Design of Online Reputation Systems:

An Economic Perspective

by

Florian Bersier

St Antony's College

University of Oxford

A thesis submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

Oxford Internet Institute

Trinity Term 2014

ABSTRACT

Design of Online Reputation Systems: An Economic Perspective

by Florian Bersier

St Antony's College, University of Oxford, Trinity Term 2014

- DPhil in Information, Communication and the Social Sciences -

---

Online reputation systems are certainly the most overlooked 'heroes' of today's social Web. While these mechanisms are a vital element of every online transaction, they have received less consideration than some of their more well-known cousins, such as recommender systems or social networks, whose success would often not have been possible and tenable without their discrete but active backing. It then follows that despite their value and importance, the implementation of current reputation mechanisms has mostly been the result of trial-and-error. Resting on an economic perspective, this thesis regroups three chapters whose frameworks and findings aim at helping mechanism designers and researchers understand key mechanisms at play and develop more efficient online reputation systems.

The first chapter examines the optimal number of ratings a reputation mechanism must make publicly available within an online marketplace in order to minimize cheating and maximize Pareto efficiency. I develop a moral hazard stage game featuring fictitious players which has the compelling property to prevent reputation effects from disappearing in the long run. I show that the number of ratings displayed by a reputation system is a fundamental predictor of market efficiency, and that the latter number should be kept minimal in order to maximize social welfare in the market – especially for economies proposing interactions with a high profit margin.

The second chapter studies how different classes of reporting behaviours commonly found online affect the reliability of a reputation mechanism. I develop an iterative stochastic approximation model which I use to construct a behavioural measure of efficiency, so-called 'reporting bias'. I demonstrate that reporting bias tends towards its maximum when raters comply with the reports left by their predecessors. Following this result, I recommend to keep the rating interface separated from the rest of the reputation system. I also find that fake ratings are particularly harmful when one type of behaviour is present in the economy and suggest to counterbalance sybil attacks by displaying pairs of contrasted ratings. Finally, I defend the use of the arithmetic mean against the median as a way to compute reputation scores.

The third chapter analyses how 5-star rating scales can lead to the formation of bimodal distributions of ratings within online marketplaces. Using a 2-time period model featuring altruistic raters, I identify the existence of a 'blind spot' of unrated transactions whose magnitude increases in the cost of rating and decreases in the number of buyers inhabiting the economy. Developing an additional model featuring Bayesian agents suffering from confirmatory bias, I show that non-binary rating scales can leave space to ambiguity and possibly wrong posteriors, even in the long run. Overall, results of the chapter hint that fine-grained rating scales best suit signalling reputation systems while coarse-grained scales should be preferred for sanctioning mechanisms. ◊

Acknowledgements

Completing a DPhil thesis is a challenging and life-changing adventure that would not be possible without the help, kindness and support of many different people.

My first debt of gratitude goes to my two supervisors, Dr Greg Taylor and Dr Nir Vulkan. I am especially grateful to Greg whose patience, guidance and assistance have been crucial in proceeding through the doctoral program and completing my dissertation.

It has been a great privilege to spend a few years at the Oxford Internet Institute both as a DPhil student and a Teaching Assistant. Its members will always remain dear to me. My time at Oxford was made enjoyable in large part due to the time spent at OII North, surrounded by other DPhil fellows and MSc students.

I gratefully acknowledge the funding sources that made my DPhil work possible. I was funded by a generous Scatcherd European Scholarship and part of my research was supported by the Economic and Social Research Council (ESRC).

My friends in the UK, France, Switzerland and other parts of the world were constant sources of laughter, joy, and support. Special thanks go to Angèle, Aurelia, Christos, Dani, Eva, Gabrielle, Isis, Kiril, Lucy, Mariana, Sophie, Themis and members of the St Antony's Boat Club, University of Oxford.

Finally, I would like to say a heartfelt and enormous thank you to my parents, Béatrice and Marcel, and my sister Bérénice. Their love, support and encouragement were the driving force that allowed me to complete this DPhil thesis. *Merci pour votre Amour et soutien.*

Contents

List of Figures

List of Tables

Chapter 1

General Introduction

1.1 Motivation

The invention and democratisation of the Web has dramatically increased the diversity and quantity of Internet-powered interactions. Now, online interactions feature predominant idiosyncrasies that represent a perfect fertilizer to maintain a large 'garden' full of Lemons à la Akerlof (1970), i.e. they are usually of 1-shot type[1], asynchronous, geographically dispersed and anonymous. It follows that interacting with anonymous and geographically dispersed online traders multiplies contracting difficulties, by making extremely difficult and costly the identification of a seller or the verification of a product's quality. Likewise, online transactions being asynchronous, a buyer may need to make a payment before she receives the product. Finally, a seller who knows he will likely not meet the same buyer in the future has incentives to misbehave since he is aware that his actions will not be known by others in the marketplace.

To limit such a natural information asymmetry and encourage cooperation, *reputation mechanisms*, also known as reputation systems, have emerged as an increasingly important component of online transactions. Nowadays, we decide to get a good or service online because we have a belief that it will correspond to our expectations. For instance, we choose to bid on an item on eBay[2] (Melnik and Alm, 2002) because the seller has a favourable feedback score or we book a table in a restaurant because of its positive ratings on Yelp[3] (Luca, 2011). Reputation systems with their large databases of transaction histories have thus emerged as a core feature of online economies, promoting trustworthy behaviours and market efficiency.

However, despite their key importance in today's Web, the practical design of current reputation systems results more from trial and error than explicit re-

[1] Resnick and Zeckhauser (2002) find that the majority of eBay trading encounters are one-shot and state that an eBay transaction is best thought of as a single interaction.

[2] eBay.com is an e-commerce website.

[3] Yelp.com is a website where users can rate local businesses such as restaurants, bars, etc.

search. Following this observation, the present work attempts to rectify – within reasonable limits – the situation by adopting an economic perspective. Specifically, this thesis proposes to examine the mechanisms and determinants necessary to the very formation of reputation in online economies.

1.2 Reputation and Reputation Systems

Before going any further, it appears relevant to provide a formal definition of what is 'reputation'.

Definition 1 (Reputation)**.** *Reputation is defined as a set of beliefs held about someone or something within a specific context.*

This definition is proposed and considered within the thesis because it clearly and concisely encapsulates the two key parameters of reputation: 'belief' and 'context'. Specifically, an economic agent, let us call him John, can enjoy a good reputation as a seller on eBay, but a bad reputation as a host on Airbnb.[4] John's reputation then depends on what other agents believe he can or cannot do (or be) within a specific context. For instance, an agent named Alice can have read that John always ships very quickly high quality products on eBay, but can have also read that John is not a friendly person and that his flat is not always tidy and clean. To form this sort of beliefs about John in online economies such as eBay and Airbnb, Alice must read John's ratings on the latter two platforms through their respective reputation systems.

Definition 2 (Reputation System)**.** *A reputation system is a mechanism that aggregates information about one's past interactions in order to facilitate the process of assessing his reputation within a specific economy.*

[4]Airbnb.com is a website where people can rent out lodging.

The primary objective of reputation systems is to promote efficient transactions in economies where cooperation is made difficult because of information asymmetries. Such asymmetries are generally divided into two (pure) categories, that is *adverse selection* and *moral hazard*. Moral hazard parametrizes an environment where Alice expects John to *do* something, and so implies hidden actions, while adverse selection represents a situation where Alice believes John to *be* something, and then implies hidden information (Cabral, 2005).

As already highlighted by Dellarocas (2005), those two types of asymmetries should not be seen as equivalent by mechanism designers since they induce a different processing of information. More particularly, a signalling mechanism can be defined as follows:

Definition 3 (Signalling Reputation System)**.** *A signalling reputation system is a mechanism that aggregates information about one's past interactions in order to signal his quality or trustworthiness within a specific economy.*

A possible illustration of a signalling reputation system is the rating mechanism used by Amazon,[5] an online marketplace where products (e.g., books) have a fixed quality which is privately known and is 'discovered' little by little thanks to the ratings published by past buyers (i.e., adverse selection). A compelling property of online signalling mechanisms is that they principally accumulate data about goods and services, while sanctioning mechanisms – whose definition follows – usually deal with individuals.

Definition 4 (Sanctioning Reputation System)**.** *A sanctioning reputation system is a mechanism that aggregates information about one's past interactions in order to sanction his non-cooperativeness within a specific economy.*

A well-known example of sanctioning mechanism is the eBay's Feedback Forum, where buyers can rate the past behaviour or trustworthiness of sellers,

[5]Amazon.com is an online retailer.

knowing that the latter sellers may have incentives to behave differently from one period to the other (i.e. moral hazard). The key mechanisms being defined, the next section proposes a concise summary of the related literature.

1.3 Background

The literature about online reputation systems can be roughly divided into three streams, i.e. (*i*) a first strand, which mainly rests on game theory and experiments, analyses the interplay between reputation and cooperation, (*ii*) a second strand, empirically-driven and focusing (for a large part) on eBay, examines the positive correlation between the selling price and a seller's reputation and (*iii*) a third and more recent strand, coming from Computing Science, focuses on the design of the mechanisms.

Cooperation From bootstrap models to Bayesian settings, economists have proposed a few different theoretical approaches to examine how reputation could enforce cooperation. One can cite the works by Milgrom and Roberts (1982), Kreps and Wilson (1982), Holmström (1999), Mailath and Samuelson (2001) and Ekmekci (2011). Besides, a certain number of experimental studies have allowed to better understand the effects of different reputation mechanisms on cooperation, e.g., Bolton et al. (2005) and Bolton et al. (2013).

Reputation Premium A considerable amount of empirical research has identified and studied the existence of a reputation premium, that is the fact that sellers with a higher reputation transact more often and for higher prices. Major works in this area come from Resnick and Zeckhauser (2002), Melnik and Alm (2002), Cabral and Hortaçsu (2010) and Diekmann et al. (2014).

Mechanism Design A different perspective on reputation which mainly comes from Computing Science, with the notable exception of Dellarocas (2005), has more recently focused on the design of online reputation mechanisms and ways of making them more trustworthy. Here, one can cite the Ph.D. theses by Mui (2003) and Jurca (2007). Resting on a systematic or an algorithmic approach, this strand of the literature is interested in proposing methods to treat and arrange existing data (i.e., ratings) in order to maximize social welfare in online markets.

The perspective followed by this thesis is closer to the first and third strands but adopts an *ex-ante* viewpoint to analyse the formation of reputation online. More particularly, this work differs from the related literature and its *ex-post* approach by focusing on the role played by the very design of a mechanism in the decisions of economic agents, from the interaction decisions to the rating ones.

1.4 Roadmap and Contributions

The contributions of this thesis are twofold. Firstly, this work presents novel methods to examine the formation of reputation within online markets. Secondly, it proposes several policy recommendations regarding the design of online reputation mechanisms. The thesis is organized as follows.

Chapter 2 examines the optimal number of ratings a reputation mechanism must display in order to maximize Pareto efficiency. It develops a repeated stage game played between one long-lived player and an infinite number of short-lived fictitious players. The analysis shows that the number of ratings made available within an online market represents a fundamental predictor of market efficiency since it determines the occurrence of cheating in the game and the equilibrium (stationary) strategy of the long-lived player. A so-called 'ratings trade-off' is identified, that is an environment where a mechanism designer faces two possible choices: either display a large number of ratings and ensure a high level of sellers' participation at the risk of encouraging low effort interactions, or display a small number of ratings and ensure a high level of effort at the risk of discouraging sellers to enter the market and form reputation. Welfare analysis shows that market efficiency suffers from the presence of noisy reporting behaviours, implying that a mechanism designer needs to increase the number of ratings displayed by the reputation system accordingly to the noise if he wants to guarantee market efficiency.

Chapter 3 analyses the effects of different reporting behaviours on the trustworthiness of a reputation system. A behavioural measure of efficiency, so-called 'reporting bias', is built following the development of an iterative stochastic approximation setting. This chapter first identifies that the usefulness of new ratings decays over time by demonstrating how reputation converges to a given score as the amount of ratings increases in the economy. A converging phenomenon is also found for the very distribution of ratings, result that presents some theoretical support for the empirically observed clustered and bimodal distribution of ratings online. Moreover, the analysis shows that when raters comply with the reports left by their predecessors, it becomes extremely difficult to guarantee the efficiency of a reputation mechanism, thus suggesting that the rating interface should be kept separated from the rest of the reputation system in order to encourage truthful reporting. Similarly, reputation manipulation is found to be especially dangerous when a herding reporting behaviour is present in an online marketplace, hinting that fake ratings should be countered and corrected by displaying pairs of contrasted ratings. The chapter ends by showing that the arithmetic mean represents a more stable and intuitive way of computing reputation scores than the median.

Chapter 4 builds upon recent empirical works to propose two models aiming at formalizing the determinants and mechanisms which can lead to U-shaped distributions of ratings. This part of the thesis presents one of the first theoretical attempts to formalize the decision of a buyer to publish a rating within an online economy. The development of a 2-time period model featuring altruistic raters allows to identify the existence of a so-called 'blind spot' of unrated transactions whose size increases in the cost of rating and decreases in altruistic concerns. A second model featuring Bayesian agents suffering from confirmatory bias is built and demonstrates that non-binary rating scales can leave space to ambiguity and possibly wrong posteriors, even in the long run. The overall results of the chapter suggest that fine-grained rating scales best suit signalling reputation systems while coarse-grained scales should be preferred for sanctioning mechanisms.

Chapter 2

Number of Ratings

How Many Ratings Should An Online Reputation System Display?

2.1 Introduction

In this chapter, I build a reputation model which takes into account the specificities of online interactions and I examine the optimal number of ratings a reputation system must disclose in order to maximize social welfare. I show that Pareto efficiency is promoted and better secured when a mechanism displays a small number of ratings. I also demonstrate that the very structure of the reputation system determines the equilibrium strategies of online sellers.

Reputation mechanisms present in online economies differ from their offline counterparts in the way information is more formally aggregated and structured. After every online transaction, a buyer may leave a positive or a negative rating depending on the outcome of the transaction, and this rating participates in forming the identity and history of the seller, which is made publicly available to all users in the economy. Due to the large size of online markets, the probability for a seller to meet the same buyer twice is extremely low, implying that a potential buyer cannot rest her decision to trust a seller on her past experience but only on the history of past ratings attached to the identity of the seller. A buyer thus forms a belief about the trustworthiness of an online seller by observing a measure summarizing the seller's performance based on his past ratings. Another structural specificity of online reputation systems is that the maximum amount of possible ratings is fixed and limited, either for technical reasons or to keep sellers' reputation scores relevant over time.

This work proposes to develop a reputation model which is consistent with the above idiosyncrasies. In particular, I assemble an infinitely repeated 2×2 moral hazard game in which (*i*) a seller's action is saved as a rating only if a specific action has been played on the buyer's side, (*ii*) a buyer's belief that the seller is trustworthy follows from the empirical frequency of seller's past trustworthy actions (i.e. positive ratings), and (*iii*) a fixed number of ratings is saved to the game's memory.

This chapter contributes to the reputation literature in a few different respects. Its first natural contribution stems from the theoretical approach it adopts to analyse online reputation mechanisms. To the best of my knowledge, this work represents one of the rare attempts to develop a reputation model adapted to the singularities of online economies, with the notable exception of papers by Dellarocas (2005) and Ekmekci (2011). Differently from the two aforementioned authors who focused on eBay-like economies, I propose a setting which features the elegant property to be generalizable to a large variety of online marketplaces, independently of their allocation mechanisms and goods types. Such a proceeding allows to derive some new normative recommendations regarding the design of online reputation systems. Specifically, I show that the number of ratings becomes a key predictor of market efficiency as the profit margin increases in an online marketplace and identify the existence of a so-called 'ratings trade-off', that is a low number of ratings tends to promote the quality of interactions over the quantity, while a high number of ratings promotes the quantity of interactions over the quality. More generally and as already highlighted by Dellarocas (2005) in the case of auction-markets, I find that reputation mechanisms which publish only a small number of ratings are more effective in promoting cooperation within online marketplaces.

The chapter also contributes to the literature about learning in games by considering a setting in which an infinitely large population of short-lived fictitious buyers is randomly matched every period with a fully rational long-lived seller. Making use of fictitious play provides a natural and intuitive basis to reproduce the way ratings are aggregated and reputation scores computed by online mechanisms. Furthermore, it allows to relax unrealistic assumptions usually made within the reputation literature about perfectly knowledgeable agents, then proposing a configuration closer to the reality of online economies where buyers often only have access to a limited set of information such as a reputation score to form their beliefs about the trustworthiness of a seller. This work thus asserts that fictitious play represents a natural and compelling learning method to use as far as online reputation models are concerned.

The chapter finally contributes to a recent strand of the literature interested in proposing new mechanisms able to prevent reputation effects from disappearing in the long run, issue formally identified by Cripps et al. (2004). In peculiar, I show that reputation formation in the game can resist time effects in the long run and that such a dynamics can be achieved without incomplete information, departing so from classical models of reputation à la Kreps and Wilson (1982) or Milgrom and Roberts (1982). This result directly follows from the observation that the structure of the reputation system determines the seller's equilibrium strategy in the game. Such findings are in line with the late literature interested in highlighting the cyclic property of reputation by bounding the memory of short-lived players or making more realistic assumptions about the (limited) rationality of economic agents. Works by Mell (2011), Ekmekci (2011), Liu (2011), Liu and Skrzypacz (2014) represent interesting and promising illustrations of this trend.[1]

This work builds on that of several previous authors. The stage game is inspired by the seminal setting presented by Fudenberg and Levine (1989), and supposes a single long-lived player who plays a simultaneous-move moral hazard game against an infinite sequence of short-lived players who play only once. As in Ellison (1997), I depart from the standard game theory literature by examining interactions between a fully rational long-lived player and a large population of evolutionary game theory's fictitious players. In particular, the behaviour

[1]Ekmekci (2011) shows that reputation effects can be preserved by proposing a model in which information about past play of the game is censored by a reputation mechanism. Liu (2011) limits the information about past play by making it costly to acquire by short-lived players. Mell (2011) develops a setting supporting non-disappearing reputation effects by introducing behavioural short-lived players whose learning rests on an adaptive play à la Young (1993), and whose decisions to trust a long-lived player are based on the drawing of a sample of past actions from the complete history of actions taken by the long-lived player.

of the short-lived players follows a boundedly rational learning process, so that the latter players myopically best respond to the long-lived player's actions (or ratings) in the previous rounds of the game. Finally, I consider a binary reputation mechanism inspired by Dellarocas (2005), this type of rating scale representing the most natural choice given the action sets of the stage game.

Chapter outline

The rest of the chapter is organized as follows. Section 2.2 sets up the model and introduces a simple reputation system. Section 2.3 presents the general reputation dynamics and results of the game. Section 2.4 proposes three extensions of the model and investigates their effects on the core findings and social welfare. Sections 2.5 discusses the set-up and outlines the implications of the results for the design of reputation mechanisms in online marketplaces.

2.2 Model

Consider an infinitely repeated moral hazard game between one long-lived player (player 1) and an infinite sequence of short-lived players. Each short-lived player lives for one period and serves as player 2 in the following stage game. Player 1 plays in every period.[2]

2.2.1 Stage Game Play

Let $A_1 = \{H, L\}$, $A_2 = \{Y, N\}$ be the set of actions available to player 1 and player 2 in the stage game. An action $a_1 \in A_1$ represents the decision of player 1 to exert a high (H) or low (L) level of effort. The action $a_2 \in A_2$ specifies whether player 2 interacts ([Y]es) or does not interact ([N]o) with player 1, respectively.

[2]For ease of exposition, I will sometimes refer to player 1 as the seller and player 2 as the buyer, although the model is also applicable to environments (such as online help forums) that use reputation mechanisms without having buyer and seller roles.

The action sets are common knowledge and the set of pure strategy profiles is denoted $A = A_1 \times A_2$. Each player in the stage game has a von Neumann-Morgenstern payoff function $\pi_i : A \to \mathbb{R}$ whose outcome is always finite. Player 1 knows both payoff functions while player 2 only knows her own. Players move simultaneously.

		Player 2: Y	Player 2: N
Player 1	H	$p - e$, ***v-p***	$-e$, 0
	L	***p***, $-p$	**0**, **0**

Table 2.1: A general 2 x 2 stage game for an online interaction.

In the stage game, e denotes the cost of effort, p denotes the price of interacting, and v represents the value of interacting with a high effort player 1 (e.g., a seller) for player 2, where $v > p > e > 0$.

I denote B_i as the set of action profiles such that player i is best responding to the action of player $j \neq i$. Then

$$B_i = \{a \in A \;:\; a_i = \underset{a_i \in A_i}{\arg\max}\; \pi_i(a_i, a_j)\}.$$

The set of Nash equilibrium outcomes is $E = B_1 \cap B_2$, where $E \subseteq B_2 \subseteq A$. It follows that the unique Nash equilibrium of our stage game is $E = \{L, N\}$, i.e. no interaction. Moreover, I denote a profile to be the unique maximand of π_1 over B_2 when

$$a' = \underset{a \in B_2}{\arg\max}\, \pi_1(a) \quad and \quad a' \notin E.$$

This represents the Stackelberg profile, where a'_1 is the Stackelberg leader action and a'_2 the Stackelberg follower action. Note that the Stackelberg convention in the game is

$$S = \{a_1', a_2'\} = \{H, Y\},$$

that is an interaction with high effort.[3] Intuitively, a_1' is the action that player 1 would take if he were able to move first or to publicly commit to taking a particular action. This so-called Stackelberg outcome yields a higher payoff for both players, but is likely to be difficult to attain since player 1 faces a strong incentive to cheat and deviate to low effort provision. The purpose of a reputation mechanism is to help players to avoid this reversion to Nash behaviour and coordinate around the Pareto superior Stackelberg convention.

Time is discrete and indexed by $t = 0, 1, 2, ...$ In each period, player 2 (the short-lived player or "she", henceforth) is anonymously matched with player 1 (the long-lived player or "he") to play the stage game of Table 2.1.[4] Let μ_t be the probability with which a short-lived player believes that the long-lived player chooses to play the high effort action H at t. The expected payoffs from playing Y and N for the short-lived player are as follows,

$$E[\pi_2(Y) \mid \mu_t] = \mu_t(v - p) - p(1 - \mu_t) = \mu_t v - p\,,$$

$$E[\pi_2(N) \mid \mu_t] = 0\,.$$

So interacting, action Y, is a *best response* for the short-lived player provided that

$$\mu_t v - p \geq 0 \quad \Leftrightarrow \quad \mu_t \geq \frac{p}{v} \in \ (0, 1). \tag{2.1}$$

[3]Player 1 is said to be the leader in the game because (i) he is the one who proposes the transaction (ii) he needs to build and keep a "reputation" for exerting high effort.

[4]We remark that the stage game features the same type of market failure than the one found in the classical Prisoner's Dilemma case, that is the unique pure Nash Equilibrium of the game (i.e., no cooperation) does not correspond to the Pareto or Stackelberg outcome (i.e. cooperation).

Put differently, the play of Y increases in the net benefit of interacting for the short-lived player and in her belief μ that the long-lived player will exert a high level of effort. In our game, this belief (μ) results solely and directly from the only piece of information a short-lived player can get about the past actions of the long-lived player, that is a measure summarizing the long-lived player's past moves which takes the form of a rating score. Then, a short-lived player behaves naively by responding optimally to the long-lived player's ratings or reputation at the time of the interaction.

2.2.2 Reputation System

In online economies, buyers assess the probability of a seller to behave well by observing the seller's reputation score.[5] Such a score can take different forms and respect different scales but always summarizes the past average trustworthiness of a seller's interactions in the marketplace. For instance, on Amazon, sellers are evaluated by buyers on a rating scale of 1 star to 5 stars, with 5 stars being the best. The feedback rating is then calculated according to the average star rating for that particular seller. On eBay, the difference between the number of positive and negative feedback ratings along with the ratio of positive ratings over the total number of ratings earned by the seller over the last 12 months are measures used to evaluate the reliability of a seller in the marketplace.[6]

In the coming lines, I propose a reputation system which reproduces the crucial idiosyncrasies of current online reputation systems. Specifically, a rating only exists if there has been an interaction between a short-lived player (e.g., a buyer) and the long-lived player (e.g., a seller). There is a fixed number of ratings made publicly visible to the players by the reputation system. Similarly to the positive Feedback percentage used by eBay, the reputation measure is quantitatively given

[5] E.g., marketplaces such as eBay, Amazon, Ricardo, etc.
[6] See Appendix A.1 for more details about eBay's reputation system.

by the ratio of high effort actions H over the overall number of interactions played by the long-lived player that have been saved to the finite memory of the reputation system. The finite memory of current reputation systems can be explained by the cost of storing large amount of data on servers and the willingness to keep reputation scores relevant or reliable over time, e.g., eBay's positive feedback score is calculated based on the ratings earned during the last 12 months, Ricardo the last 6 months.

Existence of a rating in the game Consider a binary reputation system that makes available to all the players in the game a finite number r of ratings, where $r \in \mathbb{N}_+$. A rating exists if and only if an interaction has occurred, i.e. if a short-lived player has played the trusting action Y (as shown in Table 2.2). Then, before playing, a short-lived player perfectly observes the number of times actions H and L have been chosen by the long-lived player during the last r interactions of the game. However, a short-lived player does not know in which period of the game she is called to move because she has no information about the chronology of the latter actions, especially when the number of ratings is smaller than the number of periods the game has already been played.[7]

		Player 2 (SLP)	
		Y	N
Player 1 (LLP)	H	1	0
	L	1	0

Table 2.2: Rating existence in the game where Player 1 is the long-lived player (LLP) and Player 2 a short-lived player (SLP).

In the t-th period, the available history from which a short-lived player gets access consists of the last s ratings, where $s = min\{t, r\}$ $\forall t > 0$. Let $\Omega = A_1^s$ be

[7]Note that the history about the long-lived player's past moves is truncated because of the limited memory of the reputation system, and not because of the short-lived players' ability to remember, e.g., adaptive play à la Young (1993).

the set of possible states whose time indexing of existing elements is only known by the long-lived player and the reputation mechanism. A typical member of Ω is $\omega \in \Omega$, which is a vector of the last s actions of the long-lived player saved by the reputation system after s trusting actions Y from short-lived players. Let the state in period t be denoted ω_t.

The transition rule, $\tau : \Omega \times A \to \Omega$, is that given an initial state at time $t > 0$, $\omega_t \in \Omega$, and an action by the long-lived player at time t, $a_{1t} \in \{H, L\}$, we have

$$\omega_{t+1} = \begin{cases} \omega_t & \text{if } a_{2t} = \{N\} \\ \tau(\omega_t, a_{1t}) & \text{if } a_{2t} = \{Y\} \end{cases} \tag{2.2}$$

In words, when a short-lived player chooses the distrusting action N at t, no rating is left and the history of ratings does not change from period t to period $t+1$. Conversely, when she plays the trusting action Y, the transition rule is such that:

i. If ω_t is made of r dimensions, then the left-most element (i.e. oldest rating) in ω_t is dropped;

ii. All other ratings move one place to the left;

iii. One element is added to the end (right) of the vector ω_{t+1} taking the value of a_{1t}.

Let $h_t : \Omega \to \{1, 2, ..., r\}$ be a counting function such that $h_t(\omega_t)$ indicates the number of occasions in the vector ω_t that involves the play of high effort action H. Similarly, let l_t denote the number of plays of action L present in ω_t. More particularly, the information publicly available to short-lived players in any particular state $\omega_t \in \Omega$ is given by

$$\frac{h_t(\omega_t)}{s} \equiv \frac{h_t}{s} \,, \tag{2.3}$$

simply expressed hereafter as the ratio of h high level of effort actions H over an existing history of ratings $s \leq r$.

2.2.3 Learning and Beliefs

Following the literature about learning in games and fictitious play, I consider that the frequency of cooperative actions H played by the long-lived player in the history of ratings determines the behavioural rule of the short-lived players in the game. In a process of fictitious play, a concept first introduced by Brown (1951) and first studied by Robinson (1951), players "behave as if they think they are facing a stationary, but unknown, distribution of opponents' strategies" Fudenberg and Levine (1998). More particularly, fictitious play refers to a dynamic process where at each period of the game players play a myopic pure best response to the empirical frequency distribution of their opponent's previous actions. A compelling illustration of such a forecasting rule can be found in Ellison (1997).

Specifically, assume that a short-lived player naively believes that the probability μ with which the long-lived player chooses to play the high effort action H at period t directly follows from (2.3), i.e.,

$$\mu_t = \frac{h_t}{s} \quad \forall t > 0. \tag{2.4}$$

Each short-lived player thus uses the forecast rule (2.4) to predict the strategy of the long-lived player when she is called to play in the game.[8]

Remark that such forecasting cannot be considered as fully rational, even if the update rule computing μ can be interpreted in a Bayesian or rational sense. This comes from the fact that short-lived players assume (mistakenly) that the

[8]Note that the dynamics of our behavioural rule (or updating process) has been considered as the maximum likelihood estimation (MLE) and can be made even more rigorous by interpreting it in a Bayesian sense, using a prior belief that is a Dirichlet distribution, see for instance Rezek et al. (2008).

long-lived player is following a stationary mixed strategy, and lack the cognitive ability to anticipate how their play affects the behaviour of the long-lived player. One simple way to think about this within a Bayesian setting is that a short-lived player's a-priori belief about the long-lived player's strategy may be incorrect, even though she updates correctly from this prior. In other words, short-lived players behave rationally with respect to their beliefs, but their lack of knowledge impedes them to be considered as fully rational players in the game.[9] In practice, discrepancies in online experience and cost (in terms of time) to retrieve and read information about a seller represent intuitive and realistic rationales for the buyers' (i.e. short-lived players) forecasting rule.

