industries that are more and less competitive than in the other country.

In our analysis, we ignore a very important cost attached to the exchange of confidential information: the danger of leakage of commercially sensitive information to third parties. This can be clearly an issue when cooperation involves competition agencies in countries where the law for protecting confidential information is weak or where the credibility of the agency is low. We also ignore many benefits. For instance, the exchange of information facilitates the discussion about the case and minimizes the possibility of missing an issue that needs an enforcement action. In the merger cases, it also eliminates conflicting decisions and remedies.[41]

A possible way to enrich the setup is to let each authority receive a different realization from the same random signal[42]. In this context, sharing information increases the quantity of information on which to base their decisions and allows them to better infer the level of precision used by the firm. Thus, noise becomes more costly and firms with bad average mergers will have their behavior affected. There is room for the more lenient authority to benefit from the agreement.

We compared the no-information-sharing regime with the information-sharing regime, keeping the policies of the competition authorities fixed, as we were interested in exploring the changes in the firm's behavior following the agreement. One motivation for proceeding this way is that these agreements do not usually contemplate changes in policies. Another one is that the policies are the same for national and international mergers (which are only a minority of the mergers reviewed), which

[41] A successful example of a merger involving cooperation between agencies that illustrates these points was the Holnam/Lafarge case. A waiver granted by the parties allowed the U.S. and Canadian agencies to effectively improve the coordination, which ended up in a more informed decision-making (see Valentine (2000)).

[42] This could be the case if information-sharing happens at a late state of the investigation process, where the different signal's realisations can be interpreted as the information obtained from the particular questions (interviews, questionnaires, etc.) carried out by each authority.

makes them less sensitive to the new regime. One possible direction of future work would be to enrich the model by allowing the competition agencies to strategically adapt their policies to the new regime. This would allow us to assess the total impact of the information-sharing regime, which not only includes the change in the firm's behavior but also the change in the authority's policy.

Finally, the lessons from this model can be applied to a variety of situations in industrial organization and political economy that involve binary decisions. For instance, the model can apply to the information submitted to a sectorial regulator or central bank and a competition authority (as is the case with the mergers involving banks), to interviews in a recruitment process with two interviewers, to a project approval by different departments within a firm, to reforms submitted to bicameral parliaments or, more generally, to a politician trying to gain support for a policy in front of two audiences with different opinions about the policy.

2.7 Appendix

Example 2.1 *Assume that R commits to clear the merger with probability $Z(s)$, where $Z(s)$ is the cumulative distribution of $z(s) \sim N(m,v)$ and m and v are optimally chosen by R. M chooses V to maximize the probability of clearance:*

$$\underset{V}{Max} \int_{-\infty}^{+\infty} Z(s)g(s,V)ds$$

The first order condition is:

$$\begin{aligned}
\frac{\partial}{\partial V} &= \int_{-\infty}^{+\infty} Z(s)\frac{\partial g(s,V)}{\partial V}ds \\
&= \frac{-1}{2}\int_{-\infty}^{+\infty} \frac{s-\mu}{v} z(s)g(s,V)ds \\
&= \frac{-1}{2} r(m)\int_{-\infty}^{+\infty} \frac{s-\mu}{v}\gamma(s)ds \\
&= -\frac{1}{2}\frac{\partial r(m)}{\partial m}
\end{aligned}$$

where the second and fourth line follows from integrating by parts and $r(m) \sim N(\mu, v+\eta^2+V)$ *and* $\gamma(s) \sim N\left(\frac{v\mu+m(\eta^2+V)}{v+\eta^2+V}, \frac{v(\eta^2+V)}{v+\eta^2+V}\right)$*. Hence we observe the same behavior as in Section 2.3.1: if* $\mu < m$, $\frac{\partial}{\partial V} > 0$ *and so the optimal variance is* V_H*; if* $\mu = m$, $\frac{\partial}{\partial V} = 0$ *and M is indifferent between any variance and if* $m < \mu$, $\frac{\partial}{\partial V} < 0$ *and so M chooses* V_L*. Since the behavior of the firm only depends on how* m *compares to* μ *(i.e. it is independent of* v*), R chooses the ex-post optimal* v*, that is* $v = 0$*. It is then easy to show that the optimal* m *is equal to the cut-off* $\widehat{s}^{**}$ *found in Proposition 2.2.*

Proof of Lemma 2.1. The first derivative of $P(V)$ with respect to V is:

$$\frac{\partial P(V)}{\partial V} = \int_{\widehat{s}^c}^{+\infty} \frac{\partial g(s,V)}{\partial V}ds \tag{2.4}$$

Note that $\frac{\partial g(s,V)}{\partial V} = \frac{1}{2}\frac{\partial^2 g(s,V)}{\partial s^2}$:

$$\begin{aligned}
\frac{\partial g(s,V)}{\partial V} &= \frac{1}{\sqrt{2\pi\left(\eta^2+V\right)}}\exp\left(-\frac{(s-\mu)^2}{2(\eta^2+V)}\right)\frac{(s-\mu)^2-(\eta^2+V)}{(\eta^2+V)^2} \\
\frac{\partial^2 g(s,V)}{\partial s^2} &= \frac{1}{\sqrt{2\pi\left(\eta^2+V\right)}}\exp\left(-\frac{(s-\mu)^2}{2(\eta^2+V)}\right)\frac{(s-\mu)^2-(\eta^2+V)}{2(\eta^2+V)^2}
\end{aligned}$$

Using this fact, we can integrate (2.4) by parts to obtain:

$$\frac{\partial P(V)}{\partial V} = -\frac{1}{2}\frac{\partial g(\widehat{s}^c, V)}{\partial s}$$

The optimal V is determined by the slope of $g(s, V)$ at the conjectured policy parameter, which give us Lemma 2.1. ■

Proof of Proposition 2.2. The logic of the proof is as follows. We first solve the authority's problem ignoring the constraint about the firm's behavior. Then, we find under which conditions the constraint will be binding.

Ignoring the constraint; the first order condition for a given variance $V^*(\mu, \widehat{s})$ is:

$$\frac{\partial EW(\widehat{s})}{\partial \widehat{s}} = -\frac{\widehat{s}\eta^2 + \mu V^*(\mu, \widehat{s})}{\eta^2 + V^*(\mu, \widehat{s})} g\left(\widehat{s}, V^*\left(\mu, \widehat{s}\right)\right) = 0 \tag{2.5}$$

And the second order condition is:

$$\begin{aligned}\frac{\partial^2 EW(\widehat{s})}{\partial \widehat{s}^2} &= \frac{\widehat{s}\eta^2 + \mu V^*(\mu, \widehat{s})}{\eta^2 + V^*(\mu, \widehat{s})} g\left(\widehat{s}, V^*\left(\mu, \widehat{s}\right)\right) \frac{\widehat{s} - \mu}{\eta^2 + V} \\ &\quad - \frac{\eta^2}{\eta^2 + V^*(\mu, \widehat{s})} g\left(\widehat{s}, V^*\left(\mu, \widehat{s}\right)\right) < 0\end{aligned}$$

The second order condition is locally satisfied because the first term is zero when the first order condition is satisfied and the second term is always negative.

Note that (2.5) is zero only if $\widehat{s}\eta^2 + \mu V^*(\mu, \widehat{s}) = 0$, hence $\widehat{s}^{**} = \widehat{s}^*$. When $\mu > \widehat{s}$, by Lemma 2.2, the firm chooses V_L. The signal has the maximum level of informativeness, therefore the constraint in (2.2) does not bind and the solution is equal to $\widehat{s}^*$. Conversely, when $\mu < \widehat{s}$, the firm chooses V_H and R may prefer the constrained solution $\widehat{s}^{**} = \mu$ (by Lemma 2.2, this is the minimum threshold that induces the firm to choose V_L instead of V_H) due to the resulting increase in the informativeness of the signal. In particular, the constrained solution will be chosen if the authority's objective function evaluated at this point is larger than evaluated at $\widehat{s}^*$:

$$\int_{\frac{-\mu V_H}{\eta^2}}^{+\infty} \frac{s\eta^2 + \mu V_H}{\eta^2 + V_H} g\left(s, V_H\right) ds < \int_{\mu}^{+\infty} \frac{s\eta^2 + \mu V_L}{\eta^2 + V_L} g\left(s, V_L\right) ds \tag{2.6}$$

Let us show that (2.6) is satisfied for $\mu > \widetilde{\mu}$ and not satisfied for $\mu < \widetilde{\mu}$, where $\widetilde{\mu}$ is to be defined. Integrating by parts, rewrite (2.6) as:

$$\mu\left[\frac{1}{2}-G\left(\frac{-\mu V_H}{\eta^2},V_H\right)\right]<\eta^2\left[\frac{1}{\sqrt{2\pi(\eta^2+V_L)}}-g\left(\frac{-\mu V_H}{\eta^2},V_H\right)\right]$$

We want to determine whether there exists some value of μ in the interval $(-\infty,0]$ for which (2.6) is satisfied. Note that (2.6) is trivially satisfied when $\mu=0$ as the left-hand side is zero and $g(0,V_L)-g(0,V_H)>0$. Conversely, when $\mu\to-\infty$, the left-hand side tends to $+\infty$ while the right-hand side tends to $\frac{\eta^2}{\sqrt{2\pi(\eta^2+V_L)}}$. Therefore the inequality is violated. This means that there exists at least one value of μ, $\widetilde{\mu}$, such that the expected welfare are equal:

$$\int_{\frac{-\widetilde{\mu}V_H}{\eta^2}}^{+\infty}\frac{s\eta^2+\widetilde{\mu}V_H}{\eta^2+V_H}g(s,V_H)ds=\int_{\widetilde{\mu}}^{+\infty}\frac{s\eta^2+\widetilde{\mu}V_L}{\eta^2+V_L}g(s,V_L)ds \tag{2.7}$$

Note that $\widetilde{\mu}\neq 0$. In order to show that this value is unique, we need to show that the slope of the difference of expected welfare is strictly decreasing at $\mu=\widetilde{\mu}$:

$$\left.\frac{\partial\left(\mu\left[\frac{1}{2}-G\left(\frac{-\mu V_H}{\eta^2},V_H\right)\right]-\eta^2\left[\frac{1}{\sqrt{2\pi(\eta^2+V_L)}}-g\left(\frac{-\mu V_H}{\eta^2},V_H\right)\right]\right)}{\partial\mu}\right|_{\mu=\widetilde{\mu}}$$

$$\begin{aligned}
&= \frac{1}{2}-G\left(\frac{-\mu V_H}{\eta^2},V_H\right)-\mu g\left(\frac{-\mu V_H}{\eta^2},V_H\right)\left[\frac{-V_H}{\eta^2}-1\right]+\\
&\quad \eta^2 g\left(\frac{-\mu V_H}{\eta^2},V_H\right)\left[\frac{-V_H}{\eta^2}\frac{\left(-\left(\frac{-\mu V_H}{\eta^2}-\mu\right)\right)}{\eta^2+V_H}+\frac{\left(\frac{-\mu V_H}{\eta^2}-\mu\right)}{\eta^2+V_H}\right]\\
&= \frac{1}{2}-G\left(\frac{-\mu V_H}{\eta^2},V_H\right)<0
\end{aligned}$$

This expression is negative for $\widetilde{\mu}<0$, therefore there is a unique $\widetilde{\mu}$ for which condition (2.7) holds.

The comparative statics with respect to V_L:

$$\frac{d\widetilde{\mu}}{dV_L}=\frac{\eta^2}{\left[G\left(\frac{-\widetilde{\mu}V_H}{\eta^2},V_H\right)-\frac{1}{2}\right]2(\eta^2+V_L)\sqrt{2\pi(\eta^2+V_L)}}>0$$

and with respect to V_H:

$$\frac{d\widetilde{\mu}}{dV_H} = \frac{g\left(\frac{-\widetilde{\mu}V_H}{\eta^2}, V_H\right)\left[\frac{\widetilde{\mu}^2 2\left(\eta^2+V_H\right)+\eta^4}{2\eta^2(\eta^2+V_H)}\right]}{\frac{1}{2} - G\left(\frac{-\widetilde{\mu}V_H}{\eta^2}, V_H\right)} < 0$$

noting that $\widetilde{\mu} < 0$ and that therefore $G\left(\frac{-\widetilde{\mu}V_H}{\eta^2}, V_H\right) - \frac{1}{2} > 0$ give us the signs. ■

Proof of Proposition 2.3. We compute the first and the second order conditions and show that $P^{\alpha}(V)$ is quasi-convex until $\mu = s_M$ and then quasi-concave. The first derivative of $P^{\alpha}(V)$ with respect to V is:

$$\begin{aligned}
\frac{\partial P^{\alpha}(V)}{\partial V} &= -\frac{1}{2}\left(\alpha\frac{\partial g(\widehat{s}_1, V)}{\partial s} + (1-\alpha)\frac{\partial g(\widehat{s}_2, V)}{\partial s}\right) \\
&= \frac{1}{2}\left(\alpha g(\widehat{s}_1, V)\frac{\widehat{s}_1 - \mu}{\eta^2 + V} + (1-\alpha)\, g(\widehat{s}_2, V)\frac{\widehat{s}_2 - \mu}{\eta^2 + V}\right)
\end{aligned}$$

where $g(s, V)$ is always non-decreasing at $\widehat{s}_1$ and non-increasing at $\widehat{s}_2$. When $\mu = s_M$, $g(\widehat{s}_2, V) = g(\widehat{s}_1, V) = g(\widehat{s}_i, V)$:

$$\begin{aligned}
\frac{\partial P^{\alpha}(V)}{\partial V} &= \frac{1}{2}g(\widehat{s}_i, V)\left(\alpha\frac{\widehat{s}_1 - \mu}{\eta^2 + V} + (1-\alpha)\frac{\widehat{s}_2 - \mu}{\eta^2 + V}\right) \\
&= \frac{\widehat{s}_2 - \widehat{s}_1}{4(\eta^2 + V)}g(\widehat{s}_i, V)(1 - 2\alpha)
\end{aligned}$$

where the second line follows from replacing μ by s_M. Thus, the sign will depend on α.

From $\frac{\partial P^{\alpha}(V)}{\partial V} = 0$, we find the "critical variance":

$$V(\mu, \alpha) = \frac{(\widehat{s}_2 - \mu)^2 - (\mu - \widehat{s}_1)^2}{2\ln\left(\frac{(1-\alpha)(\widehat{s}_2 - \mu)}{\alpha(\mu - \widehat{s}_1)}\right)} - \eta^2 \tag{2.8}$$

The second derivative is:

$$\begin{aligned}
\frac{\partial^2 P^{\alpha}(V)}{\partial V^2} &= \frac{\alpha\frac{\partial g(\widehat{s}_1, V)}{\partial V}(s1 - \mu) + (1-\alpha)\frac{\partial g(\widehat{s}_2, V)}{\partial V}(s2 - \mu)}{2(\eta^2 + V)} \\
&\quad - \frac{\alpha g(\widehat{s}_1, V)(s1 - \mu) + (1-\alpha)\, g(\widehat{s}_2, V)(s2 - \mu)}{2(\eta^2 + V)^2}
\end{aligned}$$

which evaluated at the first order condition, give us:

$$\begin{aligned}
\left.\frac{\partial^2 P^\alpha(V)}{\partial V^2}\right|_{\frac{\partial P^\alpha(V)}{\partial V}=0} &= \frac{\alpha g(\widehat{s}_1, V)(\widehat{s}_1-\mu)^3 + (1-\alpha) g(\widehat{s}_2, V)(\widehat{s}_2-\mu)^3}{(\eta^2+V)^2} \\
&\propto (1-\alpha)(\widehat{s}_2-\mu)^3 \left(\frac{1-\alpha}{\alpha}\frac{\widehat{s}_2-\mu}{\mu-\widehat{s}_1}\right)^{\frac{-(\widehat{s}_2-\mu)^2}{(\widehat{s}_2-\mu)^2-(\mu-\widehat{s}_1)^2}} \\
&\quad -\alpha(\mu-\widehat{s}_1)^3 \left(\frac{1-\alpha}{\alpha}\frac{\widehat{s}_2-\mu}{\mu-\widehat{s}_1}\right)^{\frac{-(\mu-\widehat{s}_1)^2}{(\widehat{s}_2-\mu)^2-(\mu-\widehat{s}_1)^2}} \qquad (2.9)
\end{aligned}$$

where the last line follows from plugging in $V(\mu,\alpha)$ from (2.8). The sign of (2.9) depends on how $\left(\frac{\widehat{s}_2-\mu}{\mu-\widehat{s}_1}\right)^3$ compares to $\frac{\widehat{s}_2-\mu}{\mu-\widehat{s}_1}$. When $\mu = s_M$, both $\left(\frac{\widehat{s}_2-\mu}{\mu-\widehat{s}_1}\right)^3$ and $\frac{\widehat{s}_2-\mu}{\mu-\widehat{s}_1}$ are equal to one, hence (2.9) is zero. When $\mu \in (\widehat{s}_1, s_M)$, $\frac{\widehat{s}_2-\mu}{\mu-\widehat{s}_1}$ is $\left(\frac{\widehat{s}_2-\mu}{\mu-\widehat{s}_1}\right)^3$, therefore, (2.9) is positive. Conversely, when $\mu \in (s_M, \widehat{s}_2)$, $\frac{\widehat{s}_2-\mu}{\mu-\widehat{s}_1}$ is larger and, thus, (2.9) is negative. ■

Proof of Corollary 2.2. First note that for $\alpha = \frac{1}{2}$, (2.3) can be rewritten as:

$$P^{\frac{1}{2}}(V) = 1 - \frac{1}{2}\left[G\left(\widehat{s}_1, V\right) + G\left(\widehat{s}_2, V\right)\right]$$

From Proposition 2.3, if $\mu \in (\widehat{s}_1, s_M)$, the optimal variance is either maximum or minimum variance. In particular, M chooses V_L whenever:

$$G(\widehat{s}_1, V_H) + G(\widehat{s}_2, V_H) > G(\widehat{s}_1, V_L) + G(\widehat{s}_2, V_L) \qquad (2.10)$$

Using the symmetry of $g(s,V)$ around μ, (2.10) becomes:

$$G(\widehat{s}_1, V_H) - G(2\mu-\widehat{s}_2, V_H) > G(\widehat{s}_1, V_L) - G(2\mu-\widehat{s}_2, V_L) \qquad (2.11)$$

Note that $2\mu - \widehat{s}_2 \leq \widehat{s}_1$ for all $\mu \in (\widehat{s}_1, s_M)$.

Define $\widetilde{s}$ as the smallest value of s at which $g(s, V_H)$ and $g(s, V_L)$ intersect:

$$\widetilde{s} = \mu - \sqrt{\frac{(V_L+\eta^2)(V_H+\eta^2)}{V_H-V_L}\ln\left(\frac{V_H+\eta^2}{V_L+\eta^2}\right)} \qquad (2.12)$$

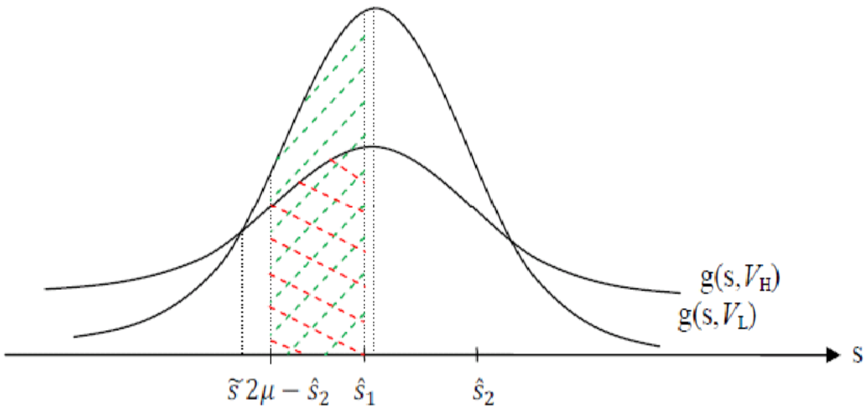

Figure 2.5: $\mu = \widehat{s}_1 + \Delta$ and $\widetilde{s} < 2\mu - \widehat{s}_2$

Rewrite $\widetilde{s}$ as $\widetilde{s} = \mu - c(V_L, V_H)$. It can be shown that $c(V_L, V_H)$ increases with V_H and decreases when V_L decrease.

We are only able to determine M's choices when $\widetilde{s} < 2\mu - \widehat{s}_2$ and when $\widehat{s}_1 < \widetilde{s}$; therefore, the conditions that follow are sufficient but not necessary.[43]

V_H will be chosen for all $\mu \in (\widehat{s}_1, s_M)$ if $\widetilde{s} < 2\mu - \widehat{s}_2$ as the area under $g(s, V_H)$ is unambiguously smaller than under $g(s, V_L)$ for the relevant range of integration (from $2\mu - \widehat{s}_2$ to $\widehat{s}_1$). See Figure 2.5 for an example. This condition implies:

$$\widehat{s}_2 - \mu < c(V_L, V_H)$$

V_L will be chosen if $\widehat{s}_1 < \widetilde{s}$, or equivalently, if:

$$c(V_L, V_H) < \mu - \widehat{s}_1.$$

Since $c(V_L, V_H) > 0$, this condition cannot hold for the smallest possible μ (i.e. $\widehat{s}_1$ in the limit); thus, when μ is close to $\widehat{s}_1$ M always chooses V_H. By continuity, there

[43] If $\widetilde{s} \in (2\mu - \widehat{s}_2, \widehat{s}_1)$, we cannot say how the area under $g(s, V_H)$ compares to the area under $g(s, V_L)$ for the relevant range of integration.

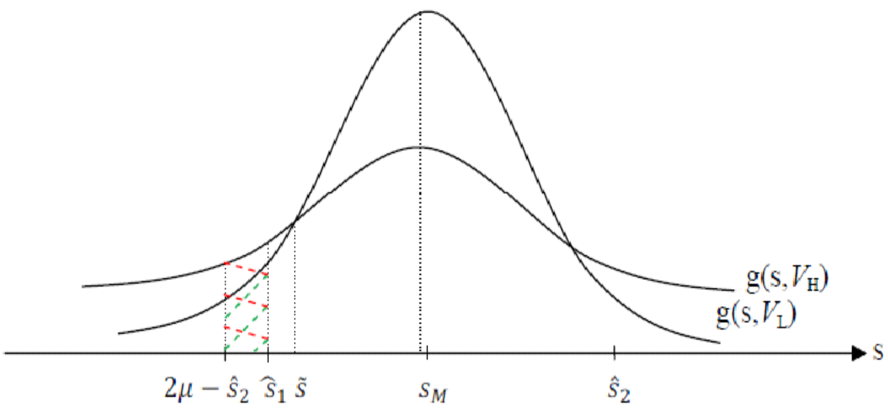

Figure 2.6: $\mu = s_M - \Delta$ and $\widehat{s}_1 < \widetilde{s}$

is a value $\overline{\mu}$ such that the firm is indifferent between V_H and V_L:

$$G(\widehat{s}_1, V_H) - G(2\overline{\mu} - \widehat{s}_2, V_H) = G(\widehat{s}_1, V_L) - G(2\overline{\mu} - \widehat{s}_2, V_L) \qquad (2.13)$$

Figure 2.6 depicts an example and shows how the area under $g(s, V_H)$ is always larger than under $g(s, V_L)$ for the relevant range of integration.

Figure 2.7 depicts the parameter space for which we can determine M's choices. In the are with triangles, M chooses V_H, in the area with inverted triangles, M chooses V_L and finally in the shaded area, M will choose either V_H or V_L, but it is not possible to draw the exact choice boundary analytically. ■

Proof of Proposition 2.4. Note that, using the continuity of $\rho(s)$, we can rewrite the objective function as: $P^c(V) = 1 - \int_{\widehat{s}_1}^{\widehat{s}_2} \rho'(s) \int_{-\infty}^{s} g(\widetilde{s}, V) d\widetilde{s} ds$. Then, the first derivative is:

$$\begin{aligned} \frac{\partial P^c(V)}{\partial V} &= -\int_{\widehat{s}_1}^{\widehat{s}_2} \rho'(s) \int_{-\infty}^{s} \frac{\partial g(\widetilde{s}, V)}{\partial V} d\widetilde{s} ds \\ &= -\frac{1}{2} \int_{\widehat{s}_1}^{\widehat{s}_2} \rho'(s) \frac{\partial g(s, V)}{\partial s} ds \end{aligned}$$

Since $\rho(s)$ has a rotational symmetry with respect to $\left(s_M, \frac{1}{2}\right)$, $\rho(s) + \rho(2s_M - s) = 1 \forall s$

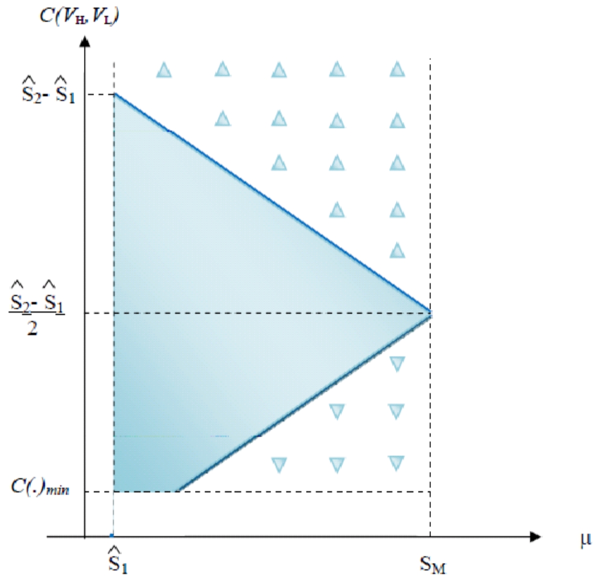

Figure 2.7: Choice of variance

and $\rho'(s) = \rho'(2s_M - s)\forall s$, where $\rho'(s)$ is even and positive. Using the second line, we can show that:

- If $\mu = s_M$ then $\frac{\partial P^c(V)}{\partial V} = 0$. This is because $\rho'(s)$ is even (and hence symmetric around s_M) and positive and $g'(s)$ is odd with s_M as the origin.

- If $\mu = \widehat{s}_1$ then $\frac{\partial P^c(V)}{\partial V} > 0$. This is because $\rho'(s)$ is positive and $g'(s) \leq 0\forall s\,[\widehat{s}_1, \widehat{s}_2]$.

- If $\mu = \widehat{s}_2$ then $\frac{\partial P^c(V)}{\partial V} < 0$. This is because $\rho'(s)$ is positive and $g'(s) \geq 0\forall s\,[\widehat{s}_1, \widehat{s}_2]$.

Now we need to prove that $\frac{\partial^2 P^c(V)}{\partial V \partial \mu} < 0\forall \mu$.

$$
\begin{aligned}
\frac{\partial^2 P^c(V)}{\partial V \partial \mu} &= -\frac{1}{2}\int_{s_1}^{s_2} \rho'(s)\,\frac{\partial^2 g(s,V)}{\partial s \partial \mu} ds \\
&= \frac{1}{2}\int_{s_1}^{s_2} \rho'(s)\,\frac{\partial^2 g(s,V)}{\partial s^2} ds
\end{aligned}
$$

since $\rho'(s)$ is positive, a sufficient condition for $\frac{\partial^2 P^c(V)}{\partial V \partial \mu} < 0$ is $\frac{\partial^2 g(s,V)}{\partial s^2} < 0$, which happens when $\mu - \sqrt{\eta^2 + V} < \{s_1, s_2\} < \mu + \sqrt{\eta^2 + V}$. The requirement of and to be close together is by no means necessary. For instance, it is easy to check that if $\rho(s) = \frac{s-\widehat{s}_1}{\widehat{s}_2-\widehat{s}_1}$, $\frac{\partial^2 P}{\partial V \partial \mu} < 0 \forall \{s_1, s_2\}$. Example 2.1 provides another instance where $s_1 = -\infty$ and $s_2 = +\infty$. ■

Bibliography

[] ANDERSON, A., and CABRAL, L. M. B., (2007), "Go for Broke or Play It Safe? Dynamic Competition with Choice of Variance", *The RAND Journal of Economics*, Vol. 38, No. 3, pp. 593-609.

[] CALZOLARI, G., and PAVAN, A., (2006), "On the optimality of privacy in sequential contracting," *Journal of Economic Theory*, Vol. 130, No. 1, pp. 168-204.

[] CASELLA, G., and BERGER, R.L., (2002), *Statistical Inference*, Duxbury Press, Pacific Grove.

[] DEWATRIPONT, M., JEWITT, I., and TIROLE, J., (1999), "The Economics of Career Concerns, Part I: Comparing Information Structures," *Review of Economic Studies*, Vol. 66, No. 1, pp. 183-198.

[] DRUGOV, M. and TROYA-MARTINEZ, M., (2012), "Vague Lies: How to Advise Consumers When They Complain," Working paper.

[] FARRELL, J. and GIBBONS, R., (1989): "Cheap Talk with Two Audiences," *American Economic Review,* Vol. 79, No. 5, pp. 1214-1223.

[] GABA, A., and KALRA, A., (1999), "Risk Behavior in Response to Quotas and Contests." *Marketing Science*, Vol. 18, No. 3, Special Issue on Managerial Decision Making, pp. 417-434.

[] GERTNER, R.,GIBBONS, R., and SCHARFSTEIN, D., (1988): "Simultaneous Signalling to the Capital and Product Markets," *The RAND Journal of Economics,* Vol. 19, No.2, pp. 173-190.

[] HOLMSTROM, B., (1982), "Managerial Incentive Schemes—A Dynamic Perspective." In Essays in Economics and Management in Honour of Lars Wahlbeck. Helsinki: Swenska Handelshogkolan.

[] ———, (1999), "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies,* Vol. 66, No. 1, pp. 169–182.

[] JOHNSON, J.P., and MYATT, D.P., (2006), "On the Simple Economics of Advertising, Marketing, and Product Design." *American Economic Review,* Vol. 96, No. 3, pp. 756-784.

