

University of Warwick institutional repository: <http://go.warwick.ac.uk/wrap>

A Thesis Submitted for the Degree of PhD at the University of Warwick

<http://go.warwick.ac.uk/wrap/56794>

This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it. Our policy information is available from the repository home page.

Library Declaration and Deposit Agreement

1. STUDENT DETAILS

Please complete the following:

Full name: Luciana Nicollier

University ID number: 0752471

2. THESIS DEPOSIT

2.1 I understand that under my registration at the University, I am required to deposit my thesis with the University in BOTH hard copy and in digital format. The digital version should normally be saved as a single pdf file.

2.2 The hard copy will be housed in the University Library. The digital version will be deposited in the University's Institutional Repository (WRAP). Unless otherwise indicated (see 2.3 below) this will be made openly accessible on the Internet and will be supplied to the British Library to be made available online via its Electronic Theses Online Service (EThOS) service.

[At present, theses submitted for a Master's degree by Research (MA, MSc, LLM, MS or MMedSci) are not being deposited in WRAP and not being made available via EThOS. This may change in future.]

2.3 In exceptional circumstances, the Chair of the Board of Graduate Studies may grant permission for an embargo to be placed on public access to the hard copy thesis for a limited period. It is also possible to apply separately for an embargo on the digital version. (Further information is available in the *Guide to Examinations for Higher Degrees by Research*.)

2.4 If you are depositing a thesis for a Master's degree by Research, please complete section (a) below. For all other research degrees, please complete both sections (a) and (b) below:

(a) Hard Copy

I hereby deposit a hard copy of my thesis in the University Library to be made publicly available to readers (please delete as appropriate) ~~EITHER~~ immediately ~~OR after an embargo period of~~ months/years as agreed by the Chair of the Board of Graduate Studies.

I agree that my thesis may be photocopied. YES / ~~NO~~ (Please delete as appropriate)

(b) Digital Copy

I hereby deposit a digital copy of my thesis to be held in WRAP and made available via EThOS.

Please choose one of the following options:

EITHER My thesis can be made publicly available online. YES / ~~NO~~ (Please delete as appropriate)

OR My thesis can be made publicly available only after.....[date] (Please give date)
YES / NO (Please delete as appropriate)

OR My full thesis cannot be made publicly available online but I am submitting a separately identified additional, abridged version that can be made available online.
YES / NO (Please delete as appropriate)

OR My thesis cannot be made publicly available online. ~~YES~~ / NO (Please delete as appropriate)

3. GRANTING OF NON-EXCLUSIVE RIGHTS

Whether I deposit my Work personally or through an assistant or other agent, I agree to the following:

Rights granted to the University of Warwick and the British Library and the user of the thesis through this agreement are non-exclusive. I retain all rights in the thesis in its present version or future versions. I agree that the institutional repository administrators and the British Library or their agents may, without changing content, digitise and migrate the thesis to any medium or format for the purpose of future preservation and accessibility.

4. DECLARATIONS

(a) I DECLARE THAT:

- I am the author and owner of the copyright in the thesis and/or I have the authority of the authors and owners of the copyright in the thesis to make this agreement. Reproduction of any part of this thesis for teaching or in academic or other forms of publication is subject to the normal limitations on the use of copyrighted materials and to the proper and full acknowledgement of its source.
- The digital version of the thesis I am supplying is the same version as the final, hard-bound copy submitted in completion of my degree, once any minor corrections have been completed.
- I have exercised reasonable care to ensure that the thesis is original, and does not to the best of my knowledge break any UK law or other Intellectual Property Right, or contain any confidential material.
- I understand that, through the medium of the Internet, files will be available to automated agents, and may be searched and copied by, for example, text mining and plagiarism detection software.

(b) IF I HAVE AGREED (in Section 2 above) TO MAKE MY THESIS PUBLICLY AVAILABLE DIGITALLY, I ALSO DECLARE THAT:

- I grant the University of Warwick and the British Library a licence to make available on the Internet the thesis in digitised format through the Institutional Repository and through the British Library via the EThOS service.
- If my thesis does include any substantial subsidiary material owned by third-party copyright holders, I have sought and obtained permission to include it in any version of my thesis available in digital format and that this permission encompasses the rights that I have granted to the University of Warwick and to the British Library.

5. LEGAL INFRINGEMENTS

I understand that neither the University of Warwick nor the British Library have any obligation to take legal action on behalf of myself, or other rights holders, in the event of infringement of intellectual property rights, breach of contract or of any other right, in the thesis.

Please sign this agreement and return it to the Graduate School Office when you submit your thesis.

Student's signature: Luciana Nicollier Date: March 13th, 2013

**Essays on Industrial Organisation:
The Role of Consumers' Generated
Information**

by

Luciana A. Nicollier

Department of Economics

University of Warwick

September 2012

A thesis submitted in partial fulfilment of
the requirements for the degree of
Doctor of Philosophy in Economics

Contents

Acknowledgments	6
Declaration	7
Abstract	8
1 Introduction	1
1.1 Modelling Complaints and Reviews	4
1.2 Free Riding and the Group-Utilitarian Assumption	6
1.3 Related Literature	11
2 Customers' Complaints and Quality Regulation	15
2.1 Introduction	15
2.2 The Model	23
2.2.1 The Consumers	25
2.2.2 The Firm's Investment Decision	28
2.2.3 Consumers' Beliefs and Expectations	31
2.3 Equilibrium	32
2.3.1 Informativeness of Complaints	37
2.3.2 Comparative Static of the Optimal Complaining Rule	40
2.4 The Repeated Game	41
2.4.1 Equilibrium	47
2.5 Conclusions	53

3	Prices, Reviews and Endogenous Information Transmission	55
3.1	Introduction	55
3.2	Basic Setup	61
3.2.1	Benchmark case: No Reviews	63
3.2.2	Reviews	64
3.2.3	Updating: Public Beliefs	67
3.2.4	Role of the reviewing rule	69
3.3	Information Transmission	70
3.4	Cost of completing a review	76
3.5	Equilibrium: Reviews and Price Discounts	80
3.5.1	Firm's pricing strategy	80
3.5.2	Optimal Reviewing Rule and Equilibrium	82
3.6	Conclusions and Further Research	86
A	Appendix to Chapter 2	89
A.1	Expected Proportion of Complaints	89
A.2	Bayesian Updating	91
A.3	Equilibria of the Repeated Game when $T \rightarrow \infty$	92
B	Appendix to Chapter 3	101
B.1	Proofs	101
B.2	Monotone Likelihood Ratio Property	110

List of Figures

2.1	Relationship between σ^* and q	29
2.2	Changes in the Optimal Cutoff as the Parameters Change	42
3.1	Reviewing Rule	84

*A mis padres y a mi amor,
que estuvieron conmigo en cada paso del camino.*

Acknowledgments

I would like to thank economic support from the British Council and the Warwick Economic Research Fellowship. I am indebted to my advisors, A. Carvajal and M. Waterson, as well as to R. Akerlof, C. Mezzetti, D. Sgroi, and A. Dhillon for inspiration advice and useful comments. For the first chapter of this thesis, I would also like to thanks participants of the Warwick Economics Lunchtime Seminar, the CRESSE 2012 Conference and the EARIE 2012 Conference.

Declaration

I submit this thesis to the University of Warwick in accordance with the requirements for the degree of Doctor of Philosophy. I declare that this thesis is my own work and has not been submitted for a degree at another university.

Luciana A. Nicollier

September 2012

Abstract

A variety of economic agents rely on information generated by the consumers when making their decisions. Not only consumers' rely on other consumers' experiences when making their buying decisions, but also some governmental agencies rely on customers' complaints to make inferences about the functioning of some markets. Little is known, however, about how this information interacts with the firms' investing and pricing decisions. A common denominator of the various types of information generated by the consumers is that its content depends on consumers' incentives to transmit information, which are not always obvious and may vary across markets and time. This thesis studies the role of the information generated by the consumers in two different contexts. The first chapter studies whether customers' complaints about the quality provided by a regulated monopolist are informative about the firm's investment decisions. The second chapter considers the pricing decision of a monopoly firm when the consumers' buying decision is based on the reviews completed by previous consumers.

The main contributions are twofold. First, by endogenising consumers decision to lodge a complaint or complete a review, I am able to derive conclusions about the informational content of consumers behaviour and about its strategic interaction with the firms decisions. Second, the thesis makes a methodological contribution because it proposes a novel way of dealing with the free riding problem that lies at the very root of the generation of information by consumers.

Chapter 1

Introduction

A variety of economic agents rely on information generated by the consumers when making their decisions. For instance, consumers' rely on other consumers' experiences when making their buying decisions,¹ and some governmental agencies rely on customers' complaints to make inferences about the quality of a service.² However, little is known about how this information interacts with the firms' investing and pricing decisions. A common denominator of the various types of information generated by the consumers is that its content depends on consumers' incentives to transmit information, which are not always obvious and may vary across markets and along time. As a result, the effect of customers' complaints and reviews on the firms' behaviour is not obvious, as neither is their impact on aggregate welfare.

This thesis explores the relationship between the information generated by the consumers and the behaviour of a monopoly firm. The first chapter studies the potential of customer complaints as a regulatory tool to induce higher invest-

¹In fact, empirical evidence suggest that consumers' reviews are an important determinant of the firms' revenues -see Luca (2011), Doyle and Waterson (2012), and Chevalier and Mayzlin (2006), among others.

²According to the European Union's report "Monitoring Consumer Markets in the European Union" (2011), the complaints measurement is considered "a key metric to evaluate the functioning of a market" (page 12).

ment in quality. The second chapter considers a monopolist's pricing strategy when consumers' information about the quality of an experience good is the one contained in the reviews completed by previous buyers. In both cases, a consumer's decision to transmit information is endogenous. As a result, they derive conclusions about the informational content of consumers' behaviour and about its strategic interaction with the firm's decisions. A common message of the two chapters is that a note of care should be taken when considering consumers' complaints and reviews. The results suggest that the widespread belief that the mere existence of consumer-generated information increases consumers' welfare may be true but it needs further qualifications.

Chapter 2, studies whether customers' complaints are informative about the investment decision of a monopoly firm in contexts in which quality is not verifiable and consumers cannot (fully) appropriate the benefits of their complaints. It proposes a psychological game in which a firm decides whether to make a costly investment that increases quality in a first order stochastic dominance sense, and the consumers decide whether to complain or not. The consumers do not observe the firm's investment. Their complaining decision is driven by the difference between the level of quality they were expecting to receive and the one they actually received.³ It is shown that the presence of this (rational) reference point may induce a positive correlation between the observed proportion of complaints and the firm's investment.

Chapter 3 studies a dynamic game in which a long-lived monopoly faces a sequence of short-lived consumers. Neither the firm nor the consumers have ex-ante information about the value of an experience good, but they learn from the reviews completed by previous buyers. If he buys, the consumer observes a

³Empirical evidence suggests that this is indeed the case. See Forbes (2008), for example.

quality realisation and may complete a review. The consumers complete reviews according to a social rule that maximises the present value of current and future consumers utility. The paper shows that the reviews induce a mean preserving spread on the posterior beliefs about the value of the good which, combined with the convexity of the utility and the profit functions, implies the reviews are valuable for both the consumers and the firm. Hence, both parties are willing to face some cost in order to increase the information available in the market. From the firm's perspective, this cost takes the form of a discount in the price offered to current consumers. As this discount has the additional effect of compensating consumers for the cost of completing reviews, it induces a reviewing rule that is more favourable to the firm (in the sense that it increases the firms expected profits). It is further shown that a necessary condition for the existence of reviews is that the firm cannot fully appropriate the surplus generated by this increased information.

The thesis also makes a methodological contribution because it proposes a novel way of dealing with the well-known free riding problem that lies at the very root of the generation of information by consumers. A common feature of the two types of consumers' generated information I study is the presence of (at least some level of) free riding incentives. When the consumer makes a complaint or when he completes a review he is essentially taking a costly action, the benefits of which he cannot (fully) appropriate.⁴ Hence, the complaining and the reviewing decisions have some similarities with an agent's decision to participate in a large election. Therefore, to tackle this difficulty, I borrow from the voting literature and I assume that consumers are *group-utilitarians*, i.e., they receive a positive payoff for acting according to a strategy that maximises their group's aggregate

⁴For example, empirical studies on eBay show that most of the times the customer is not likely to buy again from the same seller, implying that he does not receive a direct benefit from completing a review. Yet, Resnick and Zeckhauser (2002) report that 52.1% of the buyers on eBay actually provide voluntary feedback about their sellers.

utility.

The rest of this Introduction is organised as follows. The next section briefly summarises the models of chapters 2 and 3 and presents their main results. Section 1.2 discusses the role of the utilitarian assumption as a solution to the free riding problem. Finally, Section 1.3 revises the existing literature and discusses the contributions of the models proposed in this thesis.

1.1 Modelling Complaints and Reviews

Chapter 2 proposes a theoretical model to analyse the informativeness of customers' complaints when the benefits of those complaints can not be fully appropriated. A regulated monopoly decides whether to make a costly investment that increases quality. Neither the consumer nor the regulator observe the firm's investment, but consumers observe a realisation of quality that is related to investment in a first order stochastic dominance sense. After observing quality, consumers decide whether to complain or not. If a high proportion of consumers complain, the regulator punishes the firm.

The model in this chapter is a psychological game between a monopoly firm and the consumers. After observing a quality realisation, the consumers decide whether to complain by comparing the quality they received with the one they were expecting. Consumers' reference point is determined by their rational expectations.⁵ In this way, the model captures the idea that "disappointment" and "poor performance" are endogenously defined and depend on the context. The presence of a reference point in consumers' complaining decisions implies that the

⁵Forbes (2008) assumes that consumers form an unbiased expectation of the quality they will receive. With rational expectations, her empirical results imply that an increase in quality decreases the (expected) proportion of complaints only when the higher quality was not anticipated by consumers. The same is true in the model of this paper.

payoff functions of both the consumers and the firm depend not only on what they do but also on what consumers were *expecting* from the firm, which reflects the psychological aspect of the game.⁶

The main result is that complaints are not always informative about the firm's investment behaviour. Indeed, a firm might be punished despite of investment levels being high if consumers expected high quality or, on the contrary, not being punished when investing is low if consumers expected low quality. Furthermore, this lack of informativeness can be worsened by a repeated interaction between the firm and the consumers.

The final chapter of this thesis studies a firm's pricing strategy when buyers can complete reviews about the value of the product. It considers a situation in which a long lived monopoly faces a sequence of short lived consumers whose only information about the value of the product is the one contained in the reviews completed by previous buyers. After the firm chooses a product's price, the consumers decide whether to buy or not. If they buy, they may complete a review. Before buying, neither the consumers nor the firm have private information about the good's value, so the price and the previous consumers' buying decisions are not informative. However, after buying consumers observe a quality realisation that is correlated with the actual value of the product, and thus they are better informed than the firm and the future consumers. Consumers may decide to transmit this information by completing reviews.

Consumers complete reviews according to a social rule that maximises the present value of current and future consumers' utility. It is shown that customers' reviews induce a mean preserving spread on the beliefs about the value of the

⁶See, for example, Geanakoplos, Pearce, and Stacchetti (1989) and Battigalli and Dufwenberg (2009).

good. Combined with the convexity of the utility and the profit functions, the increased variability of the posterior beliefs results in the information contained in the reviews being valuable for both, the consumers and the firm. Hence, both parties are willing to face some cost in order to increase the information available in the market. From the firm’s perspective, this cost takes the form of a “discount” in the price offered to current consumers. By reducing the current price, the firm increases current (expected) demand which in turn increases the probability with which the current consumer completes reviews. As this discount has the additional effect of compensating consumers for the cost of completing reviews, it also induces a reviewing rule that is more favourable to the firm (in the sense that it increases the expected future profits in the scenario with reviews).

1.2 Free Riding and the Group-Utilitarian Assumption

The two papers that form this thesis deal with the free riding problem by assuming that consumers are *group utilitarians*. The notion of *utilitarian agents* has been proposed by Harsanyi to explain the so-called “paradox of no voting”: if voting is costly then, since the likelihood of a vote being pivotal is very small, standard game-theoretic models predict low levels of turn out (Downs 1957).

Harsanyi’s notion of group utilitarians was later formalised by Feddersen and Sandroni (2006b). This assumption is useful in the contexts I study because it constitutes a plausible explanation for consumers’ behaviour in settings in which tangible benefits accrue only if aggregate participation is high, and no one can be excluded from the benefits of group success. The application of the group utilitarian assumption to models of complaints and reviews is one of the contributions

of this thesis and it constitutes, to the best of my knowledge, the first application of this idea outside the area of Political Economy. Therefore, it is worth looking deeply into its meaning and formal implications. The next subsection defines and explains this assumption within the voting model developed by Feddersen and Sandroni (2006a, 2006b) and Coate and Conlin (2004). Then, I briefly discuss how the utilitarian assumption is used in the models of complaints and reviews proposed in this thesis.

The Utilitarian Assumption

The utilitarian assumption has been proposed as a solution to the free riding problem in the voting context. The starting point is the work by Harsanyi (1977, 1980, 1992). Harsanyi argues that voting may usefully be understood as individuals acting according to the dictates of rule-utilitarianism, and proposes a game theoretic model in which people receives a payoff for acting “ethically”. He illustrates his argument with a situation in which a fixed number of votes is needed to pass a policy that would raise aggregate utility. Each citizen faces the same cost of voting and chooses a probability of voting that, if adopted by all, would maximise aggregate utility. The key insight is that the optimal probability is between zero and one. Not everybody should stay at home, because that would mean the policy would not pass. But not everybody should vote because that would result in a surfeit of votes, imposing unnecessary costs on society. In this way, the logic of rule-utilitarianism yields an elegant theory of turnout. Harsanyi assumes that everyone does their duty, but rejects the implicit assumption that doing one’s duty always involves voting.

Harsanyi’s insight has been formalised by Feddersen and Sandroni (2006a, 2006b).⁷ They model a large election with two candidates, in which voting costs

⁷Feddersen and Sandroni (2006a) develops the conceptual and operational foundations for

vary within the population and a single vote is never pivotal. Agents have preferences over the candidates and the cost of the election. There are two types of agents: those who prefer candidate 1 and those who prefer candidate 2. Fixing the probability of winning for each candidate, all agents prefer to minimise the cost of the election. Given a preference type, a *rule* defines a cut-off point such that agents with voting costs below this threshold should vote for their favoured candidate and those with voting costs above the threshold should abstain. They assume that some agents, called *ethicals*, receive a payoff for acting according to the rule. A solution concept called “consistency” links agents’ preferences with actual behaviour in a way analogous to a Nash Equilibrium. The optimal voting rule constitutes a Nash Equilibrium of a game in which the supporters of the two parties aim to maximise the probability that their preferred candidate wins the election, net of the social cost of voting. The consistency requirement adds to the equilibrium concept the participation constraint that each agent’s payoff from ethical behaviour exceeds his voting cost.⁸ In this way, Feddersen and Sandroni’s (2006b) model shows that costly voting and strategic considerations may coexist in a formal model.⁹

Coate and Conlin (2004) apply a version of the utilitarian model to a referendum. The key difference between their model and Feddersen and Sandroni’s (2006b) is that in the former individuals follow the voting rule that maximises the

the ethical voting model used in Feddersen and Sandroni (2006b) and Coate and Conlin (2004).

⁸Feddersen and Sandroni (2006a) show that ethical voting models share a common mathematical structure with elite driven turnout models. However, the micro foundations for both types of models differ significantly. While in the ethical voting model each agent acts independently on the basis of their own assessment of what constitutes ethical action, in the elite driven models agents are provided direct incentives by the leader’s instructions. In the models of complaints and reviews, the second interpretation could mean, for example, that consumers follow the directions of some sort of Consumers’ Association.

⁹Their model predicts high levels of turn out and comparative statics that are consistent with strategic behaviour. It further delivers testable implications and predicts variations in expected turnout and margin of victory as a function of various parameters of the model, like the costs to vote, the level of disagreement within the electorate, and the importance of the election, for example.

payoffs of those on their side of the issue, while in the latter they follow the rule that (they believe) maximises aggregate utility.^{10 11} They structurally estimate a parameterised version of their group utilitarian model using data on Texas liquor referenda. The results of the empirical estimations are broadly consistent with the comparative static predictions of the model.

Apart from the applications to voting and to the complain and reviews examples of this thesis, Harsanyi’s (1980) type arguments have been proposed as a possible explanation to household’s response to conservation appeals during the California’s energy crisis in 2000 and 2001. Reiss and White (2008) find empirical evidence that consumers do respond to voluntary appeals provided the costs of a collective action failure are tangible and that the public is well aware of it. Even though their work is purely empirical, the utilitarian-type argument seems a plausible explanation for those results.¹²

Group-Utilitarians, Complaints and Reviews

In the models of Chapters 2 and 3 I assume that consumers are *group-utilitarians* in the sense of Coate and Conlin (2004).¹³ In both models, the “group” is broadly defined as “the consumers”. In Chapter 2 consumers are long lived players and so “the consumers” is the group formed by the firm’s customers. In Chapter 3 consumers are short lived and thus the firm’s potential customers

¹⁰Therefore, Coate and Conlin (2004) define agents in their model as “group utilitarians”.

¹¹Another difference between the two models is that Coate and Conlin (2004) allow the two groups to differ in the intensity of their preferences for their preferred alternative.

¹²In this case, each household faces private cost of reducing consumption, a virtually zero possibility of bringing about any tangible benefit with respect to the crisis through individual effort, and a considerable incentive to free-ride on whatever efforts are made by others.

¹³As a simplifying assumption, it is assumed in both chapters that all the consumers receive a positive payoff for acting “ethically” and so all of them could potentially lodge a complaint or complete a review. This is a minor difference with the models in Feddersen and Sandroni (2006a, 2006b), that assume that only a fraction of the population receives such a payoff. Introducing this possibility in my models would not modify the results in any significative way.

change every period; the “group” in this case is defined as all the consumers, current and future.¹⁴

In both chapters consumers are ex-ante homogeneous but ex-post heterogeneous in some dimension. In the model of Chapter 2, heterogeneity is introduced by assuming that the costs of lodging a complaint vary across agents. Therefore, a complaining rule consists in a cut off cost of complaining below which a consumer lodges a complain. This threshold maximises consumers’ aggregate utility given the regulatory rule, consumers’ prior expectations and the quality they received. In the chapter about reviews, the cost of completing a review is assumed to be the constant across consumers, but the quality observed by consumers if they buy may vary. In this case, the rule specifies which review the consumer should complete (if any) after every possible quality realisation, given the firm’s strategy.¹⁵

The assumption of group utilitarian consumers has different formal implications for the complaints and the reviews models, even though in both cases it is assumed that consumers aim to maximise the present value of consumers’ expected utility. In the first case, the specific regulatory rule I am studying implies that current complaints are a sunk cost when the firm decides its future investment level. As a result, consumers only complain to punish the firm’s current “poor performance” -i.e., they behave *as if* they were myopic. Together with the utilitarian assumption, myopic behaviour results in a game that is strategically equivalent to a game between a long lived firm and a sequence of short lived consumers. In the chapter about reviews, since consumers’ aim to maximise the

¹⁴The assumption of short lived consumers is a way of modelling the fact that consumers can not learn the value of the product from their own experience.

¹⁵Introducing heterogeneity in the costs of completing a review would not affect the results significantly. It would imply a two-dimensions rule, with one dimension related to the observed quality and another one consisting in a cutoff cost below which a consumer that observed a certain quality realisation completes a review.

present value of their group's utility, the game is strategically equivalent to a game between two long lived players.

1.3 Related Literature

This thesis relates to several strands of literature. The first strand emanates from the large literature that studies how economic agents learn from the actions of others. This includes models of worth-of-mouth communication¹⁶ and models of herding and cascades.¹⁷ Within the first group, the transmission of information is generally modelled by assuming that in every period new consumers meet (or sample) an exogenous proportion of old consumers, who tell them their experience with the product.¹⁸ In the herding models, on the other hand, the transmission of information is modelled as an externality: an agent's payoff depends on his own decision and the state of nature, but not on the actions of others. Agents take account of others' actions only because of the information revealed by them. Most of this literature aims to study the long run outcome of different variations of those processes of information transmission. What differentiates the approach in this thesis from all those papers is that I consider situations in which the agents explicitly decide whether to transmit information and which information to transmit. By making those decisions endogenous, I can study the other players' best response to the information generated by the consumers and, hence, I can get a more complete approach to the market effects of the consumers' generated

¹⁶Like the ones in Ellison and Fudenberg (1995), Banerjee and Fudenberg (2004), and Ahn and Souminen (2005), for example.

¹⁷Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992) and Smith and Sorensen (2000), Bose, Orosel, Ottaviani, and Vesterlund (2006), Bose, Orosel, Ottaviani, and Vesterlund (2008), among others.

¹⁸For example, Ellison and Fudenberg (1995) study the way in which worth-of-mouth aggregates information of individual agents. By assuming that each player hears from the current experience of a random sample of N other players (where N is an exogenous parameter), they show that the structure of the worth-of-mouth process affects the tendency of a population to display conformity or diversity.

information.

Of the many applications within the herding literature, Bose, Orosel, Ottaviani, and Vesterlund (2006) seems to be the closest one to the model presented in Chapter 3. This paper studies the pricing strategy of a monopoly firm when consumers obtain information about the value of the good from the buying decision of others. The main difference between that model and the one I present below is that in my case the transmission of information is not an externality but the consumers make an explicit decision about how much and which information to transmit.

Chapter 2 of this thesis also contributes to the extensive literature on quality provision by a monopoly firm. Starting by the seminal papers of Spence (1975) and Shesinski (1982), the literature suggests that an unregulated monopoly will over or under supply quality according to whether the marginal consumer values additional quality more or less highly than do the infra-marginal consumers on average.¹⁹ It has further been shown that regulation of service prices can compound, ameliorate or otherwise complicate the already existing market failure (see for example, Spence (1975), Mussa and Rosen (1978) and Besanko, Donnenfeld, and White (1987); Sappington (2005) surveys the literature).²⁰ The main conclusion of much of the existing literature is that when quality is not verifiable the regulator needs to have a great deal of information before even know

¹⁹The difference depends on whether quantity and quality are seen as substitutes or as complements by consumers. In the former case, consumers willingness to pay a higher price for an increase in quality decreases with the quantity that he buys (i.e., the demand curve becomes less elastic as quality increases), while in the latter the elasticity of the demand increases with quality. Thus, the monopoly is more likely to undersupply quality if quality and quantity are substitutes, and to oversupply it if they are complements.

²⁰Price ceilings that are independent of the firm's realised costs limits its incentives to supply quality, because they prevent the firm from capturing any of the incremental consumer's surpluses that would result from the higher service quality. However, as noted by Laffont and Tirole (1993), even under pure cost-of-service regulation, the regulated firm does not gain from providing costly services either, so a low perceived cost of supplying quality does not imply a high incentive to provide quality.

in which direction he should intervene. By studying the informational content of complaints, this paper considers whether the “police power” could be moved from the regulatory agency to the consumers.

This chapter is also related to the literature on reference dependent utility²¹ and with some models in marketing research on customers satisfaction.²² In both cases, it is suggested that consumers utility depends not only on the actual product quality that was received but also on whether that quality was above or below some reference level. This paper adds to the first branch of literature because, in spite of being widely accepted, the effect of that reference point on consumers’ complaining decision and on the firm’s incentives to invest have not been studied yet. It differs from the second branch in that they do not require consumers’ expectations to be rational and, as a result, they are not able to make clear predictions about the firm’s strategic response to consumers’ complaints.