The game starts at t_0 ($t = 0$) with a history of ratings $\Omega_0 = \{\emptyset\}$. Facing Ω_0, the initial history of a short-lived player is simply $h_0 = 0$ and $l_0 = 0$.[10] In the first period of the game, when the histories are empty, I suppose that a short-lived player always plays her stage game dominant action Y. The latter assumption is formalized below.

Assumption 1 (Empty Histories)**.** *A short-lived player observing Ω_0 plays the Stackelberg action Y.*

Put differently, a short-lived player is always willing to interact with a newly born long-lived player at least once, making then always possible for the long-lived player to build reputation in the game, and also ensuring the creation of a first rating in the reputation system. If this were not the case, then the game would boil down to the infinite repetition of the pure Nash equilibrium $E = \{L, N\}$ and no formation of reputation. I define the latter situation as a *locked game.*

Definition 5 (Game locking)**.** *The game is said to have become locked when the*

[9] It is worth noting that a fully knowledgeable short-lived player who has enough information to derive the long-lived player's best response to μ_t still does not know where in the state space Ω she has been called to play, and so needs to maintain a uniform prior about her possible position within the state space (uniform prior which depends on the number of ratings r).

[10] See Levin (2006) or Epelman et al. (2011) for possible illustrations of fictitious play with null histories. Informally, such a proceeding generates a fictitious state space for $t \in [0, r)$.

only strategy profile played is the one that leads to no interaction and no further reputation formation, equivalent to the pure Nash equilibrium $E = \{L, N\}$*.*

Put otherwise, once the game is locked, a short-lived player never chooses to play the trusting action Y and no farther reputation is formed, given that a rating only exists if the latter action has been chosen by a short-lived player. In practical terms, a locked game corresponds to a situation in which an agent has accumulated an irredeemably bad reputation and is forced to leave the market.

Strategies Short-lived players put equal weight on every past available play (rating) and myopically best-respond to the long-lived player's prior behaviour as summarized by his reputation score.[11] Following the forecasting rule (2.4), a short-lived player then chooses her action to maximize her payoff $\pi_2(\cdot)$,

$$B_2 = \{a : \ a_{2t} = \underset{a_2 \in A_2}{\arg\max} \ \pi_2(a_2, \mu_t)\} \ . \tag{2.5}$$

More particularly, given (2.5) and the interaction condition (2.1), a short-lived player chooses a pure strategy at $t > 0$ according to the following rule

$$a_{2t} = \begin{cases} \{Y\} & \text{if } \mu_t \geq \frac{p}{v} \\ \{N\} & \text{if } \mu_t < \frac{p}{v} \end{cases} \tag{2.6}$$

The long-lived player aims at maximizing his lifetime payoff Π_1 while discounting time at a rate $\delta \in (0, 1)$.[12] In a slight abuse of notation, let a_1 denote the path of play for the long-lived player over the length of the game,

[11]Given our learning algorithm, this myopic behaviour is optimal and players only choose pure actions at each period of the game. This is true even if they actually have in mind a mixed best response. See Hofbauer and Sandholm (2002).

[12]The time-discount factor δ parametrizes how patient the long-lived player is, where patience increases in δ. It can be understood as the probability that the game will end after each period.

$a_1 = \{a_{1,t=0}, a_{1,t=1}, a_{1,t=2}, ...\}$. Then, a_1 generates an expected lifetime payoff,

$$\Pi_1(a_1) = \sum_{t=0}^{\infty} \delta^t \pi_1(a_{1t}, a_{2t}) \ . \tag{2.7}$$

Note that while a short-lived player's strategy at period t depends on the state ω_t, the long-lived player's strategy at t depends both on ω_t and on the number of dimensions the state space can take, i.e. the number of ratings r made available by the reputation system. Specifically, the long-lived player is a sophisticated player and is so aware that his action at t may lock the game in a future period. He forms then his strategy accordingly.

2.3 A General Online Reputation Game

This section presents the general results of the reputation game. The analysis starts by examining when and how the number of ratings r displayed by a reputation mechanism can affect the frequency which the Nash action is played in the game. This proceeding allows us to study next how the long-lived player's equilibrium strategy depends on r and helps identify the existence of a trade-off that online marketplaces may face when designing their reputation systems. That is, a high number of ratings tends to favour the quantity of interactions by making the entry in the game more accessible to a less patient long-lived player while a low number of ratings promotes the quality of interactions, by deterring the play of the Nash action. I then move to the study of the long-lived player's strategy profile and demonstrate its stationarity as well as its very dependence on r. The section ends by a standard social welfare analysis. This final stage formalizes and sums the core findings of the chapter by showing that the number of ratings displayed by a reputation system becomes an essential predictor of market efficiency as the net benefit of interacting increases in the game, thus encouraging a social planner to display a minimum number of them.

2.3.1 Cheating on Reputation

In order to characterise the dynamics of reputation formation, it is necessary to understand when the long-lived player finds it worthwhile to cheat. I begin, therefore, with a result that provides a necessary and sufficient condition for the seller to be able to cheat without any consequences, scenario described as *costless cheating* henceforth.

Definition 6 (Costless cheating)**.** *The play by the long-lived player of a low level effort action L (i.e., Nash action) in the game that does not directly lead to a locking of the game is defined as costless cheating.*

Remark that if costless cheating is possible at $t > 0$, then the long-lived player never delays the play of a Nash action. Indeed, let us denote $\pi_{1t}(a)$ the long-lived player's payoff associated to an action $a \in A_1$ at period t. Following the stage game presented in Table 2.1, we have $\pi_{1t}(L) > \pi_{1t}(H)$ $\forall t$. Since $\delta \in (0, 1)$, it is direct to remark that the long-lived player always chooses to play the low effort action L as soon as he can in order to maximize his lifetime payoff if he knows the latter decision will not lock the game in the next period. By Assumption 1, we know that costless cheating is never possible (by definition) at $t = 0$ since such a strategy would imply a locking of the game at $t = 1$. Then, the opportunity (if any) for the long-lived player to cheat without risking any locking may only occur at $t > 0$.

Besides, in the case of a very patient long-lived player (i.e., $\delta \to 1$), delaying the play of Nash action is not a sustainable strategy since the number of ratings made available by the reputation system is fixed and smaller than the long-lived player's infinite lifetime horizon. Put otherwise, by waiting until later to play L, a very patient long-lived player breaks the interaction condition and locks the game. For instance, suppose a long-lived player can costlessly cheat a maximum number of x times within an history of r ratings. At the r-th period we must have $\mu = \frac{r-x}{r} \geq \frac{p}{v}$ in order to avoid any locking of the game. Now, consider the long-lived player decides to defer the play of 1 of the x L actions he could have played during the last r periods to the $(r+1)$-th period of his life. By doing so, it is straightforward to remark that he locks the game, i.e. $\mu = \frac{(r-1)-x}{r} < \frac{p}{v}$.

Knowing that the long-lived player never delays the play of a Nash action when cheating on reputation is possible, I now examine the existence condition for costless cheating within the stage game and its relationship with the number of ratings (r) displayed by the reputation system.

Proposition 1. *If $r \geq \frac{v}{v-p}$, then there exists the possibility of costless cheating in the game.*

Proof. Let h^* (l^*, respectively) be the smallest (largest) number of incidents of the play of H (L) in the available history of ratings the long-lived player can have without running any risk of reputational failure, i.e. being locked with no interaction for the remaining plays of the game. For $t > 0$, we know that the Stackelberg action Y would be a best response for the short-lived player provided that she holds a belief $\mu \geq \frac{p}{v}$. Let $t \geq r$. Following (2.3), we can write

$$\begin{aligned} \frac{h}{r} &\geq \frac{p}{v}\,, \\ h &\geq \frac{p\,r}{v} \quad \Leftrightarrow \quad l \leq l^* = r(1-\frac{p}{v})\,. \end{aligned} \tag{2.8}$$

Given (2.8), $l^* \geq 1$ if and only if $r \geq \frac{v}{v-p}$. Now let $0 < t < r$, then (2.8) can be rewritten as $l \leq r - \frac{pt}{v}$. It follows that $l^* \geq 1$ if and only if $r \geq \frac{v+pt}{v}$, which is verified for $t \geq \frac{v}{v-p}$. This completes the proof. □

Put differently, the number of times the long-lived player can safely deviate from the Stackelberg convention S increases in the number of ratings, r, available in the reputation system and in the net benefit of interacting for the short-lived player, $\pi_2 = v - p$.

The intuition is relatively simple. A larger number of ratings available in the reputation system implies that relatively more low level effort actions are tolerated by a short-lived player, given that her readiness to interact depends directly and solely on the ratio of high effort actions over the total number of ratings, or even more intuitively, if the observable history is long then a few instances of cheating are a relatively small blip in a long history of otherwise good behaviour. Similarly, a short-lived player is more prone to tolerate a relatively larger number of low-level effort actions within the history of ratings of a long-lived seller when she has 'less' to lose or, equivalently, 'more' to win by choosing to interact. More formally, consider the two limit cases for p. If $p \to 0$, then $\frac{v}{v-p} \to 1$ and costless cheating is always possible in the game. Conversely, if $p \to v$, then $\frac{v}{v-p} \to \infty$ and costless cheating is never possible in the game.

Finally, note that cheating on reputation is independent of the effort cost, e, incurred by the long-lived player. Again, the intuition is simple. A reputation system is directed towards short-lived players and so provides information regarding the payoff a short-lived player can expect by choosing to interact with the long-lived player.[13]

2.3.2 Reputation Formation

Proposition 1 established the relationship existing between the emergence of costless cheating and the number of ratings displayed by the reputation mechanism. This part of the analysis focuses on the formation of reputation in the game. Particular attention is devoted to the role played by r on the long-lived player's equilibrium strategy.

[13] Our stage game is such that a short-lived player only knows her own payoff.

2.3.2.1 Existence condition

Recall that Assumption 1 ensures that a short-lived player is always willing to interact with a newly born long-lived player in the game. Then, Lemma 7 follows.

Lemma 7. *There is formation of reputation in the game if and only if the long-lived player plays the Stackelberg action in the first period of his life.*

Proof. The proof is straightforward and comes from the observation that playing the Nash action, low level of effort L, at the first period of the long-lived player's life (t_0) will lock the game forever and no reputation will be formed given that the posterior will be $\mu_1(L) = 0 < \frac{p}{v}$ where $v > p > 0$. Then, there exists formation of reputation in the game if and only if the long-lived player starts his life by playing the Stackelberg action, high level of effort H, at t_0 which leads to $\mu_1(H) = 1 > \frac{p}{v}$. □

Lemma 7 identifies a condition for reputation formation that is consistent with the literature. Specifically, it states that the effects of reputation start to work immediately and are stronger during the initial stage of the game, when the long-lived player must 'work' hard to establish a reputation in the market. An interesting parallel can be drawn here between our setting and the one developed by Holmström (1999) in his seminal paper about career concerns. That is, an employee (long-lived player, respectively) seeks to maximize his lifetime wages (payoff) whilst minimizing his efforts, knowing that employers (short-lived players) cannot observe his ability but only the outcome of his past actions. Given this framework, Holmström shows that an employee has incentives at equilibrium to work hard at the very beginning of his career in order to build a good reputation and demonstrate his ability early in his life.

At this stage of the development and before going any further, it seems relevant to illustrate and epitomise the key concepts and dynamics we have just seen with a short example.

Example Fix $p = 1$, $v = 3$ and $r = 5$. For simplicity's sake, consider the case of an almost extremely patient long-lived player so that $\delta \to 1$. The game starts at $t = 0$. By Proposition 7, we know that the long-lived player plays H while a short-lived player plays her dominant strategy Y (see Assumption 1). Following Proposition 1, we remark that costless cheating is possible since $\frac{v}{v-p} = \frac{3}{2} < 5$. Then, the long-lived plays L at $t = 1$ while a short-lived player plays Y since $\mu_1 > \frac{1}{3}$. The same reasoning applies for $t = 2$. At $t = 3$, the long-lived player must play the Stackelberg action H if he does not want to lock the game in the next period, since playing the Nash action L would result in a belief $\mu_4 = \frac{1}{4} < \frac{1}{3}$. At the end of $t = 4$, the state space is composed of 5 ratings and, for $t > 4$, the belief μ_t then rests on the 5 latest actions of the long-lived player composing ω_t. Playing in the game is summarized in Table 2.3.

Time, t	Belief, μ_t	(h_t, l_t)	Play	Rating
0	–	$(0,0)$	(H,Y)	H
1	1	$(1,0)$	(L,Y)	L
2	$\frac{1}{2}$	$(1,1)$	(L,Y)	L
3	$\frac{1}{3}$	$(1,2)$	(H,Y)	H
4	$\frac{2}{4}$	$(2,2)$	(L,Y)	L
5	$\frac{2}{5}$	$(2,3)$	(H,Y)	H
6	$\frac{2}{5}$	$(2,3)$	(L,Y)	L
7	$\frac{2}{5}$	$(2,3)$	(L,Y)	L
8	$\frac{2}{5}$	$(2,3)$	(H,Y)	H
9	$\frac{2}{5}$	$\vdots$	$\vdots$	$\vdots$

Table 2.3: Playing in the reputation game with $r = 5$ and $\frac{p}{v} = \frac{1}{3}$.

As is standard in the literature on fictitious play, we remark that play forms a so-called 'deterministic cycle', that is (H, Y), (L, Y), (L, Y), (H, Y), (L, Y), while (h_t, l_t) converges to $(2, 3)$.

2.3.2.2 Rating system and reputation formation

Having uncovered the condition for the formation of reputation in the game, I now investigate for which value of the time-discount factor δ the long-lived player is willing to diverge from the Nash convention. Such a proceeding allows to highlight how the long-lived player forms his strategies depending on the number of ratings displayed by the reputation mechanism. I start by considering the case of a reputation system which forbids costless cheating.

Proposition 2. *If* $r < \frac{v}{v-p}$, *then*

i. reputation is formed for $\delta \in (\frac{e}{p}, 1)$,

ii. the equilibrium strategy for a sufficiently patient long-lived player involves the play of an infinite sequence of Stackelberg actions.

Proof. Let $r < \frac{v}{v-p}$. By Proposition 1, the long-lived player cannot cheat costlessy on his reputation. Playing the Stackelberg convention in every period leads to a lifetime payoff

$$\begin{aligned} \Pi_1^H &= \sum_{t=0}^{\infty} \delta^t (p - e) \ , \\ &= \frac{p - e}{1 - \delta} \ . \end{aligned} \tag{2.9}$$

Now, playing only the Nash convention, i.e. action L, leads to an immediate lock of the game and so to a lifetime payoff

$$\Pi_1^L = p \ . \tag{2.10}$$

Comparing (2.9) and (2.10), the long-lived player prefers to play the Nash action if and only if

$$p \geq \frac{p-e}{1-\delta} \quad \Leftrightarrow \quad 0 < \delta \leq \frac{e}{p} \quad \text{where} \quad p > e > 0.$$

Hence, the long-lived player chooses the Stackelberg action H in the first period of his life when his patience satisfies $\delta > \frac{e}{p}$. Now, is the converse true? That is, if $\delta \leq \frac{e}{p}$ is it necessarily the case that the long-lived player always chooses the Nash action L at t_0? Suppose an impatient long-lived player who is only willing to exert some effort in the first period of his life and plays the Nash action for the remaining of the game, implying a lifetime payoff,

$$\Pi_1^{H+L} = p - e + \delta p. \tag{2.11}$$

Then, its is straightforward to remark that (2.11) is strictly larger than (2.10) when $\delta > \frac{e}{p}$. Put otherwise, if $\delta \leq \frac{e}{p}$, the long-lived player always plays the Nash action at t_0 and no reputation is formed in the game.

Finally, by playing the Stackelberg action H at t_0, the long-lived player faces an identical decision problem in the next period and so will play the action H again, and so on and so forth. □

The long-lived player's decision to play the Stackelberg action, action H, depends on his patience δ. In a game featuring a reputation system that makes costless cheating not possible, i.e. a reputation system characterised by a number of ratings satisfying $r < \frac{v}{v-p}$, the play of H decreases in the cost of effort, e, and increases in the payment received for the interaction, p. Put differently, and quite intuitively, only a long-lived player who sufficiently values the future will be ready to play the Stackelberg action H and so to earn $(p-e)$ at every period of the game, instead of playing the Nash action L and earning a unique p in the first period.

I next turn my attention to a reputation system that makes possible costless cheating, i.e. a system characterised by a number of ratings $r \geq \frac{v}{v-p}$.

Proposition 3. *If* $r \geq \frac{v}{v-p}$, *then*

i. *reputation is formed for* $\delta \in (\underline{\delta}, 1)$ *where* $\underline{\delta} \in (\frac{e}{e+p}, \frac{e}{p}]$,

ii. *the long-lived player's expected lifetime payoff increases in the number of ratings* r.

Proof. Let $r \geq \frac{v}{v-p}$. By Proposition 1, we know that the long-lived player can costlessly deviate from the Stackelberg convention, action H, and he can play in certain periods of his life the low-effort action L without risking a locking of the game. Denote Π_1^{HL} the expected lifetime's payoff corresponding to the equilibrium strategy profile of the long-lived player when the reputation system satisfies $r \geq \frac{v}{v-p}$.

Part (i) – If costless cheating is possible in the game, the play of the Nash action L is never delayed by the long-lived player. This implies that Π_1^{HL} has for boundaries

$$\Pi_1^H < \Pi_1^{HL} < \Pi_1^L,$$
$$\frac{p-e}{1-\delta} < \Pi_1^{HL} < \frac{p}{1-\delta},$$

where the LHS of the inequality comes from (2.9) and the RHS is equivalent to the payoff associated to the infinite play of L by the long-lived player in the case of an extremely permissive or inefficient reputation system, i.e.,

$$\Pi_1^L = \sum_{t=0}^{\infty} \delta^t p,$$
$$= \frac{p}{1-\delta}.$$

I investigate for which level of patience δ reputation formation is bounded below when costless cheating is possible. Suppose a limit environment in which r and p are such that the long-lived player only plays the Stackelberg action at t_0 and costlessly cheats for the remaining length of the game. Note that such a scenario induces an extremely high number of ratings to be plausible. Then, the long-lived player's expected lifetime payoff is

$$\begin{aligned}\Pi_1^{H+\overline{L}} &= (p-e) + \sum_{t=1}^{\infty} \delta^t p, \\ &= \frac{p}{1-\delta} - e. \end{aligned} \tag{2.12}$$

Comparing (2.12) with the lifetime payoff associated to the Nash convention played from t_0 onwards, we have

$$p \geq \frac{p}{1-\delta} - e \quad \Leftrightarrow \quad 0 < \delta \leq \frac{e}{e+p},$$

and the lower bound of δ is $\underline{\delta} = \frac{e}{e+p} < \frac{e}{p}$ $\forall e, p > 0$.

Part (ii) – Following the proof of Proposition 1, we know that the number of times the long-lived player can safely play the Nash action, l^*, increases in the number of ratings r. Formally,

$$\begin{aligned} l^* = r(1-\frac{p}{v}) \quad \rightarrow \quad \frac{\partial l^*}{\partial r} &= 1 - \frac{p}{v}, \\ &= \frac{v-p}{v} > 0, \quad \text{where } v > p > 0. \end{aligned}$$

Then, for every unit increase in r, the long-lived player recovers $(\frac{v-p}{v})$ times his cost of effort e in a sequence made of r ratings. Hence, Π_1^{HL} increases in r, $\forall r \geq \frac{v}{v-p}$. □

In a game featuring a reputation system allowing costless cheating, the long-lived player's willingness to form reputation decreases in the cost of effort e, increases in the price of the interaction p, as for $r < \frac{v}{v-p}$ but now, it also increases in the number of ratings made available by the reputation system. The intuition is simple and follows from the one of Proposition 1. Indeed, the long-lived player (e.g., a seller) is more prone to enter a market when he knows that the reputation system in place is such that occasional or accidental plays of the Nash action will not ban him from the market and smash his previous effort investment decisions.

Denoting δ^{HL} the patience required for the long-lived player to form reputation when such a rating system is in place, we remark that the Stackelberg convention H is played for lower levels of patience than in games featuring rating systems forbidding the play of Nash actions L that require a level of patience denoted δ^{H}. Specifically, we have

$$\delta^{HL} \in (\frac{e}{e+p}, \delta^{H}] \quad \text{where} \quad \delta^{H} \in (\frac{e}{p}, 1) \quad \text{and} \quad v > p > e > 0. \tag{2.13}$$

Following (2.13), I identify a trade-off between the number of ratings necessary to minimize cheating in the game and the number of ratings needed to encourage interactions in the game.

Corollary 8 (The Ratings Trade-off)**.** *A reputation system that makes publicly available a small number of ratings limits cheating behaviours* <u>*but*</u> *discourages a long-lived player to form reputation, hindering then the overall number of interactions in the game.*

Put otherwise and all things being equal, locking of the game occurs for a smaller range of δ when the reputation system makes available a larger number of ratings in the economy. This is especially true when the net benefit of interacting for a short-lived player is substantial in the stage game. Such a finding calls quite naturally some comparative statics analysis.

The case of free goods ($p \to 0$). Consider a game where the price of interacting, p, and the cost of effort per interaction, e, tend both towards 0. An illustration of such a stage game could be an online marketplace proposing indivisible non-rival or anti-rival goods.[14] In this game, costless cheating is highly probable given that $\frac{v}{v-p} \to 1$. To limit the play of the Nash convention, the reputation system must then display a very small number of ratings, i.e. $r \to 1$. Making publicly available an extremely small number of ratings has for consequence to restrict the entry in the game to extremely patient long-lived players, i.e. $\delta \to 1$. Put differently, economies proposing cheap or free goods are more prone to encourage cheating behaviours (i.e. play of L) from long-lived players without endangering the risk of being punished (see Proposition 1). The intuition is relatively straightforward. Since the cost of interacting for the short-lived players is very low, they are willing to interact even in the face of a relatively bad interaction. This, in turn, means that the long-lived player can get away with frequently cheating with no real consequences. To minimize the play of the low level of effort action, the reputation system attached to the stage game must then make available a very small number of ratings, i.e. a system satisfying $r < \frac{v}{v-p}$, restricting then the entry in the reputation game to only very patient long-lived players.

The case of auctions ($p \to v$). Consider now a game where the price of interacting, p, tends towards a short-lived player's value of interacting, v. Online market using a second-price auctions mechanism to allocate indivisible goods among a large number of short-lived players, e.g. an economy à la eBay, are likely to fall closer to this end of the spectrum.[15] In this game, costless cheating

[14]Like a public good, an anti-rival good is both non-excludable and non-rival. However, it is created by private economic agents for common benefit without being motivated by pure altruism, because a single economic agent also receives benefits from the contributions of others, e.g., software. See Weber (2004).

[15]In any standard auction with symmetric private values on a finite support, the expected surplus of the winning bidder goes to zero as the number of bidders grows large. See for instance

is extremely unlikely, given that the number of ratings made available by the reputation system is finite and that $\frac{v}{v-p} \to \infty$. Put differently, since much is at stake, short-lived players (or in this case bidders) do not forgive past play of the Nash action L by the long-lived player. Aware of this, the latter player does not have any other choice than playing an infinite sequence of the Stackelberg convention if he does not want to lock the game. In this case, the trade-off is then less prominent given that costless cheating is made *structurally* impossible by the parameters of the game.

In summary, the number of ratings displayed by a reputation system influences the long-lived player's equilibrium strategy. A low r can limit the formation of reputation in the game by restricting the play of high effort actions to very patient long-lived players while a high r can encourage costless cheating or low effort strategies. The latter trade-off increases in importance as the net benefit of interacting widens in the game.

The next section analyses further the interplay existing between the number of ratings r and the long-lived player's strategy.

2.3.3 Equilibrium Strategy Profile

We know that a short-lived player decides to play the trusting action and to interact with the long-lived player only if the latter player's reputation at time t satisfies the game's interaction condition. However, a sufficiently patient long-lived player who is willing to enter the game and to form reputation cannot follow the same kind of myopic behaviour. Indeed, in order to maximize his lifetime payoff, he has to ensure that the game will not become locked in the next periods of his life because of bad ratings. It follows that the long-lived player builds his strategy profile with respect to the number of ratings featured by the

---

Vickrey (1961).

game's reputation system. Proposition 4 formalizes the argument.

Proposition 4. *The strategy profile of a sufficiently patient long-lived player is composed of an infinitely repeating sequence of r actions.*

Proof. Let $r < \frac{v}{v-p}$ and $\delta > \frac{e}{p}$. By Proposition 2, the long-lived player plays the Stackelberg strategy during the whole length of the game. The long-lived player's strategy profile is then simply made of the action H repeated over r times.

Now, let $r \geq \frac{v}{v-p}$. By Proposition 3, a sufficiently patient player enters the reputation game and maximizes his lifetime payoff by cheating on his reputation. Specifically, the optimal number of times he plays the Stackelberg strategy is given by (2.8), that is $\lceil h^* \rceil = \lceil \frac{p\,r}{v} \rceil$, and similarly the maximum number of times he plays the Nash action is respectively $\lfloor l^* \rfloor = \lfloor r(1 - \frac{p}{v}) \rfloor$ where the sum of h^* and l^* is nothing else than the number of ratings r made available by the reputation system. □

It follows from Proposition 4 that the finite number of ratings displayed by a reputation system dictates the (stationary) equilibrium strategy path of the long-lived player within the game. More particularly, a long-lived player's strategy at t only depends on the state ω of his reputation and not on which period of his life he is in. In other words, each time a specific state is visited, the long-lived player faces the same trade-off between the possibility of a high payoff today and the consequences this could have on his future payoffs and picks the same strategy, creating then a cycle over the state space Ω.

This is trivial when costless cheating is forbidden by the reputation system given the homogeneity of the finite state space Ω. If the long-lived player is not patient enough, Ω will only be made of a Nash action L, and the game will converge to the play of the stage game's pure Nash equilibrium. Conversely, for a sufficiently patient long-lived player (i.e. $\delta > \frac{e}{p}$), the state space will be only composed of Stackelberg actions H, and the game will converge to the play of the Stackelberg convention. Then, if $r < \frac{v}{v-p}$ the states remain constant over the length of the game, implying that the same strategy is played every time a specific state is visited.

The dynamics is more interesting when costless cheating is possible, i.e. $r \geq \frac{v}{v-p}$. Following Lemma 7, we know that the first action of a sufficiently patient long-lived player at t_0 must be the high effort one, H, if he wants to form reputation and not lock the game. We also know that the long-lived player will play $\lceil h^* \rceil$ times the Stackelberg action and $\lfloor l^* \rfloor$ times the Nash action within the r first periods in order to maximize his expected lifetime payoff (see Proposition 1). Consider now that we are in the $(r+1)$th period of the game. By deviating from the Stackelberg convention at t_r (i.e. $t = r$, the game starting at $t = 0$), the long-lived player inevitably locks the game, that is $\frac{h^*-1}{r} < \frac{p}{v} \Leftrightarrow \frac{pr-v}{r} < p$ which is always true given $v > p > 0$. Thus, like at t_0, the long-lived player plays again the Stackelberg action at time t_r. It is straightforward to remark that at t_{r+1}, he faces the same reputation score than at t_1 and plays the same action, and so on and so forth.

2.3.4 Number of Ratings to Display

The above analysis demonstrated that the number of ratings embedded by a reputation system determines the equilibrium strategy profile of the (sophisticated) long-lived player and helps foster the play of the Stackelberg convention over the Nash one. I now investigate the consequences of a change in r on social welfare in the game in order to determine the optimal number of ratings an online reputation system should display. I first examine how costless cheating impacts social welfare.

Lemma 9. *Social welfare is maximized when costless cheating is not possible in the game.*

Proof. I show that social welfare, or total surplus hereafter, is larger in games featuring reputation systems forbidding costless cheating than ones forgiving such behaviours. Remark first that the total surplus (TS) of a game in which reputation cannot exist (see Lemma 7) is simply zero

$$\begin{aligned} TS &= PS + CS, \\ &= \delta^0 p - p, \\ &= 0, \end{aligned}$$

where the producer surplus (PS) is the surplus for the long-lived player, and the consumer surplus (CS) represents the surplus for the short-lived players in the game.[16] Consider now the total surplus of a game in which reputation can exist. Following Table 2.1, it is straightforward to note that the per period total surplus when the Stackelberg strategy profile is played is always strictly larger than the cheating one. Hence, total surplus for the whole game, which is nothing else than the sum of the per period total surpluses, is always larger when cheating is not tolerated by the reputation system. □

In other words and as should be expected, costless cheating impedes social welfare. Suppose $r \geq \frac{v}{v-p}$. If $\delta > \frac{e}{p}$, social welfare decreases in the play of costless cheating actions and then is always smaller within this reputation mechanism than one featuring a number of ratings satisfying $r < \frac{v}{v-p}$. If $\delta \leq \frac{e}{e+p}$, no reputation is formed an we are back to a total surplus of 0. If $\delta \in (\frac{e}{e+p}, \frac{e}{p}]$, then Stackelberg actions become possible and social welfare increases in their play. We note that the range of δ which benefits of such a special condition increases in e, where $e \in (0, p)$ and tends towards its maximum for $e \to p$. However, as $e \to p$ the long-lived player's payoff associated to the Stackelberg strategy tends to 0, i.e. $(p - e) \to 0$. Besides and following the proof of Proposition 3, we know that

[16] The Nash equilibrium $E = \{L, N\}$ is not played at t_0 because of Assumption 1.

