Essays in Economic Theory

Inaugural-Dissertation

zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften

durch die

Rechts- und Staatswissenschaftliche Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Tim Julius Frommeyer

aus Bremen

Bonn 2017

Dekan:	Prof. Dr. Daniel Zimmer
Erstreferent:	Prof. Dr. Dezsö Szalay
Zweitreferent:	JProf. Francesc Dilmé, Ph.D.
Tag der mündlichen Prüfung:	21. Dezember 2016

Acknowledgements

This thesis is the result of my time as a Ph.D. student at the Bonn Graduate School of Economics and would not have been possible without the help and the support of many people. I am thankful to each and every one of them.

Foremost, I want to express deep gratitude to my first supervisor Deszö Szalay for his invaluable support, and for giving me the freedom to explore the various research questions from different areas I was interested in. I am equally thankful to my second supervisor Francesc Dilmé for his support, encouragement, and continuous feedback. From countless discussions with my supervisors I learned a lot about economics and economics research, and benefited greatly from their guidance, feedback, and suggestions throughout the years.

I want to thank the faculty members for the friendly atmosphere in Bonn, and for encouraging lively, open discussions of new ideas and projects. In particular, I am grateful to Marc Le Quement for his constant feedback and support in the development of the first chapter of this thesis, as well as to Sven Rady, Dominik Grafenhofer, and Olga Gorelkina for their encouragement, interest, and feedback.

Not only as my coauthor but also as my long-time office neighbor and good friend, I am indebted to Benjamin Schickner for his feedback, his support, and the many joyful hours we spent together.

My fellow Ph.D. students at the Bonn Graduate School of Economics deserve special mention, particularly Nazarii and Mirjam Salish, Ariel Mecikovsky, and Matthias Meier. I was very fortunate to get to know and study with so many supportive and lovable colleagues who became good friends.

Moreover, I am very grateful to Marco Köster, Martin Siering, and Peer Jansli for their continuous interest and their overwhelming support over the time of my studies.

Most importantly, I would not have been able to start or finish my Ph.D. studies without the help of my family, especially my parents. I am more than thankful for their unconditional faith and trust in me, their constant support, and for the freedom I enjoyed growing up.

Finally, I want to thank the Bonn Graduate School of Economics and the University of Bonn for financial support, and for providing an excellent research environment for Ph.D. students.

Contents

Introduction

When individuals have to make choices, the quality of their decision depends on the available information about the alternatives. For consumers, for example, who debate whether to purchase a good from one or another seller, or whether to buy a good at all, information about the sellers or the quality of the good are crucial for their decision.

Likewise, when groups have to find an agreement, aggregating information about the subject for debate is a key determinant of a well-founded decision. In jury trials, for example, where jury members have to decide on guilt or innocence of a defendant, the accuracy of their common decision depends on the information each jury member has as well as on their ability to communicate their insights to each other.

This thesis studies three theoretical models in order to analyze different microeconomic questions in which agents make strategic decisions based on incomplete information, or companies strategically influence the information of their consumers to maximize their profits.

The first chapter focuses on groups who have to make a common decision via voting, and analyzes the strategic incentives of group members to share their privately held information about the case. Specifically, Chapter 1 studies how the possibility of a binding consequence can improve the informativeness of a straw poll held prior to a decisive vote in a jury trial about a defendant's guilt or innocence when jury members hold private, incomplete information about the truth.

A common feature of the second and the third chapter is the perspective of a seller who strategically manages her reputation. The second chapter analyzes situations in which consumers are identified with the goods they consume, and the style or type a good represents impacts their buying decisions. More specifically, a monopolist repeatedly interacts with consumers of different types and her good's reputation is constituted by the types of consumers who possess the good. The analysis provides insights in how the monopolist strategically manages the composition of her clientele. In contrast to this chapter in which the seller's reputation is composed of publicly observable attributes of her customers, in Chapter 3, reputation reflects the consumer's information about the seller's ability which is her private information. An online seller of unknown logistical capabilities competes against a local retail store for buyers. The chapter analyzes how changes in information about her deliveries affect the online seller's incentives to improve the speed of her deliveries through better shipping services, and

evaluates which pieces of information benefit or harm the welfare of consumers.
The remainder of this introduction describes the theoretical models and the results of each chapter in more detail.

In the first chapter, I consider a committee of agents with heterogeneous preferences who posses private information about the unknown state of the world, and have to make a binary decision in two rounds of voting. The utility of each agent from either decision depends on the state of the world as well as on their preferences. The leading example is a jury trial in which jurors are uncertain about the guilt or innocence of a defendant. Over the course of the trial, jurors gather information and have to agree on a verdict in two voting rounds. Their preferences represent their individual level of reasonable doubt.
I investigate incentives for truthful revelation of the jurors' information in the first voting period. Coughlan (2000) shows that jury members reveal their information in a non-binding straw poll prior to a decisive vote only if their preferences are in fact homogeneous. By taking costs of time into account, I demonstrate that jury members have strictly higher incentives to reveal their private information truthfully if a decision with high levels of consensus can already be made in the straw poll. Jurors condition their decision in the decisive vote on those scenarios in which their vote is pivotal, and anticipate the influence of their information on the other jurors in the decisive vote. If straw poll is non-binding, providing false information to others hurts a juror only if all jurors evaluate information in the same way. If a straw poll can have consequences, however, providing false information can implement a less preferred verdict already before the second vote, which strengthens incentives to report information truthfully. I use these insights to show furthermore that with a potentially consequential first vote instead of a non-binding straw poll, members of all homogeneous and some heterogeneous juries are strictly better off when the requirement for early decisions is chosen carefully.

In the second chapter, which is based on joint work with Benjamin Schickner, buyers care about who else consumes the good of a monopolistic seller. In a dynamic reputation model, the long-lived seller repeatedly sells her good to short-lived buyers who differ with respect to their type. The seller's reputation depends on the types of her clientele. If mostly high types buy the good, reputation increases, and if more low types buy the good, reputation decreases. The utility a buyer receives from purchasing the good depends on the seller's reputation and on his type. The seller faces a trade-off between increasing her reputation by restricting demand to a small, but exclusive clientele and increasing her period profits by selling to a large, but broad clientele.
In a general model, we first show existence of an equilibrium and characterize the seller's value in each period. Second, we specifiy the model in a linear-uniform setup in order to analyze the dynamics of reputation and prices over time. We find that in the long run, the seller's reputation converges to a stable level. In the short run, however, we show that the durability

of the good has a significant impact on the seller's dynamic reputation management. For goods of relatively low durability, such as fashionable apparel, prices and reputation fluctuate around the long-run reputation, whereas the reputation and prices of comparatively durable goods, such as luxury watches, increase or decrease monotonically to the long-run reputation. Observations of price and reputation dynamics from fashion labels and watchmakers seem to support our theoretical results.

The third chapter examines a dynamic reputation model with two periods in which a long-lived online seller of unknown logistical ability competes against an offline retailer. In order to deliver goods to her short-lived buyers, the online seller has to choose one of two shippers which differ in their expected delivery time as well as in the fees for the seller. The overall waiting time for buyers depends on the seller's ability as well as on the chosen shipper's quality. Buyers suffer costs from traveling to the offline retailer, and purchase online if their expected overall waiting time is sufficiently short. Reputation represents the buyers' belief about the seller's ability which they update based on their past experienced waiting time for the good. In this model, the online seller's incentives to assign a fast but expensive shipper depend significantly on the information about the delivery process that buyers can observe. We compare the equilibrium outcomes of four specifications where buyers can or cannot observe the shipper's quality upon delivery of the good, and can or cannot *track and trace* the delivery process in detail. Intuitively, the ability to track the delivery allows buyers to disentangle the contributions of the seller and the shipper to their overall waiting time, and enables a good seller to signal her high ability without hiring a fast shipper. Consequently, tracking proves harmful to the buyers' welfare in most cases. In contrast, observing the chosen shipper quality upon delivery harms buyers' welfare only if the seller's initial reputation is comparatively low but is beneficial to buyers otherwise.

1

On Two-Period Committee Voting: Why Straw Polls Should Have Consequences.

1.1 Introduction

Information is particularly valuable when important and difficult decisions are pending. In many situations those decisions are not made by a single person but rather in groups, so called *committees*, in a voting procedure. University faculties typically delegate their decision on an applicant's employment to a committee and many companies, such as *Google*, proceed similarly[1]. In politics, party factions and parliaments form committees for investigations or in order to work out recommendations, and in jury trials a committee composed of representatives of the society has to decide whether a defendant is guilty or innocent of the accused crime. In any of these cases, each committee member has a distinct evaluation of the subject for debate. Therefore, a committee has a richer pool of information at its disposal than an individual decision maker and could potentially make more accurate decisions. This information, however, is dispersed. A widely used approach for coordination are *straw polls* as non-binding communication devices before the decisive vote[2].

When the United Nations Security Council discussed the candidates for the position of Secretary-General of the United Nations in 2006, it conducted a series of straw polls to determine the members opinions on the candidates. Although none of these votes had decisional power, all candidates withdrew their candidacy afterwards except for Ban Ki-moon who received the most votes in each straw poll[3]. The famous movie *12 Angry Men* from 1957

[1] "An independent committee of Googlers review feedback from all of the interviewers. This committee is responsible for ensuring our hiring process is fair and that we're holding true to our 'good for Google' standards as we grow.", Google, How We Hire.

[2] Especially in large committees straw polls are a simple and swiftly conducted communication tool when verbal deliberation is tedious and confusing. Moreover, an anonymous poll can circumvent privacy issues and still provide some communication among committee members.

[3] More importantly, Ban was the only candidate who received votes from all permanent members of the Security Council which have veto power in the decisive election process. Ban then was elected by the general Assembly on October 13, 2006.

covers the deliberation process of a jury. At the beginning of their deliberation, the jurors conduct a straw poll in order to collect their initial attitudes towards the underlying case and use the result as a starting point for the subsequent debate.
The motivation for conducting straw polls stems from the intuition that voters can harmlessly reveal their information, not being at risk of unintentionally causing a decision already. Additionally, from a series of experiments on decision making in groups, Goeree and Yariv (2011) conclude that members strongly appreciate information revealed by others. Although performed frequently and in many situations, the benefit of straw polls is disputed. According to Robert III. et al. (2000, p. 415)[4], conducting "an informal straw poll to 'test the water' (...) neither adopts nor rejects a measure and hence is meaningless and dilatory."

Condorcet (1785) was the first to formally argue in his *Jury Theorem* that voting decisions made in groups outperform those of individuals. In his setup, a jury votes once on a binary decision, each of them being objectively best in one of two possible states of the world, and the alternative with more votes is implemented. Jury members condition their voting decision only on their private information about the state, which is correct with probability higher than 0.5. Moreover, the jury members' preferences are aligned, in the sense that they would undoubtedly prefer the same (objectively best) decision if the state of the world was commonly known, e.g. to convict a guilty defendant and acquit the innocent. As the vote aggregates the individuals' information, the implemented decision is more likely to be correct than the one from an imperfectly informed individual.[5]. Condorcet assumes that agents only consider their private information when voting in a group, just as they would when deciding all alone. In strategic games however, this assumption is not necessarily satisfied. Austen-Smith and Banks (1996) as well as Feddersen and Pesendorfer (1998) formally set up a standard model and consider a homogeneous committee with commonly known and aligned preferences. Preferences can be interpreted as levels of reasonable doubt which are identical for homogeneous committee members. In other words, agents would unanimously prefer the same decision not only if the state was known but also if they faced the same information. They demonstrate that the voting behavior assumed by Condorcet constitutes an equilibrium only if the voting rule is adjusted to the agents' preferences appropriately.
Coughlan (2000) considers an extension of the standard model with committees whose members are heterogeneous with respect to their levels of reasonable doubt. Preferences are aligned but agents assess the same information differently and might disagree on the preferred decision. He studies the role of a preliminary non-binding straw poll and demonstrates that information is aggregated in equilibrium only if the agents' preferences are

[4] *Robert's Rules of Order Newly Revised* is widely used as parliamentary authority, e.g., by the US Congress, and guide for meetings and assemblies.
[5] In addition, by a law of large numbers the probability of a correct decision approaches 1 as the jury size grows large. See, for example, Piketty (1999) for an overview on the Jury Theorem. For a discussion and extensions see Ladha (1992), Miller (1986) and Young (1988).

in fact homogeneous. In the following, we refer to this insight as Coughlan's impossibility result. If an agent considers revealing his information in a straw poll, he conditions on the case where the information he discloses tips the balance in the subsequent decisive vote, that is when his vote is *pivotal*. If the committee is heterogeneous and other agents draw opposite conclusions from revealed information, disclosing information truthfully can lead to a unfavorable voting outcome for some agents. In this case, these agents can be better off with providing misinformation in the straw poll in order to manipulate the decisive vote in their favor.

In this paper, we study a modified version of Coughlan (2000)'s heterogeneous committee model in which agents additionally have preferences over the length of deliberation. More specifically, committee members incur costs of time from each round of voting and, hence, prefer to make decisions earlier. We hereby account for the deferring feature of straw polls pointed out by Robert III. et al. (2000). When committee members engage in a non-binding straw poll they always vote twice. If agents have the opportunity to circumvent a second poll whenever there is broad agreement in the first vote already, new strategic effects arise. Unlike in a non-binding straw poll, agents are then not only pivotal when their disclosed information tips the balance in the subsequent vote but also when there is a high level of consensus for one of the two alternatives within the other agents' information. In these cases an agent can maximize the probability to prevent a redundant second vote by revealing his information truthfully. As agents form beliefs about the probability of each pivotal case from their private information, consensus among the other agents for the agent's initially preferred decision is more likely than for the opposite. Hence, committee members have better incentives to reveal their information compared to a straw poll that has no consequences. In particular, even committees that are in fact heterogeneous are able to aggregate information perfectly if the straw poll is modified to implement a decision for high levels of consensus straightaway without a second vote.
In addition, we consider the committee's welfare under both a non-binding straw poll and a potentially consequential first vote. When we already allow for a decision in the first vote, we can always identify a requirement on the implementation of this vote that strictly improves the welfare of each juror of any homogeneous committee. If the requirement is strict enough, the jury will make the same decision in both setups but the ability to save costs causes a strict welfare improvement. Naturally, this positive welfare effect from saving costs is also present for heterogeneous juries. The made decisions, however, are always suboptimal for some members of heterogeneous committees. This provides a negative welfare effect for those agents. The overall welfare effect is positive for every agent as long as heterogeneity is manageable. Hence, we can show that there is always a decision threshold for the first vote for which every member of committees with bounded heterogeneity is strictly better off than with a straw poll in the first period.

Finally, we are concerned with a designer who wants to optimally set a threshold for agreement in the first period. We distinguish the case where the designer observes the committee's preferences to that where she must commit to a threshold before getting to know the agents. Facilitating early decisions saves costs of time but potentially impacts the optimality of the jury decision. The designer must solve this trade-off with her choice of the decision rule. We provide conditions for her choice to be optimal from all candidate decision rules that pareto-dominate any equilibrium outcome of setups with straw polls.