Chapter 3 is also related with several branches of literature. Economic and management literature has shown an increasing interest on different aspects of users’ reviews and online behaviour. Existing research focuses mainly on the relationship between customers’ reviews and the firm’s revenues or sales. For example, Luca (2011) uses data from Yelp.com to study the impact of customers’ reviews on restaurant revenues; he shows that higher ranking implies higher revenues. Similarly, Chevalier and Mayzlin (2006) shows that the differences between customers’ reviews posted on Barnes and Noble and those posted on Amazon are positively related to the differences in book sales on the two websites. The model proposed in Chapter 3 not only explains those empirical findings, but it also delivers implications in terms of the firm’s pricing behaviour that results from the

²¹Kahneman and Tversky (2001), Koszegi and Rabin (2006), among others.

²²See for example, Singh (1988), Zithaml, Berry, and Parasuraman (1996), Boulding, Kalra, Staelin, and Zeithaml (1993), Oliver (1977), Oliver (1980).

existence of information transmission through the reviews.

Finally, the model in Chapter 3 is related to the literature on public tests. See, for example, Gill and SgROI (2008, 2012), Lerner and Tirole (2006). Similar to this literature, the reviewing model in chapter 3 studies whether the firm would choose to test its product publicly. The reviewing model I study differs from the above mentioned ones in that the “toughness” of the test chosen by the firm in the existing literature but the firm affects it only indirectly in the my model.

Chapter 2

Customers' Complaints and Quality Regulation

2.1 Introduction

Customers' complaints constitute an important source of consumer-generated information. For instance, the European Union's (2011) report "Monitoring Consumer Markets in the European Union" states that customers' complaints constitute "a key metric to evaluate the functioning of a market" (page 12); as a result, complaints are one of the elements the study takes into consideration in order to derive conclusions about the market's performance. Despite its relevance and generalised use for policy purposes, very little is known about the informational content of customers' complaints. The conventional wisdom about the role of complaints in a market is that the smaller the number of customers that complain, the better the market performs -i.e., complaints are informative about the distortions existing in a market. However, its theoretical foundations are not clear.

Consider, for example, a firm's decision to invest in improving its customer

service department (hereafter, “quality”). Even though the consumers do not observe the firm’s investment, it is likely that they are better informed than the regulator about the quality of the service. Then, everything else constant, more complaints may be indicative of lower quality. However, there is an important caveat because consumers’ incentives to complain may vary across markets and along time for at least two reasons. First, suppose, as the empirical evidence suggests, that complaints are driven by expectations as well as by actual quality and so, that they depend (at least partially) on consumers’ “disappointment” with the quality they received (Forbes 2008).¹ As a result, the definition of what constitutes an “appropriate service level” is likely to change with the context; a higher number of complaints may then be the result not of lower quality but of higher expectations. Second, complaining is a costly action the benefits of which cannot be always fully appropriated. For instance, if future investment increases with current complaints but all the consumers benefit from the resulting higher quality, then each individual consumer would prefer others to face the cost of complaining. Hence, a smaller number of complaints may reflect a significant degree of free riding incentives and not a higher quality.

This paper studies the informativeness of customers complaints about a firm’s investment and their potential as a regulatory tool. The starting point is the assumption that, as suggested by the empirical evidence, customers’ complaints may be the result of either low quality or high expectations. As in the customer service example, the paper considers some contexts in which quality is not verifi-

¹The relationship between complaints and “disappointment” or “dissatisfaction” seems to be generally accepted in Marketing Literature -see Oliver (1977), Singh (1988) and Boulding, Kalra, Staelin, and Zeithaml (1993), among others. It also seems to be an accepted relationship among regulatory agencies. For example, OFGEM defines complaints as “any *expression of dissatisfaction* made to an organisation, related to any one or more of its products, its services or the manner in which it has dealt with any such expression of dissatisfaction” OFGEM (2008)

able² and consumers cannot (fully) appropriate the benefits of their complaints.³ It is shown that, while a regulation based on complaints may induce a higher investment, those complaints are not systematically informative about the firm's behaviour. As a result, the firm may be punished more frequently when it invests than when it does not. Furthermore, the lack of informativeness may be worsen by a repeated interaction between the firm and consumers, because it creates incentives for the firm to try to "keep expectations low".

The paper also identifies conditions under which complaints may help overcome the regulator's lack of information about the firm's investment behaviour. In particular, the results challenge the conventional wisdom that the easier it is for consumers to complain the more information is contained in these complaints. When the cost of lodging a complaint is zero, the amount of complaints becomes independent of the quality received by the consumers and so they convey no information about the firm's investment. Finally, the paper delivers comparative static results on consumers' complaining decisions that explain Forbes's (2008) empirical findings -namely, that the number of complaints decreases with actual quality and that, after controlling for actual quality, consumers complain more often when they would have expected to receive higher quality.

The paper proposes a psychological game between a monopoly firm and the consumers. A regulated monopoly decides whether to make a costly investment that increases quality. The consumers do not observe the firm's investment, but

²Quality is verifiable when it can be (costlessly) described ex ante in a contract and ascertained ex post by a court (Laffont and Tirole 1993). When quality is verifiable, the regulator can reward or punish the firm directly as a function of the level of quality. It can, for example, dictate the heating value of gas or punish an electric utility on the basis of the number and intensity of outages (Laffont and Tirole 1993). On the contrary, when quality is not verifiable it is not possible to write contracts contingent on outcomes.

³When the consumer expects to receive a direct benefit out of his complaint (like monetary compensations because of electricity shortcuts or reimbursements of incorrectly high bills), his complaining decision can be perfectly explained using standard microeconomic theory: the consumer lodges a complaint as long as the (expected) cost is below the (expected) benefit.

they observe a realisation of quality that is related to investment in a first order stochastic dominance sense. After observing quality, consumers decide whether to complain by comparing the realised quality with the one they were expecting to receive. If a high proportion of consumers complains, the firm is fined. Consumers' reference point is determined by their rational expectations.⁴ In this way, the model captures the idea that “disappointment” and “poor performance” are endogenously defined and depend on the context. The presence of a reference point in consumers' complaining decisions implies that the payoff functions of both the consumers and the firm depend not only on what they do but also on what consumers were *expecting* from the firm, which reflects the psychological aspect of the game.⁵

The combination of a reference point with a fine that depends on the number of complaints implies that “disappointed” consumers consider lodging a complaint only if by complaining they increase the probability that the firm is “punished” for its “poor performance”. However, complaining is a costly action the benefits of which the consumer cannot (fully) appropriate. Even more, since I make the simplifying assumption that there is a continuum of consumers, the model suffers from an extreme version of free riding and so, without additional assumptions, there would be no complaints in equilibrium. A similar result holds when studying consumers' incentives to participate in a large election. To tackle this difficulty I borrow from the voting literature the assumption that a fraction of the consumers are *group-utilitarians*, i.e., they receive a positive payoff for acting according to a strategy that maximises consumers' aggregate utility. Given their disappointment, consumers have preferences about the probability with which

⁴Forbes (2008) assumes that consumers form an unbiased expectation of the quality they will receive. With rational expectations, her empirical results imply that an increase in quality decreases the (expected) proportion of complaints only when the higher quality was not anticipated by consumers. The same is true in the model of this paper.

⁵See, for example, Geanakoplos, Pearce, and Stacchetti (1989) and Battigalli and Dufwenberg (2009).

the firm should be punished and the cost of complaining. These preferences are not identical across consumers because complaining costs are heterogeneous. The complaining rule maximises their (expected) aggregate utility, given their disappointment and the regulatory rule. This rule consists of a cut off cost of complaining below which a consumer lodges a complaint.⁶

An equilibrium of the complaining game satisfies three requirements.⁷ First, the firm chooses the investment level that maximises its expected profits given its beliefs about the consumers' strategy and expected quality. Second, consumers choose their complaining strategy optimally given their disappointment with the quality they received and their payoff for following the complaining rule. And third, the firm correctly anticipates consumers' expected quality, which is in turn consistent with the firm's strategy and the consumers' prior beliefs. Using this definition, the one-shot version of the model has two different equilibria: a "high quality equilibrium" in which consumers expect the firm to invest and the firm optimally invests, and a "low quality equilibrium" in which consumers do not expect the firm to invest and the firm optimally fulfils those expectations.

The paper makes a methodological contribution because it proposes a novel way of dealing with the well-known free riding problem that lies at the very root of the generation of information by consumers. The notion of *utilitarian agents* has been proposed by Harsanyi (1980) and formalised by Feddersen and Sandroni (2006a, 2006b), as a solution to the so-called "paradox of no-voting".⁸ The key insight of the utilitarian model is that the optimal probability of voting is

⁶In the context of a voting game, Feddersen and Sandroni (2006a) show that a behaviour rule profile that defines rules such that each agent decides he must follow given a proper anticipation of the behaviour of other agents can be described by cutoff points. Their result extends to the application in this paper.

⁷See Geanakoplos, Pearce, and Stacchetti (1989)

⁸If voting is costly then, since the likelihood of a vote being pivotal is very small, standard game-theoretic models predict low levels of turnout (Downs 1957).

between zero and one: not everybody should stay at home, but not everybody should vote either.⁹ This assumption is useful in the context of this paper because it constitutes a plausible explanation for consumers' behaviour in settings in which tangible benefits accrue only if aggregate participation is high, and no one can be excluded from the benefits of group success. The application in this paper is, to the best of my knowledge, the first formal application of Harsanyi's ideas outside the area of Political Economy.

The model assumes the existence of a formal regulation based on complaints. However, all that is needed for the results is that consumers' complaining decision is driven by their disappointment with the quality they received and that the firm is somehow "punished" when consumers complain (and the consumers are aware of this possibility). This situation is more general than the regulatory context I use for presentation purposes. Consider, for instance, a hotel chain that tries to verify that each of its members delivers an appropriate level of service. Clearly in this context what constitutes an "appropriate" service depends on the consumers' preferences. Thus, the chain may want to rely in customers' feedback to learn how much effort each of its members is exerting. When doing this, the chain is assuming that the feedback given by consumers can be compared across hotels and along time. The results in this paper suggest that this is not always the case.

There exist many other examples of the type of situation considered in this paper. For example, Amazon keeps record of customers' complaints about the various companies that use the platform and may prevent them from continuing

⁹If everybody stays at home, the policy will not pass (or the favourite candidate has no opportunity of winning the election), but everybody voting would result in a surfeit of votes, imposing unnecessary costs on society. In this way, the logic of rule-utilitarianism yields an elegant theory of turnout. Harsanyi (1980) assumes that everyone does their duty, but rejects the implicit assumption that doing one's duty always involves voting.

to use it if the number of customers' complaints is high enough.¹⁰ Amazon behaves in this case as a sort of "regulatory agency" that punishes the firm based on the amount of complaints. Another clear example of customers' "dissatisfaction" that was followed by a firm being punished is the decision of some major retailers to stop using the delivery services of Youdel -the biggest delivery service in the United Kingdom, outside of Royal Mail.^{11 12}

The rest of the paper is organised as follows. The reminder of this section briefly revises the existing literature and discusses the contributions made by this paper. Section 2.2 presents the details of the model and describes the players action spaces and payoff functions. It also discusses how consumers' quality expectations are formed. Section 2.3 analyses the implications of complaints for the firm's investment decision in a one shot game. This exercise is useful because it highlights most of the strategic considerations that will shape the equilibrium when the game is repeated. This section also analyses how the equilibrium proportion of complaints is affected by changes in the various parameters of the model and how informative is that proportion about the firm's investment. Section 2.4 studies a repeated version of the complaining game. Finally, section 2.5 concludes.

Related Literature

The model in this paper contributes to the extensive literature on quality provision by a monopoly. Starting with the seminal papers of Spence (1975) and Shesinski (1982), the literature suggests that an unregulated monopoly will over or under supply quality according to whether the marginal consumer values addi-

¹⁰This feature is independent of the well-known reviewing system that allows consumers and buyers to rate each other (or among them).

¹¹This includes major retailers like John Lewis, Mothercare and Matalan. (The Guardian 2012).

¹²According to The Guardian (2012) "about 5,000 customers posted messages in Amazon's online forums calling for the online retailing giant to stop using the parcel delivery company".

tional quality more or less highly than do infra-marginal consumers on average.¹³ It has further been shown that regulation of service prices can compound, ameliorate or otherwise complicate the already existing market failure (see for example, Spence (1975), Mussa and Rosen (1978) and Besanko, Donnenfeld, and White (1987); Sappington (2005) surveys the literature).¹⁴ The main conclusion of much of the existing literature is that when quality is not verifiable the regulator needs to have a great deal of information before even knowing in which direction he should intervene. By studying the informational content of complaints, this paper considers whether the “policing power” could be moved from the regulatory agency to consumers.

This paper is also related to the literature on reference dependence utility¹⁵ and with some models in marketing research on customer satisfaction.¹⁶ In both cases, it is suggested that consumers utility depends not only on the quality he actually received but also on whether that quality was above or below some reference level. This paper adds to the first branch of literature because, in spite of being widely accepted, the effect of that reference point on consumers’ complaint decisions and on the firm’s incentives to invest have not been previously studied.¹⁷ It differs from the second branch in that they do not require consumers’

¹³The difference depends on whether quantity and quality are seen as substitutes or as complements by consumers. In the former case, consumers willingness to pay a higher price for an increase in quality decreases with the quantity that he buys (i.e., the demand curve becomes less elastic as quality increases), while in the latter the elasticity of the demand increases with quality. Thus, the monopoly is more likely to undersupply quality if quality and quantity are substitutes, and to oversupply it if they are complements.

¹⁴Price ceilings that are independent of the firm’s realised costs limits its incentives to supply quality, because they prevent the firm from capturing any of the incremental consumer’s surpluses that would result from the higher service quality. However, as noted by Laffont and Tirole (1993), even under pure cost-of-service regulation, the regulated firm does not gain from providing costly services either, so a low perceived cost of supplying quality does not imply a high incentive to provide quality.

¹⁵Kahneman and Tversky (2001), Koszegi and Rabin (2006), among others.

¹⁶See for example, Singh (1988), Zithaml, Berry, and Parasuraman (1996), Boulding, Kalra, Staelin, and Zeithaml (1993), Oliver (1977), Oliver (1980).

¹⁷In a different context, Akerlof (2010) shows that norms may be followed because a failure to do so provokes anger and (potentially) punishment.

expectations to be rational and, as a result, they are not able to make clear predictions about the firm’s strategic response to consumers’ complaints.

Apart from the application in this paper and the voting literature, Harsanyi (1980)-type arguments have also been used to explain household responses to conservation appeals during the California’s energy crisis in 2000 and 2001. Reiss and White (2008) find empirical evidence that consumers do respond to voluntary appeals provided the costs of a collective action failure are tangible and that the public is well aware of it. In this case, each household faces private costs of reducing consumption, a virtually zero possibility of bringing about any tangible benefit with respect to the crisis through individual effort, and a considerable incentive to free-ride on whatever efforts are made by others. The nature of individual free-rider problems here and the lack of private incentives for electricity conservation leave largely “moral suasion”-type arguments to explain their behaviour: consumers individually wanting to “do their part” to mitigate the crisis.

2.2 The Model

This section presents a static game of quality regulation based on customers’ complaints. A regulated monopoly decides whether to make a costly investment that increases the level of quality received by the consumers in a first order stochastic dominance sense. After observing a quality realisation, the consumers may file a complaint to “inform” the regulator they received a low quality realisation. The regulatory agency is not an strategic player, it observes the proportion of customers that complained (δ) and fines the firm if that proportion is above a threshold $\bar{\delta}$. The fine equals m times the firm’s revenues, with a probability that is proportional to the level of complaints. Hence, a *regulatory rule* consists of a

pair $(\bar{\delta}, m) \in [0, 1]^2$ of parameters that are public information.

The model assumes that customers complain if *they feel disappointed with the quality they received and consider the firm should be punished for its “poor performance”*. Consumers’ disappointment is defined as the difference between the level of quality they were expecting to receive (\hat{z}) and the one they actually received (q). If they are disappointed, the consumers consider the regulator should fine the firm and so they complain in order to increase the probability with which the firm is *punished*. The fact that consumers’ prior expectations affect their complaining behaviour implies that the complaining game belongs to the class of psychological games.¹⁸

The firm faces a unit demand for its product and an exogenously given price, p (a binding price cap). Thus, its revenues are deterministic and independent of its investment decision.¹⁹ As a result, its only incentive to invest in quality is to reduce the (expected) proportion of complaints and, hence, the expected value of the fine.

The section proceeds as follows. Section 2.2.1 presents the payoff function of the consumers and discusses their complaining decision, while section 2.2.2 considers the firm’s investment decision. Finally, section 2.2.3 explains how consumers’ expectations are formed.

¹⁸Psychological games differ from standard games in that the domain of the utility function includes explicitly the beliefs a player holds about the other players’ strategies. As a result, payoffs at a given endnode are endogenous: beliefs determine the player’s utility and they are explained/predicted via some solution concept (Battigalli and Dufwenberg (2009), Geanakoplos, Pearce, and Stacchetti (1989)). In the context of this paper, this means that a given level of investment may lead to different final payoffs for different pre-play beliefs (consumers expected quality). The standard assumption is that beliefs are correct in equilibrium, and that is the condition I impose in the equilibrium definitions of sections 2.3 and 2.4.

¹⁹This implies that the firm’s investment in quality is not aimed at increasing future demand; see Shapiro (1982) for model a in which the firm’s incentives to investment are related with future demand.

2.2.1 The Consumers

There is a continuum of consumers normalised to size one. After receiving a quality draw, each consumer decides whether to lodge a complain. Hence, his action space is $\mathbb{C}_i \in \{0, 1\}$, where $\mathbb{C}_i = 1$ means consumer i files a complaint. Each consumer's utility is the sum of the consumption utility he derives from the quality he received and, if he makes a complaint, his payoff from complaining. The consumer's payoff from complaining depends on his disappointment and on his cost of complaining, but also on whether the firm is punished for its "poor performance". Each consumer i faces a cost of complaining $\sigma_i c$, where σ_i is the realisation of a random variable uniformly distributed over $[0, 1]$, and c is a positive constant. σ_i is independent of any other random variable in the model. Consumers do not observe the cost of other consumers, but do know the distribution from which they are drawn. The utility of an individual consumer i with cost $c\sigma_i$, who was expecting \hat{z} and received q is:

$$U_i(\mathbb{C}_i; q, \hat{z}, \sigma_i) = q + \theta(\hat{z} - q)\mathbf{1}_{\{\delta \geq \bar{\delta}\}}(\delta) - \mathbb{C}_i c \sigma_i \quad (2.1)$$

where $\theta \in (0, 1)$ can be interpreted as the consumer's marginal utility per unit of punished-disappointment and the indicator function $\mathbf{1}_{\{\delta \geq \bar{\delta}\}}\{\delta\} \in \{0, 1\}$ takes the value 1 if the firm is fined (i.e., if $\delta \geq \bar{\delta}$) and zero otherwise.²⁰ Implicit in the utility function is the additional assumption that consumers heterogeneity is restricted to individual costs of complaining (σ_i); this means that all the consumers have the same willingness to complain and the same intensity of preferences over

²⁰ $\theta < 1$ implies that, everything else constant, individual utility is an increasing function of quality.

quality. The assumption is relevant in that it sidesteps the question of how the burden of complaining should be shared among consumers with different intensities of preferences.

The utility function reflects the assumption that the consumer complains in order to “punish the firm’s poor performance”. The consumer receives a positive payoff only if $q < \hat{z}$ and the firm is fined. The implications in terms of complaining behaviour are twofold. First, if the realised quality is above the quality the consumer was expecting to receive he will not lodge a complain. Second, a disappointed consumer is willing to face the cost of complaining if by doing so he increases the probability with which the firm is punished.

However, individual consumers cannot appropriate the benefits of their complaints. If the firm is fined, every consumer receives a payoff $\theta(\hat{z} - q)$ independently of whether he made a complaint or not. Only those agents who actually filed a complaint ($\mathbb{C}_i = 1$) face the costs. As there is continuum of consumers, the model as defined so far suffers from an extreme version of free-riding. Hence, without additional assumptions there would be no complaints in equilibrium. To overcome this limitation, I borrow from the voting literature the assumption that consumers are “*group - utilitarians*”: they receive a positive payoff for acting according to a strategy that maximises consumers’ aggregate utility.²¹ Formally, the utilitarian assumption implies that the group’s problem is strategically equivalent to a one person decision problem with payoff function defined as consumers’ aggregate (expected) utility.²²

²¹The qualitative results would not change if only a proportion $\gamma \in (0, 1)$ of consumers were group utilitarians, as long as either γ or the distribution from which it is drawn, is public information.

²²As discussed by Feddersen and Sandroni (2006a), one possible intuition is that, if a consumer believes that all the other utilitarian agents will use the same strategy as he does himself, he will independently decide that the right strategy is the one that maximises aggregate utility. In this way, a consumer will be willing to face the cost of complaining even though he understands that his single complaint has no effect on the final outcome. However, the mathematical

Let a *social rule* σ' be a cut off that specifies a critical cost level below which a consumer makes a complaint.²³ By the law of large numbers, the proportion of complaints equals the cut off cost, i.e. $\delta = \sigma'$. The utilitarian assumption implies that the group's expected utility from following a rule σ' , when they received quality q and were expecting \hat{z} , is:²⁴

$$\mathbb{E}U(\sigma'; q, \hat{z}) = \begin{cases} q + \theta(\hat{z} - q)\sigma' - \frac{c}{2}\sigma'^2 & \text{if } \sigma' \geq \bar{\delta} \\ q - \frac{c}{2}\sigma'^2 & \text{if } \sigma' < \bar{\delta} \end{cases} \quad (2.2)$$

Consumers' complaining decision is made after they observed quality. Given their disappointment, the consumers problem is to choose the cut off rule σ^* that maximises (2.2). Thus, consumers' strategy is a mapping from their disappointment $(\hat{z} - q)$ into a cutoff point between zero and one: $\sigma(q; \hat{z}) : [0, 1]^2 \rightarrow [0, 1]$. The cut off rule that maximises consumers' expected utility, given a realisation of quality and consumers' expectations, is:²⁵

structure of the model is equivalent to the one in elite driven turnout models. In the complaining game of this paper, this second interpretation could mean, for example, that consumers follow the directions of some sort of Consumers' Associations.

²³Given a social rule σ' , a (utilitarian) consumer's action is:

$$\mathbb{C}_i(\sigma_i, \sigma') = \begin{cases} 1 & \text{if } \sigma_i < \sigma' \\ 0 & \text{otherwise} \end{cases}$$

²⁴Expectation is taken with respect to the rule σ' . The probability that an agent makes a complaint is $Prob(\sigma_i \leq \sigma') = \sigma'$. The expected cost of complaining, conditional on the consumer effectively making a complaint is $E(\sigma_i | \sigma_i \leq \sigma') = (1/\sigma') \int_0^{\sigma'} x dx = \sigma'/2$.

²⁵ $\sigma^*(q; \hat{z})$ could in principle take any value in the interval $[0, 1]$; however, it is clear from (2.2) that values of σ^* different from zero but smaller than $\bar{\delta}$ cannot be optimal. Consumers' optimisation problem can then be written as: $\text{Max}_{\sigma} \{ \mathbb{E}U(q, 0); \text{Max}_{\sigma \in [\bar{\delta}, 1]} \mathbb{E}U(q, \sigma(q; \hat{z})) \}$.

$$\sigma^*(q; \hat{z}) = \begin{cases} 1 & \text{if } q \leq \hat{z} - \psi \\ \frac{\theta(\hat{z}-q)}{c} & \text{if } \hat{z} - \psi < q \leq \hat{z} - \bar{\delta}\psi \\ \bar{\delta} & \text{if } \hat{z} - \bar{\delta}\psi < q \leq \hat{z} - \frac{\bar{\delta}\psi}{2} \\ 0 & \text{Otherwise} \end{cases} \quad (2.3)$$

where $\psi = \frac{c}{\theta}$. Given \hat{z} , the proportion of complaints induced by σ^* is decreasing in the realised quality: the smaller is q the higher is consumers' disappointment and so is the cost they are willing to face in order to have the firm punished. The optimal cut-off rule is shown in Figure 2.1. The flat regions for very low quality realisations and for $q \in (\hat{z} - \bar{\delta}\psi, \hat{z} - \frac{\bar{\delta}\psi}{2}]$ are due to the restrictions that the proportion of complaints cannot be higher than 1 in the first case, and that the probability of fine becomes zero for $\sigma^* < \bar{\delta}$ in the latter.²⁶ Finally, note that there is a “*region of tolerance*” in which consumers do not complain despite the realised quality being below \hat{z} . Within this region, the group's disappointment is not high enough to compensate the cost of a proportion of complaints equal to or greater than $\bar{\delta}$.²⁷ This result supports some arguments made in the marketing literature that define a “zone of tolerance” within which “the company is meeting customer expectations” (Singh 1988).²⁸

2.2.2 The Firm's Investment Decision


The firm is risk neutral and seeks to maximise expected profits. It can be of any of two types, “bad” (B) or “good” (G). The good firm's investment is

²⁶Consumers' optimal strategy in (2.3) implicitly assumes that when the realised quality is exactly $\hat{z} - \frac{\bar{\delta}c}{2\theta}$ consumers do complain, even though they are indifferent between complaining in a proportion $\bar{\delta}$ and not complaining at all. The exact way in which this indifference is broken does not affect the results.

²⁷Not even utilitarian consumers would follow a social rule that directs a positive proportion of consumers to complain within this region of quality realisations.

²⁸According to (Singh 1988), this region is delimited by the *desired service level* and the *adequate service level* (i.e., the level of service the customer will accept).

Figure 2.1: Relationship between σ^* and q



a binary decision $I_G \in \{L, H\}$, where H means the firm invests and delivers higher (expected) quality and L means it does not invest (and hence, it delivers low quality). The bad firm has a singleton action space $I_B \in \{L\}$. In order to keep the model tractable, it is assumed that quality is uniformly distributed over $[0, 1/2]$ if the firm is bad or if it is of the good type but it does not invest, and $q \sim U[0, 1]$ if the good firm invests.²⁹ The firm has private information about its type. To simplify notation, I is used as a shorthand notation for I_G . Investing in increasing quality costs $h > 0$. This investment cost is independent of any other cost faced by the company and it is public information.³⁰ The firm also faces a (potential) cost derived from the fine. If the proportion of complaints, δ , is above the threshold $\bar{\delta}$, the firm is fined with a probability equal to δ . Then, given a proportion of complaints δ , the good firm's profits are:

²⁹The uniform distributions simplify the exposition by allowing closed form solution for the expected proportion of complaints. However, all my findings remain true for a more general class of quality distributions as long as $F(q; I = H) \leq F(q; I = L) \forall q$ (with strict inequality for some q).

³⁰Results would not change if we assume that the regulator cannot observe the firm's costs. All that is required is that it is able to observe the proportion of complaints and the firm's revenues.

$$\Pi(I, \delta) = \begin{cases} p - \mathbf{1}_{I=H}\{I\} \cdot h - mp & \text{if } \delta \geq \bar{\delta} \text{ and the firm is fined} \\ p - \mathbf{1}_{I=H}\{I\} \cdot h & \text{otherwise} \end{cases} \quad (2.4)$$

where $\mathbf{1}_H\{I\}$ is an indicator function that takes the value of 1 if the firm invests ($I = H$) and zero otherwise; p denotes the firm's revenues and $mp\delta$ is the fine paid by the firm.