[] LEWIS, T.R., and SAPPINGTON, D.E.M., (1994) "Supplying Information to Facilitate Price Discrimination", *International Economic Review*, 35, 309-327.

[] MAIER, N., and OTTAVIANI, M., (2009), "Information Sharing in Common Agency: When is Transparency Good?," *Journal of the European Economic Association*, Vol. 7, No. 1, pp. 162-187.

[] MARSDEN, P., and WHELAN, P., (2005), "The Contribution Of Bilateral Trade Or Competition Agreements To Competition Law Enforcement Cooperation Between Canada And Costa Rica." Paper presented at the OECD Joint Group on Trade and Competition, *COM/DAF/TD(2005)51.*

[] MILGROM, P. (1981), "Good News and Bad News: Representation Theorems and Applications". *Bell Journal of Economics,* Vol.12, pp. 380-91.

[] MURIS,T. J. (2001), "Merger Enforcement in a World of Multiple Arbiters," remarks before the *Brookings Institution Roundtable on Trade & Investment*, Washington, D.C., available at http://www.ftc.gov/speeches/muris/brookings/pdf.

[] OECD (2001), "Merger enforcement and international co-operation", OECD Global Forum on Competition, *CCNM/GF/COMP/WD(2001)1,* available at http://www.oecd.org/dataoecd/40/58/2491825.pdf.

[] OECD (2003), "Information Sharing in merger control procedures: Contribution by the United States", OECD Working Party No. 3 on International Co-operation, *DAFFE/COMP/WP3/WD(2003)25,* available at http://www.ftc.gov/bc/international/docs/compcomm/2003–Information%20Sharing%20in%20Merger%20Control%20Procedures.pdf.

[] SPIEGEL, Y. and SPULBER, D., (1997), "Capital Structure with Countervailing Incentives," *The RAND Journal of Economics,* Vol. 28, No. 1, pp. 1-24.

[] TSETLIN, I., GABA, A., and WINKLER, R.L., (2004), "Strategic Choice of Variability in Multiround Contests and Contests with Handicaps," *Journal of Risk and Uncertainty,* Vol. 29, No. 2, 143-158.

[] UNCTAD (2000), "World Investment Report: Cross-Border Mergers and Acquisitions and Development," *New York and Geneva: United Nations.*

[] VALENTINE, D. A. (2000), "Cross Border Canada/US Co-operation in Investigations and Enforcement Actions", remarks before the Canada/U.S. Law Institute, Case Western Reserve University School of Law, available at http://www.ftc.gov/speeches/other/dvcrossborder.htm.

Chapter 3

Vague Lies: How to Advise Consumers When They Complain

3.1 Introduction

A buyer with limited or no experience walks into a shop in order to buy a mobile phone and sign up for a calling plan. There are many models of mobiles and different calling plans and, moreover, the buyer is not sure about her needs, which will only become fully known through the use of the mobile. The seller, who knows the product's features but may not know how they fit with the buyer's needs either (and hence the product quality of match), not only sets a price, but also provides some sales advice to the prospective buyer with the goal of convincing her to buy. In doing so, he can choose to be more or less precise and can also simply lie the buyer by exaggerating the goodness of the product. Sending a misleading sales pitch comes, however, at a cost of customers' complaints that will damage seller's reputation, draw the attention of a consumer association, trigger an action by the competition authority or even result in a litigation.

Some other goods also share these features. For example, consumer electronics, insurance, banking and medical care contracts are all characterized by the fact that neither sellers nor buyers themselves fully know the ex-ante buyers' quality of match with the product, sellers' advice is widely used and buyers learn through experience[1]. Also, these very products systematically top in the list of consumers complaints worldwide which brings about the policy relevance of the problem. For instance, the Federal Trade Commission in the United States has released a report listing top complaints consumers filed with the agency in 2009.[2] Among the top 15 places it is possible to find health care, internet services, credit cards, advance-fee loans, credit protection, banks and lenders in general and computer equipment and software. Similarly, in the United Kingdom, the Office of Fair Trading has reported (through Consumer Direct, its telephone and online service provider of information on consumer rights) that among the top 10 complaints, it is possible to find mobile phones (service agreements), mobile phones (hardware), telephone services (land line), lap-tops, notebooks and tablet PCs.[3]

We explore the seller's incentives to provide (un)biased and (im)precise advice and the resulting equilibrium communication in the following model. The seller advises the buyer about a certain product trying to increase the buyer's perceived valuation of it. The buyer's valuation for the product is determined by the quality of the match between the product characteristics (unknown to the buyer) and buyer's taste (unknown to the seller). The buyer and the seller share the same prior about quality of the match, that is, none of them has any private information.[4] The seller gives an advice (i.e., sends a signal) to the buyer about the product. The signal is the sum of the true match quality and an error term that represents frictions in the communication between the seller and the buyer.[5] Both the match quality and the error term are distributed normally and the seller secretly chooses the mean

[1] In the markets we have in mind, the seller may have private information on how his product compares to competitors but not on how it matches with each buyer's needs. Moreover, the buyer will only learn her actual needs through the use of the product.

[2] See: http://www.ftc.gov/opa/2010/02/2009fraud.shtm (accessed on the 06 December 2011).

[3] See: http://www.oft.gov.uk/shared_oft/general_policy/OFT1267.pdf pp. 21-24 (accessed on the 06 December 2011).

[4] An alternative interpretation is that the seller is facing the whole demand curve rather than an individual buyer as, for example, in Johnson and Myatt (2006).

[5] The error term can also represent the fact that the true match quality will be learned only through experience.

and publicly the variance of the error term.[6] We call them "bias" and "noise", respectively.

Upon receiving the signal, the buyer updates her beliefs about the match quality using her conjecture about the bias introduced by the seller (which has to be correct in the equilibrium). This posterior valuation of the product is the surplus generated if trade occurs and that the seller and the buyer share in a fixed proportion. This reflects the popular use of negotiation in selling banking or mobile phone contracts. While the bias unambiguously increases the perceived quality of the product and thus the seller's revenues, the effect of the noise is more subtle and depends on the particular model used.[7] We adopt a specification where the noise does not affect seller's revenues in the equilibrium, that is, when the buyer correctly anticipates the bias.

Misleading the buyers is costly since eventually, through use, they realize the true quality of the match and complain if the signal they received was much larger. These complaints trigger an action by an authority which might be a consumer protection authority (like Office of Fair Trading), a sectorial regulator (like Financial Services Authority) or the court depending on the product and the country. The authority then investigates the seller surveying a random sample of customers or sending mystery shoppers. Based on this information, it estimates the bias and determines whether there is enough evidence to conclude that the seller has misled the consumer by biasing his signal. More precisely, the authority presumes the innocence of the seller, that is, its null hypothesis is that there has been no bias.[8] It then tests whether this null hypothesis can be rejected in favor of the alternative hypothesis of a positive bias. In doing so, it uses some given significance level which can be interpreted as a standard of proof. If the seller is found guilty, he has to pay a fine that depends on the estimated bias. A larger bias always increases the costs, i.e.,

[6] We also consider the unobservable choice of variance in Section 3.5.1.

[7] See, for example, Lewis and Sappington (1994) and Johnson and Myatt (2006) where the seller always prefers to provide an extreme (i.e., maximum or minimum) amount of information.

[8] The presumption of innocence is natural in court. When it is a competition authority that punishes the seller the presumption of innocence is explained by the fact that the competition authority may need to defend its position in court if the seller decides to appeal.

the expected fine. The noise affects the costs through two channels: it decreases the probability that the seller is found guilty but increases the fine when the seller is found guilty. The total effect is U-shaped.[9]

We derive closed form solutions for the equilibrium bias and noise and find that they are complements. When the authority uses a stricter standard of proof or the match quality is more heterogenous, the seller both biases the signal more and makes it more noisy. For example, when buyers heterogeneity is higher, they put a higher weight on the seller's advice and, therefore, the seller has higher incentives to bias it. A higher noise then used since it is complementary to the bias. Introduction of punitive damages makes the seller send a less biased and more precise signal.

We also consider the case of a fixed price of the good. The goal of the seller is then to convince the marginal buyer to buy while in the baseline model it is to increase the posterior of every buyer since everybody buys there. In the special case when the price is equal to the prior mean, we derive closed-form solutions for equilibrium bias and noise and show that they are also complements. Interestingly, a higher buyer's heterogeneity decreases the equilibrium bias and noise, contrary to the baseline mode described above. When the price equals to the mean valuation of the good the only effect of a higher buyer's heterogeneity is to decrease the probability of the buyer with the average valuation, that is, the marginal buyer. Then, the marginal revenues of the bias are smaller and the seller uses it less (and correspondingly he introduces a lower noise). When the price is above (below) the mean valuation, the bias-to-noise ratio is also higher (lower). When the price is above the prior mean, the seller serves relatively high valuation consumers; the noise is relatively low since more information increases the posterior of the marginal high valuation consumer. In the case of the price below the prior mean, the marginal consumer has a relatively low valuation and providing information only decreases her posterior. Therefore, the noise is relatively high.

When two sellers sell related products and buyers can get informed from both of

[9]In many situations, especially with repeated purchases, reputational concerns also put a limit on the seller's bias. Our model is thus better suited for one-off or infrequent purchases.

them (although they are locked in and do not choose from whom to buy), there is a disciplining effect. Since each buyer attaches some weight to the other seller's signal, she puts less weight on the signal of her seller. Then, each seller has less incentives to bias his signal and the equilibrium bias and noise go down. Surprisingly, if there are also common shocks in the communication technology (that is, the error terms are correlated), the disciplining effect of the second seller can be reversed. The consumer trusts more her seller's signal since she gets the information about the error term from the second seller and, as a result, the equilibrium bias and noise may increase.

We consider a number of other extensions. In the baseline model the noise is observable by both the buyers and the authority. We consider the case of non-observable noise and show that the equilibrium has the same structure and, in particular, equilibrium bias and noise are still complements. In the baseline model the buyers are rational and are not misled in the equilibrium while in reality some consumers do appear naive, at least, to some extent (see examples in Gabaix and Laibson (2006) and Inderst and Ottaviani (2009a) among others). We consider an extension where a proportion of consumers are credulous and blindly follow the seller's advice. We find that the seller then sends a more biased and less precise signal.

3.1.1 Related literature

From the theory perspective, this model is related to the literature on career concerns (see Holmström (1999) and Dewatripont, Jewitt, and Tirole (1999)). In a two-period version of these models, the worker's expected second-period performance is his ability since he will not make any effort. The ability is not observed and firms estimate it based on the first-period worker's performance.[10] The worker then exerts some effort in the first period in order to improve his performance and trick firms into thinking that he has a higher ability, so that they pay him a larger wage.

[10] The output usually equals to the sum of the effort and the ability. See Dewatripont, Jewitt, and Tirole (1999) for a more general setup.

In the equilibrium the firms are not tricked; instead, they anticipate the worker's effort and estimate the worker's ability correctly. Our model is a signal-jamming model as the career concerns models. However, our seller not only jams the signal by manipulating its mean (what we call the "bias"), but also by changing its variance (the "noise") and, hence, its information content.

The choice of variability as a strategic variable has been considered in a large variety of setups.[11] The general conclusion of these papers is that the players that are at disadvantage tend to choose more risky strategies than those who are in a favorable position, that is, they "gamble for resurrection". In our model the seller's expected revenues are additive in the expected quality and hence his marginal incentives do not depend on it.

Johnson and Myatt (2006) consider how much product relevant information a monopolist would want to provide to its potential customers. Using the fact that information about the product rotates the demand curve, they show that the monopolist profits are convex in the information. This generalizes the result of Lewis and Sappington (1994) who showed that the monopolist supplies either maximal or minimal information. In both papers, the choice of the amount of information is costless for the seller, while in our setup this choice is endogenously linked with future fines and we also allow for the possibility of biasing the signal. Furthermore, both papers only consider the case when the amount of information is observable while we also analyze the unobservable case.[12] A different model is the one of Spiegler (2006) where firms send noisy, unbiased and costless signals of their prices

[11] For instance, Anderson and Cabral (2007) analyze a model of R&D races where there are two contestants who are allowed to choose the variance of their stochastic R&D technology. Tsetlin, Gaba, and Winkler (2004) consider the choice of variability of the performance distribution in a multi-round contest. Kräkel, Nieken, and Przemeck (2008) undertake a similar analysis in the auction context.

[12] To get an intuition of why observability matters, imagine that the product is bad on average, that is, in the absence of any information the buyer will not buy it (at a given price). To convince the buyer to buy, the seller has to generate a high realisation of the signal and wants the buyer to attach a high weight to this signal in her updating. If the noise is observable, adding noise to the signal will make the buyer attach a lower weight to the signal realization so the optimal strategy is actually to reduce the noise as much as possible (see Johnson and Myatt (2006) for the proof). However, if the noise is not observable, the buyer will not adjust the weight given to the signal to account for the noise and, therefore, the optimal strategy of the seller is to increase the noise of the signal.

as obfuscation strategies in order to soften competition. In our model, there is no competition, the signal is biased (and noisy) and costly.

The idea that introducing noise might make the bias less costly is also present in Li (2010) where the sender shares the blame for the wrong message with the noisy and possibly biased intermediary. In Blume and Board (2010) noisy messages are useful in mitigating the conflict between the sender and the receiver in a cheap-talk model.

As we discussed in the beginning, this paper is motivated by the markets where consumers often complain and, therefore, is related to the literature on consumer protection.[13] Inderst and Ottaviani (2009b), for instance, analyze contract cancellation and product return policies in a market where a more informed seller advises the buyer. Cancellation or return are costly for the seller and, therefore, the advice cannot be too misleading. Since advice in their model takes the form "to buy" or "not to buy", there is no room for the seller to use noise, only bias.

The revenue part of our model can be seen as a cheap-talk model. The interests of the buyer and the seller are opposed since the seller always wants to charge a higher price while the buyer always wants to pay a lower price. Then, there is no information transmission in the equilibrium as is well known starting from the work of Crawford and Sobel (1982).[14] Then, some exogenous (unmodelled) costs are usually added such as effort costs in career concerns models, refund costs in Inderst and Ottaviani (2009b), or unspecified lying costs in Kartik (2009). Contrary to these papers, in our model the lying costs, which are the expected fine imposed on the seller, are microfounded.

The rest of the paper is organized as follows. Section 3.2 introduces the model. Section 3.3 finds the equilibrium when the seller can price discriminate and derives comparative statics results. Section 3.4 consider the case of fixed prices. Section 3.5 contains various extensions and modifications of the baseline model of Section 3.2.

[13] See Vickers (2004) and Armstrong (2008) for an overview on this literature.

[14] See, however, Chakraborty and Harbaugh (2010) where information transmission is possible in a multidimensional cheap-talk model despite opposed preferences.

Section 3.6 concludes.

3.2 The model

Buyers approach the seller in order to get informed about his product and buy it. The seller does not know their valuations for the product while the buyers do not know the quality of the product and its features. Thus, at the beginning of the interaction, the match quality of the transaction, θ, is unknown both to the seller and the buyer. However, they know that θ is distributed as $\mathcal{N}(\mu, \sigma^2)$.[15] The production costs are normalized to zero.

3.2.1 The communication process

The seller reveals product characteristics and gives advice about the shopping decision to the buyer. In doing so, he can distort the communication strategically. The seller can exaggerate some positive features of the product and he can also be vague about them. More precisely, the seller sends an informative but possibly biased and noisy signal S which takes the following form:

$$S = \theta + \varepsilon,$$

where ε is the distortion introduced by the seller. It is distributed normally, $\varepsilon \rightsquigarrow \mathcal{N}(\beta, \eta^2)$, and both moments are controlled by the seller.[16] We refer to β as *bias* and to η^2 as *noise* and we assume that they are both chosen from a bounded interval. The signal is therefore distributed as $\mathcal{N}(\mu + \beta, \sigma^2 + \eta^2)$; denote its pdf by g. The buyer does not observe the bias since she cannot know if a certain feature is exaggerated without actually buying the product and trying it. She does observe the noise, however, since she can evaluate how precise the seller's explanations are, how much

[15] As standard in career concerns models, we adopt a setting where the two sides start with the same prior to avoid issues of signalling. See Holmström (1999) and Dewatripont, Jewitt, and Tirole (1999). See **?** for an example in the advertising literature.

[16] The might be some natural noise in the communication process. We discuss this possibility in Section 3.5.6.

into details he goes, whether there is a trial period, etc.[17]

3.2.2 Buyers' valuation and seller's revenues

The buyer values a better quality of match of the product, θ. For simplicity, the buyer's valuation of the product is linear in the match quality θ and normalized to be equal to it. Since this quality is only known through use, her valuation at the moment of purchase is the expected match quality given a particular realization s of the signal.

When the buyer observes a realization s of the signal and the noise η^2, she updates the expected quality using a conjecture about the seller's bias $\widetilde{\beta}$[18]. The posterior distribution of θ stays Normal with mean:[19]

$$E\left[\theta \mid s, \eta, \widetilde{\beta}\right] = \mu + \left(s - \mu - \widetilde{\beta}\right) \frac{\sigma^2}{\sigma^2 + \eta^2}. \tag{3.1}$$

The buyer takes the signal into account with the weight proportional to the prior variance. It can also be written as $\left(\frac{\mu}{\sigma^2} + \frac{s-\widetilde{\beta}}{\eta^2}\right) / \left(\frac{1}{\sigma^2} + \frac{1}{\eta^2}\right)$, that is, as the weighted average of the ex-ante quality of match and the ex-post signal realization (corrected for the bias), where the weights are precisions of the prior and the signal.

Instead of fixing the price before sending the signal, the seller charges the buyer some (fixed) fraction of her valuation which is normalized to 1. That is, the seller sees the effect that the realization of his signal, s, has had on the buyer and hence, he knows that the buyer is ready to pay up to $E\left[\theta \mid s, \eta^2, \widetilde{\beta}\right]$ for the product. This entails some negative payments which can be made negligible by making the average quality μ high enough.[20] It also means that the seller always sells to the consumer. This is in the spirit of the career concerns models, where a worker always gets a

[17] We discuss the case of unobservable noise in Section 3.5.1.

[18] We assume that the buyer's conjecture does not depend on the observed noise. This assumption is used in the literature, see, for example, Judd and Riordan (1994). Also, this makes the model with unobservable noise (Section 3.5.1) similar to the model of this section. It would be interesting, though, to investigate other possibilities since different equilibria are then obtained as is shown by Shelegia (2011).

[19] For the derivation of the Normal Bayes estimators see Casella and Berger (2002), p. 326, Example 7.2.16.

[20] As we will see later, μ does not affect equilibrium bias and noise (Corollary 3.1).

wage. The only reason it is done both there and in this paper is technical as it allows to integrate over the whole support of the distribution.[21]

Therefore, the seller's revenues are equal to the expected valuation of the buyer:

$$R(\beta, \eta) = \int_{-\infty}^{+\infty} E\left[\theta \mid s, \eta, \widetilde{\beta}\right] g(s)ds = \mu + \left(\beta - \widetilde{\beta}\right) \frac{\sigma^2}{\sigma^2 + \eta^2}. \tag{3.2}$$

Note that the seller extracts from the buyer the prior expected quality of match μ plus how much the buyer is misled into thinking that the product is better than it is, adjusted for the weight. Note as well, that the revenues decrease with the noise η (if $\beta > \widetilde{\beta}$) since the buyers pay less attention to the signal, i.e., place a smaller weight on it. In equilibrium, however, when the buyer's conjecture about the bias used is correct, the seller will only be able to extract the prior expected match quality μ. As a result, the noise has no effect on the revenues.

3.2.3 Estimated bias and seller's costs

When the buyers buy the product and start using it, they discover the true match quality. Some buyers, especially those with a large gap between the advice (signal) received from the seller and the true quality, will complain to a public body that we call "authority" throughout the paper. This might be a consumer association, consumer protection agency, a competition authority, an industry regulator or the court, depending on the good or service in question and the country. We assume that this authority can inflict a punishment on the seller. Depending on the nature of the authority, the punishment may be publishing a negative report, ordering to withdraw a certain advertisement, prohibiting a certain commercial practice or imposing a fine on the seller. For the sake of simplicity, we adopt the last meaning and treat the punishment as a monetary fine.

The bias, as a deliberate and conscious way to mislead consumers, is illegal and a fine is imposed on the seller if any positive bias is discovered. In the European

[21] In a different model, Spiegler (2006) also allows negative prices for technical reasons.

Union, the Unfair Commercial Practices Directive 2005/29/EC defines a commercial practice as misleading "...if it contains false information and is therefore untruthful or in any way, including overall presentation, deceives or is likely to deceive the average consumer..." (Article 6). Such practices are more generally called unfair and "...shall be prohibited" (Article 5).[22]

The noise, however, is not punished since it may come from some (unmodelled) external shocks in the communication process between the seller and the buyers. The seller may provide information that is vague and open to interpretation, so that the buyers themselves make personal, independent, errors of interpretation. Even if the seller does not bias the signal at all, there will be always some unlucky buyers that will receive a very high signal realization as compared to their true match quality (unless the communication is absolutely noiseless). For instance, a bad quality of the information at the point of purchase may be due to incompetent sales staff who have trouble giving a clear advice. If these incompetent sales staff are nonetheless objective (i.e., they do not use bias), the buyers will not be misled on average.

The authority conducts an investigation by taking a random sample of size N of buyers in order to estimate the bias introduced by the seller. It can also send "mystery shoppers" to the seller each of them reporting afterwards their experience. In the US and in the UK the competition and consumer protection authorities (FTC and OFT, respectively) routinely commission mystery shopping exercises as part of their market studies.[23]

Each mystery shopper (or buyer) $i = 1, ..., N$ reports the signal s_i received from

[22] In the common law tradition, misrepresentation is a contract law concept which means a false statement of fact made by one party to another party, which has the effect of inducing that party into the contract. For example, under certain circumstances, false statements or promises made by a seller of goods regarding the quality or nature of the product that the seller has may constitute misrepresentation. In English law, misrepresentation is regulated generally by the United Kingdom Misrepresentation Act 1967 and, in the consumer protection area, by Trade Descriptions Act 1968 and Consumer Protection Act 1987.

[23] See http://www.ftc.gov/os/2009/12/P994511violententertainment.pdf and http://www.oft.gov.uk/shared_oft/business_leaflets/credit_licences/OFT1265.pdf for recent examples (accessed on the 06 December 2011). Other tests are often conducted by courts and consumer bodies. For instance, "copy tests" are used to determine whether an advert is misleading. If enough consumers are misled, the consumer protection authority may order the advert to be withdrawn. The evidence may also come from a class action, see Issacharoff (1999) for a discussion of class actions and consumer protection.

the seller and quality θ_i.[24] In Section 3.5.5 we consider the case when only signals are observed. Having received their reports, the authority uses a statistical test to determine if the seller is guilty of having introduced the bias.[25] It computes the error terms $\varepsilon_i = s_i - \theta_i$ and estimates the bias used as:

$$\widehat{\beta} = \frac{1}{N} \sum_{i=1}^{N} \varepsilon_i \tag{3.3}$$

Since $\varepsilon_i \rightsquigarrow \mathcal{N}(\beta, \eta^2)$, this estimator is distributed as $\mathcal{N}\left(\beta, \frac{\eta^2}{N}\right)$.

There is the presumption of innocence, that is, by default the seller is assumed not to have introduced any bias, unless enough evidence is provided. In order to determine how convincing the evidence about the use of bias should be the authority uses a standard hypothesis test where the null hypothesis of no bias, $H_0 : \beta = 0$, is assessed against the alternative $H_1 : \beta > 0$. The authority constructs the statistics $\frac{\widehat{\beta}}{\eta/\sqrt{N}}$ which, under the null H_0, is distributed as $\mathcal{N}(0, 1)$. Denote z_α the threshold such that the seller is found guilty of biasing if and only if $\frac{\widehat{\beta}}{\eta/\sqrt{N}} \geq z_\alpha$, where α is the significance level of the test (i.e., the probability of incorrectly rejecting the null hypothesis).[26] A natural interpretation of z_α is the "standard of proof". A higher z_α (lower α) means that it is more difficult to reject the null, and therefore, the authority needs more evidence to be convinced. For instance, if α is 5%, then $z_\alpha \approx 1.64$, and if α is 10%, then $z_\alpha \approx 1.28$.

If the seller is found guilty, he is imposed a fine which is an increasing function of the estimated bias $\widehat{\beta}$. In this section, we take the fine to be equal to $\widehat{\beta}$ and consider

[24] For instance, the authority could provide the shopper with a list of relevant attributes of the product. Then, for each attribute, the shopper needs to report what he/she expected from the advice and what he/she really found. This could be done in relatively coarse terms (i.e. better, worse or as expected) and then the authority could aggregate the answers to obtain a measure of s_i and θ_i.

[25] Scholars in law and economics have long been modelling court decision making as probabilistic, see Polinsky and Shavell (2000) for a survey. In legal literature it is also "... now generally accepted that since all the evidence is probabilistic ... evidence should not be excluded merely because its accuracy can be expressed in explicitly probabilistic terms..." (Posner (2004), p. 370). See also Miceli (1990) and Davis (1994) for describing statistical testing in relationship with court decision making.

[26] By definition, the probability that the standard normal random variable takes a value above z_α is equal to α.

punitive damages $d\widehat{\beta}$ in Section 3.5.4. The seller's costs are the expected fine:

$$C(\beta,\eta) = \frac{\eta}{\sqrt{N}}\int_{z_\alpha}^{+\infty} z dH(z) \tag{3.4a}$$

$$= \beta\left(1-\Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right) + \frac{\eta}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right) \tag{3.4b}$$

where $z \rightsquigarrow \mathcal{N}\left(\beta\frac{\sqrt{N}}{\eta},1\right)$ and H is its cdf (it is the distribution of $\frac{\widehat{\beta}}{\eta/\sqrt{N}}$); and Φ and ϕ are cdf and pdf of the standard normal random variable, respectively. The first term in (3.4b) corresponds to how much the seller pays on average multiplied by the probability of being found guilty. The second term corrects for the selection bias as the truncation selects higher values of $\widehat{\beta}$. The crucial property of this cost function is that it is U-shaped with respect to the noise. A higher η makes it less likely for the seller to be found guilty (since the distribution of z shifts to the left, the integral in (3.4a) decreases), but increases the chances of a really large fine once he is found guilty (this integral is multiplied by a larger number). This is the next lemma. See also Figure 3.1.

Lemma 3.1 *The seller's costs (3.4b) are increasing in β and U-shaped with respect to η.*

Proof. The first derivative of (3.4b) with respect to β is

$$\frac{\partial C(\beta,\eta)}{\partial\beta} = 1-\Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right) + z_\alpha\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)$$

and it is always positive. The first derivative of (3.4b) with respect to η is

$$\frac{\partial C(\beta,\eta)}{\partial\eta} = \frac{1}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left[1-\frac{\beta\sqrt{N}}{\eta}z_\alpha\right]$$

and it is negative for $\eta < \beta\sqrt{N}z_\alpha$ and it is positive for $\eta > \beta\sqrt{N}z_\alpha$. ■

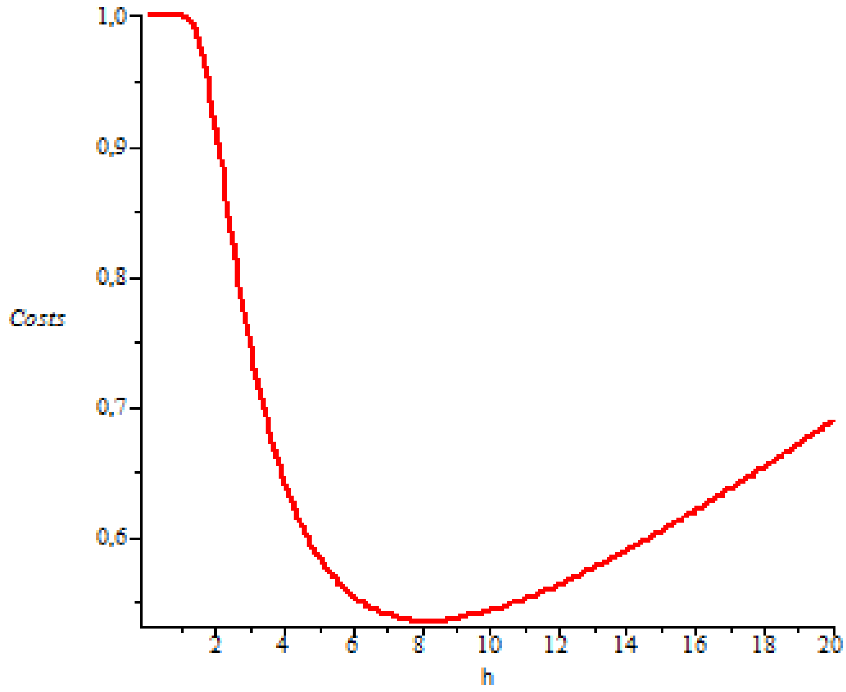

Figure 3.1: Seller's costs as a function of noise η; $\beta = 1, N = 25, z_\alpha = 1.64$.

3.3 Equilibrium

The seller maximizes his profits which are equal to the revenues (3.2) minus the costs (3.4b):

$$\Pi(\beta, \eta) = \mu + \left(\beta - \widetilde{\beta}\right) \frac{\sigma^2}{\sigma^2 + \eta^2} - \beta \left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right) - \frac{\eta}{\sqrt{N}} \phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right).$$

The equilibrium is a pair (β^*, η^*) that maximizes seller's profits $\Pi(\beta, \eta)$ given buyer's conjecture β^*. Next proposition derives the equilibrium in a closed form. For the second-order conditions to hold, we need the very mild assumption that $z_\alpha < 2.436$, that is, $\alpha > 0.75\%$, .

Proposition 3.1 *The equilibrium bias and noise are*

$$\beta^* = \frac{\sigma}{\sqrt{N}} \frac{1}{z_\alpha} \sqrt{\kappa}, \tag{3.5a}$$

$$\eta^* = \sigma\sqrt{\kappa}, \tag{3.5b}$$

where $\kappa = \frac{\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) - z_\alpha \phi\left(z_\alpha - \frac{1}{z_\alpha}\right)}{1 - \Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha \phi\left(z_\alpha - \frac{1}{z_\alpha}\right)}$.