This paper belongs to the literature on committee voting with commonly known preferences. As is standard, we use the terminology of jury trials in the following. A committee is referred to as *jury* and a committee members is called a *juror*. The jury has to decide between *conviction* or *acquittal* of a defendant who is either *guilty* or *innocent*.

Deimen et al. (2015) show that Coughlan's impossibility result can be mitigated when the information structure is enhanced and allows jurors to examine consistency of information. Hummel (2012) as well as Thordal-Le Quement and Yokeeswaran (2015) demonstrate that information aggregation can be achieved if heterogeneous committees deliberate in homogeneous subgroups first.

For the case of privately known preferences, information is aggregated in heterogeneous juries under certain conditions on the voting rule (Austen-Smith and Feddersen, 2006), the jurors' preferences and the jury size (Meirowitz, 2007; Thordal-Le Quement, 2013).

Our work is also related to information aggregation in elections. Morgan and Stocken (2008) study the informative substance of straw polls that are held prior to elections. In their setup the result of the poll influences the subsequent policy choice which constituents take into consideration strategically. They find that information can be aggregated in small polls if the electorate is sufficiently homogeneous. Piketty (2000) investigates the effect on information aggregation if jurors use elections to communicate their preference in order to influence future decisions.

Furthermore, our work is partly related to the literature on debates. Austen-Smith (1990) analyzes if debate can influence a later decision as well as its informational role. Bognar et al. (2013) study information aggregation over the course of repeated and costly negotiations whereas Damiano et al. (2010, 2012b) investigate the role of costly delay in repeated negotiations with asymmetric information. Damiano et al. (2012a) derive an optimal deadline on repeated negotiations.

The following chapter introduces the standard model of committee voting extended by costs of time from the voting process. In Chapter 3 we compare the results of Coughlan (2000) on non-binding straw polls followed by a decisive vote to a two-period setup with the opportunity for a decision in the first vote already should broad agreement occur. We show that the conditions on the jury's heterogeneity for equilibria with information aggregation in the first

period are weaker for the latter setup. In Chapter 4 we argue that a heterogeneous jury can be strictly better off with the potential for early agreement compared to a pure straw poll and we provide conditions for a designer to set the decision rule for the first period optimally. Finally, Chapter 5 concludes.
All proofs are relegated to the Appendix.

1.2 The Two-period Committee Voting Model

1.2.1 Agents/Jurors

N denotes the number and, with slight abuse of notation, the set of agents who make a binary decision. We will interpret the problem as the *conviction* (C) or *acquittal* (A) of a defendant in a jury trial. Therefore, from now on we call an agent a *juror*, the set N is a *jury* consisting of N jurors. Jurors maximize their expected utility and we assume common knowledge of rationality.

1.2.2 States and Preferences

There are two states of the world, $\omega \in \{G, I\}$. The defendant is either *guilty* $(\omega = G)$ or *innocent* $(\omega = I)$. For simplicity, we assume that both state are equally likely, that is $\Pr[\omega = G] = \Pr[\omega = I] = 1/2$. Juror j's preferences are state dependent and normalized to

$$
\begin{aligned}
U_j(C \mid \omega = G) &= U_j(A \mid \omega = I) = 0, \\
U_j(C \mid \omega = I) &= -q_j, \\
U_j(A \mid \omega = G) &= -(1 - q_j),
\end{aligned}
$$

where $q_j \in (0, 1)$. This normalization allows us to interpret q_j as a threshold probability of guilt of juror j, who prefers decision C over A if and only if his perceived probability of state G is larger or equal to q_j. Jurors with lower thresholds q_j are comparatively more biased towards C whereas jurors with higher thresholds q_j are more biased towards A. In the context of a jury trial one can also think of q_j as j's level of *reasonable doubt*. We assume that all preferences q_j are commonly known.
Without loss of generality, we sort jurors by their preferences $q_1 \leq \ldots \leq q_N$. Ex ante jurors are solely distinguished by their preferences and we will say *juror j* for a juror with preferences $q_j \in \left[q_{j-1}, q_{j+1}\right]$. Additionally, we impose that jurors incur additive costs of $c \geq 0$ on their utility from each round of voting, which represent costs of time or opportunity costs. Alternatively, one could think of impatient jurors who prefer earlier decisions to later ones.

1.2.3 Information

Prior to the decision making process each juror $j \in N$ receives an informative signal $s_j \in \{i, g\}$ about the state of the world, where

$$\Pr[s_j = i|\omega = I] = \Pr[s_j = i|\omega = G] = p \in \left(\frac{1}{2}, 1\right)$$

We refer to p as the signal's *precision*. Signals are independently drawn, privately observed and not verifiable. The signal's precision is identical and known to every juror.

1.2.4 Voting

We consider a two-period voting game of the following form. In both periods jurors vote for either A or C. Denoting by x_t the number of C-votes in period t, the decision in period $t \in \{1, 2\}$, d_t, is determined by a decision rule represented by threshold k_t. In $t = 1$, the defendant can be convicted or acquitted only if more than or equal to $N - k_1$ jurors vote for the corresponding alternative[6]. If such a majority turns out, the game ends and the jury's respective decision is implemented. Otherwise, the decision is delayed ($d_1 = D$). The number of votes for both alternatives is revealed and the jury votes again in period 2. Formally the decision rule in period 1 is given by,

$$d_1 = \begin{cases} A & \text{if } x_1 < k_1 \\ C & \text{if } N - k_1 < x_1 \\ D & \text{else.} \end{cases} \tag{1.1}$$

See Figure 1.1 for a graphical representation with the number of C-votes in $t = 1$ on the axis. If the game continues in second period after $d_1 = D$, x_1 is disclosed the jury votes again. The defendant is convicted if at least k_2 jurors vote for C and acquitted otherwise. Formally the decision rule in $t = 2$ is given by,

$$d_2 = \begin{cases} A & \text{if } x_2 < k_2 \\ C & \text{if } k_2 \leq x_2. \end{cases} \tag{1.2}$$

Figure 1.2 shows a graphical representation. In the following, we refer to a decision rule in period t by the corresponding threshold k_t, which are exogenously given to the jurors.
If k_1 is small, a decision can only be made if there is broad agreement in the first period, whereas a decision can also be made by a smaller majority if k_1 is higher. The case of $k_1 = 0$

6 We implicitly assume symmetry in the voting rule for period 1. This simplifies the analysis but does not impact the qualitative results. For the impact which costs of time have on incentives for informative voting in the first period, it will only be necessary that some super majority is needed for conviction or acquittal in the first period, but it need not be the same for both decisions.

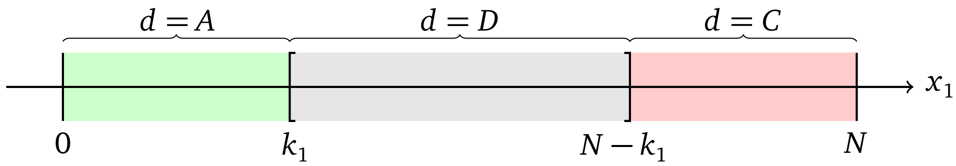

Figure 1.1: Decision rule in $t = 1$.

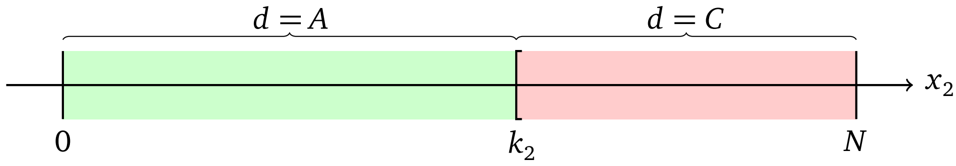

Figure 1.2: Decision rule in $t = 2$.

represents a non-binding *straw poll* in the first period which can be interpreted as a preliminary round of communication in which no decision can be made yet. In this communication, however, jurors cannot remember any information revealed but the number of jurors who prefer each alternative[7]. Straw polls are investigated in detail in Coughlan (2000). His setup is a special case of ours with $k_1 = 0$ and $c = 0$. We will use his findings, which we summarize in the following chapter, as a benchmark.

1.2.5 Notation

The posterior probability that the state is G if x of N signals indicate g is denoted by $\beta(x, N)$ and is computed according to Bayes' Rule. That is,

$$\beta(x,N) \;=\; \frac{(1-p)^{N-x}p^x}{(1-p)^{N-x}p^x + (1-p)^x p^{N-x}}.$$

When juror j enters period 1 or 2 his information is given by $\mathscr{I}_j^1 = (s_j)$, and $\mathscr{I}_j^2 = (s_j, x_1)$, respectively. We denote juror j's strategy by $\sigma_j = \left(\sigma_j^1, \sigma_j^2\right)$, where σ_j^t denotes the probability that j votes C in period t given his information $\mathscr{I}_j^t$. If j votes according to his signal, that is voting for conviction after receiving a g-signal and for acquittal otherwise, we say that j votes *informatively*. Formally, j's strategy prescribes informative voting in period t if

$$\sigma_j^t\left(\mathscr{I}_j^t\right) \;=\; \begin{cases} 1 & \text{if } s_j \;=\; g, \\ 0 & \text{if } s_j \;=\; i. \end{cases}$$

[7] Even if there is no non-binding straw poll in the first stage but no decision was made, jurors can only observe how many jurors voted for each alternative but not each jurors vote individually. This assumption does not impact our results, because each jurors signal has the same precision. If this was not the case, jurors could learn about other signal's precision, such as in Bognar et al. (2013).

When juror j votes for C in period t if and only if his perceived probability of guilt, given his information, exceeds his level of reasonable doubt q_j, we say that j votes *sincerely* in period t. Formally, j's strategy prescribes sincere voting in period t if

$$\sigma_j^t\left(\mathscr{I}_j^t\right) = \begin{cases} 1 & \text{if } \Pr_j\left[G|\mathscr{I}_j^t\right] \geq q_j, \\ 0 & \text{if } \Pr_j\left[G|\mathscr{I}_j^t\right] < q_j. \end{cases}$$

Besides their own signals, jurors access information from inferring the other jurors' signals from their strategic behavior and observed aggregated voting outcomes. As all jurors receive signals of the same precision, inferring more g-signals leads to higher posterior probability of guilt. Analogously to every juror's probability threshold q_j, we can determine thresholds on the number of g-signals that a juror needs to observe in order to prefer conviction of the defendant. In this context, λ_j^N represents the *conviction threshold* of juror $j \in N$. He prefers C over A if and only if he observes more than or equal to λ_j^N of N signals indicating g.

Definition 1.1. *Juror $j \in N$ has conviction threshold λ_j^N if*

$$\beta(\lambda_j^N - 1, N) \leq q_j \leq \beta(\lambda_j^N, N).$$

Analogously to Austen-Smith and Feddersen (2006) and Thordal-Le Quement (2013), we define a jury's *minimal diversity* as follows:

Definition 1.2. *Jury N has a minimal diversity of m if*

$$\max_{(i,j)\in N\times N} \lambda_i^N - \lambda_j^N = m.$$

A jury's minimal diversity measures its heterogeneity in terms of its jurors' conviction thresholds. If a jury N has a minimal diversity of 0, its jurors have the same conviction threshold and would agree on a decision if all private information was disclosed. This is not the case for juries with a minimal diversity larger than 0. In the following, we will call a jury *homogeneous* if $m = 0$, and *heterogeneous* otherwise.

1.2.6 Strategies and Beliefs

We are interested in conditions under which information can be aggregated and jurors in fact use revealed information in their voting strategy. Therefore, we restrict attention to the following profile of strategies. Jurors reveal their private signal by voting informatively in the first vote. If information is not congruent enough to already make a decision, jurors observe the outcome of the first vote x_1 and update their beliefs about the state of the world accordingly to $\beta(x_1, N)$. Then, every juror votes sincerely in the second period. More formally, we will derive conditions on the jurors' preferences for the following profile of strategies and beliefs to constitute a Perfect Bayesian Equilibrium. For all $j \in N$,

$$\sigma_j^1\left(s_j\right) = \begin{cases} 1 & \text{if } s_j = g, \\ 0 & \text{if } s_j = i, \end{cases}$$

and

$$\sigma_j^2\left(s_j\right) = \begin{cases} 1 & \text{if } \beta(x_1, N) \geq q_j, \\ 0 & \text{if } \beta(x_1, N) < q_j, \end{cases}$$

where jurors consistently (correctly) believe that the others vote informatively in the first stage and update their belief about the state of the world accordingly. We denote this profile of strategies and beliefs by $(\sigma, \mu) = \left(\sigma_j, \mu_j\right)_{j=1}^N$, for which we make the following observations.

- If juror j (hypothetically) observes x votes for conviction from the other jurors before he votes in the first stage, his belief about the defendant being guilty is $\beta(x, N)$ if $s_j = g$ and $\beta(x-1, N)$ otherwise.
- If no decision is made in the first stage, all jurors observe x_1 and update their belief about the defendant being guilty to $\beta(x_1, N)$ via Bayes' rule. Therefore, all jurors have the same posterior belief in the second stage.

Equilibria with these strategies and beliefs have the property that information is perfectly aggregated in the first stage. If the information is not congruent in the sense that informative voting does not yet lead to a decision in the first stage, the jurors make their decision in the second stage conditional on all available information.
This profile of strategies and beliefs features strategic behavior in the second period which is sensitive to available information. As we are interested in conditions for informative voting, strategies should take revealed information into account. Alternative strategies for the second period subgame would require jurors to vote predominately for one alternative independently of available information. Providing incentives for information revelation is both difficult and needless if information is not appreciated in the decision process. This is most dominant if the decision is always made in the second period, that is when the first period vote is a straw poll.

Before we proceed to the analysis, we make an assumption on jurors preferences.

Assumption 1.1. *For any $j \in N$, $\beta(0, N) < q_j < \beta(N, N)$ or, equivalently, $\lambda_j^N \in \{1, \ldots, N\}$.*

This assumption excludes jurors with extreme preferences who still prefer an alternative even if all signals were known and indicated the opposite state of the world. In other words, we exclude prejudiced jurors who vote for one alternative regardless of available information. In particular, this assumption ensures that the juror who is pivotal in the second vote, if all jurors vote sincerely, takes the available information into account.

1.3 Equilibrium Analysis

In this chapter, we analyze the two-period model in the following order. We start by discussing the impossibility result of Coughlan (2000) for a non-binding straw poll in the first period, that is for $k_1 = 0$. Then, we show how costs of time influence the jurors incentives to reveal information in their first period vote if $k_1 > 0$.

1.3.1 $k_1 = 0$: Non-binding First Period straw poll

Consider first a situation with a non-binding straw poll in the first period, i.e. $k_1 = 0$, as presented by Coughlan (2000). He shows that there is no equilibrium in which a jury votes informatively in the straw poll and sincerely in the decisive vote, unless all its jurors have the same conviction thresholds. In other words, the profile of strategies and beliefs (σ, μ) can only be an equilibrium for juries with a minimal diversity of 0[8].