The firm's expected cost depends on the observed proportion of complaints, which in turn depends on the realised quality, the quality the consumers were expecting to receive and their complaining strategy. When the firm makes its investment decision, it does not observe the level of quality consumers expect to receive, but it does have some beliefs about it, \hat{z} .³¹ Given those beliefs, the good firm's expected payoff is the expectation of (2.4) with respect to the probability measure over quality induced by its investment strategy. The expected proportion of complaints when the firm invests I is $\mathbb{E}_{q_I}\delta(q_I, \hat{z}, \sigma)$. The firm's optimal action depends on the trade-off between the cost of investment and the expected fine: by not investing, the firm reduces its costs by h , but it also makes it less likely that quality meets consumers' expectations, increasing the expected value of the fine. To simplify notation, denote $\pi_H(\hat{z}, \sigma) = \mathbb{E}_{q_H}\Pi(H, \delta(q_H, \hat{z}, \sigma))$ and $\pi_L(\hat{z}, \sigma) = \mathbb{E}_{q_L}\Pi(L, \delta(q_L, \hat{z}, \sigma))$. Given its beliefs about the level of quality consumers' expect to receive, the firm invests if and only if:

³¹Consumers' complaining decision depends on their prior expectations, so the firm needs to form some beliefs about them in order to decide the level of investment that maximises its profits; \hat{z} denotes the firm's belief about z .

$$\pi_H(\hat{z}, \sigma) \geq \pi_L(\hat{z}, \sigma) \iff [\mathbb{E}_{q_L} \delta(q_L, \hat{z}, \sigma) - \mathbb{E}_{q_H} \delta(q_H, \hat{z}, \sigma)] \geq \frac{h}{mp} \quad (2.5)$$

The firm's investment strategy is a function from the level of quality the firm believes consumers expect to receive (\hat{z}) to an investment level: $I : \hat{z} \rightarrow \{L, H\}$. Denote by \hat{z}^* the level of \hat{z} at which the firm is indifferent between investing and not investing; $\frac{h}{mp}$ is constant and independent of consumers' expectations, but the change in the expected proportion of complaints when the firm's investment changes is an increasing function of \hat{z} . Then, the firm's optimal strategy is a cut-off of the form:

$$I^* = \begin{cases} H & \text{if } \hat{z} \geq \hat{z}^* \\ L & \text{if } \hat{z} < \hat{z}^* \end{cases}$$

The firm's strategy is increasing in consumers' expectations. When consumers expect too much from the firm, the firm's best reply is to fulfil those expectations, as otherwise the fine becomes too heavy. However the firm also fulfils consumers prior expectations when they are low, because if consumers do not expect much, their disappointment is not very high and the (expected) proportion of complaints is not enough to compensate the cost of investment (h).

2.2.3 Consumers' Beliefs and Expectations

The level of quality consumers expect to receive is determined by their beliefs about the type and strategy of the firm, and the equilibrium condition requires

those beliefs to be correct.³² Denote by $\bar{q}_B = \int_B q \cdot f(q; B) dq$ the average quality that is delivered by the bad type of the firm, and by $\bar{q}_{G,I} = \int_{G,I} q \cdot f(q; G, I) dq$ the average quality that is delivered by the good firm if it invests I :

$$\bar{q}_{G,I} = \begin{cases} \bar{q}_{G,H} & \text{if } I = H \\ \bar{q}_{G,L} & \text{if } I = L \end{cases}$$

Then, the level of quality consumers expect to receive is:³³

$$\hat{z}_I(\tau) = \begin{cases} \tau \bar{q}_{G,H} + (1 - \tau) \bar{q}_B & \text{if } I = H \\ \tau \bar{q}_{G,L} + (1 - \tau) \bar{q}_B & \text{if } I = L \end{cases}$$

where τ is the probability consumers assign to the firm being good.

2.3 Equilibrium

The firm and the consumers choose their actions according to their prior beliefs without observing each other's action, and the consumers do not observe the type of the firm neither. The equilibrium concept I use is, therefore, Bayesian Nash Equilibrium. In this application, such an equilibrium needs to satisfy three requirements.³⁴ First, the firm chooses the investment level that maximises its expected profits given its beliefs about consumers' cut off rule and expected quality. Second, consumers choose the complaining rule optimally given their disappointment with the quality they received (i.e., given \hat{z} and q). And third, the firm correctly anticipates consumers' expected quality, which is in turn consistent with

³²This assumption rules out beliefs structures in which, for example, the consumer reduces his prior expectations so that he does not feel disappointed if the quality realisation is low. For models of belief-dependent preferences in which the agents can choose beliefs see Akerlof and Dickens (1982) or Brunnermeier and Parker (2005), for example.

³³Consumers' beliefs about the firm's strategy are they "first order beliefs", defined as a probability distribution over the firm's action space (Battigalli and Dufwenberg 2009). As I consider only pure strategies, consumers' first order beliefs assign probability one or zero to the event in which the good firm invests.

³⁴Geanakoplos, Pearce, and Stacchetti (1989)

the firm's strategy and consumers' prior about its type, τ . Definition 1 formalises the equilibrium requirements.

Definition 1. Equilibrium in the Static Game. *An equilibrium of the complaining game when $T = 1$ is a pair of strategies (I^*, σ^*) and expected qualities (\hat{z}, \hat{z}) for which the following conditions are satisfied:*

1. I^* maximises the firm's expected profits given \hat{z} and σ^*
2. σ^* maximizes consumers' expected utility given \hat{z}
3. $\hat{z} = \hat{z} = \hat{z}_{I^*}(\tau)$

The expected proportion of complaints when the good type of the firm invests is $\mathbb{E}_{q_H}(\sigma^*(q; \hat{z})) = \hat{z} - \frac{\psi}{2}$, and when it does not invest (or when the firm is bad) is $\mathbb{E}_{q_L}(\sigma^*(q; \hat{z})) = 2\hat{z} - \psi$.³⁵ Then, the cut off point in the firm's strategy is $\hat{z}^* = \frac{h}{mp} + \frac{\psi}{2}$; \hat{z}^* is determined by the magnitude of the “punishment” (mp) relative to the investment cost (h), and by consumers' relative cost of complaining ($\psi = \frac{c}{\theta}$). The less harsh the punishment or the more difficult it is for consumers to complain, the higher is \hat{z}^* and thus the higher is the \hat{z} required for the firm's optimal action to be $I = H$.

The static game has a separating and a pooling equilibrium. In the first case, the good type of the firm invests and differentiates itself from the other type with a positive probability. In the second case, the firm does not invest and so it camouflages itself with the bad type. Define a “*High Quality Equilibrium*” (HQE) as one in which the good type of the firm invests, and a “*Low Quality Equilibrium*” (LQE) as one in which it does not. Given the equilibrium definition above, a HQE exists if and only if $\pi_H(\hat{z}_H, \sigma^*) \geq \pi_L(\hat{z}_H, \sigma^*)$, while a LQE exists if and only

³⁵ $\mathbb{E}_{q_H}(\sigma^*(q; \hat{z}))$ is the expectation of consumers' optimal strategy when $q \sim U[0, 1]$, and $\mathbb{E}_{q_L}(\sigma^*(q; \hat{z}))$ is the expectation when $q \sim U[0, 1/2]$. See Appendix A.1 for details.

if $\pi_L(\hat{z}_L, \sigma^*) \geq \pi_H(\hat{z}_L, \sigma^*)$.³⁶

Proposition 1. Equilibria of the Static Game. *Given $\psi < \frac{1}{4}$ and $\tau \in (0, 1)$:*

1. *If $\hat{z}^* \geq \frac{1}{2}$ a unique low quality equilibrium exists.*
2. *If $\hat{z}^* \in (\frac{1}{4}, \frac{1}{2})$ there exists $\tau^* \in (0, 1)$ such that a unique low quality equilibrium exists for $\tau \in (0, \tau^*)$, but high and low quality equilibria coexist for $\tau \in [\tau^*, 1)$*
3. *If $\hat{z}^* \leq \frac{1}{4}$ a unique high quality equilibrium exists.*

Proof. In equilibrium the firm has correct beliefs about the level of quality consumers expect to receive, so $\hat{z} = \hat{z}$. Given that the distributions of quality are public information, $\hat{z}_H(\tau) = \frac{1}{4} + \frac{1}{4}\tau$ and $\hat{z}_L = \frac{1}{4}$.³⁷ The firm's optimal strategy depends on whether \hat{z} is greater than or smaller than \hat{z}^* . There are three possibilities:

- When $\hat{z}^* \geq 1/2$, the cost of investing in quality is high relative to the (expected) punishment, $\hat{z}_L < \hat{z}_H(\tau) \leq \hat{z}^*$, and as a result $\pi_L(\hat{z}_L, \sigma^*) > \pi_H(\hat{z}_L, \sigma^*)$ and $\pi_L(\hat{z}_H(\tau), \sigma^*) > \pi_H(\hat{z}_H(\tau), \sigma^*)$. The firm's optimal strategy is $I = L$, independently of consumers expectations, and so rational consumers do not expect something different from low quality. There is a unique low quality equilibrium.

³⁶ $\pi_L(\hat{z}_H, \sigma^*)$ and $\pi_H(\hat{z}_L, \sigma^*)$ cannot be the firm's profits in any equilibrium of the game, as they both fail to comply with the "correct beliefs" requirement of Definition 1. In both cases, the firm's actual investment differs from its beliefs about \hat{z} , meaning that either the firm has incorrect beliefs about the level of quality consumers expect to receive or that the firm's beliefs are correct but consumers expectations are not consistent with the firm's investment strategy.

³⁷When $I = L$, the good firm does not differentiate itself from the bad one. Thus, consumers' expected quality, \hat{z}_L , is independent of τ (and so is \hat{z}_L).

- If $\hat{z}^* \in (\frac{1}{4}, \frac{1}{2})$, there exists a unique τ^* such that $\hat{z}_H(\tau^*) = \hat{z}^*$. Uniqueness is given by the fact that $\hat{z}_H(\tau) \in [\frac{1}{4}, \frac{1}{2}]$ and is a monotone and increasing function of τ , while \hat{z}^* belongs to the same interval but is exogenous and independent of τ . For $\tau < \tau^*$, $\hat{z}^* > \hat{z}_H(\tau) > \hat{z}_L$, implying that $\pi_L(\hat{z}_L, \sigma^*) > \pi_H(\hat{z}_L, \sigma^*)$ and $\pi_L(\hat{z}_H(\tau), \sigma^*) > \pi_H(\hat{z}_H(\tau), \sigma^*)$. The firm's optimal strategy is $I = L \forall \hat{z}$, and there exists a unique low quality equilibrium. As τ increases so does the level of quality consumers expect to receive if they anticipate $I^* = H$. For $\tau \geq \tau^*$, $\hat{z}_H(\tau) \geq \hat{z}^* > \hat{z}_L$ and $\pi_L(\hat{z}_L, \sigma^*) > \pi_H(\hat{z}_L, \sigma^*)$ and $\pi_H(\hat{z}_H(\tau), \sigma^*) > \pi_L(\hat{z}_H(\tau), \sigma^*)$ and so there are two equilibria: the firm optimally invests if consumers expect high quality (HQE) and the firm does not invest if consumers' expected quality is \hat{z}_L (LQE).
- If $\hat{z}^* \leq 1/4$, the (expected) punishment is harsh relative to h and $\hat{z}^* \leq \hat{z}_L < \hat{z}_H(\tau) \forall \tau \in (0, 1)$. In this case $\pi_H(\hat{z}_H(\tau), \sigma^*) > \pi_L(\hat{z}_H(\tau), \sigma^*)$ and $\pi_H(\hat{z}_L, \sigma^*) > \pi_L(\hat{z}_L)$. Then, investing is the firm's optimal strategy. As consumers anticipate this, they expect high quality ($\hat{z}_H(\tau)$) and there is a unique high quality equilibrium.

□

As shown in Appendix A.1, the condition $\psi < 1/4$ is sufficient but not necessary for the results in this section and in the next ones. This condition guarantees that there exists a positive probability of complain for every quality realisation.

Proposition 1 shows how a regulatory rule based on customers complaints affects the monopoly's investment behaviour. Such a rule induces a higher investment in quality as long as the punishment is "harsh enough". In the context of this paper this requires not only that the size of the fine -the proportion of revenues lost in case of a fine (mp)- is high relative to the cost of investment (h)

but also that the consumers do transmit their dissatisfaction to the regulator. Both conditions are summarised by the parameter \hat{z}^* : a high value of \hat{z}^* reflects a reduced effectiveness of the punishment, either because the cost of investment is high relative to the fine or because consumers' relative cost of complaining is high ($\psi = \frac{c}{\theta}$). Therefore, as \hat{z}^* increases the game moves towards a low quality equilibrium.

The quality consumers expect to receive is higher when they believe the firm is investing, but also when they assign a higher probability to the firm being of the good type -i.e., consumers expect more from a good firm. As a result, the firm's payoff in a HQE is a decreasing function of τ : the more convinced consumers are that they are facing a good firm, the more they expect and so the higher is the (expected) proportion of complaints and the lower are the firm's (expected) profits. Furthermore, the closer is τ to one, the smaller is the size of the fine required to induce investment (m).

It is worth noting, however, that the change in the set of equilibria resulting from the introduction of the fine does not necessarily imply an increase in total welfare. The (average) quality in a low quality equilibrium is the same that would be delivered without the regulatory rule. The firm's expected profits, however, are smaller after the introduction of the regulation because it faces a positive probability of fine. An equivalent statement about the change in consumers' welfare with and without the regulatory rule is less clear because it is assumed that they derive some positive utility from complaining.

In a high quality equilibrium the firm optimally invests because the cost of investing is smaller than the fine it would have to pay if a low realisation of quality results in a high proportion of complaints. Even though the level of quality

consumers receive in this case is higher than without the regulation, the cost of that quality exceeds the cost of investing by the expected fine (because there is a positive probability of fine). As this creates an inefficiency, the result is only a “second best” result. Furthermore, the cost of delivering a higher quality is increasing in consumers’ expectations and so the more convinced are consumers that the firm is of the good type, the higher is the cost of the quality increase.

2.3.1 Informativeness of Complaints

The regulatory agency may be interested in punishing the firm more harshly when it is not investing. However, the proportion of complaints observed by the regulator reflects consumers’ disappointment with the quality they received and not necessarily the quality itself. This section identifies conditions under which a higher proportion of complaints reflects both a higher disappointment and a lower investment - i.e., when $\mathbb{E}_{q_L}(\sigma^*(q; \hat{z}_L)) \geq \mathbb{E}_{q_H}(\sigma^*(q; \hat{z}_H))$. If this is the case, I say that complaints are *informative* about the equilibrium being played.

Given \hat{z} , the expected proportion of complaints is lower in a separating than in a pooling equilibrium because the probability of high quality realisations is higher when the firm invests -i.e., $\mathbb{E}_{q_L}(\sigma^*(q; \hat{z})) \geq \mathbb{E}_{q_H}(\sigma^*(q; \hat{z}))$. However, in equilibrium consumers have correct beliefs about the firm’s strategy and they modify their expectations accordingly. Because consumers expectations are higher in a high quality equilibrium, it is not clear whether they will complain more when the firm is not investing. Lemma 1 presents the conditions under which complaints are informative in the one shot game. In order to study the informativeness of complaints, I focus on the set of parameters for which a high quality and a low quality equilibria coexist.

Lemma 1. *Given $\hat{z}^* \in (\frac{1}{4}, \frac{1}{2})$, complaints are informative about the equilibrium being played if and only if the following conditions hold:*

1. $\tau \in (\tau^*, 1 - 2\psi)$, where τ^* is such that $\hat{z}_H(\tau^*) = \hat{z}^*$
2. $\psi \in (0, \frac{1}{2} - \frac{h}{mp})$

Proof. First, note that the definition of informativeness of complaints is based in the existence of multiple equilibria, so the analysis is restricted to $\tau > \tau^*$ and $\hat{z}^* \in (\frac{1}{4}, \frac{1}{2})$. $\mathbb{E}_{q_L}(\sigma^*(q; \hat{z}_L)) > \mathbb{E}_{q_H}(\sigma^*(q; \hat{z}_H(\tau)))$ if and only if $2\hat{z}_L - \hat{z}_H(\tau) > \frac{\psi}{2}$ (where $\psi = \frac{c}{\theta}$). Given the distributions of quality in each case, $\hat{z}_L = \frac{1}{4}$ and $\hat{z}_H(\tau) = \frac{1}{4} + \frac{1}{4}\tau$, which is an increasing function of τ . Then, the expected proportion of complaints is higher in a LQE than in a HQE for values of $\tau \in (\tau^*, 1 - 2\psi)$, and part (1) of the Lemma holds.

The second part of the Lemma implies that for complaints to be informative, complaining must be “neither too cheap nor too costly”. When the relative cost of complaining, ψ , is zero the expected utility in (2.2) is maximised when every consumers complains if $\hat{z} > q$ (because $\sigma^* = 1$ maximises the probability that the firm is fined) and when nobody complains if $\hat{z} \leq q$.³⁸ Then, the proportion of complaints becomes constant and independent of consumers’ disappointment. Finally, when complaining is very costly, complaints are not informative because the set of τ ’s determined in the previous paragraph is empty. $\tau \in (\tau^*, 1 - 2\psi)$ is not an empty set if $1 - 2\psi \geq \tau^*$. From Proposition 1, $\tau^* = 4\hat{z}^* - 1$, where $\hat{z}^* = \frac{h}{mp} + \frac{\psi}{2}$. Then, the second condition in the Lemma implies that complaints are informative only for values of ψ in the set $(0, \frac{1}{2} - \frac{h}{mp})$ -which is smaller than the set induced by the condition that $\hat{z}^* \leq 1/2$. \square

The Lemma shows that complaints are not always a good signal of the firm’s

³⁸Note that $\psi = 0$ may be the result of either $c = 0$ or $\theta \rightarrow \infty$.

investment. Consumers complaints are not informative of the equilibrium being played whenever the change in their disappointment between the low and the high quality equilibria is driven by a change in their expectations and not by a change in the (average) quality being delivered by the firm. When consumers are reasonably convinced that the firm is “good” (high τ), \hat{z} increases more than the (average) realisations of quality and so complaints are (on expectations) higher in an equilibrium in which the firm invests. In this case, complaints are informative about how disappointed consumers are with the quality they received but not about the firm’s investment. As a result, the firm might be punished more harshly when it is investing than when it is not.

The condition $\tau \leq 1 - 2\psi$ means that, given τ , the informativeness of complaints decreases if ψ increases -i.e., if the cost of making a complaint is higher relative to consumers’ willingness to complain. Hence, the informativeness of complaints depends also on how easy it is for consumers to complain. The result in the Lemma shows that if complaining is too costly, the level of disappointment required for consumers to be willing to face the cost of “informing” the regulator is too high and so the regulator observes only a small proportion of complaints -i.e., consumers do not complain enough so as to transmit information to the regulator. On the other extreme, if $c = 0$ the optimal social norm becomes independent of the size of the difference between expected and realised quality, and the (expected) proportion of complaints is the same in both equilibria.³⁹ When complaining is very cheap, the proportion of consumers that lodge a complain is so high that complaints become meaningless.

The limited informativeness of complaints is due to the fact that consumers’ complaining decision does not depend solely on the realisation of quality but also

³⁹When $c = 0$, $\hat{z}^* = \frac{h}{mp}$ and the firm’s incentives to invest depend solely on the relative magnitude of the investment cost and the fine.

on their prior expectations. As a result, there is not a unique relationship between the proportion of complaints and the firm's investing behaviour. However, the existence of a reference point is a necessary condition for the existence of a positive proportion of complaints in any equilibrium of the game.

2.3.2 Comparative Static of the Optimal Complaining Rule

The optimal rule in (2.3) can be used to derive predictions about the way in which complaints depend on the exogenous variables in both, consumers' individual utility and the regulatory rule. Those predictions are summarised in the following properties. Figure 2.2 presents the changes in the optimal complaining rule resulting from the properties below. In all the cases, the continuous line represents consumers' optimal strategy before the change and the dashed line is their optimal strategy after it.

Property 1. *The (expected) proportion of complaints is increasing in consumers' prior expectations.*

A higher \hat{z} increases consumers' disappointment with every realisation of quality (q), and so consumers are willing to accept the higher social cost that results from an increase in the cutoff point. This effect is showed in Figure 2.2a.

Property 2. *Given consumers' disappointment, the optimal social rule decreases when $\psi = \frac{c}{\theta}$ increases.*

An increase in c or a decrease in θ implies that the relative utility consumers derive from complaining is reduced. As a result, the optimal proportion of complaints for any given quality realisation is smaller. The only exception are very

low realisations for which it is still optimal to direct every ethical agent to complaint.

Property 3. *An increase in $\bar{\delta}$ makes it more costly for consumers to punish the firm when quality realisations are relatively high.*

A higher $\bar{\delta}$ makes it more difficult for consumers to meet the regulator's requirements and so it increases the group's cost of punishing the firm for high levels of quality, when the payoff of complaining is relatively low. For realisations of quality that induce a $\sigma^* > \bar{\delta}$, the change in $\bar{\delta}$ does not affect the optimal cut off rule.


2.4 The Repeated Game

In this section I explore how the repetition of the game affects the firm's incentives to invest in quality and the informativeness of consumers' complaints. When the complaining game is played repeatedly, consumers' beliefs are updated at the end of every period and so their reference point changes over time. The firm's strategy depends on consumers current expectations, but also on how today's investment affects the level of quality they expect to receive in the future: higher investment in a given period reduces that period's expected fine but it increases consumers' future expectations and so it increases the probability of being fined in the future. In this way, the repetition of the game generates incentives for the good firm to induce particular beliefs in the consumers.⁴⁰ The main result is that a repeated interaction between the consumers and the firm may reduce both the firm's incentives to invest and the informativeness of complaints. The later is due to the fact that the adverse effect on investment is stronger for values


⁴⁰The psychological aspect of the game means that current actions affect future play like in standard dynamic games, but they also affect players' beliefs and, because beliefs affect payoffs, current actions affect future payoff for any possible action.

Figure 2.2: Changes in the Optimal Cutoff as the Parameters Change


(a) Increase in \hat{z}



(b) Increase in ψ



(c) Increase in $\bar{\delta}$



of the parameters for which complaints would have been informative in the one shot game: by not investing in any of the T periods, the firm keeps consumers' expectations (and the proportion of complaints) low, and so the regulator can not infer the lack of investment from the proportion of consumers that complain.

The public information in period t is the history of past quality realisations and the proportion of consumers' complaints, $h^t = (q_1, \sigma_1; q_2, \sigma_2; \dots; q_{t-1}, \sigma_{t-1})$ for $t \geq 2$. The set of public histories is then $\mathcal{H} = [0, 1]^{2(T-1)}$. When making a complaining decision, consumers know the public history up to that moment but they also have private information about the level of quality they expect to receive (and the level they expected in any previous period). The set of their private histories is then defined as $\mathcal{H}_C = \bigcup_{t=1}^T H_C^t$, where $H_C^t = [0, 1]^{3(t-1)}$. Consumers' optimal rule maximises the present value of their expected utility, $\sigma_k^* \in \arg \max \sum_{k=t}^T \beta^{k-t} \mathbb{E}U(\sigma_k; q_k, \hat{z}_k)$ and so, their strategy in the repeated game is a sequence of complaining decisions $\{\sigma_k\}_{k=t}^T$, each of which maps their private history and the current period quality into a cutoff between zero and one, $\sigma_t : H_C^t \times [0, 1] \rightarrow [0, 1]$.⁴¹

The firm has information about past quality and complaints, but also about its own past actions and its beliefs about consumers' expectations. Then, a private history for the good type of the firm includes both the public history and the history of its investment decisions and beliefs. The firm's private history up to period t is $H_G^t = \{L, H\}^{t-1} \times [0, 1]^{3(t-1)}$ and the set of all possible private

⁴¹Consumers make their period- t complaining decision after observing the realisation of q_t , so this last quality realisation also forms part of the information they have when deciding how strongly to complain. Also note that individual consumers have private information about their costs of complaining and actions. However, the group utilitarian assumption implies that their strategy is the same that would result if there were only one "big" consumer. As a result, the only relevant information is the distribution of the costs of complaining (which is public information). The latter holds because, as there is a continuous of anonymous consumers, each of them can do no better than myopically follow the complaining rule. See Maliath and Samuelson (2006).

histories for the firm is $\mathcal{H}_G = \bigcup_{t=1}^T H_G^t$. The strategy of the good type of the firm in period t is a sequence $\{I_k\}_{k=t}^T$ that assigns, in each period, an investment level for any possible private history, $I_t : H_G^t \rightarrow \{L, H\}$. The firm's investment decision in period t (I_t^*) maximises the present value of its profits, which is the expectation with respect to the probability distribution induced by the current investment; this implies that given a history H_G^t , I_t^* solves:⁴²

$$\max_{I_t \in \{L, H\}} \pi_{I_t}(\hat{z}_t, \sigma_t^*) + \mathbb{E}_{I_t}[\sum_{k=t+1}^T \beta^{k-t} \pi_{I_k}(\hat{z}_k, \sigma_k^*)] \quad (2.6)$$

The equilibrium of the repeated game requires players behaviour to be optimal in every period given their beliefs about the other players' type and strategy but also their understanding of the way in which current behaviour affects future payoffs. At the end of every period consumers update their beliefs about the type of the firm and form some expectations about the level of quality they should receive in the following period. The equilibrium requires consumers' beliefs to be correct in the sense of being consistent with the firm's strategy in the repeated game. Beliefs about the type of the firm are required to be consistent in a bayesian way. At the beginning of the game consumers assign a probability τ_1 to the firm being good. At the end of each period they update that probability using Bayes' rule (and the firm's strategy). It is worth mentioning that the probability distributions I am assuming imply that quality realisations between zero and 1/2 can not be off the equilibrium path, while realisations higher than 1/2 can be off the equilibrium path but they are fully revealing of the firm's type. Definition 2 formalises the requirements for an equilibrium.

⁴²Recall that $\pi_I(\hat{z}, \sigma) = \mathbb{E}_{q_I} \Pi(I, \delta(q_I, \hat{z}, \sigma))$

Definition 2. *An equilibrium of the complaining game when $T > 1$ is a sequence of strategies $\{I_t^*, \sigma_t^*\}_{t=1}^T$ and expected qualities $\{\hat{z}_t, \hat{z}_t\}_{t=1}^T$ such that:*

1. *For each $h_G^t \in H_G^t$, I_t^* maximises the present value of the firm's expected profits.*
2. *For each $h_G^t \in H_G^t$, σ_t^* maximises the present value of consumers expected utility.*
3. *For each $h_G^t \in H_G^t$ the firm has correct beliefs about the level of quality consumers expect ($\hat{z}_t = \hat{z}_t$) and consumers expectations are consistent with the equilibrium strategies in the repeated game and Bayes' Rule.*
4. *For each h_G^t in which every $q_k \leq 1/2$, consumers' beliefs about the type of the firm are updated according to Bayes' Rule and the firm's strategy; otherwise, $\tau_{t+1} = 1$ and consumers expect the firm to invest.*

The definition states that consumers' optimal strategy in the repeated game maximises the sum of their current and future expected utility. However, the specific regulatory rule I am studying implies that the firm is punished in the same period in which complaints occur. Hence, if the consumers complaining strategy is a function only of their beliefs (i.e., if consumers use Markov strategies), current complaints do not affect the firm's future behaviour. As rational consumers anticipate this, $\mathbb{E}(q_{t+1}; \sigma_t) = \mathbb{E}(q_{t+1})$ and $\mathbb{E}(\hat{z}_{t+1}; \sigma_t) = \mathbb{E}(\hat{z}_{t+1})$. This results in consumers behaving *as if* they were myopic: consumers complain only to punish the firm's current poor performance and so $\sigma_t^*(q_t; \hat{z}_t) = \sigma^*(q_t; \hat{z}_t)$. This result is presented in Lemma 2.