Proof. See Appendix. ■

The noise does not affect the revenues in the equilibrium since the buyer is not misled and pays the prior expected match quality μ.[27] Thus, for any bias, the noise can be seen as minimizing the costs for that bias. As the cost function is U-shaped with respect to the noise (Lemma 3.1) this yields an interior solution $\frac{\beta\sqrt{N}}{\eta} = \frac{1}{z_\alpha}$.[28] The bias increases the revenues by changing the signal distribution in the first-order stochastic dominance sense; the marginal effect of the bias on the revenues is the weight of the signal in the buyer's posterior, $\frac{\sigma^2}{\sigma^2+\eta^2}$. The marginal cost of the bias depends only on the ratio $\frac{\beta}{\eta}$ which is pinned down from the cost-minimization with respect to the noise. A closed-form solution can then easily be found.

The comparative statics results are summarized in the next corollary.

Corollary 3.1 *The comparative statics of the equilibrium bias and noise (3.5a-3.5b) is the following:*

- $\frac{\partial \beta^*}{\partial \mu} = \frac{\partial \eta^*}{\partial \mu} = 0$;
- $\frac{\partial \beta^*}{\partial \sigma^2}, \frac{\partial \eta^*}{\partial \sigma^2} > 0$;
- $\frac{\partial \beta^*}{\partial z_\alpha}, \frac{\partial \eta^*}{\partial z_\alpha} > 0$;
- $\frac{\partial \beta^*}{\partial N} < 0$ *and* $\frac{\partial \eta^*}{\partial N} = 0$.

Given that a change in any parameter does not affect the equilibrium bias and noise in opposite directions, we say that they are *complements in the equilibrium.*

The average quality μ does not affect the bias and the noise. Since the seller sells for any realization of the signal and the buyer's valuation is linear in the quality, the seller's expected revenues are additive in μ. His marginal incentives then do not depend on μ.

When the prior is less precise, that is, σ^2 is larger, it becomes more profitable to mislead the consumer because more attention is paid to the signal. Hence, the bias

[27] A higher noise increases the revenues if $\beta < \widetilde{\beta}$ and decreases them if $\beta > \widetilde{\beta}$ since it decreases the signal weight in the buyer's posterior.

[28] The cost function is homogeneous of degree 1 in β and η; therefore, the marginal cost functions are homogeneous of degree 0 in β and η, that is, they depend only on the ratio $\frac{\beta}{\eta}$.

is used more. The costs are unaffected by σ^2 per se. However, as the bias increases, a higher noise should be used to bring the "standardized" bias $\frac{\beta\sqrt{N}}{\eta}$ down to its optimal level $\frac{1}{z_\alpha}$. Also, the noise is proportional to σ^2 since it is the ratio of the two that determines the weights in the buyer's posterior (3.1).

The variance of the prior σ^2, which is the buyers' heterogeneity, can be thought of as affected by the seller at an earlier stage through the product design. In Johnson and Myatt (2006) the relationship between the buyers' heterogeneity and the informativeness of the signal is the opposite: the benefits of giving precise information are higher if the buyers differ largely in their tastes and, therefore, more idiosyncratic products are complemented by detailed advertising and marketing activities. Therefore, if litigation is possible and the revenues are not affected by the noise, we should expect more heterogeneity in the product design to come together with larger bias and noise. Note that heterogeneity in the product design is a salient feature of the markets we described in the introduction.

A stricter (i.e., higher) standard of proof z_α means that it is more difficult to convict the seller, so the bias becomes cheaper. However, it also increases the noise that minimizes the costs for any given bias. Indeed, since there are less chances to convict the seller, the selection bias problem mentioned in Section 3.2.3 becomes less of an issue. The larger noise decreases the marginal revenues of the bias because the buyers pay less attention to the advice. Overall, we find that this decrease in the marginal revenues is always dominated by the decrease in the marginal cost, and as a result, the bias always increases following an increase in z_α.

When the authority increases the sample size N, the seller should counteract this increase in the precision of the bias estimation by increasing the noise (i.e., the cost minimizing noise shifts to the right). A larger noise decreases the marginal revenues of the bias (through a lower weight on the advice in the buyer's updating) and the bias is used less as a result. Since the seller biases less his advice, he also needs less noise to "hide" his misleading practice. It turns out that the noise decreases by enough to compensate the original increase which leads to the quite surprising result that the equilibrium noise does not depend on N.

Let us conclude this section with the following observation. In our model, setting a sufficiently high sample size N or a sufficiently low standard of proof z_α (or a sufficiently high punitive damages d that we investigate in Section 3.5.4) drives equilibrium noise and bias to zero. We would not, however, overestimate practical implications of this result for two reasons. First, this might not be the case if there is some natural and irreducible noise in communication between the buyer and the seller (see Section 3.5.6), as there will be then some lower bound on the equilibrium noise. Second, a richer model (which is beyond the scope of this paper and is an interesting topic for future work) is needed for a study of the optimal policy. Such a model will have physical costs of sending N mystery shoppers and welfare costs of punishing innocent sellers. Then, allowing for positive noise and bias in the equilibrium might be optimal.

3.4 Fixed price

In many cases the price of a product is fixed and the buyer decides whether to buy it or not at that given price. It is an interesting question (for future work) how the pricing strategy interacts with the choice of bias and noise. Assume for now that the price, p, is exogenously fixed at a pre-advice stage.[29] The buyer decides to buy if her posterior valuation of the product (3.1) is higher than p. The seller's revenues are then

$$R(\beta,\eta) = p\Pr\left\{\frac{\left(s-\widetilde{\beta}\right)\sigma^2+\mu\eta^2}{\sigma^2+\eta^2}\geq p\right\} = p\left(1-G\left(p+\widetilde{\beta}+(p-\mu)\frac{\eta^2}{\sigma^2}\right)\right)$$

and the seller's costs are unchanged.

Consider first the case where the price is exogenously fixed at the mean level μ.[30]

[29] This scenario also corresponds to the case where the price is fixed by the marketing division that operates independently from the selling division.

[30] We make this assumption for technical reasons because when $p=\mu$ the noise continues to have no effect on the revenues in equilibrium, just as in the case of price discrimination of Section 3.3. Note that the average price that the seller charges to the consumer in Section is 3.3 μ.

Proposition 3.2 *When $p = \mu > 0$, the equilibrium bias and noise are:*

$$\beta_p^* = \frac{\sqrt{\max\left\{\frac{\mu^2}{2\pi\lambda^2} - \sigma^2, 0\right\}}}{z_\alpha\sqrt{N}} \tag{3.6a}$$

$$\eta_p^* = \sqrt{\max\left\{\frac{\mu^2}{2\pi\lambda^2} - \sigma^2, 0\right\}} \tag{3.6b}$$

where $\lambda = 1 - \Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)$.[31]

Proof. See Appendix.[32] ■

The seller is trying to convince the marginal buyer. As the noise is observable, it changes the identity of the marginal buyer. In the equilibrium, however, the noise does not affect the seller's revenues since the buyer is not misled. Thus, as in the baseline model of Section 3.2, the noise can be seen as minimizing the costs for any given bias which pins down the bias-to-noise ratio, $\frac{\beta\sqrt{N}}{\eta} = \frac{1}{z_\alpha}$. The effect of the bias is to sell to the marginal buyer, i.e., to the one who receives signal $s = \bar{s}$. In the equilibrium the marginal buyer is the one who has the mean valuation μ and, therefore the marginal effect of bias on the revenues is μ times the probability of the mean valuation. Since the latter depends on the noise while the marginal cost of the bias depends only on the ratio $\frac{\beta}{\eta}$, a closed-form solution can then easily be found.

The comparative statics results are summarized in the next corollary.

Corollary 3.2 *The comparative statics of the equilibrium bias and noise (3.6a-3.6b) is the following:*

- $\frac{\partial\beta_p^*}{\partial\mu}, \frac{\partial\eta_p^*}{\partial\mu} > 0$;
- $\frac{\partial\beta_p^*}{\partial\sigma^2}, \frac{\partial\eta_p^*}{\partial\sigma^2} < 0$;
- $\frac{\partial\beta_p^*}{\partial z_\alpha}, \frac{\partial\eta_p^*}{\partial z_\alpha} > 0$;[33]

[31] Then κ of Proposition 3.1 can be written as $\kappa = \frac{1-\lambda}{\lambda}$.

[32] A sufficient (but not necessary) condition for the second order conditions to hold is $z_\alpha \geq 1.09$ (i.e., $\alpha \leq 13.79\%$).

[33] A sufficient (but not necessary) condition for $\frac{\partial\beta_p^*}{\partial z_\alpha} > 0$ is $z_\alpha \geq 1.29$ (i.e., $\alpha \leq 9.85\%$). See the proof for the exact condition.

- $\frac{\partial \beta_p^*}{\partial N} < 0$ *and* $\frac{\partial \eta_p^*}{\partial N} = 0$.

Proof. See Appendix. ■

Given that a change in any parameter does not affect the equilibrium bias and noise in opposite directions, we say that they are still *complements in the equilibrium*.

The average quality now affects the bias and the noise unlike the baseline model of Section 3.2. A higher μ increases both the posterior expected quality and the price, so the probability of selling remains unchanged. However, an increase in μ (through the increase in p) also increases the revenue per unit sold. The increase in the profitability of the sale induces the seller to increase the bias, as this increases the probability of selling. Following the increase in the bias, the noise increases too to minimize the costs for this higher bias.

The comparative statics with respect to the variance of the prior, σ^2, have the opposite sign than in Section 3.3 (see Corollary 3.1). The intuition is the following. While σ^2 still enters the buyer's posterior, it does not affect her decision to buy. Indeed, since the price equals to the mean valuation, the buyer buys if the signal exceeds the mean (after debiasing) and does not buy otherwise. Thus, the only effect of a higher σ^2 is to decrease the probability of the marginal buyer, that is, the buyer with the average valuation μ. Then, the marginal revenues of the bias are smaller and the seller uses it less (and correspondingly he introduces a lower noise).

When the standard of proof, z_α, increases, the seller can afford to use a larger noise as the probability of ending up convicted and paying a large fine decreases. A stricter standard of proof, also decreases the marginal cost of biasing the advice. However, as in the comparative statics of Section 3.3, the larger noise also decreases the marginal revenue of the bias although through a different mechanism. As with σ^2, a larger noise decreases the probability of the marginal consumer and hence makes biasing the advice less profitable. However, if z_α is large enough, the decrease of the marginal revenues is smaller than the one of marginal costs and the bias increases.

The intuition for the comparative statics of the sample size N is the same as in Section 3.3, with the difference that the decrease in the marginal revenues of bias following an increase in the noise is due to a decrease in the probability of the marginal consumer (and not to a decrease in the attention paid to the advice).

Finally, it is easy to check that if the ratio $\frac{\mu^2}{2\pi\sigma^2}$ is large enough (i.e., $\frac{\mu^2}{2\pi\sigma^2} > \lambda$), then $\beta_p^* > \beta^*$ and $\eta_p^* > \eta^*$. This is because, as shown in Corollary 3.2, a larger μ and/or a smaller σ^2 increases the marginal benefit of biasing the advice.

Consider now the case $p \neq \mu$. While we cannot obtain closed-form solutions we still can shed light on the bias-to-noise ratio in the equilibrium.

Proposition 3.3 *When $p > (<)\mu$, then $\frac{\beta_p^*}{\eta_p^*} > (<)\frac{\beta^*}{\eta^*}$, where β^* and η^* are defined in (3.6a) and (3.6b).*

Proof. See Appendix. ■

The only difference with the previous setup where $p = \mu$ is that now the noise has two effects on the marginal revenues: like the bias, it affects the decision of the marginal consumer (in the direction that depends on how the marginal consumer is related to the mean of the signal) but, since the noise is observable, it also changes the identity of the marginal consumer. In the equilibrium, a higher noise brings the buyer's posterior closer to the mean which decreases the seller's revenues when $p > \mu$ since the marginal consumer has the valuation above the mean and increases them when $p < \mu$ since the marginal consumer has the valuation below the mean. As a result, when $p > \mu$ the bias-to-noise ratio will be larger than the one in Section 3.3 and smaller when $p < \mu$. Actually, the analysis of the effects of the noise on the seller's revenues becomes essentially the one of Johnson and Myatt (2006). The case $p > \mu$ corresponds to their "niche market", where the seller only serves the high valuation consumers and provides them with a lot of information, and the case $p < \mu$ corresponds to their "mass market" where the seller serves most consumers and gives them few information. However, contrary to their setup where noise is costless, in our model the noise also affects the seller's costs and, therefore, an extreme noise is not optimal.

3.5 Extensions

3.5.1 Unobservable noise

In some cases it is more reasonable to assume that the noise chosen by the seller is not observed by the buyer. For instance, the seller can increase the noise by failing to explain hidden costs that the buyer may incur just once he uses the service or some features that the buyer did not think about or expected to use significantly (such as a bank overdraft). The unobserved noise can also come from the small script that is typically attached to this type of contracts. These clauses are usually not read in the moment of the purchase and they can be more or less precise. In all these examples, the buyer cannot assess at the moment of receiving the signal how precise it is.

The buyer updates her beliefs in a way similar to (3.1):

$$E\left[\theta \mid s, \widetilde{\eta}, \widetilde{\beta}\right] = \mu + \left(s - \mu - \widetilde{\beta}\right) \frac{\sigma^2}{\sigma^2 + \widetilde{\eta}^2},$$

where $\widetilde{\eta}$ is the buyer's conjecture about the noise used by the seller (that has to be correct in the equilibrium). The seller's revenues are equal to the expected valuation that the seller extract from the buyers:

$$R(\beta, \eta) = \int_{-\infty}^{+\infty} E\left[\theta \mid s, \widetilde{\eta}, \widetilde{\beta}\right] g(s) ds = \mu + \left(\beta - \widetilde{\beta}\right) \frac{\sigma^2}{\sigma^2 + \widetilde{\eta}^2}.$$

Note that the seller extracts from the buyer the expected quality of match plus how much the buyer is misled into thinking that the product is good. The only difference with (3.2) is that the buyer uses her conjecture about the noise $\widetilde{\eta}$ since the noise is not observable. The noise η does not affect the revenues then. In equilibrium, when the buyer's conjecture about the bias used is correct, the seller extracts the expected match quality μ as before.

If buyers do not observe the noise, the authority does not observe it either and has to estimate it as

$$\widehat{\eta^2} = \frac{1}{N-1}\sum_{i=1}^{N}\left(\varepsilon_i - \widehat{\beta}\right)^2,$$

where $\widehat{\beta}$ is the estimated bias (3.3). This estimator is distributed as $\widehat{\eta^2} \rightsquigarrow \chi^2(N-1)$. The authority constructs statistics $t = \frac{\widehat{\beta}}{\widehat{\eta}/\sqrt{N}}$ which, under the null hypothesis of seller's innocence $\beta = 0$ has $t-$distribution with $N-1$ degrees of freedom. The authority rejects H_0 if $\frac{\widehat{\beta}}{\widehat{\eta}/\sqrt{N}} \geq t_\alpha$, where α is the significance level of the test.[34] However, for a given bias β, the true distribution of $\frac{\widehat{\beta}}{\widehat{\eta}/\sqrt{N}}$ is non-central $t-$distribution with $N-1$ degrees of freedom and non-centrality parameter $ncp = \beta\frac{\sqrt{N}}{\eta}$ (denote its cdf as T_{N-1}^{ncp}).[35] The seller's costs are

$$C(\beta,\eta) = \frac{\eta}{\sqrt{N}}\int_{t_\alpha}^{+\infty} t dT_{N-1}^{ncp}(t). \tag{3.7}$$

The only difference with costs (3.4a) is that the normal distribution there is now replaced with non-central $t-$distribution. A higher β increases ncp and, therefore, the distribution T_{N-1}^{ncp} moves to the right and the integral in (3.7) goes up. A higher η has two effects, as before: it makes it less likely for the seller to be found guilty (the integral in (3.7) decreases since ncp is lower), but increases the chances of a really large fine once he is found guilty (the integral in (3.7) is multiplied by a higher number). It is easily shown that costs (3.7) are decreasing at $\eta = 0$ and increasing at $\eta \to +\infty$ and, hence, minimized for some interior η.

The equilibrium then has a similar structure to the one in Section 3.3 (in particular, equilibrium bias and noise are *complements*) and is found in a similar way. Now, in the equilibrium $\beta = \widetilde{\beta}$ and $\eta = \widetilde{\eta}$. The first-order conditions of the seller's

[34] By definition, the probability that a random variable following $t-$distribution with $N-1$ degrees of freedom takes a value above t_α is equal to α.

[35] Indeed,

$$\Pr\left\{\frac{\widehat{\beta}}{\widehat{\eta}/\sqrt{N}} \geq t_\alpha\right\} = \Pr\left\{\frac{\frac{\widehat{\beta}-\beta}{\eta/\sqrt{N}} + \frac{\beta}{\eta/\sqrt{N}}}{\sqrt{\frac{N-1}{N-1}\frac{\widehat{\eta}}{\eta}}} \geq t_\alpha\right\} = \Pr\left\{\frac{\mathcal{N}(0,1) + \frac{\beta}{\eta/\sqrt{N}}}{\sqrt{\frac{\chi^2_{N-1}}{N-1}}} \geq t_\alpha\right\} = 1 - T_{N-1}^{ncp}(t_\alpha),$$

where the last equality is the definition of the non-central $t-$distribution.

profit maximization in the equilibrium are

$$\frac{\partial \Pi\left(\beta,\eta\right)}{\partial \beta} = \frac{\sigma^2}{\sigma^2+\eta^2} - \frac{\partial}{\partial ncp}\int_{t_\alpha}^{+\infty} tdT_{N-1}^{ncp}\left(t\right) = 0, \tag{3.8a}$$

$$\frac{\partial \Pi\left(\beta,\eta\right)}{\partial \eta} = -\frac{1}{\sqrt{N}}\int_{t_\alpha}^{+\infty} tdT_{N-1}^{ncp}\left(t\right) + \frac{\beta}{\eta}\frac{\partial}{\partial ncp}\int_{t_\alpha}^{+\infty} tdT_{N-1}^{ncp}\left(t\right) = 0. \tag{3.8b}$$

Condition (3.8b) depends only on $ncp = \beta\frac{\sqrt{N}}{\eta}$ (after multiplying it by $\sqrt{N}$) and, therefore, in the equilibrium the ratio $\frac{\beta}{\eta}$ is chosen to minimize the costs with respect to the noise. Then, the noise is found from (3.8a), in which the second term, the equilibrium marginal costs of bias $\frac{\partial}{\partial ncp}\int_{t_\alpha}^{+\infty} tdT_{N-1}^{ncp}\left(t\right)$, does not depend on bias and noise.

With observable noise, $\beta\frac{\sqrt{N}}{\eta}$ is also chosen to minimize the costs. Then, η is found from $\frac{\partial\Pi(\beta,\eta)}{\partial\beta} = 0$. With unobservable noise, it is the buyer's conjecture $\widetilde{\eta}$ that makes the condition $\frac{\partial\Pi(\beta,\eta)}{\partial\beta} = 0$ satisfied. Since the equilibrium $\widetilde{\eta}$ is equal to η, all the difference comes from the fact that the authority statistics has normal distribution in the observable case and non-central $t-$distribution here. It also means that, when N increases and costs (3.7) converge to the costs in the baseline case (3.4a) since the non-observability of the noise matters less and less for the authority's estimation, the equilibrium then also converges to the equilibrium in the baseline case (3.5a-3.5b).

Finally, when the noise is unobserved, there is no off-equilibrium events. The buyer observes only a signal realization and any realization is compatible with any choice of bias and noise by the seller. The fact that the equilibrium when the noise is unobservable looks similar to the case when the noise is observable is one of the main reasons why we focused on the constant conjecture $\widetilde{\beta}\left(\eta\right)$ in the baseline model (see fn. 18).

3.5.2 Informational externalities between sellers

Suppose that, while still buying from a given seller, the buyer also observes a signal from another seller. For example, the buyer may want to buy only a certain brand of a mobile phone, say, for aesthetic or compatibility reasons, but she can learn about some new technical features from the mobile phones of other brands as well.[36] Sellers cannot distinguish between buyers who buy from them and those who only get informed.

There are two sellers, seller 1 and seller 2, each sending a signal with their respective bias and noise. The buyer has the match quality $\theta_1 \rightsquigarrow \mathcal{N}(\mu_1, \sigma_1^2)$ with seller 1 and $\theta_2 \rightsquigarrow \mathcal{N}(\mu_2, \sigma_2^2)$ with seller 2. The correlation between θ_1 and θ_2 is ρ_θ and the correlation between distortions ε_1 and ε_2 is ρ_ε. The former describes how similar the two products are. Somebody may be determined to buy an iPhone but knows that many functional features are similar in the phones of other brands. The latter reflects some random changes in the buyer's mood or attitude towards. A buyer may have a bad moment and she would then be critical to any information she receives from the sellers.

Seller i's signal s_i is distributed as $\mathcal{N}(\mu_i + \beta_i, \sigma_i^2 + \eta_i^2)$. Having observed signals from both sellers, s_1 and s_2, the buyer who buys from seller 1 computes her expected match quality in the following way (see Appendix for the derivation):

$$E\left[\theta_1 \mid s_1, s_2, \eta_1, \eta_2, \widetilde{\beta}_1, \widetilde{\beta}_2\right] = \mu_1 + \frac{1+\frac{\eta_2^2}{\sigma_2^2} - \rho_\theta\left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)}{\left(1+\frac{\eta_2^2}{\sigma_2^2}\right)\left(1+\frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2}\left(s_1 - \mu_1 - \widetilde{\beta}_1\right)$$
$$+ \frac{\rho_\theta \frac{\eta_1^2}{\sigma_1\sigma_2} - \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_2^2}}{\left(1+\frac{\eta_2^2}{\sigma_2^2}\right)\left(1+\frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2}\left(s_2 - \mu_2 - \widetilde{\beta}_2\right). \tag{3.9}$$

The two sellers move simultaneously. In the equilibrium, when deciding on the bias and noise, each of them knows the bias and noise of the other but not the realization of the other seller's signal. Also, the buyer has correct conjectures about the biases of both sellers. If seller 1 decides to deviate from the equilibrium to some bias β_1 (when seller 2 does not) his expected revenues are

[36] See Ivanov (2011) for a persuasion model (in which the bias cannot be used) where the buyer can choose from which seller to buy.

$$
\begin{aligned}
R_1(\beta_1, \eta_1) &= \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} E\left[\theta_1 \mid s_1, s_2, \eta_1, \eta_2, \widetilde{\beta}_1, \widetilde{\beta}_2\right] dG(s_1)\, dG(s_2) \\
&= \mu_1 + \frac{1 + \frac{\eta_2^2}{\sigma_2^2} - \rho_\theta\left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)}{\left(1+\frac{\eta_2^2}{\sigma_2^2}\right)\left(1+\frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2}\left(\beta_1 - \widetilde{\beta}_1\right). \quad (3.10)
\end{aligned}
$$

The costs are still given by (3.4b).[37]

Proposition 3.4 *With informational externalities between the two sellers, when correlation is imperfect, $|\rho_\theta| < 1$ and $|\rho_\varepsilon| < 1$, there is a unique symmetric (pure-strategy) equilibrium*

$$
\begin{aligned}
\beta_i^e &= \tfrac{1}{z_\alpha}\tfrac{\eta_i^e}{\sqrt{N}}, \\
\eta_i^e &= \tfrac{\sigma_i}{\sqrt{2}}\sqrt{(\kappa-1)\tfrac{1-\rho_\theta\rho_\varepsilon}{1-\rho_\varepsilon^2} + \sqrt{(\kappa-1)^2\left(\tfrac{1-\rho_\theta\rho_\varepsilon}{1-\rho_\varepsilon^2}\right)^2 + 4\kappa\tfrac{1-\rho_\theta^2}{1-\rho_\varepsilon^2}}}
\end{aligned} \quad (3.11)
$$

where κ is defined in Proposition 3.1.

If the match qualities are perfectly correlated, $|\rho_\theta| = \pm 1$, there is always equilibrium $\beta_1^e = \beta_2^e = \eta_1^e = \eta_2^e = 0$. If and only if $\kappa > 1$, there is another equilibrium[38]

$$
\begin{aligned}
\beta_i^e &= \frac{1}{z_\alpha}\frac{\eta_i^e}{\sqrt{N}}, \\
\eta_i^e &= \sigma_i\sqrt{\frac{\kappa-1}{1\pm\rho_\varepsilon}}.
\end{aligned}
$$

If the shocks are perfectly correlated, $|\rho_\varepsilon| = \pm 1$, if and only if $\kappa < 1$, there is a unique symmetric (pure-strategy) equilibrium

$$
\begin{aligned}
\beta_i^e &= \frac{1}{z_\alpha}\frac{\eta_i^e}{\sqrt{N}}, \\
\eta_i^e &= \sigma_i\sqrt{\frac{\kappa}{1-\kappa}(1\pm\rho_\theta)}.
\end{aligned}
$$

[37] Despite the correlation between ε_1 and ε_2, the estimator of the bias is still the sample average. See Casella and Berger (2002), p. 358, ex. 7.18.

[38] A higher κ means a higher standard of proof z_α (and a lower α); $\kappa = 1$ corresponds to $z_\alpha \approx 1.70$ (and $\alpha \approx 4.5\%$).

Proof. See Appendix. ■

The effects of the correlations ρ_θ and ρ_ε are sometimes ambiguous when they are both quite high. If ρ_ε is close to zero, then the effect of ρ_θ is to decrease the equilibrium value of bias and noise. This comes from the fact that when ρ_θ increases, the buyer puts a lower weight on the signal of her seller and a higher one on the signal of the other seller, as can be seen from (3.9). Then, each seller has less incentives to bias his signal and, correspondingly, he introduces less noise.

When ρ_θ is close to zero, an increase in ρ_ε has an opposite effect on the buyer's posterior. Indeed, the buyer trusts the signal from her seller more since she knows something about the realization of the error term through the correlation between ε_1 and ε_2. Then, the seller has more incentives to bias his signal and add noise. It can be easily checked that β_i^e and η_i^e in (3.11) are higher than β^* and η^* in (3.5a-3.5b) for $\rho_\theta = 0$ and $\rho_\varepsilon \neq 0$. Thus, the presence of another seller who sells an unrelated good but with correlated shocks in the communication process decreases welfare. The seller cannot commit not to bias his signal more and, while the consumers are not misled in the equilibrium, he ends up paying a higher fine.[39]

When the match quality of one seller is perfectly correlated with the one of the other, $|\rho_\theta| = \pm 1$, the situation is slightly different. If one of the sellers does not use any noise, he perfectly reveals his match quality and, therefore, the match quality of the other seller. The latter then cannot mislead the buyer and chooses zero bias and noise. There might be another equilibrium where both sellers do use some bias and noise and which is the limit of the equilibrium with imperfect correlation (3.11).

3.5.3 Credulous buyers

In the baseline model, the buyers are rational and anticipate correctly the bias introduced by the seller. Then, the bias (and the noise) do not harm the buyers in the equilibrium: each of them pays on average μ which is the average match quality.

However, there is extensive empirical evidence about the existence of credulous

[39] See Meyer and Vickers (1997) for a dynamic model of implicit incentives where informational externalities between two agents also may reduce welfare.

buyers. For instance, brokers are usually paid by the consumers but also receive compensation from the lenders - which influence the broker to offer the consumer loan terms or products that are not in the consumer's interest. However, according to the FTC (2008), "many consumers purportedly view mortgage brokers as trusted advisors who shop for the best loan for the consumer" (p. 16). See other examples in Gabaix and Laibson (2006) and Inderst and Ottaviani (2009a) among others.

In this Section, we let some buyers to be credulous and blindly believe the seller's signal. More precisely, a share of credulous buyers c does not understand that the seller might send a biased and imprecise signal. Thus, these consumers blindly follow the signal.

The seller's expected revenues are then

$$R(\beta,\eta) = \mu + (1-c)\left(\beta - \widetilde{\beta}\right)\frac{\sigma^2}{\sigma^2+\eta^2} + c\beta.$$

Proposition 3.5 *When there is share c of credulous consumers, the equilibrium bias and noise are*

$$\begin{aligned} \beta_c^* &= \frac{\sigma}{\sqrt{N}}\frac{1}{z_\alpha}\sqrt{\kappa_c}, \\ \eta_c^* &= \sigma\sqrt{\kappa_c}, \end{aligned}$$

where $\kappa_c = \frac{\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) - z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)}{1-\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right) - c}$.[40]

They both increase with c.

Proof. Similar to the proof of Proposition 3.1. ■

As expected, equilibrium bias increases with the share of credulous consumers since marginal returns to bias are higher. The equilibrium noise also increases since it is complementary to the bias.

[40] If $1-\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right) - c \leq 0$, the seller's marginal costs in the equilibrium are smaller then his marginal revenues and, thus, he introduces an infinite bias.

3.5.4 Punitive damages

Punitive (or exemplary) damages are used often in common law countries. In the US, both the frequency and the magnitude of punitive damages verdicts has increased dramatically in recent years (Sunstein, Hastie, Payne, Schkade, and Viscusi (2002)). We can model them by assuming that the seller is fined by the amount $d\widehat{\beta}$ when he is found guilty, $d \geq 1$.[41] The cost function (3.4b) is then multiplied by d. Following the same steps as in the baseline model, we obtain the equilibrium bias and noise.

Proposition 3.6 *When there are punitive damages d, the equilibrium bias and noise are*

$$\begin{aligned}\beta_d^* &= \frac{\sigma}{\sqrt{N}}\frac{1}{z_\alpha}\sqrt{\max\{\kappa_d, 0\}},\\ \eta_d^* &= \sigma\sqrt{\max\{\kappa_d, 0\}},\end{aligned}$$

where $\kappa_d = \frac{1-d\left[1-\Phi\left(z_\alpha-\frac{1}{z_\alpha}\right)+z_\alpha\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)\right]}{d\left[1-\Phi\left(z_\alpha-\frac{1}{z_\alpha}\right)+z_\alpha\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)\right]}$.