In the decisive voting period $t = 2$, jurors vote sincerely by taking revealed information into account. Each juror $j \in N$ conditions his vote on the situation in which he is pivotal, that is when his vote actually decides upon the defendant's acquittal or conviction. In any other case his vote does not affect the decision and thus, his expected payoff from voting A or C is equal. As the jurors maximize their expect utility with their votes, they condition on the unique situation that affects their expected utility. Recall that jurors learn the outcome of the (informative) straw poll before casting their vote in $t = 2$. Fully informed, every $j \in N$ updates his posterior probability of guilt to $\beta(x_1, N)$ and votes for conviction if and only if

$$\beta(x_1, N) \geq q_j \quad \Leftrightarrow \quad x_1 \geq \lambda_j^N,$$

which coincides with sincere voting. If jurors vote informatively in the straw poll, sincere voting is straightforward part of any equilibrium strategy which is sensitive to revealed information.

Having established the equilibrium strategies in $t = 2$, we can now consider incentives for informative voting in the straw poll. The juror who is actually pivotal in the decisive vote is juror k_2. When k_2 votes for C sincerely, every $j < k_2$ will do so as well. When k_2 prefers A given the revealed information, every $j > k_2$ has the same preference. Therefore, whichever decision k_2 prefers in $t = 2$ will be implemented. Although no decision can be made in the straw poll, the

[8] Cf. Proposition 5 in Coughlan (2000) where he differentiates three cases. Apart from the one mentioned, the other cases are ruled out by Assumption 1.1. Information does not influence the jury's final decision in those cases and information revelation is trivially an equilibrium behavior.

vote has an impact on the information influencing the decisive vote later. Given the strategy profile, jurors anticipate that k_2 with conviction threshold $\lambda_{k_2}^N$ is pivotal in the second period. As a consequence, juror j is pivotal in the straw poll, if his vote in $t = 1$ influences k_2'th information to swing the pivotal vote in $t = 2$ to either C or A. This is the case if $\lambda_{k_2}^N - 1$ of the other $N - 1$ jurors informatively vote C in the straw poll. For j to vote informatively in $t = 1$ as well, he has to prefer C after he receives a g-signal and A after an i-signal. In the event that j is actually pivotal in the first period in the above sense, he faces $\lambda_{k_2}^N - 1$ g-signals from the other jurors, resulting in $\lambda_{k_2}^N$ g-signals in total if $s_j = g$, and $\lambda_{k_2}^N - 1$ g-signals in total if $s_j = i$. Therefore, j votes informatively if and only if

$$\beta\left(\lambda_{k_2}^N - 1, N\right) \leq q_j \leq \beta\left(\lambda_{k_2}^N, N\right) \quad \Leftrightarrow \quad \lambda_j^N = \lambda_{k_2}^N, \tag{1.3}$$

that is if he has the same conviction threshold as juror k_2.

Since this argument is the same for any juror of the jury, all jurors must have the same conviction threshold as k_2 to sustain an equilibrium with informative voting in the straw poll. In other words, informative voting in a straw poll cannot be part of an equilibrium strategy for any heterogeneous jury. In contrast, suppose a juror j's conviction threshold conflicts with the one of k_2, i.e., $\lambda_j^N \neq \lambda_{k_2}^N$. Given that all other jurors vote informatively, providing information truthfully in the straw poll implements an undesirable decision for j. By misinforming, however, he could improve the jury's decision from his point of view which represents a profitable deviation.

Note that costs from voting play no role for the jurors' incentives to vote informatively in a straw poll. They always vote twice and cannot avoid costs with their behavior in the first period. Therefore, this impossibility result for non-binding first period straw polls is independent of assuming costs of time.

1.3.2 $k_1 > 0$: Allowing for early agreements

Having seen the difficulties to incentivize information aggregation in a straw poll that is followed by a decisive vote, we turn towards two-period voting setups that already allow for agreement in the first stage. Formally, we consider $k_1 > 0$, so that the defendant can be convicted or acquitted in $t = 1$ according to decision rule (1.1), that is, if $x_1 > N - k_1$, or $x_1 < k_1$ respectively. For example with $k_1 = 1$, a decision in the first vote can be made unanimously and the final vote follows only if there is no unanimous agreement in the first period. Thereby jurors can avoid entering the second period and save costs of time. As a result, incentives for informative voting are influenced. Jurors trade off the influence of their vote on the information of the pivotal juror in period 2 against the opportunity that their vote causes an earlier decision and saves costs of time. This trade-off is solved in favor of informative voting in the first period. While the effect in the case of being pivotal as in a straw poll remains, voting informatively in the first period increases the probability that if an earlier decision is made it is the juror's preferred one. As a result, juries vote informatively even if their minimal diversity is larger than 0.

Proposition 1.1. *Suppose each juror j has voting costs of $c \geq 0$ and $k_1 > 0$. The profile of strategies and beliefs (σ, μ) constitutes a Perfect Bayesian Equilibrium if and only if preferences satisfy one of the following conditions:*

1. *For $\lambda^N_{k_2} \in \{1, \dots, k_1\}$,*

$$\beta(k_1 - 1, N) - c \cdot \alpha(k_1 - 1) \;\leq\; q_j \;\leq\; \beta(k_1, N) + c \cdot \gamma(k_1) \quad \forall j \in N, \tag{1.4}$$

2. *for $\lambda^N_{k_2} \in \{k_1 + 1, \dots, N - k_1\}$,*

$$\beta(\lambda^N_{k_2} - 1, N) - c \cdot \alpha(\lambda^N_{k_2} - 1) \;\leq\; q_j \;\leq\; \beta(\lambda^N_{k_2}, N) + c \cdot \gamma(\lambda^N_{k_2}) \quad \forall j \in N, \tag{1.5}$$

3. *for $\lambda^N_{k_2} \in \{N - k_1 + 1, \dots, N\}$,*

$$\beta(N - k_1, N) - c \cdot \gamma(k_1) \;\leq\; q_j \;\leq\; \beta(N - k_1 + 1, N) + c \cdot \alpha(k_1 - 1) \quad \forall j \in N, \tag{1.6}$$

where

$$\alpha(x) = \binom{N-1}{x}^{-1}\binom{N-1}{k_1-1}\frac{(2p-1)\left[(1-p)^{k_1-1}p^{N-k_1} - (1-p)^{N-k_1}p^{k_1-1}\right]}{(1-p)^x p^{N-x} + (1-p)^{N-x}p^x} > 0,$$
$$\gamma(x) = \binom{N-1}{x-1}^{-1}\binom{N-1}{k_1-1}\frac{(2p-1)\left[(1-p)^{k_1-1}p^{N-k_1} - (1-p)^{N-k_1}p^{k_1-1}\right]}{(1-p)^x p^{N-x} + (1-p)^{N-x}p^x} > 0.$$

The terms $\alpha(x)$ and $\gamma(x)$ represent weighted conditional probabilities for those pivotal scenarios in which voting informatively saves costs without changing the jury's final decision for decision rule k_1. Note that the impossibility result of Coughlan (2000) follows immediately for $k_1 = 0$ or $c = 0$[9].

In order to provide intuition why heterogeneous juries aggregate information once the first vote can have consequences also, we highlight the changes on incentives compared to the previously discussed case of a straw poll in which each juror faces the same trade-off in the unique case of being pivotal. Jurors with different preferences solve this trade-off differently and some prefer to misinform the others in a straw poll in order to manipulate k_2's belief and make him implement a superior decision.

[9] The original result Coughlan (2000) also covers the case where preferences are such that the pivotal juror in $t = 2$ always prefers C or A independently of revealed information. As discussed, we neglect these cases by Assumption 1.1, as informative voting is trivially equilibrium behavior for those. We focus instead on juries for which a straw poll does not always aggregate information.

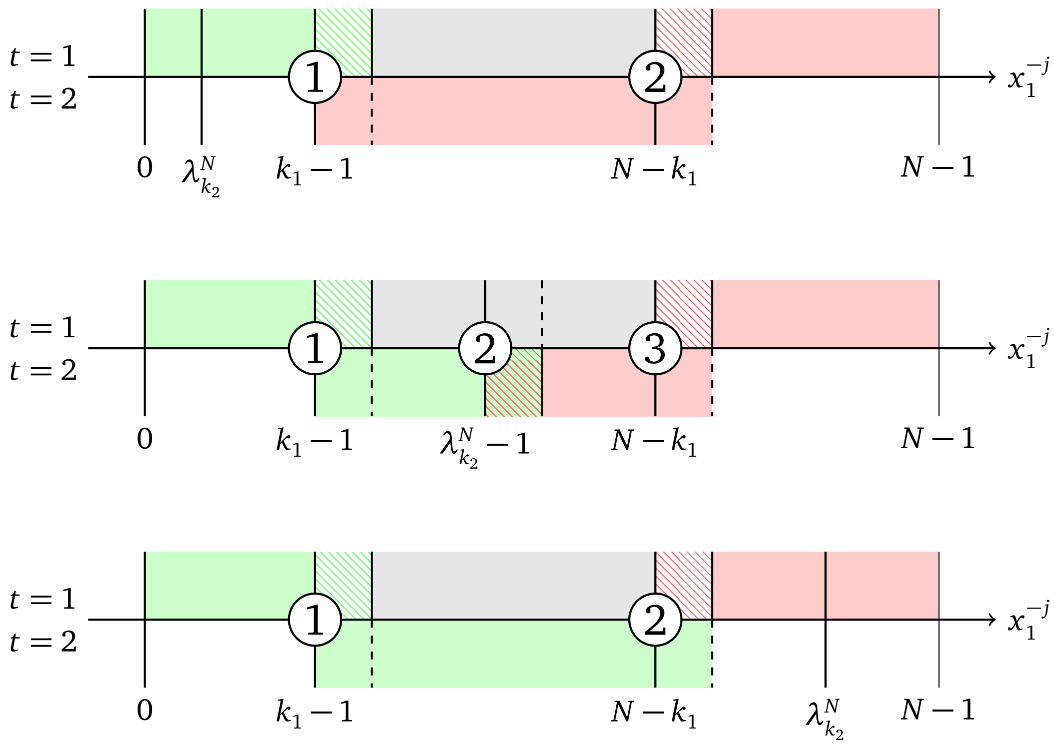

Figure 1.3: Pivotal scenarios in $t = 1$ for condition (1.4), (1.5), and (1.6).

Let us treat the case where $\lambda_{k_2}^N \in \{N - k_1 + 1, \ldots, N\}$ first, which corresponds to condition (1.5). In this case, the decision of juror k_2 in period 2 is not predetermined if the vote takes place. That is, the k_2 could decide for either A or C in $t = 2$ depending on $x_1 \in \{k_1, \ldots, N - k_1\}$ that leads to a second vote. Consider some juror $j \in N$ and suppose that the other jurors vote according to (σ, μ) , i.e., informatively in the first and sincerely in the second period. When the first vote can already have consequences, 2 more scenarios arise in which jurors are pivotal compared to a straw poll. Denote by x_{-j} the number of (informative) C votes of jury N except juror j. There are now three pivotal scenarios for j who faces the following trade-offs.

1. $x_{-j} = k_1 - 1$: If j votes C, the decision is delayed and $d_2 = A$ at additional costs c. If j votes A the defendant is acquitted immediately.

2. $x_{-j} = \lambda_{k_2}^N - 1$: If j votes C, the decision is delayed and $d_2 = C$ at additional costs c. If j votes A, the decision is delayed and $d_2 = A$ at additional costs c.

3. $x_{-j} = N - k_1$: If j votes C, the defendant is convicted immediately. If j votes A the decision is delayed and $d_2 = C$ at additional costs c.

See Figure **??** for a graphical representation. Scenario 2 corresponds to the unique pivotal scenario each juror faces in a straw poll and this trade-off is solved in favor of voting C if and only if condition (1.3) is satisfied, i.e., j has the same conviction threshold as k_2. In Scenario

1, regardless of j's vote, the jury will decide in favor of A in the second period. By voting A in the first period, however, j can cause this decision earlier and thus save costs of time. Voting A is strictly better in Scenario 1. An analogous argument holds in Scenario 3. Here, the jury's decision will be C regardless of j's vote in the first period, which he can implement earlier by voting C. Therefore, voting C is strictly better in Scenario 3.

Now, whereas the likelihood of Scenario 2 is untouched, it is more likely for j that Scenario 1 occurs than Scenario 3 if he receives an i-signal, and Scenario 3 appears more likely if he receives a g-signal[10]. In expectation, voting A becomes more attractive after j observes an i-signal and voting C gains attraction otherwise. As these effects add up, the incentives for informative voting improve. Moreover, the intervals in which preferences q_j of all $j \in N$ have to be located to vote informatively expand linearly in c with factors $\alpha(\cdot)$ and $\gamma(\cdot)$ representing the weighted cost saving effect due to Scenario 1 and 3. As a result incentives are provided,even for heterogeneous juries, to aggregate information truthfully in the first vote for any $c > 0$.

Complementing this, condition (1.4) and (1.6) cover the cases in which the pivotal juror's decision in the second period is always C, or A respectively, if reached. In these cases there are only two pivotal scenarios. If the second vote results always in C, then both decisions about an earlier agreement and about the jury's final decision coincide so that both Scenario 1 and 2 occur together if $k_1 - 1$ of the remaining jury votes C in the first period. Conversely, if the pivotal jurors always votes A in the second period, Scenarios 2 and 3 occur together by an analogous argument. See Figures **??** and 1.3 for a graphical representations. Besides that, the argument is similar as for the first case; if the first vote is potentially consequential, the opportunity to save costs of time improve any juries incentives to vote informatively in the first period.

The improvement of juror's incentives to vote informatively depends on the ability to avoid costs of time without changing the jury's final verdict. In order to achieve this effect, it is essential to make the first period vote potentially consequential and to allow for the presence of costs of time. Its strength mirrors in the three conditions of Proposition 1.1 and depends on the positive factors $\alpha(\cdot)$ and $\gamma(\cdot)$ as well as on the costs parameter c. As an immediate consequence from Proposition 1.1 we can observe that juries of arbitrary heterogeneity vote informatively in equilibrium if costs of time are high enough.[11].

---

[10] For this argument we only need that the number of votes of the other jurors to trigger Scenario 1 is smaller than $N/2$, and for Scenario 3 larger than $N/2$ respectively. This is assured by assuming symmetry in $k_t < N/2$. Symmetry, however, is not necessary but simplifies the analysis dramatically.

[11] Note as well that informative voting is trivially equilibrium behavior in any circumstances if p approaches 1. As signals are (almost) perfect, information disparity vanishes and jurors vote according to their (almost) perfect signal.

1.4 Welfare

In the previous chapter we showed that making the first vote consequential provides incentives for heterogeneous juries to reveal their information. We are now concerned with the effect of this change on the jury's welfare. We begin by providing a short numerical example to stress the extent by which requirements on the juries' heterogeneity are relaxed due to potential early agreement before going on to discuss resulting welfare effects.