$\delta \to \frac{e}{e+p}$ supposes a limit environment with an extremely large number of ratings and an almost infinite play of cheating actions whose associated surplus is 0 in the game. Hence, as $e \to p$ and costless cheating is possible, social welfare tends towards 0. It remains to consider the inverse situation, that is $e \to 0$. Here, $\frac{e}{e+p} \to \frac{e}{p}$ and it becomes not costly in terms of δ to enforce Pareto efficiency.

Quite naturally, this results induces that a reputation system should discourage the play of the Nash convention in the game. Lemma 10 follows.

Lemma 10. *Conditional on $r < \frac{v}{v-p}$, any strategy that maximizes the long-lived player's lifetime payoff also maximizes social welfare in the game.*

Proof. The proof is direct and follows from Lemma 9 and Proposition 2. By Lemma 9, we know that social welfare is maximized for $r < \frac{v}{v-p}$. By Proposition 2, we know that the equilibrium strategy of a sufficiently (insufficiently, respectively) patient long-lived player is to indefinitely play the Stackelberg (Nash) convention when the number of ratings is such that costless cheating is not possible. For $\delta > \frac{e}{p}$, this is trivial to remark that social welfare attains its maximum when the Stackelberg convention is played. For $r < \frac{v}{v-p}$ and $\delta \leq \frac{e}{p}$, playing the Nash convention is the only way to limit losses on the consumer's side (short-lived players) and so is Pareto optimal. □

To maximize social welfare (also market efficiency henceforth), a social planner or reputation system designer must ensure that costless cheating is not sustainable in the game. Without loss of generality, we remark that the strategy that maximizes the time-discounted long-lived player's payoff is also the one that maximizes market efficiency in the game when r forbids safe deviations from the Stackelberg convention. It directly follows that a mechanism designer who seeks to maximize social welfare needs to display a minimum of ratings.

Proposition 5. *A social planner maximizes social welfare by fixing $r^* = 1$.*

Proof. The proof is as follows. By Lemma 9, we know that social welfare in the game is always larger and tends towards its maximum when the reputation system forbids costless cheating, that is a reputation system which makes publicly available a number of ratings $r < \frac{v}{v-p}$. Following Proposition 2, we know that the long-lived player's decision to play well in a game characterised by such a reputation system and to form reputation is independent of r and solely depends on the cost of effort e and interaction's price p, $\delta > \frac{e}{p}$. Hence, supposing a game in which costless cheating is made impossible for a large range of r, the decision to increase the number of ratings displayed by the reputation system does not extend the range of δ for which the long-lived player is willing to enter the reputation game.

It remains to show that the minimum number of ratings supported by the reputation system, that is $r = 1$, is also the optimal number of ratings that should be displayed within the game. Following Proposition 4, we know that the equilibrium strategy of a sufficiently patient long-lived player is to play indefinitely the Stackelberg action H when $r < \frac{v}{v-p}$, implying quite trivially the infinite repetition of a stationary strategy profile of size one. □

Proposition 5 states that market efficiency is maximized when a reputation system only displays the most recent rating in a stage game featuring perfect monitoring. It is worth noting that such a result rests on our previous findings and observations. Indeed, Lemma 9 showed that social welfare decreases in costless cheating implying that long-lived player's strategies featuring costless cheating should be discouraged by the reputation mechanism. Proposition 1 demonstrated that costless cheating is not possible when the number of ratings satisfies $r < \frac{v}{v-p}$. Proposition 2 determined that if $r < \frac{v}{v-p}$, then the long-lived player's equilibrium strategy involves the play of an infinite sequence of Stackelberg actions, provided that his patience satisfies $\delta > \frac{e}{p}$. In this situation, we thus noticed that the long-lived player's willingness to form reputation does not depend on r. Furthermore, Proposition 4 established the stationary property of the long-lived player's equilibrium strategies. The two latter results lead to the conclusion that the condition $r < \frac{v}{v-p}$ can be simplified to $r = 1$ without loss of generality.[17]

To complete and illustrate Proposition 5, consider the following remarks which focus on two different types of marketplaces.

Remark 11. *For $p \to v$, the number of ratings r has a limited effect on market efficiency.*

When $\frac{v}{v-p} \to \infty$, costless cheating becomes impossible and the unique strategy for a sufficiently patient long-lived player is to play high effort action H, the Stackelberg action. Given that the state space of the latter player is only composed of actions H, a reputation system can pick any number of ratings ($r \in \mathbb{N}_+$) to maximize welfare in the game, the minimum being the stationary value $r = 1$. Put differently, in marketplaces where $p \to v$, a mechanism that only publishes the outcome of a seller's single most recent interaction ($r = 1$) is capable of generating the same levels of cooperation and social welfare as more complex systems that compile large numbers of interactions. Interestingly and as already highlighted by Dellarocas (2005), the long-lived player's decision to stick to the Stackelberg convention in such markets (e.g., second-price auctions) could provide some theoretical background to explain the remarkably small fraction of negative ratings empirically observed in online marketplaces à la eBay, thus contradicting the literature explaining such a polyannaism by the poor quality of reputation systems in place in those economies.[18]

[17]Besides stationarity and δ independent of r for $r < \frac{v}{v-p}$, an intuitive explanation for $r^* = 1$ instead of $r < \frac{v}{v-p}$ would be to suppose that v can vary among short-lived players, $\forall v > p$. For instance, by considering $\frac{v}{v-p}$ to be uniformly distributed in $(1, n)$, n being potentially very large.

[18]For instance, Resnick and Zeckhauser (2002) report that, based on a dataset of 36,233

Remark 12. *For $p \to 0$ or $v \to \infty$, the number of ratings r has a substantial effect on market efficiency.*

When $\frac{v}{v-p} \to 1$, costless cheating can be sustained for almost any r and social planners must make available a minimal number of ratings if they want to maximize total surplus in the game, choosing then $r = 1$.[19] This policy has for main consequence to restrict the entry in the reputation game to only extremely patient long-lived players, limiting so the use of the mechanism but not its usefulness. In marketplaces featuring such properties, we note that the trade-off first identified in Corollary 8 between the quality of the interactions (low r, high δ) and the quantity of interactions (high r, low δ) is hence especially prominent.

In short and as already hinted by Corollary 8, the number of ratings displayed by a reputation system becomes an essential predictor of market efficiency as the net benefit of interacting increases in the game. This result will be further analysed and expanded in the coming section.

2.4 Extensions

I develop in this section variants of the model presented at the beginning of the chapter. A first extension identifies how the original setting is subject

randomly chosen completed eBay transactions, less than 1% of ratings left by buyers are negative or neutral. They speculated at the time that such a low proportion could be explained by fear of retaliation, i.e. the fear that the seller would retaliate by leaving negative feedback for a buyer after receiving a negative one. In Spring 2008, eBay changed its feedback rules, making then impossible for sellers to leave negative feedback for buyers. Following this policy change, Saeedi et al. (2013) find that the percentage of positive feedback left by eBay buyers has increased even more. While this finding can be counter-intuitive at first sight, a plausible explanation emphasized by the authors is that in the absence of retaliation, sellers – especially the ones proposing lower quality interactions – lost a tool to control the market outcome by intimidating the buyers, and had either to exert more effort per transaction to maintain their reputation or leave the marketplace.

[19] Note that $r = 1$ is a design policy which only applies to our current setting in which the long-lived player always plays his maximizing lifetime payoff strategy. Section 2.4.3 develops the setting further by acknowledging the possibility that players may have a lapse in reason or judgement.

to 'whitewashing', that is the fact to clean a bad reputation by re-entering a marketplace under a new identity. Following this observation, I propose to make our reputation system robust to costless identity changes by introducing a 1-time entry fee. A second extension relaxes the assumption that a rating is created after each single interaction in the game, and examines how our results evolve when incomplete reporting is inserted into the model. A third and final extension considers that the reputation system may feature imperfect monitoring and shows that a social planner can mitigate the effects of erroneous or wrong rating decisions by increasing the number of ratings displayed by the reputation mechanism.

2.4.1 Whitewashing (Cheap Pseudonyms)

Changing identity online[20] is a relatively simple and cheap process. Nowadays, it may require just a few keystrokes and an email address to enter an online marketplace under a specific username or pseudonym.[21] The efficiency of an online reputation system thus depends on its ability to prevent 'opportunistic' sellers from leaving and re-entering the market with a clean and fresh reputation as soon as they receive a bad rating.

In this section, I briefly examine how the model can be extended in order to prevent whitewashing. I first show that the long-lived player maximizes his expected lifetime payoff by always playing the Nash action when he can freely enter the marketplace under new identities. Then, I propose to introduce an entry fee or bail in order to restrict the use of cheap pseudonyms and improve the efficiency of the reputation mechanism.

[20]This process is also called 'whitewashing', hereafter.
[21]See for instance Resnick and Zeckhauser (2002).

2.4.1.1 Problem formulation

Consider that costless cheating is not possible in the game thanks to a small number of ratings displayed by the reputation system.[22] The long-lived player then has the choice at every period t of his life to either form reputation and play the Stackelberg action (H) or lock the game and play the Nash action (L). Furthermore, suppose that the long-lived player is free to leave and re-enter the game as many times as he wishes. Then Lemma 13 follows.

Lemma 13. *The long-lived player maximizes his expected lifetime payoff by whitewashing his reputation.*

Proof. The proof is direct. By playing the Stackelberg action during his whole lifetime, the long-lived player earns

$$\begin{aligned}\Pi_1^H &= \sum_{t=0}^{\infty} \delta^t (p - e), \\ &= \frac{p - e}{1 - \delta}.\end{aligned}$$

Playing the Nash action after a Stackelberg action irremediably leads to a locking of the game and thus an expected lifetime payoff smaller than Π_1^H.

Now, if the long-lived player adopts a 'whitewashing' strategy and always plays the Nash action by exiting the game at the end of each period and re-entering it under a new clean identity at the beginning of every t, then he earns

$$\begin{aligned}\Pi_1^L &= \sum_{t=0}^{\infty} \delta^t p, \\ &= \frac{p}{1 - \delta},\end{aligned}$$

where $\Pi_1^L > \Pi_1^H \; \forall \delta \in (0, 1)$. □

[22] As shown previously in this chapter, social welfare is maximized when costless cheating is not possible. Note that the findings of this section do not change if costless cheating is tolerated.

Lemma 13 establishes that the current reputation system is vulnerable to whitewashing, thus impeding both the formation of reputation and social welfare in the game. This result is especially critical since it is independent of δ, implying that even a sufficiently patient long-lived player finds it optimal to adopt a Nash strategy.

In order to prevent the use of cheap pseudonyms and guarantee market efficiency, a social planner needs to perturb the mechanism so that a newly-born long-lived player never finds it optimal to adopt a whitewashing behaviour. This can be achieved by considering a setting in which the long-lived player starts the game from a state which features the smallest possible payoff among all states. The coming lines present the changes it induces for our original model.

2.4.1.2 Entry fee

Following Friedman and Resnick (2000), suppose there exists a fixed 1-time fee $z > 0$ the long-lived must pay at t_0 in order to enter the game. The key question for a mechanism designer is to fix 'z' such that it prevents whitewashing while still encouraging the play of the Stackelberg action in the stage game.

Proposition 6. *If $z > \frac{e}{\delta}$, then whitewashing is prevented and market efficiency preserved.*

Proof. Suppose costless cheating is not possible and fix $z > 0$. By playing the Nash action L during his whole lifetime, the long-lived player's payoff now reduces to

$$\begin{aligned}\Pi_1^{L,z} &= \sum_{t=0}^{\infty} \delta^t (p - z), \\ &= \frac{p - z}{1 - \delta}.\end{aligned}$$

It is direct to remark that if $z = p$, then $\Pi_1^{L,z} = 0$.

In contrast, by playing the Stackelberg action H, the long-lived player's payoff becomes

$$\begin{aligned}\Pi_1^{H,z} &= p - e - z + \sum_{t=1}^{\infty} \delta^t (p-e), \\ &= \frac{p-e}{1-\delta} - z,\end{aligned}$$

where $\Pi_1^{H,z} > \Pi_1^{L,z}$ $\forall \delta > \frac{e}{z} \Leftrightarrow z > \frac{e}{\delta}$. □

Following Proposition 6, we remark that the condition on z shares some evident similarities with the one identified by Proposition 2, where the long-lived player's equilibrium strategy involves the play of H if $\delta > \frac{e}{p}$. Hence, a possible and reasonable choice of z for a social planner would be to make the reputation mechanism robust to whitewashing by charging the long-lived player an entry fee equal to the revenue of his first interaction, i.e. $z = p$.

More generally, the core intuition for Proposition 6 is that a new seller must prove his worth when entering a marketplace and so pay a kind of trust fee. In large online markets such as eBay, the fee z can be parametrized by the price premium that is the difference existing between the auction's price of a new seller and a high-reputation one. In smaller online marketplaces, registration can be non-free and new entrants may have to pay a fixed fee. For instance, the underground (or 'Deep Web') and now defunct online marketplace 'Silk Road' prevented whitewashing by asking for a fixed fee whenever a new seller wanted to register.

2.4.2 Incomplete Reporting

The previous analysis considered a setting in which short-lived players always submit a rating after an interaction, i.e. $\mathbb{P}(r|Y) = 1$. However, this one-to-one relationship is not verified empirically, mainly because online ratings can be viewed as public goods.[23] Indeed, their creation is (time) costly but benefits the entire economy, implying that economic agents are more incline to free-ride on the contributions of others.

Using a laboratory experiment, Lafky (2014) demonstrates that the cost of rating plays a major role in influencing rating behaviour, a rating being more often submitted when the observed quality of an interaction – either good or bad – outweighs the rating cost. Gathering data on eBay, Resnick and Zeckhauser (2002) find that 51.7% of buyers leave feedback after a transaction with a seller.[24] Constructing a panel of eBay transactions, Cabral and Hortaçsu (2010) note that 40.7% of those transactions resulted in a feedback.[24] Dellarocas et al. (2004) study the drivers and dynamics of buyer participation in eBay's feedback system. The aforementioned authors observe that a buyer rates a seller 68% of the time and demonstrate through a probit model that this high rate of voluntary participation can be explained by the combined effects of (impure) altruism, self-interest, and reciprocation.[24] In a similar study, Diekmann et al. (2014) find an even higher rate of participation and show that up to 85% of eBay transactions are rated by buyers.[24] Their results reinforce the idea that reciprocity and altruism (especially for inexperienced sellers) both play a major role in the decision of a buyer to submit a rating.

As illustrated by the works mentioned above, most of the empirical studies have explored online reputation systems using data coming from eBay or tried to reproduce the fundamentals of the latter economy in experiments. However, eBay features idiosyncrasies that cannot be generalized to every online marketplace, e.g., goods are allocated through auctions and are mainly indivisible. Following this observation, I scraped and gathered data during Spring 2012 from

[23] See Dellarocas et al. (2004).

[24] To note that these findings all rest on data collected before 2008, when eBay changed its two-sided feedback system to make it less sensitive to retaliation by forbidding sellers from giving negative feedback on buyers. Bolton et al. (2013) describe this new reputation system as a "one-sided feedback system that still allows for positive reciprocity".

three online marketplaces which all propose non-rival goods[25] for free or a very small (fixed) price, namely Google Play[26], the Chrome Webstore[27], and AMO (Mozilla Add-ons marketplace)[28]. Results show that the probability to get a rating after an interaction is on average significantly lower in those economies than in markets à la eBay. Specifically, I find that only up to 5% of transactions are rated.

Type of good	Non-rivalrous	Rivalrous
Example of marketplace	Google Play	eBay
Probability of a rating	**1% to 5%** [29]	**40% to 85%** [30]

Table 2.4: Probability for a 'buyer' to submit a rating after a transaction.

The dramatic difference in rating probabilities between those two types of online marketplaces (and their respective goods) suggests that there may exist a positive relationship between the price of a good and the probability to report a rating.

2.4.2.1 Existence of a rating in the game

To incorporate such empirical evidence into the model, suppose that a rating now exists after an interaction (i.e. action Y) with probability $\phi \in [0, 1]$, where ϕ is time-independent (as shown in Table 2.5). It follows that the transition function τ takes an additional parameter ϕ, i.e. $\tau(\omega_t, a_{1t}, \phi)$, and the long-lived player now expects to play a number $\frac{r}{\phi}$ of interactions to get r ratings in the system. Indeed, let

[25] E.g., browser extensions and mobile applications.
[26] https://play.google.com
[27] https://chrome.google.com/webstore
[28] https://addons.mozilla.org
[29] See Appendix A.2.
[30] Cabral and Hortaçsu (2010) (40.7%), Resnick and Zeckhauser (2002) (51.7%), Dellarocas et al. (2004) (68%) and Diekmann et al. (2014) (65% - 85%).

		Player 2	
		Y	N
Player 1	H	ϕ	0
	L	ϕ	0

Table 2.5: Existence of a rating in the game.

X_1 be the number of interactions (trials) to the first rating (success),

X_2 be the number of interactions to the second rating,

X_3 be the number of interactions to the third rating,

$\vdots$

Each X_i is a geometric variable with success factor ϕ, so $\mathbb{E}(X_i) = \frac{1}{\phi}\ \forall i$. Let N be the number of interactions to the r-th success, where N is a negative binomial random variable, $N = \sum_{i=1}^{r} \mathbb{E}(X_i)$. Since the expectation of a sum of random variables is the sum of their expectations, we have

$$\begin{aligned} \mathbb{E}(N) &= \mathbb{E}(\sum_{i=1}^{r} X_i)\,, \\ &= \sum_{i=1}^{r} \mathbb{E}(X_i)\,, \\ &= \frac{r}{\phi}\,. \end{aligned} \tag{2.14}$$

The coming analysis starts by investigating the effects of the measure ϕ on the formation of reputation and the occurrence of costless cheating in the game. Then, I examine how incomplete reporting affects the long-lived player's equilibrium strategy in the game. Finally, I check whether this new setting engenders some changes for the optimal number of ratings an online reputation system should display.

2.4.2.2 Reputation formation

I start by identifying a new condition for reputation formation in the game, in addition to Lemma 7.

Remark 14. *A necessary condition for exertion of effort is* $\phi > 0$.

If $\phi = 0$, no rating is left in the game and then no reputation is formed, implying that costless cheating is the only convention played following Assumption 1. In other words, the existence of reputation and its formation now depend on the probability ϕ to get a rating after an interaction. Lower the ϕ, smaller is the chance to get a rating about the long-lived player's past move. I now examine the effects of ϕ on the occurrence of costless cheating in the game.

2.4.2.3 Costless cheating and rating creation

Proposition 7. *Proposition 1 holds* $\forall \phi > 0$.

Proof. The proof is straightforward and follows directly from the one of Proposition 1. Let μ_t be the long-lived player's reputation at time $t > 0$. Suppose that the reputation system makes publicly available a number of ratings $r < \frac{v}{v-p}$. Then $\forall \phi > 0$, the probability to lock the game at $(t+1)$ after the play of the Nash action L is always larger than 0, which is not compatible with our definition of costless cheating. □

The existence condition of costless cheating in the game remains the same and is independent of ϕ as long as $\phi > 0$. Put differently, for $r < \frac{v}{v-p}$, the long-lived player always incurs the risk to lock the game at $t+1$ with a probability $\phi > 0$ if he decides to play the low effort action at t. Knowing this, the long-lived player cannot expect to costlessly cheat on his reputation when a reputation system displays such a number of ratings.

2.4.2.4 Equilibrium strategies

Surely, the introduction of ϕ influences the decision of the long-lived player to form reputation in the game and to play the Stackelberg or high effort action. As in section 2.3.2.2, I start the analysis by considering the case of a reputation system which forbids costless cheating.

Proposition 8. *If* $r < \frac{v}{v-p}$*, then*

i. reputation is formed for $\delta \in (\overline{\delta}, 1)$ *where* $\overline{\delta} = \frac{e}{e+\phi(p-e)} > \frac{e}{p}$ $\forall \phi < 1$*,*

ii. the equilibrium strategy for a sufficiently patient long-lived player involves the play of an infinite sequence of Stackelberg actions.

Proof. Let $r < \frac{v}{v-p}$. By Proposition 1, the long-lived player cannot cheat costlessy on his reputation. Playing the Stackelberg convention in every period leads to a lifetime payoff Π_1^H,

$$\begin{aligned}\Pi_1^H &= \sum_{t=0}^{\infty} \delta^t (p-e) \ , \\ &= \frac{p-e}{1-\delta} \ . \end{aligned} \tag{2.15}$$

Now, playing only the Nash convention implies a probability to lock the game after each action L equal to ϕ, which leads to an expected lifetime payoff Π_1^L,

$$\begin{aligned}\mathbb{E}(\Pi_1^L) &= \sum_{t=0}^{\infty} \delta^t (1-\phi)^t p \ , \\ &= \frac{p}{1-\delta(1-\phi)} \ . \end{aligned} \tag{2.16}$$

Comparing (2.15) and (2.16), the long-lived player prefers to play the Nash action if and only if

$$\frac{p}{1-\delta(1-\phi)} \geq \frac{p-e}{1-\delta} \quad \Leftrightarrow \quad 0 < \delta \leq \frac{e}{e+\phi(p-e)} \equiv \overline{\delta} \quad \text{where} \quad p > e > 0\ . \tag{2.17}$$

Hence, the long-lived player chooses the Stackelberg action H in the first period of his life when his patience satisfies $\delta > \overline{\delta}$. Now, is the converse true? That is, if $\delta \leq \overline{\delta}$ is it necessarily the case that the long-lived player always chooses the Nash action L at t_0? Suppose an impatient long-lived player who is only willing to exert some effort in the first period of his life and plays the Nash action for the remaining of the game, implying an expected lifetime payoff,

$$\begin{aligned} \mathbb{E}(\Pi_1^{H+\overline{L}}) &= p - e + \sum_{t=1}^{\infty} \delta^t (1-\phi)^t p, \\ &= \frac{p}{1-\delta(1-\phi)} - e. \end{aligned} \tag{2.18}$$

Then, it is straightforward to remark that (2.18) is strictly smaller than (2.16) $\forall\delta$. Put differently, if $\delta \leq \overline{\delta}$ and $r < \frac{v}{v-p}$, the long-lived player expects to maximize his lifetime payoff by playing the Nash action from t_0 onwards. □

In words, the long-lived player's decision to play the Stackelberg convention decreases in the cost of effort, e, but increases in the payment received for the interaction, p, and in the probability to get rated, ϕ. Interestingly, the level of patience needed to ensure such a strategy from the long-lived player is now higher than in our previous setting, *ceteris paribus* and $\forall\phi < 1$.

The intuition is simple. When the chance to get rated after an interaction in the market decreases, the long-lived player's incentives to cheat from t_0 onwards and so to earn higher revenues in the first periods of his life increase. Divergence from the Nash convention is then more difficult to implement for low levels of ϕ, meaning that effort investment decisions will be limited to higher levels of patience δ.

Now, consider the case $r \geq \frac{v}{v-p}$.

Proposition 9. *If* $r \geq \frac{v}{v-p}$*, then*

i. reputation is formed for $\delta \in (\underline{\delta}, 1)$ *where* $\underline{\delta} \in (\frac{1}{2}(\frac{2e+\phi(p-e)}{e(1-\phi)} - \sqrt{\frac{\phi(4ep+\phi(e-p)^2)}{e^2(\phi-1)^2}}), \frac{e}{e+\phi(p-e)}]$,

ii. the long-lived player's expected lifetime payoff increases in the number of ratings r.

Proof. Proof is similar to the one of Proposition 3, except for the lower bound analysis. Now and as already seen previously, the lifetime expected payoff of playing the Nash action from t_0 onwards is $\Pi_1^L = \frac{p}{1-\delta(1-\phi)}$. Let r and p be such that the long-lived player only plays the Stackelberg action at t_0 and costlessly cheats for the remaining length of the game without risking a locking of the game, then his expected lifetime payoff can be expressed as

$$\begin{aligned} \Pi_1^{H+\overline{L}} &= (p-e) + \sum_{t=1}^{\infty} \delta^t p, \\ &= \frac{p}{1-\delta} - e. \end{aligned} \tag{2.19}$$

Comparing (2.19) with the lifetime payoff associated to the Nash convention played from t_0 onwards, we have

$$\frac{p}{1-\delta(1-\phi)} \geq \frac{p}{1-\delta} - e, \tag{2.20}$$

and simplifying (2.20) w.r.t. δ, we find

$$0 < \delta \leq \frac{1}{2}(\frac{2e+\phi(p-e)}{e(1-\phi)} - \sqrt{\frac{\phi(4ep+\phi(e-p)^2)}{e^2(\phi-1)^2}}) \geq \frac{e}{e+p} \quad \forall e, p > 0. \tag{2.21}$$

This completes the proof. $\square$

As previously, the introduction of ϕ does not change the core findings and comparative statics analysis presented in section 2.3.2.2. However, the interplay existing between ϕ and δ is especially compelling. Specifically, as the probability to get rated decreases, then the level of patience required to induce a Stackelberg strategy in the game increases, even when the number of ratings is such that costless cheating is 'tolerated'. Intuitively, a low ϕ implies that a large number of Stackelberg actions is needed in order to ensure the formation of a good reputation in the game. Now, such strategies are costly in terms of effort and only a long-lived player who values sufficiently the future is willing to play a higher proportion of high effort actions, especially when he knows that he can cheat and that the probability to be 'punished' is very low.

Such observations lead to an interesting finding. That is, the trade-off first identified in Corollary 8 decreases in importance as $\phi \to 0$. It is indeed straightforward to remark that $\underline{\delta} \to \overline{\delta}$ as $\phi \to 0$. Hence, when the probability to report a rating after an interaction is low in a marketplace, increasing the number of ratings made publicly available by the reputation system does not encourage a long-lived player (e.g. seller) with a lower patience to enter the economy and form reputation.

I conclude this subsection by focusing on the optimal strategy of a long-lived player whose patience satisfies the conditions identified in Propositions 8 and 9.

Proposition 10. *If the long-lived player is sufficiently patient to enter the game, then his optimal strategy profile is stationary and composed of $r\phi^{-1}$ actions on average.*

Proof. The proof follows the same arguments and procedure than the one of Proposition 4. The only difference is that in our new setting, the long-lived player expects to play (on average) $\frac{r}{\phi}$ times to fill the state space of r actions – see equation (2.14). □

Put otherwise, the introduction of ϕ has no effect on the stationary property of the long-lived player's optimal strategy. However, incomplete reporting limits the applicability of reputation systems to economies with larger stage-game profit margins and higher levels of patience.

2.4.2.5 Number of ratings to display

As in section 2.3, I end the analysis by focusing on social welfare. The key question here is to see if our previous findings in terms of market efficiency hold when the game features incomplete reporting. Proposition 11 directly follows.

Proposition 11. *Proposition 5 holds* $\forall \phi > 0$.

Proof. It is straightforward to remark that Lemme 9 and 10 hold in our new setting, the very structure of the stage game remaining the same. By Proposition 7, we know that the condition for the existence of costless cheating did not change neither, $\forall \phi > 0$. Propositions 8 and 10 established that both the form and stationary property of the long-lived player's strategy profile were unchanged. Hence, the result of Proposition 5 applies to a setting featuring incomplete reporting as long as the probability to get rated is larger than 0. □

In other words, the original results and conclusions about the interplay existing between r and market efficiency do not change when ϕ is added to the model since the condition for costless cheating remains the same in our extended setting. However, it is worth noting that the play by the long-lived player of the Stackelberg strategy decreases in ϕ, as seen in Proposition 8. Hence, a higher δ may be required in order to secure an equivalent level of social welfare in the game when $0 < \phi < 1$.

2.4.3 Imperfect Monitoring

This section adds imperfect monitoring to the stage game and discusses the implications of such a structural change for the reputation mechanism's efficiency. The analysis starts by presenting the key change we need to make in order to have imperfect monitoring within the game.

2.4.3.1 A reputation system with imperfect monitoring

Contrarily to the original setting, I now consider that short-lived players differ in their capability of interpreting the exact level of effort chosen by the long-lived player when interacting with him. More particularly and to keep the analysis intuitive, I suppose that there exists a fraction $\rho \in (0, 1)$ of short-lived players who never find out what action was taken by the long-lived player and then view it as a low effort action L, even if the long-lived player chose the high effort action H. Let ρ be independent of time t and set of actions A.

In practice, the measure ρ could epitomise differences in tastes or online experience among short-lived players (e.g., buyers). It could also parametrize accidental third-party errors, lost parcels, bugs.

The next lines identify and discuss the challenges set by the introduction of the measure ρ.