They both decrease with d.

Proof. Similar to the proof of Proposition 3.1. ■

When $d = 1$, we obtain our baseline case results (3.5a-3.5b). When the punishment becomes more severe, that is, d increases, there is less bias and less noise. If d is high enough, the seller is completely deterred from misleading the buyers.

3.5.5 Only signals are observed[42]

In the baseline model above we assumed that the authority knows the quality θ_i obtained by each mystery shopper (see Section 3.2.3). However, the authority may not have such precise information though it still may know more about them than

[41] We abstract from the fact that in practice it is not very clear how the amount of punitive damages awarded depends on the harm made. Sunstein, Hastie, Payne, Schkade, and Viscusi (2002) say that there "...is inability to explain ... various punitive damages verdicts on a rational basis" (p. 2).

[42] The third possibility, namely, that only quality of the match is observed (or can be used in court) does not help here: since all consumers buy for all signal realizations, the ex post distribution of quality is the same as the prior one. This would be different if consumers buy only above a certain threshold as in Section 3.4.

about the general population of buyers. For example, mystery shoppers are likely to be more homogenous, say, young and educated, than the general population of buyers. To see the effect of information about mystery shoppers on the equilibrium, suppose that their match qualities are less heterogenous, that is, they are distributed with variance $\sigma_m^2 \leq \sigma^2$. The authority observes only signals s_i obtained by mystery shoppers which are distributed according to $\mathcal{N}(\mu + \beta, \sigma_m^2 + \eta^2)$. If $\sigma_m^2 = 0$, we are back to the baseline model. If $\sigma_m^2 = \sigma^2$, the authority does not know anything about mystery shoppers.

The estimated bias is

$$\widehat{\beta} = \frac{1}{N}\sum_{i=1}^{N} s_i - \mu.$$

The authority tests $H_0 : \beta = 0$ against $H_1 : \beta > 0$ and punishes the seller only when H_0 is rejected. The estimator $\widehat{\beta}$ is distributed as $\widehat{\beta} \rightsquigarrow \mathcal{N}\left(\beta, \frac{\sigma_m^2+\eta^2}{N}\right)$. In a way similar to (3.4b), the seller's costs can be written as

$$C(\beta, \eta) = \beta\left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2 + \eta^2}}\right)\right) + \frac{\sqrt{\sigma_m^2 + \eta^2}}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2 + \eta^2}}\right). \tag{3.12}$$

The seller's problem, to maximize revenues (3.2) minus costs (3.12), is similar to the one in Section 3.3. However, the condition guaranteeing that it is locally concave has to be more restrictive. Intuitively, a higher σ_m^2 makes the costs less responsive to the bias (i.e. it is harder to fine the seller). It is possible then that the seller prefers to introduce an infinite bias coupled with infinite noise. Introducing punitive damages d as in Section 3.5.4 restores the concavity.

Proposition 3.7 *If only signals are observed (and there are punitive damages d), the equilibrium bias and noise are*

- *If* $\kappa_d \geq 0$,

$$\begin{aligned}
\beta_s^* &= \frac{\sqrt{\sigma_m^2 + \sigma^2\kappa_d}}{\sqrt{N}}\frac{1}{z_\alpha}, \\
\eta_s^* &= \sigma\sqrt{\kappa_d}.
\end{aligned}$$

The equilibrium bias is higher, $\beta_s^* > \beta_d^*$, *and the noise is the same,* $\eta_s^* = \eta_d^*$, *as in the case when both signals and match qualities are observed.*

The second-order conditions are satisfied if $z_\alpha < 2.436$ *and* $\left(1 - \frac{\sigma_m^2}{\sigma^2}\right) d = 1$.[43]

- *If* $\kappa_d < 0$, $\eta_s^* = 0$ *and* β_s^* *is found from*

$$1 - d\left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sigma_m}\right) + z_\alpha \phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sigma_m}\right)\right) = 0 \tag{3.13}$$

Proof. See Appendix. ■

A higher heterogeneity of mystery shoppers, σ_m^2, makes it more difficult to convict the seller and the equilibrium bias increases.[44] The equilibrium noise does not change. This perhaps counterintuitive result has the same explanation as the one of the equilibrium noise not depending on the sample size N in Corollary 3.1. The minimum costs with respect to the noise are obtained when $\frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}} = \frac{1}{z_\alpha}$. The noise is then found from $\frac{\sigma^2}{\sigma^2+\eta^2} = d\left(1 - \Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)\right)$, where the right hand side is the equilibrium marginal costs of bias that do not depend on σ_m^2.

3.5.6 Natural noise in the communication

Suppose there is a natural noise in the communication v^2. For example, the communication is always imperfect unless the seller makes it better, possibly, at some effort costs. Then, the error term ε is distributed as $\varepsilon \rightsquigarrow \mathcal{N}(\beta, v^2 + \eta^2)$. Let us consider two polar cases. First, suppose that the effort costs are zero, that is, the seller can set η^2 equal to $-v^2$ making communication perfect at no cost. The seller's problem with respect to noise can be thought of as choosing the total variance $v^2 + \eta^2$ and its equilibrium level will be the same as in baseline model (3.5b), that is, $\sqrt{v^2 + \eta^{*2}} = \sigma\sqrt{\kappa}$. The equilibrium bias will also be the same as in the baseline model (3.5a).

---

[43] We use $z_\alpha < 2.436$ because we have this condition in Proposition 3.1. Requiring a lower treshold for z_α allows to decrease the necessary d.

[44] In case of $\kappa_d > 0$ this is clear from the expression of β_s^*. In case $\kappa_d < 0$ it holds since the derivative of (3.13) with respect to σ_m^2 is positive.

The opposite case is when the effort costs are so high that the seller never chooses a negative η^2. The seller's problem with respect to noise is choosing the total variance $v^2 + \eta^2$ with the constraint $\eta^2 \geq 0$. If the (unconstrained) equilibrium noise η^* (3.5b) is higher than v^2, then this constraint is not binding and the equilibrium is not affected. Otherwise, the equilibrium becomes $\eta^* = 0$ (that is, the variance of the error term is at its lowest level v^2) and the equilibrium bias is found from the first-order condition (3.14a) when $\eta = v$.

3.5.7 One-stage procedure used by the authority

The procedure used by the authority in the baseline model of Section 3.2 is probably best interpreted as a two-stage procedure. At the first stage, the authority determines if the seller is guilty, and if it is the case, the authority fixes the punishment at the second stage. For this reason, the fine depends only on the estimated bias but not on the noise.

If the authority were to determine both if the seller is guilty and the size of the fine at the same stage, it then might account for the noise when fixing the fine, that is, to fine the seller by $\frac{\widehat{\beta}}{\eta/\sqrt{N}}$. The costs are then

$$C(\beta, \eta) = \int_{z_\alpha \frac{\eta}{\sqrt{N}}}^{+\infty} \widehat{\beta} \frac{\sqrt{N}}{\eta} f\left(\widehat{\beta}\right) d\widehat{\beta} = \beta \frac{\sqrt{N}}{\eta} \left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right) + \phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right).$$

The costs now depend only on the ratio $\frac{\beta}{\eta}$. A higher noise decreases both the probability of conviction and the fine whenever it is imposed. Therefore, the seller will always introduce maximum possible noise. This observation leads to the next proposition.

Proposition 3.8 *If the authority adjusts the fine for the noise, there is no information transmission in the equilibrium.*

Proof. The first derivative of the profit function with respect to the noise is

$$\frac{\partial \Pi}{\partial \eta} = -\left(\beta - \widetilde{\beta}\right) \frac{2\eta\sigma^2}{\left(\sigma^2+\eta^2\right)^2} + \frac{\beta\sqrt{N}}{\eta^2}\left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right) + z_\alpha \phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right)$$

In the equilibrium the buyers are not misled, $\beta = \widetilde{\beta}$, and this derivative is positive. The seller introduces then infinite noise, $\eta = +\infty$, and the buyers ignore the seller's signal. ■

3.6 Conclusion

In this paper we investigated seller's incentives to provide (un)biased and (un)informative advice and the resulting equilibrium communication. We found that the biasing the advice and making it more noisy are complements: the seller employs says either an exact truth or a vague lie. For example, a higher buyers' heterogeneity, a higher standard of proof employed by the authority and a higher share of credulous consumers make the signal sent by the seller to the buyers more biased and less precise.

An interesting direction for future work is to characterize the optimal policy of the authority. In the model, a policy is described by three parameters: the number of consumers sampled N, the standard of proof used when establishing if the seller is guilty z_α and the size of the fine in relation to the estimated bias d. We have taken these parameters as given and found that the comparative statics are intuitive: a higher N, a lower z_α and a higher d reduce the equilibrium bias. This poses the question of why the authority does not set them up at the level that would completely deter the seller from biasing the signal. The reason is that there are of course costs of a strict policy. Surveying a lot of consumers, that is, a high N, is costly in monetary terms; a lower standard of proof z_α makes conviction of innocent sellers more likely and higher punitive damages d make innocent sellers pay a higher fine. Extending the model to incorporate these costs may produce relevant and realistic recommendations on the optimal policy.

3.7 Appendix

Proof of Proposition 3.1. The first-order conditions of the seller's problem with respect to β and η, respectively, are

$$\begin{aligned}
\frac{\sigma^2}{\sigma^2+\eta^2} - \left(1-\Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right) + z_\alpha\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right) &= 0 \quad (3.14a)\\
-\left(\beta-\widetilde{\beta}\right)\frac{2\eta\sigma^2}{\left(\sigma^2+\eta^2\right)^2} - \frac{1}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left[1-\frac{\beta\sqrt{N}}{\eta}z_\alpha\right] &= 0 \quad (3.14b)
\end{aligned}$$

In the equilibrium $\beta = \widetilde{\beta}$ and from (3.14b) $\frac{\beta\sqrt{N}}{\eta} = \frac{1}{z_\alpha}$. Plug this into (3.14a) to obtain β^* and then η^*.

To check the second-order conditions, differentiate (3.14a) with respect to β to obtain $\frac{\partial^2\Pi}{\partial\beta^2}$ and (3.14b) with respect to η and β to obtain $\frac{\partial^2\Pi}{\partial\eta^2}$ and $\frac{\partial^2\Pi}{\partial\eta\partial\beta}$, respectively:

$$\begin{aligned}
\frac{\partial^2\Pi}{\partial\beta^2} &= -\frac{\sqrt{N}}{\eta}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left(1-\frac{\beta\sqrt{N}}{\eta}z_\alpha + z_\alpha^2\right)\\
\frac{\partial^2\Pi}{\partial\eta^2} &= -\left(\beta-\widetilde{\beta}\right)\frac{\sigma^2-3\eta^2}{\left(\sigma^2+\eta^2\right)^3} - \frac{\beta^2\sqrt{N}}{\eta^3}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left(1-\frac{\beta\sqrt{N}}{\eta}z_\alpha + z_\alpha^2\right)\\
\frac{\partial^2\Pi}{\partial\eta\partial\beta} &= -\frac{2\eta\sigma^2}{\left(\sigma^2+\eta^2\right)^2} + \frac{\beta\sqrt{N}}{\eta^2}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left(1-\frac{\beta\sqrt{N}}{\eta}z_\alpha + z_\alpha^2\right)
\end{aligned}$$

In the equilibrium $\beta = \beta^* = \widetilde{\beta}$ and $\frac{\beta^*\sqrt{N}}{\eta^*} = \frac{1}{z_\alpha}$ so these derivatives become

$$\begin{aligned}
\frac{\partial^2\Pi}{\partial\beta^2} &= -\frac{\sqrt{N}}{\eta^*}\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)z_\alpha^2 < 0\\
\frac{\partial^2\Pi}{\partial\eta^2} &= -\frac{\beta^{*2}\sqrt{N}}{\eta^{*3}}\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)z_\alpha^2 < 0\\
\frac{\partial^2\Pi}{\partial\eta\partial\beta} &= -\frac{2\eta^*\sigma^2}{\left(\sigma^2+\eta^{*2}\right)^2} + \frac{\beta^*\sqrt{N}}{\eta^{*2}}\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)z_\alpha^2
\end{aligned}$$

Check that the determinant of the Hessian is positive:

$$
\begin{aligned}
\frac{\partial^2\Pi}{\partial\beta^2}\frac{\partial^2\Pi}{\partial\eta^2}-\left(\frac{\partial^2\Pi}{\partial\eta\partial\beta}\right)^2 &= \frac{\beta^{*2}N}{\eta^{*4}}\phi^2\left(z_\alpha-\frac{1}{z_\alpha}\right)z_\alpha^4-\left(\frac{2\eta^*\sigma^2}{(\sigma^2+\eta^{*2})^2}\right)^2 \\
&\quad+\frac{4\eta^*\sigma^2}{(\sigma^2+\eta^{*2})^2}\frac{\beta^*\sqrt{N}}{\eta^2}\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)z_\alpha^2-\frac{\beta^{*2}N}{\eta^{*4}}\phi^2\left(z_\alpha-\frac{1}{z_\alpha}\right)z_\alpha^4 \\
&= \frac{4\sigma^2}{(\sigma^2+\eta^{*2})^2}\left(\frac{\beta^*\sqrt{N}}{\eta^{*2}}\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)z_\alpha^2-\frac{\eta^*\sigma^2}{(\sigma^2+\eta^{*2})^2}\right) \\
&= \frac{4\sigma^2}{\eta^*(\sigma^2+\eta^{*2})^2}\left(\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)z_\alpha-\frac{\kappa}{(1+\kappa)^2}\right)
\end{aligned}
$$

It is positive for $z_\alpha < 2.436$, i.e., $\alpha > 0.75\%$.

These are only local second-order conditions. In our candidate equilibrium, i.e., if the seller's bias and buyer's conjecture are (3.5a) and the noise is (3.5b), the seller earns μ minus the costs. Then, he might prefer to deviate in the following way: stop the communication by introducing infinite noise, in which case the buyer disregards the signal and the revenues are still μ, and introduce an infinite negative bias in order to drive the costs to zero.[45] This deviation is not feasible because the seller chooses both bias and noise from a bounded interval. Usually, indeed, the seller is obliged to provide some minimal amount of information about the product; in other words, there is an upper bound on η. While a small negative bias can be interpreted as "modesty", an infinite one is clearly unrealistic. ■

Proof of Proposition 3.2. The first-order conditions of the seller's problem with respect to β and η, respectively, are

$$
\mu g\left(\mu+\widetilde{\beta}\right)-\left(1-\Phi\left(z_\alpha-\frac{\beta\sqrt{N}}{\eta}\right)+z_\alpha\phi\left(z_\alpha-\frac{\beta\sqrt{N}}{\eta}\right)\right) = 0 \quad (3.15a)
$$

$$
-\mu\eta g(\mu+\widetilde{\beta})\frac{\beta-\widetilde{\beta}}{\sigma^2+\eta^2}-\frac{1}{\sqrt{N}}\phi\left(z_\alpha-\frac{\beta\sqrt{N}}{\eta}\right)\left[1-\frac{\beta\sqrt{N}}{\eta}z_\alpha\right] = 0 \quad (3.15b)
$$

In the equilibrium, $\beta=\widetilde{\beta}$. From (3.15b) $\frac{\beta\sqrt{N}}{\eta}=\frac{1}{z_\alpha}$ as in Section 3.3. Plug this into (3.15a) to obtain β_p^* and then η_p^*.

[45] For a negative or zero bias the costs are strictly increasing in the noise, see the proof of Lemma 3.1.

To check the second-order conditions, we differentiate (3.15a) with respect to β and η to obtain $\frac{\partial^2\Pi}{\partial\beta^2}$ and $\frac{\partial^2\Pi}{\partial\eta\partial\beta}$, respectively and (3.15b) with respect to η to obtain $\frac{\partial^2\Pi}{\partial\eta^2}$:

$$\begin{aligned}
\frac{\partial^2\Pi}{\partial\beta^2} &= \mu g(\mu+\widetilde{\beta})\frac{\widetilde{\beta}-\beta}{\sigma^2+\eta^2} - \frac{\sqrt{N}}{\eta}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left(1+z_\alpha\left(z_\alpha-\frac{\beta\sqrt{N}}{\eta}\right)\right)\\
\frac{\partial^2\Pi}{\partial\eta^2} &= -\mu g(\mu+\widetilde{\beta})\frac{\beta-\widetilde{\beta}}{\sigma^2+\eta^2}\left[1+\eta^2\left(\frac{\left(\widetilde{\beta}-\beta\right)^2-3\left(\sigma^2+\eta^2\right)}{\left(\sigma^2+\eta^2\right)^2}\right)\right]\\
&\quad -\frac{\beta^2}{\eta^2}\frac{\sqrt{N}}{\eta}\phi\left(z_\alpha-\frac{\beta\sqrt{N}}{\eta}\right)\left[1+z_\alpha\left(z_\alpha-\frac{\beta\sqrt{N}}{\eta}\right)\right]\\
\frac{\partial^2\Pi}{\partial\eta\partial\beta} &= \mu\eta g(\mu+\widetilde{\beta})\frac{\left(\widetilde{\beta}-\beta\right)^2-\left(\sigma^2+\eta^2\right)}{\left(\sigma^2+\eta^2\right)^2}+\frac{\beta\sqrt{N}}{\eta^2}\phi\left(z_\alpha-\frac{\beta\sqrt{N}}{\eta}\right)\left(1+z_\alpha\left(z_\alpha-\frac{\beta\sqrt{N}}{\eta}\right)\right)
\end{aligned}$$

In equilibrium, $\beta=\beta_p^*=\widetilde{\beta}$ and $\frac{\beta_p^*\sqrt{N}}{\eta_p^*}=\frac{1}{z_\alpha}$ so these derivatives become:

$$\begin{aligned}
\frac{\partial^2\Pi}{\partial\beta^2} &= -\frac{z_\alpha^2\sqrt{N}}{\eta_p^*}\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)<0\\
\frac{\partial^2\Pi}{\partial\eta^2} &= -\frac{z_\alpha\beta_p^*}{\eta_p^{*2}}\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)<0\\
\frac{\partial^2\Pi}{\partial\eta\partial\beta} &= -\mu\eta_p^* g(\mu+\beta_p^*)\frac{1}{\sigma^2+\eta_p^{*2}}+\frac{z_\alpha}{\eta_p^*}\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)
\end{aligned}$$

The determinant of the Hessian, $\frac{\partial^2\Pi}{\partial\beta^2}\frac{\partial^2\Pi}{\partial\eta^2}-\left(\frac{\partial^2\Pi}{\partial\eta\partial\beta}\right)^2$, is equal to

$$\begin{aligned}
&\frac{z_\alpha^2\sqrt{N}}{\eta_p^*}\frac{z_\alpha\beta_p^*}{\eta_p^{*2}}\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)^2-\left(\frac{z_\alpha}{\eta_p^*}\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)-\mu\eta_p^* g(\mu+\beta_p^*)\frac{1}{\sigma^2+\eta_p^{*2}}\right)^2\\
&= 2\frac{\mu g(\mu+\beta_p^*)}{\sigma^2+\eta_p^{*2}}z_\alpha\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)-\left(\frac{\mu\eta_p^* g(\mu+\beta_p^*)}{\sigma^2+\eta_p^{*2}}\right)^2\\
&= \frac{\mu}{\sqrt{2\pi\left(\sigma^2+\eta_p^{*2}\right)}\left(\sigma^2+\eta_p^{*2}\right)}\left[2z_\alpha\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)-\frac{\eta_p^{*2}\lambda}{\sigma^2+\eta_p^{*2}}\right]\\
&= \frac{2\pi\sigma^2\lambda^3-\mu^2\left[1-\Phi\left(z_\alpha-\frac{1}{z_\alpha}\right)-z_\alpha\phi\left(z_\alpha-\frac{1}{z_\alpha}\right)\right]}{\mu\sqrt{2\pi}\left(\sigma^2+\eta_p^{*2}\right)^{\frac{3}{2}}}
\end{aligned}$$

Its sign is the one of the numerator. In Figure 3.2 we plot the parameter range

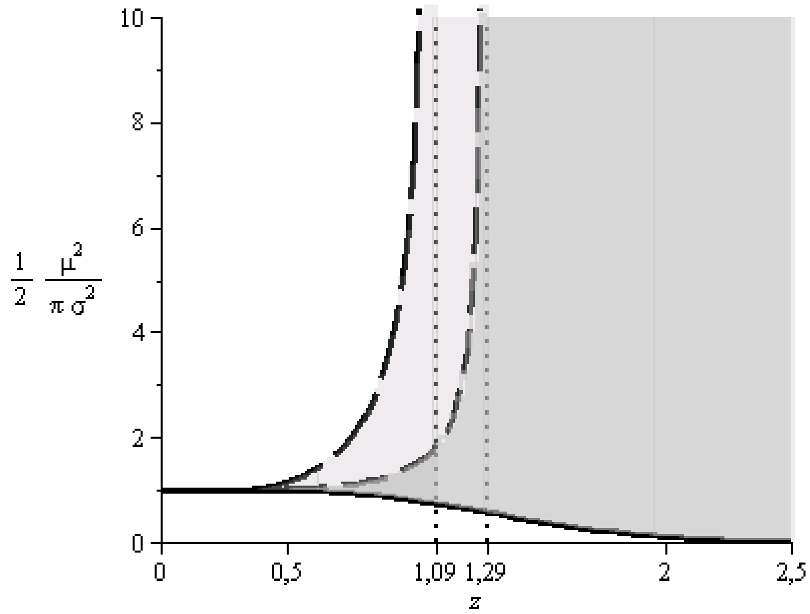

Figure 3.2: In both shaded areas the second-order conditions are satisfied and β_p^* and η_p^* are positive, that is, $\frac{\mu^2}{2\pi\lambda^2} - \sigma^2 > 0$. In the dark shaded area $\frac{\partial \beta_p^*}{\partial z_\alpha}$ is positive.

for which it is positive and, therefore, the second-order conditions are satisfied. In particular, they are satisfied if $z_\alpha \geq 1.09$ (i.e., $\alpha \leq 13.79\%$), in which case $1 - \Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) - z_\alpha \phi\left(z_\alpha - \frac{1}{z_\alpha}\right) < 0$.

■

Proof of Corollary 3.2. Consider the effect of z_α. Since $\frac{\partial \lambda}{\partial z_\alpha} = -z_\alpha^2 \phi\left(z_\alpha - \frac{1}{z_\alpha}\right) < 0$, $\frac{\partial \eta_p^*}{\partial z_\alpha} > 0$. The effect on the equilibrium bias is more involved:

$$\frac{\partial \beta_p^*}{\partial z_\alpha} = \frac{\sqrt{\frac{\mu^2}{2\pi\lambda^2} - \sigma^2}}{z_\alpha^2 \sqrt{N}} \left[\frac{\mu^2 z_\alpha^3 \phi\left(z_\alpha - \frac{1}{z_\alpha}\right)}{\lambda\left(\mu^2 - 2\pi\lambda^2\sigma^2\right)} - 1 \right].$$

Since $\frac{\mu^2}{\sigma^2 2\pi} > \lambda^2$ for the bias to be positive, $\frac{\partial \beta}{\partial z_\alpha}$ is positive if:

$$2\pi\lambda^3\sigma^2 > \mu^2 \left[\lambda - z_\alpha^3 \phi\left(z_\alpha - \frac{1}{z_\alpha}\right)\right].$$

In Figure 3.2 we plot the parameter range for which this inequality holds. In particular, it holds if $z_\alpha \geq 1.29$ (i.e., $\alpha \leq 9.85\%$) in which case $\lambda - z_\alpha^3 \phi\left(z_\alpha - \frac{1}{z_\alpha}\right) < 0$.

The remaining comparative statics are straightforward. ■

Proof of Proposition 3.3. The first-order conditions of the seller's problem with respect to β and η, respectively, are

$$pg\left(\overline{s}\right) - \left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right) + z_\alpha \phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right) = 0 \quad (3.17a)$$

$$-p\eta g(\overline{s})\left[\frac{\beta - \widetilde{\beta}}{\sigma^2 + \eta^2} + \frac{p-\mu}{\sigma^2}\right] - \frac{1}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left[1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha\right] = 0 \quad (3.17b)$$

where $\overline{s} = p + \widetilde{\beta} + (p-\mu)\frac{\eta^2}{\sigma^2}$. In the equilibrium, $\beta = \widetilde{\beta}$. If $p > (<)\mu$ then $\frac{1}{\sqrt{N}z_\alpha} < (>)\frac{\beta}{\eta}$ in order for (3.17b) to hold. ■

Derivation of the buyer's posterior in Section 3.5.2. See Theil (1971), ch. 4.7, for the details of the derivation of posterior in the multivariate normal case.

Let us write the covariance matrix of θ_1, s_1 and s_2. Using

$$\begin{aligned}
Cov\left(\theta_1, s_1\right) &= Cov\left(\theta_1, \theta_1 + \varepsilon_1\right) = Var\left(\theta_1\right) = \sigma_1^2 \\
Cov\left(\theta_1, s_2\right) &= Cov\left(\theta_1, \theta_2 + \varepsilon_2\right) = Cov\left(\theta_1, \theta_2\right) = \rho_\theta\sigma_1\sigma_2 \\
Cov\left(s_1, s_2\right) &= Cov\left(\theta_1 + \varepsilon_1, \theta_2 + \varepsilon_2\right) = Cov\left(\theta_1, \theta_2\right) + Cov\left(\varepsilon_1, \varepsilon_2\right) = \rho_\theta\sigma_1\sigma_2 + \rho_\varepsilon\eta_1\eta_2
\end{aligned}$$

the matrix is

	θ_1	s_1	s_2
θ_1	σ_1^2	σ_1^2	$\rho_\theta\sigma_1\sigma_2$
s_1	σ_1^2	$\sigma_1^2 + \eta_1^2$	$\rho_\theta\sigma_1\sigma_2 + \rho_\varepsilon\eta_1\eta_2$
s_2	$\rho_\theta\sigma_1\sigma_2$	$\rho_\theta\sigma_1\sigma_2 + \rho_\varepsilon\eta_1\eta_2$	$\sigma_2^2 + \eta_2^2$

Denote the covariance matrix of s_1 and s_2 as Σ and find its inverse:

$$\Sigma^{-1} = \frac{1}{\det\left(\Sigma\right)}\begin{pmatrix} \sigma_2^2 + \eta_2^2 & -\rho_\theta\sigma_1\sigma_2 - \rho_\varepsilon\eta_1\eta_2 \\ -\rho_\theta\sigma_1\sigma_2 - \rho_\varepsilon\eta_1\eta_2 & \sigma_1^2 + \eta_1^2 \end{pmatrix},$$

where $\det(\Sigma) = (\sigma_2^2+\eta_2^2)(\sigma_1^2+\eta_1^2) - (\rho_\theta\sigma_1\sigma_2+\rho_\varepsilon\eta_1\eta_2)^2$. The expectation of θ_1 conditional on s_1 and s_2, $E\left[\theta_1 \mid s_1, s_2, \eta_1, \eta_2, \widetilde{\beta}_1, \widetilde{\beta}_2\right]$, equals to

$$
\begin{aligned}
& \mu_1 + \begin{pmatrix} \sigma_1^2 & \rho_\theta\sigma_1\sigma_2 \end{pmatrix} \Sigma^{-1} \begin{pmatrix} s_1-\mu_1-\widetilde{\beta}_1 \\ s_2-\mu_2-\widetilde{\beta}_2 \end{pmatrix} \\
= \; & \mu_1 + \frac{1}{\det(\Sigma)} \begin{pmatrix} \sigma_1^2(\sigma_2^2+\eta_2^2) - \rho_\theta\sigma_1\sigma_2(\rho_\theta\sigma_1\sigma_2+\rho_\varepsilon\eta_1\eta_2) \\ -\sigma_1^2(\rho_\theta\sigma_1\sigma_2+\rho_\varepsilon\eta_1\eta_2) + \rho_\theta\sigma_1\sigma_2(\sigma_1^2+\eta_1^2) \end{pmatrix}^{-1} \begin{pmatrix} s_1-\mu_1-\widetilde{\beta}_1 \\ s_2-\mu_2-\widetilde{\beta}_2 \end{pmatrix} \\
= \; & \mu_1 + \frac{\sigma_1^2(\sigma_2^2+\eta_2^2) - \rho_\theta\sigma_1\sigma_2(\rho_\theta\sigma_1\sigma_2+\rho_\varepsilon\eta_1\eta_2)}{(\sigma_2^2+\eta_2^2)(\sigma_1^2+\eta_1^2) - (\rho_\theta\sigma_1\sigma_2+\rho_\varepsilon\eta_1\eta_2)^2} \left(s_1-\mu_1-\widetilde{\beta}_1\right) \\
& + \frac{-\sigma_1^2\rho_\varepsilon\eta_1\eta_2 + \rho_\theta\sigma_1\sigma_2\eta_1^2}{(\sigma_2^2+\eta_2^2)(\sigma_1^2+\eta_1^2) - (\rho_\theta\sigma_1\sigma_2+\rho_\varepsilon\eta_1\eta_2)^2} \left(s_2-\mu_2-\widetilde{\beta}_2\right) \\
= \; & \mu_1 + \frac{1+\frac{\eta_2^2}{\sigma_2^2} - \rho_\theta\left(\rho_\theta+\rho_\varepsilon\frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)}{\left(1+\frac{\eta_2^2}{\sigma_2^2}\right)\left(1+\frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta+\rho_\varepsilon\frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2} \left(s_1-\mu_1-\widetilde{\beta}_1\right) \\
& + \frac{\rho_\theta\frac{\eta_1^2}{\sigma_1\sigma_2} - \rho_\varepsilon\frac{\eta_1\eta_2}{\sigma_2^2}}{\left(1+\frac{\eta_2^2}{\sigma_2^2}\right)\left(1+\frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta+\rho_\varepsilon\frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2} \left(s_2-\mu_2-\widetilde{\beta}_2\right).
\end{aligned}
$$

■

Proof of Proposition 3.4.