1.4.1 Inferiority of Straw Polls

The potential for early agreement allows jurors to avoid the costly second vote whenever agreement among them is broad enough. Although this new possibility drives the improvement of incentives for informative voting, it is not immediately clear that jurors are better-off compared to a non-binding straw poll. In the latter setup the conviction threshold of the pivotal juror in the second vote pins down the unique pivotal scenario for all jurors. This is not necessarily the case anymore if $k_1 > 0$ is chosen relatively large. Either (1.4) or (1.6) prescribe the conditions for an informative equilibrium in this case. Both have the common feature that the minimal number of C votes required in the first period to finally convict the defendant is pinned down by the threshold for early agreement, k_1, and not by the jurors conviction thresholds. If both are too far apart, some relatively extreme homogeneous juries do not vote informatively in the first period any longer. In a straw poll, however, information is fully aggregated and the optimal decision from each jurors point of view is implemented. If saved costs are not enough to make up for this loss these juries are in fact worse off with $k_1 > 0$.

Example 1.1. *As an example for such a case consider a homogeneous jury of $N = 20$ jurors with a common conviction threshold of $\lambda = 1$. This jury votes informatively in a non-binding straw poll but might not if early agreement is possible with $k_1 = 5$. As $\lambda < k_1$, the jury will convict the defendant whenever the second vote happens. A juror who is pivotal in the first period when 4 out of the other 19 jurors vote C informatively prefers to vote C instead of A, even if he receives an i-signal. Voting A in this situation saves costs but leads to an undesirable acquittal because $\lambda < 4 = k_1 - 1$. The incentives to save costs had to be unreasonably large in order to justify voting A after receiving an i-signal.*

There is, however, always a voting rule for the first period that overcomes this issue such that a consequential first period vote does not alter a homogeneous jury's behavior in equilibrium. As a result, each such jury is strictly better off when it can avoid costs of time. Moreover, we can extend this insight to some heterogeneous juries. We show that the outcome of an informative equilibrium in setups with a consequential first period vote outperforms the upper bound of any equilibrium in straw poll setups for those juries.

Proposition 1.2. *Fix some k_2 and consider a jury N with bounded heterogeneity. For some $k_1 > 0$, every juror $j \in N$ is strictly better off in the equilibrium constituted by (σ, μ) than in any equilibrium of any setup with a straw poll ($k_1 = 0$) in the first period.*

We proceed in four steps. First, we show that any jury that votes informatively in a straw poll in equilibrium does the same for some decision rule $k_1 > 0$ in the first period. This is ensured if k_1 is chosen small enough. Second, we argue that these homogeneous juries are strictly better off under $k_1 > 0$. For small values of k_1, the jury's decision does not change compared to a straw poll but for some realizations of signals, it is made earlier. As this saves costs, jurors are better off in expectation. In the third step, we establish an upper bound on the jurors' utilities from a setup with a straw poll in the first period. In fact, the best jurors can achieve from a straw poll setup is their preferred decision in the second vote under full information. This, however, is the outcome of (σ, μ) for homogeneous juries. Finally, we demonstrate that there are indeed some heterogeneous juries, that can improve upon this benchmark for some $k_1 > 0$. Unlike in the case of homogeneous juries, the final decision is not optimal for all jurors in this case. Nevertheless, these inefficiencies are compensated by saved costs in expectation, if their level of heterogeneity is not too high.
Note that the extent of heterogeneity that is compatible with this pareto improvement depends on the cost saving effect which is identical to each juror. A higher value of c strictly improves welfare for any $k_1 > 0$. Hence, a stronger costs saving effect can compensate higher inefficiencies from imperfect decisions caused by higher levels of heterogeneity. Additionally, note that the welfare criterion used is very strict. We require juries to improve upon the upper bound of each juror from a straw poll. In considering weaker welfare measures, more heterogeneous juries could improve upon straw poll setups as well.

A non-binding straw poll in the first vote fails to provide the right incentives for heterogeneous juries to reveal their information truthfully, as Coughlan (2000) pointed out. We showed in the previous chapter that this can be accomplished if the first period vote can have consequences. Now, we establish that straw polls not only hurt incentives for information aggregation in a two-period voting setup but also lower the jurors expected utility compared to setups in which juries can already make a decision in the first period. In fact, homogeneous juries that reveal their information truthfully in a straw poll in equilibrium are strictly better off if the first vote allows for an earlier decision. One can conclude that straw polls rarely achieve information aggregation and if they do, a setup with opportunities for earlier agreement performs strictly better.

1.4.2 Designer's Choice of $k_1 > 0$

As a final step we consider a designer who can change the voting rule from a straw poll in the first vote to $k_1 > 0$. In many situations there are legal requirements or regulations on the majority that are needed for a decision which a designer cannot influence. We account for that by treating k_2, the majority rule of the final vote, as given and consider $k_1 < k_2$. If a decision is made early, the majority in favor is at least as large as legally required. In the following, we discuss which threshold for agreement the designer should set when the jury's preferences are known or unknown to her.

Let us consider the latter case first, where either the designer does not know the jurors' preferences or she has to commit to a decision rule for the first vote before the jury members are announced. In this case, the designer cannot observe the exact value of c either, which is part of the jurors' preferences. However, it is known that some costs $c > 0$ are present. For the reasons discussed previously, allowing for a decision in the first period increases the jurors welfare but if k_1 is chosen too high some relative extreme homogeneous juries are worse off. The following result follows immediately from Proposition 1.2.

Proposition 1.3. *Suppose the designer cannot observe the jurors preferences. She can increase the welfare of any jury that votes informatively in a straw poll by setting* $k_1 = 1$.

As discussed, a straw poll accomplishes information aggregation only if the jury is homogeneous. For this case however, there is no disagreement among the jurors if all receive the same signal. Hence, implementing the unanimous decision earlier makes every juror better off and does not alter the jurors voting behavior. If the designer chose a $k_1 > 1$ there is are homogeneous juries (with extreme preferences) that would have revealed their information in a straw poll and subsequently made a decision unanimously but do not vote informatively anymore in the first vote which can already implement a decision. If their costs c, which are unknown to the designer, are small then these juries are worse off.

Now suppose that the designer knows the jurors' preferences and can adjust the voting rule in the first period to increase their welfare accordingly. In order to do so she considers all $k_1 > 0$ that induce informative voting in the first period for this jury and we call the set of all such values $\mathcal{K}_1^N$. In the next step we isolate a subset of $\mathcal{K}_1^N$, which we call $\widetilde{\mathcal{K}_1^N}$. For any $k_1 \in \widetilde{\mathcal{K}_1^N}$, the jury votes informatively in $t = 1$ and it's welfare is strictly higher compared to the upper bound of setups with a straw poll in the first round. Let us assume that the jury's heterogeneity is bounded from above such that $\widetilde{\mathcal{K}_1^N}$ is non-empty. From Proposition 1.2, we know that this in indeed the case for small degree of heterogeneity. Finally, we provide conditions to identify a $k_1 \in \widetilde{\mathcal{K}_1^N}$ which maximizes the jury's welfare.

Proposition 1.4. *Suppose the designer can observe the jurors preferences whose degree of heterogeneity is bounded from above. Then, the optimal* k_1 *for which* (σ, μ) *is an equilibrium is given as follows:*

1. *If* $\lambda_{k_2}^N \in \{k_1, \ldots, N - k_1 + 1\}$, *set*

$$k_1 = \max \widetilde{\mathcal{K}_1^N}.$$

2. *If* $\lambda^N_{k_2} < \frac{N}{2}$ *and* $\left\{k_1 \mid \lambda^N_{k_2} < k_1\right\} \subseteq \widetilde{\mathcal{K}}^N_1$, *set*

$$k_1 = \max\left\{k_1 \in \widetilde{\mathcal{K}}^N_1 \mid q_1 \geq \beta(k_1 - 1, N) - 2c \vee \lambda^N_{k_2} \in \left\{k_1, \dots, N - k_1 + 1\right\}\right\}$$

3. *If* $\lambda^N_{k_2} > \frac{N}{2}$ *and* $\left\{k_1 \mid \lambda^N_{k_2} > N - k_1 + 1\right\} \subseteq \widetilde{\mathcal{K}}^N_1$, *set*

$$k_1 = \max\left\{k_1 \in \widetilde{\mathcal{K}}^N_1 \mid q_N \leq \beta(N - (k_1 - 1), N) + 2c \vee \lambda^N_{k_2} \in \left\{k_1, \dots, N - k_1 + 1\right\}\right\}$$

Intuitively, whenever the choice from all candidates $k_1 \in \widetilde{\mathcal{K}}^N_1$ does not influence the final decision, then increasing k_1 as high as possible saves the most costs and, hence, is optimal to choose. If an additional increase in k_1 impacts the jury's decision, the designer must trade off the saved costs against the repercussions from interfering with the decision. The conditions above reflect this trade-off: The designer wants to choose a high $k_1 \in \widetilde{\mathcal{K}}^N_1$ in order to decrease expected costs for the jurors but not impact the jury's decision to an extend that outweighs the saved costs.

1.5 Concluding Remarks

This paper contributes to the literature on two period committee voting by additionally considering costs of time. We show that it is beneficial to grant an option for early agreement to the jurors as opposed to letting them engage in a straw poll first. Unlike in the straw poll setup of Coughlan (2000), even heterogeneous juries aggregate their information perfectly in the first vote once if it can already have consequences. Moreover, not only every homogeneous jury but also some heterogeneous juries are strictly better off in terms of expected utility when the jurors can make the decision earlier. We demonstrated how the ability to save costs of time positively influences the jurors incentives to aggregate information. Finally, we showed how a designer can profit from our insights when she can set the voting rule for the first period based on her information on the jurors' preferences.

The results are derived in a simple framework. There are numerous possible ways of identifying robustness of the effects on incentives for information aggregation from making the first vote consequential. Naturally, introducing a more sophisticated information structure or uncertainty about other players' types immediately come to mind. Moreover, the discussion on the designer's optimal choice gives rise to the consideration of endogenous costs. Higher costs increase incentives to vote informatively in setups which allow for agreement early. However, they negatively impact the jurors welfare. Even now the simple framework from this paper allows us to identify interesting effects that arise from introducing consequences in the first vote already.

This considerations can lead to interesting policy implications. We show that straw polls are not only dilatory, the reason why Robert III. et al. (2000) judged them "*meaningless*", but can also be easily outperformed in the ability to aggregate information if they can have consequences.

1.A Appendix

Proof of Proposition 1.1

We proceed by backward induction.
$t=2$: Suppose jurors vote informatively in $t=1$ and x_1 is the revealed number of C-votes from that period. Consistent with informative voting, any $j \in N$ believes with probability 1 that x_1 represents the amount of g-signals among the jurors. Accordingly, all $j \in N$ update their beliefs with Bayes' rule about the state of the world being G consistently to $\beta(x_1, N)$ for all $j \in N$, and for any signal s_j. Given the other jurors vote sincerely, j conditions his vote on being pivotal and votes C if and only if

$$\begin{aligned} -(1-\beta(x_1,N))q_j &\geq -\beta(x_1,N)\left(1-q_j\right) \\ \Leftrightarrow q_j &\leq \beta(x_1,N). \end{aligned}$$

That is, j votes sincerely in $t=2$.

$t=1$: The received signal determines the jurors' prior beliefs. Any $j \in N$ attaches probability p to the state being G whenever $s_j = g$, and $1-p$ otherwise. The jurors anticipate the outcome of the vote in $t=2$, if it is reached, for any realization of x_1. There are three cases to distinguish.

1. Suppose $\lambda_{k_2}^N \in \{1, \ldots, k_1\}$. In this case $d_2 = C$ because $\lambda_{k_2}^N \leq k_1$. Denote by $\widetilde{\sigma} = \left(\sigma_{-j}, \left(\widetilde{\sigma}_j^1, \sigma_j^2\right)\right)$ the jury's strategy profile where all jurors vote sincerely in $t=2$, and in $t=1$ all but j vote informatively and j deviates to vote contrarily to his received signal, i.e., A, if $s_j = g$ and C, if $s_j = i$. We compute j's expected utilities in $t=1$ from voting informatively, that is sticking to σ, as well as from deviating to $\widetilde{\sigma}$, that is providing misinformation.

$$\begin{aligned} EU_j\left[\sigma \mid \mu, s_j = g\right] =& \sum_{x=0}^{k_1-2}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}\left(-(1-q_j)\right) + \sum_{x=k_1-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}(-c) \\ &+ \sum_{x=k_1-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}(-c-q_j) + \sum_{x=N-k_1}^{N-1}\binom{N-1}{x}(1-p)^{x-1}p^{N-x+1}(-q_j). \\ EU_j\left[\widetilde{\sigma} \mid \mu, s_j = g\right] =& \sum_{x=0}^{k_1-1}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}\left(-(1-q_j)\right) + \sum_{x=k_1}^{N-k_1}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}(-c) \\ &+ \sum_{x=k_1}^{N-k_1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}\left(-c-q_j\right) + \sum_{x=N-k_1+1}^{N-1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}(-q_j). \end{aligned}$$

Juror j prefers to vote informatively after receiving a g-signal if and only if

$$EU_j\left[\sigma \mid \mu, s_j = g\right] - EU_j\left[\widetilde{\sigma} \mid \mu, s_j = g\right] \geq 0.$$

This condition is equivalent to

$$
\begin{aligned}
&(1-p)^{N-k_1}p^{k_1}(-c) + (1-p)^{k_1}p^{N-k_1}(-c-q_j) + (1-p)^{N-k_1+1}p^{k_1-1}(-q_j) \\
\geq &(1-p)^{N-k_1}p^{k_1}\left(-(1-q_j)\right) + (1-p)^{k_1-1}p^{N-k_1+1}(-c) + (1-p)^{N-k_1+1}p^{k_1-1}(-c-q_j)
\end{aligned} \tag{1.7}
$$

$$
\Leftrightarrow (1-p)^{k_1}p^{N-k_1}q_j - (1-p)^{N-k_1}p^{k_1}(1-q_j) \;\leq\; c(2p-1)\left[(1-p)^{k_1-1}p^{N-k_1} - (1-p)^{N-k_1}p^{k_1-1}\right] \tag{1.8}
$$

However, (1.8) is an equivalent reformulation of $q_j \leq \beta(k_1, N) + c \cdot \gamma(k_1 - 1)$. As a result, all jurors vote informatively after receiving a g-signal if and only if $q_j \leq \beta(k_1, N) + c \cdot \gamma(k_1 - 1)$ for any $j \in N$.
Analogously, given the other jurors vote informatively, a juror who receives an i-signal has the following expected utilities:

$$
\begin{aligned}
EU_j\left[\sigma \mid \mu, s_j = i\right] = &\sum_{x=0}^{k_1-2}\binom{N-1}{x}(1-p)^{N-x}p^x\left(-(1-q_j)\right) + \sum_{x=k_1-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^{N-x}p^x(-c) \\
&+ \sum_{x=k_1-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^x p^{N-x}(-c-q_j) + \sum_{x=N-k_1}^{N-1}\binom{N-1}{x}(1-p)^x p^{N-x}(-q_j) \\
EU_j\left[\widetilde{\sigma} \mid \mu, s_j = i\right] = &\sum_{x=0}^{k_1-1}\binom{N-1}{x}(1-p)^{N-x}p^x\left(-(1-q_j)\right) + \sum_{x=k_1}^{N-k_1}\binom{N-1}{x}(1-p)^{N-x}p^x(-c) \\
&+ \sum_{x=k_1}^{N-k_1}\binom{N-1}{x}(1-p)^x p^{N-x}\left(-c-q_j\right) + \sum_{x=N-k_1+1}^{N-1}\binom{N-1}{x}(1-p)^x p^{N-x}(-q_j)
\end{aligned}
$$

Analogously to the above case, $q_j \leq \beta(k_1, N) + c \cdot \gamma(k_1 - 1)$ is equivalent to

$$
EU_j\left[\sigma \mid \mu, s_j = i\right] - EU_j\left[\widetilde{\sigma} \mid \mu, s_j = i\right] \geq 0
$$

for all $j \in N$. Thus, jurors vote informatively after receiving an i-signal if and only if $q_j \leq \beta(k_1, N) + c \cdot \gamma(k_1 - 1)$ for all $j \in N$.