2.4.3.2 Challenges

Suppose that short-lived players are not aware that ρ exists in the game.[31] Then, the interaction condition (2.1) remains unchanged, implying that Proposition 1 holds. Now, it is straightforward to remark that opportunities for the long-lived player to costlessly cheat on his reputation decrease in ρ, all other things being equal. Furthermore, the very existence condition for reputation for-

[31] That is, short-lived players do not know that a proportion ρ of them misinterpret long-lived player's actions.

mation highlighted by Lemma 7 in the game is endangered as well since playing the Stackelberg action does not guarantee a high rating anymore. It follows that a social planner or mechanism designer who applies the findings of our perfect monitoring setting to this new framework cannot ensure an optimal situation in terms of social welfare. Indeed, displaying only the latest rating can deter a long-lived player to enter the reputation game and to adopt the Stackelberg strategy when the reputation system does not take into account that imperfect monitoring lies within the game. Proposition 12 formalizes the latter argument.

Proposition 12. *If $r < \frac{v}{v-p}$, then the long-lived player's equilibrium strategy is to play the*

i. Stackelberg action if $\delta \in (\overline{\delta}, 1)$ where $\overline{\delta} = \frac{e}{p(1-\rho)} > \frac{e}{p}$ $\forall \rho$,

ii. Nash action if $\delta \in (0, \overline{\delta}] \vee (1-\rho) < \frac{e}{p}$.

Proof. Let $r < \frac{v}{v-p}$, forbidding so costless cheating. Playing the Stackelberg action H an indefinite number of time leads to an expected lifetime payoff

$$\begin{aligned}\mathbb{E}[\Pi_1^H] &= \sum_{t=0}^{\infty}(1-\rho)^t\delta^t(p-e),\\ &= \frac{p-e}{1-\delta(1-\rho)},\end{aligned}$$

while playing the Nash action leads to a lifetime payoff $\Pi_1^L = p$. The long-lived player prefers to play the Nash action as long as $p \geq \frac{p-e}{1-\delta(1-\rho)}$ which is equivalent to $0 < \delta \leq \frac{e}{p(1-\rho)}$. □

In short, the play of the Stackelberg strategy decreases in the game as ρ increases, *ceteris paribus*. The intuition is straightforward. When there is a high chance to earn a bad rating in the marketplace independently of the effort invested, it becomes extremely difficult for a social planner to ensure that a seller will have much incentive to behave well.

2.4.3.3 Optimal design implementation

Following Proposition 12, a mechanism designer then needs to counterbalance the negative impact that ρ induces for social welfare by modifying the structure of the reputation system. Given our model, the most natural lever at his disposal is r. Proposition 13 thus presents the optimal number of ratings a reputation system should display when a fraction ρ of erroneous ratings lies in the game.

Proposition 13. *Social welfare is maximized for* $r^* = \frac{v(1-\rho)}{v(1-\rho)-p}$.

Proof. Suppose a social planner wishes to approach the level of participation (in terms of δ) present within a perfect monitoring setting, see Proposition 2. Then, he needs to make sure that the reputation mechanism takes into account the existence of ρ. Let $\mathbb{E}[h]$ ($\mathbb{E}[l]$, respectively) denote the expected number of ratings associated to the play of actions H (L), and $\mathbb{E}[h^*]$ ($\mathbb{E}[l^*]$) the smallest (largest) number of incidents of the play of H (L) in the available history of ratings the long-lived player can expect to earn without running any risk of reputational failure. Fix $t \geq r$. Then, using a similar reasoning than the proof of Proposition 1, we have

$$\mu \equiv \frac{\mathbb{E}[h](1-\rho)}{r} \geq \frac{p}{v},$$

which implies

$$\begin{aligned}\mathbb{E}[l] \leq \mathbb{E}[l^*] &= r - \frac{pr}{v(1-\rho)}, \\ &\leq r(1 - \frac{p}{v(1-\rho)}).\end{aligned}$$

Hence, $\mathbb{E}[l^*] > 1$ if $r > \frac{v(1-\rho)}{v(1-\rho)-p}$. In contrast, if $r < \frac{v(1-\rho)}{v(1-\rho)-p}$ and as $r \to 1$, the probability to lock the game independently of the long-lived player's action increases and thus restricts the play of the Stackelberg strategy to higher δ, as shown in Proposition 12. The same trivial reasoning than the one presented in the proof of Proposition 1 applies for $0 < t \leq r$. □

In order to approach the level of social welfare attainable in a setting featuring perfect monitoring, a mechanism designer must increase the number of ratings made publicly available by the reputation system when imperfect monitoring emerges in the game to compensate the now underestimated long-lived player's reputation μ. In particular, the optimal number of ratings to display (r^*) increases in the proportion of ratings of type ρ.

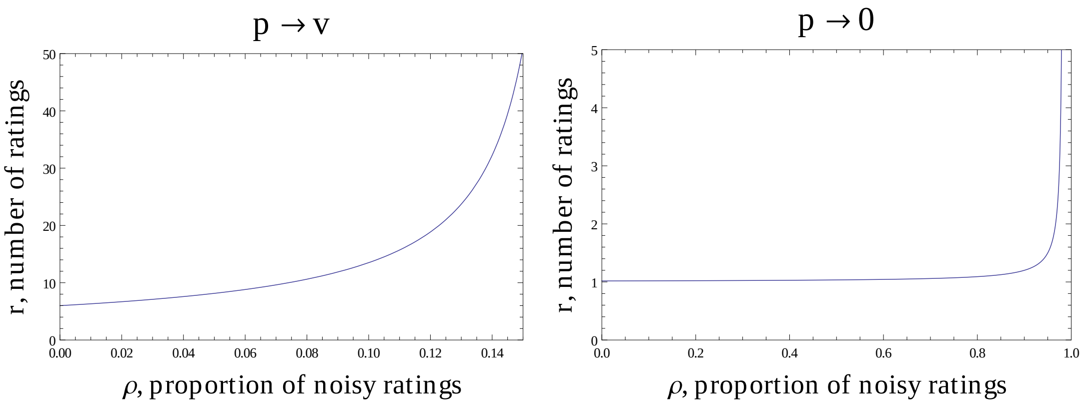

Figure 2.1: Number of ratings to display in order to maximize market efficiency and the play of the Stackelberg convention in function of the proportion of noisy ratings reported in the game. Left graph depicts a situation where $p \to v$ (e.g., auctions). Right graph depicts $p \to 0$ (e.g., free goods). Notice that the two graphs use different scales.

As shown in Figure 2.1, the consequences of ρ on r^* gain in importance as the short-lived player's net benefit of interacting increases in the game. Specifically, the same level of ρ requires a significantly larger increase in the number of ratings disclosed by the reputation mechanism when $p \to v$ than when $p \to 0$. Such an observation hints that the countermeasures to adopt in order to mitigate the negative effects of imperfect monitoring on market efficiency vary in intensity depending on the stage game's profit margins.

2.5 Discussion and Conclusion

In this chapter, I have proposed and developed a model of reputation adapted to the idiosyncrasies of online economies. My approach has differed from the standard reputation literature by relaxing the unrealistic assumptions usually made about the infinite knowledge and computational abilities of players. In particular, I have put forward a moral-hazard stage game in which short-lived players form beliefs about a long-lived player's strategy and choose their actions with respect to these beliefs. The introduction of such a 'belief-based' learning rule, so-called fictitious play, has provided a means to parametrize the specific forecasting rule in place in online marketplaces, where the reputation of a seller comes from the aggregation of his past actions into a bounded space of histories. The analysis has demonstrated that the model's information structure features the compelling property to prevent reputation effects from disappearing in the long run and has allowed to focus on the crucial role played by the amount of information made available through the reputation mechanism, that is the number of ratings.

I have shown that the latter number determines the occurrence of cheating in the game and the equilibrium strategy of the long-lived player. In peculiar, I have been able to identify the 'stationarity' of the long-lived player's equilibrium strategy profile and its dependence on the very structure of the reputation mechanism. A few interesting parallels can be drawn here with some similar results coming from the recent related literature. Mell (2011) found that the optimal strategy of his long-lived player was stationary and identified cycles in reputation using a moral-hazard setting featuring adaptive learning à la Young (Young, 1993). Liu and Skrzypacz (2014) built a reputation model with limited record keeping resting on incomplete information and examined how limited records engender stationary perfect Bayesian equilibria.

The chapter has also provided the reader with some conclusive normative insights concerning the design of online reputation systems. The first and principal finding is that the number of ratings displayed by a reputation mechanism is a fundamental predictor of market efficiency. Specifically, the analysis has shown that the number of ratings has to be kept minimal in order to maximize the efficiency of the mechanism and guarantee a high level of social welfare in the game. It has been highlighted that such a policy recommendation should especially be observed by online marketplaces proposing interactions with a high profit margin. Indeed, the latter economies are more sensitive to the so-called 'ratings trade-off', i.e. a situation where a mechanism designer faces two possible choices: either display a large number of ratings and ensure a high level of sellers' participation at the risk of encouraging low effort interactions, or display a small number of ratings and ensure a high level of effort at the risk of discouraging sellers to enter the market and form reputation. In contrast, when the profit margin is low in the stage game, then high effort interactions are more easily guaranteed by the very structure of the game and are less dependent on the reputation mechanism's design. Such a result provides a theoretical explanation to the pollyannaism which has been often spotted and commented by the empirical literature studying auction marketplaces such as eBay.

By extending the original model, I have been able to present a few additional results which convey a clear normative interest. Namely, fixing an appropriate entry fee can effectively deter whitewashing and generate a more trustworthy reputation system. Besides, the efficiency of a reputation mechanism and the optimal number of ratings it should display are neither altered nor changed when the possibility of incomplete reporting is added to the initial setting. However, seller's participation may decrease when creation of ratings is low in the economy, thus possibly implying a lower level of market efficiency. Finally, social welfare suffers from the presence of noisy reporting behaviours or imperfect monitoring, and a mechanism designer then needs to rectify the situation by increasing the number of ratings displayed by the reputation system. The latter recommendation appears particularly important for economies proposing interactions characterized by a low profit margin, e.g., auctions.

This work represents a possible start for future research about the design of online reputation mechanisms, both theoretically and empirically. It appears especially interesting to investigate further how the long-lived player's effort investment decision may affect a short-lived player's willingness to report a rating after an interaction, thus making completely endogenous the existence of a rating in the game.

Chapter 3

Reporting Bias

How Do Reporting Behaviours Affect the Efficiency of An Online Reputation System?

3.1 Introduction

In this chapter, I propose a theoretical approach which allows to study how reporting behaviours affect the efficiency of an online reputation system. I develop a stochastic setting that makes possible the analysis of different behavioural models and reproduces the way ratings are aggregated online. The efficiency of a reputation mechanism is assessed thanks to a measure, so-called *reporting bias*, which identifies and computes the dispersion of ratings around the true average quality of a seller.

Nowadays, it has become utterly common to rely on online ratings contributed by others to judge the quality of products or services and trustworthiness of individuals. We book restaurants (Luca, 2011), watch movies (Dellarocas et al., 2007), install software or purchase goods (Melnik and Alm, 2002) after having considered their respective online feedback. Due to their increasing importance and use, it appears fundamental to ensure that the information relayed by online ratings are reliable and unbiased. Recent empirical and experimental results, however, raise serious doubts about the ability of current reputation systems to supply truthful ratings.

Using a large-scale randomized experiment on a social news aggregation website, Muchnik et al. (2013) show that past ratings generate a significant bias in individual reporting behaviour. In particular, the authors find that social influence encourages the apparition of herding (also *collective* hereafter) mechanisms which have the potential to substantially distort the dynamics and formation of reputation. Analysing hotel reviews from TripAdvisor[1], Talwar et al. (2007) find that past ratings influence future reports by creating some prior expectation concerning the service quality. Specifically, they identify that the subjective perception of a user is affected by the gap existing between the user's prior ex-

[1]TripAdvisor.com is a website proposing reviews of travel-related content such as hotels.

pectation and the observed performance of the service, and show that the latter disparity may give users incentives to correct the reputation estimate. Similarly to the authors of the aforementioned study, Nagle and Riedl (2013) identify the existence of a *correcting* mechanism and find that disagreement in prior movie ratings leads to a higher propensity to post a rating.

By designing an experiment which replicates an online economy with a 1-sided reputation system, Lafky (2014) finds that the cost of rating leads to a U-shaped distribution of reports and produces average ratings which are not representative of true quality. On the same vein, Hu et al. (2009) show that a considerable proportion of products on Amazon feature a bimodal or U-shaped distribution of ratings, creating then a discrepancy between the reputation score of a specific product and its actual quality. The authors propose a 'Brag and Moan' model in which ratings are only created when the utility gained from a product – drawn from a normal distribution – falls outside a median interval.

Luca and Zervas (2013) show that reputation manipulation is widespread in online economies, and highlight that almost 20% of ratings are detected as fake by Yelp's algorithm.[2] They find that fake ratings tend to be more extreme than other reports in the market, and are published by users with less established reputations. In a different study, Anderson and Simester (2014) find that approximately 5% of product ratings on a large e-commerce website are reported by customers with no record of ever purchasing the product they are rating. Their results suggest that those specific ratings are significantly more negative than average.

This chapter contributes to the stream of literature studying the bias engen-

[2]Yelp.com is a website where users can review and rate local businesses such as restaurants, bars, etc. The parameters used by Yelp's filtering algorithm are trade secrets but results by Mukherjee et al. (2013) suggest that the platform may apply a behavioural approach based on its internal data for each of its users, e.g., number of reviews published within a given timeframe, % of positive reviews, review length, content similarity, IP addresses and geo-locations, session logs, mouse and click gestures, social network interactions, etc.

dered through online reporting behaviours by adopting a theoretical perspective, and thus departs from the usual (empirical) practice. I develop an iterative stochastic approximation model à la Robbins and Monro (1951) where, at each period of time, a short-lived buyer experiences an i.i.d. level of quality from a unique long-lived seller and must rate the transaction based on her own observation of seller's quality and the current reputation score of the seller in the economy. I show that the seller's reputation converges to a specific score and I rest on such a finding to build a measure able to assess the underlying reporting bias present in the market.

I then employ this novel measure to test the effects of common online reporting behaviours on the efficiency of a reputation mechanism. I demonstrate that collective reporting behaviours have the potential to more seriously endanger the reliability of a reputation system than correcting ones. Once a robustness analysis completed, I propose to use my measure of reporting bias to examine the dynamics of reputation in a few different empirically observed *scenari* such as sybil attacks or costly ratings, and to test the relevance of the median as a means to compute the reputation aggregate statistic.

Normative findings about the design of an online reputation system naturally follow. Specifically, results suggest that online reputation mechanisms should rest on two different interfaces in order to limit the emergence of non-linear reporting behaviours, that is one interface for the reading and appreciation of reputation and another 'blank' or reputation-free interface for reporting purposes. To counterbalance the risks associated to reputation manipulation, it appears recommended to display pairs of recent and contrasted ratings side-by-side to promote a 'natural' correction of reputation scores. Marketplaces which feature a high cost of rating should prefer coarse-grained rating scales in order to better contain reporting bias. Finally, while the median appears as an intuitive solution to lessen the impact of outliers on the computation of reputation scores, it is less adapted than the mean in online environments characterized by bimodal distributions of ratings and high-variance reporting behaviours such as correcting mechanisms.

Chapter outline

The remainder of the chapter is organised as follows. The next section introduces the setting. Section 3.3 studies the evolution of reputation in the long-run and builds a measure to assess the efficiency of a reputation system, so-called reporting bias. Section 3.4 analyses the effects that different reporting behaviours may have on the magnitude of the reporting bias in an online marketplace. Sections 3.5 and 3.6 extend further the setting to assess its robustness and empirical pertinence. Section 3.7 concludes.

3.2 The Setting

The setting involves an online marketplace where, in each period of time, a monopolist long-lived seller ("he") transacts with one of multiple risk-neutral short-lived buyers ("she"). The marketplace embeds a reputation system which makes possible the aggregation of ratings for the seller after each transaction. Buyers live for one period and enter the market sequentially. More particularly, at every period $t > 0$, one buyer enters the economy, observes the seller's reputation at t, transacts with the seller for a price equal to his reputation score[3], and rates the outcome of the transaction at the end of her life.

Let Q be a random variable and $q_t \sim Q$, $t = 1, 2, ..T$ be T independent,

[3]For simplicity's sake, I consider a model in which the price of a transaction equals the buyer's expectation of success as parametrized by the seller's reputation score. Specifically, buyers are risk-neutral and offer a price equal to their willingness to pay, valuing a high quality transaction at 1 and a low quality one at 0. As an illustration, suppose a Bayesian setting in which the price of interacting p only depends on a buyer's belief μ that the seller is good. Then, p is simply given by $p = \mu \cdot 1 + (1 - \mu) \cdot 0 = \mu$. This corresponds to a marketplace short on the seller's side, in which several buyers wish to transact with the same seller; a situation commonly found online. See Bar-Isaac (2003) or Bar-Isaac and Tadelis (2008).

identically distributed (i.i.d.) realizations of Q sharing the same cumulative distribution function $F : \mathbb{R} \rightarrow [0, 1]$. Let Q parametrize the level of effort invested by the seller and let its realizations correspond to the level of quality experienced or observed by buyers. Consider that the observation of Q varies across periods and buyers depending on their experience, tastes, or because of third-party errors. The *reputation score* of the seller after t transactions is defined as

$$\mu_t = \frac{1}{t} \sum_{i=1}^{t} r_i, \tag{3.1}$$

where r_t is the rating published by a buyer at the end of t. To capture the essence of *reporting behaviours* found online, a buyer decides to publish a rating r_t based on the quality q_t she observed while transacting and the seller's reputation score or aggregate statistic μ_{t-1} she witnessed before transacting. More particularly, the creation of ratings in the marketplace rests on the two following assumptions.

Assumption 2 (First Rating). *For $t = 1$, $r_1 = q_1$.*

The seller starts his life at $t = 1$ with an empty reputation history ($\mu_0 = \emptyset$), implying that the value of the first rating, r_1, only depends on the outcome of the first transaction, q_1. For subsequent periods of time, past ratings form the seller's reputation and this existing reputation may influence future ratings.

Assumption 3 (Reporting Behaviour). *$\forall t > 1$, a rating is given by the buyer's reporting behaviour Φ,*

$$r_t = \Phi(q_t, \mu_{t-1}), \tag{3.2}$$

where Φ is a function non-decreasing in q_t.

To illustrate Assumption 3, suppose a buyer gets a lower level of quality than what the seller's reputation was signalling. She could then leave a very negative rating attempting to correct the seller's reputation in future periods. Similarly, if the outcome earned after transacting is higher than what was expected, a buyer could be incline to publish a very positive rating. We could also imagine a buyer unsure of what rating to give after a transaction and publishing a rating in accordance with the current reputation of the seller.[4] This leads us to the following definition.

Definition 15 (Reporting Types)**.** *Reporting* Φ *is defined as* truthful *if* $\Phi(q,\mu) = q$ *and* biased *if* $\Phi(q,\mu) \neq q$.

Put otherwise, when the rating left by a buyer does not correspond to the level of quality she observed while transacting, reporting is said to be biased.[5] Biased reporting leads to insincere ratings. The two next assumptions complete the presentation of the setting.

Assumption 4 (Consistency)**.** $\Phi(q,q) = q$.

Assumption 5 (Continuity)**.** $F(\cdot)$ *is Lipschitz continuous. There exists a Lipschitz constant* $K_\Phi > 0$ *such as*

$$\mid \Phi(q,\mu) - \Phi(q,\mu') \mid \leq K_\Phi \mid \mu - \mu' \mid, \quad \forall \mu, \mu' \in [0,1]. \tag{3.3}$$

Assumption 4 articulates a trivial but necessary condition to ensure that the reporting structure remains consistent over time. Assumption 5 imposes continuity conditions to simplify the analysis.

[4] As long as the non-decreasing property of $\Phi(q,\mu)$ in q (the observed transaction's quality) is respected. Formally, let $q' > q$, then we must have $\Phi(q',\mu) \geq \Phi(q,\mu)$, $\forall q, q' \in [0,1]$.

[5] I.e. buyers take into account the current reputation of the seller when publishing their ratings.

3.3 Long-Run Reputation Dynamics

The section starts by characterizing the evolution of reputation over time. I then move to the analysis of convergence for both the reputation aggregate statistic and the distribution of ratings. Such a proceeding allows to build a measure that will be used in the remaining of the chapter to evaluate the overall efficiency of the reputation system. Consequences of convergence for the ratings' dispersion are discussed next. I finally illustrate the findings of the section by introducing mechanisms which portray the principal reporting behaviours found online.

3.3.1 Reputation dynamics

I commence the analysis by formally expressing the update rule governing the formation of reputation in the marketplace.

Lemma 16. *Seller's reputation evolution over time can be approximated as*

$$\mu_{t+1} = \mu_t + \frac{1}{t+1}\,(r_{t+1} - \mu_t). \tag{3.4}$$

Proof. The proof relies on an application of the strong law result for stochastic sums or Strong Law of Large Numbers (SLLN). For $\{q_1, q_2, ...\}$ being a sequence of i.i.d. random variables, we have a reporting function $\Phi(q_i, \mu_{i-1})$ that generates an i.i.d. random sequence of ratings $R = \{r_1, r_2, ...\}$. I denote α_t the small decreasing step-size parameter or learning rate, i.e. $\alpha_t = \frac{1}{t} \in (0, 1]$. Given (3.1) and the step-size α_t, it is straightforward to remark that $\mu_t \to \mathbb{E}[R]$ almost surely. Following this observation and using a recursive argument, reputation at t can be estimated as

$$\begin{aligned}\mu_t &= (1-\alpha_t)\mu_{t-1} + \alpha_t \Phi(q_t, \mu_{t-1}),\\ &= (1-\alpha_t)\mu_{t-1} + \alpha_t r_t,\\ &= \mu_{t-1} + \alpha_t (r_t - \mu_{t-1}),\\ &= \mu_{t-1} + \frac{1}{t}(r_t - \mu_{t-1}),\end{aligned} \tag{3.5}$$

where (3.5) is consistent with SLLN, the learning rate α satisfying $\sum_{i=1}^{\infty} \alpha_i = \infty$ and $\sum_{i=1}^{\infty} \alpha_i^2 < \infty$. □

Lemma 16 is quite intuitive and highlights how the seller's future reputation depends on his current reputation and on the next rating he will earn. Interestingly, we note that the impact of buyers' ratings on the reputation of the seller decays over time, indicating a possible convergence to a fixed reputation score.

3.3.2 Convergence in probability

Proposition 14 directly follows from the previous observation and identifies a fixed score towards which seller's reputation converges over time.

Proposition 14. *$\mu_t \to \mu^*$ w.p.1 when $t \to \infty$.*

Proof. The proof is a generalization of the stochastic approximation algorithm. Time is discrete, implying that the stochastic process can be written as

$$\mu_{t+1} - \mu_t = \alpha_{t+1} V_{t+1}, \quad \forall t > 0,$$

where μ_t takes its value in the Euclidean space, V_{t+1} is a random variable and $\alpha_{t+1} \in (0, 1]$ is the decreasing learning rate. Specifically, μ_t represents the reputation score of the seller changing over time and

$$
\begin{aligned}
V_{t+1} &= r_{t+1} - \mu_t, \\
&= \Phi(q_{t+1}, \mu_t) - \mu_t, \\
&\equiv v(q_{t+1}, \mu_t).
\end{aligned}
$$

At each period, a new quality level (q_{t+1}) is observed that causes previous reputation score (μ_t) to be updated according to an 'algorithm' characterized by the function $v(\cdot)$. $\{q_{t+1}\}$ being a sequence of i.i.d. random variables, $v(\cdot)$ can be written in the form

$$
\begin{aligned}
v(q_{t+1}, \mu_t) &= \mathbb{E}[v(q_{t+1}, \mu_t) \mid \mu_t] + \xi_t, \\
&= w(\mu_t) + \xi_t,
\end{aligned}
$$

where the mean field $w(\cdot)$ is the function of μ we want to cancel and ξ_t is a random perturbation with 0-mean noise. Following Robbins and Monro (1951), we know that μ_t converges almost surely to a fixed value μ^* if the three conditions below are satisfied:

C1) $\sum_{i=1}^{\infty} \alpha_i = \infty$ and $\sum_{i=1}^{\infty} \alpha_i^2 < \infty$;

C2) $\sigma^2(\mu) \leq \lambda + \gamma \mid \mu \mid^2$, where $\sigma^2(\mu)$ is the variance of μ and λ and γ are some finite constants;

C3) $\exists \mu^* : \sup_{\frac{1}{\varepsilon} \geq |\mu - \mu^*| \geq \varepsilon} (\mu - \mu^*) w(\mu) < 0, \forall \varepsilon > 0.$

(C1) directly follows from the decreasing form of our learning rate α. (C2) follows from the fact that μ is bounded. More formally,

$$
\sigma^2(\mu) \leq 1 \quad \text{and} \quad \lim_{t \to \infty} \frac{1}{t} \sum_{i=1}^{t} \xi_i = 0.
$$

For (C3), fix $\varepsilon > 0$. If $\mu^* \geq \varepsilon$ and $\mu \in [0, \mu^* - \varepsilon]$, then

$$\begin{aligned}(\mu - \mu^*)w(\mu) &\overset{(a)}{\leq} (\mu - \mu^*)w(\mu^* - \varepsilon),\\ &\overset{(b)}{\leq} -\varepsilon w(\mu^* - \varepsilon),\end{aligned}$$

where inequality (a) comes from the fact that $w(\mu)$ is decreasing on $[0, 1]$, and (b) directly follows from $\mu \in [0, \mu^* - \varepsilon]$. Now, suppose $\mu^* \leq 1 - \varepsilon$ and $\mu \in [\mu^* + \varepsilon, 1]$, then

$$\begin{aligned}(\mu - \mu^*)w(\mu) &\leq (\mu - \mu^*)w(\mu^* + \varepsilon),\\ &\leq \varepsilon w(\mu^* + \varepsilon).\end{aligned}$$

Finally, we have

$$\sup_{\mu \in [0,1]:|\mu - \mu^*| \geq \varepsilon} (\mu - \mu^*)w(\mu) \leq -\varepsilon \, min\{w(\mu^* - \varepsilon), -w(\mu^* + \varepsilon)\} < 0, \quad \forall \varepsilon > 0.$$

Hence, (C3) is satisfied and proof is complete. □

Proposition 14 states that for any reporting behaviour, the seller's reputation will always converge with probability 1 to a specific value in the long-run. More particularly, as the number of ratings increases, the average seller's reputation score will be such that $\mu_t \approx \mu^*$.

Note that when reporting is biased, μ^* does not correspond to the seller's average quality $\mathbb{E}[Q]$. The difference between the two latter measures remains constant over time and thus captures entirely the *reporting bias* which may exist in the marketplace. A formal definition of the aforementioned bias follows.

Definition 17 (Reporting Bias)**.** *The reporting bias present in the marketplace is defined as the difference between the long-term average reputation score of the* *seller and the average transactions' quality proposed by the seller, that is*

$$\Delta :=\mid \mu^* - \mathbb{E}[Q] \mid \in [0,1]. \tag{3.6}$$

It results, quite directly, that the reporting bias represents a natural candidate to assess the reputation system's efficiency, a fully efficient mechanism revealing $\mu^* = \mathbb{E}[Q]$, i.e. $\Delta = 0$. More particularly, μ^* being the solution of $\mathbb{E}[\Phi(Q,\mu)] = \mu$, Δ characterizes the information bias caused by a reporting behaviour Φ.[6]

Convergence in probability being established, I now consider the case of convergence in distribution (denoted $\rightarrow_d$) in order to ensure that buyers adopt a given reporting behaviour Φ^* in the long-run, where $\Phi^* := \Phi(\cdot, \mu^*)$.

3.3.3 Convergence in distribution

Since seller's reputation always converges to a fixed score in the long-run, the next natural question is to investigate whether a converging phenomenon is also observed for the distribution of ratings published by buyers.

Proposition 15. $r_t \rightarrow_d \Phi^*$ *when* $t \rightarrow \infty$.