Seller 1 maximizes revenues (3.10) minus costs (3.4b) taking η_2 as given. In the equilibrium $\beta_1 = \widetilde{\beta}_1$ and, as in the proof of Proposition 3.1, $\frac{\beta_1\sqrt{N}}{\eta_1} = \frac{1}{z_\alpha}$. The first-order condition with respect to the bias becomes

$$
\frac{1+\frac{\eta_2^2}{\sigma_2^2} - \rho_\theta\left(\rho_\theta+\rho_\varepsilon\frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)}{\left(1+\frac{\eta_2^2}{\sigma_2^2}\right)\left(1+\frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta+\rho_\varepsilon\frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2} = 1-\Phi\left(z_\alpha-\frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha-\frac{1}{z_\alpha}\right). \quad (3.18)
$$

Note that (3.18) is symmetric in "normalized" noises $\frac{\eta_1}{\sigma_1}$ and $\frac{\eta_2}{\sigma_2}$. To find symmetric equilibria, denote $\widetilde{\eta} = \frac{\eta_1}{\sigma_1} = \frac{\eta_2}{\sigma_2}$ and rewrite (3.18) as

$$
\left(1-\rho_\varepsilon^2\right)\widetilde{\eta}^4 - (\kappa-1)(1-\rho_\theta\rho_\varepsilon)\widetilde{\eta}^2 - \kappa\left(1-\rho_\theta^2\right) = 0
$$

which is a quadratic equation in $\widetilde{\eta}^2$ (unless $\rho_\varepsilon = \pm 1$) (κ is defined in Proposition

3.1). It has two roots of opposite signs and the positive root is

$$\widetilde{\eta}^2 = \frac{(\kappa-1)(1-\rho_\theta\rho_\varepsilon) + \sqrt{(\kappa-1)^2(1-\rho_\theta\rho_\varepsilon)^2 + 4\kappa(1-\rho_\theta^2)(1-\rho_\varepsilon^2)}}{2(1-\rho_\varepsilon^2)}.$$

Considering the cases of $\rho_\varepsilon = \pm 1$ and $\rho_\theta = \pm 1$ is straightforward.

We could not prove the absence of asymmetric equilibria; however, we have not encountered them in many numerical examples that we plotted.[46]

The second-order conditions are difficult to check but they are satisfied, by continuity, at least for ρ_ε and ρ_θ close to zero. ■

Proof of Proposition 3.7. The proof is similar to the one of Proposition 3.1. The first-order conditions of the seller's problem with respect to β and η, respectively, are

$$\frac{\sigma^2}{\sigma^2+\eta^2} - d\left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}\right) + z_\alpha\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}\right)\right) \tag{3.19a}$$

$$-\left(\beta-\widetilde{\beta}\right)\frac{2\eta\sigma^2}{\left(\sigma^2+\eta^2\right)^2} - d\frac{1}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}\right)\left[1 - \frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}z_\alpha\right]\frac{\eta}{\sqrt{\sigma_m^2+\eta^2}} \tag{3.19b}$$

In the equilibrium $\beta = \widetilde{\beta}$ and from (3.19b) $\frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}} = \frac{1}{z_\alpha}$. Plug this into (3.19a) to obtain η_s^*. If $\kappa_d \geq 0$, then $\eta_s^* = \sigma\sqrt{\kappa_d}$ and so $\beta_s^* = \frac{\sqrt{\sigma_m^2+\eta_s^{*2}}}{\sqrt{N}}\frac{1}{z_\alpha} = \frac{\sqrt{\sigma_m^2+\sigma^2\kappa_d}}{\sqrt{N}}\frac{1}{z_\alpha}$. If $\kappa_d < 0$, $\eta_s^* = 0$ and (3.19a) determines β_s^*.

To check the second-order conditions, differentiate (3.19a) with respect to β to

[46] To find asymmetric equilibria (or to prove that they do not exist), denote $\widetilde{\eta}_1 = \frac{\eta_1}{\sigma_1}$ and $\widetilde{\eta}_2 = \frac{\eta_2}{\sigma_2}$. The reaction curve $\widetilde{\eta}_1(\widetilde{\eta}_2)$ can be found from rewriting (3.18) as

$$\widetilde{\eta}_1^2\left(1+\widetilde{\eta}_2^2\left(1-\rho_\varepsilon^2\right)\right) + (\kappa-1)\widetilde{\eta}_1\widetilde{\eta}_2\rho_\theta\rho_\varepsilon - \kappa\left(1-\rho_\theta^2+\widetilde{\eta}_2^2\right) = 0.$$

Equilibria are its intersections with the (symmetric) reaction curve $\widetilde{\eta}_2(\widetilde{\eta}_1)$.

obtain $\frac{\partial^2 \Pi}{\partial \beta^2}$ and (3.19b) with respect to η and β to obtain $\frac{\partial^2 \Pi}{\partial \eta^2}$ and $\frac{\partial^2 \Pi}{\partial \eta \partial \beta}$, respectively:

$$\begin{aligned}
\frac{\partial^2 \Pi}{\partial \beta^2} &= -d\frac{\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}\right)\left(1-\frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}z_\alpha + z_\alpha^2\right) \\
\frac{\partial^2 \Pi}{\partial \eta^2} &= -\left(\beta-\widetilde{\beta}\right)\frac{\sigma^2-3\eta^2}{\left(\sigma^2+\eta^2\right)^3} \\
&\quad -d\frac{\beta^2\sqrt{N}}{\left(\sigma_m^2+\eta^2\right)^{\frac{3}{2}}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}\right)\left(1-\frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}z_\alpha + z_\alpha^2\right)\frac{\eta^2}{\sigma_m^2+\eta^2} \\
&\quad -d\frac{1}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}\right)\left[1-\frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}z_\alpha\right]\frac{\sigma^2}{\left(\sigma_m^2+\eta^2\right)^{\frac{3}{2}}} \\
\frac{\partial^2 \Pi}{\partial \eta \partial \beta} &= -\frac{2\eta\sigma^2}{\left(\sigma^2+\eta^2\right)^2}+d\frac{\beta\sqrt{N}}{\sigma_m^2+\eta^2}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}\right)\left(1-\frac{\beta\sqrt{N}}{\sqrt{\sigma_m^2+\eta^2}}z_\alpha + z_\alpha^2\right)\frac{\eta}{\sqrt{\sigma_m^2+\eta^2}}
\end{aligned}$$

Let us first show that second-order condition is satisfied when $\kappa_d = \frac{1-d\lambda}{d\lambda} < 0$, where $\lambda = 1-\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)$. For this we need $\frac{\partial^2 \Pi}{\partial \beta^2} < 0$, that is, $1-\frac{\beta^*\sqrt{N}}{\sigma_m}z_\alpha + z_\alpha^2 > 0$. Take d such that (3.19a) yields $\eta^* = 0$, that is, (3.19a) is written as $1-d\lambda = 0$. This is the highest d for which we have an interior equilibrium, that is, when $\frac{\beta^*\sqrt{N}}{\sqrt{\sigma_m^2+\eta^{*2}}} = \frac{1}{z_\alpha}$. Increase d. The optimal noise η^* is still zero, and so we have a problem with one variable β. Since differentiating (3.19a) with respect to d gives $-\lambda < 0$, β^* decreases with d. Then, $\frac{\beta^*\sqrt{N}}{\sigma_m} < \frac{1}{z_\alpha}$. This implies $z_\alpha + \frac{1}{z_\alpha} - \frac{\beta^*\sqrt{N}}{\sigma_m} > 0$ and so $\frac{\partial^2 \Pi}{\partial \beta^2} < 0$.

Now consider the interior equilibrium. We have $\beta = \beta^* = \widetilde{\beta}$ and $\frac{\beta^*\sqrt{N}}{\sqrt{\sigma^2+\eta^{*2}}} = \frac{1}{z_\alpha}$ so the second cross-derivatives become

$$\begin{aligned}
\frac{\partial^2 \Pi}{\partial \beta^2} &= -d\frac{\sqrt{N}}{\sqrt{\sigma_m^2+\eta^{*2}}}\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)z_\alpha^2 < 0 \\
\frac{\partial^2 \Pi}{\partial \eta^2} &= -d\frac{\beta^{*2}\eta^{*2}\sqrt{N}}{\left(\sigma_m^2+\eta^{*2}\right)^{\frac{5}{2}}}\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)z_\alpha^2 < 0 \\
\frac{\partial^2 \Pi}{\partial \eta \partial \beta} &= -\frac{2\eta^*\sigma^2}{\left(\sigma^2+\eta^{*2}\right)^2}+d\frac{\beta^*\eta^*\sqrt{N}}{\left(\sigma_m^2+\eta^{*2}\right)^{\frac{3}{2}}}\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)z_\alpha^2
\end{aligned}$$

Check that the determinant of the Hessian is positive (omit the argument in

$\phi\left(z_\alpha - \frac{1}{z_\alpha}\right))$

$$\begin{aligned}
\frac{\partial^2\Pi}{\partial\beta^2}\frac{\partial^2\Pi}{\partial\eta^2} - \left(\frac{\partial^2\Pi}{\partial\eta\partial\beta}\right)^2 &= d^2\frac{\beta^{*2}\eta^{*2}N}{\left(\sigma_m^2+\eta^{*2}\right)^3}\phi^2 z_\alpha^4 - \left(\frac{2\eta^*\sigma^2}{\left(\sigma^2+\eta^{*2}\right)^2}\right)^2 \\
&\quad +2\frac{2\eta^*\sigma^2}{\left(\sigma^2+\eta^{*2}\right)^2}d\frac{\beta^*\eta^*\sqrt{N}}{\left(\sigma_m^2+\eta^{*2}\right)^{\frac{3}{2}}}\phi z_\alpha^2 - d^2\frac{\beta^{*2}\eta^{*2}\sqrt{N}}{\left(\sigma_m^2+\eta^{*2}\right)^3}\phi^2 z_\alpha^4 \\
&= \frac{4\eta^{*2}\sigma^2}{\left(\sigma^2+\eta^{*2}\right)^2}\left(d\frac{\beta^*\sqrt{N}}{\left(\sigma_m^2+\eta^{*2}\right)^{\frac{3}{2}}}\phi z_\alpha^2 - \frac{\sigma^2}{\left(\sigma^2+\eta^{*2}\right)^2}\right) \\
&= \frac{4\eta^{*2}\sigma^2}{\left(\sigma^2+\eta^{*2}\right)^2}\left(d\frac{1}{\sigma_m^2+\eta^{*2}}\phi z_\alpha - \frac{\sigma^2}{\left(\sigma^2+\eta^{*2}\right)^2}\right) \\
&= \frac{4\eta^{*2}\sigma^2}{\left(\sigma^2+\eta^{*2}\right)^2}\left(d\frac{1}{\sigma_m^2+\sigma^2\frac{1-d\lambda}{d\lambda}}\phi z_\alpha - \frac{d^2\lambda^2}{\sigma^2}\right) \\
&= \frac{4\eta^{*2}\sigma^2}{\left(\sigma^2+\eta^{*2}\right)^2}\left(d\frac{d\lambda}{\sigma_m^2 d\lambda+\sigma^2\left(1-d\lambda\right)}\phi z_\alpha - \frac{d^2\lambda^2}{\sigma^2}\right) \\
&= \frac{4\eta^{*2}}{\left(\sigma^2+\eta^{*2}\right)^2}d^2\lambda\left(\frac{\phi z_\alpha}{\frac{\sigma_m^2}{\sigma^2}d\lambda+1-d\lambda} - \lambda\right) \\
&= \frac{4\eta^{*2}}{\left(\sigma^2+\eta^{*2}\right)^2}\frac{d^2\lambda}{\frac{\sigma_m^2}{\sigma^2}d\lambda+1-d\lambda}\left(\phi z_\alpha - \lambda\left(\frac{\sigma_m^2}{\sigma^2}d\lambda+1-d\lambda\right)\right)
\end{aligned}$$

Since we consider $\sigma_m^2 \leq \sigma^2$, a higher d decreases $\frac{\sigma_m^2}{\sigma^2}d\lambda+1-d\lambda$ and thus makes the determinant higher. In the baseline model, the sign of the determinant is determined by $\phi z_\alpha - \frac{\kappa}{(1+\kappa)^2}$ (see the proof of Proposition 3.1). Note that $\frac{\kappa}{(1+\kappa)^2} = \lambda\left(1-\lambda\right)$ and, therefore, the sign of the determinant here is the same as in the baseline model if and only if $\frac{\sigma_m^2}{\sigma^2}d\lambda+1-d\lambda = 1-\lambda$, that is, $\left(1-\frac{\sigma_m^2}{\sigma^2}\right)d=1$. ■

Bibliography

ANDERSON, A., AND L. M. B. CABRAL (2007): "Go for broke or play it safe? Dynamic competition with choice of variance," *The RAND Journal of Economics*, 38(3), 593–609.

ARMSTRONG, M. (2008): "Interactions between Competition and Consumer Policy," *Competition Policy International*, 4(1), 96–147.

BLUME, A., AND O. BOARD (2010): "Intentional Vagueness," http://www.pitt.edu/~ojboard/papers/vagueness.pdf.

CASELLA, G., AND R. L. BERGER (2002): *Statistical inference*. Duxbury Press, Belmont.

CHAKRABORTY, A., AND R. HARBAUGH (2010): "Persuasion by Cheap Talk," *American Economic Review*, 100(5), 2361–82.

CRAWFORD, V. P., AND J. SOBEL (1982): "Strategic Information Transmission," *Econometrica*, 50(6), 1431–1451.

DAVIS, M. L. (1994): "The Value of Truth and the Optimal Standard of Proof in Legal Disputes," *Journal of Law, Economics, and Organization*, 10(2), 343–359.

DEWATRIPONT, M., I. JEWITT, AND J. TIROLE (1999): "The Economics of Career Concerns, Part I: Comparing Information Structures," *The Review of Economic Studies*, 66(1), 183–198.

FTC (2008): "Before the Board of Governors of the Federal Reserve System: In the Matter of Request for Comments on Truth in Lending, Proposed Rule Docket No. R-1305," Available at http://www.federalreserve.gov/SECRS/2008/April/20080424/R-1305/R-1305_1349_1.pdf (accessed on the 06 December 2011).

Gabaix, X., and D. Laibson (2006): "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *The Quarterly Journal of Economics*, 121(2), 505–540.

Holmström, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169–182.

Inderst, R., and M. Ottaviani (2009a): "Misselling through Agents," *The American Economic Review*, 99(3), 883–908.

——— (2009b): "Sales Talk, Cancellation Terms, and the Role of Consumer Protection," http://didattica.unibocconi.it/mypage/upload/48832_20111118_025320 _REFUNDS.PDF.

Issacharoff, S. (1999): "Group Litigation of Consumer Claims: Lessons from the U.S. Experience," *Texas International Law Journal*, 34, 135–150.

Ivanov, M. (2011): "Information Revelation in Competitive Markets," *Economic Theory*, forthcoming.

Johnson, J. P., and D. P. Myatt (2006): "On the Simple Economics of Advertising, Marketing, and Product Design," *The American Economic Review*, 96(3), 756–784.

Judd, K. L., and M. H. Riordan (1994): "Price and Quality in a New Product Monopoly," *The Review of Economic Studies*, 61(4), 773–789.

Kartik, N. (2009): "Strategic Communication with Lying Costs," *Review of Economic Studies*, 76(4), 1359–1395.

KRÄKEL, M., P. NIEKEN, AND J. PRZEMECK (2008): "Risk Taking in Winner-Take-All Competition," Bonn Econ Discussion Paper bgse7_2008.

LEWIS, T. R., AND D. E. M. SAPPINGTON (1994): "Supplying Information to Facilitate Price Discrimination," *International Economic Review*, 35(2), 309–327.

LI, W. (2010): "Peddling Influence through Intermediaries," *The American Economic Review*, 100(3), 1136–1162.

MEYER, M. A., AND J. VICKERS (1997): "Performance Comparisons and Dynamic Incentives," *The Journal of Political Economy*, 105(3), 547–581.

MICELI, T. J. (1990): "Optimal Prosecution of Defendants Whose Guilt Is Uncertain," *Journal of Law, Economics, and Organization*, 6(1), 189–201.

POLINSKY, A. M., AND S. SHAVELL (2000): "The Economic Theory of Public Enforcement of Law," *Journal of Economic Literature*, 38(1), 45–76.

POSNER, R. A. (2004): *Frontiers of Legal Theory*. Harvard University Press, Cambridge.

SHELEGIA, S. (2011): "Quality Choice of Experience Goods," http://homepage.univie.ac.at/sandro.shelegia/Personal/Research_files/quality choice.pdf.

SPIEGLER, R. (2006): "Competition over agents with boundedly rational expectations," *Theoretical Economics*, 1(2), 207–231.

SUNSTEIN, C., R. HASTIE, J. PAYNE, D. SCHKADE, AND K. VISCUSI (2002): *Punitive Damages: How Juries Decide*. University Of Chicago Press, Chicago.

THEIL, H. (1971): *Principles of Econometrics*. North-Holland, Amsterdam.

TSETLIN, I., A. GABA, AND R. L. WINKLER (2004): "Strategic Choice of Variability in Multiround Contests and Contests with Handicaps," *Journal of Risk and Uncertainty*, 29(2), 143–158.

Vickers, J. (2004): "Economics for Consumer Policy," in *Proceedings of the British Academy*, vol. 125. Oxford University Press, Oxford.

Chapter 4

Vertical relational contracts and trade credit

4.1 Introduction

In developing countries, firms have to deal with a number of constraints that are not present in developed countries and that necessarily interfere in the way firms compete and interact with each other along the vertical chain. Among other adversities, they face weak institutions, underdeveloped legal system and capital markets, corruption and non-prosecuted criminal offences such as robberies. This makes the conclusions of the standard literature on industrial organization less useful to understand the functioning of markets in these countries.

As can be seen from Figure 4.1, the number of competition authorities in developing countries has rocketed from under ten in 1990 to almost eighty in 2008. Therefore, understanding the industrial organization of developing countries has become a timely issue as these young antitrust authorities need to base their policies on theories that fit their environments rather than the ones created for developed countries.

The goal of this paper is to incorporate some stylized facts of developing countries into the analysis of vertical relations. We want to explore how these stylized facts

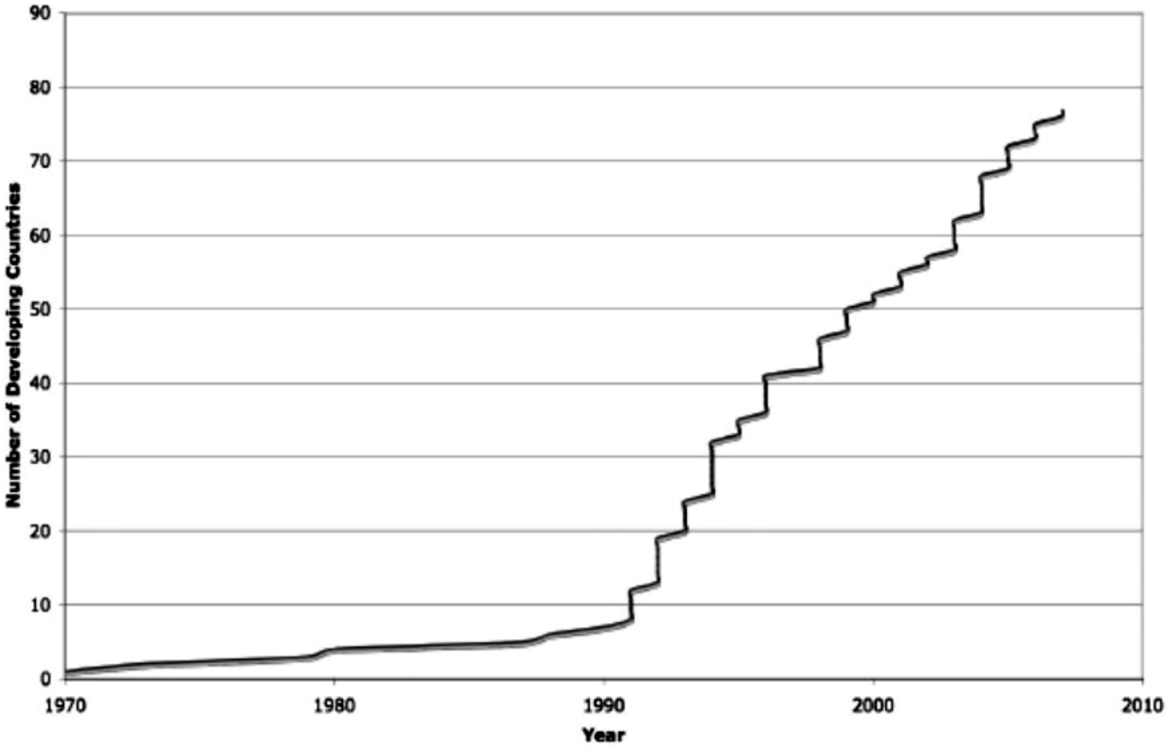

Figure 4.1: Time-line of adopting competition laws in the developing world. *Source: Waked (2010)*

affect firms' behavior and market outcomes and to generate testable predictions.

The first stylized fact we want to consider is the underdevelopment of legal systems. An important consequence of the firms not being able to resort to the rule of law to enforce their contracts is the need to rely on relational contracts, where the agreement is self-enforced thanks to the gains from future repeated trade. McMillan and Woodruff (1999a and 1999b) find evidence of relational contracting using surveys from 259 non-state firms in Vietnam. They show how prior experience with a trading partner and prior information gathering on that partner are associated with increased provision of credit. Johnson, McMillan and Woodruff (2002) document similar findings in post-communist countries.

Indeed, supplier's credit (the temporal dissociation between delivery and payment) is particularly common due to the underdevelopment of financial markets. Among others, McMillan and Woodruff (1999a and 1999b) and Fafchamps (1997 and 2000) point out the important role played by supplier credit[1] as compared to

[1] Trade credit is normally defined as the dissociation between invoicing and payment and it is

bank credit in developing countries. This is the second feature we want to include in our model.

The third and last ingredient that we want to incorporate is the existence of adverse shocks that make the firms unable to honour their agreements, such as a warehouse being robbed or unexpected low revenues[2]. These adverse shocks are not always observed by the other contracting party which creates a moral hazard (with hidden knowledge) problem, as the downstream firms may not return the trade credit and pretend that was a shock.[3]

With our analysis, we want to address the following questions. How is the contract offered by the manufacturer affected by the trade credit and the asymmetric information problem? What are the resulting market distortions?

Although we have framed the introduction in a developing country setting, the results are also useful to understand inter-firm relations in developed countries, where trade credit is also very common, with firms that operate in the shadow economy, that do not declare all their transactions or whose trade volume is small and does not justify the cost of going to the courts. In all these instances, firms are left to rely on relational contracts.

Another application of this framework are black markets like drug trade. In these markets, obviously, contracts cannot be legally enforced, traders are usually credit constrained and are also subject to many unobservable shocks. Since little is known of these markets due to their undercover aspect, these results are valuable for shedding some light in the understanding of these markets.

In the model that we consider there is an upstream firm that supplies a good and offers trade credit to a cashless downstream firm. To fix ideas, imagine that the upstream firm ("she") is a manufacturer and the downstream firm ("he") is a retailer. For instance, it can be the case that the manufacturer's machinery can be

also widely used in developed countries.

[2]Fafchamps et al. (forthcoming) document the volatility of sales and profits of microenterprises in Ghana and show how this volatility is not created by measurment errors.

[3]Note that the shocks may be observable and the trading party who offers the trade credit may still not be able to determine whether the other party is unable to fullfil the agreement because of the shock - an expensive or unavailable legal process is needed to have such information disclosed.

used as collateral and this makes her less credit constrained than the retailer. The lack of access to the financial system for the retailer may also be explained by the retailer operating in the shadow economy. The manufacturer proposes a quantity forcing contract of the form of a quantity and a repayment[4] and the retailer pays back to the manufacturer only after having sold the good. However, a shock may occur before the payment, making him unable to repay either part or the whole amount. The manufacturer cannot observe if the failure of payment is due to the shock or to the retailer keeping the money[5] and she punishes him by terminating the relational contract for a number of periods.

We first solve a simple example (Section 4.3), where the shock makes the retailer lose all the revenues. Therefore, the retailer can either payback all the credit or make no payment at all (if he received the shock). Then we introduce a second demand into the example. In this extended example (Section 4.4), the retailer places the order in the market and receives high or (if the shock occurs) low revenues. In this setup, the retailer can then either return the whole credit, or if unlucky, give back only part of it (we also allow for the possibility that the retailer just walks away with all the revenues). Finally, we analyze a general model where there is a continuum of states of demand (Section 4.5).

In a perfect world with symmetric information (i.e. observable demand shocks), both firms maximize their joint profits and no punishment is used. The quantity sold in the market is smaller than in the absence of shocks because the firms take into account the "effective" marginal cost which not only accounts for the production cost but also for the likelihood of the shocks.

When demand shocks are not observed by the manufacturer, trade credit, by postponing the payment until the shock is realized, creates a source of asymmetric information. If contracts can be legally enforced and the retailer has access to

[4] Since the quantity is delivered before any private information becomes known to the retailer, it is not used for sorting purposes. As a result, using a two-part tariff or a more complex nonlinear scheme is equivalent to choosing a quantity and a total repayment.

[5] In order to find out, the manufacturer would need to use (if existent) a legal system that is too coslty or too corrupt.

unlimited funds, it is possible to replicate the first best outcome by asking the retailer to repay a fixed payment (equal to the expected revenues) regardless of the state. The first best is no longer achievable, however, when either the retailer has access to unlimited funds but contracts cannot be legally enforced or contracts can be enforced but the retailer is cashless. This paper explores the consequences of the first departure, although we have also (unnecessarily) assumed that the retailer is cashless to make the setup closer to a developing country environment (See Fafchamps (1997 and 2000)). Section 4.4.1 illustrate the persistence of the distortions even when the cash constraint does not bind.

From the analysis of Example 1, it emerges that in the optimal contract, the manufacturer distorts the quantity downwards. The reason why she does not choose the profit maximizing quantity and then shares the extra profits with the retailer through a lower repayment is because the efficient quantity generates large revenues which increases the retailer's incentives to steal them. In order to curb these incentives, the manufacturer needs to increase the termination period which is also costly for her. Thus, despite having enough available instruments to set the quantity double marginalization occurs. The downward distortion in the quantity is entirely driven by the retailer's impatience. Moreover, it is also present in Example 2 and the general model and hence it is one of the general characteristics of the relational contract in our setting.

In Example 2, we find that if the high and the low demands are very similar (i.e. the shock is mild), then the manufacturer prefers to charge the same amount in both situations and to terminate the contract forever only if there is no repayment. As the low demand shrinks (i.e. the shock increases), then the limited liability constraint starts binding and the manufacturer asks for a larger repayment in the high demand state than in the low demand state. In order to implement this contract, however, she needs to interrupt the relationship with the retailer for a positive but finite number of periods following a low repayment. To minimize the punishment imposed, she asks for the maximum possible repayment in the low state. When the limited liability constraint binds, termination is used and the outcome is bounded away from efficiency and the downward distortion persist in the quantity even as the firms become arbitrarily patient.

In the general model of Section 4.5, we find that the optimal contract offered by the manufacturer resembles a debt contract. In particular, if the states of demand are good enough (above some threshold optimally chosen by the manufacturer), the manufacturer ask for the same repayment and never terminates the contract. The larger this fixed repayment, the smaller the quantity distortion imposed by the manufacturer. Otherwise, the manufacturer asks for the highest possible repayment (which in the absence of retail costs is equal to the revenues) and punishes for a number of periods. Asking for the maximum possible repayment is optimal as it allows the manufacturer to soften the termination policy while keeping the retailer's incentives unaffected. Furthermore, the smaller the repayment, the larger the termination period. Indeed, the manufacturer makes the retailer value the relationship by distributing rents in the good times. The retailer always keeps a minimum amount of rents (equal to the current period revenues) due to the fact that he can always choose not to return the credit. If the value of the future (i.e. the discount factor) is large enough, then the manufacturer actually prefers to leave even more rents to the retailer. This is because trying to extract these extra rents would imply larger termination periods which are very costly for the manufacturer as she also loses future trade. In this last case, the contract does not depend on the particular discount factor. As in the Examples, since the limited liability constraint binds for the bad states, punishments are used in equilibrium and the quantity remains distorted even as the discount factor becomes arbitrarily close to 1.

4.1.1 Related literature

This paper belongs to the literature on inter-firm relational contracts with asymmetric information, where incentives are usually given by discretionary payments conditional on good performance (see Malcomson 2010). In this setup, however, the use of trade credit makes the transfer of continuation payoffs unfeasible and instead incentives are provided by the threat of termination, and hence by burning value.

Levin (2003) is the first paper to introduce moral hazard or adverse selection in a principal-agent relational contract in order to identify in which way contract self-enforceability reduces the provision of high-powered incentives. In our model, however, there is a problem of moral hazard with hidden knowledge, where the retailer's actions are observable (whether he repaid or not) but not the information on which their are based (whether he received a shock and of which size).[6] This asymmetric information is reminiscent of the model of Green and Porter (1984) where an oligopoly is colluding in a market with noisy prices. When a low price is observed, firms do not know with certainty if this bad outcome is due to a market shock or a firm deviating to a larger quantity. As in our model, incentives to collude are created by value burning (i.e. the threat of triggering a price war if the observed price is too low).