2. Suppose $\lambda_{k_2}^N \in \{k_1 + 1, \ldots, N - k_1\}$. This implies $d_2 = C$ if $x_1 \geq \lambda_{k_2}^N$, and $d_2 = A$ otherwise. Juror j's expected utilities are computed as follows in this case:

$$
\begin{aligned}
EU_j\left[\sigma \mid \mu, s_j = g\right] = &\sum_{x=0}^{k_1-2}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}\left(-(1-q_j)\right) + \sum_{x=k_1-1}^{\lambda_{k_2}^N-2}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}\left(-(1-q_j)-c\right) \\
&+ \sum_{x=\lambda_{k_2}^N-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}(-c) + \sum_{x=k_1-1}^{\lambda_{k_2}^N-2}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}(-c) \\
&+ \sum_{x=\lambda_{k_2}^N-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}(-q_j-c) + \sum_{x=N-k_1}^{N-1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}(-q_j)
\end{aligned}
$$

$$\begin{aligned}EU_j\left[\tilde{\sigma} \mid \mu, s_j = g\right] =& \sum_{x=0}^{k_1-1}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}\left(-(1-q_j)\right) + \sum_{x=k_1}^{\lambda_{k_2}^N-1}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}\left(-(1-q_j)-c\right)\\ &+ \sum_{x=\lambda_{k_2}^N}^{N-k_1}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}(-c) + \sum_{x=k_1}^{\lambda_{k_2}^N-1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}(-c)\\ &+ \sum_{x=\lambda_{k_2}^N}^{N-k_1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}(-q_j-c) + \sum_{x=N-k_1+1}^{N-1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}(-q_j)\end{aligned}$$

As before, we reformulate $q_j \leq \beta(\lambda_{k_2}^N, N) + c \cdot \gamma(\lambda_{k_2}^N)$ to the equivalent expression

$$\begin{aligned}&\binom{N-1}{k_1-1}(1-p)^{N-k_1}p^{k_1}(-c) + \binom{N-1}{k_1-1}(1-p)^{k_1}p^{N-k_1}(-c) + \binom{N-1}{\lambda_{k_2}^N-1}(1-p)^{\lambda_{k_2}^N}p^{N-\lambda_{k_2}^N}(-q_j)\\ \geq&\binom{N-1}{\lambda_{k_2}^N-1}(1-p)^{N-\lambda_{k_2}^N}p^{\lambda_{k_2}^N}\left(-(1-q_j)\right) + \binom{N-1}{N-k_1}(1-p)^{k_1-1}p^{N-k_1+1}(-c)\\ &+\binom{N-1}{N-k_1}(1-p)^{N-k_1+1}p^{k_1-1}(-c),\end{aligned} \tag{1.9}$$

which in turn is equivalent to

$$EU_j\left[\sigma \mid \mu, s_j = g\right] - EU_j\left[\tilde{\sigma} \mid \mu, s_j = g\right] \geq 0.$$

Therefore, no juror has a profitable deviation from strategy profile σ after receiving a g-signal if and only if $q_j \leq \beta(\lambda_{k_2}^N, N) + c \cdot \gamma(\lambda_{k_2}^N)$ for all $j \in N$.
Analogously, a juror who receives an i-signal has the following expected utilities:

$$\begin{aligned}EU_j\left[\sigma \mid \mu, s_j = i\right] =& \sum_{x=0}^{k_1-2}\binom{N-1}{x}(1-p)^{N-x}p^{x}\left(-(1-q_j)\right) + \sum_{x=k_1-1}^{\lambda_{k_2}^N-2}\binom{N-1}{x}(1-p)^{N-x}p^{x}\left(-(1-q_j)-c\right)\\ &+ \sum_{x=\lambda_{k_2}^N-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^{N-x}p^{x}(-c) + \sum_{x=k_1-1}^{\lambda_{k_2}^N-2}\binom{N-1}{x}(1-p)^{x}p^{N-x}(-c)\\ &+ \sum_{x=\lambda_{k_2}^N-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^{x}p^{N-x}(-q_j-c) + \sum_{x=N-k_1}^{N-1}\binom{N-1}{x}(1-p)^{x}p^{N-x}(-q_j)\end{aligned}$$

$$\begin{aligned}EU_j\left[\tilde{\sigma} \mid \mu, s_j = i\right] =& \sum_{x=0}^{k_1-1}\binom{N-1}{x}(1-p)^{N-x}p^{x}\left(-(1-q_j)\right) + \sum_{x=k_1}^{\lambda_{k_2}^N-1}\binom{N-1}{x}(1-p)^{N-x}p^{x}\left(-(1-q_j)-c\right)\\ &+ \sum_{x=\lambda_{k_2}^N}^{N-k_1}\binom{N-1}{x}(1-p)^{N-x}p^{x}(-c) + \sum_{x=k_1}^{\lambda_{k_2}^N-1}\binom{N-1}{x}(1-p)^{x}p^{N-x}(-c)\\ &+ \sum_{x=\lambda_{k_2}^N}^{N-k_1}\binom{N-1}{x}(1-p)^{x}p^{N-x}(-q_j-c) + \sum_{x=N-k_1+1}^{N-1}\binom{N-1}{x}(1-p)^{x}p^{N-x}(-q_j)\end{aligned}$$

We reformulate $\beta(\lambda_{k_2}^N - 1, N) - c \cdot \alpha(\lambda_{k_2}^N - 1) \leq q_j$ equivalently to

$$
\begin{aligned}
&\binom{N-1}{k_1-1}(1-p)^{N-k_1+1}p^{k_1-1}(-c) + \binom{N-1}{k_1-1}(1-p)^{k_1-1}p^{N-k_1+1}(-c) \\
+&\binom{N-1}{\lambda_{k_2}^N-1}(1-p)^{\lambda_{k_2}^N-1}p^{N-\lambda_{k_2}^N+1}(-q_j) \\
\leq&\binom{N-1}{\lambda_{k_2}^N-1}(1-p)^{N-\lambda_{k_2}^N+1}p^{\lambda_{k_2}^N-1}\left(-(1-q_j)\right) + \binom{N-1}{N-k_1}(1-p)^{k_1}p^{N-k_1}(-c) \\
+&\binom{N-1}{N-k_1}(1-p)^{N-k_1}p^{k_1}(-c) \qquad (1.10)
\end{aligned}
$$

which is in turn equivalent to

$$EU_j\left[\sigma \mid \mu, s_j = i\right] - EU_j\left[\widetilde{\sigma} \mid \mu, s_j = i\right] \geq 0.$$

We can now follow that all jurors prefer to vote informatively in the first period after receiving an i-signal if and only if $\beta(\lambda_{k_2}^N - 1, N) - c \cdot \alpha(\lambda_{k_2}^N - 1) \leq q_j$ for all $j \in N$.

3. Suppose $\lambda_{k_2}^N \in \left\{N - k_1, \ldots, N\right\}$. That implies $d_2 = A$ because $\lambda_{k_2}^N \geq N - k_1 + 1$. We compute j's expected utilities from playing the equilibrium strategy σ and the previously defined deviation $\widetilde{\sigma}$ as follows:

$$
\begin{aligned}
EU_j\left[\sigma \mid \mu, s_j = g\right] =& \sum_{x=0}^{k_1-2}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}(-(1-q_j)) + \sum_{x=k_1-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}\left(-(1-q_j)-c\right) \\
&+ \sum_{x=k_1-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}(-c) + \sum_{x=N-k_1}^{N-1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}\left(-q_j\right) \\
EU_j\left[\widetilde{\sigma} \mid \mu, s_j = g\right] =& \sum_{x=0}^{k_1-1}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}(-(1-q_j)) + \sum_{x=k_1}^{N-k_1}\binom{N-1}{x}(1-p)^{N-x-1}p^{x+1}\left(-(1-q_j)-c\right) \\
&+ \sum_{x=k_1}^{N-k_1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}(-c) + \sum_{x=N-k_1+1}^{N-1}\binom{N-1}{x}(1-p)^{x+1}p^{N-x-1}(-q_j)
\end{aligned}
$$

Reformulating $q_j \leq \beta(N - k_1 + 1, N) + c \cdot \alpha(k_1 - 1)$ yields

$$
\begin{aligned}
&(1-p)^{N-k_1}p^{k_1}\left(-(1-q_j)-c\right) + (1-p)^{k_1}p^{N-k_1}(-c) + (1-p)^{N-k_1+1}p^{k_1-1}(-q_j) \\
\geq&(1-p)^{N-k_1}p^{k_1}\left(-(1-q_j)\right) + (1-p)^{k_1-1}p^{N-k_1+1}\left(-(1-q_j)-c\right) + (1-p)^{N-k_1+1}p^{k_1-1}(-c) \qquad (1.11)
\end{aligned}
$$

Analogously to the previous cases, (1.11) is equivalent to

$$EU_j\left[\sigma \mid \mu, s_j = g\right] - EU_j\left[\widetilde{\sigma} \mid \mu, s_j = g\right] \geq 0.$$

Therefore, all jurors vote informatively after receiving a g-signal if and only if $q_j \leq \beta(N - k_1 + 1, N) + c \cdot \alpha(k_1 - 1)$ for all $j \in N$.

Analogously, we compute expected utilities for the case where $s_j = i$.

$$\begin{aligned}EU_j\left[\sigma \mid \mu, s_j = i\right] = &\sum_{x=0}^{k_1-2}\binom{N-1}{x}(1-p)^{N-x}p^x(-(1-q_j)) + \sum_{x=k_1-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^{N-x}p^x\left(-(1-q_j)-c\right)\\ &+\sum_{x=k_1-1}^{N-k_1-1}\binom{N-1}{x}(1-p)^x p^{N-x}(-c) + \sum_{x=N-k_1}^{N-1}\binom{N-1}{x}(1-p)^x p^{N-x}\left(-q_j\right)\\ EU_j\left[\tilde{\sigma} \mid \mu, s_j = i\right] = &\sum_{x=0}^{k_1-1}\binom{N-1}{x}(1-p)^{N-x}p^x(-(1-q_j)) + \sum_{x=k_1}^{N-k_1}\binom{N-1}{x}(1-p)^{N-x}p^x\left(-(1-q_j)-c\right)\\ &+\sum_{x=k_1}^{N-k_1}\binom{N-1}{x}(1-p)^x p^{N-x}(-c) + \sum_{x=N-k_1+1}^{N-1}\binom{N-1}{x}(1-p)^x p^{N-x}(-q_j)\end{aligned}$$

From $q_j \geq \beta(N-k_1, N) - c\cdot\gamma(k_1-1)$ we derive the equivalent formulation

$$\begin{aligned}&(1-p)^{N-k_1+1}p^{k_1-1}\left(-(1-q_j)-c\right) + (1-p)^{k_1-1}p^{N-k_1+1}(-c) + (1-p)^{N-k_1}p^{k_1}\left(-q_j\right)\\ \leq &(1-p)^{N-k_1+1}p^{k_1-1}(-(1-q_j)) + (1-p)^{k_1}p^{N-k_1}\left(-(1-q_j)-c\right) + (1-p)^{N-k_1}p^{k_1}(-c)\end{aligned} \tag{1.12}$$

Again, (1.12) is equivalent to $q_j \geq \beta(N-k_1, N) - c\cdot\gamma(k_1-1)$. As

$$EU_j\left[\sigma \mid \mu, s_j = i\right] - EU_j\left[\tilde{\sigma} \mid \mu, s_j = i\right] \geq 0$$

is equivalent to (1.12) in turn, jurors vote informatively after receiving an i-signal if and only if $q_j \geq \beta(N-k_1, N) - c\cdot\gamma(k_1-1)$ for all $j \in N$.

Summing up, (1.4) - (1.6) characterize sufficient and necessary conditions for the profile $\left(\sigma_j, \mu_j\right)_{j=1}^N$ to constitute an equilibrium. □

Proof of Proposition 1.2

Fix some k_2 for each of the following steps.
Step 1: Any jury that votes informatively in a straw poll in equilibrium does so as well in a setup with some $k_1 > 0$.
Consider $k_1 = 1$.
We show that if jury N votes informatively in $t = 1$ in equilibrium under $k_1 = 0$ then it votes informatively in $t = 1$ in equilibrium as well for $k_1 = 1$.
If jury N votes informatively under $k_1 = 0$ their preferences have a minimal diversity of 0. That is, for some $\lambda \in \{1, \ldots, N\}$,

$$\beta(\lambda-1, N) \leq q_j \leq \beta(\lambda, N) \quad \forall j \in N. \tag{1.13}$$

Depending on λ there are three cases to consider:

(i) $\lambda \in \{2, \ldots, N-1\}$: For $k_1 = 1$, (σ, μ) constitutes an equilibrium if condition (1.5) holds. This is implied by (1.13).

(ii) $\lambda = 1$: We know that

$$\beta(0,N) - c \cdot \alpha(0) \;\leq\; q_j \quad \forall j \;\in\; N, \tag{1.14}$$

by Assumption 1.1. Moreover, from (1.13) we know

$$q_j \;\leq\; \beta(1,N) \;\leq\; \beta(1,N) + c \cdot \gamma(1) \quad \forall j \;\in\; N. \tag{1.15}$$

The bounds (1.14) and (1.15) coincide with those from (1.4) which establish informative voting in equilibrium in $t = 1$ for jury N.

(iii) $\lambda = N$: Analogously to (b), we establish the bounds from (1.6). By Assumption 1.1 we have

$$q_j \;\leq\; \beta(N,N) \;\leq\; \beta(N,N) + c \cdot \alpha(0) \quad \forall j \;\in\; N, \tag{1.16}$$

and (1.13) yields the lower bound

$$\beta(N-1,N) - c \cdot \gamma(1) \;\leq\; q_j \quad \forall j \;\in\; N. \tag{1.17}$$

By combining both we establish condition (1.6) for all $j \in N$ so that voting informatively in $t = 1$ is an equilibrium behavior.