Proof. Define $\Phi^* := \Phi(q_t, \mu^*)$. We want to show that $\mid r_t - \Phi^* \mid \rightarrow_p 0$. Fix $\varepsilon > 0$, then

$$\begin{aligned} \mathbb{P}(\mid r_t - \Phi^* \mid \geq \varepsilon) &= \mathbb{P}(\mid \Phi(q_t, \mu_{t-1}) - \Phi^* \mid \geq \varepsilon), \\ &\overset{(a)}{\leq} \frac{1}{\varepsilon}\mathbb{E}[\mid \Phi(q_t, \mu_{t-1}) - \Phi^* \mid], \end{aligned} \tag{3.7}$$

where (a) is the Markov's inequality linking probabilities to expectations. Gen-

[6]The relationship existing between Φ and Δ is examined in section 3.4.

eralizing the Portmanteau Lemma, we look for the upper bound of (3.7), i.e.,

$$\begin{aligned}\mathbb{E}[\mid \Phi(q_t,\mu_{t-1}) - \Phi^* \mid] &\leq \mathbb{P}(q_t \in [\mu_{t-1} \wedge \mu^*, \mu_{t-1} \vee \mu^*]) \\ &\quad + \mathbb{E}[\mid \Phi(q_t,\mu_{t-1}) - \Phi^* \mid \mathbf{1}\{q_t \notin [\mu_{t-1} \wedge \mu^*, \mu_{t-1} \vee \mu^*]\}], \\ &\overset{(b)}{\leq} \underbrace{\mathbb{E}[F(\mu_{t-1} \vee \mu^*)] - \mathbb{E}[F(\mu_{t-1} \wedge \mu^*)]}_{\text{A}} + K_\Phi \underbrace{\mathbb{E}[\mid \mu_{t-1} - \mu^* \mid]}_{B},\end{aligned}$$

where (b) results from Assumption 5, K_Φ being the Lipschitz constant corresponding to the reporting function Φ. Applying the Lebesgue's Dominated Convergence Theorem and Proposition 3.5, we see that $\mathbb{E}[\mu_{t-1}] \to \mathbb{E}[\mu^*]$ as $t \to \infty$, implying that $B \to 0$ when $t \to \infty$. In addition, $F(\cdot)$ being continuous, the mapping theorem for convergence in distribution induces that $A \to 0$ as well. As a direct result, we remark that $\mathbb{P}(\mid r_t - \Phi^* \mid \geq \varepsilon) \to 0$, and then $\mid r_t - \Phi^* \mid \to_p 0$. □

In words, the distribution of ratings left by buyers is expected to stabilize over time, converging to Φ^*. This result hints that attempts to manipulate the reputation of a seller can be successful and their imprints vary depending on the form of the reporting structure, especially if manipulation occurs early in the seller's life.[7]

Besides, Proposition 15 provides some theoretical support to explain the existence of U-shaped or J-shaped distributions of ratings empirically observed within online reputation systems.[8] Specifically, for a certain class of reporting behaviours $\Phi(\cdot)$, ratings have a tendency to put mass away from μ^*, thus forming U-shaped distributions. Corollary 18 formalizes the latter argument.

Corollary 18. *If* $\mid \Phi^* - \mu^* \mid > \mid q - \mu^* \mid \; \forall q, \mu^* \in [0,1]$, *then the ratings tend to move away from* μ^*.

[7]This observation will be developed further in section 3.6.

[8]The empirical literature identified that only very negative (e.g., 1-star grade) or very positive (e.g., 5-star grade) ratings coexisted in online marketplaces. See for instance the "moan and groan" model proposed by Hu et al. (2009) or the experimental results of Lafky (2014).

Consider that reporting in the marketplace is such that buyers feature a higher predisposition to leave positive (negative, respectively) ratings when they get a higher (lower) level of quality than expected. Then, the reporting behaviour follows the structure below, i.e.,

$$\Phi(q,\mu) = \begin{cases} \beta q & \text{if } q < \mu \\ (1-\beta) + \beta q & \text{if } q \geq \mu \end{cases} \tag{3.8}$$

where $\beta \in [0,1]$ parametrizes the granularity of Φ.[9] Quite intuitively, we see that as $\beta \to 0$, the precision of the mechanism decreases and the ratings move away from μ^*, forming denser clusters at the two extrema of the spectrum $[0,1]$.[10] Put otherwise, rating scales resting on a limited granularity do not cope well with converging reporting behaviours, thus promoting the disclosure of extreme ratings. While previous works demonstrated that bimodal distributions of ratings could be explained by the *cost of rating* in an online market (Lafky, 2014), the latter reasoning suggests that the role played by the architecture of a reputation system and the reporting behaviours it encourages should not be disregarded and belittled.

I conclude the convergence analysis with an application of the central limit theorem, making then possible the estimation of the deviations existing between the seller's reputation at t, μ_t, and the average long-term reputation score μ^*.

Proposition 16. *$\sqrt{t}(\mu_t - \mu^*) \to_d \mathcal{N}(0,\sigma^2)$ when $t \to \infty$, where $\mathcal{N}(0,\sigma^2)$ means a normal distribution with mean 0 and variance $\sigma^2 = \frac{Var(\Phi^*)}{-(2w'(\mu^*)+1)}$.*

Proof. The proof is an application of the central limit theorem (CLT). For a detailed demonstration, the reader is invited to refer to Chung (1954) or Hodges

[9] For simplicity's sake, the precision of a reporting system is considered to result directly from its granularity.

[10] For its minimum value $\beta = 0$, the reporting mechanism takes a binary form. Note that in practice, the granularity of rating scales used by online reputation systems is often limited, e.g. Amazon employs a 5-star rating system, eBay a quasi-binary one [Positive, Neutral, Negative].

and Lehmann (1956). □

Proposition 16 characterizes the deviations of ratings from μ^* through the variance σ^2. Specifically, it allows to rank different reporting behaviours according to the dispersion they induce within the reputation system. We remark that dispersion increases in the variance of the mechanism Φ itself (the numerator) and decreases in the denominator, which is nothing else that the first derivative of the mean field, where $w'(\mu^*) \leq 0 \ \forall \mu \in [0, 1]$. Put differently, Proposition 16 provides a way to approximate the speed of convergence necessary for a reporting behaviour Φ to reach μ^*.

3.3.4 Reporting Behaviours

At this stage of the development, it seems opportune to introduce reporting mechanisms able to capture the diversity of rating behaviours occurring in 1-sided online reputation systems. To make the analysis tractable, reporting behaviours are divided into three distinct classes, whose respective definitions follow.

Definition 19 (Truthful)**.** *A reporting behaviour is defined as* truthful *if the rating published by a buyer is equal to the transaction's level of quality she observed, i.e.* $\Phi(q, \mu) = q$.

Definition 20 (Correcting)**.** *A reporting behaviour is defined as* correcting *if* $\Phi(q, \mu)$ *is non-increasing in* μ*, the seller's reputation aggregate statistic.*

To illustrate Definition 20, suppose $\mu' > \mu$, then a correcting reporting behaviour induces that $\Phi(q, \mu') \leq \Phi(q, \mu) \ \forall \mu, \mu' \in [0, 1]$. In order to make the latter definition more intuitive, the reader can conceptualize a correcting reporting behaviour as a buyer's decision to correct the seller's reputation by publishing a rating whose magnitude depends on the difference existing between the observed transaction's quality and the expected one (Talwar et al., 2007). For instance, consider the reporting function (3.8) first introduced to epitomise Corollary 18, that is

$$\Phi(q,\mu) = \begin{cases} \beta q & \text{if } q < \mu \\ (1-\beta) + \beta q & \text{if } q \geq \mu \end{cases} \tag{3.9}$$

where $\beta \in [0,1]$. Before we go any further, note that when $\beta = 1$, reporting is truthful and a rating corresponds to the observed quality of the transaction, as defined by Definition 19. However, when $\beta < 1$, a correcting reporting behaviour takes place. Specifically, as $\beta \to 0$, buyers' reporting behaviour becomes less sincere and more concerned by correcting the seller's reputation, meaning that buyers publish more extreme ratings. The limit case being $\beta = 0$, a situation which corresponds to a binary mechanism where ratings are either positive (1) or negative (0).

Definition 21 (Collective)**.** *A reporting behaviour is defined as* collective *if* $\Phi(q,\mu)$ *is non-decreasing in* μ*, the seller's reputation aggregate statistic.*

Suppose again $\mu' > \mu$, then a collective reporting behaviour now implies that $\Phi(q,\mu') \geq \Phi(q,\mu)\ \forall \mu, \mu' \in [0,1]$. Following Definition 21, a typical collective reporting behaviour Φ could be represented as follows

$$\Phi(q,\mu) = \beta q + (1-\beta)\mu, \tag{3.10}$$

where $\beta \in (0,1)$. As an illustration, fix $\beta = \frac{1}{2}$. In this situation, a buyer who observes a transaction's quality equal to $\frac{2}{5}$ (e.g., 2 stars out of 5) while the seller's current reputation score is $\frac{4}{5}$ will report a rating of $\frac{3}{5}$ instead of a lower grade, adjusting then her appreciation in accordance with the collective ratings left by her predecessors (Muchnik et al., 2013).

Figures 3.1 and 3.2 depict the evolution of the seller's reputation μ_t over

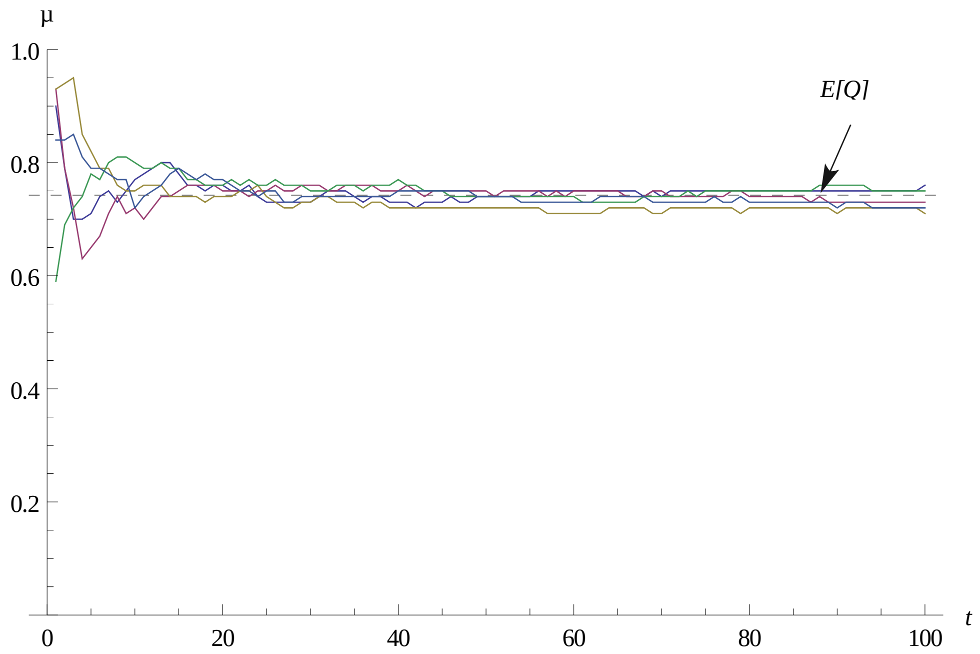

(a) $\Phi(q, \mu) = q$

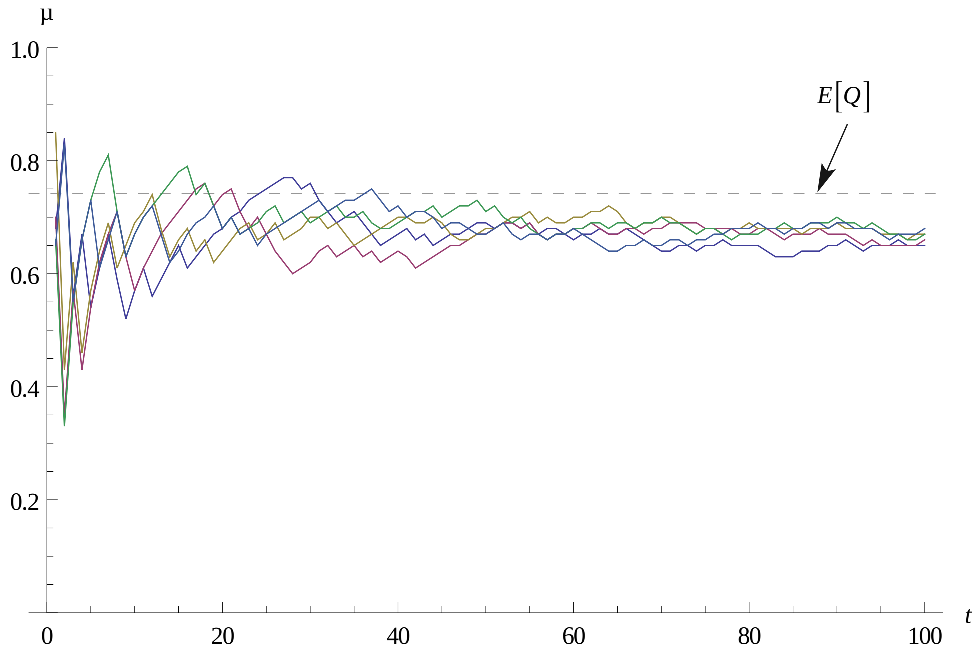

(b) $\Phi(q, \mu) = \begin{cases} 0 & \text{if } q < \mu \\ 1 & \text{if } q \geq \mu \end{cases}$

Figure 3.1: Evolution of the seller's reputation μ_t over time for the correcting reporting behaviour (3.9) with $Q \sim \mathcal{N}(\frac{4}{5}, \frac{1}{5})$. The top figure depicts $\beta = 1$. The bottom figure depicts $\beta = 0$.

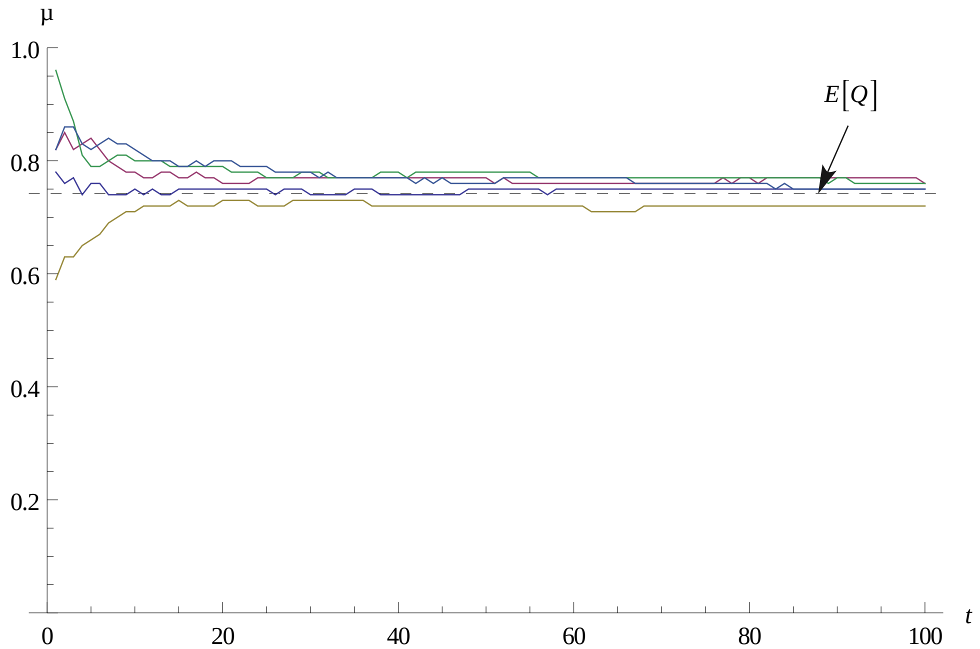

(a) $\Phi(q,\mu) = \frac{q}{2} + \frac{\mu}{2}$

Figure 3.2: Evolution of the seller's reputation μ_t over time for the collective reporting behaviour (3.10) with $Q \sim \mathcal{N}(\frac{4}{5}, \frac{1}{5})$ and $\beta = \frac{1}{2}$.

time (100 first periods) for the two reporting mechanisms (3.9) and (3.10), where the cumulative distributive function $F(\cdot)$ is a truncated normal distribution, i.e. $Q \sim (\frac{4}{5}, \frac{1}{5})$ with $Q \in [0, 1]$.

As predicted by Proposition 14, it is straightforward to observe that all the random walks converge to a limit point μ^* close to $\mathbb{E}[Q]$, where the mean of Q is represented by a dashed horizontal line in the different graphs.[11]

Figure 3.1a) depicts a truthful reporting behaviour. It is direct to remark that the reporting bias Δ is almost 0 and that convergence towards the limit reputation score μ^* occurs faster than for the two other mechanisms detailed in Figures 3.1b) and 3.2a).

Figure 3.1b) represents a correcting reporting behaviour with a binary rating scale. Besides a significant reporting bias, it is interesting to note that such a Φ induces a limit reputation score which steadily underestimates the true expected

[11] I consider a truncated normal distribution to respect $Q \in [0, 1]$, then $\mathbb{E}[Q] \simeq 0.74251 < \frac{4}{5}$.

quality of the seller, $\mathbb{E}[Q]$.

Finally, Figure 3.2a) characterizes a mild collective reporting behaviour. We observe that reporting bias is contained and convergence towards μ^* relatively smooth and swift.

3.4 Reporting Bias

This part of the analysis focuses on the reporting bias, Δ. I first investigate which types of behaviours lead to a reporting bias in the marketplace. Then, I look for the maximum rating bias attainable and the consequences it can have for the reputation system's efficiency.

3.4.1 Existence Condition

To understand how reporting behaviours may affect the reliability of a reputation system, we must first examine which classes of Φ produce a bias larger than 0. This is the very purpose of Proposition 17.

Proposition 17. *If Φ is non-linear, then there exists a reporting bias, $\Delta > 0$.*

Proof. Suppose seller's reputation score converges to μ^* for any distribution $F(\cdot)$. Then the reporting bias Δ becomes

$$\begin{aligned}\Delta &=\mid \mu^* - \mathbb{E}[Q] \mid, \\ &\stackrel{(a)}{=}\mid \mathbb{E}[\Phi(Q,\mu^*)] - \mathbb{E}[Q] \mid, \\ &\stackrel{(b)}{=} 0,\end{aligned}$$

where (a) results from the continuity of $F(\cdot)$ which implies the continuity of

$\mathbb{E}[\Phi(\cdot,\mu)]$ w.r.t. μ, and (b) is true if and only if

$$\begin{aligned}\mathbb{E}[\Phi(Q,\mu^*)] &= \mathbb{E}[\Phi(Q,\mathbb{E}[Q])],\\ &\overset{(c)}{=} \mathbb{E}[Q].\end{aligned}$$

For equality (c) to be verified, we must have

$$\mathbb{E}[\Phi(Q,\mathbb{E}[Q])] \overset{(d)}{=} \Phi(\mathbb{E}[Q],\mathbb{E}[Q]),$$

which follows from Assumption 4. Jensen showed that equality (d) only holds if Φ is linear. Hence, for any non-linear function Φ, the reporting bias Δ will be larger than 0. □

In words, for any non-linear reporting behaviour, ratings are not truthful and reporting bias is present in the marketplace. To illustrate Proposition 17, consider the reporting function Φ introduced at (3.9). Given this behaviour, $w(\mu)$ can be expressed as

$$\begin{aligned}w(\mu) &= \mathbb{E}[\Phi(Q,\mu)] - \mu,\\ &= \mathbb{E}[\beta Q\ \mathbf{1}\{Q<\mu\} + (\beta Q + 1 - \beta)\ \mathbf{1}\{Q \geq \mu\}] - \mu,\\ &= \mathbb{E}[\beta Q + (1-\beta)\ \mathbf{1}\{Q\geq\mu\}] - \mu,\\ &= \beta\mathbb{E}[Q] + (1-\beta)\mathbb{P}(Q\geq\mu) - \mu. \end{aligned} \tag{3.11}$$

Let $\beta = 1$, then (3.11) becomes $w(\mu) = \mathbb{E}[Q] - \mu$. In this case, Φ is linear and it is trivial to remark that $\mu^* = \mathbb{E}[Q]$.[12] Hence and quite intuitively, truthful reporting behaviours do not generate any reporting bias.

[12] μ^* is the solution of $\mathbb{E}[\Phi(Q,\mu)] = \mu$.

Now, let $\beta = 0$, (3.11) simplifies to

$$\begin{aligned} w(\mu) &= \mathbb{P}(Q \geq \mu) - \mu \\ &\equiv \overline{F}(\mu) - \mu, \end{aligned}$$

where $\overline{F}(\mathbb{E}[Q]) \neq \mathbb{E}[Q]$, implying thus $\Delta = \mid \mu^* - \mathbb{E}[Q] \mid > 0$.[13] This simple example adumbrates that the structure and properties of a reporting behaviour play an important role in determining the scope of the information bias existing within the reputation system.

3.4.2 Maximum Bias

Proposition 17 showed that the linearity of Φ prevents the formation of a reporting bias within the marketplace. As should be expected, such a finding suggests that truthful reporting behaviours propose a sound basis to build efficient reputation systems. I now examine the maximum reporting bias (denoted Δ^{max}) non-linear reporting behaviours can induce. For the sake of clarity and coherence, the analysis differentiates between collective and correcting behaviours.

Proposition 18. $\Delta^{max} = 1$ *for collective reporting behaviours.*

Proof. Suppose a collective reporting behaviour of the following form

$$\Phi(q, \mu) = \beta q + (1 - \beta) \min(q, \mu),$$

where $\beta \in (0, 1)$. Let $Q \sim \text{Beta}(x, 1)$, implying a mean $\mathbb{E}[Q] = \frac{x}{x+1}$ with $x > 0$. Fix $Z := \min(Q, \mu)$ and denote $\mathcal{F}_Z$ its corresponding σ-algebra. In order to look for the maximum reporting bias Δ^{max}, I compute $w(\mu)$ corresponding to Φ, that

[13] Note that this result holds in general, the only exception being Q uniformly distributed in $[0, 1]$.

is

$$
\begin{aligned}
w(\mu) &= \mathbb{E}[\Phi(Q,\mu)] - \mu, \\
&= \beta\mathbb{E}[Q] + (1-\beta)\mathbb{E}[Z \mid \mathcal{F}_Z] - \mu, \\
&= \beta\mathbb{E}[Q] + (1-\beta)\mathbb{E}[Q\mathbf{1}\{Q \le \mu\} \mid \mathcal{F}_Z] + \mu(1 - \mathbb{E}[\mathbf{1}\{Q \le \mu\} \mid \mathcal{F}_Z]) - \mu, \\
&= \beta\frac{x}{x+1} + (1-\beta)(\int_0^\mu q f_Q(q)dq + \mu(1 - \mathbb{P}[Q \le \mu]) - \mu, \\
&= \beta\frac{x}{x+1} + (1-\beta)(\frac{x}{x+1}q^{x+1}\Big|_0^\mu + \mu(1-\mu^x)) - \mu, \\
&= \beta\frac{x}{x+1} + (1-\beta)(\mu - \frac{\mu^{x+1}}{x+1}) - \mu. \qquad (3.12)
\end{aligned}
$$

As $x \to \infty$, it is direct to observe that $\mathbb{E}[Q] \to 1$, implying then $\mu^* \to 0$ for $\beta \to 0$. Hence, the maximum reporting bias attainable is $\Delta^{max} = 1$. □

Proposition 19. $\Delta^{max} = \frac{1}{4}$ *for correcting reporting behaviours.*

Proof. I show that the maximum reporting bias engendered by correcting reporting behaviours cannot be larger than $\frac{1}{4}$. To do so, I focus on the most extreme version of the mechanism presented at 3.9, that is

$$
\Phi(q,\mu) = \begin{cases} 0 & \text{if } q < \mu \\ 1 & \text{if } q \ge \mu \end{cases}
$$

and I determine the general bound on Δ by dividing the latter reporting function into two pieces in order to look for its lower and upper bounds, that is

$$
\Phi_0(q,\mu) = \begin{cases} 0 & \text{if } q < \mu \\ q & \text{if } q \ge \mu \end{cases} \qquad \Phi_1(q,\mu) = \begin{cases} q & \text{if } q \le \mu \\ 1 & \text{if } q > \mu \end{cases}
$$

Part (i) - Suppose $\mathbb{E}[Q] > \mu^*$. Then, $\mid \mathbb{E}[Q] - \mu^* \mid \le \mathbb{E}[Q] - \mu_0^*$, where μ_0^* can be written as follows

$$\begin{aligned} \mu_0^* &= \mathbb{E}[\Phi_0(Q, \mu_0^*)], \\ &= \mathbb{E}[Q \, \mathbf{1}\{Q \geq \mu_0^*\}], \\ &= \mathbb{E}[Q] \, \mathbb{P}(Q \geq \mu_0^*), \\ &\leq 1 - \mathbb{P}(Q < \mu_0^*). \end{aligned} \tag{3.13}$$

Computing the reporting bias Δ corresponding to Φ_0 leads to

$$\begin{aligned} \mathbb{E}[Q] - \mu_0^* &= \mathbb{E}[Q] - \mathbb{E}[\Phi_0(Q, \mu_0^*)], \\ &= \mathbb{E}[Q] - \mathbb{E}[Q \, \mathbf{1}\{Q \geq \mu_0^*\}], \\ &= \mathbb{E}[Q] \, \mathbb{P}(Q < \mu_0^*), \\ &\leq \mu_0^* \, \mathbb{P}(Q < \mu_0^*), \\ &\overset{(a)}{\leq} \mu_0^*(1 - \mu_0^*), \\ &\leq \frac{1}{4} \quad \forall \mu \in [0, 1]. \end{aligned}$$

(a) : thanks to (3.13), we know that $\mathbb{P}(Q < \mu_0^*) \leq 1 - \mu_0^*$.

Part (ii) - Now, suppose $\mathbb{E}[Q] \leq \mu^*$. Then, $\mid \mathbb{E}[Q] - \mu^* \mid \leq \ \mu_1^* - \mathbb{E}[Q]$, where μ_1^* can be written as follows

$$\begin{aligned} \mu_1^* &= \mathbb{E}[\Phi_1(Q, \mu_1^*)], \\ &= \mathbb{E}[Q \, \mathbf{1}\{Q \leq \mu_1^*\} + \ \mathbf{1}\{Q > \mu_1^*\}], \\ &= \mathbb{E}[Q] \, \mathbb{P}(Q \leq \mu_1^*) + \mathbb{P}(Q > \mu_1^*), \\ &\geq \mathbb{P}(Q > \mu_1^*). \end{aligned} \tag{3.14}$$

Then, the reporting bias Δ generated by Φ_1 can be computed as

$$\begin{aligned}
\mu_1^* - \mathbb{E}[Q] &= \mathbb{E}[\Phi_1(Q, \mu_1^*)] - \mathbb{E}[Q], \\
&= \mathbb{E}[Q]\,\mathbb{P}(Q \leq \mu_1^*) + \mathbb{P}(Q > \mu_1^*) - \mathbb{E}[Q], \\
&= \mathbb{E}[Q](1 - \mathbb{P}(Q > \mu_1^*)) + \mathbb{P}(Q > \mu_1^*) - \mathbb{E}[Q], \\
&= \mathbb{P}(Q > \mu_1^*)(1 - \mathbb{E}[Q]), \\
&\overset{(b)}{\leq} \mu_1^*(1 - \mu_1^*), \\
&\leq \frac{1}{4} \quad \forall \mu \in [0, 1].
\end{aligned}$$

(b) : by (3.14), $\mathbb{P}(Q > \mu_1^*) \leq \mu_1^*$ and $\mathbb{E}[Q] > \mu^*$ by assumption. This completes the proof. □

Proposition 18 identifies that some collective behaviours may lead to completely inefficient reputation systems, the limit reputation score μ^* tending to the opposite value of the seller's genuine quality in some cases. Proposition 19 shows that the maximum reporting bias possible when correcting behaviours are in place is better contained and cannot be larger than $\frac{1}{4}$.

The intuition is relatively simple. When buyers comply with ratings left by their predecessors, there is almost no chance to get a rating in the reputation system capable of bringing back the seller's reputation close to its true score. Conversely, when a correcting behaviour prevails in the market, correction occurs as soon as the seller's reputation score diverges too significantly from its legitimate one, and the extent of the reporting bias is thus better controlled and contained within the reputation mechanism.

3.5 Robustness

The convergence analysis of section 3.3 provided a way to assess the efficiency of a reputation system by introducing the concept of reporting bias, Δ. Section 3.4 identified the existence condition for the aforementioned bias and examined its magnitude conditional upon the reporting behaviour in place in the marketplace. In this section, I check whether the measure of efficiency Δ remains relevant and reliable when the setting evolves and complexifies. A first development stages a marketplace with multiple sellers and a second development examines the case of heterogeneous reporting behaviours coexisting within the same marketplace.

3.5.1 Sellers of Different Quality

Following Propositions 18 and 19, we know that reporting behaviours can induce large reporting biases in the reputation system. The purpose of the coming lines is to ensure that even if Δ is substantial in an online market, a seller of a better quality always enjoys an average reputation score higher in the long-run than a seller of a lower quality.

To check for the latter monotonicity in reputation, I modify the original setting by considering that the marketplace can now be inhabited by more than one seller. More particularly, I study the case of two sellers with different distributions of Q.