Lemma 2. *In any Markovian Equilibrium the consumers behave as if they were myopic.*

Proof. Consider the firm's optimal action in the last period. In period T its optimal strategy is a cut off analogous to the one in the one-shot game: the firm invests as long as $\hat{z}_T > \hat{z}^*$ and it does not invest if the inequality is reversed. In equilibrium, $\hat{z}_T = \hat{z}_T$, which depends on consumers' beliefs about the firm's type (τ_T) and investment strategy. τ_T is a function of past quality realisations (through Bayesian updating) and consumers beliefs about the firm's strategy reflect common knowledge of the strategy profile. Therefore, the firm's investment in period T is determined by past realisations of quality but it is not affected by the fact that the firm was fined in previous periods. A similar argument explains why past fines (and hence past complaints) do not affect current or future investment in periods before the last one. As consumers anticipate the firm's best response, they do not expect current complaints to affect future quality. Consumers' problem in the repeated game, $\max_{\sigma_k} \sum_{k=t}^T \beta^{k-1} \mathbb{E}U(\sigma_k; q_k, \hat{z}_k)$ is then equivalent to $\sum_{k=t}^T \beta^{k-1} \max_{\sigma_k} \mathbb{E}U(\sigma_k; q_k, \hat{z}_k)$ and so $\sigma_t^*(q_t; \hat{z}_t) = \sigma^*(q_t; \hat{z}_t)$. \square

The intuition behind this result is that, as the firm is punished in the same period in which complaints occur, past fines (and hence past complaints) become a sunk cost when the firm decides its current (and future) investment. Rational consumers understand that future quality (and future firm's behaviour) is not affected by current complaints and so their complaining strategy is the myopic best response to the quality realisation they received, given their prior expectations -i.e., they complain in order to punish the firm's current "poor performance", and not to affect its future behaviour. An important consequence of Lemma 2 is that without the presence of the reference point in consumers' utility function, the optimal complaining strategy would be $\sigma_t^* = 0$ for every $t = 1, 2, \dots, T$, because not even utilitarian consumers would receive a positive payoff from complaining.

The formal implication of consumers' myopic behaviour is that the repeated game is strategically equivalent to a game in which a long-lived firm faces a se-

quence of short-lived consumers, each of which plays only once but observes all previous realisations of quality and complaints.

2.4.1 Equilibrium

The definition of informativeness of complaints I introduced in section 2.3 is based on the existence of multiple equilibria, as it compares the (expected) proportion of complaints in a low quality equilibrium with that in a high quality equilibrium. Therefore, in this section I consider only the set of parameters for which high and low quality equilibria coexist in the stage game and I study how that set is affected by the repetition of the game. This means that I focus on values of $\hat{z}^* \in [1/4, 1/2]$ and $\tau_1 \geq \tau^*$ (see Proposition 1). Recall that $\hat{z}^* = \frac{h}{mp} + \frac{\psi}{2}$ summarises the main parameters determining the strength of the punishment, namely, the relative size of the fine and consumers' cost of complaining. In a way analogous to the one in section 2.3, I say that there is a *high quality equilibrium* (HQE) when the firm invests in every period, and a *low quality equilibrium* (LQE) when it does not invest in any period. Furthermore, as two periods are enough to prove the main results, this section presents only the case in which $T = 2$. The case in which $T \rightarrow \infty$ is presented in Appendix A.3.⁴³ Proposition 2 summarises the main result of the repeated game.

Proposition 2. *Let τ^* be as defined in Proposition 1. Given $\hat{z}^* \in [1/4, 1/2]$, $\psi < \frac{1}{4}$ and $T = 2$, there exist $\tau^{**} \leq 1 - 2\psi$ such that for every $\tau_1 \in (\tau^*, \tau^{**})$, the one shot game has a high quality and a low quality equilibrium, but the repeated game has only a low quality equilibrium. As a result, complaints are less informative in the repeated game than in the one shot game.*

⁴³The infinite horizon game shows that the results of this Section do not depend upon the existence of a final period.

The proof of this proposition is divided in three parts. Lemmas 3 and 4 below characterise the set of parameters for which a high and a low quality equilibria exist in the two-period game, while Lemma 5 relates those results to the degree of informativeness of complaints.

Lemma 3. *Given $\hat{z}^* \in (1/4, 1/2)$ and $T = 2$, there exists an equilibrium in which the firm invests in both periods for every $\tau_1 > \tau^{**}$, for some $\tau^{**} \in (0, 1)$.*

Proof. Denote by s a strategy in which the firm invests in the first period and it invest in the second period only if $\tau_2 > \tau^*$. The value for the firm from following strategy s in the first period is:⁴⁴

$$\begin{aligned} V^s(\tau_1) &= \pi_{H,1}(\hat{z}_{H,1}(\tau_1), \sigma^*) + \beta \left[\frac{1}{2} \pi_{H,2}(\hat{z}_{H,2}(1), \sigma^*) + \right. \\ &\quad \left. + \frac{1}{2} [\Pr(\tau_2 > \tau^*) \pi_{H,2}(\hat{z}_{H,2}(\tau_2), \sigma^*) + \Pr(\tau_2 \leq \tau^*) \pi_{L,2}(\hat{z}_{L,2}, \sigma^*)] \right] \end{aligned}$$

Under the equilibrium strategy, $\tau_2 = \frac{\tau_1}{2-\tau_1}$ after $q_1 \leq 1/2$ and $\tau_2 = 1$ after $q_1 > 1/2$. For $\tau_2 > \tau^*$, the consumers expect the firm to invest in the second period and so the level of quality they expect is either $\hat{z}_{H,2}(1)$ or $\hat{z}_{H,2}(\tau_2)$; for $\tau_2 < \tau^*$, their expected quality is $\hat{z}_{L,2}$.

If the firm deviates in the first period, the realisation of quality is smaller than $1/2$ for sure; consumers do not detect the deviation, but the (low) quality realisation induces them to reduce the probability they assign to the firm being good and to lower their second period expectations accordingly. Hence, the value of deviating in the first period but following strategy s in the second one is:

⁴⁴From Lemma 2, the consumers' complaining strategy is the same in every period. Therefore, I denote it by σ^* and not by σ_t^* .

$$\begin{aligned}
V^d(\tau_1) = & \pi_{L,1}(\hat{z}_{H,1}(\tau_1), \sigma^*) + \beta \left[\Pr(\tau_2 > \tau^*) \pi_{H,2}(\hat{z}_{H,2}(\tau_2), \sigma^*) + \right. \\
& \left. + \Pr(\tau_2 \leq \tau^*) \pi_{L,2}(\hat{z}_{L,2}, \sigma^*) \right]
\end{aligned}$$

By looking at the game backwards, it can be shown that for τ_1 high enough the investment strategy s , together with expected qualities $\hat{z}_{H,1}(\tau_1)$, $\hat{z}_{H,2}(\tau_2)$ for $\tau_2 > \tau^*$ and $\hat{z}_{L,2}$ for $\tau_2 \leq \tau^*$, constitute an equilibrium of the two-period game. Consider the second period. As this is the last period of the game it is equivalent to a one shot game with prior τ_2 and so, from Proposition 1, the firm has no incentives to deviate from s . The firm follows s in the first period if $V^s(\tau_1) \geq V^d(\tau_1)$. The difference $[V^s(\tau_1) - V^d(\tau_1)]$ is monotone, increasing and continuous in τ_1 , $V^s(0) < V^d(0)$ and $V^s(1) > V^d(1)$. Then, there exists a unique prior belief $\bar{\tau}$ such that $V^s(\bar{\tau}) = V^d(\bar{\tau})$, and so for any $\tau_1 \geq \bar{\tau}$ and $t \in \{1, 2\}$ the firm has no incentives to deviate from s .

The fact that the firm follows strategy s may or may not result in $I_2 = H$. The firm invests in the second period if and only if consumers expect so and $\tau_2 > \tau^*$. This second condition holds for $\tau_1 \geq \bar{\tau} = \frac{2\tau^*}{1+\tau^*}$. In this case, $\Pr(\tau_2 > \tau^*) = 1$ and the firm invests in the first period too if $\tau_1 \geq \bar{\tau} = 4\hat{z}^* - 1 + \frac{\beta}{2}(1 - \frac{\bar{\tau}}{2-\bar{\tau}}) = \tau^* + \frac{\beta}{2}(1 - \tau_2)$.⁴⁵ Then, there exists an equilibrium in which the firm invests in both periods if and only if $\tau_1 \geq \tau^{**} = \max\{\bar{\tau}, \bar{\bar{\tau}}\}$.⁴⁶ Consumers' expected qualities in equilibrium are $\hat{z}_1 = \hat{z}_{H,1}(\tau_1)$ and $\hat{z}_2 = \hat{z}_{H,2}(1)$ if $q_1 > 1/2$ and $\hat{z}_2 = \hat{z}_{H,2}(\tau_2)$ if $q_1 \leq 1/2$. \square

⁴⁵Substituting the firm's expected profits, this expression becomes $V^s(\tau_1) - V^d(\tau_1) = -h + mp(\hat{z}_{H,1}(\tau_1) - \frac{\psi}{2}) - \frac{\beta}{2}mp[\hat{z}_{H,2}(1) - \hat{z}_{H,2}(\tau_2)] > 0$. As both $\hat{z}_{H,1}(\tau_1)$ and $\hat{z}_{H,2}(\tau_2)$ are increasing in τ_1 , while $\hat{z}_{H,2}(1)$ and $\hat{z}_{L,2}$ are independent of τ_1 , the difference is increasing in τ_1 .

⁴⁶Note that while $\bar{\tau} > \tau^*$ and $\bar{\bar{\tau}} \geq \tau^*$, $\bar{\tau} \geq \bar{\bar{\tau}}$ depending on how patient is the firm. In particular, there exists $\hat{\beta}$ such that when $\beta > \hat{\beta}$ $\bar{\tau} > \bar{\bar{\tau}}$ while if $\beta \leq \hat{\beta}$ the opposite is true. To see that this is the case, note that $\tau^{**} = \tau^* < \frac{2\tau^*}{1+\tau^*}$ if $\beta = 0$, τ^{**} is strictly increasing in β and $\tau^{**} > \frac{2\tau^*}{1+\tau^*} > \tau^*$ when $\beta = 1$.

Lemma 4. *Given $\hat{z}^* \in [1/4, 1/2]$ and $T = 2$, there exists an equilibrium in which the firm does not invest in any of the two periods.*

Proof. The second period is the last period of the game. From Proposition 1, if consumers expected quality in the second period is \hat{z}_L , the firm's best reply is $I_2 = L$ as long as $\hat{z}^* \geq 1/4$. On the equilibrium path, $I_1 = L$ and so a low quality realisation in the first period is not informative about the firm's type. Hence, $\tau_2 = \tau_1$ and consumers expect low quality in the second period too.

Given $\hat{z}^* \geq 1/4$, the firm has no incentives to deviate either in the first or in the second period. If the firm deviates in the first period (i.e., it invests when consumers were expecting \hat{z}_L), that period's (expected) profits are smaller because the smaller fine does not compensate the higher cost of investment. The (expected) second period profits do not increase with the deviation either. After $I_1 = 1$, there exists a positive probability that $q_1 \leq 1/2$. In this case, the consumers do not detect the firm's deviation and second period profits are not affected by the deviation. There also exists a positive probability of a high quality realisation that reveals the firm's type. In this case, $\hat{z}_{2,H}(1)$ and the firm's second period profits are smaller under the deviation than under the equilibrium path. Then, $(I_1^*, I_2^*) = (0, 0)$ together with $\hat{z}_1 = \hat{z}_2 = \hat{z}_L$ and $\hat{z}^* \in (1/4, 1/2)$ constitute an equilibrium when $\hat{z}^* \geq 1/4$, and the Lemma is proved. \square

It can be further shown that the low quality equilibrium holds for a wider set of parameters. In particular, it can be shown that the firm's expected profits under the deviation are smaller than under the equilibrium strategy as long as

$$\hat{z}^* > \frac{1}{2(2+\beta)}.^{47}$$

⁴⁷The value for the firm of following the non-investment strategy is $\pi_{L,1}(\hat{z}_{L,1}, \sigma^*)(1+\beta)$, and the value of deviating in the first period is $\pi_{H,1}(\hat{z}_{L,1}, \sigma^*) + \frac{\beta}{2}[\pi_{H,2}(\hat{z}_{H,2}(1), \sigma^*) + \pi_{L,2}(\hat{z}_{L,2}, \sigma^*)]$. The first expression is greater than the latter if and only if $\hat{z}^* > \frac{1}{2(2+\beta)}$. Thus, if the firm is very impatient ($\beta = 0$), we are back in the one shot game and the LQE holds for $\hat{z}^* > \frac{1}{4}$, but if $\beta \rightarrow 1$, the firm will not invest in any of the two periods as long as $\hat{z}^* > \frac{1}{6}$. Therefore, the region of parameters in which the firm does not invest is greater in the repeated game.

Lemma 5. *Customers complaints are less informative in the repeated game than in the one shot game.*

Proof. Lemmas 3 and 4 show that when $T = 2$ and $\hat{z}^* \in (1/4, 1/2)$, there exists an equilibrium in which the firm does not invest in any of the two periods for any $\tau_1 \in (0, 1)$, but an equilibrium in which the firm invests in $t = 1, 2$ exists for $\tau_1 \geq \tau^{**}$. Proposition 1 shows that, for the same set of values of \hat{z}^* , the static game also has a low quality equilibrium for any τ_1 , and a high quality equilibrium for $\tau_1 \geq \tau^*$. As long as $\tau_1 < 1$, $\tau^{**} > \tau^*$ and so there exists a set $\tau_1 \in (\tau^*, \tau^{**})$, for which the HQE exists in the one shot game but not in the repeated game. $\tau^{**} - \tau^* = \frac{\beta}{2}(1 - \tau_2)$, which implies that the more patient is the firm and the smaller is the initial τ_1 , the greater is the set of values of τ_1 for which the repetition of the game eliminates the high quality equilibrium.

In the one shot game, complaints are informative if $\tau_1 \in (\tau^*, 1 - 2\psi)$. As shown in Lemma 1, the lower bound on τ_1 is the minimum level at which the firm is willing to invest, while the upper bound results from imposing the condition that the (expected) proportion of complaints is higher in an equilibrium in which the firm is not investing than in one in which $I = 1$. Imposing the same condition to each period of the repeated game results in complaints being informative about the firm's investment only if $\tau_1 \in (\tau^{**}, 1 - 2\psi)$. This guarantees that $\mathbb{E}_{q_H}(\sigma^*(q; \hat{z}_{1,H}(\tau_1))) < \mathbb{E}_{q_L}(\sigma^*(q; \hat{z}_L))$ and that $\mathbb{E}_{q_H}(\sigma^*(q; \hat{z}_{2,H}(\tau_2))) < \mathbb{E}_{q_L}(\sigma^*(q; \hat{z}_{2,L}))$.⁴⁸ Therefore, the set (τ^*, τ^{**}) for which the equilibrium with high quality ceases to exist in the repeated game contains values of τ_1 for which complaints would have been informative if $T = 1$. Hence, the repetition of the game also reduces the degree of informativeness of complaints. \square

Lemma 5, together with Lemmas 3 and 4, shows that the set of τ_1 's for which the firm invests in equilibrium is reduced by the repetition of the game, which

⁴⁸ $\tau_2 < \tau_1$, so $\tau_1 < 1 - 2\psi$ implies $\tau_2 < 1 - 2\psi$.

in turn reduces the informativeness of complaints. The intuition behind this result is that when the firm invests it faces the risk that consumers find out its type. When that happens in the last period (or when the game is played only once) it does not affect the firm's continuation value and so it optimally invests if the reduction of the (expected) fine compensates the cost of investment. However, when consumers find out the type of the firm before the end of the game, they increase the level of quality they expect to receive in the future, reducing the firm's future profits. The cost of reducing the current fine is higher in the repeated game as it includes not only the cost of a higher (current) investment but also the cost of smaller (expected) future profits. The change in the firm's expected profits due to higher consumers expectations is a decreasing function of τ_1 and so the firm's incentives to keep consumers' expectations low are higher for smaller values of τ_1 . When τ_1 is small, the inter temporal trade-off is more relevant for its investment decision than the intra temporal trade-off and the firm's optimal action is to keep future expected quality low by not investing today. On the contrary, when the probability consumers assign to the firm being of the good type is very high, there is only a small scope to "manage" consumers expectations, what makes the intra temporal trade-off more relevant. In this case, the firm's profit maximisation strategy is to invest if consumers expect high quality and not to invest if they expect so. The relevance of the inter temporal trade off also depends on how patient is the firm: the higher β the more value the firm assigns to future profits and the more it cares about keeping consumers expectations low.⁴⁹

The appendix A.3 shows that, the same as in the finitely repeated game, when $T \rightarrow \infty$, the existence of a high quality equilibrium depends on the scope of the firm to manage consumers' future expectation, which is positively related to β and

⁴⁹In the case in which the game is repeated two periods, the lower bound of τ goes up from $4\hat{z}^* - 1$ to $4\hat{z}^* - 1 + \frac{\beta}{2}(1 - \tau_2)$.

inversely related to τ_1 . Furthermore, as discussed in this section, the less scope to manage consumers expectations the less informative are complaints about the equilibrium being played (because the higher is the proportion of complaints in a high quality equilibrium).

2.5 Conclusions

The model in this paper considers the role of customer complains as a regulatory tool in contexts in which quality is relatively well-perceived by the consumers, but it is very costly for the regulator to observe it. In line with some empirical evidence, it is assumed that consumers complain because they feel “disappointed” with the level of quality they received and consider the firm should be punished for its “poor performance”. The paper studies the firm’s incentives to invest and the informativeness of customers complaints in such a context.

It is shown that complaints are informative when complaining is neither too cheap nor too costly and when consumers assign a relatively low probability to the firm being good. The presence of a rational reference point in consumer’s complaining decision implies that “disappointment” and “poor performance” are endogenously determined. As a result, the proportion of complaints may be higher in a high quality equilibrium than in a low quality one if consumers believe they are facing a good firm. Hence, the regulator may observe more complaints when the firm is investing than when it is not and may punish the firm more harshly in the first case. The paper further shows that the degree in which complaints are informative is reduced by the repetition of the game. It is shown that when the agency uses a regulatory rule as the one analysed in this paper, consumers’ optimal behaviour in the repeated game is the myopic best response to the quality they received, given they prior expectations. This behaviour creates

the conditions for the firm to try to “keep expectations low” making complaints less informative in the repeated game than in the one shot game. As a result, the set of parameters for which complaints are informative about the firm’s behaviour is reduced when the game is played repeatedly.

The results in this paper provide insights that may be useful in interpreting consumers’ complaints (or other type of feedback) in a variety of settings. The context considered here is one in which consumers receive no direct benefit out of their complaints and thus their only reason to lodge a complaint is to transmit their dissatisfaction. However, it is likely that in another settings consumers do appropriate at least partially the benefit of their complaints. In this case, complaints may be explained by a combination of reasons, one of which could be dissapointment. Therefore, a note of care should be taken when interpreting consumers’ complaints also in those settings.

Chapter 3

Prices, Reviews and Endogenous Information Transmission

3.1 Introduction

Online reviews of products, services or business are an increasingly important source of consumers information about experience goods, i.e., goods whose quality is learned only after consumption. Recent empirical evidence suggests that reviews are also becoming a more relevant determinant of the firm's revenues, either because of their impact in the quantity demanded or because consumers are willing to pay a sort of "reputation premium" for products or services that have good reviews.¹

Closely related, though less documented, is the practice of offering important online discounts and then ask the buyers to complete reviews.² This practice seems to be increasingly used by recently established (or recently refurbished)

¹See, for example Luca (2011), Chevalier and Mayzlin (2006). As shown by Doyle and Waterson (2012) and Resnick, Zeckhauser, Swanson, and Lockwood (2006), among others, the effect of reviews on revenues also seems to be present in the case of online auctions.

²Price discounts are offered through a variety of web pages, like groupon.com or voucher-codes.com, for example.

small firms, like restaurants or pubs.³ It also seems to be an increasingly common practice among restaurants or hotels in touristic areas. A common denominator to both situations is that the firm does not have complete information about the demand function it faces and so it values the information consumers can provide it. At the same time, being experience goods provided by new small firms or firms located far away from the consumer, it is unlikely that the consumer has ex ante information about the value of the firm's product or service.

Despite the growing importance of customers' reviews, their role in the firm's pricing decisions has not been studied. This paper proposes a dynamic model to investigate how prices and reviews affect each other. It considers a situation in which a long lived firm faces a sequence of short lived consumers whose only information about the value of the product is the one contained in the reviews completed by previous consumers. As in the examples above, it is further assumed that the firm does not know the actual value of the product either.⁴ After buying the product, the consumer observes a quality realisation and decides which review to complete (if any). The model assumes that consumers complete reviews in order to maximise the joint expected utility of current and future (potential) buyers.

The results offer an explanation for those price discounts based on the value of the information contained in the reviews. It is shown that the information generated by the reviews is valuable for both, the consumers and the monopoly. As a result, the consumers and the firm "share" the cost of generating information. It is further shown that consumers are willing to complete reviews only if it is not too costly and the firm cannot appropriate all the surplus generated by the

³The list of business that resort to this type of practices is considerably long. Apart from restaurants and pubs, it seems to be a common practice among recently established hairdressers, beauty saloons and various entertainment-related firms.

⁴As a result, the model in this paper is closer to a screening model, in the sense that the firm uses the price not to signal the quality of its product but to induce consumers to "transmit" information to the firm about it.

increased information. In this way, the incentives of the firm are “align” to those of the second period consumers.

The existence of the reviews induces a mean preserving spread on the agents’ beliefs about the value of the good. As the posterior beliefs form a martingale with respect to the reviews completed by previous consumers, reviews do not affect the expected value of the posterior beliefs, but do increase their variability. Combined with the convexity of the indirect utility and the profit functions, the increased variability of the posterior beliefs results in the information contained in the reviews being valuable for both, the consumers and the firm. Hence, both parties are willing to face some cost in order to increase the information available in the market.

The paper shows that, from the firm’s perspective, the cost of the information contained in the reviews takes the form of a “discount” in the price offered to current consumers. It is widely believed that the firm’s decision to offer price discounts is due to an intention of “getting a good review”. The result in this paper offers an alternative explanation. By reducing the current price, the firm increases current (expected) demand which in turn increases the probability with which the current consumer completes reviews. As this discount has the additional effect of compensating consumers for the cost of completing reviews, it also induces a reviewing rule that is more favourable to the firm (in the sense that it increases the expected future profits in the scenario with reviews).

From the perspective of the consumers, the price discount behaves as a “subsidy” to the reviewing activity and thus it has an effect similar to a reduction in the cost of completing reviews. It is further shown that the reviewing rule chosen by consumers is “softer” the lower is the cost of completing reviews. In this

paper, a reviewing rule is softer than another rule if the posterior belief resulting from a bad review is higher and the one following a good review is lower. A softer reviewing rule has a higher positive impact on the firm's future profits because it induces a mean preserving spread relative to a tougher rule. As a result, the firm offers a higher price discount when it is easier for consumers to complete reviews (and thus, when it is more interested in increasing the probability with which the consumers complete reviews).

The paper shows that a necessary condition for the existence of reviews is that the firm cannot extract all the surplus generated by the increased information. Since the behaviour of the consumers and the firm changes according to the observed reviews, the informational content of the reviews has a positive value for both. As a result, the incentives of the firm are aligned with those of the consumers in the sense that both prefer the existence of a reviewing system over a situation with no information transmission. Consumers complete reviews in order to increase the sum of current and future consumers' expected utility. Hence, if the firm could appropriate all the surplus consumers would not complete reviews: completing reviews is a costly activity, then not even utilitarian consumers are willing to complete reviews if by doing so they do not improve the utility of those in their group (the consumers, in this case).

Before analysing the results in more detail, it is important to note that when the consumer completes a review he is taking a costly action, the benefits of which he cannot (fully) appropriate.⁵ Hence, the reviewing decision has some similarities with an agent's decision to contribute to the provision of a public good.⁶ The

⁵In a way analogous to the situations I study, empirical studies on eBay show that most of the times the customer is not likely to buy again from the same seller, implying that he does not receive a direct benefit from completing a review. Yet, Resnick and Zeckhauser (2002) report that 52.1% of the buyers on eBay actually provide voluntary feedback about their sellers.

⁶Since I make the simplifying assumption that there is a continuum of consumers, the model in this paper suffers from an extreme version of free riding. Therefore, the standard result of

free riding incentives in this context are analogous to the ones that originate the “paradox of no voting”. Thus, to tackle this difficulty I borrow from the voting literature and I assume that consumers are *group-utilitarians*, i.e., they receive a positive payoff for acting according to a strategy that maximises consumers’ aggregate utility.

From a formal perspective, the utilitarian assumption implies that the game is strategically equivalent to a two persons game, in which both players are long lived. Therefore, the proposed reviewing game becomes analogous to a situation of a bilateral monopoly, in which the firm is the only potential “buyer” of information and the group of consumers are the only potential “suppliers”. The equilibrium results suggest that the cost of completing reviews allocates the surplus created by that information between the firm and the consumers.

This paper is related with the large literature that studies how agents learn from the actions of others. See for example Bikhchandani, Hirshleifer, and Welch (1992), Smith and Sorensen (2000) and Bose, Orosel, Ottaviani, and Vesterlund (2006), Kremer, Mansour, and Perry (2012).⁷ In most part of that literature, the transmission of information is an externality: agents’ actions carry information about their private signals, and so other agents can learn from those actions. The model in this paper adds to that literature because it endogenises consumers decision to transmit information. Knowing that the previous consumer bought is informative about his preferences over quality, but not about the actual value of the good. However, after observing a realisation of quality, the consumer may

suboptimal provision obtained in public good games apply to the games analysed in this thesis in a very extreme way, resulting in no complaints/reviews in equilibrium. See Osborne (2004).

⁷Kremer, Mansour, and Perry (2012) offer a normative analysis of a situation in which agents may learn from the actions of others. They show that perfect information sharing through internet does not always support an optimal outcome. This result is due to the fact that information is a public good that is both produced and consumed by the same agents. Then, a note of care should be taken when considering the agents’ incentives to explore and produce new information.

decide to complete a review, i.e., the consumer explicitly decides whether to transmit information and which information to transmit (which review to complete).

This paper is also related to the literature on strategic information transmission and to the literature on public tests. When there is no cost of completing reviews, the results in this paper are similar to that in Crawford and Sobel (1982), in that the optimal set of messages is maximum because the preferences of the “sender” (current consumers) and the “receiver” (future consumers) are aligned. Furthermore, as the cost of completing reviews increases, the preferences become less aligned. However, the reviewing model proposed in this paper differs from the standard model of strategic transmission of information in that there is more than one “receiver”, namely the second period consumers and the firm.

The model is also related to the literature on public tests.⁸ Gill and SgROI (2012) study a framework in which a firm can have its product publicly tested before launch and tests vary in their toughness. They show that the firm always prefers to have its product tested and that it will choose a test that is either very tough or very soft. From the firm’s perspective, consumers’ reviews also constitute a “public test” about its product, and I get a similar result to Gill and SgROI’s (2012) in the sense that the firm always prefer the existence of reviews. However, the characteristics of the test in the model presented in this paper are chosen by the consumers (and only indirectly affected by the firm).⁹ This allows me to derive conclusions about the price the firm is willing to pay for tests with different degrees of “toughness”.

The rest of the paper is organised as follows. The next section describes the

⁸See Gill and SgROI (2008, 2012) and Lerner and Tirole (2006).

⁹Comparing the results of the reviewing model of this paper and the public test model as regards the toughness of the test is not as clear cut. However, the results suggest that in the model presented here the firm’s preferred test is neither the softest possible nor the toughest.

model, defines the reviewing rule and discusses its role on the public updating of beliefs. Sections 3.3 and 3.4 develop the building blocks for the equilibrium analysis of Section 3.5. Section 3.3 studies the optimal reviewing rule would information transmission be free, while Section 3.4 looks at how those result change when there exists a positive cost of completing reviews. Section 3.6 concludes. Appendix B.1 contains all the proofs that are not in the text.