Step 2: Any juror of a jury that votes informatively in a straw poll in equilibrium is strictly better off in the equilibrium constituted by (σ, μ) with some $k_1 > 0$.
We compare the jurors' ex-ante expected utilities from a non-binding straw poll ($k_1 = 0$) and a potentially consequential first period vote with $k_1 = 1$. Note that in both cases the same juror k_2 is pivotal in the second vote. Also by Assumption 1.1, $\lambda_{k_2} \notin \{0, N+1\}$. For any λ_{k_2}, the final decision is the same for both $k_1 = 0$ and $k_1 = 1$. But in some cases, namely if all jurors receive the same signal, the process in terminated is the first period already under $k_1 = 1$, whereas jurors have to vote again after a straw poll at additional costs c. Therefore, any juror of that homogeneous jury is better off in the equilibrium (σ, μ) with $k_1 = 1$ than in a straw poll setup.
Step 3: With a straw poll in the first period, no juror can get higher expected utility than in the equilibrium with an informative straw poll, i.e., (σ, μ) with $k_1 = 0$.
Informative voting in a straw poll requires a homogeneous jury, where every juror has the same conviction threshold. In equilibrium, jurors reveal their information truthfully and agree on a decision under full information unanimously, because the decision is optimal for each (homogeneous) juror. There are no other sources that impact utilities, in particular ju-

rors cannot agree earlier and save costs. A setup that ensures jurors always the optimal decision from their points of view given full information can not be improved upon. Therefore, the jurors' utility levels in this equilibrium will serve in the following as an upper bound on the jurors' expected utilities in any equilibrium in a straw poll setup.

This upper bound on jurors' expected utilities for any equilibrium of setups with straw polls is given for each $j \in N$ by

$$\sum_{x=0}^{\lambda_j^N-1} \binom{N}{x} (1-p)^{N-x} p^x \left(-\left(1-q_j\right)\right) + \sum_{x=\lambda_j^N}^{N} \binom{N}{x} (1-p)^x p^{N-x} \left(-q_j\right) - 2c. \tag{1.18}$$

Step 4: Even some heterogeneous juries are strictly better off in the informative equilibrium of a voting setup that admits agreement in the first period than in any equilibrium of a straw poll setup.

Note that his statement is true for homogeneous juries by Step 2 and 3.

Now consider a heterogeneous jury which satisfies one of the conditions of Proposition 1.1 for $k_1 = 1$, so that the jurors vote informatively in the first period. The pivotal juror in the second vote is k_2 and the other jurors' conviction thresholds differ from $\lambda_{k_2}^N$ at most by 1. We show that any juror's expected utility is strictly higher in the informative equilibrium under $k_1 = 1$ than the upper bound (1.18) established in Step 3.

(i) $\lambda_j^N = \lambda_{k_2}^N$: In equilibrium, j's preferred decision given full information is always implemented by juror k_2 and in some cases he saves costs of time. This is strictly better than the upper bound of any equilibrium in a straw poll setup.

(ii) $\lambda_j^N = \lambda_{k_2}^N - 1$: For a simpler notation we set $\lambda_{k_2}^N \equiv \lambda$ and consider jurors with preferences q_j such that

$$q_j > \beta(\lambda-1,N) - c \cdot \min\left\{\alpha(\lambda-1), 2\binom{N}{\lambda-1}^{-1} \frac{(1-p)^N + p^N}{(1-p)^{N-\lambda+1}p^{\lambda-1} + (1-p)^{\lambda-1}p^{N-\lambda+1}}\right\},$$

where $\alpha(\lambda-1)$ is defined as in Proposition 1.1 for $k_1 = 1$. Note that these jurors vote informatively in equilibrium in the first vote. We compare the expected utility in the informative equilibrium with $k_1 = 1$ to the upper bound of any straw poll setup (1.18), in which every juror votes twice but his preferred decision given full information is made in the second vote for sure. When early agreement is possible, jurors might save costs of time but the jury's decision is sub optimal for jurors with $\lambda_j < \lambda_{k_2}^N$ if the number of g-signals among all jurors is between λ_j and $\lambda_{k_2}^N$.

The net effect on expected utilities from early agreement of the gain by cost saving and the loss by a sub optimal decision is positive for jurors with preferences as specified above, because

$$2c \cdot \left[(1-p)^N + p^N\right] + \binom{N}{\lambda-1}\left[(1-p)^{N-\lambda+1}p^{\lambda-1}(-(1-q_j)) - (1-p)^{\lambda-1}p^{N-\lambda+1}(-q_j)\right] > 0$$
$$\Leftrightarrow q_j > \beta(\lambda-1,N) - 2c \cdot \binom{N}{\lambda-1}^{-1} \frac{(1-p)^N + p^N}{(1-p)^{N-\lambda+1}p^{\lambda-1} + (1-p)^{\lambda-1}p^{N-\lambda+1}}.$$

(iii) $\lambda_j^N = \lambda_{k_2}^N + 1$: Consider jurors with preferences q_j such that

$$q_j < \beta(\lambda,N) + c \cdot \min\left\{\gamma(\lambda), 2\binom{N}{\lambda}^{-1} \frac{(1-p)^N + p^N}{(1-p)^{N-\lambda}p^{\lambda} + (1-p)^{\lambda}p^{N-\lambda}}\right\},$$

where $\gamma(\lambda)$ is defined as in Proposition 1.1 for $k_1 = 1$. Analogously to the previous case, the net effect on expected utilities from early agreement of the gain by cost saving and the loss by a sub optimal decision is positive for these jurors, because

$$2c \cdot \left[(1-p)^N + p^N\right] + \binom{N}{\lambda}\left[(1-p)^{\lambda}p^{N-\lambda}(-q_j) - (1-p)^{N-\lambda}p^{\lambda}(-(1-q_j))\right] > 0$$
$$\Leftrightarrow q_j < \beta(\lambda,N) - 2c \cdot \binom{N}{\lambda}^{-1} \frac{(1-p)^N + p^N}{(1-p)^{N-\lambda}p^{\lambda} + (1-p)^{\lambda}p^{N-\lambda}}.$$

□

Proof of Proposition 1.3

We know from the proof Proposition 1.2 that juries which vote informatively with $k_1 = 1$ are strictly better off than with a straw poll in the first period. In addition, any higher value of k_1 does not provide incentives to vote informatively to any homogeneous jury, like a straw poll would. Suppose the designer sets $k_1 = 2$ and consider a homogeneous jury with common conviction threshold of $\lambda_j^N = 1$ for all $j \in N$, i.e.,

$$\beta(0,N) \leq q_j \leq \beta(1,N) \quad \forall j \in N.$$

By (1.4), jurors vote informatively in the first vote with $k_1 = 2$ if and only if

$$\beta(1,N) - c \cdot \alpha(1) \leq q_j \leq \beta(2,N) + c \cdot \gamma(2) \quad \forall j \in N.$$

As c is unknown to the designer, she cannot rule out that

$$q_j \in [\beta(0,N), \beta(1,N) - c),$$

in which case this jury is worse off compared to a straw poll. This argument holds for any higher $k_1 > 0$ as well. □

Proof of Proposition 1.4

We proceed in three steps.
Step 1: Definition of $\mathcal{K}_1^N$.
Denote by $\mathcal{K}_1^N$ the set of all $k_1 < N/2$ for which jury N votes informatively in the first vote.

Definition 1.3. $k_1 \in \mathcal{K}_1^N$ *if and only if for all* $j \in N$, q_j *satisfies one of the conditions* (1.4), (1.5) *or* (1.6) *of Proposition 1.1 for* k_1.

The set $\mathcal{K}_1^N$ is non-empty if the jury's heterogeneity is bounded from above by an according value which is assumed.
Step 2: Definition of $\widetilde{\mathcal{K}}_1^N$.
Denote by $\widetilde{\mathcal{K}}_1^N$ the set of all $k_1 \in \mathcal{K}_1^N$ which make all jurors $j \in N$ (weakly) better off compared to the best equilibrium outcome of a straw poll in the first period. Recall from the proof of Proposition 1.2, the best equilibrium for any juror in a setup with a straw poll is that of a homogeneous jury, where information is aggregated in the straw poll and every jury agrees to the jury's decision in $t = 2$. The expected utility in that equilibrium for any $j \in N$ is given by (1.18).
The expected utility of any $j \in N$ from (σ, μ) and $k_1 > 0$ depends on the relative position of k_1 and $\lambda_{k_2}^N$. We have to distinguish three cases.

(i) For $\lambda_{k_2}^N < k_1$, j's expected utility from (σ, μ) is given by

$$\sum_{x=0}^{k_1-1} \binom{N}{x} (1-p)^{N-x} p^x \left(-(1-q_j)\right) + \sum_{x=k-1}^{N} \binom{N}{x} (1-p)^x p^{N-x} \left(-q_j\right)$$
$$-\left(1 + \sum_{x=k_1}^{N-k_1} \binom{N}{x} \left((1-p)^{N-x} p^x + (1-p)^x p^{N-x}\right)\right) c. \tag{1.19}$$

(ii) For $\lambda_{k_2}^N \in \{k_1, \ldots, N-k_1\}$, j's expected utility from (σ, μ) is given by

$$\sum_{x=0}^{\lambda_{k_2}^N-1} \binom{N}{x} (1-p)^{N-x} p^x \left(-(1-q_j)\right) + \sum_{x=\lambda_{k_2}^N}^{N} \binom{N}{x} (1-p)^x p^{N-x} \left(-q_j\right)$$
$$-\left(1 + \sum_{x=k_1}^{N-k_1} \binom{N}{x} \left((1-p)^{N-x} p^x + (1-p)^x p^{N-x}\right)\right) c. \tag{1.20}$$

(iii) For $\lambda_{k_2}^N > N - k_1$, j's expected utility from (σ, μ) is given by

$$\sum_{x=0}^{N-k_1} \binom{N}{x} (1-p)^{N-x} p^x \left(-\left(1-q_j\right)\right) + \sum_{x=N-k+1}^{N} \binom{N}{x} (1-p)^x p^{N-x} \left(-q_j\right)$$
$$-\left(1 + \sum_{x=k_1}^{N-k_1} \binom{N}{x} \left((1-p)^{N-x} p^x + (1-p)^x p^{N-x}\right)\right) c. \tag{1.21}$$

We can now derive conditions for which the equilibrium (σ, μ) with $k_1 > 0$ is (weakly) better for any $j \in N$ than the best equilibrium with a straw poll.

(i) For $\lambda_{k_2}^N < k_1$, (1.19) $\geq$ (1.18) for all $j \in N$ if and only if

$$q_j \geq \frac{\sum_{x=\lambda_j^N}^{k_1-1} \binom{N}{x} (1-p)^{N-x} p^x}{\sum_{x=\lambda_j^N}^{k_1-1} \binom{N}{x} ((1-p)^{N-x} p^x + (1-p)^x p^{N-x})} - 2c \cdot \frac{\sum_{x=0}^{k_1-1} \binom{N}{x} \left((1-p)^{N-x} p^x + (1-p)^x p^{N-x}\right)}{\sum_{x=\lambda_j^N}^{k_1-1} \binom{N}{x} ((1-p)^{N-x} p^x + (1-p)^x p^{N-x})}$$
$$\forall j \in N. \tag{1.22}$$

(ii) For $\lambda_{k_2}^N \in \left\{k_1, \ldots, N - k_1\right\}$, (1.20) $\geq$ (1.18) for all $j \in N$ if and only if

$$q_j \geq \frac{\sum_{x=\lambda_j^N}^{\lambda_{k_2}^N-1} \binom{N}{x} (1-p)^{N-x} p^x}{\sum_{x=\lambda_j^N}^{\lambda_{k_2}^N-1} \binom{N}{x} ((1-p)^{N-x} p^x + (1-p)^x p^{N-x})} - 2c \cdot \frac{\sum_{x=0}^{k_1-1} \binom{N}{x} \left((1-p)^{N-x} p^x + (1-p)^x p^{N-x}\right)}{\sum_{x=\lambda_j^N}^{\lambda_{k_2}^N-1} \binom{N}{x} ((1-p)^{N-x} p^x + (1-p)^x p^{N-x})}, \tag{1.23}$$

for $j \in N$ with $\lambda_j^N < \lambda_{k_2}^N$, and

$$q_j \leq \frac{\sum_{\lambda_{k_2}^N}^{\lambda_j^N-1} \binom{N}{x} (1-p)^{N-x} p^x}{\sum_{\lambda_{k_2}^N}^{\lambda_j^N-1} \binom{N}{x} ((1-p)^{N-x} p^x + (1-p)^x p^{N-x})} + 2c \cdot \frac{\sum_{x=N+k_1+1}^{N} \binom{N}{x} \left((1-p)^{N-x} p^x + (1-p)^x p^{N-x}\right)}{\sum_{\lambda_{k_2}^N}^{\lambda_j^N-1} \binom{N}{x} ((1-p)^{N-x} p^x + (1-p)^x p^{N-x})}, \tag{1.24}$$

for $j \in N$ with $\lambda_j^N > \lambda_{k_2}^N$.

(iii) For $\lambda_{k_2}^N > N - k_1$, (1.21) $\geq$ (1.18) for all $j \in N$ if and only if

$$q_j \leq \frac{\sum_{x=N-k_1+1}^{\lambda_j^N-1} \binom{N}{x} (1-p)^{N-x} p^x}{\sum_{x=N-k_1+1}^{\lambda_j^N-1} \binom{N}{x} ((1-p)^{N-x} p^x + (1-p)^x p^{N-x})} + 2c \cdot \frac{\sum_{x=N+k_1+1}^{N} \binom{N}{x} \left((1-p)^{N-x} p^x + (1-p)^x p^{N-x}\right)}{\sum_{x=N-k_1+1}^{\lambda_j^N-1} \binom{N}{x} ((1-p)^{N-x} p^x + (1-p)^x p^{N-x})}$$
$$\forall j \in N. \tag{1.25}$$

Note that (1.20) $\geq$ (1.18) holds for all $j \in N$ with $\lambda_j^N = \lambda_{k_2}^N$.

From here we can define the set $\widetilde{\mathscr{K}}_1^N$ formally.

Definition 1.4. $k_1 \in \widetilde{\mathscr{K}}_1^N$ *if and only if*

- $k_1 \in \mathscr{K}_1^N$, *and*
- *either* (1.22)*, or* (1.23) *and* (1.24)*, or* (1.25) *holds for* k_1.

Note, that $\widetilde{\mathscr{K}_1^N}$ is non-empty whenever the jury's heterogeneity is bounded from above sufficiently as assumed.

Step 3: Conditions for the optimal choice of k_1 *from* $\widetilde{\mathscr{K}_1^N}$.

For any $k_1 \in \widetilde{\mathscr{K}_1^N}$, (σ, μ) is an equilibrium and all jurors $j \in N$ are better off than in any equilibrium in setups with a straw poll in the first period. We now prove that the optimal $k_1 \in \widetilde{\mathscr{K}_1^N}$ is determined as in the Proposition.