Proposition 20. *Suppose a marketplace with 2 sellers (X and Y) where $Q_X \sim F_X(\cdot)$ and $Q_Y \sim F_Y(\cdot)$. If $Q_X < Q_Y$, then $\mu^*_X < \mu^*_Y$.*

Proof. Suppose there exist two sellers, called X and Y. Let $Q_X < Q_Y$ denote the fact that Q_X is stochastically smaller than Q_Y (i.e., $F_X \neq F_Y$) and $\mu^*_X := \mu^*(\Phi, F_X)$, $\mu^*_Y := \mu^*(\Phi, F_Y)$. We want to prove that $\mu^*_X < \mu^*_Y$. The proof follows by contradiction, i.e. we want $\mu_X^* > \mu_Y^*$ and fix then

$$\varepsilon = \frac{\mu_X^* - \mu_Y^*}{2} > 0. \tag{3.15}$$

By continuity of $F(\cdot)$, we know that

$$\mu_X^* = \mathbb{E}[\Phi(Q_X, \mu_X^*)],$$
$$\mu_Y^* = \mathbb{E}[\Phi(Q_Y, \mu_Y^*)].$$

Given (3.15), we have

$$\begin{aligned}\mathbb{E}[\Phi(Q_X, \mu_X^* - \varepsilon)] &\overset{(a)}{>} \mu_X^* - \varepsilon \\ &\overset{(b)}{=} \mu_Y^* + \varepsilon \\ &\overset{(c)}{>} \mathbb{E}[\Phi(Q_Y, \mu_Y^* + \varepsilon)] \\ &\overset{(d)}{>} \mathbb{E}[\Phi(Q_X, \mu_Y^* + \varepsilon)].\end{aligned}$$

Inequalities (a) and (c) directly follow from

$$\mathbb{E}[\Phi(Q_X, \mu_X^*)] - \mathbb{E}[\Phi(Q_X, \mu_X^* - \varepsilon)] < \varepsilon,$$

for collective reporting behaviours, and from

$$\mathbb{E}[\Phi(Q_X, \mu_X^* - \varepsilon)] > \mathbb{E}[\Phi(Q_X, \mu_X^*] = \mu_X^* > \mu_X^* - \varepsilon,$$

for correcting reporting behaviours. (b) results from (3.15). (d) is a consequence of Φ non-decreasing in Q, where $Q_X < Q_Y$. Given (b), that is $\mu_Y^* - \varepsilon = \mu_X^* + \varepsilon$, we should have $\mathbb{E}[\Phi(Q_X, \mu_X^* - \varepsilon)] = \mathbb{E}[\Phi(Q_X, \mu_Y^* + \varepsilon)]$ which is contradicted by (d). $\square$

Proposition 20 demonstrates that (1st-order) stochastic dominance is respected by reporting behaviours Φ, identifying so monotonicity in reputation. Put otherwise, if a seller X proposes a transaction's quality lower in distribution than a seller Y, then it is not possible for seller X to get a higher reputation score than Y in the long-run, and this result holds even if the reporting bias is large in the market.

3.5.2 Heterogeneous Reporting Behaviours

The previous analysis considered a setting in which only one type of reporting behaviour Φ was possible in the marketplace. Consider now that two different types of reporting behaviours coexist within the same economy. More particularly, let a proportion $\lambda \in [0,1]$ of buyers rate the seller using a mechanism Φ_λ whilst the remaining buyers follow a behaviour dictated by $\Phi_{1-\lambda}$. I denote μ^*_Λ the average long-run reputation score corresponding to this mixed setting.

Proposition 21. *$\mu_t \to \mu^*_\Lambda$ w.p.1 when $t \to \infty$.*

Proof. Suppose $\lambda \in [0,1]$ independent of Q, and define the heterogeneous reporting behaviour in place in the marketplace as follows

$$\Phi_\Lambda(Q,\mu) := \lambda\Phi_\lambda(Q,\mu) + (1-\lambda)\Phi_{1-\lambda}(Q,\mu).$$

It is direct to show that $w_\lambda(\mu)$ and $w_{1-\lambda}(\mu)$ are both strictly decreasing on $[0,1]$,

$$w_\lambda(\mu) = \mathbb{E}[\Phi_\lambda(Q,\mu)] - \mu,$$
$$w_{1-\lambda}(\mu) = \mathbb{E}[\Phi_{1-\lambda}(Q,\mu)] - \mu,$$

which implies that the heterogeneous mean field $w_\Lambda(\mu)$ is decreasing as well. The conditions for convergence first identified in the proof of Proposition 14 are thus all satisfied. □

Proposition 21 establishes that convergence towards a specific reputation score in the long-run subsists when the marketplace is inhabited by buyers with heterogeneous reporting behaviours. This result thus guarantees that the reporting bias Δ remains a viable measure to assess the efficiency of the reputation system within a mixed behavioural setting.[14]

The next logical question is to investigate the respective weights our two reporting behaviours λ and $(1-\lambda)$ hold in the marketplace. Proposition 22 follows.

Proposition 22. *Suppose that a proportion λ of buyers follows a correcting reporting behaviour while the remaining $(1-\lambda)$ buyers follow a collective behaviour, then*

$$\mid \mu_\Lambda^* - \mu_\lambda^* \mid \leq (1-\lambda) \mid \mu_\lambda^* - \mu_{1-\lambda}^* \mid \quad \forall \lambda \in [0,1]. \tag{3.16}$$

Proof. Define $\overline{\mu}_\Lambda := \lambda \mu_\lambda^* + (1-\lambda)\mu_{1-\lambda}^*$ and suppose that $\mu_\lambda^* \geq \mu_{1-\lambda}^*$, implying $\overline{\mu}_\Lambda \in [\mu_{1-\lambda}^*, \mu_\lambda^*]$. Then, we can write

$$\begin{aligned}
\mathbb{E}[\Phi_\Lambda(Q,\overline{\mu}_\Lambda)] &= \lambda\,\mathbb{E}[\Phi_\lambda(Q,\overline{\mu}_\Lambda)] + (1-\lambda)\mathbb{E}[\Phi_{1-\lambda}(Q,\overline{\mu}_\Lambda)], \\
&\overset{(a)}{\geq} \lambda\,\mathbb{E}[\Phi_\lambda(Q,\mu_\lambda^*)] + (1-\lambda)\mathbb{E}[\Phi_{1-\lambda}(Q,\mu_{1-\lambda}^*)], \\
&= \overline{\mu}_\Lambda.
\end{aligned} \tag{3.17}$$

(a) : the correcting behaviour $\Phi_\lambda(q,\mu)$ is non-increasing while the collective behaviour $\Phi_{1-\lambda}(q,\mu)$ is non-decreasing $\forall q \in [0,1]$.

Following (3.17), we have $\mathbb{E}[\Phi_\Lambda(Q,\overline{\mu}_\Lambda)] \geq \overline{\mu}_\Lambda$. Then, $w(\overline{\mu}_\Lambda) = \mathbb{E}[\Phi_\Lambda(Q,\overline{\mu}_\Lambda)] - \overline{\mu}_\Lambda \geq 0$ and given that the mean field $w(\cdot)$ decreases in μ, we need $\mu_\Lambda^* \geq \overline{\mu}_\Lambda$ to

[14] The reporting bias becomes $\Delta = \mid \mu_\Lambda^* - \mathbb{E}[Q] \mid$.

get $w(\mu^*_\Lambda) = 0$, that is

$$\mathbb{E}[\Phi_\Lambda(Q, \mu^*_\Lambda)] = \mu^*_\Lambda \geq \overline{\mu}_\Lambda.$$

Consider now that $\mu^*_\lambda \leq \overline{\mu}_\Lambda \leq \mu^*_{1-\lambda}$, then

$$\begin{aligned} \mathbb{E}[\Phi_\Lambda(Q, \overline{\mu}_\Lambda)] &= \lambda\, \mathbb{E}[\Phi_\lambda(Q, \overline{\mu}_\Lambda)] + (1-\lambda)\mathbb{E}[\Phi_{1-\lambda}(Q, \overline{\mu}_\Lambda)], \\ &\leq \lambda\, \mathbb{E}[\Phi_\lambda(Q, \mu^*_\lambda)] + (1-\lambda)\mathbb{E}[\Phi_{1-\lambda}(Q, \mu^*_{1-\lambda})], \\ &= \overline{\mu}_\Lambda. \end{aligned} \tag{3.18}$$

Following (3.18), $w(\overline{\mu}_\Lambda) = \mathbb{E}[\Phi_\Lambda(Q, \overline{\mu}_\Lambda)] - \overline{\mu}_\Lambda \leq 0$ and $\mu^*_\Lambda \leq \overline{\mu}_\Lambda$. This completes the proof. □

If the rating published by every single buyer had a comparable effect on μ^*_Λ, then the limit reputation score would be quite trivially given by $\lambda\mu^*_\lambda + (1-\lambda)\mu^*_{1-\lambda}$. Proposition 22 demonstrates that the situation is not so balanced in general, correcting reporting behaviours tending to have a greater influence on μ^*_Λ than other classes of behaviours.

3.6 Extensions

This section proposes to build upon the general setting in order to study how different *scenari* highlighted by the literature can impact the efficiency of a reputation mechanism. A first extension makes sybil attacks possible and analyses how reputation measures are sensible to fake ratings. A second extension introduces costly ratings and examines the effects such a change induces for the reporting bias and the choice of a rating scale. A final extension assesses the relevance of an alternative measure to compute the aggregate statistic μ by using the median instead of the mean.

3.6.1 Sybil Attacks (Fake Ratings)

An increasing amount of anecdotal[15] and empirical evidence suggests that reputation manipulation through fake ratings is substantial in online marketplaces. For instance, Luca and Zervas (2013) highlight that nearly one out of five submitted ratings are classified as fake by Yelp's algorithm. I take into account such a phenomenon in the coming lines by supposing that a seller can now hire a third party service to generate a proportion γ of ratings.

Proposition 23. *(i) μ^* increases in the amount of positive fake ratings γ and (ii) decreases in the amount of negative ones.*

Proof. Part (i) – Suppose the seller can buy or publish a proportion γ of the overall ratings making his reputation. Then, a rating at t, r_t is equal to 1 with probability γ and is given by the reporting mechanism $\Phi(q_t, \mu_{t-1})$ with probability $(1-\gamma)$. More formally, we have

$$r_t \equiv \Phi^\gamma(q_t, \mu_{t-1}) = \begin{cases} 1 & \text{with probability } \gamma \\ \Phi(q_t, \mu_{t-1}) & \text{with probability } 1-\gamma \end{cases} \tag{3.19}$$

Remember that $w(\mu) = \mathbb{E}[\Phi(Q,\mu)] - \mu$, then

$$\begin{aligned} w^\gamma(\mu) &= \mathbb{E}[\Phi^\gamma(Q,\mu)] - \mu, \\ &= \gamma + (1-\gamma)\mathbb{E}[\Phi(Q,\mu)] - \mu, \\ &= \gamma + (1-\gamma)(w(\mu)+\mu) - \mu, \\ &= \gamma(1-\mu) + (1-\gamma)w(\mu). \end{aligned} \tag{3.20}$$

We know that $w(\mu) = 0$ only admits one solution in $[0,1]$, implying the same

[15]In February 2004, due to a software error, Amazon's Canadian website revealed the identities of thousands of people who had anonymously posted book reviews on the United States site. It turned out that an important fraction of those reviews were written by the books' own publishers, authors and competitors. (Harmon, 2004)

behaviour for $w^\gamma(\mu) = 0$. Hence, (3.20) is equal to 0 if and only if

$$\frac{\gamma}{1-\gamma} = -\frac{w(\mu)}{1-\mu}, \tag{3.21}$$

where the LHS of (3.21) increases in γ, and the RHS increases in μ, $w(\cdot)$ being strictly decreasing on $[0, 1]$ with $w(0) \geq 0$ and $w(1) \leq 0$.

Part (ii) – Now, suppose a seller Y can buy a proportion γ of the overall ratings making the reputation of seller X, his competitor within the marketplace. The reporting mechanism becomes

$$r_t \equiv \Phi^\gamma(q_t, \mu_{t-1}) = \begin{cases} 0 & \text{with probability } \gamma \\ \Phi(q_t, \mu_{t-1}) & \text{with probability } 1-\gamma \end{cases} \tag{3.22}$$

Then,

$$\begin{aligned} w^\gamma(\mu) &= \mathbb{E}[\Phi^\gamma(Q, \mu)] - \mu, \\ &= 0 + (1-\gamma)\mathbb{E}[\Phi(Q, \mu)] - \mu, \\ &= (1-\gamma)(w(\mu) + \mu) - \mu, \\ &= (1-\gamma)w(\mu) - \gamma\mu. \end{aligned} \tag{3.23}$$

Reducing to

$$\frac{\gamma}{1-\gamma} = \frac{w(\mu)}{\mu}, \tag{3.24}$$

where the LHS of (3.24) increases in γ, and the RHS decreases in μ. This completes the proof. □

Proposition 23 shows that the reputation score of a seller in the long-run (μ^*) is sensible to manipulation, then suggesting that fake ratings have the potential to significantly harm the reliability of a reputation system independently of the reporting behaviour(s) inhabiting the economy. Specifically, the latter aggregate statistic moves in the direction of the manipulation, and larger the amount of manipulated ratings γ, greater the consequences for the seller's reputation in the long-term.[16]

Ceteris paribus, we expect the effects of reputation manipulation to be more critical for collective reporting behaviours than for correcting ones, especially if manipulation occurs early in the seller's life. Indeed, the probability to have a rating in the reputation system capable of bringing back the seller's reputation close to its true quality is small when a collective reporting behaviour lies in the marketplace. Hence, reputation manipulation is both cheaper in terms of γ and more effective when this type of reporting behaviour exists.

Following the latter remark and given Propositions 18 and 19, it appears interesting for a mechanism designer to encourage correcting reporting behaviours within a marketplace facing a high risk of sybil attacks. The key idea behind the latter recommendation is that such behaviours are better able to contain the combined and harmful effects of reputation manipulation and collective reporting, thus limiting reporting bias.[17] In practice, a possible design recommendation would be to develop an online reputation system which would display side-by-side one positive and one negative rating, thus making less likely the occurrence of collective rating behaviours and promoting more truthful reporting decisions even in markets featuring fake ratings.

[16] Put differently, even within a marketplace populated by correcting raters, it is never possible to increase a seller's reputation by publishing fake negative ratings and expecting an important correcting response afterwards.

[17] $\Delta^{max} = 1$ for collective reporting behaviours (see Proposition 18) while $\Delta^{max} = \frac{1}{4}$ for correcting reporting behaviours (see Proposition 19).

3.6.2 Rating Scale and Rating Cost

Publishing a rating is a (time) costly process. Think of actions such as logging into an online marketplace or entering a rating into a reputation system. Following this observation, I now suppose that a buyer incurs a fixed cost $c > 0$ whenever she publishes a rating within the market. Empirical evidence (Talwar et al., 2007) suggests that the decision of a buyer to submit a rating depends on the gap existing between the transaction's quality the buyer observes and the quality she expected. I take into account the latter observation in the coming analysis by considering a reporting behaviour Φ in which a rating is published whenever the latter gap exceeds the cost of rating, that is

$$\Phi(q, \mu) = \begin{cases} q & \text{if } \mid q - \mu \mid \geq c. \end{cases}$$

Proposition 24 follows.

Proposition 24. *The reporting bias Δ increases in the cost of rating c.*

Proof. Fix $c > 0$. Consider the two limit rating scales possible in a market, that is a continuous scale (subscripted $_\infty$) and a binary one (subscripted $_{01}$).

$$\Phi_\infty(q, \mu) = \begin{cases} q & \text{if } \mid q - \mu \mid \geq c \end{cases} \qquad \Phi_{01}(q, \mu) = \begin{cases} 0 & \text{if } q - \mu < c \\ 1 & \text{if } q - \mu \geq c. \end{cases}$$

I first study the continuous situation, looking for its associated average reputation score in the long-run, μ^*_∞.

$$\begin{aligned} w_\infty(\mu) &= \mathbb{E}[\Phi_\infty(Q, \mu)] - \mu, \\ &= \mathbb{E}[Q\ \mathbf{1}\{\mid Q - \mu \mid \geq c\}] - \mu, \\ &= \mathbb{E}[Q]\ \mathbb{P}(\mid Q - \mu \mid \geq c) - \mu. \end{aligned} \tag{3.25}$$

Following (3.25), we have $\mu^*_\infty = \mathbb{E}[Q]\ \mathbb{P}(\mid Q - \mu^*_\infty \mid \geq c)$. It is straightforward to remark that μ^*_∞ increases in the probability to get a rating in the system, probability which decreases in c, *ceteris paribus*. Hence, the reporting bias $\Delta = \mid \mu^*_\infty - \mathbb{E}[Q] \mid$ decreases in μ^*_∞ and so increases in c. I now focus on the binary situation.

$$\begin{aligned} w_{01}(\mu) &= \mathbb{E}[\Phi_{01}(Q, \mu)] - \mu, \\ &= \mathbb{E}[\mathbf{1}\{Q - \mu \geq c\}] - \mu, \\ &= \mathbb{P}(Q - \mu \geq c) - \mu. \end{aligned} \tag{3.26}$$

Then, $\mu^*_{01} = \mathbb{P}(Q - \mu^*_{01} \geq c)$. Similar comparative statics findings apply for the binary system regarding the role of c. This completes the proof. □

Put otherwise, the cost of rating plays a crucial role in the efficiency of a reputation system, higher c larger will be Δ. Such a finding provides a theoretical basis to support the recent experimental and empirical results identified by the related literature. For instance, Lafky (2014) finds by developing a laboratory experiment that the introduction of a rating cost (even small) has a significant effect on reporting behaviours, leading to fewer and more polarized ratings, and causing inaccurate average ratings for moderate quality products.

The cost of rating thus appears as a crucial parameter for mechanism designers, and innovative ways to minimize it further either technically or through incentive schemes represent a promising avenue for future research. Speaking of which, Proposition 3.25 carries an interesting (normative) finding, as summarized by Corollary 22.

Corollary 22. *If $c \in (0, 1)$ and $\mathbb{E}[Q] < 1$, a fined-grained rating scale faces a larger reporting bias Δ than a coarse-grained one.*

Proof. The proof directly follows from the one of Proposition 24. We know that the reporting behaviour resting on an infinite rating scale leads to $\mu^*_\infty = \mathbb{E}[Q]\,\mathbb{P}(| Q - \mu^*_\infty |\geq c)$, while the one using a binary rating scale converges to $\mu^*_{01} = \mathbb{P}(Q - \mu^*_{01} \geq c)$. Given that $\mathbb{E}[Q] \in [0,1]$ is the same in both systems, we have $\mu^*_\infty \leq \mu^*_{01}$ and $\mu^*_\infty \to \mu^*_{01}$ as $\{\mathbb{E}[Q], c\} \to 1$. □

In short, Corollary 22 establishes a relationship between the granularity of a rating scale and the cost of rating. More particularly, it shows that reporting bias is always larger in reputation mechanisms proposing a high granularity of ratings than ones resting on a low or binary rating scale, $\forall c > 0$. This result thus suggests that coarse-grained rating scales may more effectively control and limit the reporting bias induced by the cost of rating.

A natural explanation for this result is that the range of quality levels that are not graded increases in c.[18] In other words, the cost of rating encourages polarizing reporting behaviours, then causing distorted average reputation scores and larger Δ.[19]

3.6.3 Mean versus Median

A recent strand of the literature coming from Computing Science advocates the use of the median instead of the arithmetic mean to calculate the aggregate statistic μ, see for instance Faltings et al. (2010) or Garcin et al. (2013). The core motivation behind such a switch is that the median is less sensible to outliers and so would provide a more stable and robust way to aggregate and rank ratings. While the aforementioned argument sounds both appealing and intuitive, different results in this chapter tend to indicate that using the median does not represent a viable panacea.

Firstly and quite directly, an aggregate statistic resting on the median would

[18]This observation will be further developed in chapter 4.

[19]A few simulations depicting the evolution of μ w.r.t. c for Φ_∞ and Φ_{01} are presented in Appendix 4.

only make sense for reputation systems which use fine-grained rating scales. Now, Corollaries 18 and 22 showed that such reputation systems tend to feature a high level of polarization, explained in part by the existence of correcting behaviours and the cost of rating. Thus, fine-grained scales are subject to bimodal reporting behaviours and using the median instead of the mean would not improve the efficiency of a reputation system, even the contrary (Mosteller and Tukey, 1977).

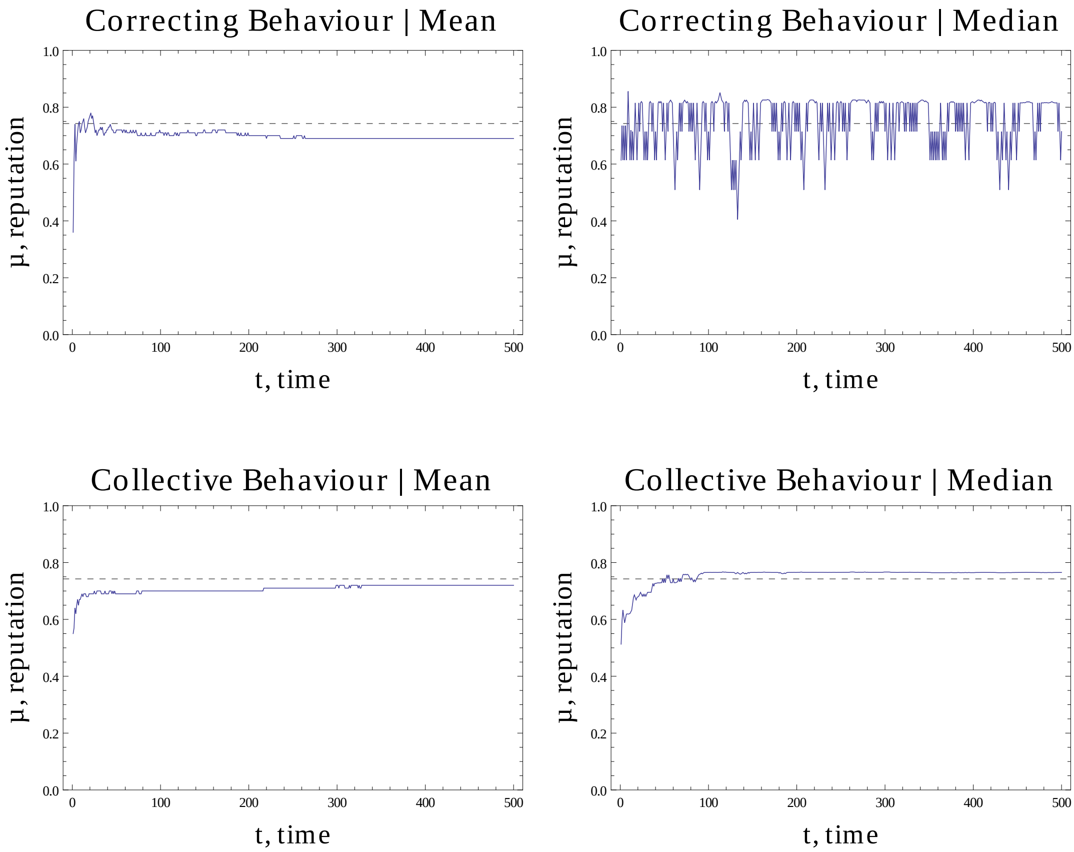

Figure 3.3: Evolution of the seller's reputation μ_t over time for the correcting reporting behaviour (3.9) (top row) and the collective reporting behaviour (3.10) (bottom row) where $Q \sim \mathcal{N}(\frac{4}{5}, \frac{1}{5})$ and $\beta = \frac{1}{2}$. Left column computes the aggregate statistic μ using the mean. Right column uses the median.

Secondly and less intuitively, correcting reporting behaviours feature a higher variance than collective ones, by definition. As shown in Figure 3.3, convergence to μ^* then becomes extremely difficult to attain whenever the median is used to compute a seller's reputation score (μ_t at time t) instead of the mean[20], i.e.,

$$\mathbb{P}[\{r_1, ...r_t\} \leq \mu_t] \geq \frac{1}{2},$$
$$\mathbb{P}[\{r_1, ...r_t\} \geq \mu_t] \geq \frac{1}{2}.$$

More particularly, we observe that the reporting bias remains very high and does not decrease over time.[21] Note that the latter observation does not hold for collective reporting behaviours. However, Proposition 22 showed that correcting behaviours have a larger effect on μ^* than collective ones. It follows that computing μ by using the median should not be advised when correcting reporting behaviours are present in a marketplace since this would damage the efficiency of the reputation mechanism.

Interestingly, the latter reasoning provides some theoretical support to the scarce empirical evidence about the topic: McGlohon et al. (2010) examine several methods to aggregate and rank 5-star ratings from diverse online markets and find that the median performs poorly in comparison to others. They also highlight that the arithmetic mean performs as well or better than a variety of more sophisticated statistics.

3.7 Conclusion

In this chapter, I have proposed a measure, so-called reporting bias, able to assess how different reporting behaviours influence the efficiency of an online reputation system. I have developed an iterative stochastic approximation setting which has made possible the analysis of different models of behaviour while reproducing the way ratings are aggregated online. Following recent findings from the empirical and experimental related literature, I have introduced two specific

[20]See equation (3.1).

[21]See Appendix A.4.3 for a detailed characterization and simulation.

online reporting behaviours – that is 'correcting' and 'collective' behaviours – and studied their effects on the evolution of a seller's reputation over time.

I have demonstrated that the usefulness of new ratings decays over time and that reputation converges to a fixed score as the number of ratings increases in the marketplace. The latter convergence has provided the necessary support for the development of our reporting bias measure. I have also been able to demonstrate that the distribution of ratings left by buyers stabilizes over time, thus presenting some theoretical support for the empirically observed clustered and bimodal distribution of online ratings.

Resting on the aforementioned general results regarding the dynamics of reputation, I have next turned my attention to the analysis of reporting bias. I have showed that when collective reporting behaviours are in place in an online market, that is when buyers report ratings by complying with the ones left by their predecessors, it becomes difficult for a reputation system to bring back the seller's reputation close to its true score, and this can lead to totally inefficient reputation mechanisms. In contrast, the maximum reporting bias generated by buyers who attempt to correct the reputation of a seller remains more limited and so endangers to a lesser extent the trustworthiness of a reputation system. Besides, the robustness analysis has revealed that correcting reporting behaviours have a greater impact than collective ones on the evolution of the seller's reputation in the long run.

The stochastic model presented in this chapter has provided a novel basis to test different empirical schemes and to draw normative conclusions about the design of online reputation systems. I have demonstrated that reporting bias is minimized and mechanism's efficiency maximized when truthful reporting inhabits the economy. A first design recommendation would then be to keep the rating or reporting interface separated from the one that summarizes a seller's reputation in order to encourage truthful ratings and limit the emergence of any correcting or collective stimuli. I have also been able to identify how reputation manipulation through sybil attacks can decrease the efficiency of a reputation system, especially when collective behaviours are present. To restrict the effects of fake ratings, it would appear (again) useful to develop rating interfaces able to encourage truthful or correcting reports in order to 'naturally' wipe out fake ones. Of course, the latter proposal should only be seen as a complement to current filtering algorithms. I have also examined the effects of the rating cost on the reporting bias and found that the efficiency of a reputation mechanism is better guaranteed for coarse-grained rating scales than fine-grained ones as the cost of rating increases in an online marketplace. Finally, I have defended the use of the arithmetic mean against the median, showing that the aggregation and computation of reputation scores with the median would not represent a suitable choice given the underlying bimodal distribution of ratings and the highly probable existence of correcting behaviours in online marketplaces.

This chapter represents a very early attempt to understand how reporting behaviours affect the formation of reputation in online markets and there is still more work to be done. From a modelling viewpoint, it would be especially interesting to propose a framework capable of examining behaviours within 2-sided reputation systems and to assess the reporting bias added by the possibility of bilateral ratings.

Chapter 4

Rating Scales

The Good, the Bad, and the Ambiguous...

4.1 Introduction

This chapter examines how fine-grained rating scales in general, and 5-star scales in particular, can lead to the formation of bimodal distributions of ratings within online marketplaces. Two theoretical frameworks are proposed. A first model focuses on the role played by the cost of rating in the generation of U-shaped distributions. A second model analyses the effects of online experience.

Recent anecdotal evidence from YouTube[1] captures well the challenge faced by reputation mechanism designers regarding the choice of a rating scale. In September 2009, the Google owned video service moved from a standard 5-star rating scale to a binary one (i.e., 'thumbs up/thumbs down') after having observed that the majority of videos had a five-star rating, as shown in Figure 4.1.

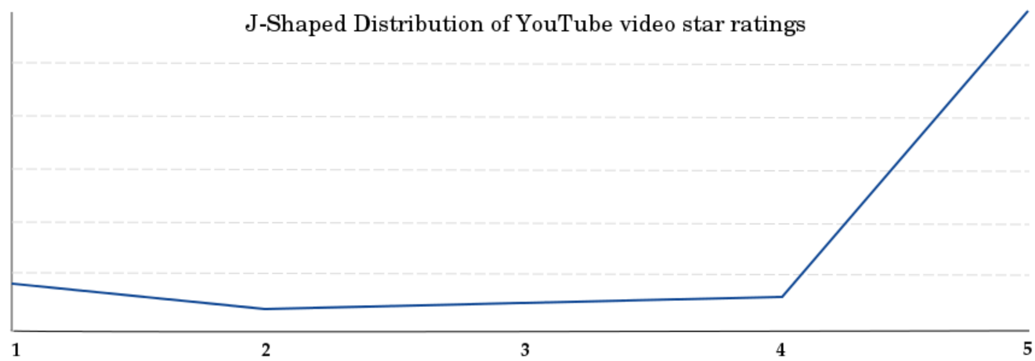

Figure 4.1: *'Seems like when it comes to ratings it's pretty much all or nothing.'* – Shiva Rajaraman, YouTube Product Manager, September 2009.[2]

Now, YouTube is not an isolated case and U-shaped or J-shaped distributions of ratings have already been identified in the literature. Dellarocas and Narayan (2006) show that the distribution of user ratings for movies on Yahoo! Movies[3] is seriously bimodal, almost 50% of the posted ratings are 5-star and approximately 20% 1-star. Hu et al. (2009) find that Amazon ratings of books or CDs follow a

[1] YouTube.com is a popular video-sharing website.
[2] http://youtube-global.blogspot.co.uk/2009/09/five-stars-dominate-ratings.html
[3] https://www.yahoo.com/movies

J-shaped distribution. Talwar et al. (2007) highlight the bimodal distribution of hotel reviews on TripAdvisor. More recently, Lafky (2014) designs an experiment replicating an online marketplace and makes rating costly. The latter author finds that the presence of a rating fee generates a U-shaped distribution of ratings. Finally, the utility-based setting developed by Anderson (1998) to examine the relationship between buyer satisfaction and their engagement in word-of-mouth activities also suggests that a U-curve distribution function of buyer opinions may form.