3.2 Basic Setup

A risk-neutral monopolist sells a good of unknown value to a sequence of consumers. The value of the good, v , can be high (H) or low (L), with $H > L \geq 0$. Nature selects v once and forever at the beginning of the game. Neither the firm nor the consumers observe it, but they have a prior belief $\lambda \in (0, 1)$ about the good being high value.

There is a finite sequence of risk neutral buyers, each of which has a (potential) unit demand and lives for one period.¹⁰ Consumers' preferences over quality are random and change every period. Consumer t 's valuation of quality is γ_t , where each γ_t is independently distributed $U[0, 1]$, and it is independent of any other random variable in the model. At the beginning of the period, the consumer learns his valuation for quality, which is not observed by the firm. The assumption that consumers are short lived implies that an individual consumer cannot learn the value of the good from his personal experience. Furthermore, the assumptions about γ_t together with the fact that the consumer has no private information about the value of the good before buying, imply that his buying decision contains no information about v either. As a result, the only information a

¹⁰Alternatively, it can be considered that every period there is a continuum of identical consumers normalised to size one, who live during one period.

consumer has about the value of the product before buying is the one contained in the reviews completed by previous buyers.

As two periods are enough to present the main results, I consider only the case with $T = 2$. The sequence of events is as follows. Given their prior beliefs about v , a first generation of consumers choose the reviewing rule \mathcal{R} that will be followed if they buy the good. Then, the consumer observes γ_1 and p_1 and decides whether to buy or not. If he buys, he observes a realisation of quality, $q_1 \in [q^0, q^K] \in \mathbb{R}_+$. This quality realisation is distributed conditional on the actual value of the good, $q_1 \sim F_j(q)$ with $j \in \{L, H\}$ and $\mathbb{E}(q; H) = H$ and $\mathbb{E}(q; L) = L$.¹¹ It is further assumed that no quality realisation is fully revealing of the product's value and that monotone likelihood ratio property holds, so $f_H(q)/f_L(q)$ is increasing in q .

At the beginning of the second period, a new generation of consumers observe the *reviews* completed by previous buyers, their preferences over quality and the price offered by the monopolist in the previous period. The previous buyer may have completed a review $i \in \{G, N, B\}$; where G means he completed a “Good” review, N means he did not complete a review (“remain silent”) and B that he completed a “Bad” review. The consumer in period two uses this information together with the knowledge of the reviewing rule used by the previous consumers to update his beliefs about the probability of the good being high value (λ').¹²

In any of the two periods, the consumer's payoff from buying is $\gamma_t q_t - p_t$, while his payoff from not buying is zero. Thus, he buys if and only if the expected pay-

¹¹This assumption simplifies the notation and the algebra, but does not affect the results. All that is needed for the results is that $\mathbb{E}(q; H) > \mathbb{E}(q; L)$, which is implied by increasing monotone likelihood ratio.

¹²As reviews are public, consumers' beliefs about the value of the good are “public beliefs”.

off from buying is positive: $\gamma_t \mathbb{E}_{\lambda_t}(q_t) \geq p_t$.¹³ Without loss of generality, we can assume $H - L = 1$, and so $\mathbb{E}(q_t; \lambda_t) = \lambda_t H + (1 - \lambda_t)L = L + \lambda_t$. Given λ_t and the price, the consumer buys if $\gamma_t \geq \frac{p_t}{\lambda_t + L}$.

From the firm's perspective, $P(\gamma_t \geq \frac{p_t}{\lambda_t + L})$ plays the role of the demand function: given γ_t , a higher price decreases expected demand, while a higher belief about the good being of high value increases demand. The firm does not observe γ_t but it knows the distribution from which it is drawn. Given consumer's prior belief and the price, expected profits in period t are:

$$\pi(p_t; \lambda_t) = (p_t - c)P\left(\gamma_t \geq \frac{p_t}{\lambda_t + L}\right) = (p_t - c)\left(\frac{\lambda_t + L - p_t}{\lambda_t + L}\right) \quad (3.1)$$

where $c < L$ is the constant marginal cost of production.¹⁴

In what follows, I denote consumers' prior and posterior beliefs after observing a review i by λ and λ'_i , respectively.

3.2.1 Benchmark case: No Reviews

If there are no reviews, the possibility of transmitting information does not exist and so there is no updating of beliefs and $\lambda' = \lambda$. The firm's optimal pricing strategy consists in offering in every period the price that maximises static profits, i.e., the price that solves:

$$\text{Max}_p \pi(p; \lambda) = (p - c)\left(\frac{\lambda + L - p}{\lambda + L}\right)$$

¹³The weak inequality implies that if indifferent, the consumer buys.

¹⁴This assumption implies that the monopolist is willing to sell for every $\lambda_t \in [0, 1]$. If $c \in (L, H)$ the monopolist would prefer to stop selling for some $\lambda_t > 0$.

which implies:

$$\hat{p}(\lambda) = \frac{\lambda + L + c}{2} \quad (3.2)$$

Given λ , L and c , the maximum is unique because the profit function is strictly concave with respect to the price.¹⁵ In this case, the firm's maximum expected profits in any period t are: $\hat{\pi}(\lambda) = \pi(\hat{p}; \lambda) = \frac{(\lambda + L - c)^2}{4(\lambda + L)}$.

The consumer buys the good if $\gamma_t \mathbb{E}(q_t; \lambda) \geq p_t$. Given λ and \hat{p} , the probability that consumer t buys is the probability that $\gamma_t \geq \frac{\hat{p}(\lambda)}{\lambda + L}$, and his expected utility is:

$$\begin{aligned} \hat{u}(\lambda) &= u(\lambda, \hat{p}) \\ &= P\left(\gamma_t \geq \frac{\hat{p}(\lambda)}{\lambda + L}\right) \left[\mathbb{E}_\gamma\left(\gamma_t | \gamma_t \geq \frac{\hat{p}(\lambda)}{\lambda + L}\right) (\lambda + L) - \hat{p}(\lambda) \right] \\ &= \frac{[\lambda + L - c]^2}{8(\lambda + L)} \end{aligned} \quad (3.3)$$

3.2.2 Reviews

The consumers complete reviews in order to maximise the sum of current and future consumers' net (expected) utility. Consumers are utilitarians, and so they are willing to follow the social norm that maximises the group's expected utility, as long as it is not too costly.¹⁶ Consider a rule that determines two thresholds of quality realisations, $\bar{q}(p_1, \lambda)$ and $\underline{q}(p_1, \lambda)$, such that if the first-period consumer receives a quality draw greater than or equal to $\bar{q}(p_1, \lambda)$ he completes a good review, and if he receives $q_1 \leq \underline{q}(p_1, \lambda)$ he completes a bad review. Finally, if he receives a quality in between the thresholds, he completes no review. Denote by $\mathcal{R}(p_1, \lambda) = \{\underline{q}(p_1, \lambda), \bar{q}(p_1, \lambda)\}$ the reviewing rule followed by consumers in period 1. To simplify notation, I use \underline{q} , \bar{q} and \mathcal{R} as shorthand notation for $\underline{q}(p_1, \lambda)$,

¹⁵ $\frac{\partial^2 \pi}{\partial p^2} = -\frac{2}{\lambda + L} < 0$.

¹⁶ Feddersen and Sandroni (2006b).

$\bar{q}(p_1, \lambda)$ and $\mathcal{R}(p_1, \lambda)$, respectively.

The expected utility a period- t consumer derives from buying the good at price p_t , given λ_t , is $u(\lambda_t, p_t) = \mathbb{E}_\gamma \left(\gamma_t | \gamma_t > \frac{p_t}{\lambda_t + L} \right) (\lambda_t + L) - p_t$. The total utility of a first period consumer also depends on the cost of completing reviews and on the impact of his buying and reviewing decisions on the expected utility of the second-period consumer. Then, the utility of the first period consumer is:

$$U(\mathcal{R}; \lambda, p_1) = P\left(\gamma_1 \geq \frac{p_1}{\lambda + L}\right) [u(\lambda, p_1) - \Psi(\mathcal{R}; h, \lambda)] + \quad (3.4)$$

$$+ P\left(\gamma_1 \geq \frac{p_1}{\lambda + L}\right) \sum_{i \in \{B, N, G\}} u(\lambda'_i, p_2(\lambda'_i)) + \quad (3.5)$$

$$+ P\left(\gamma_1 < \frac{p_1}{\lambda + L}\right) u(\lambda, p_2(\lambda))$$

where $\Psi(\mathcal{R}; h, \lambda)$ is the expected cost of completing a review.¹⁷ The second line is the expected utility of a period-2 consumer when the previous consumer bought the good and completed review $i \in \{G, N, B\}$ according to the rule \mathcal{R} .¹⁸ The last line is the expected utility when the previous consumer did not buy. In this case, there is no updating of beliefs and $\lambda' = \lambda$.

The existence of the review system induces a sequential game between the firm and the consumers. Neither the firm nor the consumers know the actual value of the product, but the firm chooses its price knowing the reviewing rule consumers are going to follow. Consumers' problem in the first period is to choose the reviewing rule \mathcal{R} that maximises the sum of current and future consumers' expected utility, given their prior beliefs, their understanding of how future con-

¹⁷This cost function is studied in detail in Section 3.4, where I look deeply into the effects of the costs of completed reviews on the optimal reviewing rule.

¹⁸ λ'_i is a shortcut for $\lambda'_i(\lambda; \mathcal{R})$

sumers will interpret the reviews and a proper anticipation of the firm's pricing strategy. The firm, on the other hand, behaves as a "Stackelberg follower": given the reviewing rule followed by consumers, it chooses the pricing strategy that maximises the sum of current and future expected profits. Therefore, the price offered by the firm in the first period is a best response to the consumers' reviewing rule. As a result, an equilibrium of the reviewing game is defined as a pair of strategies $\{\mathcal{R}^*, p_1^*\}$ such that, \mathcal{R}^* maximises (3.4) and p_1^* maximises the present value of the firm's profits, given \mathcal{R}^* .

It is worth noting that at the moment in which the consumers and the firm choose their actions (\mathcal{R} and p_1) they have no more information about the actual value of the good than the one that is publicly available. Therefore, neither the reviewing rule nor the price are informative about the probability of the good being high value.

It becomes apparent from expression (3.4) that the price offered by the firm in period one affects the probability that the current consumer buys the good and, as a consequence, it affects the probability that current consumers transmit information to future consumers (and to the firm itself) through the reviews. The consumers' reviewing rule determines not only what information is transmitted (in the sense of which review is observed by the second-period consumer), but also which inferences future consumers (and the firm) draw from the observed reviews. Both elements affect future consumers' willingness to pay for the good.

The problems of how much information is transmitted and which information is transmitted induce different tradeoffs for the agents. Therefore, I analyse the two problems separately before solving for the equilibrium strategies of the firm and the consumers.

The paper proceeds as follows. The remainder of this section analyses the updating of beliefs after each possible review and the role of the rule in that updating. The next section studies the impact of reviews in the firm's profits and the consumers' utility without taking into account the costs of transmitting information. Section 3.4 analyses how those results change when the cost of completing reviews is taken into account. Finally, section 3.5 studies the equilibrium reviewing rule and pricing strategy.

3.2.3 Updating: Public Beliefs

At the beginning of the second period the consumers (and the firm) use the reviews completed by past consumers to update their beliefs about the good being high value. The reviewing rule divides the space of quality realisations into three intervals, determining which realisations induce which reviews. Therefore, the reading the agents do into the reviews is a function of \mathcal{R} .¹⁹

After observing a review, and given the history up to that point, the consumer and the firm use Bayes' Rule to update their beliefs about the good being high value. This beliefs are "public" in the sense that they are entirely based on public information. As a result, after observing a review i both, the firm and the consumers assign the same probability to $v = H$. When observing a good review, consumers know the realisation of quality received by the previous consumer was higher than or equal to \bar{q} . As a result, their updated belief is:

$$\lambda'_G(\lambda; \mathcal{R}) = \frac{\lambda \int_{\bar{q}}^{q^K} f_H(q) dq}{\lambda \int_{\bar{q}}^{q^K} f_H(q) dq + (1 - \lambda) \int_{\bar{q}}^{q^K} f_L(q) dq} \quad (3.6)$$

¹⁹When the consumer in the second period updates his beliefs about the value of the good, γ_1 and p_1 are already known and so the reviewing rule can be considered as given when analysing the posterior beliefs.

Analogously, after observing a bad review the consumer knows that the previous consumer received a quality realisation equal to or below the threshold \underline{q} ; his beliefs after a bad review are:

$$\lambda'_B(\lambda; \mathcal{R}) = \frac{\lambda \int_{\underline{q}^0}^{\underline{q}} f_H(q) dq}{\lambda \int_{\underline{q}^0}^{\underline{q}} f_H(q) dq + (1 - \lambda) \int_{\underline{q}^0}^{\underline{q}} f_L(q) dq} \quad (3.7)$$

Finally, if the first period consumer does not complete a review, it might be because he did not buy the good or because he bought and received a quality realisation within the no reviewing region. In the first case, the consumer in the second period has nothing to learn from the absence of review, so $\lambda' = \lambda$. In the second case, the absence of review is informative about the quality realisation being somewhere “in the middle”. The updating of beliefs in the latter case is:

$$\lambda'_N(\lambda; \mathcal{R}) = \frac{\lambda \int_{\underline{q}}^{\bar{q}} f_H(q) dq}{\lambda \int_{\underline{q}}^{\bar{q}} f_H(q) dq + (1 - \lambda) \int_{\underline{q}}^{\bar{q}} f_L(q) dq} \quad (3.8)$$

A good review increases the probability the agents assign to the good being high value, and so it constitutes “good news” in the sense that $\lambda'_G \in (\lambda, 1)$.²⁰ A bad review has the opposite effect: as it reflects a low quality realisation, it reduces the agents beliefs; thus, a bad review is “bad news” and $\lambda'_B \in (0, \lambda)$. Finally, when there are no reviews but the previous consumer bought the good, beliefs about the good being high value may increase or decrease depending on the conditional distributions of quality. However, λ'_N is always higher than the beliefs after observing a bad review, because it is an indication of a quality realisation above \underline{q} , and it is always smaller than their beliefs after observing a good

²⁰See Milgrom (1981).

review. The next two claims summarise these effects.

Claim 1. *Given \mathcal{R} and $\lambda \in (0, 1)$, good reviews are always good news about the value of the good being high, while bad reviews are always bad news. No reviews may be either good or bad news.*

Claim 2. *For every rule \mathcal{R} , no reviews is better than a bad review but worst than a good review: $\lambda'_G \geq \lambda'_N \geq \lambda'_B$.*

3.2.4 Role of the reviewing rule

The reviewing rule is chosen by consumers before deciding whether to buy the good or not and so, it is chosen without having more information about the product's value than the one that is publicly available. As consumers are not better informed than future consumers or the firm when choosing \mathcal{R} , the reviewing rule itself contains no information about v .

However, the rule does affect the beliefs of an agent that observes the reviews. A higher \bar{q} means that it requires a higher quality realisation to get a good review. As getting a good review is more difficult, consumers assign a higher probability to the good being of high value the higher is \bar{q} . Analogously, the higher is \underline{q} the more likely it is that the firm gets a bad review, so a bad review is less damaging for higher values of \underline{q} . The thresholds of the reviewing rule also affect the inferences made after observing no reviews: observing that the previous consumer completed no reviews (given that he bought the good) is better news about the quality realisation he received the higher are \underline{q} and \bar{q} , because they imply that the consumer remained silent for higher quality realisations. These intuitions are summarised in Claim 3.

Claim 3. *For any $\lambda \in (0, 1)$, λ'_G is increasing in \bar{q} , and λ'_B is increasing in \underline{q} . Given that the previous consumer bought the good, the beliefs after observing that he completed no reviews is an increasing function of both, \bar{q} and \underline{q} .*

3.3 Information Transmission

This section studies the effect of reviews on the expected payoffs of the firm and the consumers when the costs of transmitting information are not taken into account -i.e., when the cost of completing a review is zero and the price of the first period is fixed. Isolating the effects of information transmission from its costs is useful in that it highlights the strategic considerations that will shape the equilibrium of the game.

Consider the last period of the game. The optimal reviewing rule has been determined at the beginning of the previous period, and it is thus given by the time the consumer observes a review. Furthermore, the consumer and the firm know whether the previous consumer bought the good or not, so the analysis can be conditional on the previous consumer having bought. After observing a review $i \in \{G, N, B\}$, the firm's optimal action is to offer the price that maximises its static profits given the observed review, $\hat{p}(\lambda'_i)$, and hence the expected utility of the second period consumer is $\hat{u}(\lambda'_i)$.²¹ In this context, the problem faced by the first period consumer is to choose the reviewing rule that maximises the expected utility of the next consumer. Denote $V(\mathcal{R}; \lambda) = \sum_i P(i; \mathcal{R}, \lambda) \hat{u}(\lambda'_i)$. Then, the consumers' problem is:

$$\max_{\{\underline{q}, \bar{q}\}} V(\mathcal{R}; \lambda) \tag{3.9}$$

²¹See Section 3.2.1.

A necessary condition for the existence of such a rule is that the information transmitted through the reviews increases the expected utility of second period consumers.²² Whether this is the case or not depends on the curvature of the utility function because, as shown by the next claim, beliefs form a martingale.

Claim 4. *For every reviewing rule \mathcal{R} , and for every $\lambda \in (0, 1)$, beliefs form a martingale, i.e. $\mathbb{E}(\lambda'; \lambda, \mathcal{R}) = \lambda$.*

Proof. Conditional on the previous consumer having bought the good:²³,

$$\begin{aligned}
\mathbb{E}(\lambda'; \lambda, \mathcal{R}) &= \sum_{i \in \{G, N, B\}} P(i; \mathcal{R}, \lambda) \lambda'_i \\
&= \sum_{i \in \{G, N, B\}} P(i; \mathcal{R}, \lambda_t) \frac{\lambda \int_i f_H(q) dq}{P(i; \mathcal{R}, \lambda)} \\
&= \lambda \left[\int_{\bar{q}_t}^{q^K} f_H(q) dq + \int_{\underline{q}_t}^{\bar{q}_t} f_H(q) dq + \int_{q^0}^{\underline{q}_t} f_H(q) dq \right] \\
&= \lambda
\end{aligned}$$

□

The next Proposition shows that both the consumers and the firm prefer a rule in which reviews are completed with positive probability.

Proposition 3. *The information contained in the reviews increases the expected payoff of both the firm and the consumers.*

Proof. The consumers' (expected) utility is a convex function of λ' ; thus, by Jensen inequality and the martingale property of the beliefs, it is higher when there is some information transmission:

²²Otherwise, consumers would receive no benefit from the reviews and so not even utilitarian consumers would be willing to complete reviews.

²³It is shown in the Appendix that the martingale property also holds if the expectation is not conditional in the previous consumer having bought the good.

$$\begin{aligned}
\mathbb{E}_i(u; \mathcal{R}, \lambda) &= \sum_{i \in \{G, N, B\}} P(i; \mathcal{R}, \lambda) \hat{u}(\lambda'_i) \\
&\geq \hat{u} \left(\sum_{i \in \{G, N, B\}} P(i; \mathcal{R}, \lambda) \lambda'_i \right) \\
&= \hat{u}(\lambda)
\end{aligned}$$

where $\hat{u}(\lambda)$ is the expected utility of a second period consumer when there is no updating of beliefs.

A similar analysis holds for the firm's profits. Given that the previous consumer bought the good, the firm's expected profits when consumers can submit reviews are:

$$\mathbb{E}_i(\pi; \mathcal{R}, \lambda) = \sum_{i \in \{G, N, B\}} P(i; \mathcal{R}, \lambda) \hat{\pi}(\lambda'_i)$$

Since the profit function is convex, the martingale property of the beliefs and Jensen's inequality imply:²⁴

$$\begin{aligned}
\mathbb{E}_i(\pi; \mathcal{R}, \lambda) &= \sum_{i \in \{G, N, B\}} P(i; \mathcal{R}, \lambda) \hat{\pi}(\lambda'_i) \\
&\geq \hat{\pi} \left(\sum_{i \in \{G, N, B\}} P(i; \mathcal{R}, \lambda) \lambda'_i \right) \\
&= \hat{\pi}(\lambda)
\end{aligned}$$

□

Proposition 3 shows that, when its costs are not considered, the possibility of transmitting information through the reviews increases the (expected) payoff of both the consumers and the firm. The reviews completed by the consumers

²⁴Convexity of the expected profits is shown in Appendix B.1.

are a function of the quality realisations they observed, which in turn are correlated with the actual v . Therefore, the reviews are informative about the value of the good. The informativeness of the reviews increases expected profits and utility because it allows the firm and the consumers to adjust the price and the willingness to pay to a better approximation of v . It is important to note that this “*alignment*” of the incentives of the firm and the consumers holds because the monopolist cannot fully appropriate the additional surplus generated by the transmission of information. In the context of this paper, if the firm could appropriate all the surplus, leaving second period consumers indifferent between buying and not buying for every observed review, second period consumers’ would be indifferent between receiving or not the information contained in the reviews. As a result, first period consumers would be indifferent between completing reviews or not when it is costless, but they would not complete reviews when there is a positive cost of doing it.

In order to determine the existence of an optimal reviewing rule that first period consumers are willing to follow, it is useful to look at the optimal amount of messages they would chose to use when completing reviews is free.²⁵ The next Lemma shows that, when the first consumer buys the good, the expected utility of second period consumers is maximised by using all the available messages.²⁶ Proposition 4 uses the result in the Lemma to show the existence of a reviewing rule that maximises the expected utility of second period consumers.

Lemma 6. *Assume that there is no cost of transmitting information. If there*

²⁵When there is no cost of completing a review, the three available messages (G , N and B) have the same unit cost and, given that the previous consumer bought, the three are informative about v .

²⁶The implication of the Lemma is that, if we consider a system with M available messages and the cost of completing any two messages is the same, consumers will always choose a reviewing rule that assigns positive probability to all the M messages as this increases the precision of the information received by second period consumers. As discuss later in this section, this result is similar to the one in Crawford and Sobel (1982), and it is related to the fact that the preferences of the “sender” and the “receiver” are aligned.

exists an optimal reviewing rule that consumers are willing to follow, then it assigns positive probability to all the available messages.

The proof of the Lemma is in Appendix B.1. It shows that the addition of a third message induces a mean preserving spread with respect to the case in which there are only two messages. Given that the utility function is convex in the prior, this means that consumers always prefer using three messages instead of two. The introduction of a third message induces a finer partition of the set of quality realisations. As a result, the level of information received by the consumers in the second period is higher and so they can adjust their behaviour to a better approximation of the actual v . The firm's expected profits are convex in the belief too; hence, its payoff is also higher when consumers use all the available messages.

When there are no costs of completing reviews, the preferences of first and second period consumers are perfectly aligned. Therefore, the result in Lemma 6 is analogous to the one in Crawford and Sobel (1982). They show that the more similar are the preferences of the sender and the receiver, the larger is the maximal number of reports in equilibrium. The model in this paper differs from the standard model of strategic information transmission in that it has two “receivers” of the information, the future consumers and the firm. It is worth noting, however, that the firm is only an “indirect” receiver, because consumers aim when completing reviews is to transmit information to future consumers. The price offered by the firm in the first period may affect consumers' choice of the optimal amount of messages. I explore this possibility in Section 3.5.1.

An immediate implication of Lemma 6, is that there exists a rule, characterised by $\bar{q} < q^K$, $\underline{q} > q^0$ and $\bar{q} > \underline{q}$ that maximises the expected payoff of second-period consumers.

Proposition 4. *Expected second period utility is maximised for some rule $\mathcal{R}_u = \{\underline{q}_u, \bar{q}_u\}$ such that $\underline{q}_u \in (0, \bar{q}_u)$ and $\bar{q}_u \in (\underline{q}_u, q^K)$.*

Proof. The Proposition results from the fact that the rule maximises a continuous function over a non-empty and compact set. The fact that the consumers prefer a rule that uses three messages over a rule that uses only two means that their expected utility increases as \underline{q} moves away from q^0 and \bar{q} moves away from q^K . \square

Furthermore, as shown in the next Lemma, this rule also maximises the firm's expected profits. As a result, the existence of the reviews aligns the incentives of the firm and the consumers.

Lemma 7. *The reviewing rule that maximises second period consumers' expected utility, \mathcal{R}_u , also maximises the firm's expected second period profits.*

Proof. After observing a review $i \in \{G, N, B\}$ the firm optimally sets the second period price at $\hat{p}_2(\lambda'_i)$. At this price, the expected utility of a second period consumer is $\hat{u}_2(\lambda'_i) = \frac{(\lambda'_i + L - c)^2}{8(\lambda'_i + L)}$ and the firm's (maximum) expected profit is $\hat{\pi}_2(\lambda'_i) = \frac{(\lambda'_i + L - c)^2}{4(\lambda'_i + L)}$. As defined before, $V(\mathcal{R}; \lambda) = \sum_{i \in \{G, N, B\}} P(i; \mathcal{R}) \hat{u}_2(\lambda'_i)$ and so $\mathbb{E}_i(\pi_2; \mathcal{R}, \lambda) = 2V(\mathcal{R}; \lambda)$. Then, if $\{\underline{q}_u, \bar{q}_u\}$ maximises $V(\cdot)$, it also maximises $\mathbb{E}_i(\pi_2; \mathcal{R}, \lambda)$. \square

The more accurate the information the firm has about the value of v , the better it can adjust its second period price and thus, the higher are its expected second period profits. At the beginning of the second period both the firm and the consumers have the same information about value of the good. Hence, a rule that maximises the information available to second period consumers also maximises the one available to the firm.

Corollary 1. *The existence of information transmission through reviews aligns the incentives of the consumers and the firm.*

3.4 Cost of completing a review

The previous section showed that there exists a reviewing rule $\mathcal{R}_u = \{\underline{q}_u, \bar{q}_u\}$ that consumers are willing to follow when completing reviews is costless and the first period price is given. When the cost of completing a review is taken into account, the three messages available to the consumers (G , N and B) are not equivalent anymore. Conditional on the previous consumer having bought the good, the three available messages are informative about the quality realisation observed by the previous consumer, but while a good or a bad review have a positive cost, not completing a review is costless. When the firm's response in terms of first period price is not considered, the rule is not affected by (and does not affect) the probability that the first consumer buys. This section considers the effect of the cost of completing reviews when the first consumer bought the good. The main result is that, as long as the unit cost of completing a review is not very high, a rule that uses the three available messages is still optimal, but the set of quality realisations for which consumers do not complete reviews increases with the cost.

The cost of completing one review is $h > 0$; the total expected cost given a reviewing rule \mathcal{R} is h times the probability of completing either a good or a bad review. Then, the total expected cost, $\Psi(\mathcal{R}; h, \lambda)$, is given by the following expression:

$$\begin{aligned}
\Psi(\mathcal{R}; h, \lambda) &= h [P(B; \mathcal{R}, \lambda) + P(G; \mathcal{R}, \lambda)] \\
&= h \left[\lambda \left(\int_{q_0}^{\underline{q}} f_H(q) dq + \int_{\bar{q}}^{q_K} f_H(q) dq \right) + (1 - \lambda) \left(\int_{q_0}^{\underline{q}} f_L(q) dq + \int_{\bar{q}}^{q_K} f_L(q) dq \right) \right]
\end{aligned} \tag{3.10}$$

The first two terms in the second line define the probability of completing a good or a bad review conditional on the true value being H , while the last two terms measure the expected cost of completing a review conditional on the true value of the good being L . Consumers face the cost of completing a review only if they buy the good.