1. If for all $k_1 \in \widetilde{\mathscr{K}_1^N}$ it holds that $\lambda_{k_2}^N \in \{k_1, \ldots, N - k_1\}$, it is straightforward from comparing (1.20) and (1.18), that it is optimal to set
$$k_1 \;=\; \max \widetilde{\mathscr{K}_1^N}.$$

2. If $\lambda_{k_2}^N < \frac{N}{2}$ and $\lambda_{k_2}^N < k_1$ for some $k_1 \in \widetilde{\mathscr{K}_1^N}$ then the designer faces a trade-off. By the previous argument a natural candidate is
$$k_1 \;=\; \max\left\{k_1 \;\mid\; \lambda_{k_2}^N \;\in\; \left\{k_1, \ldots, N - k_1\right\}\right\} \;\equiv\; k^m.$$
Increasing k_1 by 1 does decrease the expected costs but changes the jury's decision. For any $k_1 \geq \lambda_{k_2}^N$, an increase to $k_1 + 1$ leads to $d_1 = A$ instead of $d_1 = D$ and $d_2 = C$ for $x_1 = k_1$. Therefore, in this case each juror saves expected costs of
$$2c \cdot \binom{N}{k_1}\left((1-p)^{N-k_1} p^{k_1} + (1-p)^{k_1} p^{N-k_1}\right)$$
and, because a false judgment can be avoided if $\omega = I$, each juror additionally saves in expectation
$$q_j \cdot \binom{N}{k_1}(1-p)^{N-k_1} p^{k_1}.$$
On the other hand, if $\omega = I$ a false judgment is enacted which yields in expectation a loss of
$$(1 - q_j) \cdot \binom{N}{k_1}(1-p)^{k_1} p^{N-k_1}.$$
Expected gains are higher than expected losses from an increase of k_1 to $k_1 + 1$ for each juror if and only if
$$q_j \;\geq\; \beta(k_1, N) - 2c \quad \forall j \;\in\; N. \tag{1.26}$$

Moreover, note that it can not be the case that an increase from $k_1 + 1$ to $k_2 + 2$ is profitable to all jurors but not from k_1 to $k_1 + 1$. Suppose to the contrary that this was the case. Using (1.26) would yield the contradiction $\beta(k_1, N) - 2c \geq q_j \geq \beta(k_1 + 1, N) - 2c$.

Therefore, the following choice of $k - 1$ is optimal. Set the highest k_1 for which $\lambda^N_{k_2} \in \{k_1, \ldots, N - k_1\}$ if (1.26) is not satisfied for this value. If it is, however, set the highest k_1 for which $k_1 - 1$ does satisfy (1.26). Formally, set

$$k_1 \ = \ \max\left\{k_1 \ \in \ \widetilde{\mathcal{K}}^N_1 \ \mid \ q_1 \ \geq \ \beta(k_1 - 1, N) - 2c \vee \lambda^N_{k_2} \ \in \ \left\{k_1, \ldots, N - k_1 + 1\right\}\right\}.$$

As condition (1.26) is most binding for q_1, it suffices to consider juror 1 only.

3. If $\lambda^N_{k_2} > \frac{N}{2}$ and $\lambda^N_{k_2} > N - k_1$ for some $k_1 \in \widetilde{\mathcal{K}}^N_1$ then an analogous argumentation to the previous case applies. The difference is that an increase from $k_1 \geq N - \lambda^N_{k_2}$ to $k_1 + 1$ now changes the decision for $x_1 = N - k_1$. As a result the designer faces an adjusted trade-off for each juror. An increase from k_1 to $k_1 + 1$ yields expected gains of

$$2c \cdot \binom{N}{k_1}\left((1-p)^{N-k_1}p^{k_1} + (1-p)^{k_1}p^{N-k_1}\right) + (1 - q_j) \cdot \binom{N}{k_1}(1-p)^{k_1}p^{N-k_1},$$

whereas expected losses are

$$q_j \cdot \binom{N}{k_1}(1-p)^{N-k_1}p^{k_1}.$$

Expected gains are higher than expected loss if and only if

$$q_j \ \leq \ \beta(N - k_1, N) - 2c \quad \forall j \ \in \ N. \tag{1.27}$$

By the same argument as before, it is therefore optimal for the designer to set

$$k_1 \ = \ \max\left\{k_1 \ \in \ \widetilde{\mathcal{K}}^N_1 \ \mid \ q_N \ \leq \ \beta(N - (k_1 - 1), N) + 2c \vee \lambda^N_{k_2} \ \in \ \left\{k_1, \ldots, N - k_1 + 1\right\}\right\}$$

□

2

Target Mass or Class? Dynamic Reputation Management with Heterogeneous Consumption Externalities

2.1 Introduction

This chapter is based on joint work with Benjamin Schickner. The reputation of several goods is substantially influenced by the types of their consumers. This effect is particularly pronounced for luxury and fashion goods, such as exclusive watches and fashionable apparel. When considering buying such a good, consumers not only take into account the utility from the good itself but also whether they want to be associated with its clientele. *"Good-looking people attract other good-looking people*", as Mike Jeffries, at that time CEO of *Abercrombie & Fitch*, states in a controversial interview in 2006.[1] Economically, past buyers impose an externality on new buyers. This externality can be positive or negative depending on whether the association with the good's clientele is desirable for new buyers. Sellers seem to be aware of their consumers' concerns and strategically try to target a certain clientele. In the same interview, Mike Jeffries emphasizes: *"(...) we want to market to cool, good-looking people. We don't market to anyone other than that.*" He goes on to say: *"Are we exclusionary? Absolutely. Those companies that are in trouble are trying to target everybody: young, old, fat, skinny. But then you become totally vanilla. You don't alienate anybody, but you don't excite anybody, either.*" These quotes exemplify how sellers manage their reputation by targeting a specific clientele, in Jeffries' words "good-looking people", in order to attract new consumers who wish to be associated with this clientele by possessing the same good. In turn, new consumers influence the good's reputation. How a seller manages reputation depends on the characteristics of the market and the good and may differ substantially across markets and goods.

In this paper, we present a dynamic model to study how a monopolist optimally manages

[1] See Denizet-Lewis (2006).

her reputation by targeting specific clienteles in order to maximize profits. In particular, we are interested in how the underlying characteristics of the market affect the evolution of reputation, demand, and prices over time.

The reputation of a seller has several facets. Traditionally, a good's reputation is linked to its intrinsic properties such as quality. As these properties are usually private information of the seller, we refer to this facet of reputation as *private reputation*. Goods with high private reputation are believed to be of better quality, for example, than goods with low private reputation. However, reputation also has a public facet. Often, buyers do not only derive utility from the good's intrinsic properties but also from what the good symbolizes to others. In itself this *public reputation* has many aspects. For example, a seller forms her public reputation through advertisement. In this paper, we study the effect of the seller's clientele on her public reputation.

For most goods both facets of reputation are present. However, the importance of one or the other may vary across goods, time, and cultures. In this paper, we focus solely on the public reputation a seller derives from her clientele. Whenever referring to reputation in the following, we mean reputation in this sense unless explicitly stated otherwise.

In our model, a profit-maximizing monopolist repeatedly offers a good to a continuum of heterogeneous, short-lived buyers. Buyers are ordered according to their type which describes their effect on the seller's reputation. Selling to buyers of higher types improves the seller's reputation, and selling to buyers of lower types decreases the seller's reputation. A buyer's willingness to pay increases in the seller's reputation and in their type.[2] Therefore, the seller can choose her clientele through her pricing strategy and thereby manage her reputation. For any price, buyers purchase the good only if their type exceeds a cutoff type.

In the first part of the paper, we consider a general model in which types are drawn from a general, continuous distribution and a buyer's utility is quasilinear. Reputation tomorrow is a function of reputation today and today's cutoff type, which characterizes the current clientele. We impose, first, that reputation is persistent, i.e., ceteris paribus a higher reputation today yields a higher reputation tomorrow, and, second, that reputation is increasing in today's clientele, i.e., selling to an exclusive group of high-type buyers increases reputation, whereas selling to a broader group of buyers with heterogeneous types decreases reputation.

Within this general framework, we start by establishing existence of a Markov perfect equilibrium and characterize the seller's value in any equilibrium. We then argue that the seller's value is increasing in reputation which implies that reputation is beneficial for the seller. In each period, the seller solves an intertemporal trade-off. On the one hand, selling to a small,

[2] Intuitively, the higher a buyer's type, the more reputation improves after he purchases the good. The buyer anticipates his influence and, hence, his willingness to pay increases in his type.

exclusive clientele increases her reputation and, hence, her future profits. On the other hand, limiting demand in this way decreases her profits today.
Conventionally, high prices increase the seller's profit by increasing revenue. In our model, in addition, the seller sets high prices to target an exclusive clientele in order to improve her reputation. In other words, high prices prevent buyers with lower types from purchasing the good and, hence, protect the seller's reputation.

To obtain explicit results regarding the reputation dynamics, we specialize our setup in the second part of the paper. We assume that a buyer's utility is linear and that the distribution of types is uniform. Reputation tomorrow is a convex combination of today's reputation and the cutoff type. We associate the weight of the convex combination on current-period reputation with the good's durability. Intuitively, if the durability of the good is higher, buyers possess the good for a longer period of time. Consequently, the good is longer associated with their type, and the influence of past buyers' types, captured by today's reputation, on tomorrow's reputation is comparatively high. Conversely, if the durability of the good is lower, new buyers constitute a significant fraction of the seller's clientele. Thus, the influence of today's buyers' types, characterized by today's cutoff type, on tomorrow's reputation is relatively high. An example for a market that is characterized by comparatively high durability is the market for watches, whereas, for example, the market for fashionable apparel is characterized by comparatively low durability. With these adjustments, we obtain a linear-quadratic setup. We determine the seller's value and policy function in closed form which makes the setup tractable for a more explicit analysis.

Next, we study optimal reputation dynamics, in particular, their dependence on the good's durability. First, we show that reputation always converges to a long-run reputation. Although it is optimal in the short run for the seller to target different clienteles, this result implies that it is optimal in the long run to target a fixed clientele and maintain a constant reputation. In contrast to the private reputation literature, e.g. Holmström (1999), Cripps et al. (2004), and Cripps et al. (2007), reputation is not a short run phenomenon. Even in the long run, the seller trades off the benefit of increasing her reputation against realizing higher current-period profits. The long-run reputation is increasing in the discount factor and decreasing in the good's durability.
Second, convergence behavior towards the long-run reputation is substantially different for goods with different durability. If the durability of the good is below a threshold, reputation oscillates towards the long-run reputation. A period of high reputation is followed by reputation of low reputation and vice versa. If the durability is above the threshold, reputation and price dynamics are monotone. If the initial reputation is high, reputation decreases monotonically to the long-run reputation, and, in contrast, increases monotonically to the long-run reputation if reputation is initially low.

Despite the substantially different convergence behavior of the cases described above, we identify an underlying monotonicity in the degree of fluctuations across these cases. To this end, we determine an appropriate measure for fluctuations, the normalized distance between two subsequent reputation levels, and show that fluctuations are monotonically decreasing in durability, i.e., the higher the durability of the good the less reputation fluctuates over time.

Our model predicts substantial fluctuations in reputation and prices for goods with low durability and relatively stable, monotone reputation dynamics for goods with high durability. There are many factors that drive price and reputation dynamics. The relative importance of these factors may vary significantly across markets. Nevertheless, our findings seem to be in line with anecdotal evidence. As an example for a durable good, consider the Swiss watchmaker *Rolex*. *Business Insider* documents how prices of *Rolex* watches have steadily increased over the last sixty years, both in absolute terms and as measured as a proportion of average income. At the same time its reputation seems to have improved constantly: *"(...) today's Submariner, the tool-watch of yesterday, has transformed into an internationally recognized status symbol (...)."*[3] As an example for a good with low durability, it is insightful to come back to *Abercrombie & Fitch*. During its reputational high before and at the time of Jeffries' interview, customers were queuing in front its stores. Simultaneously, *Abercrombie & Fitch* was expanding considerably. The German newspaper *FAZ* notes in an article that the brand has "lost its coolness" since then, which some analysts attribute to the fact that it has become too widespread. Further, the authors observe that *Abercrombie & Fitch* is currently cutting back its network of stores.[4] In a similar vein, consider the rapid rise and decline of the fashion label *Ed Hardy*. At its reputational height, many celebrities wore *Ed Hardy* clothes. According to a *CNN* article, the designer himself attributes the subsequent fall of the brand to the fact that " *(...) widespread licensing aspired to make the brand more accessible to people at every price point.*" An analyst observes that, as a result, *Ed Hardy* became *"very trailer park*" and states that *"they made it too unexclusive."*[5] Similar upward and downward fluctuations in reputation due to an expansion and contraction of clientele are documented for the fashion labels *Burberry* and *Louis Vuitton*.[6]

Finally, we argue that public reputation provides another rationale for planned obsolescence. Intuitively, if a seller starts with a low reputation, she can improve her reputation by including higher types into her clientele more quickly if the durability of the good is low. This reputational explanation of planned obsolescence complements the traditional

---

[3] See Bredan (2015).
[4] See Lindner and Löhr (2015).
[5] See Alabi (2013).
[6] See also Lindner and Löhr (2015).

demand-driven explanation.

The rest of the paper is organized as follows. After discussing related literature in the next section, we introduce the model in Section 2.2. In Section 2.3, we derive the existence result and the trade-off between building reputation and current-period profits in the general model. In Section 2.4, we proceed by studying reputation dynamics in the specified model. In Section 2.5, we discuss the implications on planned obsolescence. Section 2.6 concludes.

2.1.1 Related Literature

In the private reputation literature, the seller's reputation represents the market's belief about her unknown type, productivity, or quality. In Kreps et al. (1982), Mailath and Samuelson (2001), and Cripps et al. (2004), reputation reflects the market's belief that the seller is a competent type who strategically chooses her effort level instead of a behavioral type who always exerts the same effort. In Holmström (1999), Tadelis (1999), Board and Meyer-ter Vehn (2013), and Dilmé (2016), reputation is the buyers' belief about the seller's productivity or quality which influences their utility from purchasing the seller's good. Whereas in these models buyers update their beliefs based on openly observable information, it is costly to acquire information about the seller's quality in Liu (2011) as well as in Lee and Liu (2013). Overall, the models in the private reputation literature are dynamic and, typically, a monopolistic, long-living seller faces a sequence of myopic, short-lived buyers. For a more detailed overview, see, for example, Bar-Isaac and Tadelis (2008).
To illustrate the difference to the public reputation studied in our paper, consider the classical example of a restaurant from the private reputation literature. The restaurant's private reputation describes the customers' belief that food and service are of high quality. Public reputation describes the clientele a customer is associated with when visiting the restaurant. This could be, for example, students or politicians and business men.

The literature on consumption externalities distinguishes the *bandwagon effect*, the *snob effect*, and the *Veblen effect* (Veblen 1899, Leibenstein 1950). The bandwagon effect and the snob effect describe the case when demand increases or decreases, respectively, if others consume the same good. The Veblen effect refers to conspicuous consumption, that is, when demand for a good increases in its price. In our model, there are bandwagon effects with respect to high-type buyers and snob effects with respect to low-type buyers. Becker (1991) and Becker and Murphy (1993) consider consumers who care about who else possesses a good and find snob and bandwagon effects. Bagwell and Bernheim (1996) as well as Amaldoss and Jain (2005a) show that Veblen effects can arise when consumers signal wealth from the consumption of conspicuous goods or when consumers have social needs, such as desire for prestige. Grilo et al. (2001) and Amaldoss and Jain (2005b) examine optimal pricing of conspicuous goods in static duopoly competition.