In Levin's model, the agent undertakes a costly action and the principal rewards him with a fixed fee (to which she can commit) plus a discretionary bonus (that can be positive or negative). The fixed fee enables the parties to share the total surplus in an arbitrary way, and hence their objective is to maximize it (subject to the self-enforceability constraint) and whether is the principal or the agent who offers the contract does not matter. Another consequence of this assumption is that there is no loss of generality in restricting attention to stationary contracts where incentives are provided by punishing and rewarding the agent with transfers that depend on the particular outcome, but not of the time.[7] If deviation occurs, there is

[6] This problem is also known as post-contractual adverse selection, where the type of the agent (which is associated with the size of the shock received, which determines how much of the trade credit he is able to give back) becomes known to the downstream firm after having signed the contract. The Spence-Mirrlees condition does not hold as repaying the credit is equally costly for all types. Since the retailer is credit constraint, he cannot repay more than his revenues. Hence, it is the trade credit constraint that links the type (how much revenues the retailer has) with the repayment.

[7] This is because, following a low outcome, the principal can punish the agent by: (1) moving to a future state that is worse for the agent than the present one; or (2) to pay less to the agent today. Since in this setup both instruments are perfect substitutes, there is no loss of generality in restricting attention to the second option only.

no need to terminate the contract as the principal can ask for a transfer that leaves the agent indifferent between staying in the relation and the outside option.

Trade credit bundled with limited liability imposes important restrictions on the contract that the principal can use. Loosely speaking, trade credit is a large up-front payment to the agent equal to some uncertain revenues. However, contrary to Levin's fixed fee, this payment is not independent of the repayment amount and hence it affects the total surplus (indeed, the size of the trade credit will depend on the quantity sold). As a result, who offers the contract makes a difference and exploring how the relational contract changes when it is offered by the retailer is left for future work. The quantity of the good is delivered to the retailer before the private information is learnt and hence it cannot be used as an instrument for sorting. Moreover, the repayment is always positive[8] and the principal has no other means to reward the agent. Due to the limited liability, the manufacturer cannot punish the agent in monetary terms. As a result termination is necessary to give the agent incentives to repay and stationary contracts may not be optimal. Indeed, Fong and Li (2010) show in their simplified moral hazard version of Levin (2003) how the introduction of limited liability (i.e. a minimum wage needs to be paid) makes stationary contracts non-optimal. As in Green and Porter (1984), we have imposed the use of a stationary two-state (trading and termination) T-stage Markov contract. In the Appendix, we show that in Example 1, the manufacturer does not want to offer a second chance to the retailer by selling a potentially different quantity in the period following no repayment for a large parameter space[9]. Keeping contracting with the retailer after no repayment, although it reduces the inefficiency following a shock, it also increases the incentives not to repay in the first period which will push the manufacturer to distort downwards this quantity even more. Indeed, the optimal punishment in the Green and Porter model involves choosing larger quantities than in the stage-game Nash equilibrium in order to allow equilibria that

[8] In terms of Levin (2003), the bonus is always negative and it is paid by the agent when the outcome is high (i.e. there is no shock) rather than when it is low.

[9] For the parameter specification for which the manufacturer gives a second opportunity to the retailer, we still do not know whether the manufacturer earns more profits than in Example 1.

are "more collusive"[10]. Moreover, the strongly symmetric equilibria of Abreu, Pearce and Stacchetti (1986), which only use value burning to create incentives, show how one can restrict attention to equilibria with the bang-bang property which, under some conditions, implement the maximum pure perfect public equilibrium payoffs using continuation values drawn exclusively from the maximum and minimum pure perfect public equilibrium payoffs.

This paper is closely related to the literature where an entrepreneur is wealth constrained and needs to be financed by an investor, who cannot observe the resulting cashflows of the investment. Hence, the entrepreneur can potentially divert or steal them and incentives to repay are given by liquidating the entrepreneur's assets or by threatening to withhold the investment in the second (and last) period. The early papers in this literature are by Bolton and Scharfstein (1990) and Hart and Moore (1998). Faure-Grimaud (2000) and Povel and Raith (2004a, 2004b) extend the basic two-state model to the case where the cashflows are distributed continuously on a bounded interval. Debt contracts are optimal in this environment because they minimize the probability of inefficient liquidation while inducing the entrepreneur to repay.

Debt contracts[11] are also known to be optimal in the literature of costly state verification (see Townsend (1979) and Gale and Hellwig (1985)), where the cashflows of an investment are also unobserved to the investor but the parties can agree to carry on a costly audit to make the realized cashflow verifiable. If the entrepreneur repays the agreed fixed payment, no inspection takes place; otherwise, an inspection is carried out and the entrepreneur needs to repay all the realized cashflows. Asking for the maximum repayment is optimal as this relaxes the participation constraint of the investor, who can then reduce the repayment and hence the audit set.

Indeed, the probability of liquidation (or not refinancing) and the probability of inspection play the same role as the length of termination imposed by the manufac-

[10] See Mailath and Samuelson (2006) page 349.

[11] Innes (1990) finds that debt contracts are optimal in an environement with moral hazard with limited liability. This is because a debt contract gives the best incentives as it makes the agent residual claimant in the good times and penalises him in the bad times by extracting all the surplus.

turer. First, it relaxes the incentive compatibility of the retailer. Second, like the auditing or liquidating, imposing the termination policy is costly for the manufacturer, because the relation is always profitable (and hence socially valuable).

The more closely related paper to ours is Povel and Raith (2004a)[12]. The authors allow for the size of the investment to be chosen by the entrepreneur as well as how much to repay and the probability of liquidating. Their focus is on showing that the optimal contract is still a debt contract even when the investment choice (as well as the resulting cashflows) is not observed by the investor. They also find that the entrepreneur under-invests as compared to the first best as this decreases the fixed repayment and hence the likelihood of defaulting which may result in an inefficient liquidation. In the same way, the manufacturer of our model (the "investor") chooses to sell less output than the first best to the retailer to soften the (inefficient) termination policy. Our contribution is to endogenize the future value that accrues to the retailer (i.e. the "entrepreneur") if he does not default, as it corresponds to the potential profits generated within the relationship. This allows us to explore the impact of the future on the contract. For instance, when the parties value the future enough, the optimal quantity and the repayment do not depend on the discount factor or the observation that the impatience of the retailer is the only driver of the downward distortion in the quantity.

The goal of this paper is not to explain why and how much trade credit is offered by the manufacturer[13]. We take this decision as given, and rather we explore how trade credit affects the different contract characteristics, such as late payment penalty, quantity and price of the good sold. The main prediction of this analysis is that the use of trade credit does have an important impact on the market outcome.

[12] Faure-Grimaud (2000) and Povel and Raith (2004b) focus on the effect that financial constraints have on the choice of output when the firm is competing a la Cournot with another (financially unconstrained) firm.

[13] Answers to these questions can, for instance, be found in Cunat (2007) who shows how suppliers of services and differentiated goods are more willing to sell on credit than suppliers of standarized goods because they may be harder to replace and hence the downstream firm is more reluctant to default. Also, Smith (1987) finds that trade credit may be a consequence of an agency problem. Indeed, if there is quality variation in the good supplied, the downstream firm may be more reluctant to pay before having had the time to inspect the good.

In particular, the quantity sold in the market is expected to be lower than the efficient one in order to give incentives to the retailer to repay without having to use too severe termination policies. In the setup of Green and Porter (1984), when quantities are chosen from a sufficiently fine grid of points, a similar result emerges. In particular, firms "produce quantities larger than the monopoly output to reduce the incentives to deviate from equilibrium play, which in turn allows equilibrium punishments to be less severe. Because punishments actually occur in equilibrium, this reduced severity is valuable". (Mailath and Samuelson (2006), p. 353). Another implication of this result is that as the firms become arbitrarily patient, the outcome is still bounded away from efficiency. Intuitively, the imperfection in the monitoring implies that on the equilibrium path, the manufacturer must face low continuation payoffs with positive probability. Because the retailer is credit constrained, it is impossible to use inter-temporal transfers of payoffs to maintain efficiency while still providing incentives.

Finally, Buehler and Gartner (forthcoming) also explore the use of relational contracts within two vertically related firms, however their question is very different. They show, how in a vertical relationship where the manufacturer has private information about the manufacturing cost, recommending a retail price may be necessary in order to maximize total profits. The reason is that in order to induce the manufacturer to report the true cost, he needs to be the residual claimant of the costs. This can result in the wholesale price being non-monotonic in the cost and hence the need for recommended retail prices.

The paper proceeds as follows. Section 4.2 introduces the model. Section 4.3 explores a simple example where an unlucky retailer may lose all his revenues before paying to the manufacturer. Section 4.4 extends this example to the case where the retailer faces to possible demands: one is high and the other is low. Section 4.5 generalizes the previous examples to the case where the retailer faces a continuum of demand states. Section 4.6 discusses the empirical implications of the analysis. Finally, Section 4.7 concludes. Tedious computations are relegated to the Appendix.

4.2 The Setup

A manufacturer (M) and a retailer (R) have the opportunity to trade at dates $t = 0, 1, 2, ...$ In each period, M produces a good at a marginal cost $c > 0$ and needs a retailer to market the product to the final consumer. The retailer can sell the good at no cost but he is completely credit constrained and needs to be fully financed by M in order to sell. As a result, M offers trade credit to R, who will then pay M back after selling the good but within the same period (so no interest rate is charged).

In order to remove distortions coming from the manufacturer not having enough instruments to determine the final quantity, we let the manufacturer offer a quantity forcing contract. When contracts are not legally enforced and the retailer does not repay, there are not many instruments available to the manufacturer to punish for misbehavior. For instance, audits or liquidation of the retailer's assets cannot be implemented. Therefore, the manufacturer is left to use the threat of a T period termination as a mean to provide incentives. McMillan and Woodruff (1999a and 1999b) find that such retaliation occurs in Vietnam, although it is not as forceful as one would expect in a standard repeated game framework. An alternative interpretation to the periods of termination is to trade in less profitable terms (for instance by diminishing the quality of the good).[14]

We denote by $0 < \delta < 1$ the discount factor and we assume that in the periods of no trade, both firms get a constant outside option which is normalized to 0.

The timing is summarized in Figure 4.2. In each period, the manufacturer offers a contract to the retailer. The retailer rejects or accepts and if he accepts, he places an order in the market. Then an *iid* shock is realized (and this information is only observed by the retailer) which determines the size of the revenues. Finally, R decides how much to repay and the contract is terminated for a number of periods if it is specified in the contract.

[14] See Baker, Gibbons and Murphy (1994) for an example in the employer-employee relationship framework.

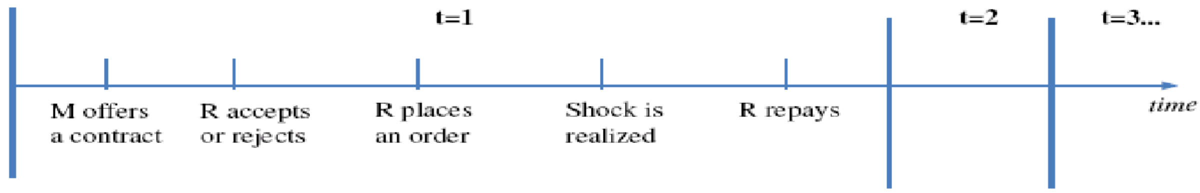

Figure 4.2: Timing of the game

The effect of the shock is to add randomness to the revenues of the retailer. We can interpret this shock as an uncertainty with the respect to the willingness to pay of final consumers. For instance, the retailer could be placing the quantity offered by the manufacturer in the market and some periods he is paid a high price for it and other periods a low price. Another interpretation is the one of an adverse shock whereby demand is certain but either the goods or the revenues are stolen now and then (for instance, by an organized crime group). Finally, the setup could also represent a situation where the uncertainty refers to how many units of a non-perishable (where it is not possible to give back the unsold units)[15] or perishable good are demanded every period in the market.

4.3 Example 1

As an example, we consider the case where the retailer faces a demand $p(q)$ and with probability s, there is a shock that makes him lose all the revenues, so he can either repay the whole credit or nothing. The manufacturer offers a quantity forcing contract $\{q, D, T\}$ where q is the quantity, D the repayment and T periods of termination following the lack of payment.

In a benchmark situation with no asymmetric information and law enforcement, the parties would maximize their joint one-period profits:

$$\max_q (1-s)\,p(q)q - cq$$

[15] If the good was non-perishable and the retailer could give back the unsold units, the problem would become trivial as the source of asymmetric information would disappear.

The first best quantity is determined by:

$$q^{FB} : p'(q)q + p(q) = \frac{c}{1-s} = \widetilde{c}$$

where $\widetilde{c}$ is the effective marginal cost, which accounts for the likelihood of the shock. This is the relevant marginal cost against which we will make comparisons. The parties want to sell less than in the absence of the shock because with some probability they will not receive the revenues but incur the production costs anyway.

With the relational contract, since the manufacturer offers the contract, she maximizes her profits ensuring that the retailer has incentives to repay. Let π_R denote the retailer's present discounted value of selling the good and repaying to the manufacturer from date t on.

$$\pi_R = (1-s)\left[p(q)q - D + \delta\pi_R\right] + s\delta^{T+1}\pi_R$$

The previous equation says that with probability $1-s$ there is no shock, and hence the retailer receives the revenues and pays back the credit to the manufacturer; in which case she renews the relational contract and hence, the game remains in this cooperative phase in the next period. However, with probability s, there is a shock that "destroys" the revenues, the retailer cannot pay back the credit and the game moves to the termination phase in the next period. In this case, the retailer will earn again π_R only after the end of the punishment phase of T periods.

In a similar way, let π_M be the present discounted value for the manufacturer:

$$\pi_M = (1-s)\left[-cq + D + \delta\pi_M\right] + s\left[-cq + \delta^{T+1}\pi_M\right]$$

With probability $1-s$, R repays so the relationship move on to the next period and with the complementary probability, R cannot repay and M terminates the contract for T periods.

The problem of M is to choose $\{q, D, T\}$ that maximizes π_M subject to the incentive compatibility constraint (4.1) and R's choice of q[16]:

$$\pi_M = \frac{(1-s)\,D - cq}{1-\delta\,(1-s) - s\delta^{T+1}}$$

s.t.

$$p(q)q - D + \delta\pi_R \geq p(q)q + \delta^{T+1}\pi_R \tag{4.1}$$

Condition (4.1) reflects the following retailer's trade-off: if he does not pay back, he keeps D but this will automatically trigger the termination phase, which yields valuation π_R only after T periods. For the retailer to pay back, the manufacturer needs to ensure that tomorrow's gains from not being terminated are larger than the payment today: $\delta\left(\pi_R - \delta^T\pi_R\right) \geq D$.

Note that an increase in D decreases the expected payoff of repaying directly and indirectly through a decrease in π_R (LHS of (4.1)). It also decreases the expected payoff of not repaying (RHS of (4.1)) but only indirectly and to less extent (i.e. discounted by δ^T). As a result, to minimize the rents given to R, M can choose D to make (4.1) bind, which allows us to rewrite it as follows:

$$D = \frac{1-\delta^T}{1-\delta^{T+1}}\delta\,(1-s)\,p(q)q$$

The tightness of this restriction depends on the discount factor δ, the probability of the shock s, as well as on the profitability of the market, $p(q)$. The incentive compatibility constraint imposes an upper bound on how much M can ask from R. In particular, the payment today should be smaller than what R can steal tomorrow if there is no shock, adjusted by the coefficient $f(T) = \frac{1-\delta^T}{1-\delta^{T+1}}$. Note that $f(T)$ is increasing in T and that at $T = 0$ it takes value 0. Therefore, (4.1) is not satisfied if the manufacturer never punishes. Similarly, the tougher the punishment, the easier it is for M to satisfy the constraint.

[16]Note that the participation constraint is always satisfied as in expectation R can ensure a minimum amount of rents, $(1-s)p(q)q$, by walking away with the generated revenues.

The manufacturer chooses q so that:

$$p'(q)q + p(q) = \frac{c}{\frac{1-\delta^T}{1-\delta^{T+1}}\delta(1-s)^2} \tag{4.2}$$

Therefore, the marginal costs (RHS) in (4.2) are larger than $\widetilde{c}$, and hence the resulting quantity is smaller than in the benchmark.

Proposition 4.1 *Despite the use of a quantity forcing contract, there exists double marginalization. Furthermore, the smaller the length of the punishment, the larger the double marginalization.*

The intuition is as follows: imagine that M were to offer the surplus maximizing q, and share the extra generated profits by charging a lower D. The reason why this is not optimal is because the resulting larger revenues increase the retailer's incentives to steal them. In order to curb these incentives, M needs to make the termination policy tougher but this is also costly for her in terms of lost future trade. Hence, the manufacturer prefers to distort the wholesale price at the cost of not maximizing the total profits in order to relax the incentive compatibility constraint. Thus, there is a trade-off between increasing the length of the punishment and distorting the wholesale price above the effective marginal cost.

When the manufacturer is deciding for how long to stop contracting with the retailer following a no repayment, she is facing the following trade-off: on one hand, by increasing the punishment length, she decreases the incentives of the retailer to steal, relaxing in this way the incentive compatibility constraint. On the other hand, she increases the inefficiency following a shock as she terminates profitable trading with the unlucky retailer.

Lemma 4.1 *The manufacturer terminates with the retailer forever if $s < \frac{(1-\delta)^2}{\delta^2}$. Otherwise, depending on the parameters, she may also terminate forever (for instance when $s \to 1$), or set a positive and finite T (for instance when $\delta \to 1$).*

Proof. See the Appendix. ■

Finally, it is worth highlighting that it is the impatience of the retailer the one that is creating the downward distortion in the quantity. Indeed, if the discount factor for R were different from the one for M, then it would be retailer's discount factor appearing in condition (4.2). In contrast, both discounts factors would affect the choice of the termination policy.

4.4 Example 2

This example extends the previous one by decreasing the severity of the shock; namely, the retailer faces a low demand $p_L(q)$ with probability s and a high demand $p_H(q)$ with probability $(1-s)$, where $p_L(q) \leq p_H(q) \forall q$. The manufacturer chooses the quantity q, together with a repayment D_H if demand is high and a repayment D_L if it is low. In order to give incentives to repay the true amount, M terminates the contract for T periods following D_L and forever following no payment[17]. Finally, let the expected revenues be $R(q) = [(1-s)\, p_H(q) + s p_L(q)]\, q$.

In the first best, the parties maximize their joint one-period profits, $R(q) - cq$, so the resulting optimal quantity is determined by:

$$q^{FB} : R'(q) = c$$

and since there is no asymmetric information, no termination policy is imposed.

As an example, assume that $p_H(q) = p(q)$ and $p_L(q) = lp(q)$, where $0 < l < 1$.[18] The optimal quantity is determined by:

$$p'(q)q + p(q) = \frac{c}{1 - s + ls} = \widetilde{c}$$

and hence, as in Example 1, the new effective marginal cost is larger. As the size of

[17] This is because M knows that, provided $p_L(q) > 0$ for the chosen q, R can repay some amount, so if he does not that means that he is stealing and hence inflicts the maximum punishment. It is optimal to impose the maximum punishment because it helps to relax the incentive compatiblity constraint and it is not imposed in equilibrium.

[18] We use this multiplicative functional form of the demand function very often in the paper because of its simplicity. This functional form is, for instance, used in Green and Porter (1984).

the shock, $1-l$, decreases; the downward quantity distortion decreases.

With the relational contract, and as we did in Example 1, we can define the present discounted value of the manufacturer:

$$\pi_M = \frac{(1-s)\,D_H + sD_L - cq}{1-\delta\,(1-s) - s\delta^{T+1}}$$

and the one of the retailer:

$$\pi_R = \frac{R(q) - [(1-s)\,D_H + sD_L]}{1-\delta\,(1-s) - s\delta^{T+1}}$$

With two state-demand uncertainty, M needs to ensure that the three incentive compatibility constraints are satisfied. The first one says that the retailer who faces a high demand does not want to pretend that he is facing a low one:

$$p_H(q)q - D_H + \delta\pi_R \geq p_H(q)q - D_L + \delta^{T+1}\pi_R \qquad (IC_H)$$

For this constraint to be satisfied, the profits in the high demand state plus tomorrow's value of staying in the relationship should be larger than the profits when making a D_L payment but being punished for T periods.

The second constraint says that a low demand retailer does not want to walk away with all the revenues:

$$p_L(q)q - D_L + \delta^{T+1}\pi_R \geq p_L(q)q \qquad (IC_L)$$

or, equivalently, the future expected value of being within the non-cooperative phase (rather than being terminated forever) should be larger than the low demand payment.

Finally, the last constraint says that a high demand retailer does not want to keep all the revenues:

$$p_H(q)q - D_H + \delta\pi_R \geq p_H(q)q \qquad (IC_H^{Global})$$

In other words, the future expected value of being within the cooperative phase (rather than being terminated forever) should be larger than the high demand payment.

Note that if IC_H and IC_L are satisfied, then IC_H^{Global} is also satisfied. Since an $H-type$ retailer never wants to walk away with all his revenues, we can ignore IC_H^{Global}.

In addition to the incentive compatibility constraints, the retailer's limited liability requires the payment done by R to be smaller than his revenues. That is, $p_L(q)q \geq D_L$ for the low state and $p_H(q)q \geq D_H$ for the high state, respectively. We call them LLs.

Using the above information, and since the manufacturer directly chooses q, M's problem becomes:

$$\max_{q,D_H,D_L,T} \pi_M = \frac{(1-s)\,D_H + sD_L - cq}{1-\delta\,(1-s) - s\delta^{T+1}}$$

s.t.

$$\delta\left(\pi_R - \delta^T \pi_R\right) \geq D_H - D_L \qquad (IC_H)$$

$$\delta^{T+1}\pi_R \geq D_L \qquad (IC_L)$$

$$p_H(q)q \geq D_H \qquad (LL_H)$$

$$p_L(q)q \geq D_L \qquad (LL_L)$$

Inspecting π_M, we see that M always wants to increase D_H and D_L as much as possible, therefore at least two constraints need to bind to bound D_H and D_L from above. The limited liability constraint in the high state, LL_H, is never binding because M has to give away rents in "the good times" to make the relationship valuable for R, so he has incentives to repay and stay within the relationship. Finally, the case where only IC_L and LL_L bind is not possible either. This is because an increase in the punishment period not only decreases π_M but also makes IC_L more binding. Therefore, M chooses no punishment ($T = 0$) and as a result IC_H is violated.

Therefore, we are left with three possible combinations of binding constraints: $IC_L - IC_H$, $IC_H - LL_L$ and $IC_L - IC_H - LL_L$. We deal with each of them in turn relegating tedious computations to the Appendix. After characterizing the three regimes, we construct a numerical example using the linear demand to determine under which parameter specification each of the regime occurs.

4.4.1 IC_H and IC_L bind

If the revenues obtained in the low demand state are nonetheless relatively large, LL_L will not be binding. Then IC_H and IC_L determine the payments in each state as a function of q and T: $D_H = \delta R(q)$ and $D_L = \delta^{T+1} R(q)$. The manufacturer asks to repay in the high state what R can steal in expected terms in the next period if they keep trading (i.e. tomorrow's discounted expected revenues) and in the low state what he can steal after the termination period ends.

Plugging D_H and D_L into π_M, M is left to maximize:

$$\max_{q,T} \pi_M = \frac{\left((1-s)\,\delta + s\delta^{T+1}\right) R(q) - cq}{1 - \delta\,(1-s) - s\delta^{T+1}}$$

Note that T always decreases the objective function, hence M never punishes following D_L: $T = 0$, and therefore she can only ask for the same payment in both states: $D_H = D_L = \delta R(q)$. The LL_L not binding implies $\frac{\delta(1-s)}{1-\delta s} p_H(q) < p_L(q)$ and a lower δ or a larger s makes it easier for this condition to be satisfied.[19] Finally, the optimal quantity is defined by: $R'(q) = \frac{c}{\delta}$. As in Example 1, there is a downward quantity distortion in the quantity with respect to the benchmark. In order to compare this distortion with the one in Example 1, let us assume that $p_H(q) = p(q)$ and $p_L(q) = lp(q)$, where $0 < l < 1$. Then the previous condition becomes:

$$q : p'(q)q + p(q) = \frac{c}{\delta\,(1 - s + ls)} > \widetilde{c} \tag{4.3}$$

The quantity distortion is lower that the one in Example 1 (condition (4.2)) because,

[19]Note that if LL_L is satisfied, LL_H is also satisfied because the last price is higher for any q.

since the repayment is the same for both states, the manufacturer only needs to prevent the retailer from stealing the revenues.[20]. Finally, a larger δ or a smaller s increases the quantity sold and hence the total repayment.

4.4.2 IC_H, IC_L and LL_L bind

When the difference between the high and low demand becomes too large, the limited liability constraint, LL_L, starts binding, which determines the payment in the low state as a function of the quantity: $D_L = p_L(q)q$. The incentive constraints determine the high demand payment and the punishment period also as a function of q: $D_H = \delta R(q)$ and $T = \frac{\ln\left(\frac{p_L(q)q}{R(q)}\right)}{\ln\delta} - 1$. Therefore, a difference in repayment, $D_L < D_H$, requires M to terminate with R following a low payment for a positive number of periods. In order to be in this case, the shock size should be large enough: $p_L(q) < \frac{\delta(1-s)}{1-\delta s} p_H(q)$.

Plugging the above information into M's objective function, M is left to maximize:

$$\max_q \pi_M = \frac{\delta\left(1-s\right)R(q) + sp_L(q)q - cq}{1-\delta\left(1-s\right) - \frac{sp_L(q)q}{R(q)}}$$

The first order condition can be found in the Appendix. If we use the multiplicative functional form again, then l should be small enough so the differences in demands are large: $l < \frac{\delta(1-s)}{1-\delta s}$ and the quantity is determined by:

$$q : p'(q)q + p(q) = \frac{c}{\delta\left(1-s\right)\left(1-s+sl\right)+sl} > \widetilde{c} \tag{4.4}$$

Since $0 < l < 1$ the quantity downward distortion is larger than in the case where only the incentive compatibility constraints bind. This is because the manufacturer, in addition to preventing an $L-type$ retailer from stealing the revenues, she also needs to prevent an $H-type$ retailer from pretending to be an $L-type$. For this particular example, the termination policy is: $T = \frac{\ln\left(\frac{l}{1-s+sl}\right)}{\ln\delta} - 1$ where $0 < T <$

[20] That is, $\frac{c}{\delta(1-s+ls)} < \frac{c}{\frac{1-\delta^T}{1-\delta^{T+1}}\delta(1-s)^2} \forall T$.

$+\infty$[21].

Finally, we can see from (4.4) that q increases following an increase in δ. Since the quantity increases, so does the repayment amounts D_H and D_L. To enforce these larger repayments, M increases the length of the termination period. When s increases, q decreases as well as both repayment amounts and T (see Appendix).[22]

4.4.3 IC_H and LL_L bind

When the demands are still different enough so that the limited liability constraint is still binding, but when the probability of the shock is low or the value of the future is large, then the incentive compatibility constraint for the $L-type$ stops binding, that is, the retailer does not want to walk away with the small revenues. To be precise, we are in this case when: $p_L(q) < \frac{\delta^{T+1}(1-s)}{1-s\delta^{T+1}} p_H(q)$, which requires $T < +\infty$ to be satisfied. As in the previous regime, the limited liability determines the low demand payment: $D_L = p_L(q)q$. The payment in the high demand state is determined by IC_H: $D_H = \frac{\delta\left(1-\delta^T\right)R(q)+(1-\delta)p_L(q)q}{1-\delta^{T+1}}$.

Using this information, M maximizes:

$$\max_{q,T} \pi_M = \frac{\frac{\delta(1-s)\left(1-\delta^T\right)R(q)+\left(1-\delta(1-s)-s\delta^{T+1}\right)p_L(q)q}{1-\delta^{T+1}} - cq}{1-\delta\left(1-s\right)-s\delta^{T+1}}$$

and again the first order conditions are relegated to the Appendix.[23] Using the multiplicative example, the high demand payment becomes:

$$D_H = \frac{\left(1-\delta^T\right)\delta\left(1-s+sl\right)+\left(1-\delta\right)l}{1-\delta^{T+1}} p(q)q$$

[21] T is positive because $\frac{l}{1-s+sl} < 1$ and $\frac{l}{1-s+sl} < \delta$ as l is constrained by $l < \frac{\delta(1-s)}{1-\delta s}$. Moreover, it is never infinity because $\delta < 1$.

[22] Note that LL_H is always satisfied as it requires:

$$\frac{\delta s}{1-\delta(1-s)} p_L(q) < p_H(q)$$

[23] Note that TC_H is always satisfied as it requires:

$$p_L(q) < p_H(q)$$

and that the quantity is determined by:

$$p'(q)q + p(q) = \frac{c}{\frac{\delta(1-s)^2\left(1-\delta^T\right)+\left(1-\delta(1-s)\left(1-s\left(1-\delta^T\right)\right)-s\delta^{T+1}\right)l}{1-\delta^{T+1}}} > \widetilde{c} \tag{4.5}$$

where q and T are the optimal choices. Finally, we are in this case when $l < \frac{\delta^{T+1}(1-s)}{1-s\delta^{T+1}}$.

Setting $l = 0$, it is easy to check that this solution corresponds to the one in Example 1. Since in Example 1 the low demand is equal to zero, M cannot ask for a repayment in this state and hence cannot distinguish between R keeping the money or being in the low state. This case is equivalent to Example 1 because the IC_L is not binding, that is, an $L-type$ retailer has no incentives to walk away with his revenues (so this possibility is irrelevant). Note that because of the existence of a positive low demand, the manufacturer can sell a larger quantity q and ask for a larger D_H.

4.4.4 Comparison of regimes

From the previous analysis, it emerges the same conclusion that we found in Example 1: the manufacturer, even though she has enough instruments to set the quantity, chooses to sell less than the efficient amount in each of the three previous regimes. However, this example is useful to shed light on the reasons for such a downward distortion of q.