This chapter contributes to the existing literature in several ways. Firstly, it builds upon recent empirical works to propose two generalisable models aiming at formalizing the determinants and mechanisms which can lead to bimodal distributions of ratings. In peculiar, this work represents one of the first theoretical attempts to formalize the decision of a buyer to publish a rating within an online marketplace. Secondly, the chapter explicitly identifies the existence of a 'blind spot' of unrated transactions, and following empirical evidence, establishes that the amount of neutral ratings lying within an online economy increases in the cost of rating, decreases in altruistic concerns and centres around the prior expectation buyers hold before interacting. Thirdly, this work highlights how fine-grained rating scales promote mid-range unclear ratings and may lead to inefficient reputation mechanisms, passed a certain proportion of ambiguous ratings. Specifically, the development of a model featuring Bayesian agents suffering from confirmatory bias allows to show that the ambiguity made possible by 5-star rating scales can impact the way posteriors are updated and possibly result in wrong posteriors in the long run. Finally, the overall results of the chapter suggest that 5-star rating scales or fine-grained rating scales are more adapted to online marketplaces featuring an important level of adverse selection (i.e., signalling mechanisms) while coarse-grained or binary rating scales should be preferred for economies characterised by moral hazard (i.e., sanctioning mechanisms).

Chapter outline

The remainder of the chapter is organized as follows. Section 4.2 develops a 2-period model and considers that rating a seller at the end of a transaction has an associated cost. Section 4.3 builds an economy inhabited by Bayesian buyers who suffer from confirmatory bias. Last section concludes.

4.2 Rating Cost

Chapter 3 proposed an early study of the interplay existing between the cost of rating and the distribution of reports. The analysis showed that the reporting bias was increasing in the rating cost, and even suggested that the efficiency of a reputation mechanism was better preserved by using a binary rating scale when the cost of reporting was high in the economy. Following the experimental results of Lafky (2014), this section investigates further the relationship between the cost of rating and the distribution of reports. I develop a 2-period model where a potentially large population of buyers can interact with one of two sellers.[4] I demonstrate that a buyer's decision to publish a rating follows a cut-off process whose discontinuity depends on the cost of rating. In particular, I identify that the range of transactions which remain unrated increases in the latter cost, thus generating a kind of 'blind spot' in the marketplace.

4.2.1 Framework

An online marketplace is inhabited by n short-lived buyers, $\mathcal{B} = \{b_1, b_2 ... b_n\}$. At every period of time $t = \{1, 2\}$, a new seller enters the economy. The game starts at t_1 ($t = 1$) when a 2-period lived seller, s_1, is born and chooses his lifetime level of quality $q_1 \in [0, 1]$. The choice of quality q_i is private information and has

[4] A 2-period setting is considered here for analytical convenience, and makes possible the study of intertemporal decisions in a very similar way than atemporal ones.

an associated cost function $\gamma(q_i)$, where $\gamma(0) = 0$, $\gamma'(q_i) > 0$, $\gamma''(q_i) > 0$ and $\lim_{q_i \to 1} \gamma(q_i) = \infty$, for every seller i.[5]

Let buyer b_1 (also the first-buyer, henceforth) be the first to transact in the game at t_1 whilst the $n-1$ remaining buyers all transact in the second period at t_2. For analytical convenience and without loss of generality, let us call them second-buyers and suppose they only discover that the game ends at the end of their period, t_2, after having transacted with a seller.[6]

Once s_1 has chosen his quality q_1, b_1 transacts with him while considering $q_0 > 0$ as a satisfactory or fair level of quality. Once the transaction completed, the first-buyer observes the quality q_1 and may pay a cost $c \in [0,1)$ to publish a rating $r \in \mathbb{R}$ about s_1's true quality. If b_1 decides to report a rating, all other buyers learn it and use this information at t_2. Otherwise, the other buyers have no way to know which seller transacted with b_1 and end up in the same situation than b_1 at t_1. In both cases, all remaining buyers $\{b_2 ... b_n\}$ choose one of the two sellers inhabiting the market at t_2.

4.2.2 Strategies

The coming lines briefly present the strategies of the different inhabitants of our economy, starting with the sellers.

Sellers A seller i maximizes his expected lifetime payoff, that is

$$\Pi_i^{\mathcal{S}} = n_i(\pi - \gamma(q_i)), \tag{4.1}$$

where n_i is the number of buyers who interact with s_i during his lifetime and $\pi > 0$ is the per-buyer profit.

---

[5] As it is standard in economics, we consider the cost function to be increasing and convex in quality. See Tirole (1988).

[6] An equivalent setting would be to consider that the 2 sellers $\{s_1, s_2\}$ are born at t_1 and choose both their lifetime quality at the beginning of the game.

Buyers Buyers are risk-neutral and homogeneous in terms of preferences. They thus maximize their 1-period payoffs. Specifically, b_1 maximizes

$$\Pi_1^{\mathcal{B}} = \phi q_1 + U_R(q_0, q_1, n) - I_R c, \tag{4.2}$$

where ϕ is the marginal benefit of transacting, I_R is an indicator function for whether the first-buyer reported a rating or not, and $U_R(\cdot)$ is the utility earned by b_1 for publishing a rating in the online marketplace.[7] First-buyer's readiness to rate a seller can be explained by altruistic concerns (Fehr and Schmidt, 2006), monetary or 'karma' rewards provided by the economy (Miller et al., 2005) or direct reciprocity (Bolton et al., 2013). Following empirical evidence by Zhang and Zhu (2011) about the effects of group size on public good provisions within an online market such as Wikipedia[8], I suppose that b_1's altruistic concerns increases in the number of buyers n populating the marketplace and following Talwar et al. (2007) that the utility from rating U_R increases in the difference between the expected quality and the experienced one. As is standard in the literature, I assume that U_R is monotonic. For analytical convenience and based on the latter observations, the utility function of rating in the game takes the following form

$$U_R(q_0, q_1, n) = \ln(n) \mid q_1 - q_0 \mid . \tag{4.3}$$

If the fist-buyer publishes a rating, her preferences are in line with those of second-buyers. Then, second-buyers decide to transact with s_1 if the latter seller earned a high rating from b_1, or do not transact with s_1 and pick the new seller, s_2, if s_1 got a low rating. If no rating is left, second-buyers don't know whether b_1 has transacted with s_1 at t_1 and face the same game than b_1. Every second-buyer

[7] For simplicity's sake, prices are normalized to 0 in order to avoid any interpretation of a transaction's price as a signal of its quality.

[8] Zhang and Zhu (2011) show that contributions to Chinese Wikipedia by non-blocked Chinese editors decreased by 42.8% after the government's decision to block access to the collaborative encyclopedia in mainland China, in October 2005.

maximizes

$$\Pi^{\mathcal{B}}_{-1} = \phi q_i. \tag{4.4}$$

Remark that the game ends at t_2, implying that the payoff of a second-buyer's type does not include utility from rating. Likewise, if the first-buyer's payoff was equivalent to (4.4), then it is straightforward to notice that the unique equilibrium in the game would be $q_i = 0$ and no rating, $\forall c > 0$.[9]

4.2.3 Rating Cost and Ratings Distribution

I now analyse the consequences induced by the rating cost c on the reputation system's use and players' strategies. The setting considers an infinitely granular rating scale, implying that b_1 has the possibility to report the exact level of quality q she observed while transacting with s_1.[10] Proposition 25 presents and formalizes the main result of this section.

Proposition 25. *If $c > 0$, then*

i. The range of unrated quality levels increases in c.

ii. The first-buyer plays a cut-off strategy for rating, and reports a negative rating if $q \in [0, \underline{q})$, a positive rating if $q \in (\overline{q}, 1]$ and no rating if

$$q \in [\underline{q}, \overline{q}] \equiv [q_0 - \frac{c}{\ln(n)}, q_0 + \frac{c}{\ln(n)}]. \tag{4.5}$$

Proof. The first buyer decides to report a rating when $c < U_R(q_0, q, n)$. Then, b_1

[9] A parallel can be drawn here with the pure Nash equilibrium of chapter 2.

[10] To simplify notations and because the analysis focuses on the equilibrium strategies of the first seller, q_1 is simply denoted q henceforth.

publishes a low or negative rating if

$$c < \ln(n)(q_0 - \underline{q}) \quad \Leftrightarrow \quad \underline{q} < q_0 - \frac{c}{\ln(n)}.$$

Likewise, b_1 publishes a high or positive rating if

$$c < \ln(n)(\overline{q} - q_0) \quad \Leftrightarrow \quad \overline{q} > q_0 + \frac{c}{\ln(n)}.$$

It follows that the range of quality levels for which a seller remains unrated after a transaction can be expressed as $q \in [\underline{q}, \overline{q}] \equiv [q_0 - \frac{c}{\ln(n)}, q_0 + \frac{c}{\ln(n)}]$. It is straightforward to remark that as $c \to 1$, the gap between $\underline{q}$ and $\overline{q}$ widens $\forall n > 1$. Conversely, if $c \to 0$ then $\underline{q} \to \overline{q}$, $\forall n > 1$. □

Proposition 25 identifies the idiosyncratic rating strategy of the first buyer and represents one of the first theoretical attempts to formalize the decision of a buyer to report a rating within an online marketplace. We remark that the discontinuity of her strategies weakens as the cost of rating tends towards zero. In particular, the cost of rating determines the range of quality levels which are not rated in the marketplace. If $c = 0$, all transactions are rated. In contrast, if $c > 0$, the proportion of unrated qualities grows as c increases. Such a finding helps understand how bimodal distributions can emerge when reputation mechanisms adopt fine-grained rating scales. In particular, the 'blind spot' of unrated quality levels widens in c, implying that a high cost of rating can transform a continuous rating scale into a quasi-binary one.

It is also worth noting that altruism, parametrized in our setting by n, may help mitigate the effect of c on the reporting decision. In other words and *ceteris paribus*, the use of mid-range grades increases in the population size of buyers who will read the seller's reputation. It appears that the latter configuration is more probable in signalling reputation economies than sanctioning ones, thus hinting that fine-grained rating scales could better suit signalling reputation systems. This comes from the observation that online economies featuring a high level of adverse selection (i.e., signalling systems) are generally proposing goods or services which have the potential to support a higher number of interactions than economies featuring a high level of moral hazard (i.e., sanctioning systems).[11]

Going back to the essence of Proposition 25, we notice that the focal point for every rating is q_0. Hence Corollary 23 follows.

Corollary 23. *The range of positive (negative, respectively) ratings available in the marketplace decreases (increases) in* q_0*.*

The quality expected by b_1 determines the nature of the rating to come, if any. More particularly, higher is q_0, smaller will be the range of positive grades available to b_1, as shown in Figure 4.2.

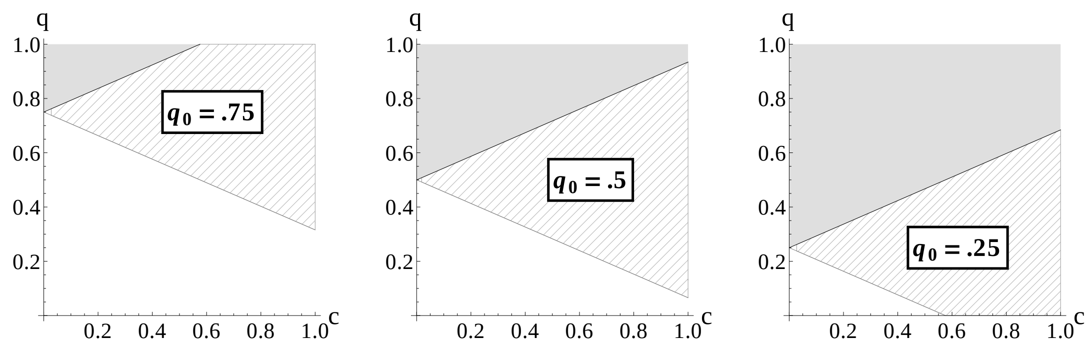

Figure 4.2: Proportion of unrated quality levels, $q \in [q_0 - \frac{c}{\ln(n)}, q_0 + \frac{c}{\ln(n)}]$ (ordinate), in function of the cost of rating, c (abscissa). Hatched area represents the 'blind spot' or unrated levels of quality. Light-gray area depicts the positive ratings and blank area the negative ratings. Left graph depicts $q_0 = \frac{3}{4}$, center graph $q_0 = \frac{1}{2}$ and right graph $q_0 = \frac{1}{4}$. All graphs use $n = 10$.

This finding masks a compelling insight. That is, the pollyannaism empirically observed in online marketplaces could be partly explained by the buyers' high

[11]Intuitively, suppose that the same book is available on Amazon and eBay. Then, one unique reputation profile will be created for the book on Amazon while it will be available from different sellers on eBay.

expectations, in addition to the rating cost. Note that this result also shares some similarities with what the related empirical literature refers to as 'self-selection bias', see for instance Kramer (2007) or Brandes et al. (2013). In short, buyers tend to purchase products or services they like, and then in turn, rate them highly.

I next turn my attention to the equilibrium strategies of the first seller, s_1. Given the previous results and the identification of the so-called 'blind spot', it appears especially interesting to examine when it is optimal for the seller to choose a level of quality that leaves him unrated. Proposition 26 follows.

Proposition 26. *The first seller's equilibrium strategy is to choose the minimum 'unrated' level of quality* $\underline{q}$ *if*

$$\frac{\lambda\pi(n-1)}{2n-\lambda(n+1)} < \gamma(\underline{q}) < \frac{\pi(n-1)}{n+1}, \tag{4.6}$$

where $\lambda = \frac{\gamma(\underline{q})}{\gamma(\overline{q}+\varepsilon)} \in (0,1)$ *and* $\varepsilon > 0$*, which is verified for*

i. $\forall n > 1$ *when* $\lambda \in (0, \frac{1}{2}]$*,*

ii. $\forall \lambda < \frac{n}{n+1}$ *when* $\lambda \in (\frac{1}{2}, 1)$*.*

Proof. The proof follows from the observation that s_1 has three potentially payoff maximizing strategies, i.e. $q = 0$, $q = \underline{q}$ and $q = \overline{q}+\varepsilon \equiv \widehat{q} \in (\overline{q}, 1]$, where ε is small and positive. By Proposition 25, we know that higher c, lower the probability to get a rating in the marketplace. The cost of rating being common knowledge in the game and buyers homogeneous, each seller selects a level of quality that maximizes his expected lifetime payoff. A seller's equilibrium strategy then depends on the parameters of the game. Let us first express the lifetime payoff of a seller who selects a level of quality $q = 0$, that is

$$\Pi_1^{\mathcal{S}}(0) = \pi.$$

Likewise, selecting $q = \widehat{q}$ leads to a lifetime payoff

$$\Pi_1^{\mathcal{S}}(\widehat{q}) = n(\pi - \gamma(\widehat{q})).$$

Finally, choosing $q = \underline{q}$ gives an expected lifetime payoff

$$\begin{aligned} \mathbb{E}[\Pi_1^{\mathcal{S}}(\underline{q})] &= \pi - \gamma(\underline{q}) + \frac{n-1}{2}(\pi - \gamma(\underline{q})), \\ &= \frac{n+1}{2}(\pi - \gamma(\underline{q})). \end{aligned} \tag{4.7}$$

For analytical convenience, let $\gamma(\underline{q}) = \lambda\gamma(\widehat{q})$, where $\lambda \in (0,1)$. Then (4.7) can be rewritten as

$$\mathbb{E}[\Pi_1^{\mathcal{S}}(\underline{q})] = \frac{n+1}{2}(\pi - \lambda\gamma(\widehat{q})).$$

It follows that a seller's equilibrium strategy is to pick a high quality $\widehat{q}$ if $\Pi_1^{\mathcal{S}}(\widehat{q}) > \Pi_1^{\mathcal{S}}(0)$ and $\Pi_1^{\mathcal{S}}(\widehat{q}) > \Pi_1^{\mathcal{S}}(\underline{q})$. Solving w.r.t. $\gamma(\widehat{q})$, this is true if $\gamma(\widehat{q}) < \frac{\pi(n-1)}{n}$ and $\gamma(\widehat{q}) < \frac{\pi(n-1)}{2n-\lambda(n+1)}$. Likewise, he does not invest in quality (i.e., $q = 0$) if $\Pi_1^{\mathcal{S}}(0) > \Pi_1^{\mathcal{S}}(\widehat{q})$ and $\Pi_1^{\mathcal{S}}(0) > \Pi_1^{\mathcal{S}}(\underline{q})$, which is verified for $\gamma(\widehat{q}) > \frac{\pi(n-1)}{n}$ and $\gamma(\widehat{q}) > \frac{\pi(n-1)}{\lambda(n+1)}$. Finally, he selects the 'unrated' quality $\underline{q}$ if $\Pi_1^{\mathcal{S}}(\underline{q}) > \Pi_1^{\mathcal{S}}(0)$ and $\Pi_1^{\mathcal{S}}(\underline{q}) > \Pi_1^{\mathcal{S}}(\widehat{q})$, conditions which are satisfied for

$$\begin{aligned} \frac{\pi(n-1)}{2n-\lambda(n+1)} &< \gamma(\widehat{q}) < \frac{\pi(n-1)}{\lambda(n+1)}, \\ \frac{\lambda\pi(n-1)}{2n-\lambda(n+1)} &< \gamma(\underline{q}) < \frac{\pi(n-1)}{n+1}, \end{aligned} \tag{4.8}$$

where the strict inequality (4.8) holds for $\{n > 1 \wedge \lambda \in (0, \frac{1}{2}]\} \vee \{\lambda < \frac{n}{n+1} \wedge \lambda \in (\frac{1}{2}, 1)\}$. This completes the proof. □

It is straightforward to observe that the decision of s_1 to select the minimum unrated quality level $\underline{q}$ at the beginning of his life depends on the cost differential

λ and the number of buyers n. Independently of n, the probability for the first-seller to choose $\underline{q}$ as his equilibrium strategy tends to 1 as $\lambda \to 0$. The latter strategy becomes independent of λ as $n \to \infty$, implying $\frac{n}{n+1} \to 1$.

Put differently, by featuring a rating cost $c > 0$, a reputation mechanism can only guarantee a level of quality $\underline{q}$ when the economy is large and/or when the seller's investment necessary to get a positive rating, $\gamma(\widehat{q})$, is too important in comparison to the one needed to get a neutral rating (i.e., no rating). Interestingly and following Proposition 25, we know that the level of quality $\underline{q}$ decreases in c, thus implying that a high rating cost tends to encourage low quality interactions.[12] However, we also know that a high n can increase altruistic concerns in the marketplace and then decrease artificially the rating cost, c. Hence, we identify the existence of a compensating effect between c and n.

4.3 Online Experience and Confirmation Bias

In this section, I turn my attention to the role played by online experience in the generation of U-shaped distributions of ratings. In particular, I consider a specific cognitive bias, known as confirmation bias (or confirmatory bias), which can lead individuals to misinterpret new information when their past experience generated a strong belief about a state of the world (Yoon and Gurhan-Canli, 2012).

The key idea behind my approach is to consider that fine-grained rating scales enable buyers to use mid-range grades to report ambiguous or unclear levels of quality. The section introduces long-lived Bayesian buyers who learn under ambiguity and form their beliefs about the possible true state of the world by interacting each period of their lives in the same online market. Such a proceeding allows to examine how online experience and the confirmatory bias it induces can

[12]Note that for $c = 0$ and $n \to \infty$, s_1 would always select the high quality $\widehat{q}$, provided that his net benefit of interacting is larger than 0, that is $\pi > \gamma(\widehat{q})$.

lead to bimodal distributions of ratings while there are more than two possible levels of quality (Zimper and Ludwig, 2009).

4.3.1 Framework

An online marketplace features a reputation system which rests on a 5-star rating scale. The economy is inhabited by a large population of risk-neutral short-lived sellers and long-lived buyers who interact at every time period $t \in \{0, 1, 2, ...\}$.[13] More particularly, buyers transact an infinitely number of periods while sellers transact only once. Such a change in configuration allows to focus on the transactions' history of buyers and is made possible by the introduction of a setting featuring adverse selection. In this new framework, reputation is thus formed and attached to the marketplace and not the sellers, specifically. Once a transaction has been completed, a buyer reports a rating about the transaction's quality by choosing one of the possible grades. For analytical convenience, I consider henceforth that the 5-star rating scale reduces to a 3-category measure in which the quality of a transaction is rated as 'Bad', 'Good' or 'Ambiguous' – as shown in Figure 4.3.

$$\{\underbrace{1,2}_{\text{Bad}} \quad \overbrace{3}^{\text{Ambiguous}} \quad \underbrace{4,5}_{\text{Good}}\} \quad \vee \quad \{\underbrace{1}_{\text{Bad}} \quad \overbrace{2,3,4}^{\text{Ambiguous}} \quad \underbrace{5}_{\text{Good}}\}$$

$$\Downarrow$$

$$\{\ B,\ A,\ G\ \}$$

Figure 4.3: Possible mappings to transform a 5-star rating scale into a 3-category grading system.

Suppose that at $t = 0$, Nature chooses a state $\eta = \{B, G\}$ which determines whether the marketplace attracts a majority of Bad or Good sellers. After a

[13] As in previous chapters, sellers are male and buyers are female. The size of the population is supposed large enough on both sides to support a repeated random matching protocol.

transaction, a buyer i experiences three possible levels of quality, $q_i \in \{b, a, g\}$. Let 'b' indicate a bad or low transaction's quality, 'g' a good or high quality, and 'a' an ambiguous or unclear level of quality.

Quality levels observed by buyers are independent over time. More particularly, suppose that a transaction's quality is interpreted as ambiguous with a probability $\alpha \in [0, 1)$ independent of the state η, and as unambiguous with probability $1 - \alpha$. In online marketplaces such as Amazon or eBay, α could epitomize actions not controlled by a seller which may make more difficult for buyers to discern the genuine quality of a transaction, e.g., database errors, third-party errors (payment gateways), shipping delays. Let the probability to experience a good (bad, respectively) transaction's quality be $\beta > \frac{1}{2}$ when $\eta = \{G\}$ ($\eta = \{B\}$, resp.). Buyers all share a common prior $\mu_0 \in (0, 1)$ that $\eta = \{G\}$ when they enter the marketplace and update their beliefs in a Bayesian manner after each transaction with a seller.

The coming analysis proposes to explain the bimodal distribution of ratings observed empirically by adding the possibility of confirmation bias to the framework. More particularly, I consider fully Bayesian buyers whose interpretation of ambiguous levels of quality evolves over time depending on their respective experiences or transactions' histories. I show that polarization emerges naturally as buyers' experience grows. Such a setting then helps understand the crucial role played by a-priori beliefs in the use of ambiguous grades and identifies when ambiguity can make a reputation system inefficient.

Definition 24 (Reputation Mechanism Efficiency)**.** *A reputation mechanism is defined as* efficient *if the buyers' posteriors converge with probability* 1 *to the true state* η *independently of the prior* μ_0.

Before going any further, let us highlight the interplay existing between the measure α and the occurrence of ambiguous ratings in the model.

Remark 25. *The number of ambiguous ratings decreases in α.*

Such a remark is trivial but it encloses an interesting intuition. As the noise due to external causes decreases in the marketplace (e.g., number of third-party services), buyers are able to assess in a clearer way the genuine quality of sellers and then report acuter ratings.

4.3.2 Confirmation bias

Confirmation bias is present within an economy when agents (in our case, buyers) update their posteriors in a way that makes more likely the confirmation of previous held beliefs. In their seminal paper, Rabin and Schrag (1999) develop a theoretical model to analyse confirmatory bias within a Bayesian framework. They show that if confirmation bias exists, then a buyer is under or over-confident when she interprets new signals in comparison to a standard Bayesian updater. Specifically, this bias may lead to wrong posteriors because buyers tend to disregard new contradictory information. For instance, suppose that a buyer experienced a large number of good quality transactions in her first $t-1$ periods, then she believes that $\eta = \{G\}$. Because of the confirmation bias, the latter buyer may neglect a bad quality signal in the next period of her life and not update her posterior accordingly.

The coming lines propose to adopt and adapt the notion of confirmatory bias to our framework. Differently from Rabin and Schrag (1999), I consider that buyers are still fully Bayesian rational agents and update their posteriors accordingly. However, confirmation bias is introduced in the rating decision. Specifically, the rating given by a buyer i after observing an unclear level of quality at time t results directly from i's posterior at t (denoted $\mu_{i,t}$) that the state is $\eta = \{G\}$. Without loss of generality, let the rating scale be uniformly divided into three ranges. Specifically, let a buyer i experiencing a transaction of type 'a' at t

give a bad rating 'B' when $\mu_{i,t} \in [0, \frac{1}{3})$, an ambiguous rating '$A$' when $\mu_{i,t} \in [\frac{1}{3}, \frac{2}{3})$ and a good rating 'G' when $\mu_{i,t} \in [\frac{2}{3}, 1]$. Note that unambiguous quality levels $\{b, g\}$ lead to unambiguous ratings $\{B, G\}$, and are thus independent of $\mu_{i,t}$. Proposition 27 follows quite naturally.

Proposition 27. *Polarization of ratings increases in the buyer's experience.*

Proof. The proof results from the Law of Large Numbers (Doob, 1949). A rational agent in the Bayesian-sense has a posterior that identifies almost surely one of the 2 states η in place in the market as $t \to \infty$ (i.e., as the number of interactions becomes large enough). This results holds given that ambiguous outcomes are uninformative in our framework, their existence being independent of η. Hence, a Bayesian buyer (updater) remembers all the realizations of her past transactions and is able to overlook the polysemous ones. Since different buyers may experience different levels of quality at the same period of time, the chronology of their transactions' histories will differ, leading to different posteriors and polarization as $t \to \infty$. □

Proposition 27 establishes that a buyer's posterior converges with probability 1 to one of the two states $\eta = \{B, G\}$ as the number of transactions increases, implying that her ratings of ambiguous quality levels follow the same direction than her posterior as $t \to \infty$. In other words, higher is the online experience of a buyer, lower is the probability that she reports ambiguous ratings. The intuition is relatively simple. By transacting repeatedly in the marketplace, a buyer learns little by little the level of quality she can expect from sellers. Then, if her past transactions lead her to think that $\eta = \{G\}$, she will be more prone to report a rating in the higher spectrum of the scale whenever she experiences an ambiguous or unclear level of quality in the next period, and vice-versa.

Example To illustrate our Bayesian updating process and the polarization it may induce, let us consider a specific example. Let the true state be $\eta = \{G\}$. The probability to experience a (unambiguous) good transaction's quality is $\beta = \frac{2}{3}$ in the marketplace, and there exists a state-independent probability $\alpha = \frac{1}{4}$ to experience an ambiguous quality whenever a transaction occurs. Suppose that a buyer i starts her life with a prior $\mu_0 = \frac{1}{2}$. When faced with a good quality in the first period, i's posterior becomes

$$\begin{aligned}\mu_{i,1} &= \frac{\mu_0\beta}{\mu_0\beta + (1-\mu_0)(1-\beta)},\\ &= \frac{\frac{1}{2}\times\frac{2}{3}}{\frac{1}{2}\times\frac{2}{3}+\frac{1}{2}\times\frac{1}{3}},\\ &= \frac{2}{3}.\end{aligned}$$

Now suppose that after t transactions in the marketplace, buyer i has a posterior of $\mu_{i,t} = \frac{7}{9}$. At $t+1$, she transacts with a new seller. The transaction completed, i observes an ambiguous quality of a. Given that $\mu_{i,t} > \frac{2}{3}$, then buyer i will report a good rating.

Our Bayesian setting leads to another interesting observation. That is, the prior μ_0 plays a prominent role in the use of a fine-grained rating scale. Hence the following remark:

Remark 26. *The prior μ_0 determines the use of the 'ambiguous' grade(s) within a reputation system.*

Let $\eta = \{G\}$, $\alpha = \frac{1}{3}$ and $\mu_0 = \frac{1}{2}$, that is a setting in which buyers enter the marketplace with a prior that puts equal weights on the two possible states. Then, it is straightforward to remark that there exists a higher probability for a buyer to report an ambiguous rating early in her life when her prior is $\mu_0 = \frac{1}{2}$ than if it were closer to 1, *ceteris paribus*. Specifically, larger is the mismatch between μ_0 and η, greater will be the proportion of ambiguous ratings reported in the market. In practice, such a scenario seems especially plausible for new online marketplaces.

I now propose to examine how the ambiguity α can impact the reputation mechanism's efficiency. Proposition 28 follows.