The expected cost of completing a review is an increasing function of h and \underline{q} , and a decreasing function of \bar{q} . A higher value of \bar{q} reduces the probability that the consumer completes a good review and so it reduces the expected cost:

$$\frac{\partial \Psi(\mathcal{R}; h, \lambda)}{\partial \bar{q}} = h \frac{\partial P(G; \mathcal{R}, \lambda)}{\partial \bar{q}} = -h[\lambda f_H(\bar{q}) + (1 - \lambda) f_L(\bar{q})] < 0 \tag{3.11}$$

Analogously, a higher \underline{q} increases the probability of getting a quality realisation low enough so as to complete a bad review, which increases the expected cost:

$$\frac{\partial \Psi(\mathcal{R}; h, \lambda)}{\partial \underline{q}} = h \frac{\partial P(B; \mathcal{R}, \lambda)}{\partial \underline{q}} = h[\lambda f_H(\underline{q}) + (1 - \lambda) f_L(\underline{q})] > 0 \tag{3.12}$$

When completing reviews is costless, first period consumers would be indifferent between any two rules that induce the same expected utility for the consumers in the second period. This is not true anymore when the cost of completing reviews is taken into account: would there exist two rules that deliver the same value of $V(\cdot)$, first period consumers would now prefer the one that induces the smaller probability of completing either a good or a bad review. As a result, the optimal rule when completing reviews is not free makes a more extensive use of the “cheap message”, N .

Proposition 5. *When there exists a positive cost h of completing a review, there exists an optimal rule $\mathcal{R}_c = \{\underline{q}^c, \bar{q}^c\}$ such that $\underline{q}^c < \underline{q}_u$ and $\bar{q}^c > \bar{q}_u$.*

Proof. From Proposition 6, there exists a reviewing rule $\mathcal{R}_u = \{\underline{q}_u, \bar{q}_u\}$ that maximises consumers' (expected) payoff when completing a review is costless. Now, consider how those thresholds change when the cost of completing reviews, $\Psi(\mathcal{R}; h, \lambda)$, is taken into account. Given p_1 , the reviewing rule \mathcal{R} affects the expected utility of the first period consumer through the cost of completing reviews and through its impact on future consumers' expected utility. The consumer's problem in this case is:

$$\max_{\{\underline{q}, \bar{q}\}} V(\mathcal{R}; \lambda) - \Psi(\mathcal{R}; h, \lambda)$$

Taken derivatives with respect to \bar{q} and \underline{q} , the first order conditions are:

$$\frac{\partial V(\mathcal{R}; \lambda)}{\partial \bar{q}} - \frac{\partial \Psi(\mathcal{R}; h, \lambda)}{\partial \bar{q}} = 0 \quad (3.13)$$

$$\frac{\partial V(\mathcal{R}; \lambda)}{\partial \underline{q}} - \frac{\partial \Psi(\mathcal{R}; h, \lambda)}{\partial \underline{q}} = 0 \quad (3.14)$$

When $\mathcal{R} = \mathcal{R}_u = \{\underline{q}_u, \bar{q}_u\}$, the expression in (3.13) is positive because $\frac{\partial V(\mathcal{R}_u; \lambda)}{\partial \bar{q}} = 0$ as \mathcal{R}_u maximises the second period's expected utility, but the derivative of $\Psi(\cdot)$ with respect to \bar{q} is negative for every $\bar{q} \in [q_0, q_K]$. Therefore, when the cost of completing reviews is considered, the optimal cut-off quality for good reviews, \bar{q}^c , must be higher than \bar{q}_u . Analogously, the optimal cut-off quality for bad reviews, \underline{q}^c , is below \underline{q}_u when the cost of completing reviews is taken into account. In this case, (3.14) evaluated at $\{\underline{q}, \bar{q}\} = \{\underline{q}_u, \bar{q}_u\}$ is negative: $\frac{\partial V(\mathcal{R}_u; \lambda)}{\partial \underline{q}} = 0$ and the cost of completing reviews is an increasing function of \underline{q} . As a result, $\underline{q}^c < \underline{q}_u$. □

When completing a review is costly, consumers face a trade off because transmitting more accurate information increases future consumers' expected payoff,

but it reduces the payoff of current consumers. As the intermediate message is informative but costless, they solve the trade off by making a more extensive use of this “free message”. As a result, $\bar{q}^c - \underline{q}^c > \bar{q}_u - \underline{q}_u$. The reviewing rule used by consumers when h is positive may be considered “more tough” than \mathcal{R}_u . When consumers rely more extensively on the cheapest message (N), the set of quality realisations after which the first consumer completes a good review becomes smaller and so does the set after which he completes a bad review. As extreme reviews become less likely, the updating they induce becomes more extreme: a good review constitutes better news (has a higher positive impact on beliefs and profits) and a bad review is more damaging in the sense that it induces a higher reduction of the beliefs -i.e., $\lambda'_B(\lambda; \mathcal{R}_c) < \lambda'_B(\lambda; \mathcal{R}_u)$ and $\lambda'_G(\lambda; \mathcal{R}_c) > \lambda'_G(\lambda; \mathcal{R}_u)$.²⁷

Whether it is still optimal for consumers to use the three available messages depends on the unit cost of completing reviews, h . As shown above, when the cost of completing a review, h is taken into account, the thresholds in consumers’ optimal reviewing rule become closer to the extremes. However, more accurate information is still better for second period consumers. Therefore, as shown in Lemma 8, as long as h is not very high, consumers still prefer a rule such that $\bar{q}_c < q_K$ and $\underline{q}_c > q_0$.

Lemma 8. *There exists $\underline{h}(\lambda, \bar{q}^c)$ and $\bar{h}(\lambda, \underline{q}_c)$ such that for every $\lambda \in (0, 1)$ and for every $h < \min\{\underline{h}(\lambda, \bar{q}^c), \bar{h}(\lambda, \underline{q}_c)\}$, the optimal reviewing rule implies $\underline{q}^c > q^0$ and $\bar{q}^c < q^K$.²⁸*

²⁷The distribution of posterior beliefs induced by the rule \mathcal{R}_c dominates stochastically of second order the one induced by \mathcal{R}_u . As both distributions have the same mean (by the martingale property of the beliefs) this means that the distribution induced by \mathcal{R}_u is a mean preserving spread of the one induced by \mathcal{R}_c . Combined with the convexity of the utility function with respect to the beliefs, this implies that $V(\mathcal{R}_c, \lambda) < V(\mathcal{R}_u, \lambda)$.

²⁸A formal proof of this result in Appendix B.1.

3.5 Equilibrium: Reviews and Price Discounts

This section solves for the equilibrium reviewing rule and price when $T = 2$. Given that neither the firm nor the consumers know the actual value of v , the equilibrium concept is similar to a subgame perfect equilibrium. The consumers move first and so their strategy is a reviewing rule \mathcal{R} . The firm chooses the first period price after consumers have chosen the rule and thus, its strategy assigns a price to every possible reviewing rule chosen by the consumers. The game is solved backwards: subsection 3.5.1 looks at the optimal first period price and subsection 3.5.2 presents the optimal reviewing rule and the equilibrium of the game, using the results from previous sections as building blocks.

3.5.1 Firm's pricing strategy

The firm's strategy in the two-period game consists of a first period price that maximises the sum of current and future profits, given a reviewing rule \mathcal{R} and its own optimal pricing behaviour in the last period.²⁹ In the second period the firm will charge the optimal static price, given the public beliefs about the value of v -i.e., $p_2^* = \hat{p}_2(\lambda'_i)$ if the first period consumer bought the good and completed review i , and $p_2^* = \hat{p}_2(\lambda)$ if he did not buy. From Proposition 3, the information generated by the reviews increases the monopolist's expected second period profits. Therefore, the firm has incentives to reduce the first period price if by doing so it increases the probability that the first period consumer buys the good and completes reviews. As a price discount reduces first period's profits, the firm faces a standard trade off between current and future profits. The firm's

²⁹The same is true when $T > 2$. I look at the two periods case because it is enough to derive the main intuitions.

optimisation problem, given a reviewing rule \mathcal{R} is:³⁰

$$\max_{p_1} \Pi(p_1; \lambda, \mathcal{R}) = \pi_1(p_1; \lambda) + P(\gamma_1 \geq \frac{p_1}{\lambda + L}) \mathbb{E}_i(\hat{\pi}_2(\lambda'_i)) + P(\gamma_1 < \frac{p_1}{\lambda + L}) \hat{\pi}_2(\lambda) \quad (3.15)$$

The first term is the first period expected profit, $\pi_1(p_1; \lambda) = P(\gamma_1 \geq \frac{p_1}{\lambda + L})(p - c)$. The second term is the expected second period profit if the first consumer buys the good and so there is some transmission of information between periods - $\hat{\pi}_2(\lambda'_i) = \pi_2(\hat{p}_2; \lambda'_i)$. The last term is the firm's expected second period profit when the first period consumer does not buy (and so $\lambda' = \lambda$). From Proposition 3, $\mathbb{E}_i(\hat{\pi}_2(\lambda'_i)) \geq \hat{\pi}_2(\lambda)$. This maximisation problem results in an optimal first period price:

$$\begin{aligned} p_1^*(\mathcal{R}; \lambda) &= \frac{\lambda + L + c - [\mathbb{E}_i(\hat{\pi}_2(\lambda'_i)) - \hat{\pi}_2(\lambda)]}{2} \\ &= \hat{p}_1(\lambda) - \frac{[\mathbb{E}_i(\hat{\pi}_2(\lambda'_i)) - \hat{\pi}_2(\lambda)]}{2} \end{aligned} \quad (3.16)$$

which is smaller than the static optimal price by an amount that depends on the increase in the (expected) future profits induced by the reviews.³¹ Denote the first period profits induced by price $p_1^*(\mathcal{R}; \lambda)$ by $\pi_1^*(\lambda; \mathcal{R})$.

The price function in (3.16) is the firm's best reply to consumers' reviewing rule. The difference in expected second period profits, $[\mathbb{E}_i(\hat{\pi}_2(\lambda'_i)) - \hat{\pi}_2(\lambda)]$, is the value for the firm of the information contained in the reviews. The firm is no better informed than consumers are about v and so the price it chooses in the first period is not a signal about the actual value of the good. From the firm's perspective, the role of $(\hat{p}_1 - p_1^*)$ is to assign probabilities between two possible states of the world: one in which the first period consumer buys the good, and so

³⁰Recall that λ'_i is a shortcut for $\lambda'_i(\lambda; \mathcal{R})$.

³¹The firm's maximisation problem considers the case in which the monopolist's discount factor is equal to one. Considering a discount factor $\delta \in (0, 1)$ would result in a first period price closer to the static optimal price but would not change the main intuitions.

he completes reviews, and another state in which the first period consumer does not buy (and so $\lambda' = \lambda$). The firm's expected profits are higher in the first case, but increasing the probability of that state requires a reduction of p_1 that reduces (expected) first period profits. Therefore, $[\hat{\pi}_1(\lambda) - \pi_1^*(\lambda; \mathcal{R})] > 0$ constitutes a measure of the “price” paid by the firm in order to get information about v .³² It is apparent from expression (3.16) that the price discount the firm is willing to offer depends on the benefit it expects to receive from the information contained in the reviews, which in turns depends on the rule chosen by consumers (\mathcal{R}).

3.5.2 Optimal Reviewing Rule and Equilibrium

The consumers' problem is to choose a reviewing rule that, given a proper anticipation of the firm's pricing strategy, maximises the sum of the utilities of first and second periods' consumers. Therefore, their strategy is a mapping from their prior beliefs (λ) and cost of completing reviews (h) into a pair of thresholds $\{\underline{q}, \bar{q}\} \in [q_0, q_K]^2$. Their maximisation problem is as follows:

$$\max_{\{\underline{q}, \bar{q}\}} P\left(\gamma_1 \geq \frac{p_1(\mathcal{R}; \lambda)}{\lambda + L}\right) [u_1(p_1; \lambda) - \Psi(\mathcal{R}; h, \lambda) + \mathbb{E}_i(\hat{u}_2(\lambda'_i))] + P\left(\gamma_1 < \frac{p_1(\mathcal{R}; \lambda)}{\lambda + L}\right) \hat{u}_2(\lambda) \quad (3.17)$$

where $p_1(\mathcal{R}; \lambda)$ is the price function in (3.16). The expression above shows that the reviewing rule chosen by the consumers affects, through its effect on p_1 , the probability that the first period consumers buys and, if he buys, his expected utility. The rule also affects the expected utility of the second period consumer because the reading the consumers do into the reviews affects their valuation for the good, which in turn determines the equilibrium values of $p_2(\cdot)$ and $u_2(\cdot)$.

³²It can also be considered as the price paid by the firm in order to have its product “tested” by consumers.


From the consumers' perspective, the smaller first period price compensates part of the cost of completing reviews and so they optimally move the reviewing rule closer to the one that maximises $V(\cdot)$.³³ As shown in Lemma 7, such a rule also increases the firm's expected second period profits, which compensates the profits lost because of the discount. These intuitions are formalised in the following Equilibrium Proposition.

Proposition 6. *Equilibrium. Given $\lambda \in (0, 1)$ and $h < \min\{\underline{h}, \bar{h}\}$, there exists a reviewing rule $\mathcal{R}^* = \{\underline{q}^*, \bar{q}^*\}$ and a first period price $p_1^*(\mathcal{R}^*, \lambda) < \hat{p}_1(\lambda)$, such that \mathcal{R}^* is the consumers' best response to $p_1^*(\mathcal{R}^*, \lambda)$ and $p_1^*(\mathcal{R}^*, \lambda)$ is the best reply of the firm to $\mathcal{R}^*(\cdot)$. Furthermore, the equilibrium reviewing rule is such that $\underline{q}^* \in (\underline{q}^c, \underline{q}_u]$ and $\bar{q}^* \in [\bar{q}_u, \bar{q}^c)$.*

Proof. The conditions about λ and h guarantee that the consumers' problem has an interior solution. The result in the Proposition can be proved by contradiction. Considered the graph in Figure 3.1. The figure presents the thresholds of the rules \mathcal{R}_c and \mathcal{R}_u ; as shown in Proposition 5, $\bar{q}_c > \bar{q}_u$ and $\underline{q}_c < \underline{q}_u$. Consider a rule such that \mathcal{R} in the Figure. This rule's thresholds are $\bar{q} < \bar{q}_u$ and $\underline{q} > \underline{q}_u$ (i.e., it is a "softer" rule than \mathcal{R}_u). Starting from rule \mathcal{R}_u , moving the thresholds according to the rule \mathcal{R} reduces the expected utility of second period consumers, $V(\cdot)$. Furthermore, as this rule implies a higher probability of both, good and bad reviews relative to \mathcal{R}_u and \mathcal{R}_c , it also increases the cost of completing reviews. As a result, $V(\mathcal{R}, \lambda) - \Psi(\mathcal{R}; h, \lambda) < V(\mathcal{R}_u, \lambda) - \Psi(\mathcal{R}_u; h, \lambda) < V(\mathcal{R}_c, \lambda) - \Psi(\mathcal{R}_c; h, \lambda)$. \mathcal{R} would also imply a smaller expected utility for first period consumers: from Lemma 7, $\mathbb{E}(\pi_2; \hat{p}_2(\lambda'_i))$ is maximised when the reviewing rule is \mathcal{R}_u . Then, for any other rule the firm offers a smaller price discount, which reduces the expected utility of first period consumers. Then, starting from \mathcal{R}_u consumers have no in-

³³As discussed below, the fact that the price discount has a similar effect (from consumers' point of view) to a reduction in h does not mean that reducing p_1 and reducing h are substitutes from the firm's perspective.

Figure 3.1: Reviewing Rule



centives to move to a “softer” rule like \mathcal{R} .

On the other hand, consider a rule like \mathcal{R}' in the figure, with $\underline{q}' < \underline{q}_c$ and $\bar{q}' = \bar{q}_c$. Recall that \mathcal{R}_c maximises the expected utility of second period consumers net of the cost of completing reviews. Therefore, $V(\mathcal{R}', \lambda) - \Psi(\mathcal{R}'; h, \lambda) < V(\mathcal{R}_c, \lambda) - \Psi(\mathcal{R}_c; h, \lambda)$. Starting from \mathcal{R}_c , this new rule reduces the firm’s expected second period profits when there are reviews, and thus it implies a smaller discount in the price of the first period. Therefore, the expected utility of first and second period consumers is smaller with a rule like \mathcal{R}' than with \mathcal{R}_c and so consumers have no incentives to move to a rule tougher than \mathcal{R}_c . A similar analysis applies in the case of a rule like \mathcal{R}'' , for which $\underline{q}' = \underline{q}_c$ but $\bar{q}' \geq \bar{q}_c$. Thus, consumers best response to the firm’s pricing strategy must be a rule \mathcal{R}^* such that $\underline{q}^* \in (\underline{q}^c, \underline{q}_u]$ and $\bar{q}^* \in [\bar{q}_u, \bar{q}^c)$ -i.e., a rule whose thresholds lie within the dotted part of the quality line of Figure 3.1.

□

Proposition 4 showed that there exists a reviewing rule \mathcal{R}_u that maximises consumers expected second period utility when the cost of completing reviews is not taken into account. The firm’s expected profits in the second period with reviews are also maximised by that rule. However, when consumers take into account the cost of competing reviews they move to a rule that induces smaller probabilities of completing good and bad reviews. As a result, a bad review has a more damaging impact on the firm’s profits and a good review has a higher

impact on the posterior (though it has a smaller probability). The firm is willing to pay a higher “price” for the first rule than for the second one and, as the price discount “subsidises” part of the cost of completing reviews, it induces first period consumers to choose a reviewing rule that is closer to the one that increases the expected payoff of both, the firm and the second period consumers.

Discussion

Role of h . The discount the firm is willing to offer to first period consumers increases as h decreases. The smaller is the cost of completing reviews the closer is the rule chosen by consumers to the one preferred by the firm (\mathcal{R}_u). As shown before, this implies that, the value of the information contained in the reviews increases and so the firm’s best response is to reduce p_1 in order to increase the probability that the first consumer buys. In the limit in which $h = 0$, the price discount ($\hat{p}_1 - p_1^*$) is maximum. This result implies that, contrary to some widespread beliefs, the price discount is not a substitute of a smaller cost of completing reviews, but it is instead its *complement*. It also suggests that as h decreases the burden of the “cost” of transmitting information moves from the consumers towards the firm. It is worth noting that a smaller h increases the expected payoff of both, the consumers and the firm.

A Bilateral Monopoly. The equilibrium results presented above can be easily associated with a situation of bilateral monopoly. If current and future consumers are considered as a “group”, they may be regarded as the only “suppliers” of the information contained in the reviews while the firm is the unique potential “buyer”. The firm pays the consumers for the information in the reviews by reducing the first period price. Furthermore, the “price” the firm is willing to pay is higher the higher is the positive impact of the reviews on the firm’s future profits.

Hence, the cost of completing reviews is one of the elements that allocates the surplus between the two sides of the market. As h decreases, the reviewing rule “offered” by the consumers gets closer to the one preferred by the firm, and so the firm is willing to pay a higher price.

3.6 Conclusions and Further Research

This paper proposes a dynamic game to explain how the reviews completed by consumers about the quality of an experience good and the pricing strategy of a monopoly firm affects each other. The results suggest that information is valuable for both, the firm and the consumers and so they are willing to “share” the costs of generating that information.

An important result of the paper is that a necessary condition for the existence of reviews in equilibrium is that the firm cannot appropriate all the surplus generated by the information in the reviews. As a consequence, second period consumers are not indifferent between observing the reviews or not and so first period consumers are willing to complete reviews (and the firm is willing to “pay” for them).

It is worth noting that the linearity of the demand function used in the model does not allow for the possibility that second period consumers do not get part of the surplus. However, the results may be affected by a change in the demand function. Under the assumptions made in this paper, would the firm be able to leave second period consumers indifferent between buy and not buying for every observed review, first period consumers would have less (or none) incentives to complete reviews. The firm would like to reduce the price of the second period,

but it would face a commitment problem: in any finite game, after a review is observed the firm would have incentives to extract all the surplus. Solving backwards, consumers would not complete reviews if they complete reviews in order to increase the expected utility of future consumers.

The implications of this possibility are matter of future research. One possible explanation is that the reason we observe reviews is not related with consumers' intentions to maximise current and future consumers' (expected) payoff, as was assumed in this paper. Therefore, the implications of alternative assumptions about why consumers complete reviews (like anger, punishment or reciprocity, for example) should be considered.

Appendices

Appendix A

Appendix to Chapter 2

A.1 Expected Proportion of Complaints

The expected proportion of complaints depends on the realisation of quality, which depends on the firm's type and investment. Taking expectations of the consumers' optimal complaining rule with respect to the quality distribution induced by investment I , the expected proportion of complaints is:

$$\begin{aligned}\mathbb{E}_{q_I}(\sigma^*(q; \hat{z})) &= [P_{q_I}(\sigma^*(q; \hat{z}) = 1) * 1 + \\ &+ P_{q_I}(\sigma^*(q; \hat{z}) = \frac{\theta}{c}(\hat{z} - q))\mathbb{E}_{q_I}(\sigma^*(q; \hat{z}) | \sigma^*(q; \hat{z}) \in (\bar{\delta}, 1)) + \\ &+ P_{q_I}(\sigma^*(q; \hat{z}) = \bar{\delta}) * \bar{\delta}]\end{aligned}$$

Using consumer' optimal strategy in (2.3), the expectation can be rewritten in terms of the realised level of quality:

$$\begin{aligned}\mathbb{E}_{q_I}(\sigma^*(q; \hat{z})) &= [P_{q_I}(q \leq \hat{z} - \frac{c}{\theta}) + \\ &+ P_{q_I}(\hat{z} - \frac{c}{\theta} \leq q \leq \hat{z} - \frac{\bar{\delta}c}{\theta}) * \\ &* \mathbb{E}_{q_I}[\sigma^*(q; \hat{z}) | \hat{z} - \frac{c}{\theta} \leq q \leq \hat{z} - \frac{\bar{\delta}c}{\theta}] + \\ &+ P_{q_I}(\hat{z} - \frac{\bar{\delta}c}{\theta} \leq q \leq \hat{z} - \frac{\bar{\delta}c}{2\theta}) * \bar{\delta}]\end{aligned}\tag{A.1}$$

where:

$$\mathbb{E}_{q_I}[\sigma^*(q; \hat{z}) | \hat{z} - \frac{c}{\theta} \leq q \leq \hat{z}_t - \frac{\bar{\delta}c}{\theta}] = \frac{\theta}{c} (\hat{z} - E_{q_I}[q | \hat{z} - \frac{c}{\theta} \leq q \leq \hat{z} - \frac{\bar{\delta}c}{\theta}])$$

$$\mathbb{E}_{q_I}[q | \hat{z} - \frac{c}{\theta} \leq q \leq \hat{z} - \frac{\bar{\delta}c}{\theta}] = \frac{1}{P_{q_I}(\hat{z} - \frac{c}{\theta} \leq q \leq \hat{z} - \frac{\bar{\delta}c}{\theta})} \int_{\hat{z} - \frac{c}{\theta}}^{\hat{z} - \frac{\bar{\delta}c}{\theta}} xf(x)dx$$

The expected proportion of complaints is then a function of the distribution of quality, which in turns depend on the type and investment decision of the firm. Denote $\psi = \frac{c}{\theta}$. For a good firm which invests $q \sim U[0, 1]$. The probabilities in the above expressions, become:

$$P_{I=H}(q \leq \hat{z} - \psi) = \hat{z} - \psi$$

$$P_{q_H}(\hat{z} - \psi \leq q \leq \hat{z} - \bar{\delta}\psi) = \psi(1 - \bar{\delta})$$

$$P_{q_H}(\hat{z} - \bar{\delta}\psi \leq q \leq \hat{z} - \frac{\bar{\delta}\psi}{2}) = \frac{\bar{\delta}\psi}{2}$$

$$\mathbb{E}_{q_H}(q | \hat{z} - \psi < q < \hat{z} - \bar{\delta}\psi) = \hat{z} - \frac{c}{2\theta}(1 + \bar{\delta})$$

$$\mathbb{E}_{q_H}[\sigma^*(q; \hat{z}_t) | \hat{z} - \psi \leq q \leq \hat{z} - \bar{\delta}\psi] = \frac{1 + \bar{\delta}}{2}$$

All the above probabilities are equal than or greater than zero if $\tau > 4\psi - 1$. Substituting these results in (A.1), the expected proportion of complaints when the good type of the firm invests becomes $\mathbb{E}_{q_H}(\sigma^*(q; \hat{z})) = \hat{z} - \frac{\psi}{2}$.

When the good type of the firm does not invest, or when the firm is bad, $q \sim U(0, 1/2)$. The above probabilities in this case are:

$$P_{q_L}(q \leq \hat{z} - \psi) = 2(\hat{z} - \psi)$$

$$P_{q_L}(\hat{z} - \psi \leq q \leq \hat{z} - \bar{\delta}\psi) = 2\psi(1 - \bar{\delta})$$

$$P_{q_L}(\hat{z} - \bar{\delta}\psi \leq q \leq \hat{z} - \frac{\bar{\delta}\psi}{2}) = \bar{\delta}\psi$$

$$\mathbb{E}_{q_L}(q | \hat{z} - \psi \leq q \leq \hat{z} - \bar{\delta}\psi) = \hat{z} - \frac{\psi}{2}(1 + \bar{\delta})$$

$$\mathbb{E}_{q_L}[\sigma^*(q; \hat{z}) | \hat{z} - \psi \leq q \leq \hat{z} - \bar{\delta}\psi] = \frac{1 + \bar{\delta}}{2}$$

The expected proportion of complaints when a good firm does not invest or when the firm is of the bad type is $\mathbb{E}_{q_L}(\sigma^*(q; \hat{z})) = 2\hat{z} - \psi$.

$\psi = \frac{c}{\theta} \geq \frac{1}{4}$, guarantees that all the above probabilities are greater than zero. Then, the condition in Propositions 1 and 2 is sufficient but not necessary.

A.2 Bayesian Updating

When consumers expect the good type of the firm to invest, a low realisation of quality increases the evidence in favour of the firm being bad. Bayesian updating implies:

$$\begin{aligned}
\tau_1 = P(G|q_1 < \frac{1}{2}) &= \frac{P(q_1 < \frac{1}{2}|G)P(G)}{P(q_1 < \frac{1}{2}|G)P(G) + P(q_1 < \frac{1}{2}|B)P(B)} \\
&= \frac{\frac{\tau}{2}}{\frac{\tau}{2} - (1 - \tau)} \\
&= \frac{\tau}{2 - \tau}
\end{aligned} \tag{A.2}$$

If consumers do not expect the good type of the firm to invest, a low realisation of quality does not give them any additional information about the firm's type and so there is no updating of beliefs.

In both cases, a high (expected or unexpected) realisation of quality is fully revealing of the firm type, and implies $\tau_1 = 1$.

A.3 Equilibria of the Repeated Game when $T \rightarrow$

∞

This Appendix studies how the results of Section 2.4 extend to the case in which the number of periods goes to infinity. It is shown that, even when $T \rightarrow \infty$, a regulation based on customers complaints induces an equilibrium with positive probability of investment as long as the firm is not extremely patient, consumers' prior about the firm being good is relatively high and the punishment is harsh enough. The informativeness of customers complaints is also analysed. The main results of the Appendix are summarised in the following Proposition:

Proposition 7. *Given $\hat{z}^* < 1/2$ and $\beta \in (0, 1)$ there exists τ^{***} such that if $\tau_1 > \tau^{***}$ a high quality and low quality equilibria coexist for all T . Customers complaints become less informative as T increases.*

In order to prove this Proposition, I first show the existence of a high quality

(separating) and a low quality (pooling) equilibrium. The following Lemma will be useful in proving the existence of those equilibria. It shows that, once consumers know for sure the firm is good, they expect the firm to invest in every period and the firm's best reply is to invest.¹

Lemma 9. *Given $\hat{z}_t = \hat{z}_H(1)$ and $\hat{z}^* < 1/2$, there exists a unique subgame perfect equilibrium in which $I_t = 1$ for all t .*

Proof. Let $\tau_t = 1$. The firm's best reply in the one-shot game is $I_t = H$ (see sub-section 2.2.2). Furthermore, $\tau_t = 1$ implies that $\tau_k = 1$ for any $k \geq t$. Then, the game becomes a repetition of the one shot game in which consumers expect to receive $\hat{z}_{H,k}(1)$ in every period $k \geq t$, and thus the firm optimally invests in any period $k \geq t$.