Proposition 4.2 *As δ becomes arbitrary large, the outcome is bounded away from efficiency only if the limited liability constraint binds.*

Proof. See Appendix. ■

Indeed, as the firms become arbitrarily patient, the downward distortion of q remains only in the regimes where LL_L binds. Intuitively, when LL_L does not bind, M is better off asking for the same quantity regardless of the demand realization to avoid using a termination policy to separate the types, as this practice destroys surplus. When the retailer values more the future, the manufacturer can offer a larger

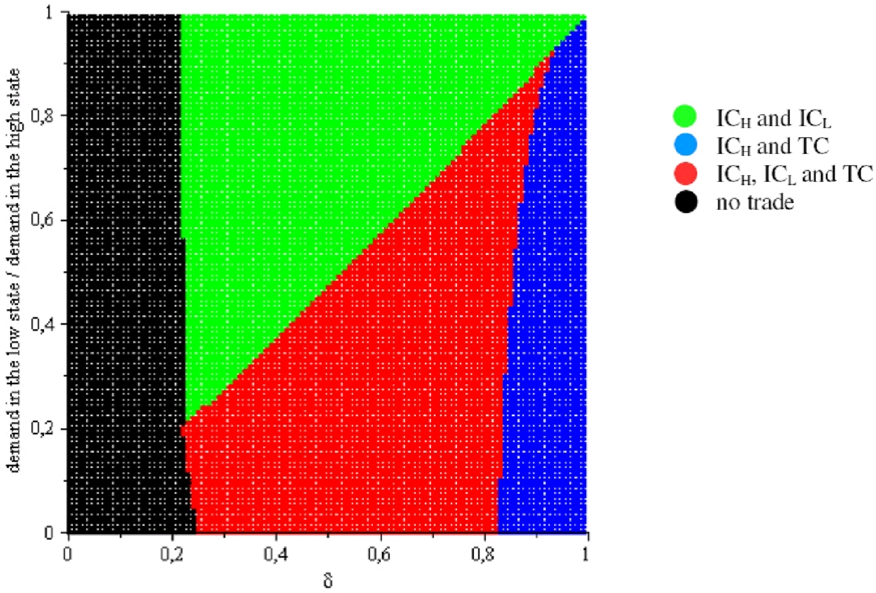

Figure 4.3: Optimal regime when $s = 0.1, a = 10, b = 1$ and $c = 2$

quantity and ask for a larger repayment because the retailer has more incentives to repay and stick with the relationship in the future. However, as the size of the shock increases, it becomes unprofitable to ask for such a low repayment regardless of the revenues. Hence, M is pushed to ask for a larger payment in "good times". Since the retailer is credit constrained and inter-temporal transfers of payoffs are impossible to use, termination occurs in equilibrium, which bounds q away from efficiency for no matter which δ.

Note that for any T, $\frac{\delta^{T+1}(1-s)}{1-s\delta^{T+1}} \leq \frac{\delta(1-s)}{1-\delta s}$, therefore from the previous sections we know that if $0 < l \leq \frac{\delta^{T+1}(1-s)}{1-s\delta^{T+1}}$ we are in case $IC_H + LL_L$, if $\frac{\delta^{T+1}(1-s)}{1-s\delta^{T+1}} < l \leq \frac{\delta(1-s)}{1-\delta s}$ we are in case $IC_H + IC_L + LL_L$ and if $\frac{\delta(1-s)}{1-\delta s} < l < 1$ we are in case $IC_H + IC_L$. Let us assume that the demand is linear and equal to: $p_H(q) = a - bq$ and $p_L(q) = l\,(a - bq)$, where $0 < l < 1$. Figures 4.3-4.5 depict which regime yields the largest profits for M in the space δ and l. The difference between figures is the probability of the shock s.

Figure 4.3 depicts the case where the likelihood of the shock is low and equal to $s = 0.1$. The line that separates the green and the red area is defined by $l = \frac{\delta(1-s)}{1-\delta s}$ and the line that separates the red and the blue area is $l = \frac{\delta^{T+1}(1-s)}{1-s\delta^{T+1}}$ where T is the

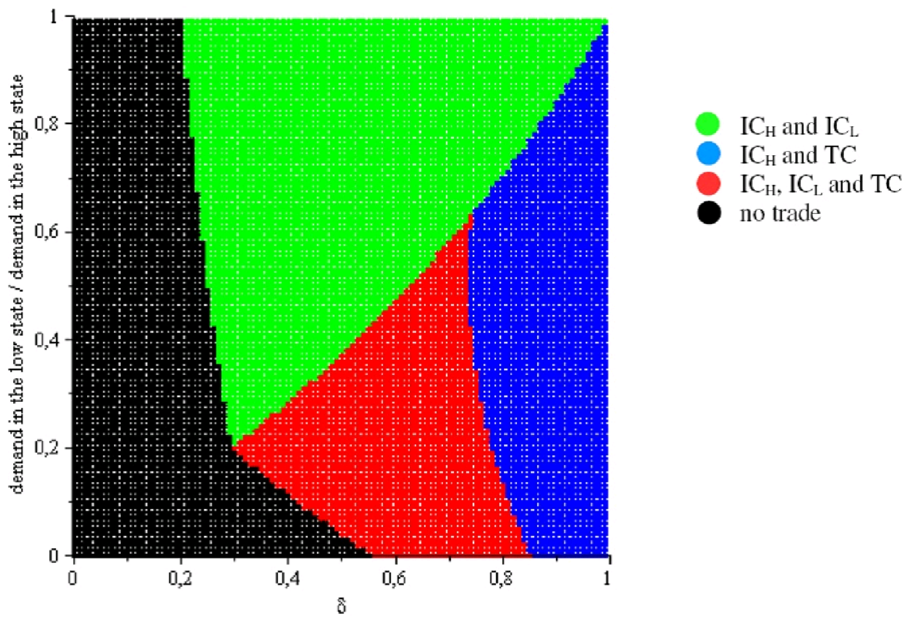

Figure 4.4: Optimal regime when $s = 0.4, a = 10, b = 1$ and $c = 2$

optimal termination policy. When the value of the future is small (i.e. δ is small) no contract can be implemented (and hence no trade occurs).[24] If the demands in both states are similar (i.e. l large and above $\frac{\delta(1-s)}{1-\delta s}$), then M asks for the same payment and hence does not need to use any termination policy. As the low demand becomes smaller (l decreases below $\frac{\delta(1-s)}{1-\delta s}$), then LL_L starts binding and M asks for a lower payment in the low state than in the high state ($D_L < D_H$) and punishes R following the low payment for a finite number of periods. Finally, as δ increases, the value of the relationship increases and IC_L stops binding (when $l < \frac{\delta^{T+1}(1-s)}{1-s\delta^{T+1}}$), that is, the $L - type$ retailer does not have incentives to walk away with the revenues.

Figure 4.4 and Figure 4.5 depict a shock of likelihood $s = 0.4$ and $s = 0.9$, respectively. Note that as s increases, less contracts can be self-enforced (for low l and δ). Since the shock is more likely and the manufacturer does not want to punish an unlucky retailer, the regime $IC_H + IC_L$, where the payment is the same for both states and there is no termination, becomes more common. Finally, when δ is high, M can ask for a larger payment from R (i.e. we are in the case $IC_H + LL_L$) thanks to the increased value of the future.

[24] Although this is out of the scope of this model, the no trade situation could also correspond to a manufacturer that vertically integrates downwards and serves the consumers.

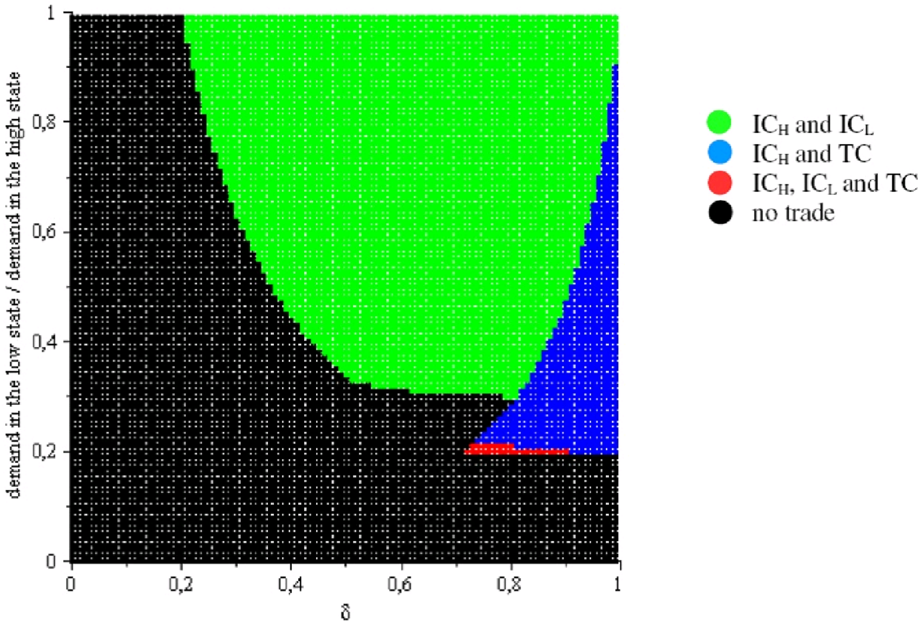

Figure 4.5: Optimal regime when $s = 0.9, a = 10, b = 1$ and $c = 2$

4.5 The Model

In this Section, we generalize the examples presented above by exploring the scenario where the retailer faces a continuum of demand states. In particular, the demand takes the following form: $p(q; s)$ where s is a random variable distributed on the interval $[0, \overline{s}]$. We denote by $h(s)$ and $H(s)$ the density and cumulative distribution functions, respectively. In contrast with the examples above, a larger s means a better state of demand for a given q (that is, $\frac{\partial p(q;s)}{\partial s} > 0$). As in the examples, M offers a quantity q, a repayment $D(\widetilde{s}, q)$ and a termination policy $T(\widetilde{s}, q)$, for each particular shock reported $\widetilde{s}$. Denote the revenues for a given state s as $R(q; s) = p(q; s)q$ and the expected revenues as $R(q) = \int_0^{\overline{s}} R(q; s)h(s)ds$.

4.5.1 Benchmark

The firms maximize their one-period joint profits:

$$\max_q R(q) - cq$$

The first best quantity is determined by:

$$q^{FB}:\int_0^{\overline{s}}\left[\frac{\partial p(q;s)}{\partial q}q+p(q;s)\right]h(s)ds=c$$

If the demand has the multiplicative functional form: $p(q;s)=sp(q)$, the optimal quantity is then given by

$$p'(q)q+p(q)=\frac{c}{\int_0^{\overline{s}}sh(s)ds}$$

and hence we observe a reduction in quantity as in Examples 1 and 2. Now, the effective marginal cost is $\widetilde{c}=\frac{c}{E(s)}$.

4.5.2 Relational contract

In this Section we consider the case of relational contracts. We first find the structure of an incentive compatible contract in this context and then we proceed to solve the problem of the manufacturer.

Incentive compatible contract

Since the quantity is chosen before the state of the demand is realized, it is not used for sorting purposes. Thus, the quantity is fixed at the time that the retailer is reporting how much he can repay. Then, for a given state of the demand s and a given quantity q, R chooses a report $\widetilde{s}$ to maximize his profits:

$$\pi_R(\widetilde{s};s)=\underbrace{R(q;s)-D(\widetilde{s},q)}_{\text{Today's payoff of reporting }\widetilde{s}}+\underbrace{\delta^{T(\widetilde{s},q)+1}\pi_R}_{\text{Expected payoff of reporting }\widetilde{s}}$$

where π_R is the expected discounted profits from staying in the relationship. Note that if the retailer were not credit constrained, the choice about which $\widetilde{s}$ to report would not depend on the true s as it is equally costly for any type of the retailer to report any $\widetilde{s}$. In other words, there is no sorting condition in this model, $\frac{\partial^2\pi_R(\widetilde{s};s)}{\partial\widetilde{s}\partial s}=0$. This feature makes this setup different from Levin (2003). Indeed, repaying is like the "burning money" story of advertising[25], where expending resources in advertising is equally costly for firms with different levels of efficiency. Because R is credit constrained, however, he cannot repay a larger $D(\widetilde{s}, q)$ than his actual revenues $R(q; s)$:

$$D(\widetilde{s}, q) \leq R(q; s) \quad \forall s, \widetilde{s} \tag{4.6}$$

The limited liability condition (4.6) links the choice of the report with the true state.

For a given q, the manufacturer has two instruments to deter R from under-reporting the state of the demand: following a low $\widetilde{s}$, she can either increase the payment today $D(\widetilde{s}, q)$ or increase the length of the termination period $T(\widetilde{s}, q)$, which decreases R's value of the future trade. Given that increasing $T(\widetilde{s}, q)$ is also costly for M (because she loses future trade), whenever it is possible, M will set the payment today as high as possible: $D(\widetilde{s}, q) = R(\widetilde{s}, q)$.

In equilibrium, however, M cannot extract all the revenues from R in all the states, as this would make the relationship worthless to the retailer (i.e., $\pi_R = 0$). There must be a report s^* for which M does not ask for all the revenues and, consequently, does not terminate, $T(s^*, q) = 0$[26]:

$$\pi_R(s^*; s) = R(q; s) - R(q; s^*) + \delta \pi_R \tag{4.7}$$

Hence for $\widetilde{s} \geq s^*$, M can only offer the contract: $T(\widetilde{s}, q) = 0$ and $D(\widetilde{s}, q) = R(q; s^*)$, because she can no longer decrease the punishment period to compensate R for a larger repayment. The following Lemma summarizes the structure of the contract:

Lemma 4.2 *The manufacturer offers the following repayment quantities:*

$$D(\widetilde{s}, q) = \begin{cases} R(q; \widetilde{s}) \text{ if } \widetilde{s} < s^* \\ R(q; s^*) \text{ if } \widetilde{s} \geq s^* \end{cases}$$

[25] See, for instance, Nelson (1974).

[26] If M were to terminate for some periods following the report s^*, she would strictly prefer to increase the repayment to decrease the termination length (and hence to reduce the inefficiency generated by the no-trade situation).

and length of contract termination:

$$T(\widetilde{s},q)=\begin{cases} T(\widetilde{s},q) \ \textit{if}\ \widetilde{s}<s^{*} \\ 0 \ \ \textit{if}\ \widetilde{s}\geq s^{*} \end{cases}$$

where $T(\widetilde{s},q)$ *is to be defined.*

Let us proceed to lay down the conditions under which this contract induces truth-telling. Since $\frac{\partial^2 \pi_R(\widetilde{s};s)}{\partial \widetilde{s}\partial s}=0$, let us denote by $u(\widetilde{s})$ the part of R's payoff that does not depend on his type s:

$$u(\widetilde{s})=-D(\widetilde{s},q)+\delta^{T(\widetilde{s},q)+1}\pi_R \tag{4.8}$$

The independence between the incentives to report a demand state and the actual demand state makes the task of inducing truth-telling quite simple. Intuitively, if it were feasible, the retailer would always report the demand state $\widetilde{s}$ with the highest $u(\widetilde{s})$, regardless of the true s. Therefore, to retailer will report the truth if the rents are not dependent on s.

Proposition 4.3 *To induce truth-telling from the retailer, the manufacturer gives him a non-negative rent that does not depend on the state of the demand (i.e., his type):* $u(\widetilde{s})\geq 0$ *and* $u'(\widetilde{s})\mid_{\widetilde{s}=s}=0\forall s$.

Proof. See Appendix. ■

Since the rent does not depend on the report, let us redefine $u(\widetilde{s})=u$. Using Proposition 4.3, equation (4.7) becomes $u=-R(q;s^{*})+\delta\pi_R$ when $u>0$. We can combine it with (4.8) to find:

$$\delta^{T(s,q)}=\frac{\delta\pi_R-R(q;s^{*})}{\delta\pi_R}+\frac{R(q;s)}{\delta\pi_R} \tag{4.9}$$

Equation (4.9) defines $T(\widetilde{s},q)$ from Lemma 4.2. Note that when R reports $s=0$, if $u>0$, then the M terminates for a finite number of periods, while if $u=0$, she terminates the contract forever. In the last case, equation (4.9) becomes $\delta^{T(s,q)} = \frac{R(q;s)}{\delta\pi_R}$. In both cases, a larger report is coupled with a shorter punishment period.

Note that, as we mentioned in Section 4.1.1, the contract offered by the manufacturer resembles to what it is known in the corporate finance literature as a debt contract. Debt contracts leave no money to the borrower in the bad states while making the borrower the residual claimant in the good states. The debt contract is optimal in this model because it minimizes the inefficiency associated with the termination (by trading-off larger repayments for lower termination periods in the default states) while inducing the retailer to repay.

Before solving the problem of the manufacturer, let us find π_R using Lemma 4.2:

$$\pi_R = \int_0^{s^*} \delta^{T(s,q)+1}\pi_R h(s)ds + \int_{s^*}^{\overline{s}} \left[R(q;s) - R(q;s^*) + \delta\pi_R\right] h(s)ds$$

Therefore, the retailer gives back all his revenues if the shock is smaller than s^* and has the contract terminated for $T(s,q)$ periods. Otherwise, he repays the constant amount $R(q;s^*)$ and keeps trading with the manufacturer in the next period. Using condition (4.9), π_R can be simplified to:[27]

$$\pi_R = \int_0^{\overline{s}} R(q;s)h(s)ds - R(q;s^*) + \delta\pi_R$$

Consider a first best world where the retailer is not credit constrained. Then an optimal contract could be a quantity (chosen so that the expected revenues minus the production costs are maximized) and a fixed up-front payment from the retailer. If the manufacturer has all the bargaining power, then this payment would be equal to the expected revenues (and hence the manufacturer would extract all the rents, $\pi_R = 0$) and the retailer would never have his contract terminated.[28] We can interpret the above equation in these terms, where the retailer keeps the expected

[27] If $u = 0$, then

$$\pi_R = \int_0^{\overline{s}} R(s,q)h(s)ds$$

[28] Note that an ex-ante equivalent possibility is for the manufacturer to ask for the obtained revenues in all the states.

profits minus a fixed payment, $R(q;s^*)$, and he never has his contract terminated. The difference from the first best world would come in terms of a different quantity and a smaller repayment amount in order to leave some rents to the retailer ($\pi_R > 0$) so it is worth for him to stay in relationship. Solving for π_R yields

$$\pi_R = \frac{R(q) - R(q;s^*)}{1-\delta} \tag{4.10}$$

Hence, for the expected discounted profits to be positive, it needs to be the case that the revenues at s^* are smaller than the expected revenues.

Finally, using (4.10), the incentive compatibility condition of Proposition 4.3 becomes:

$$u = \frac{\delta R(q) - R(q;s^*)}{1-\delta} \geq 0 \tag{4.11}$$

In order to ensure the repayment, the manufacturer needs to guarantee that the retailer obtains at least the difference between what he can steal tomorrow if he stays in the relationship and the maximum repayment that he can potentially do today.

The problem of the Manufacturer

Using Lemma 4.2, the profits of the manufacturer are:

$$\pi_M = \int_0^{s^*} \left[R(q;s) + \delta^{T(s,q)+1}\pi_M\right] h(s)ds + (1 - H(s^*))\left[R(q;s^*) + \delta\pi_M\right] - cq$$

that is, the manufacturer receives all the revenues in the bad states ($s < s^*$) and terminates the contract with the retailer for $T(s,q)$ periods. Otherwise, the manufacturer charges a fixed repayment, $R(q;s^*)$, and never terminates the contract. Finally, the manufacturer incurs the production costs.

When choosing s^*, M faces the following trade-off: a larger s^* allows M to extract a larger expected repayment from the retailer today (and hence π_R decreases); however, by (4.9), and in order to keep the incentives unchanged for all other types, this comes at a cost of a longer termination period.[29] On the other hand, an increase in the quantity q leads to an increase in the total payment as well as an increase in the production costs. Since it also leads to an increase of π_R, the effect on the termination policy is less clear and depends on the particular form of the demand function.

Let us plug (4.9) in the objective function and single out π_M:

$$\pi_M = \pi_R \frac{\int_0^{s^*} R(q;s)h(s)ds + (1 - H(s^*))\, R(q;s^*) - cq}{R(q) - \left(\int_0^{s^*} R(q;s)h(s)ds + (1 - H(s^*))\, R(q;s^*)\right)}$$

where π_R is defined in (4.10). The expected discounted profits of the manufacturer can be rewritten as the product of the retailer's earnings and the manufacturer's expected earnings today, divided by the loss in the surplus due to the relational contract.

The problem of the manufacturer is then to choose s^* and q to maximize π_M subject to (4.11). After replacing π_R, the problem of the manufacturer is[30]

$$\max_{q,s^*} \pi_M = \frac{\int_0^{s^*} R(q;s)h(s)ds + (1 - H(s^*))\, R(q;s^*) - cq}{(1-\delta)\frac{R(q) - \left(\int_0^{s^*} R(q;s)h(s)ds + (1-H(s^*))R(q;s^*)\right)}{R(q) - R(q;s^*)}} \tag{4.12}$$

$$\text{s.t.} \quad \frac{\delta R(q) - R(q;s^*)}{1-\delta} \geq 0$$

Note that as $\delta \to 1$, the constraint (4.11) does not bind. When this is the case, the optimal s^* and q do not depend on δ. Conversely, if (4.11) binds, both choices depend on δ. In particular, and in line with Example 1 and 2, the fixed repayment

[29] Indeed, after substituting for $\pi_R(s^*, q)$, the sign of the derivative is negative:

$$\frac{\partial \delta^{T(s,q)}}{\partial s^*} = -\frac{1-\delta}{\delta} \frac{\partial R(s^*, q)}{\partial s^*} \frac{R(q) - R(s,q)}{(R(q) - R(s^*, q))^2} < 0 \ \forall s \leq s^*$$

where the inequality follows from the fact that in order to have $\pi_R(s^*, q) > 0$, it needs to be the case that $R(q) > R(s^*, q)$.

[30] Note that the retailer's participation constraint is never binding as he always has the option of walking away with the current revenues $R(q;s)$ by reporting $\widetilde{s} = 0$ and hence not repaying anything. Thus, R receives at least the current revenues:

$$\pi_R(\widetilde{s}; s) = R(q;s) + \delta^{T(0,q)+1} \pi_R \geq 0 \forall s$$

that the manufacturer asks for the types above s^* will be equal to what the retailer expects to obtain tomorrow if he does not repay: $R(q;s^*) = \delta R(q)$. Therefore, a larger discount factor allows for a larger s^* that the manufacturer can impose on the retailer. Also, when (4.11) binds, an increase in s^* does not unambiguously increase the expected repayment because the fixed repayment for larger types $R(q;s^*)$ is constrained by $\delta R(q)$ and hence remains unchanged. Since the expected repayment remains unchanged, the termination policy does not change either. Thus, the only instrument available for the manufacturer is the choice of q. Note as well that the manufacturer will terminate with the retailer forever if there is no repayment at all.

The following Proposition summarizes this discussion and characterizes the optimal choices of the manufacturer.

Proposition 4.4 *The optimal s^* and q are determined by*

$$(1 - H(s^*))\,\pi_R - H(s^*)\pi_M \frac{E(s) - E\left(s \mid s \leq s^*\right)}{E(s) - s^*} = \widehat{\lambda} \tag{4.13a}$$

$$\pi_R \frac{\frac{\partial R(q;s^*)}{\partial q} - c}{\frac{\partial R(q;s^*)}{\partial q} - \delta\frac{\partial R(q)}{\partial q}} + \frac{\int_0^{s^*}\left(\frac{\partial R(q;s)}{\partial q} - \frac{\partial R(q;s^*)}{\partial q}\right)h(s)ds}{\frac{\partial R(q;s^*)}{\partial q} - \delta\frac{\partial R(q)}{\partial q}}\left(\pi_R + \pi_M\right) \tag{4.13b}$$
$$+\pi_M\left(\frac{\int_0^{s^*}\left(R(q;s^*) - R(q;s)\right)h(s)ds}{R(q) - R(q;s^*)}\,\frac{\frac{\partial R(q)}{\partial q} - \frac{\partial R(q;s^*)}{\partial q}}{\frac{\partial R(q;s^*)}{\partial q} - \delta\frac{\partial R(q)}{\partial q}}\right) = \widehat{\lambda}$$

where $\widehat{\lambda}$ is the shadow cost of the incentive compatibility constraint adjusted by the second best loss of revenues, $\widehat{\lambda} = \lambda\left(R(q) - \int_0^{s^*} R(q;s)h(s)ds - (1 - H(s^*))\,R(q;s^*)\right)$. If δ is large enough $\left(\frac{R(q;s^*)}{R(q)} < \delta\right)$, then $\widehat{\lambda} = 0$ and the optimal s^* and q are independent of δ.

Proof. See Appendix. ■

Since the solution depends on the particular demand function, in what follows, we illustrate these results with an example. For simplicity and comparability with the previous examples we use again the multiplicative demand function.

Example

Let us consider the case where $p(q;s) = sp(q)$. Then the retailer's profits are: $\pi_R = \frac{(E(s)-s^*)p(q)q}{1-\delta}$, and thus, in order for $\pi_R > 0$, the threshold s^* needs to be smaller than the expected shock. The problem of the manufacturer (4.12) becomes:

$$\begin{aligned} &\max_{q,s^*} \pi_M = \frac{\widehat{E}(s,s^*)p(q)q - cq}{(1-\delta)\,\frac{E(s)-\widehat{E}(s,s^*)}{E(s)-s^*}} \\ \text{s.t.} \quad &\frac{\delta E(s) - s^*}{1-\delta} \geq 0 \end{aligned} \tag{4.14}$$

where $\widehat{E}(s,s^*) = H(s^*)\, E(s \mid s \leq s^*) + (1 - H(s^*))\, s^*$ is the expected shock that matters to M in terms of the repayment. Indeed, for $s \leq s^*$ the manufacturer gets all the revenues repaid while for $s > s^*$ she gets the revenues of the state s^*. Clearly, for any $s^* < \overline{s}$, $\widehat{E}(s,s^*) < E(s)$, which leads to the downward distortion of q compared to the benchmark. The incentive compatibility constraint (4.11) becomes (4.14). It does not bind if s^* is smaller than the discounted value of the expected shock tomorrow, which happens for a large δ.

Figure 4.6 depicts the current and future profits of the manufacturer as a function of the realization of the shock for the case where the shock is distributed like a Uniform on $[0,1]$.

The payment today is the current revenues, $R(q;s)$, up to s^*, and a fixed payment, $R(q;s^*)$, from s^* onwards. The future discounted benefits from maintaining the relationship are constant and equal to $\delta\pi_M$ for s larger than s^*, and $\delta^{T(s,q)+1}\pi_M$ for s smaller than s^*.

Before we state the first order conditions, let us understand the trade-offs that the manufacturer is facing when choosing s^* and q. If the incentive compatibility constraint (4.14) binds, then s^* is determined by it so M is left to choose only q. In particular, M sells the quantity that maximizes her profits taking into account the repayment that she is expected to obtain and that it is determined by s^*.

When the constraint (4.14) does not bind, M can also choose s^*. A larger s^*, on

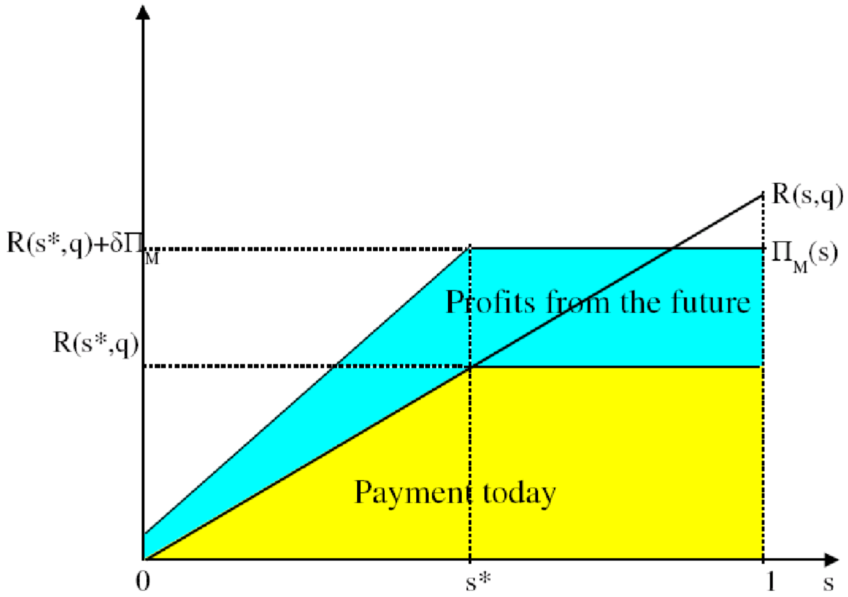

Figure 4.6: $p(s,q) = sp\,(q)\,,\quad s \rightsquigarrow U\,[0,1]$

one hand, increases the expected repayment[31] and hence π_M. On the other, because it also decreases π_R and R's incentives to repay, it increases the inefficiency due to tougher punishments[32].

Corollary 4.1 *If $p(q;s) = sp(q)$, then the optimal s^* and q are determined by the following first order conditions when $\delta E(s) > s^*$:*

$$q : p'(q)q + p(q) = \frac{c}{\widehat{E}(s,s^*)} \tag{4.15}$$

$$s^* : \frac{p(q)}{c} = \frac{H\,(s^*)}{1 - H(s^*)} \frac{E(s) - E\,(s \mid s \leq s^*)}{\widehat{E}(s,s^*)\,(E\,(s \mid s \geq s^*) - s^*) - E(s)\,(E(s) - s^*)} \tag{4.16}$$

Otherwise, $s^* = \delta E(s)$ and q is determined by (4.15). s^* and q are strategic complements.

Note that (4.15) and (4.16) do not contain δ. Thus if the value of the future is large, the constraint (4.14) does not bind, and the choice of q and s^* is not affected by δ. In this case, the manufacturer would set the first best quantity only

[31] Indeed, $\frac{\partial \widehat{E}(s,s^*)}{\partial s^*} > 0$.

[32] This is reflected in the denominator of π_M, which is increasing in s^* because $\frac{\partial \widehat{E}(s,s^*)}{\partial s^*} < 1$.

if $s^* = \overline{s}$[33], which is not possible as this makes the retailer's profits, π_R, zero. Therefore, Proposition 4.2 applies and the outcome is bounded away from efficiency even for δ arbitrarily close to 1. This is because in order to make the relationship profitable for the manufacturer, the limited liability constraint (4.6) binds for some states (i.e. $0 < s^*$[34]). This in turn implies that the termination is used in equilibrium and thus the surplus is bounded away from efficiency.