Proposition 28. *If* $\alpha > 1 - \frac{\beta^{-1}}{2}$, *then the reputation mechanism becomes inefficient.*

Proof. Suppose that $\eta = \{G\}$. By Proposition 27, we know that $\mu_{i,t}$ converges with probability 1 to one of the two possible states. Here, an efficient mechanism should guarantee convergence towards $\mu_{i,t} \to 1$. Denote β^* (β^{**}, respectively) the probability that a buyer i experiences a quality level confirming her belief that one state is more likely when in fact the other state is true (when in fact it is true). We thus have $\beta^* = (1-\beta) + \alpha\beta$ and $\beta^{**} = \beta + \alpha(1-\beta)$. For the posterior to converge to 0 instead of 1 as $t \to \infty$, we must have $\beta^* > \frac{1}{2}$ and the result follows

$$(1-\beta) + \alpha > \frac{1}{2} \quad \Leftrightarrow \quad \alpha > 1 - \frac{1}{2\beta}.$$

Symmetric proof holds. □

In other words and as should be expected, the efficiency of a reputation system decreases in the amount of ambiguous transactions featured by the marketplace. Specifically, if a non-trivial proportion of experiences is open to interpretation, then there exists a risk for buyers to not update their posteriors in the right direction of η. This being equivalent to say that the reputation mechanism becomes totally inefficient.[14]

[14] A parallel can be drawn here with the results of chapter 2 concerning the imperfect monitoring case.

4.4 Conclusion

In this chapter, I have proposed two theoretical (and generalisable) frameworks in order to analyse and explain the formation of U-shaped distributions of ratings in online markets.

Resting on recent empirical and experimental works, I have built a 2-period model which represents an early attempt to formalize the decision of a buyer to report a rating within an online marketplace. I have considered altruistic concerns and correcting behaviours as the major predictors of a buyer's rating decision, and identified a blind spot of unrated transactions whose magnitude increases in the rating cost and decreases in the number of buyers inhabiting the economy. At equivalent cost of rating, the analysis has suggested that fine-grained rating scales would better suit online markets featuring signalling reputation mechanisms because such economies tend to attract more readers per reputation profile.

Considering Bayesian agents suffering from confirmatory bias, I have then been able to demonstrate that the ambiguity made possible by 5-star rating scales can impact the way beliefs are updated and possibly lead to wrong posteriors in the long run, depending on the amount of underlying ambiguity present in the market.

Finally, this chapter has highlighted the crucial role played by prior expectations, showing that they are the focal points of the blind spot and determine the use of the ambiguous grades within a rating scale. Following the latter results, it appears particularly interesting to develop experiments able to disentangle the effects of prior beliefs on the choice of a rating.

Appendix A

Appendices

A.1 eBay Feedback Reputation System

The online marketplace 'eBay.com' facilitates trade between buyers and sellers, principally through 2nd-price auction. Once a transaction is completed, the buyer can give the seller a grade of +1 (positive), 0 (neutral), or −1 (negative) while the seller can only give the buyer a positive or a neutral grade. Both the buyer and the seller may supplement their rating with any textual comments. In eBay, several aggregates are used to measure the reputation and experience of an agent, including (*i*) the difference between the number of positive and negative feedback ratings, (*ii*), the percentage of positive feedback ratings, (*iii*) the date of registration, and (*iv*) a summary of the most recent feedback ratings.

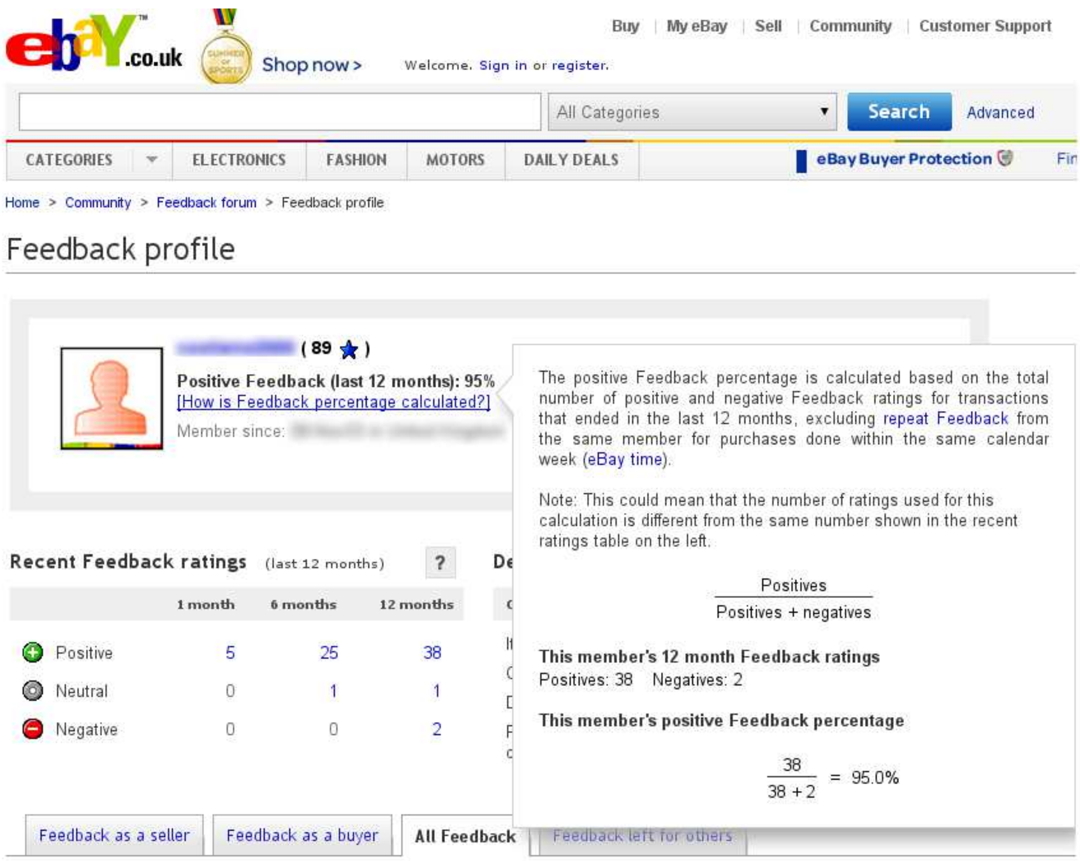

Figure A.1: The eBay Feedback percentage is calculated by dividing all positive Feedback earned in the last 12 months by all positive, neutral and negative Feedback.

Note that the measure of reliability (*ii*) is similar to the one presented in chapter 2, that is

$$\mu_t = \frac{h_t}{r},$$

where h_t is the number of positive Feedback present in the seller's history at time t and r the total number of Feedback earned by an eBay seller during the previous 12-month period, as shown in Figure A.1.

A.2 Descriptive Statistics

This appendix presents summary descriptive statistics about approximatively 85% (after cleaning) of the browser extensions and mobile applications available in three different online marketplaces proposing non-rival goods and featuring 1-sided reputation systems, that is

- Chrome Webstore (14'136 extensions)
- Add-ons Mozilla Marketplace, AMO (7'013 extensions)
- Google Play (1'008 applications)

Data has been collected at the beginning of Spring 2012 using two self-developed scrapers written in PHP and resting on parallel cURL requests.

Datasets are freely downloadable online either in .csv or .dta formats at:

florianbersier.com/research/dphil/data.zip [2.1MB]

Table A.1: Summary Statistics Chrome Webstore

Category (users)	# Users	# Ratings	Grade	Pr rating (%)	# Characters
100 - 999	369.4887	11.7634	4.236338	3.653181	82.94984
	235.8706	24.7641	.7222572	3.552058	77.53932
1,000 - 9,999	3,091.996	49.97911	4.277119	1.674447	88.26009
	2,101.232	158.3217	.558666	1.752254	64.4834
10,000 - 99,999	28,882.47	284.4329	4.358282	.965671	89.81798
	20,893.2	615.0735	.4099862	1.106244	39.9074
$\geq$ 100,000	294,664.7	2,073.118	4.340241	.6419882	83.86436
	231,623.9	3,568.758	.4018165	.5287151	33.0717
Overall	7,119.535	70.46858	4.258646	2.8388	85.08675
	43,645.89	513.5715	.6558563	3.154984	71.13017

For each category, the first row indicates means and the second row standard deviations from the mean.

Table A.2: Summary Statistics Add-ons Mozilla Marketplace (AMO)

Category (users)	**# Users**	**# Ratings**	**Grade**	**Pr rating (%)**	**# Characters**
100 - 999	376.4224	6.073591	4.141298	1.929978	185.5366
	242.0573	8.813073	.958699	3.056695	161.8616
1,000 - 9,999	3,510.64	21.84059	4.121151	.7849486	180.1061
	2,331.218	24.65895	.7541231	.9549886	91.41162
10,000 - 99,999	31,515.87	78.80901	4.099533	.3008095	187.0488
	21,699.84	102.7656	.650461	.3416028	108.1392
≥ 100,000	591,125.1	361.0877	4.105263	.1006772	180.0346
	1,293,243	445.8478	.6690875	.0947447	51.64757
Overall	28,220.69	35.29045	4.12571	1.233574	183.6689
	274,754.5	118.1932	.8312412	2.288428	130.7109

For each category, the first row indicates means and the second row standard deviations to the mean.

Table A.3: Summary Statistics Google Play

Category (users)	# Ratings	Grade	Pr rating (%)	# Characters
100 - 999	22.4	4.54	5.984037	120.4971
	19.85699	.427785	.55343	38.02792
1,000 - 9,999	153.6329	4.172152	4.173531	105.3528
	166.555	.5954823	1.720941	37.49467
10,000 - 99,999	1610.418	4.298316	1.328385	111.2507
	1853.139	.4264449	1.143161	37.54264
$\geq$ 100,000	30981.56	4.223792	.323146	99.59065
	142081	.4090582	.442384	43.99083
Overall	32324.35	4.254464	.92156	105.5373
	121386.5	.4238424	1.410506	41.762

For each category, the first row indicates means and the second row standard deviations to the mean.

A.3 Rating Scale, Rating Cost and Reporting Bias

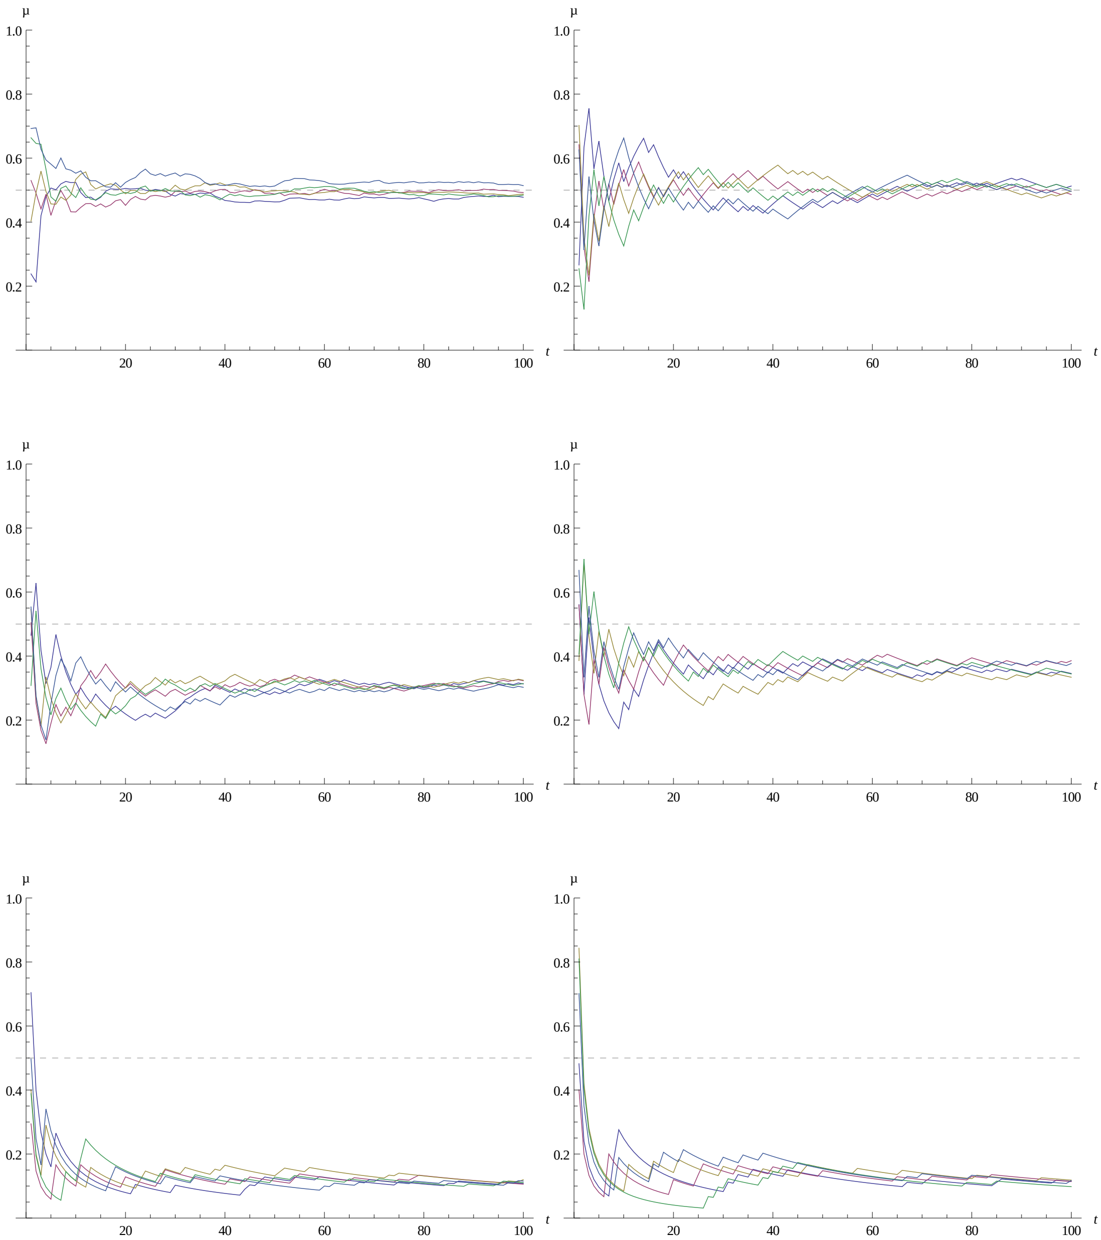

Figure A.2: Evolution of the seller's reputation μ_t over time for the reporting behaviours introduced in section 3.6.2 with $Q \sim \mathcal{N}(\frac{1}{2}, \frac{1}{5})$. Left column depicts an infinitely-grained rating scale with corresponding behaviour $\Phi_\infty(q, \mu)$. Right column depicts a binary rating scale with corresponding behaviour $\Phi_{01}(q, \mu)$. Top row figure depicts a cost of rating $c = 0$, middle row depicts $c = 0.2$ and bottom row $c = 0.6$, $\forall c \in [0, 1]$.

A.4 Mean versus Median

This appendix proposes a characterization of the method used to get the results of section 3.6.3 and Figure 3.3.

A.4.1 Reporting Behaviours

Following Assumption 2, we know that the seller starts his life at $t = 1$ with an empty reputation history ($\mu_0 = \emptyset$), meaning that the first rating only depends on the observed quality of the first transaction, i.e., $r_1 = q_1$. For subsequent periods of time (i.e., $\forall t > 1$), the existence of a rating depends on the reporting function Φ (see Assumption 3). Untruthful reporting behaviours are then divided into two classes: correcting and collective. A correcting reporting behaviour is given by equation (3.9):

$$\Phi(q, \mu) = \begin{cases} \beta q & \text{if } q < \mu \\ (1 - \beta) + \beta q & \text{if } q \geq \mu \end{cases}$$

while a collective reporting behaviour is given by equation (3.10):

$$\Phi(q, \mu) = \beta q + (1 - \beta)\mu.$$

A.4.2 Reputation Measures

The reputation score of the seller after t transactions computed thanks to the *mean* is denoted $\overline{\mu}_t$ and follows equation (3.1), that is

$$\overline{\mu}_t = \frac{1}{t} \sum_{i=1}^{t} r_i.$$

The reputation score of the seller after t transactions computed thanks to the *median* is denoted $\widehat{\mu}_t$, and is given by

$$\mathbb{P}[\{r_1, ...r_t\} \leq \widehat{\mu}_t] \geq \frac{1}{2},$$
$$\mathbb{P}[\{r_1, ...r_t\} \geq \widehat{\mu}_t] \geq \frac{1}{2}.$$

A.4.3 Simulations

Suppose that Q follows a truncated normal distribution, i.e., $Q \sim \mathcal{N}(\frac{4}{5}, \frac{1}{5})$ with $Q \in [0, 1]$ and $\beta = \frac{1}{2}$. We then have $\mathbb{E}[Q] \simeq 0.74251 < \frac{4}{5}$. Table A.4 presents the evolution of the seller's reputation for his first 12 transactions w.r.t. the mean and median, when the economy is inhabited by correcting raters.

t	q_t	$r_t = \Phi(q_t, \overline{\mu}_{t-1})$	$\overline{\mu}_t$	$r_t = \Phi(q_t, \widehat{\mu}_{t-1})$	$\widehat{\mu}_t$
1	0.945346	0.945346	0.945346	0.945346	0.945346
2	0.62016	0.31008	0.627713	0.31008	0.627713
3	0.789941	0.894971	0.716799	0.894971	0.894971
4	0.695434	0.347717	0.624528	0.347717	0.621344
5	0.491163	0.245582	0.548739	0.245582	0.347717
6	0.735211	0.867606	0.601883	0.867606	0.607661
7	0.791722	0.895861	0.64388	0.895861	0.867605
8	0.884134	0.942067	0.681154	0.942067	0.881288
9	0.651529	0.325765	0.641666	0.325765	0.867605
10	0.951604	0.975802	0.675079	0.975802	0.881288
11	0.579524	0.289762	0.640051	0.289762	0.867605
12	0.723109	0.861555	0.658509	0.361555	0.61458
13	0.557238	⋮	⋮	⋮	⋮

Table A.4: Mean versus Median: Evolution of a seller's reputation with correcting raters, where $Q \sim \mathcal{N}(\frac{4}{5}, \frac{1}{5})$ and $\beta = \frac{1}{2}$.

Similarly, Table A.5 presents the evolution of the seller's reputation for his first 12 transactions w.r.t. the mean and median, when the economy is inhabited by collective raters.

t	q_t	$r_t = \Phi(q_t, \overline{\mu}_{t-1})$	$\overline{\mu}_t$	$r_t = \Phi(q_t, \widehat{\mu}_{t-1})$	$\widehat{\mu}_t$
1	0.945346	0.945346	0.945346	0.945346	0.945346
2	0.62016	0.782753	0.86405	0.782753	0.86405
3	0.789941	0.826995	0.851698	0.826995	0.826995
4	0.695434	0.773566	0.832165	0.761215	0.804874
5	0.491163	0.661664	0.798065	0.648019	0.782753
6	0.735211	0.766638	0.792827	0.758982	0.771984
7	0.791722	0.792275	0.792748	0.781853	0.781853
8	0.884134	0.838441	0.79846	0.832993	0.782303
9	0.651529	0.724994	0.790297	0.716916	0.781853
10	0.951604	0.87095	0.798362	0.866728	0.782303
11	0.579524	0.688943	$0.788415x$	0.680913	0.781853
12	0.723109	0.755762	0.785694	0.752481	0.771534
13	0.557238	⋮	⋮	⋮	⋮

Table A.5: Mean versus Median: Evolution of a seller's reputation with collective raters, where $Q \sim \mathcal{N}(\frac{4}{5}, \frac{1}{5})$ and $\beta = \frac{1}{2}$.

Bibliography

AKERLOF, G. A. (1970): "The Market for Lemons," *Quarterly Journal of Economics*, 84, 488–500.

ANDERSON, E. T. (1998): "Customer satisfaction and word of mouth," *Journal of Service Research*, 1:1, 5–17.

ANDERSON, E. T. AND D. I. SIMESTER (2014): "Reviews Without a Purchase: Low Ratings, Loyal Customers, and Deception," *Journal of Marketing Research*, 51:3, 249–269.

BAR-ISAAC, H. (2003): "Reputation and Survival: learning in a dynamic signalling model," *Review of Economic Studies*, 70, 231–251.

BAR-ISAAC, H. AND S. TADELIS (2008): "Seller Reputation," *Foundations and Trends in Microeconomics*, 4, 273–351.

BOLTON, G., B. GREINER, AND A. OCKENFELS (2013): "Engineering Trust: Reciprocity in the Production of Reputation Information," *Management Science*, 59:2, 265–285.

BOLTON, G., E. KATOK, AND A. OCKENFELS (2005): "Cooperation among strangers with limited information about reputation," *Journal of Public Economics*, 89:8, 1457–1468.

BRANDES, L., D. GODES, AND D. MAYZLIN (2013): "Controlling for Self-Selection Bias in Customer Reviews," *Working Paper*.

Brown, G. W. (1951): "Iterative Solutions of Games by Fictitious Play," *Activity Analysis of Production and Allocation. Cowles Commission for Research in Economics, Monograph No. 13, John Wiley and Sons, Inc., New York.*

Cabral, L. (2005): "The Economics of Trust and Reputation: A Primer," *Working Paper. New York University and CEPR.*

Cabral, L. and A. Hortaçsu (2010): "The Dynamics of Seller Reputation: Evidence from eBay," *Journal of Industrial Economics*, 58, 54–78.

Chung, K. L. (1954): "On a Stochastic Approximation Method," *The Annals of Mathematical Statistics*, 25, 463–483.

Cripps, M. W., G. J. Mailath, and L. Samuelson (2004): "Imperfect Monitoring and Impermanent Reputations," *Econometrica*, 72, 407–432.

Dave, C. and K. W. Wolfe (2004): "On confirmation bias and deviations from Bayesian updating," *Working Paper.*

Dellarocas, C. (2005): "Reputation Mechanism Design in Online Trading Environments with Pure Moral Hazard," *Information Systems Research*, 16, 209–230.

Dellarocas, C., M. Fan, and C. A. Wood (2004): "Self-interest, reciprocity, and participation in online reputation systems," *MIT Sloan Working Papers.*

Dellarocas, C. and R. Narayan (2006): "A Statistical Measure of a Populations Propensity to Engage in Post-Purchase Online Word-of-Mouths," *Statistical Science*, 21:2, 277–285.

Dellarocas, C., X. M. Zhang, and N. F. Awad (2007): "Exploring the Value of Online Product Reviews in Forecasting Sales: The Case of Motion Pictures," *Journal of Interactive Marketing*, 21:4, 23–45.

Diekmann, A., B. Jann, W. Przepiorka, and S. Wehrlia (2014): "Reputation Formation and the Evolution of Cooperation in Anonymous Online Markets," *American Sociological Review*, 79, 69–85.

Doob, J. (1949): "Application of the theory of martingales," *Le Calcul des Probabilités et ses Applications, Colloques Internationaux du Centre National de la Recherche Scientifique*, 13:2327.

Ekmekci, M. (2011): "Sustainable reputations with rating systems," *Journal of Economic Theory*, 146, 479–503.

Ellison, G. (1997): "Learning from personal experience: One rational guy and the justification of myopa," *Games and Economic Behavior*, 19:2, 180–210.

Epelman, M., A. Ghate, and R. L. Smith (2011): "Sampled Fictitious Play for Approximate Dynamic Programming," *Comput. Oper. Res.*, 38, 1705–1718.

Faltings, B., F. Garcin, A. Talwar, and R. Jurca (2010): "Reporting incentives and biases in online review forums," *ACM Trans. Web 42, Article 5 (April 2010).*

Fehr, E. and K. M. Schmidt (2006): "The Economics of Fairness, Reciprocity and Altruism: Experimental Evidence and New Theories," *In Handbook on the Economics of Giving, Reciprocity and Altruism*, 1, 615–691.

Friedman, E. J. and P. Resnick (2000): "The Social Cost of Cheap Pseudonyms," *Journal of Economics and Management Strategy*, 10, 173–199.

Fudenberg, D. and D. K. Levine (1989): "Reputation and equilibrium selection in games with a patient player," *Econometrica*, 57, 759–778.

——— (1998): "The Theory of Learning in Games," *The MIT Press*, 29–50.

GARCIN, F., L. XIA, AND F. BOI (2013): "How Aggregators Influence Human Rater Behavior?" *ACM Conference on Electronic Commerce - Workshop on Social Computing and User Generated Content.*

HARMON, A. (2004): "Amazon glitch unmasks war of reviewers," *the New York Times, February 14,2004.*

HODGES, J. L. AND E. L. LEHMANN (1956): "Two Approximations to the Robbins-Monro Process," *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, 95–104.

HOFBAUER, J. AND W. H. SANDHOLM (2002): "On the Global Convergence of Stochastic Fictitious Play," *Econometrica*, 70, 2265–2294.

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

HU, N., J. ZHANG, AND P. A. PAVLOU (2009): "Overcoming the J-shaped distribution of product reviews," *Commun. ACMA - View of Parallel Computing*, 52:10, 144–147.

JURCA, R. (2007): "Truthful Reputation Mechanisms for Online Systems," *Ph.D. thesis, Swiss Federal Institute of Technology, Lausanne.*

KRAMER, M. (2007): "Self-Selection Bias in Reputation Systems," *Trust Management*, 238, 255–268.

KREPS, D. M. AND R. WILSON (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253–279.

LAFKY, J. (2014): "Why Do People Rate? Theory and Evidence on Online Ratings." *Games and Economic Behavior*, 87, 554–570.

Levin, J. (2006): "Learning in Games," *Game Theory and Economic Applications, Stanford University.*

Liu, Q. (2011): "Information Acquisition and Reputation Dynamics," *Review of Economic Studies*, 78, 1400–1425.

Liu, Q. and A. Skrzypacz (2014): "Limited Records and Reputation Bubbles," *Journal of Economic Theory*, 151, 2–29.

Luca, M. (2011): "Reviews, Reputation, and Revenue: The Case of Yelp.com," *Harvard Business School Working Paper, No. 12-016.*

Luca, M. and G. Zervas (2013): "Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud," *Working Paper.*

Mailath, G. and L. Samuelson (2001): "Who wants a good reputation?" *Review of Economic Studies*, 68:2, 415–441.

McGlohon, M., R. Zach, and N. Glance (2010): "Star Quality: Aggregating Reviews to Rank Products and Merchants," *International Conference on Weblogs and Social Media.*

Mell, A. (2011): "Re-Thinking Reputation," *Working Paper*, Nuffield College, University of Oxford.

Melnik, M. and J. Alm (2002): "Does a seller's reputation matter? evidence from ebay auctions," *Journal of Industrial Economics*, 50:3, 337–350.

Milgrom, P. and J. Roberts (1982): "Predation, Reputation, and Entry Deterrence," *Journal of Economic Theory*, 27, 280–312.

Miller, N., P. Resnick, and R. Zeckhauser (2005): "Eliciting Informative Feedback: The Peer-Prediction Method," *Management Science*, 51:9, 1359–1373.

Mosteller, F. and J. Tukey (1977): "Reading, Mass, Addison-Wesley Pub Co," *The Annals of Mathematical Statistics*.

Muchnik, L., S. Aral, and S. J. Taylor (2013): "Social Influence Bias: A Randomized Experiment," *Journal of Industrial Economics*, 341:6146, 647–651.

Mui, L. (2003): "Computational Models of Trust and Reputation: Agents, Evolutionary Games, and Social Networks," *Ph.D. thesis, Massachusets Institute of Technology*.

Mukherjee, A., V. Venkataraman, L. Bing, and N. S. Glance (2013): "What Yelp Fake Review Filter Might Be Doing?" *ICWSM,, The AAAI Press*.

Nagle, F. and C. Riedl (2013): "Online Word of Mouth and Product Quality Disagreement," *Harvard Business School Working Paper, No. 13-091*.

Rabin, M. and J. L. Schrag (1999): "First impressions matter: A model of confirmatory bias," *Quaterly Journal of Economics*, 114(1), 37–82.

Resnick, P. and R. Zeckhauser (2002): "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System," *The Economics of the Internet and E-Commerce*, 11.

Rezek, I., D. S. Leslie, S. Reece, S. J. Roberts, A. Rogers, R. K. Dash, and N. R. Jennings (2008): "On Similarities between Inference in Game Theory and Machine Learning." *J. Artif. Intell. Res. (JAIR)*, 33, 259–283.

Robbins, H. and S. Monro (1951): "A Stochastic Approximation Method," *The Annals of Mathematical Statistics*, 22:3, 400–407.

Robinson, J. (1951): "An iterative method of solving a game," *Annals of Mathematics*, 54, 296–301.

Saeedi, M., Z. Shen, and N. Sundaresan (2013): "The Value of Feedback: An Analysis of Reputation System," *Working Paper.*

Sun, M. (2012): "How Does the Variance of Product Ratings Matter?" *Management Science*, 58:4, 696–707.

Talwar, A., R. Jurca, and B. Faltings (2007): "Understanding User Behavior in Online Feedback Reporting," *Proceedings of the ACM Conference on E-Commerce (EC'07).*

Tirole, J. (1988): "The Theory of Industrial Organization," *MIT Press.*

Vickrey, W. (1961): "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8–37.

Weber, S. (2004): "The Success of Open Source," *Harvard University Press*, ISBN 978-0-674-01292-9.

Yoon, Yeosunand Sarial-Abi, G. and Z. Gurhan-Canli (2012): "Effect of Regulatory Focus on Selective Information Processing," *Journal of Consumer Research*, 39:1, 93–110.

Young, H. P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57–84.

Zhang, X. M. and F. Zhu (2011): "Group Size and Incentives to Contribute: A Natural Experiment at Chinese Wikipedia," *American Economic Review*, 101:4, 1601–15.

Zimper, A. and A. Ludwig (2009): "On attitude polarization under Bayesian learning with non-additive beliefs," *Journal of risk and uncertainty*, 39:2, 181–212.