Consider the firm's incentives to deviate to L in a period $s \geq t$. Because consumers' already know the type of the firm, the deviation does not affect the level of quality consumers' expect to receive in any period after s , and thus it does not affect the firm's future profits. However, the deviation reduces the firm's current profits as it induces a lower quality realisation (on expectations). Then, given $\tau_t=1$, there is no profitable one-period deviation.

□

Claim 5. Investment in Equilibrium. *Given $\hat{z}^* \leq 1/2$ and $\beta \in (0, 1)$, there exists τ^{***} such that there is an equilibrium in which the firm invests in every period t in which τ_t is above τ^{***} and it does not invest otherwise.*

Proof. Denote by $s(\underline{\tau})$ a firm's investment strategy in which:

$$I = \begin{cases} H & \text{if } \tau \geq \underline{\tau} \\ L & \text{otherwise} \end{cases}$$

¹See Definition 2.

The level of quality consumers expect to receive when the firm follows this strategy is:

$$\hat{z} = \begin{cases} \hat{z}_H(\tau) & \text{if } \tau \geq \underline{\tau} \\ \hat{z}_L & \text{otherwise} \end{cases}$$

The value for the firm of following the strategy $s(\underline{\tau})$ depends only on τ . For $\tau < \underline{\tau}$, this value is:

$$V_{s(\underline{\tau})}(\tau) = \frac{\pi_L(\hat{z}_L, \sigma^*)}{1 - \beta} \quad (\text{A.3})$$

while for $\tau \geq \underline{\tau}$ this value is:

$$\begin{aligned} V_{s(\underline{\tau})}(\tau) &= \pi_H(\hat{z}_H(\tau), \sigma^*) + \beta \left[\frac{\pi_H(\hat{z}_H(1), \sigma^*)}{2(1 - \beta)} + \frac{1}{2} \Pr(\tau' \geq \underline{\tau}) V_{s(\underline{\tau})}(\tau') + \right. \\ &\quad \left. + \frac{1}{2} \Pr(\tau' < \underline{\tau}) \frac{\pi_L(\hat{z}_L, \sigma^*)}{(1 - \beta)} \right] \end{aligned} \quad (\text{A.4})$$

where $\Pr(\tau' > \underline{\tau}) = 1$ if consumers' beliefs about the type of the firm after a low quality realisation is greater than $\underline{\tau}$: $\tau' = \frac{\tau}{2 - \tau} \geq \underline{\tau}$.

There exists a unique function, $V_{s(\underline{\tau})}^*(\tau)$, that solves (A.4). Consider the set $C([\underline{\tau}, 1])$ of bounded, continuous and weakly decreasing functions of τ , and define

an operator T on $C([\underline{\tau}, 1])$ as:

$$\begin{aligned} TV_{s(\underline{\tau})}(\tau) = & \pi_H(\hat{z}_H(\tau), \sigma^*) + \beta \left[\frac{\pi_H(\hat{z}_H(1), \sigma^*)}{2(1-\beta)} + \frac{1}{2} \Pr(\tau' \geq \underline{\tau}) V_{s(\underline{\tau})}(\tau') + \right. \\ & \left. + \frac{1}{2} \Pr(\tau' < \underline{\tau}) \frac{\pi_L(\hat{z}_L, \sigma^*)}{(1-\beta)} \right] \end{aligned} \quad (\text{A.5})$$

Standard arguments can be used to show that a function $V \in C([\underline{\tau}, 1])$, solves (A.4) if and only if it is a fixed point of T . The (expected) fine in any period t is a weakly increasing function of consumers' expected quality and so the firm's current profits are weakly decreasing in τ .² Furthermore, the per period return function is bounded and continuous in τ .³ Then, $TV_{s(\underline{\tau})}(\tau)$ is also bounded, continuous and weakly decreasing in τ , and so T maps $C([\underline{\tau}, 1])$ into itself. It is straightforward to show that Blackwell's sufficient conditions of monotonicity and discounting apply to the operator T . Hence, T is a contraction mapping and so it has a unique fixed point, $V_{s(\underline{\tau})}^*(\tau)$,⁴ which is a bounded, continuous and weakly decreasing function of τ .

If the firm deviates in the first period, but it attaches to the original strategy thereafter, its expected payoff is:

$$V_{s(\underline{\tau})}^d(\tau) = \begin{cases} \pi_L(\hat{z}_H(\tau), \sigma^*) + \beta \left[\Pr(\tau' > \underline{\tau}) V_{s(\underline{\tau})}^*(\tau') + \Pr(\tau' \leq \underline{\tau}) \frac{\pi_L(\hat{z}_L, \sigma^*)}{(1-\beta)} \right] & \text{if } \tau \geq \underline{\tau} \\ \pi_H(\hat{z}_L, \sigma^*) + \frac{\beta}{2(1-\beta)} [\pi_L(\hat{z}_L, \sigma^*) + \pi_H(\hat{z}_H(1), \sigma^*)] & \text{otherwise} \end{cases} \quad (\text{A.6})$$

Denote by $f_{s(\underline{\tau})}(\tau)$ the difference between the firm's expected payoffs when it follows strategy $s(\underline{\tau})$ and when it deviates. When $\tau \geq \underline{\tau}$, $f_{s(\underline{\tau})}(\tau)$ is a bounded and continuous function of τ because both $V_{s(\underline{\tau})}^*(\tau)$ and $V_d(\tau)$ are bounded and continuous. Furthermore, for all $\tau \geq \underline{\tau}$, $f_{s(\underline{\tau})}(\tau)$ is weakly increasing in τ and it

²For any $I \in \{L, H\}$, $\pi_I(\hat{z}_I(\tau), \sigma^*)$ is a weakly decreasing function of τ (strictly decreasing for $\tau < 1$), and $\pi_I(\hat{z}_I, \sigma^*)$ is constant in τ .

³ $V_{s(\underline{\tau})}(\tau') = \frac{\pi_L(\hat{z}_L, \sigma^*)}{(1-\beta)}$ if $\tau = \underline{\tau}$. Then, $V_{s(\underline{\tau})}(\tau)$ is continuous in $\underline{\tau}$.

⁴By the Contraction Mapping Theorem.

is given by:

$$\begin{aligned}
f_{s(\underline{\tau})}(\tau) &= V_{s(\underline{\tau})}^*(\tau) - V_{s(\underline{\tau})}^d(\tau) \\
&= -h + mp\left(\hat{z}_H(\tau) - \frac{\psi}{2}\right) + \beta \left[\frac{\pi_H(\hat{z}_H(1), \sigma^*)}{2(1-\beta)} - \right. \\
&\quad \left. - \frac{1}{2} \Pr\left(\frac{\tau}{2-\tau} > \underline{\tau}\right) V_{s(\underline{\tau})}^*(\tau) - \frac{1}{2} \Pr\left(\frac{\tau}{2-\tau} \leq \underline{\tau}\right) \frac{\pi_L(\hat{z}_L, \sigma^*)}{1-\beta} \right]
\end{aligned} \tag{A.7}$$

The first term of (A.7) is an increasing function of τ because $\hat{z}_H(\tau)$ is increasing in τ . The second term is also weakly increasing in τ . When $\tau > \frac{\tau}{2-\tau} > \underline{\tau}$, $\Pr\left(\frac{\tau}{2-\tau} > \underline{\tau}\right) = 1$ and the second term is weakly increasing in τ because $V_{s(\underline{\tau})}^*(\tau)$ is weakly decreasing in τ . If τ is such that $\tau \geq \underline{\tau} \geq \frac{\tau}{2-\tau}$, $\Pr\left(\frac{\tau}{2-\tau} \leq \underline{\tau}\right) = 1$ and the second term of (A.7) becomes independent of τ . Recall that when $\tau' < \underline{\tau}$, $V_{s(\underline{\tau})}^*(\tau') = \frac{\Pi(L, \hat{z}_L)}{2(1-\beta)}$. Then, equation (A.7) can be written as:

$$f_{s(\underline{\tau})}(\tau) = -h + mp\left(\hat{z}_H(\tau) - \frac{\psi}{2}\right) + \beta \left[\frac{\pi_H(\hat{z}_H(1), \sigma^*)}{2(1-\beta)} - \frac{1}{2} V_{s(\underline{\tau})}^*(\tau') \right]$$

The strategy $s(\underline{\tau})$ constitutes an equilibrium strategy if and only if $f_{s(\underline{\tau})}(\tau) \geq 0$ for every $\tau \in [0, 1]$. Since $f_{s(\underline{\tau})}(\tau)$ is weakly increasing in τ for every $\tau \geq \underline{\tau}$, it suffices to show that (1) $f_{s(\underline{\tau})}(\underline{\tau}) \geq 0$ and (2) $f_{s(\underline{\tau})}(\tau) > 0$ for every $\tau < \underline{\tau}$. Consider the first case:

$$f_{s(\underline{\tau})}(\underline{\tau}) = -h + mp\left(\hat{z}_H(\underline{\tau}) - \frac{\psi}{2}\right) - \frac{\beta}{2(1-\beta)} [\pi_H(\hat{z}_H(1), \sigma^*) - \pi_L(\hat{z}_L, \sigma^*)]$$

which is positive as long as $\hat{z}_H(\underline{\tau}) \geq \hat{z}^*(1 + \frac{\beta}{2(1-\beta)})$. Denote by τ^{***} the value of τ for which $f_{s(\underline{\tau})}(\underline{\tau}) = 0$ -i.e., the τ for which $\hat{z}_H(\tau) = \hat{z}^*(1 + \frac{\beta}{2(1-\beta)})$. Since $\hat{z}_H(\tau)$

is increasing in τ , $\hat{z}_H(\tau) \geq \hat{z}^*(1 + \frac{\beta}{2(1-\beta)})$ for all $\tau \geq \tau^{***}$.

Let $\underline{\tau} \geq \tau^{***}$; given strategy $s(\underline{\tau})$, the firm does not have a profitable deviation. Suppose, to the contrary, that for some $\underline{\tau} \geq \tau^{***}$, $f_{s(\underline{\tau})}(\underline{\tau}) < 0$; then $\hat{z}_H(\underline{\tau}) < \hat{z}^*(1 + \frac{\beta}{2(1-\beta)})$, and so $\underline{\tau} < \tau^{***}$ which is a contradiction. Hence (1) holds.

Finally, consider (2). If $\tau < \underline{\tau}$, $f_{s(\underline{\tau})}(\tau)$ is defined as the difference between (A.3) and the second line of (A.6), and is equal to:

$$f_{s(\underline{\tau})}(\tau) = h - mp \left(\hat{z}_L - \frac{\psi}{2} \right) + \frac{\beta}{2(1-\beta)} \left(h + mp \frac{\psi}{2} \right)$$

which is positive as long as $\hat{z}_L \leq \hat{z}^*(1 + \frac{\beta}{2(1-\beta)})$.⁵ Furthermore, this condition is independent of τ and so it implies that for every $\tau \in [0, \underline{\tau})$, $f_{s(\underline{\tau})}(\tau) \geq 0$ and so the firm is not willing to deviate from strategy $s(\underline{\tau})$.

Thus, given any $\underline{\tau} \geq \tau^{***}$, there exists an equilibrium in which the firm follows strategy $s(\underline{\tau})$, and the Claim is proved. \square

Claim 6. Pooling equilibrium *Given $\tau \in [0, 1)$, there exists an equilibrium in which $I_t^* = 0$, for all t .*

Proof. In a pooling equilibrium the good type of the firm does not invest and consumers anticipate this behaviour; therefore, quality realisations below $1/2$ are not informative about the firm's type. On equilibrium path, consumers' beliefs about the type of the firm are:

$$\tau_t = \begin{cases} \tau_0 & \text{if } q_k \leq 1/2 \text{ for every } k \leq t-1 \\ 1 & \text{if } q_k > 1/2 \text{ for some } k \leq t-1 \end{cases}$$

where τ_0 is the probability consumers assign to the firm being good at the beginning of the game. Denote by $s(1)$ an strategy in which the firm invests only

⁵Recall that $\hat{z}^* = \frac{h}{mp} + \frac{\psi}{2}$.

when $\tau = 1$. The firm's value from following this strategy is:

$$V_{s(1)}(\tau) = \frac{\pi_L(\hat{z}_L, \sigma^*)}{1 - \beta} = \frac{p - mp\mathbb{E}_{q_L}(\sigma^*(q; \hat{z}_L))}{1 - \beta} \quad (\text{A.8})$$

The firm's incentives to deviate depend on the value of \hat{z}^* . Consider the case in which $\hat{z}^* > 1/4$. In this case, the firm has neither short run nor long run incentives to deviate from the non-investing strategy. If in any period t consumers expect to receive low quality (\hat{z}_L) and the firm deviates, there exists a positive probability that a high quality realisation reduces the (expected) proportion of complaints and the expected fine. However, this reduction in the fine does not compensate the cost of investment (h) and so period- t profits are reduced too, implying that the firm has a short run incentive to fulfil consumers' low quality expectations.⁶ Long run considerations also prevent the firm from deviating. If the firm invests in period t there exists a positive probability that consumers do not find out the deviation ($q_t < 1/2$); in this case the firm's current profits are smaller and future profits remain unchanged because $\tau_{t+1} = \tau_t = \tau_0$. There is also a positive probability that a high quality realisation reveals the firm's type: if $q_t > 1/2$, $\tau_{t+1} = 1$ and $\hat{z}_k = \hat{z}_H(1) \forall k > t$. As shown in Lemma 9, the firm's best response is to invest in every period after k , which induces a continuation value smaller than $V_{s(1)}(\tau)$.⁷ Therefore, when $\hat{z} > 1/4$ and $\tau < 1$ there exists a pooling equilibrium in which the firm never invests and consumers expect low quality.

When $\hat{z} < 1/4$, the firm has short run incentives to invest. In this case,

⁶Recall from Section 2.3 that when $\hat{z}^* > 1/4$, low quality is an equilibrium of the stage game -i.e., the stage game has an equilibrium in which $I = 0$ and consumers' expect low quality (\hat{z}_L).

⁷The firm's expected profits in a low quality equilibrium of the stage game are $\pi_L(\hat{z}_L, \sigma^*) = p - mp(2\hat{z}_L - \psi)$, while its profits in a high quality equilibrium if $\tau = 1$ are $\pi_H(\hat{z}_H(\tau), \sigma^*) = p - h - mp(\hat{z}_H(1) - \frac{\psi}{2})$. There exists no value of $\hat{z}^* > 0$ for which the latter profits are higher than the former ones. As there is no updating of beliefs, the expected profits of the repeated game are $\frac{\pi_L(\hat{z}_L, \sigma^*)}{1 - \beta} > \frac{\pi_H(\hat{z}_H(\tau), \sigma^*)}{1 - \beta}$.

the punishment is harsh relative to the cost of investment and so current profits increase if the firm deviates. The same as before, if the firm deviates in period t there is a positive probability that consumers find out its type (and hence expect high quality from $t + 1$ onwards) and a positive probability that they do not detect the firm's deviation. The value for the firm if it deviates from $s(1)$ in the current period, but follows it since tomorrow onwards is:

$$V_{s(1)}^d(\tau) = \pi_H(\hat{z}_L, \sigma^*) + \frac{\beta}{2} \left[\frac{\pi_L(\hat{z}_L, \sigma^*)}{1 - \beta} + \frac{\pi_H(\hat{z}_H(1), \sigma^*)}{1 - \beta} \right] \quad (\text{A.9})$$

The firm will attach to the non-investment strategy as long as the continuation value in (A.8) is greater than the one above. A necessary and sufficient condition is:

$$\hat{z}_L \leq \hat{z}^* \left[1 + \frac{\beta}{2(1 - \beta)} \right] \quad (\text{A.10})$$

□

Claim 7. *Given $\hat{z}^* < 1/2$ and $\beta \in (0, 1)$, there exist $\tau \in (\tau^*, \tau^{***})$ such that the game has a HQE and a LQE when $T = 1$ but there exists a unique LQE when $T \rightarrow \infty$. As a result, complaints are less informative when $T \rightarrow \infty$.*

Proof. From Proposition 1, there exist $\tau^* = 4\hat{z}^* - 1$ such that for every $\tau \geq \tau^*$ the one shot game has a high quality equilibrium. Analogously, from Claim 5, there exist $\tau^{***} = 4\hat{z}^*(1 + \frac{\beta}{2(1-\beta)}) - 1$ such that for $\tau > \tau^{***}$ and $T \rightarrow \infty$ there exists an equilibrium in which the firm invests in every period. As $\frac{\beta}{2(1-\beta)} > 0$, $\tau^{***} \geq \tau^*$, and so for $\tau \in (\tau^*, \tau^{***})$ the firm invests in equilibrium when $T = 1$ but it does not when $T \rightarrow \infty$.

The same as when the game is repeated an infinite number of times, the set of $\tau \in (\tau^*, \tau^{***})$ corresponds to values of the parameters for which complaints are informative in the one shot game. When $T = 1$, complaints are informative if $1 - 2\psi \geq \tau \geq 4\hat{z}^* - 1$, which implies $1 - 2\psi \geq \tau \geq \tau^*$. In each period of the

infinitely repeated game this condition becomes $1 - 2\psi \geq \tau_t \geq \tau^{***}$. The set of values of τ for which the second condition holds is smaller because $\tau^{***} \geq \tau^*$. \square

Appendix B

Appendix to Chapter 3

B.1 Proofs

Claim 1

For any rule \mathcal{R} , good reviews increase the probability consumers assign to the value of the good being high: $\lambda'_G(\lambda; \mathcal{R}) \geq \lambda \iff \int_{\bar{q}}^{q^K} f_H(q) dq \geq \int_{\bar{q}}^{q^K} f_L(q) dq$. This is true because increasing likelihood ratio implies first order stochastic dominance.¹

Analogously, $\lambda'_B(\lambda; \mathcal{R}) \leq \lambda \iff \int_0^{\bar{q}} f_H(q) dq \leq \int_0^{\bar{q}} f_L(q) dq$, which is also implied by first order stochastic dominance.

Finally, observing that the previous consumer bought the good but he did not complete a review is bad news about v if $\lambda'_N(\lambda; \mathcal{R}) \leq \lambda$, or $\int_{\underline{q}}^{\bar{q}} f_H(q) dq \leq \int_{\underline{q}}^{\bar{q}} f_L(q) dq$. Whether this inequality holds or not depends on the quality realisation at which $f_H(q) = f_L(q)$.² To see this, denote by \hat{q} the crossing point of the distributions and consider two extreme cases:

¹First order stochastic dominance implies $F_H(q) \leq F_L(q)$ for every $q \in [q_0, q_K]$.

²As $F_H(q)$ dominates $F_L(q)$ in terms of the likelihood ratio, $f_H(q)$ and $f_L(q)$ cross only once.

- If $\hat{q} \leq \underline{q}$, then $f_H(\underline{q}) \geq f_L(\underline{q})$. Increasing MLRP implies that the ratio $f_H(q)/f_L(q)$ is increasing in q , then $f_H(\underline{q}) \geq f_L(\underline{q})$ implies $f_H(q) \geq f_L(q)$ for every $q \in (\underline{q}, \bar{q})$ and so $\int_{\underline{q}}^{\bar{q}} f_H(q) dq \geq \int_{\underline{q}}^{\bar{q}} f_L(q) dq$ and $\lambda'_N(\lambda; \mathcal{R}) \geq \lambda$.
- On the other extreme, consider the case in which $\hat{q} \geq \bar{q}$. In this case, $f_H(\bar{q}) \leq f_L(\bar{q})$ and by increasing MLRP, this implies $f_H(q) \leq f_L(q)$ for every $q \in (q, \bar{q})$. As a result, $\lambda'_N(\lambda; \mathcal{R}) \leq \lambda$.

As a result, no reviews are bad news when the crossing of the quality distributions is very high relative to the upper bound of the social norm, but they become more and more good news the closer is the crossing point to the lower bound of the social norm.

Claim 2

$\lambda'_G(\lambda; \mathcal{R}) \geq \lambda'_N(\lambda; \mathcal{R})$ for every social norm if and only if $\frac{\int_{\underline{q}}^{\bar{q}} f_L(q) dq}{\int_{\underline{q}}^{\bar{q}} f_H(q) dq} \geq \frac{\int_{\bar{q}}^{q^K} f_L(q) dq}{\int_{\bar{q}}^{q^K} f_H(q) dq}$ for every \mathcal{R} . This expression can be written as:

$$\frac{1 - F_H(\bar{q})}{F_H(\bar{q}) - F_H(\underline{q})} \geq \frac{1 - F_L(\bar{q})}{F_L(\bar{q}) - F_L(\underline{q})} \quad (\text{B.1})$$

The inequality holds because first order stochastic dominance implies that the numerator of the left hand side is greater than that of the right hand side, while $F_H(\bar{q}) - F_H(\underline{q}) \leq F_L(\bar{q}) - F_L(\underline{q})$.

A similar argument can be used to show that $\lambda'_N(\lambda; \mathcal{R}) \geq \lambda'_B(\lambda; \mathcal{R})$. When the previous consumer bough the good, observing no reviews results in a higher posterior than observing a bad review if and only if:

$$\frac{F_H(\bar{q}) - F_H(\underline{q})}{F_H(\underline{q})} \geq \frac{F_L(\bar{q}) - F_L(\underline{q})}{F_L(\underline{q})} \quad (\text{B.2})$$

for every rule $\{\bar{q}, q\}$. To see that this inequality holds, note that both numerators

are increasing functions of \bar{q} , but the left hand side increases at a rate $f_H(\bar{q})$ while the right hand side increases at a lower rate $f_L(\bar{q})$.³ When $\bar{q} \rightarrow \underline{q}$, condition (B.2) becomes $\frac{F_L(\underline{q})}{f_L(\underline{q})} \geq \frac{F_H(\underline{q})}{f_H(\underline{q})}$, which holds because increasing monotone likelihood property implies reverse hazard rate dominance.⁴ As \bar{q} increases, the denominator of the left hand side of (B.2) increases faster than that of the right hand side and so, given any $\underline{q} > q^0$, condition (B.2) holds for any $\bar{q} \in (\underline{q}, q^K]$.

Claim 3

$$\frac{\partial \lambda'_G(\lambda; \underline{q}, \bar{q})}{\partial \bar{q}} = \frac{\lambda(1-\lambda)[f_L(\bar{q}) \int_{\bar{q}}^{q^K} f_H(q) dq - f_H(\bar{q}) \int_{\bar{q}}^{q^K} f_L(q) dq]}{[\lambda \int_{\bar{q}}^{q^K} f_H(q) dq + (1-\lambda) \int_{\bar{q}}^{q^K} f_L(q) dq]^2} \geq 0 \quad (\text{B.3})$$

The denominator is the probability of observing a good review squared, so it is positive. The sign of the numerator depends on the sign of $[f_L(\bar{q}) \int_{\bar{q}}^{q^K} f_H(q) dq - f_H(\bar{q}) \int_{\bar{q}}^{q^K} f_L(q) dq]$, which is positive as long as the distribution of quality conditional on H dominates in hazard rate sense the one conditional on L .⁵ As hazard rate dominance is implied by increasing monotone likelihood ratio, the numerator is positive and the result in the Claim holds.

Consumers' beliefs after observing a bad review are increasing in \underline{q} :

$$\frac{\partial \lambda'_B(\lambda; \underline{q}, \bar{q})}{\partial \underline{q}} = \frac{\lambda(1-\lambda)[f_H(\underline{q}) \int_{\underline{q}}^q f_L(q) dq - f_L(\underline{q}) \int_{\underline{q}}^q f_H(q) dq]}{[\lambda \int_{\underline{q}}^q f_H(q) dq + (1-\lambda) \int_{\underline{q}}^q f_L(q) dq]^2} \geq 0 \quad (\text{B.4})$$

The denominator is the probability of observing a bad review squared, so it is positive. The numerator is positive as long as $f_H(\underline{q}) \int_{\underline{q}}^q f_L(q) dq \geq f_L(\underline{q}) \int_{\underline{q}}^q f_H(q) dq$,

³Increasing monotone likelihood implies that $f_H(q)/f_L(q)$ is an increasing function of q .