When the value of the future is low enough, $\delta < \frac{s^*}{E(s)}$, then the constraint (4.14) binds, in which case s^* and q do depend on δ. Then an efficient quantity will require $\delta > 1$[35], which is not possible either. Finally, note that, since $\frac{\partial \widehat{E}(s,s^*)}{\partial s^*} > 0$, the downward distortion in the quantity decreases with s^*, regardless of whether the constraint binds. This is because, as in the previous examples, a larger quantity needs to be associated with a tougher punishment and hence a larger s^*. Therefore, q and s^* are strategic complements.

Using Corollary and condition (4.9), it is possible to derive the termination policy implemented by the manufacturer. In particular, if (4.14) binds:

$$T(s,q) = \begin{cases} \frac{\ln\left(\frac{s}{E(s)}\right)}{\ln\delta} - 1 & \text{if } s < \delta E(s) \\ 0 & \text{if } \delta E(s) \leq s \end{cases}$$

If (4.14) does not bind:

$$T(s,q) = \begin{cases} \frac{\ln\left(\frac{\delta E(s) - s^* + (1-\delta)s}{E(s) - s^*}\right)}{\ln\delta} - 1 & \text{if } s < s^* \\ 0 & \text{if } s^* \leq s \end{cases}$$

where s^* is defined by (4.16).

Finally, in order to compare the general model with the numerical example in Section 4.4, we assume the linear demand function $p(q;s) = s\,(a - bq)$ and that s is distributed uniformly on the interval $[0,1]$[36]. For $a = 10$, $b = 1$ and $c = 2$,

[33] Because only then: $E(s) = \widehat{E}(s,s^*)$.

[34] Otherwise, M would produce no quantity because $\widehat{E}(s,s^*) = 0$.

[35] Indeed, $\delta = \frac{\overline{s}}{E(s)} > 1$.

[36] The first order conditions, (4.15) and (4.16), for this assumption in the distribution of the

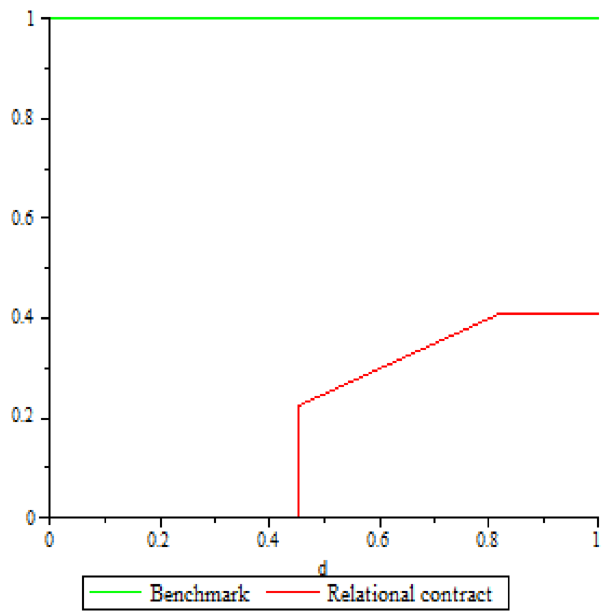

Figure 4.7: s^* when $a = 10$, $b = 1$ and $c = 2$

the optimal $s^* = 0.41$ and $q = 1.92$ if δ is at least 0.82. Otherwise, s^* and q are increasing in the discount factor. If $\delta < 0.45$, then there is no trade. Figures 4.7 and 4.8 depict s^* and q as a function of δ for this particular example.

4.6 Empirical implications

This analysis suggests that we might observe fewer contracts renewed, lower quantity sold and greater wholesale markups when demand uncertainty is higher. Demand uncertainty could potentially be proxied by macroeconomic volatility (across places). Another approximation to demand uncertainty could be the differences in the ease of verification. For instance, demand shocks could be easier to verify with non-perishable goods as compared to perishable ones. The distance between wholesale and retail markets could also be used to proxy the difference in the ease of verification as the upstream firm may have more information about the state of the demand the

shock are:

$$\begin{aligned} q \quad &: \quad p'(q)q + p(q) = \frac{c}{s^*\left(1 - \frac{s^*}{2}\right)} \\ \frac{p(q)}{c} \quad &= \quad \frac{s^*}{s^*\left(1 - \frac{s^*}{2}\right)(1 - s^*) - \left(\frac{1}{2} - s^*\right)} \end{aligned}$$

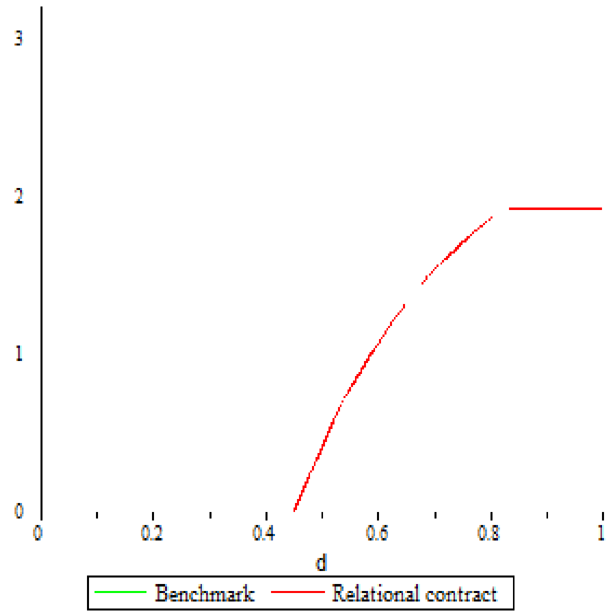

Figure 4.8: q when $a = 10$, $b = 1$ and $c = 2$

closer she is to the retail market. Finally, another possibility could be the presence or absence of ethnic or other links between wholesalers and retailers which would make the upstream or downstream firm more informed.

If we interpret the "no trade" situations of the model as the upstream firm directly selling to the consumer, then we expect to see more vertical integration when demand uncertainty is higher.

Finally, it would be interesting to test whether higher discount rates (proxied for instance, by larger interest rates) or greater revenues are associated with fixed repayments.

4.7 Conclusions

The goal of this paper has been to study the impact of trade credit and adverse selection on the vertical relational contract between an upstream and a downstream firm. Indeed, trade credit limits the available instruments that can be used within a relational contract and this affects the form of the emerging contract.

Contrary to Levin (2003), the upstream firm does not want to maximize the joint surplus, even though more surplus could be used to credibly punish and reward the downstream firm. This is because maximizing the joint surplus would give the downstream firm more incentives to steal which would require larger termination periods. Since the termination policy is also costly for the upstream firm (because of the resulting loss of future trade), the upstream firm prefers to distort the quantity downwards, away from the efficient quantity, in order to decrease the revenues held by the downstream firm. Because of this, we observe double marginalization despite the use of two-part tariffs. It is left for future work to determine in which way the contract would change if the upstream firm could not perfectly choose the final quantity sold in the market.

Not renewing the contract for a number of periods is costly for the upstream firm and this also has two clear implications in terms of the contract offered. First, if the revenues are large enough (above the state of the demand s^*), the upstream firm will never terminate the contract and ask for a fixed repayment. Otherwise, the upstream firm will ask for the highest possible repayment (equal to the generated revenues) and punish for a number of periods. Furthermore, the smaller the repayment, the larger the termination period. The choice of the threshold state s^* is linked to the choice of the quantity sold to the retailer. In particular, the larger this threshold, the smaller will be the downward distortion of the quantity because the larger is the share accrued by the upstream firm.

Finally, since the use of trade credit does not allow the firms to share the joint profits in an arbitrary way (i.e. using fixed transfers), the contract is expected to change depending on whether it is the upstream or downstream firm making the offer. It is left for future work to explore in which particular way the contract would change if it were offered by the downstream firm.

In our analysis, we have assumed that the quantity offered by the downstream firm does not change depending on the past repayment history. This framework would be suitable for industries where it is very costly to adjust the production quantity from one period to the other. In the future, it would be interesting to explore the form of non-stationary contracts and determine in which particular way the upstream firm will increase or decrease the quantity in each period as a function of the previous period repayment.

4.8 Appendix

Remark about non-stationarity. To get a sense about how restrictive it is the assumption of not letting the manufacturer offer a second opportunity to the retailer following an nonpayment, let us explore the following setup: imagine that the manufacturer starts selling the quantity q to the retailer and asks him to repay D_q. If the retailer does not pay back, then he gets a second chance where he is offered a quantity x and a repayment D_x. If after this second opportunity, the retailer repays, the manufacturer offers again $\{q, D_q\}$; otherwise, he is punished for T periods, after which the trading starts again with the contract $\{q, D_q\}$. Denote by π_{Rq} and π_{Rx} the present discounted value of the terms of trade $\{q, D_q\}$ and $\{x, D_x\}$, respectively. The retailer repays in the first period if:

$$(1-s)\left[p(q)q - D_q + \delta\pi_{Rq}\right] + s\delta\pi_{Rx} \geq (1-s)\,p(q)q + \delta\pi_{Rx} \qquad ((IC_1))$$

and in the second period if:

$$(1-s)\left[p(x)x - D_x + \delta\pi_{Rq}\right] + s\delta^{T+1}\pi_{Rq} \geq (1-s)\,p(x)x + \delta^{T+1}\pi_{Rq} \qquad ((IC_2))$$

Note that these constraints limit the repayment that M is expected to receive in each state respectively: $D_q \leq \delta\left(\pi_{Rq} - \pi_{Rx}\right)$ and $D_x \leq \delta\left(1-\delta^T\right)\pi_{Rq}$, where $\pi_{Rq} = (1-s)\,\frac{p(q)q - D_q + s\delta[p(x)x - D_x]}{1-\delta(1-s) - s\delta\left(\delta(1-s)+s\delta^{T+1}\right)}$ and $\pi_{Rx} = (1-s)\,\frac{(1-\delta(1-s))[p(x)x - D_x] + \left(\delta(1-s)+s\delta^{T+1}\right)[p(q)q - D_q]}{1-\delta(1-s) - s\delta\left(\delta(1-s)+s\delta^{T+1}\right)}$.

The respective present discounted values for the manufacturer are defined by:

$$\pi_{Mq} = (1-s)\left(D_q + \delta\pi_{Mq}\right) + s\delta\pi_{Mx} - cq$$

and:

$$\pi_{Mx} = (1-s)\left(D_x + \delta\pi_{Mq}\right) + s\delta^{T+1}\pi_{Mq} - cx$$

The problem of the manufacturer is then:

$$
\begin{aligned}
\underset{q,x,T}{Max}\ \pi_{Mq} &= \frac{(1-s)\,D_q - cq + s\delta\left[(1-s)\,D_x - cx\right]}{1-\delta(1-s) - s\delta\left[\delta(1-s)+s\delta^{T+1}\right]} \\
& s.t. \\
D_q &= \delta\left(\pi_{Rq}-\pi_{Rx}\right) && (IC_1) \\
D_x &= \delta\left(1-\delta^T\right)\pi_{Rq} && (IC_2)
\end{aligned}
$$

therefore introducing the second chance, decreases the incentives to repay in the first period (see IC_1), but reduces the inefficiency following the shock (see π_{Mq}). Plugging in IC_1 and IC_2 into the objective function, we obtain the following first order condition for x:

$$
\frac{\partial \pi_{Mq}}{\partial x} = -\frac{1-\delta - s\delta^2\left(1-\delta^T\right)}{A+BC}(1-s)^2\left[p'(x)x+p(x)\right] - sc = 0
$$

where $A = \delta^2(1-s)^2(1-\delta)\left(1-\delta^T\right)$, $B = 1-\delta(1-s) - s\delta^{T+2}$ and $C = 1 - \delta^2(1-s) - s\delta^{T+2}$. Note that $\frac{\partial \pi_{Mq}}{\partial x} < 0$ if $1-\delta-s\delta^2\left(1-\delta^T\right) > 0$, which is obviously true for $T=0$. Since the expression is decreasing in T, the lower bound is provided when $T=+\infty$. It easy to check that $\frac{\partial \pi_{Mq}}{\partial x} < 0$ when $\delta < 0.62$ and if $\delta \geq 0.62$, then when $\frac{1}{\delta}\left(\frac{1}{\delta}-1\right) > s$. In the shaded area of Figure 4.9 the manufacturer offers $x=0$. Note that this is for $T=0$ and that the shaded area will increase if one consider the optimal T. ■

Proof of Lemma 4.1. The first order condition for the length of the punishment is:

$$
\frac{\partial \pi_M}{\partial T} = \frac{-\delta^{T+1}\ln\delta}{\left(1-\delta+s\delta\left(1-\delta^T\right)\right)^2}\left[\frac{(1-\delta)^2 - s\delta^2\left(1-\delta^T\right)^2}{\left(1-\delta^{T+1}\right)^2}(1-s)^2 p(q) + sc\right] q
$$

This derivative is always positive at $T=0$. It tends to 0^+ as $T\to+\infty$ if $s < \frac{(1-\delta)^2}{\delta^2}$, therefore, in the white region of Figure 4.10 M terminates forever. If $s > \frac{(1-\delta)^2}{\delta^2}$, M

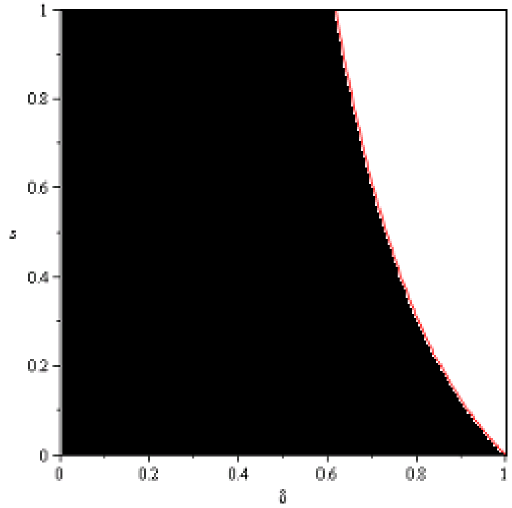

Figure 4.9: The shaded area is the minimum region in which the optimal $x = 0$

still terminates forever if:

$$\frac{c}{\delta\left(1-s\right)^2} > \left[\frac{\delta^2 s - \left(1-\delta\right)^2}{\delta s}\right] p(q)$$

Since by (4.2), $\frac{c}{\delta(1-s)^2} < p(q)$, and because $\frac{\delta^2 s-(1-\delta)^2}{\delta s} \in [0,1]$, for some parameter values, this inequality may be true (for instance when $s \to 1$). When the inequality does not hold (for instance when $\delta \to 1$), the optimal termination period is positive and finite. Since the second derivative (evaluated at $\frac{\partial \pi_M}{\partial T} = 0$) is negative, the solution is a maximum.

$$\frac{\partial^2 \pi_M}{\partial T^2} = -\frac{\delta^{2(T+1)}\left(\ln\delta\right)^2\left(1-s\right)^2\left(\left(1-\delta\right)^2+\delta s\left(1-\delta^T\right)\left(1-\delta+1-\delta^{T+1}\right)\right)}{\left(1-\delta+s\delta\left(1-\delta^T\right)\right)^2\left(1-\delta^{T+1}\right)^4} p(q)q$$

■

Operations for Example 2. When the three constraints bind, the first order condition for q is:

$$\frac{\partial \pi_M}{\partial q} = \frac{\left(\delta\left(1-s\right)R'(q)+s\left(p_L'(q)q+p_L(q)\right)-c\right)\left(R(q)\left(1-\delta\left(1-s\right)\right)-sp_L(q)q\right)}{R(q)\left(1-\delta\left(1-s\right)-s\frac{p_L(q)q}{R(q)}\right)^2}$$
$$+\frac{s\left(\delta\left(1-s\right)R(q)+sp_L(q)q-cq\right)\left(p_L'(q)q+p_L(q)-p_L(q)q\frac{R'(q)}{R(q)}\right)}{R(q)\left(1-\delta\left(1-s\right)-s\frac{p_L(q)q}{R(q)}\right)^2} = 0$$

in Ex1

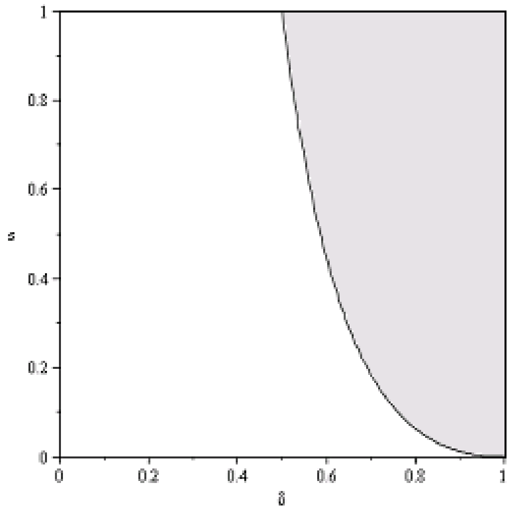

$_1$9.png

Figure 4.10: In the white region the optimal T is termination forever

substituting $p_L(q)$ by $lp(q)$ and $p_H(q)$ by $p(q)$ give us condition (4.4).The comparative statics of s are:

$$\frac{\partial^2 \pi_M}{\partial q \partial s} \propto -\frac{c\left(\delta\left[1-s+sl+(1-s)(1-l)\right]-l\right)}{\left(\delta(1-s)(1-s+sl)+sl\right)^2}$$

therefore, q decreases following an increase in s if:

$$\delta > \frac{l}{1-s+(1-s)(1-l)+sl}$$

which is always true because the RHS is increasing in l, and the inequality holds at the upper bound of l, $\frac{\delta(1-s)}{1-\delta s}$. By inspection, it is easy to see that $\frac{\partial D_H}{\partial s} < 0$, $\frac{\partial D_L}{\partial s} < 0$ and $\frac{\partial T}{\partial s} < 0$.

For the case where IC_H and LL_L bind, the first order condition for q is:

$$\frac{\partial \pi_M}{\partial q} = \frac{\delta(1-s)\left(1-\delta^T\right)R'(q)+\left(1-\delta(1-s)-s\delta^{T+1}\right)\left(p_L'(q)q+p_L(q)\right)}{1-\delta^{T+1}} - c$$

and for T is:

$$\frac{\partial \pi_M}{\partial T} = \frac{-\delta^{T+1}\ln\delta\left[\frac{\left[(1-\delta)^2-s\delta^2\left(1-\delta^T\right)^2\right](1-s)R(q)-\left(1-\delta(1-s)-s\delta^{T+1}\right)^2 p_L(q)q}{\left(1-\delta^{T+1}\right)^2}+scq\right]}{\left(1-\delta(1-s)-s\delta^{T+1}\right)^2}$$

Note that $\frac{\partial \pi_M}{\partial T}$ evaluated at $T = 0$ is:

$$\frac{\partial \pi_M}{\partial T} = \frac{-\delta \ln \delta \left[(1-s)\, R(q) - p_L(q)q + scq\right]}{(1-\delta)^2}$$

which is positive if the expression in brackets is positive. Plugging in the upper bound of $p_L(q)$ given by IC_L, we obtain: $\frac{1-s-\delta}{1-s\delta}(1-s)p_H(q) + sc > 0$. Thus a sufficient condition for $T > 0$ is $1 - s > \delta$. ■

Proof of Proposition 4.2. By inspection, the RHS of (4.3) and (4.4) are larger than $\widetilde{c}$. Using l'Hopital's rule, the RHS of (4.5) tends to $\frac{(1-s)^2 T + (1-s+Ts(1-s)+(T+1)s)l}{T+1}$ as $\delta \to 1$. Noting that it is increasing in T and that it is larger than $\widetilde{c}$ at $T = 0$ completes the proof. ■

Proof of Proposition 4.3. Because of condition (4.6), a retailer can lie only downwards. To avoid the retailer underreporting his type, M needs to ensure that $u'(\widetilde{s}) \geq 0$, that is, the payoff $u(\widetilde{s})$ increases with reported $\widetilde{s}$. However, M cannot design a contract that satisfies $u'(\widetilde{s}) > 0$ for $s \in [s^*, \overline{s}]$ because she has no means to increase $u(\widetilde{s})$ (i.e., she is already not terminating with the retailer and asking for a smaller repayment $R(q; s\prime) < R(q; s^*)$ would make types $s \in [s', s^*]$ overreport their types). Furthermore, for $s \in [0, s^*)$ the contract cannot satisfy $u'(\widetilde{s}) > 0$ either, because M is already extracting the maximum payment and increasing the payoff with $\widetilde{s}$ would imply a costly (for M) increase in the termination period for smaller types. Indeed, M wants to set $T(s, q)$ as low as possible as long as there is truth-telling. Finally, to prevent R walking away with the earnings, M need to ensure that $u(\widetilde{s})$ is non-negative, which completes the proof. ■

Proof of Proposition 4.4. The Lagrangian function is:

$$\mathcal{L} = \frac{\int_0^{s^*} R(q;s)h(s)ds + (1 - H(s^*))\, R(q;s^*) - cq}{(1-\delta)\left(1 + \frac{H(s^*)R(q;s^*) - \int_0^{s^*} R(q;s)h(s)ds}{R(q) - R(q;s^*)}\right)} - \lambda\left(R(q;s^*) - \delta R(q)\right)$$

where λ is the shadow cost associated with the constraint. The first order conditions are:

$$
\begin{aligned}
\frac{\partial \mathcal{L}}{\partial s^*} &= \pi_R \frac{\partial R(q;s^*)}{\partial s^*} \frac{1-H(s^*)}{R(q)-\int_0^{s^*} R(q;s)h(s)ds-(1-H(s^*))R(q;s^*)} \\
&\quad -\pi_M \frac{\partial R(q;s^*)}{\partial s^*} \frac{\frac{H(s^*)R(q)-\int_0^{s^*} R(q;s)h(s)ds}{R(q)-R(q;s^*)}}{R(q)-\int_0^{s^*} R(q;s)h(s)ds-(1-H(s^*))R(q;s^*)} \\
&\quad -\lambda \frac{\partial R(q;s^*)}{\partial s^*} = 0 \\
\frac{\partial \mathcal{L}}{\partial q} &= \pi_R \frac{\int_0^{s^*} \frac{\partial R(q;s)}{\partial q} h(s)ds + (1-H(s^*)) \frac{\partial R(q;s^*)}{\partial q} - c}{R(q)-\int_0^{s^*} R(q;s)h(s)ds-(1-H(s^*))R(q;s^*)} \\
&\quad -\pi_M \frac{H(s^*)\frac{\partial R(q;s^*)}{\partial q} - \int_0^{s^*} \frac{\partial R(q;s)}{\partial q} h(s)ds - \frac{\left(H(s^*)R(q;s^*)-\int_0^{s^*} R(q;s)h(s)ds\right)\left(\frac{\partial R(q)}{\partial q}-\frac{\partial R(q;s^*)}{\partial q}\right)}{R(q)-R(q;s^*)}}{R(q)-\int_0^{s^*} R(q;s)h(s)ds-(1-H(s^*))R(q;s^*)} \\
&\quad -\lambda \left(\frac{\partial R(q;s^*)}{\partial q} - \delta \frac{\partial R(q)}{\partial q} \right)
\end{aligned}
$$

and

$$\lambda \left(R(q;s^*) - \delta R(q) \right) = 0$$

Setting the first two conditions equal to zero, give us equation (4.13a) and (4.13b), respectively.

We assume that the second order conditions hold. Note that the optimal s^* cannot be in a corner as this would give zero profits to either the manufacturer (if $s^* = 0$) or to the retailer (if $s^* = 1$). We assume that the demand and the cost function are such that production takes place, i.e., $0 < q < +\infty$. ■

Proof of Corollary 4.1. Suppose that the constraint does not bind, the first order conditions are:

$$\frac{\partial \pi_M}{\partial q} = \frac{\widehat{E}(s,s^*)\left(p'(q)q+p(q)\right)-c}{(1-\delta)\frac{E(s)-\widehat{E}(s,s^*)}{E(s)-s^*}} = 0$$

$$
\begin{aligned}
\frac{\partial \pi_M}{\partial s^*} &= (1-H(s^*)) \frac{E(s)-s^*}{E(s)-\widehat{E}(s,s^*)} p(q)q \\
&\quad + H(s^*) \frac{E(s \mid s \le s^*) - E(s)}{\left(E(s)-\widehat{E}(s,s^*)\right)^2} \left(\widehat{E}(s,s^*)p(q)q - cq\right) = 0
\end{aligned}
$$

The cross derivative is:

$$\begin{aligned}\frac{\partial^2 \pi_M}{\partial s^* \partial q} &= (1 - H(s^*))\,\frac{E(s) - s^*}{E(s) - \widehat{E}(s, s^*)}\,(p'(q)q + p(q)) \\ &\quad + H\left(s^*\right)\frac{E\left(s \mid s \leq s^*\right) - E(s)}{\left(E(s) - \widehat{E}(s, s^*)\right)^2}\left(\widehat{E}(s, s^*)\left(p'(q)q + p(q)\right) - c\right) = 0\end{aligned}$$

Using $\frac{\partial \pi_M}{\partial q} = 0$, it is easy to see that $\frac{\partial^2 \pi_M}{\partial s^* \partial q} > 0$. ■

Bibliography

[] ABREU, D., PEARCE, D., and STACCHETTI, E., (1986), "Optimal Cartel Equilibria With Imperfect Monitoring", *Journal of Economic Theory*, Vol. 39, pp. 251-269.

[] BAKER, G., R. GIBBONS, and MURPHY, K. J., (1994) "Subjective Performance Measures in Optimal Incentive Contracts," *The Quarterly Journal of Economics*, Vol.109(4), pp. 1125-1156.

[] BOLTON, P. and SCHARFSTEIN, D.S., (1990), "A Theory of Predation Based on Agency Problems in Financial Contracting." *American Economic Review*, Vol. 80, pp. 93-106.

[] BUEHLER, S., and GARTNER D.L., (forthcoming), "Making Sense of Non-Binding Retail-Price Recommendations", *American Economic Review.*

[] CUNAT, V. M., (2007), "Trade Credit: Suppliers as Debt Collectors and Insurance Providers", *Review of Financial Studies*, Vol. 20, pp. 491-527.

[] FAFCHAMPS, M., (1997), "Trade credit in Zimbabwean manufacturing", *World Development*, Vol. 24 No. 3, pp. 795-815.

[] FAFCHAMPS, M., (2000), "Ethnicity and credit in African manufacturing", *Journal of Development Economics*, Vol. 61, pp. 205-235.

[] FAFCHAMPS, M., MCKENZIE, D., QUINN, S., and WOODRUFF, C., (forthcoming), "Using PDA consistency checks to increase the precision of profits and sales measurement in panels," *Journal of Development Economics.*

[] FAURE-GRIMAUD, A., (2000), "Product Market Competition and Optimal Debt Contracts: The Limited Liability Effect Revisited." *European Economic Review*, Vol. 44, pp.1823-1840.

[] FONG, Y-F. and Li, J., (2010) “Relational Contracts, Efficiency Wages, and Employment Dynamics”, mimeo.

[] GALE, D. and HELLWIG, M., (1985), "Incentive-Compatible Debt Contracts: The One-Period Problem." *Review of Economic Studies*, Vol. 52, pp. 647-663.

[] GREEN, E.J., and PORTER R.H., (1982), "Noncooperative collusion under imperfect price information", *Econometrica*, Vol. 52, No. 1 (Jan., 1984), pp. 87-100.

[] HART, O. and MOORE, J., (1998), "Default and Renegotiation: A Dynamic Model of Debt." *Quarterly Journal of Economics*, Vol. 113, pp.1-41.

[] INNES, R.D., (1990), "Limited Liability and Incentive Contracting with Ex-ante Action Choices." *Journal of Economic Theory*, Vol. 52, pp. 45-67.

[] JOHNSON, S., MCMILLAN, J., and WOODRUFF, C., (2002), "Courts and Relational Contracts", Journal of Law, Economics and Organization, Vol. 18, pp. 221-77.

[] LEVIN, J., (2003), "Relational Incentive Contracts", *American Economic Review*, Vol. 93, No. 3, pp. 835-857(23).

[] MAILATH, G., and SAMUELSON, L., (2006): *Repeated Games and Reputations: Long-Run Relationships*. New York: Oxford University Press.

[] MALCOMSOM, J., (2010): "Relational incentive contracts", Chapter for Handbook of Organizational Economics (edited by Robert Gibbons and John Roberts), to be published by Princeton University Press

[] MCMILLAN, J. and WOODRUFF, C., (1999a), "Interfirm Relationships and Informal Credit in Vietnam", *Quarterly Journal of Economics*, Vol.114, pp. 1285-1320.

[] MCMILLAN, J. and WOODRUFF, C., (1999b), "Dispute Prevention Without Courts in Vietnam," *Journal of Law, Economics, and Organization*, No.15, pp. 637-58.

[] NELSON, P., (1974), "Advertising as information," *Journal of Political Economy*, Vol. 81, pp. 729-754.

[] POVEL, P. and RAITH, M., (2004a), "Optimal Debt with Unobservable Investments." *RAND Journal of Economics*, Vol. 35, pp. 599-616.

[] POVEL, P. and RAITH, M., (2004b), "Financial Constraints and Product Market Competition: Ex Ante vs. Ex Post Incentives." *International Journal of Industrial Organization*, Vol. 22, pp. 917-949.

[] SMITH, J., (1987), "Trade Credit and Informational Asymmetry", *Journal of Finance*, Vol. 42, pp. 863-69.

[] TOWNSEND, R.M., (1979), "Optimal Contracts and Competitive Markets with Costly State Verification." *Journal of Economic Theory*, Vol. 21, pp. 265-293.

[] WAKED, D., (2010), "Antitrust Enforcement in Developing Countries: Reasons for Enforcement & Non-Enforcement using Resource-Based Evidence", mimeo.