⁴Using L'Hopital's rule:

$$\lim_{\bar{q} \rightarrow \underline{q}} \frac{F_L(\bar{q}) - F_L(\underline{q})}{F_H(\bar{q}) - F_H(\underline{q})} = \frac{f_L(\underline{q})}{f_H(\underline{q})}$$

⁵Hazard rate dominance implies $f_L(\bar{q})[1 - F_H(\bar{q})] \geq f_H(\bar{q})[1 - F_L(\bar{q})]$. See ?.

which holds by inverse hazard rate dominance.⁶

A similar analysis shows that the updating after observing no reviews (when the previous consumer bought the good) is also an increasing function of the thresholds of the social rule:

$$\frac{\partial \lambda'_N(\lambda; \underline{q}, \bar{q})}{\partial \bar{q}} = \frac{\lambda(1-\lambda)[f_H(\bar{q}) \int_{\underline{q}}^{\bar{q}} f_L(q) dq - f_L(\bar{q}) \int_{\underline{q}}^{\bar{q}} f_H(q) dq]}{[\lambda \int_{\underline{q}}^{\bar{q}} f_H(q) dq + (1-\lambda) \int_{\underline{q}}^{\bar{q}} f_L(q) dq]^2} \geq 0 \quad (\text{B.5})$$

$$\frac{\partial \lambda'_N(\lambda; \underline{q}, \bar{q})}{\partial \underline{q}} = \frac{\lambda(1-\lambda)[f_L(\underline{q}) \int_{\underline{q}}^{\bar{q}} f_H(q) dq - f_H(\underline{q}) \int_{\underline{q}}^{\bar{q}} f_L(q) dq]}{[\lambda \int_{\underline{q}}^{\bar{q}} f_H(q) dq + (1-\lambda) \int_{\underline{q}}^{\bar{q}} f_L(q) dq]^2} \geq 0 \quad (\text{B.6})$$

Proposition 3

Convexity of the profit function

If the firm offers the optimal static price, $\hat{p}(\lambda'_i)$, its expected profits are $\pi(\lambda'_i) = \frac{(\lambda'_i + L - c)^2}{4(\lambda'_i + L)}$. Taking derivatives with respect to λ'_i :

$$\frac{\partial \pi(\lambda'_i)}{\partial \lambda'_i} = \frac{(\lambda'_i + L - c)(\lambda'_i + L + c)}{4(\lambda'_i + L)^2} > 0 \text{ for every } L > c$$

$$\frac{\partial^2 \pi(\lambda'_i; \mathcal{R})}{\partial \lambda'^2_i} = \frac{c^2}{4(\lambda'_i + L)^3} > 0 \text{ for every } c > 0$$

Convexity of the utility function

The expected utility in any period t , given the consumer's prior λ'_i and a price p is:

$$\begin{aligned} u(\lambda'_i; \mathcal{R}) &= P(\gamma > \frac{p}{\lambda'_i + L}) \left[\mathbb{E}(\gamma | \gamma > \frac{p}{\lambda'_i + L})(\lambda'_i + L) - p \right] \\ &= \left[\frac{\lambda'_i + L - p}{\lambda'_i + L} \right] \left[\frac{\lambda'_i + L + p}{2} - p \right] \\ &= \frac{(\lambda'_i + L - p)^2}{2(\lambda'_i + L)} \end{aligned}$$

⁶Reverse hazard rate dominance implies $f_L(\underline{q})F_H(\underline{q}) \leq f_H(\underline{q})F_L(\underline{q})$. See ?

where I used the conditional expectation of γ :

$$\mathbb{E}(\gamma|\gamma \geq \frac{p}{\lambda+L}) = \frac{1}{1 - \frac{p}{\lambda+L}} \left[\int_{\frac{p}{\lambda+L}}^1 x dx \right] = \frac{\lambda + L + p}{2(\lambda + L)}$$

Taking partial derivatives with respect to λ'_i :

$$\frac{\partial u(\lambda'_i; \mathcal{R})}{\partial \lambda^i} = \frac{(\lambda'_i + L + p)(\lambda'_i + L - p)}{2(\lambda'_i + L)^2}$$

$$\frac{\partial^2 u(\lambda'_i; \mathcal{R})}{\partial \lambda_i'^2} = \frac{p^2}{(\lambda'_i + L)^3}$$

The second expression is positive for every $p > 0$, while the first one is positive as long as $p < L + c$. Then, consumers' utility function is increasing and convex with respect to λ'_i for every positive price at which some consumer is willing to buy. In particular, it is increasing and convex in the prior when $p = \hat{p}(\lambda'_i) = \frac{\lambda'_i + L + c}{2}$. The expected utility of a consumer who observed review i , when the second period price is $\hat{p}(\lambda'_i)$ is:⁷

$$u(\lambda'_i; \mathcal{R}) = \frac{[\lambda'_i + L - c]^2}{8(\lambda'_i + L)}$$

Taking derivatives with respect to λ'_i :

$$\frac{\partial u(\lambda'_i; \mathcal{R})}{\partial \lambda^i} = \frac{(\lambda'_i + L - c)(\lambda'_i + L + c)}{8(\lambda'_i + L)^2} > 0 \text{ for every } i \in \{G, N, B\} \text{ and } L > c$$

$$\frac{\partial^2 u(\lambda'_i; \mathcal{R})}{\partial \lambda_i'^2} = \frac{c^2}{4(\lambda'_i + L)^3} > 0 \text{ for every } i \in \{G, N, B\} \text{ and } c > 0$$

Then, whichever the review completed by the previous consumer and the price

⁷Given $p = \hat{p}(\lambda'_i)$, the probability that the consumer buys the good is:

$$P(\gamma \geq \frac{\hat{p}}{\lambda'_i + L}) = \frac{\lambda'_i + L - p^*}{\lambda'_i + L} = \frac{\lambda'_i + L - c}{2(\lambda'_i + L)}$$

and the expected value of γ conditional on the consumer buying is:

$$\mathbb{E}(\gamma|\gamma \geq \frac{\hat{p}}{\lambda'_i + L}) = \frac{1}{1 - \frac{\hat{p}}{\lambda'_i + L}} \left[\int_{\frac{\hat{p}}{\lambda'_i + L}}^1 x dx \right] = \frac{3\lambda'_i + 3L + c}{4(\lambda'_i + L)}$$

set by the monopolist, the expected utility of the second consumer is increasing and convex in λ'_i .

Claim 4

The martingale property of the beliefs also holds when the previous consumer bought the good with some probability in $(0, 1)$:

$$\begin{aligned}
\mathbb{E}(\lambda_{t+1} | \lambda, \underline{q}, \bar{q}) &= P\left(\gamma \geq \frac{p}{\lambda + L}\right) \sum_{i \in \{G, N, B\}} P(i) \lambda_{t+1}^i + P\left(\gamma < \frac{p}{\lambda + L}\right) \lambda \\
&= P\left(\gamma \geq \frac{p}{\lambda + L}\right) \lambda \left[\int_{\bar{q}}^{q^K} f_H(q) dq + \int_{\underline{q}}^{\bar{q}} f_H(q) dq + \int_{q^0}^{\underline{q}} f_H(q) dq \right] + P\left(\gamma < \frac{p}{\lambda + L}\right) \lambda \\
&= P\left(\gamma \geq \frac{p}{\lambda + L}\right) \lambda + P\left(\gamma < \frac{p}{\lambda + L}\right) \lambda \\
&= \lambda
\end{aligned}$$

Lemma 6

Assume that the first period consumer bought the good and that completing reviews is not costly. Consider two alternative social rules: one that uses two messages and another one that uses three messages. Each rule determines a distribution of posterior beliefs with mean λ (because of the martingale property of beliefs). Given that consumers' payoff is convex in λ' , they prefer the rule with three messages over the one with two messages if and only if the second distribution of posterior beliefs is a mean preserving spread of the first one.⁸

Consider first a norm such that consumers complete a bad review if $q_1 \leq \hat{q}$ and a good review otherwise. Denote by λ^- the beliefs of second period consumers after observing a bad review and by λ^+ their beliefs after observing a good review. Denoting by $F(\cdot)$ the cumulative distribution of λ' induced by this norm, then: $F(\lambda^-) = P(\lambda' \leq \lambda^-) = P(q \leq \hat{q})$, $F(\lambda^+) = P(\lambda' \leq \lambda^+) = 1$. The expected value

⁸Rothschild and Stiglitz (1970).

of λ' under this rule is:

$$\begin{aligned}
\mathbb{E}(\lambda'|\lambda, \hat{q}) &= P(\lambda^-)\lambda^- + P(\lambda^+)\lambda^+ \\
&= [\lambda \int_{q_0}^{\hat{q}} f_H(q) dq + (1-\lambda) \int_{q_0}^{\hat{q}} f_L(q) dq] \frac{\lambda \int_{q_0}^{\hat{q}} f_H(q) dq}{\lambda \int_{q_0}^{\hat{q}} f_H(q) dq + (1-\lambda) \int_{q_0}^{\hat{q}} f_L(q) dq} + \\
&\quad + [\lambda \int_{\hat{q}}^{q^K} f_H(q) dq + (1-\lambda) \int_{\hat{q}}^{q^K} f_L(q) dq] \frac{\lambda \int_{\hat{q}}^{q^K} f_H(q) dq}{\lambda \int_{\hat{q}}^{q^K} f_H(q) dq + (1-\lambda) \int_{\hat{q}}^{q^K} f_L(q) dq} \\
&= \lambda [\int_{q_0}^{\hat{q}} f_H(q) dq + \int_{\hat{q}}^{q^K} f_H(q) dq] = \lambda
\end{aligned} \tag{B.7}$$

Consider an alternative rule in which first period consumers can send three different messages, G , N and B . They complete a bad review if the quality realisation was below a threshold \underline{q} , complete no reviews if $q_1 \in (\underline{q}, \bar{q})$ and they complete a good review if $q_1 \geq \bar{q}$. Denote by λ^B , λ^N and λ^G consumers beliefs after observing a bad review, no review or a good review, respectively. Denote by $H(\cdot)$ the distribution of second period beliefs induced by this rule. Then, $H(\lambda^B) = P(\lambda' \leq \lambda^B) = P(q_1 \leq \underline{q})$, $H(\lambda^N) = P(\lambda_2 \leq \lambda^N) = P(q_1 \leq \bar{q})$ and $H(\lambda^G) = P(\lambda' \leq \lambda^G) = 1$. As shown in Claim 4, $\mathbb{E}(\lambda'|\lambda, \underline{q}, \bar{q}) = \lambda$.

As both distributions of beliefs have the same mean and consumers' payoff function is convex in the beliefs, they will prefer the distribution induced by the second rule over the one induced by the first rule if the second one is a mean preserving spread of the first one or, equivalently, if $F(\cdot)$ dominates stochastically of second order $H(\cdot)$. A sufficient condition for this to be true is:

$$\int_0^{\lambda'} H(t) dt \geq \int_0^{\lambda'} F(t) dt \text{ for every } \lambda' \in [0, 1] \tag{B.8}$$

Consider the case in which $\hat{q} = \bar{q}$.⁹ In this case, $\lambda^N > \lambda^- > \lambda^B$ (see below) and $\lambda^+ = \lambda^G$. To show second order stochastic dominance it is necessary to show

⁹A similar analysis can be done by assuming any other value of $\hat{q} \in [q^0, q^K]$ and the conclusions would not change.

that condition (B.8) holds for every possible value of λ' . To start with, note that for $\lambda' < \lambda^-$ the distribution of beliefs under $H(\cdot)$ accumulates more mass than under $F(\cdot)$ because $F(\lambda' < \lambda^-) = 0$ while $H(\lambda' < \lambda^-) = H(\lambda^B) > 0$. The probability of observing a value of $\lambda' \in (\lambda^-, \lambda^N)$, on the other hand, is greater under $F(\cdot)$. However, as $P(\lambda^B)[\lambda^- - \lambda^B] = [P(\lambda^-) - P(\lambda^B)](\lambda^N - \lambda^-)$, where $P(\lambda^-) - P(\lambda^B) = P(\lambda^N)$,¹⁰ both distributions accumulate the same mass for every $\lambda' > \lambda^N$. Then:

- For $\lambda' < \lambda^-$: $\int_0^{\lambda'} H(t)dt > \int_0^{\lambda'} F(t)dt$,
- For $\lambda' < \lambda^N$, $\int_0^{\lambda'} H(t)dt > \int_0^{\lambda'} F(t)dt$.
- For $\lambda' \geq \lambda^N$: $\int_0^{\lambda'} H(t)dt = \int_0^{\lambda'} F(t)dt$.

Then, the distribution of beliefs induced by the two-messages rule second order stochastically dominates the one induced by the three-messages rule. Together with the fact that both distributions have the same mean, this implies that $H(\cdot)$ is a mean preserving spread of $F(\cdot)$ and so consumers expected payoff is greater under the last rule. A similar result can be obtained for $\hat{q} = \underline{q}$ or for any other $\hat{q} \in (0, 1)$.

The results above assume that $\lambda^N \leq \lambda^- \leq \lambda^B$. Now I show that those assumptions are correct.

$\lambda^- \leq \lambda^B \iff \frac{\int_{\underline{q}}^{\bar{q}} f_L(q)dq}{\int_{\underline{q}}^{\bar{q}} f_H(q)dq} \geq \frac{\int_{\underline{q}}^{\hat{q}} f_L(q)dq}{\int_{\underline{q}}^{\hat{q}} f_H(q)dq}$, which can be written as $\frac{F_H(\hat{q})}{F_H(\underline{q})} \geq \frac{F_L(\hat{q})}{F_L(\underline{q})}$. As we are assuming $\hat{q} = \bar{q}$, this is the same as $\frac{F_H(\bar{q})}{F_H(\underline{q})} \geq \frac{F_L(\bar{q})}{F_L(\underline{q})}$ which implies $\lambda^G(\bar{q}, \underline{q}) \geq \lambda^B(\bar{q}, \underline{q})$.

$\lambda^- \geq \lambda^N \iff \frac{F_L(\bar{q}) - F_L(\underline{q})}{F_H(\bar{q}) - F_H(\underline{q})} \geq \frac{F_L(\hat{q})}{F_H(\hat{q})}$. Using the fact that $\hat{q} = \bar{q}$, the previous condition becomes $\frac{F_L(\bar{q}) - F_L(\underline{q})}{F_H(\bar{q}) - F_H(\underline{q})} \geq \frac{F_L(\bar{q})}{F_H(\bar{q})}$, which holds because it implies

¹⁰Given $\hat{q} = \bar{q}$, $P(\lambda^-) = P(\lambda^G)$, and $P(\lambda^G) - P(\lambda^B) = P(\lambda^N)$ by construction.

$$\lambda^G(\bar{q}, \underline{q}) \geq \lambda^N(\bar{q}, \underline{q}).$$

Lemma 8

The result in Lemma 8 is an immediate implication of the results in the next two claims:

Claim 8. *For every $\lambda \in (0, 1)$, $\bar{q} \in (q^0, q^K]$ there exists $\underline{h}(\lambda, \bar{q})$ such that $\frac{\partial V(\mathcal{R}, \lambda)}{\partial \underline{q}}|_{\underline{q}=q^0} - h \frac{\partial P(B; \mathcal{R}, \lambda)}{\partial \underline{q}}|_{\underline{q}=q^0} > 0$ for all $0 < h \leq \underline{h}(\lambda, \bar{q})$*

Claim 9. *For every $\lambda \in (0, 1)$, $\underline{q} \in [q^0, q^K)$ there exists $\bar{h}(\lambda, \underline{q})$ such that $\frac{\partial V(\mathcal{R}, \lambda)}{\partial \bar{q}}|_{\bar{q}=q^K} - h \frac{\partial P(G; \mathcal{R}, \lambda)}{\partial \bar{q}}|_{\bar{q}=q^K} < 0$ for all $0 < h \leq \bar{h}(\lambda, \underline{q})$.*

To see that the result in Claim 8 holds, consider what happens when $\underline{q} \rightarrow q^0$. From Lemma 6, $\frac{\partial V(\mathcal{R}, \lambda)}{\partial \underline{q}}|_{\underline{q}=q^0} > 0$ for every $\lambda \in (0, 1)$ because when the cost of completing reviews is not taken into account, consumers' expected utility is greater when the reviewing rule uses the three available messages. On the other hand, $\frac{\partial P(B; \mathcal{R}, \lambda)}{\partial \underline{q}}|_{\underline{q}=q^0} = \lambda f_H(q^0) + (1 - \lambda) f_L(q^0)$. If the distribution of quality realisations is such that $f_H(q^0) = f_L(q^0) = 0$, then $\frac{\partial P(B; \mathcal{R}, \lambda)}{\partial \underline{q}}|_{\underline{q}=q^0} = 0$ and the result in the Claim holds. If the distribution of quality realisations has fat tails and $f_H(q^0) > 0$ and $f_L(q^0) > 0$, then the result in the Claim holds as long as there exists $h > 0$ such that

$$\underline{h}(\lambda, \bar{q}) \leq \frac{\frac{\partial V(\mathcal{R}, \lambda)}{\partial \underline{q}}|_{\underline{q}=q^0}}{\frac{\partial P(B; \mathcal{R}, \lambda)}{\partial \underline{q}}|_{\underline{q}=q^0}} = \frac{\frac{\partial V(\mathcal{R}, \lambda)}{\partial \underline{q}}|_{\underline{q}=q^0}}{\lambda f_H(q^0) + (1 - \lambda) f_L(q^0)} \quad (\text{B.9})$$

which holds because both, the numerator and the denominator are positive for every $\lambda \in (0, 1)$ and for every $\bar{q} \in (q^0, q^K]$.

An analogous argument can be used to prove Claim 9. In this case the condition for consumers to prefer $\bar{q} < q^K$ is that $\frac{\partial V(\mathcal{R}, \lambda)}{\partial \bar{q}}|_{\bar{q}=q^K} - h \frac{\partial P(G; \mathcal{R}, \lambda)}{\partial \bar{q}}|_{\bar{q}=q^K} < 0$.

From Lemma 6, as \bar{q} moves away from q^K , consumers' expected utility $V(\cdot)$ increases; as a result, $\frac{\partial V(\mathcal{R}, \lambda)}{\partial \bar{q}}|_{\bar{q}=q^K} < 0$. Furthermore, $\frac{\partial P(G; \mathcal{R}, \lambda)}{\partial \bar{q}}|_{\bar{q}=q^K} = -[\lambda f_H(q^K) + (1 - \lambda)f_L(q^K)] \leq 0$, with strict inequality if the distribution of quality realisations has “fat tails”. If $f_H(q^K) = f_L(q^K) = 0$, the condition in the claim holds for every $\lambda \in (0, 1)$ and for every $h > 0$. Otherwise, if the distribution of quality realisations assigns positive probability to the tails, the condition in the claim becomes:

$$\bar{h}(\lambda, \underline{q}) \leq -\frac{\frac{\partial V(\mathcal{R}, \lambda)}{\partial \bar{q}}|_{\bar{q}=q^K}}{\frac{\partial P(G; \mathcal{R}, \lambda)}{\partial \bar{q}}|_{\bar{q}=q^K}} = -\frac{\frac{\partial V(\mathcal{R}, \lambda)}{\partial \bar{q}}|_{\bar{q}=q^K}}{\lambda f_H(q^K) + (1 - \lambda)f_L(q^K)} \quad (\text{B.10})$$

which is positive because, as mentioned above, the numerator is negative and the denominator is positive.

B.2 Monotone Likelihood Ratio Property

For any $\theta_0, \theta_1 \in \Theta$ such that $\theta_1 \geq \theta_0$, MLRP implies $\frac{1-F(\bar{q}|\theta_1)}{F(\bar{q}|\theta_1)} \geq \frac{1-F(\bar{q}|\theta_0)}{F(\bar{q}|\theta_0)}$. The claim below shows why this is the case.

Claim 10. *$f(\cdot|\theta)$ satisfies MLRP, then for any $\theta_0, \theta_1 \in \Theta$ such that $\theta_1 \geq \theta_0$ we have $\frac{1-F(\bar{q}|\theta_1)}{F(\bar{q}|\theta_1)} \geq \frac{1-F(\bar{q}|\theta_0)}{F(\bar{q}|\theta_0)}$ and $\frac{1-F(\underline{q}|\theta_1)}{F(\underline{q}|\theta_1)} \geq \frac{F(\underline{q}|\theta_0)}{F(\underline{q}|\theta_1)}$.*

Proof. The family of densities $f(\cdot|\theta)$ satisfies the monotone likelihood ratio property if for all $q_1 \geq q_0$ and $\theta_1 \geq \theta_0$ we have:

$$f(q_1|\theta_1)f(q_0|\theta_0) \geq f(q_0|\theta_1)f(q_1|\theta_0) \quad (\text{B.11})$$

Integrating both sides of this expression over q_0 from 0 (lower bound of q) to q_1 :

$$\int_0^{q_1} f(q_1|\theta_1)f(q_0|\theta_0)dq_0 \geq \int_0^{q_1} f(q_0|\theta_1)f(q_1|\theta_0)dq_0$$

$$f(q_1|\theta_1)F(q_1|\theta_0) \geq F(q_1|\theta_1)f(q_1|\theta_0)$$

Let $a = q_1$, we have:

$$\frac{f(a|\theta_1)}{f(a|\theta_0)} \geq \frac{F(a|\theta_1)}{F(a|\theta_0)} \quad (\text{B.12})$$

Integrating both sides of (B.11) with respect to q_1 , from q_0 to 1 (upper bound of q):

$$\int_{q_0}^1 f(q_1|\theta_1)f(q_0|\theta_0)dq_1 \geq \int_{q_0}^1 f(q_0|\theta_1)f(q_1|\theta_0)dq_1$$

Let $a = q_0$, we have:

$$\frac{1 - F(a|\theta_1)}{1 - F(a|\theta_0)} \geq \frac{f(a|\theta_1)}{f(a|\theta_0)} \quad (\text{B.13})$$

Combining (B.12) and (B.13), we have:

$$\frac{1 - F(a|\theta_1)}{1 - F(a|\theta_0)} \geq \frac{F(a|\theta_1)}{F(a|\theta_0)}$$

This result holds for any $a \in [0, 1]$. In particular: $\frac{1-F(\bar{q}|\theta_1)}{1-F(\bar{q}|\theta_0)} \geq \frac{F(\bar{q}|\theta_1)}{F(\bar{q}|\theta_0)}$ and $\frac{1-F(\underline{q}|\theta_1)}{1-F(\underline{q}|\theta_0)} \geq \frac{F(\underline{q}|\theta_1)}{F(\underline{q}|\theta_0)}$.

□

Bibliography

- AHN, I., AND M. SOUMINEN (2005): “Word-of-Mouth Communication and Community Enforcement,” *Journal of Political Economy*, 113(5), 1146–62.
- AKERLOF, G., AND W. T. DICKENS (1982): “The Economic Consequences of Cognitive Dissonance,” *American Economic Review*, 3(72), 307–19.
- AKERLOF, R. (2010): “Punishment, Compliance and Anger in Equilibrium,” *Unpublished - Job Market Paper*.
- BANERJEE, A., AND D. FUNDENBERG (2004): “Worth-of-mouth Learning,” *Games and Economic Behavior*, (46), 1–22.
- BANERJEE, A. V. (1992): “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, (107), 797–817.
- BATTIGALLI, P., AND M. DUFWENBERG (2009): “Dynamic Psychological Games,” *Journal of Economic Theory*, (144), 1–35.
- BESANKO, D., S. DONNENFELD, AND L. WHITE (1987): “Monopoly and Quality Distortion: Effects and Remedies,” *The Quarterly Journal of Economics*, 104(4), 743–768.
- BIKHCHANDANI, A. S., D. HIRSHLEIFER, AND I. WELCH (1992): “A Theory of Fads, Fashion, Custom and Cultural Change as Information Cascades,” *Journal of Political Economy*, (100), 992–1026.

- BOSE, S., G. OROSEL, M. OTTAVIANI, AND L. VESTERLUND (2006): “Dynamic Monopoly Pricing and Herding,” *The RAND Journal of Economics*, 37(4), 910–28.
- (2008): “Monopoly Pricing in the Binary Herding Model,” *Economic Theory*, 37(2), 203–41.
- BOULDING, W., A. KALRA, A. STAELIN, AND V. ZEITHAML (1993): “A Dynamic Process Model of Service Quality: From Expectations to Behavioral Intentions,” *Journal of Marketing Research*, 30(1), 7–27.
- BRUNNERMEIER, M., AND J. PARKER (2005): “Optimal Expectations,” *American Economic Review*, 4(95), 1092–118.
- CHEVALIER, J., AND D. MAYZLIN (2006): “The Effect of Worth of Mouth on Sales: Online book Reviews,” *Journal of Marketing Research*, 43(3), 345–54.
- COATE, S., AND M. CONLIN (2004): “A Group rule: Utilitarian Approach to Voter Turnout. Theory and Evidence,” *The American Economic Review*, 94(5), 1476–1504.
- CRAWFORD, V. P., AND J. SOBEL (1982): “Strategic Information Transmission,” *Econometrica*, 50(6), 1431–51.
- DOWNS, A. (1957): *An Economic Theory of Democracy*. New York: Harper and Row.
- DOYLE, C., AND M. WATERSON (2012): “Your call: eBay and demand for the iPhone 4,” *International Journal of the Economics of Business (forthcoming)*.
- ELLISON, G., AND D. FUDENBERG (1995): “Word-of-Mouth Communication and Social Learning,” *The Quarterly Journal of Economics*, 110(1), 93–125.
- EUROPEAN UNION (2011): “Monitoring Consumer Markets in the European Union,” Discussion paper.

- FEDDERSEN, T., AND A. SANDRONI (2006a): “The Calculus of Ethical Voting,” *International Journal of Game Theory*, (35), 1–35.
- (2006b): “A theory of Participation in Elections,” *The American Economic Review*, 96(4), 1271–1282.
- FORBES, S. (2008): “The Effect of Service Quality and Expectations on Customer Complaints,” *Journal of Industrial Economics*, 56(1), 190–213.
- GEANAKOPOLOS, J., D. PEARCE, AND E. STACCHETTI (1989): “Psychological Games and Sequential Rationality,” *Games and Economic Behavior*, (1), 160–79.
- GILL, D., AND D. SGROI (2008): “Sequential Decisions with Tests,” *Games and economic Behavior*, (63), 663–78.
- (2012): “The optimal choice of Pre-Launch Reviewer,” *Journal of Economic Theory*, 147(3), 1247–1260.
- HARSANYI, J. (1977): “Morality and the Theory of Rational Behavior,” *Social Research*, 44(4), 623–56.
- (1980): “Rule Utilitarianism, Rights, Obligations and the Theory of Rational Behaviour,” *Theory and Decision*, (12), 115–133.
- (1992): “Game and Decision Theoretic Models in Ethics,” in *Handbook of Game Theory with Economic Applications*, ed. by R. Aumann, and S. Hart, vol. I, pp. 669–707. Amsterdam: Elsevier Science, North-Holland.
- KAHNEMAN, D., AND A. TVERSKY (2001): “Loss Aversion in Riskless Choice: A reference-Dependent Model,” *Quarterly Journal of Economics Working Paper*, (106), 1039–61.
- KOSZEGI, R., AND M. RABIN (2006): “A model of reference-Dependent preferences,” *Quarterly Journal of Economics*, 4(121), 1133–66.

- KREMER, I., Y. MANSOUR, AND M. PERRY (2012): “Implementing the Wisdom of the Crowd,” *Working Paper*.
- LAFFONT, J., AND J. TIROLE (1993): *A Theory of Incentives in Procurement and Regulation*. Cambridge: MIT Press.
- LERNER, J., AND J. TIROLE (2006): “A Model of Forum Shopping,” *American Economic Review*, (96), 1091–1113.
- LUCA, M. (2011): “Reviews, Reputation and Revenue: The case of Yelp.com,” *Harvard Business School, Working Paper 12-016*, (146), 479–503.
- MALIATH, G., AND L. SAMUELSON (2006): *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press.
- MILGROM, P. R. (1981): “Good News and Bad News: Representation theorems and applications,” *The Bell Journal of Economics*, 12(2), 380–91.
- MUSSA, M., AND S. ROSEN (1978): “Monopoly and Product Quality,” *Journal of Economic Theory*, (18), 301–317.
- OFGEM (2008): “Complaints Handling Standards,” Discussion paper.
- OLIVER, R. (1977): “Effect of Expectations and Disconfirmation on Post Exposure Product Evaluation: An alternative Interpretation,” *Journal of Applied Psychology*, (62), 480–486.
- (1980): “A Cognitive Model of the Antecedents and Consequences of Satisfaction Decisions,” *Journal of Marketing Research*, (17), 460–469.
- OSBORNE, M. (2004): *An Introduction to Game Theory*. Oxford University Press.
- REISS, P., AND M. WHITE (2008): “What changes energy consumption? Prices and public pressures,” *RAND Journal of Economics*, 39(3), 636–663.

- RESNICK, ZECKHAUSER, SWANSON, AND LOCKWOOD (2006): “The value of reputation on eBay: A controlled Experiment,” *Experimental Economics*, 9(2), 79–101.
- RESNICK, P., AND R. ZECKHAUSER (2002): “Trust among strangers in Internet Transactions: Empirical Analysis of eBay’s Reputation System. In M.E. Baye (ed.),” *The Economics of Internet and E-Cpmmerce*, pp. 127–57.
- ROTHSCHILD, M., AND J. E. STIGLITZ (1970): “Increasing Risk I: A Definition,” *Journal of Economic Theory*, 2(3), 225–43.
- SAPPINGTON, D. (2005): “Regulating Service Quality: A Survey,” *Journal of Regulatory Economics*, 27(2), 123–154.
- SHAPIRO, C. (1982): “Information, Product Quality and Seller Reputation,” *The Bell Journal of Economics*, 13(1), 20–35.
- SHEINSKI, E. (1982): “Price, Quality and Quantity Regulation in Monopoly Situations,” *Economica, New Series*, 43(170), 127–137.
- SINGH, J. (1988): “Consumer Complaint Intentions and Behaviour: Definitional and Taxonomical Issues,” *The Journal of Marketing*, 52(1), 93–107.
- SMITH, L., AND P. SORENSEN (2000): “Pathological Outcomes of Observational Learning,” *Econometrica*, 68(2), 371–98.
- SPENCE, M. (1975): “Monopoly, Quality and Regulation,” *Bell Journal of Economics*, 6(2), 417–429.
- THE GUARDIAN (2012): “Amazon deluged with complaints over Yodel’s poor delivery service,” .
- ZITHAML, V., L. BERRY, AND A. PARASURAMAN (1996): “The Behavioral Consequences of Service Quality,” *Journal of Marketing*, 60(2), 31–46.