In our paper, we combine the dynamic approach of the private reputation literature with consumption externalities. Pesendorfer (1995) and Hashimoto and Matsubayashi (2014) also consider the problem of a monopolistic seller in a dynamic model with consumption externalities. Pesendorfer (1995) examines fashion cycles in a model where consumers of different types buy a good to distinguish themselves from each other. The good's reputation stems from its ability to signal a buyer's type in a secondary marriage market. Over time, more consumers buy the good such that the signaling effect and, hence, reputation and prices decrease monotonically. When the price has dropped low enough, the seller introduces a new, initially exclusive product line. As opposed to the focus on optimal replacement of a product in Pesendorfer (1995), we examine the effects of changes in the good's characteristics on the seller's reputation management. Moreover, considering the reputation of a good, we obtain richer dynamics. In our model, reputation is not always monotonically decreasing. Depending on initial reputation and the good's durability, reputation can monotonically increase, monotonically decrease, or oscillate. Hashimoto and Matsubayashi (2014) analyze optimal pricing of a monopolistic seller in a dynamic model with either positive or negative consumption externalities. Consumers' utilities depend on past sales which the seller anticipates when solving her pricing problem. For the case where consumption externalities are negative, their results are in line with ours. If the consumers' utilities only depend on recent sales, reputation oscillates downwards, whereas reputation decreases monotonically if consumers discount past sales less strongly. In contrast to their paper, however, in our model some consumers exert positive consumption externalities whereas others exert negative externalities. Furthermore, we provide a full characterization of the seller's optimal pricing strategy, varying the characteristics of the market. Moreover, we model the reputation channel more explicitly through which buyers influence others' utilities from buying the good.

In contrast to social consumption externalities as studied in our and the above papers, another branch of literature studies technological consumption externalities. These arise when the utility of a good increases with the size of its user base, such as telephones or social networks. Katz and Shapiro (1985) as well as Lee and Mason (2001) analyze price competition between firms under positive or negative network effects, and Dhebar and Oren (1985), Bensaid and Lesne (1996), and Gabszewicz and Garcia (2008) consider dynamic pricing strategies of a monopolist.

Further related is Rayo (2013) who studies a monopolist selling a good that signals social status in a screening model as well as the literature on peer groups (see, for example, Board 2009), and the literature on scarcity (see, for example, Stock and Balachander 2005).

2.2 Model

Time is discrete and infinite, $t \in \{0, \ldots\}$. A single long-lived seller (she) repeatedly offers a good to a unit mass of short-lived buyers (he).

Seller. In each period t, the long-lived seller sets a price p_t for the good. She has no production costs and discounts future payoffs with discount factor $\delta \in (0,1)$.

Buyers. Every period-t buyer is characterized by his type $\theta \in [0,1]$ which is distributed according to a continuous distribution function $F(\theta)$ with support $[0,1]$. The types of buyers who purchase the good in periods $0, 1, \ldots, t-1$ determine the *reputation* of the good at time t which we denote by $\lambda_t \in [0,1]$. A period-t buyer with type θ who purchases the good with reputation λ_t at price p_t receives utility

$$u(\theta, \lambda_t) - p_t,$$

where $u : [0,1]^2 \to \mathbb{R}_{\geq 0}$ is strictly increasing in both arguments and continuous. We do not explicitly model buyers as players in the game. Notice that there exists a cutoff type $\theta^\dagger(\lambda_t, p_t) \in [0,1]$ for every reputation λ_t and every price p_t such that $u(\theta^\dagger(\lambda_t, p_t), \lambda_t) = p_t$.[7] We assume that all buyers with type $\theta \geq \theta^\dagger(\lambda_t, p_t)$ purchase the good and all buyers with type $\theta < \theta^\dagger(\lambda_t, p_t)$ do not purchase the good. We justify this assumption with the following reasoning. Assume that each buyer decides whether to buy the good, and normalize buyers' utility from not buying the good to zero. Consequently, a period-t buyer with type θ purchases the good only if

$$u(\theta, \lambda_t) - p_t \;\geq\; 0. \tag{2.1}$$

Recall that u is strictly increasing in the first argument. If it is optimal for a buyer with type θ to buy the good then it is optimal for any buyer with type $\theta' \geq \theta$ to buy the good as well. Analogously, if it is optimal for a buyer with type θ not to buy the good then it is also optimal for any buyer with type $\theta' \leq \theta$ not to buy the good. The cutoff is given by the type for which (2.1) holds with equality.[8] Thus, the seller's demand in period t is $1 - F(\theta^\dagger(\lambda_t, p_t))$. We refer to the set of buyers who purchase the good as the seller's *clientele*.

Reputation Transition. Initially, the seller's reputation is $\lambda_0 \in [0,1]$. It evolves dynamically depending on the types of buyers who purchase the good. We assume that reputation satisfies the following two properties. First, reputation is persistent in the sense that, ceteris paribus,

---

7 For $p_t < u(0, \lambda_t)$ we set $\theta^\dagger(\lambda_t, p_t) = 0$, and for $p_t > u(1, \lambda_t)$ we set $\theta^\dagger(\lambda_t, p_t) = 1$.

8 In light of this reasoning, our assumption to not model buyers as players is mainly for notational convenience. This is a common assumption in the private reputation literature, see, for example, Board and Meyer-ter Vehn (2013). Essentially, we only assume that the buyer with the cutoff type purchases the good.

a higher reputation today yields a higher reputation tomorrow. Second, the more exclusive the seller's clientele today, the higher her reputation, that is, reputation increases in today's cutoff type. Formally, we assume that reputation evolves according to

$$\lambda_{t+1} = \phi(\lambda_t, \theta^\dagger(\lambda_t, p_t)), \tag{2.2}$$

where $\phi : [0,1]^2 \to [0,1]$ is strictly increasing and continuous in both arguments.[9]

Information and Timing. At time t, the seller knows the set of agents who bought the good in periods $0, 1, \ldots, t-1$ as well as the corresponding prices and reputation levels. The timing of the game in each period t is as follows. First, the seller sets a price p_t. Then, period-t buyers arrive, buy the good or not, and leave the market. Last, reputation updates according to (2.2).

Histories and Payoffs. Let $p^t = (p_0, \ldots, p_{t-1})$ be the history of prices up to time t. For any t and initial reputation λ_0, p^t determines a history of reputation levels $\lambda^t(p^t) = (\lambda_0, \lambda(p^1), \ldots \lambda(p^t))$ through (2.2). The seller's history at the start of period t is given by $h^t = (p^t, \lambda^t(p^t))$ and $h^0 = \lambda_0$. Further, denote the set of all possible histories at the start of period t by $\mathscr{H}^t$.

Fix any history h^t, the seller's continuation payoff from a sequence of prices $(p_s)_{s=t}^\infty$ is

$$\sum_{s=t}^{\infty} \delta^{s-t} p_s (1 - F(\theta^\dagger(\lambda(p^s), p_s)). \tag{2.3}$$

Strategies and Equilibrium. A (behavioral) pure strategy of the seller is a collection of functions $\rho = (\rho_t)_{t=0}^\infty$, where

$$\begin{aligned} \rho_t \ : \quad \mathscr{H}^t &\longrightarrow \mathbb{R}_{\geq 0}, \\ h^t &\longmapsto p_t. \end{aligned}$$

The strategy is Markovian if ρ_t is a function of λ_t only, for all t.

Definition 2.1. *A strategy $\rho^* = (\rho_t^*)_{t=0}^\infty$ of the seller constitutes*

(i) a Nash equilibrium (NE) if it maximizes (2.3) *at h^0.*

(ii) a subgame perfect Nash equilibrium (SPNE) if it maximizes (2.3) *at any h^t, for any t.*

(iii) a Markov perfect equilibrium (MPE) if it is Markovian and a subgame perfect equilibrium.

---

9 We consider a deterministic reputation transition in order to obtain a clean comparison of reputation dynamics.

2.3 Equilibrium Existence and Value of Reputation

We start by establishing existence of an equilibrium.

Proposition 2.1. *There exists a MPE. In any NE the seller's value is $V(\lambda_0)$, where $V(\lambda)$ is the unique solution to*

$$V(\lambda) = \sup_{p \in \mathbb{R}_{\geq 0}} \left\{ p(1 - F(\theta^{\dagger}(\lambda, p)) + \delta V\left(\phi(\lambda, \theta^{\dagger}(\lambda, p))\right)\right\}. \tag{2.4}$$

This and all subsequent proofs are relegated to the Appendix. The seller's problem is to choose an infinite sequence of prices in order to maximize her discounted sum of profits. A strategy of the seller which solves this problem is a Nash equilibrium. In Proposition 2.1, we show that the value of the seller is characterized by Bellman equation (2.4). A policy function, corresponding to a solution of the Bellman equation, induces a strategy that is Markovian. Thus, solving the seller's problem through (2.4) yields a Markov perfect equilibrium. We establish existence of a value function which solves (2.4) and of a corresponding policy function. As any Markov perfect equilibrium is also a Nash equilibrium, Proposition 2.1 particularly implies existence of a Nash equilibrium. The proof draws on classical results from the literature on dynamic programming.
We proceed by characterizing the seller's value function in more detail.

Proposition 2.2. *Reputation is valuable for the seller, that is, $V(\lambda)$ is increasing.*

This result shows that a higher reputation entails a higher value and is thus better for the seller. Specifically, a higher reputation benefits the seller in two ways. First, the buyers' willingness to pay increases in reputation because their utility function is increasing in reputation. Therefore, a seller with a higher reputation can sell to the same group of buyers at a higher price. Second, as reputation is persistent, a higher reputation today implies a higher reputation tomorrow, for a fixed group of buyers. As a result, the first effect carries over to subsequent periods. The proof of Proposition 2.2 exploits these two effects using a mimicking argument. A high-reputation seller can always induce the same proportion of buyers to purchase the good as a low-reputation seller at higher prices. Consequently, the high-reputation seller can earn higher profits.[10]

Moreover, Proposition 2.2 reveals that there is an intertemporal trade-off between building a reputation for the future and realizing profits in the current period. On the one hand, the seller wants to set a price to increase her next-period reputation in order to get a higher continuation value. Considering only this effect, it is optimal for the seller to set a high price such that only a small exclusive group of high-type buyers purchases the good. This drives

[10] The fact that the value function is increasing does not follow from classical sufficient conditions in the literature because the per-period payoff function is not necessarily increasing in reputation.

up the seller's next-period reputation. On the other hand, the seller wants to set a price that maximizes current-period profits. If the seller sets a high price only a small number of buyers purchase the good. In order to maximize current-period profits, it is better for the seller to set an intermediate price, the static monopoly price, at which a larger set of buyers purchase the good. In turn, however, a larger set of buyers contains a more diverse pool of types. Therefore, setting the static monopoly price comes at the cost of driving down the seller's next-period reputation compared to the case where the seller targets a small, exclusive group of high-type buyers. To maximize overall profits, the seller has to balance these two countervailing effects. Consequently, the seller's optimal price today is higher than the static monopoly price to protect her reputation tomorrow.

In other words, the seller trades off targeting mass, that is a large group of buyers with diverse types, and targeting class, that is a small but exclusive group of buyers with high types.

2.4 Reputation Dynamics

So far, we have shown existence of an equilibrium and that reputation is valuable for the seller in a general model. In this section, we are interested in reputation dynamics. In particular, we study the reputation's long-run behavior and its dynamic evolution over time. To characterize the dynamic properties of the model, we need to impose additional structure which allows us to obtain closed-form solutions for the seller's value and policy function.

Model Adjustments. In the following, we consider the case where types are uniformly distributed, and a buyer's utility from buying the good is linear in λ_t and θ, specifically

$$u(\theta, \lambda_t) = \theta + \lambda_t. \tag{2.5}$$

This implies that for a given cutoff $\theta_t^\dagger$ the good's price is $\lambda_t + \theta_t^\dagger$, and demand is $1 - \theta_t^\dagger$. Further, we assume the following reputation transition,

$$\lambda_{t+1} = \alpha\lambda_t + (1-\alpha)\theta_t^\dagger, \tag{2.6}$$

where $\alpha \in [0, 1]$. Reputation evolves according to a convex combination of today's reputation and cutoff type. The latter serves as a proxy for the clientele of the current-period.

We interpret α as a measure for the good's physical durability. A good that is more durable is used longer and, hence, also associated longer with its buyers' types. A high value of α corresponds to a good with high durability whose buyers' types are identified with it for a relatively long time. Therefore, new buyers only account for a small fraction of goods in circulation and reputation depends less on their types. In turn, however, their types impact the good's reputation for a long time. Examples for such goods are expensive watches. In contrast, a low value of α represents a good with low durability. It is in circulation for a relatively short period of time and, hence, reputation depends more on the current period's clientele. Fashionable apparel is an example for such a good.

Intuitively, in our model, we keep the size of the group of buyers who determine the good's reputation constant. One could think of a reference group of buyers whose types are associated with the good. This keeps the influence of each buyer on reputation constant over time and simplifies the analysis. Keeping this in mind, we can motivate the specific form of (2.6) more formally if we imagine that the good becomes unusable with probability $1-\alpha$. The types of buyers in the reference who do not use the good any longer are replaced by those of current-period buyers. Therefore, the fraction $1-\alpha$ of the good's reputation is determined by current-period buyers' types, represented by the proxy $\theta^{\dagger}$, and the fraction α by past-period buyers' types. See also the discussion following Proposition 2.3.

Value and Policy Function. In this framework, we derive closed-form solutions for the seller's value function and the seller's policy function. We start by establishing a preliminary result.

Lemma 2.1. *The value function $V(\lambda)$ is continuous in α, δ and λ.*

This result equips us with the properties we need to characterize the value function for the uniform-linear model.

So far, the seller's strategy determines a price for any reputation level. There is a one-to-one relation between the cutoff type and the price. Hence, instead of setting a price, it is convenient to think of the seller's strategy as choosing cutoff types in the following.

Proposition 2.3.

(i) The value function $V(\lambda)$ is quadratic and continuously differentiable.

(ii) The corresponding unique policy function $\theta^*(\lambda)$ is of the form

$$\theta_1\lambda + \theta_0, \tag{2.7}$$

for some θ_0 and $\theta_1 \in (-1, 0)$.

This result can be seen from two perspectives. Technically, we derive a closed-form solution for the uniform-linear model which turns out to be linear-quadratic, i.e., the policy function is linear, and the value function is quadratic. Economically, we show that the policy function is decreasing in reputation. In other words a seller with high reputation chooses a lower cutoff type and, thus, sells to a larger group of buyers than a seller with low reputation. To gain some intuition for this result, it is instructive to consider the seller's current-period profits. The effect of $\theta^{\dagger}$ on current-period profits is twofold. On the one hand, a higher $\theta^{\dagger}$ corresponds to a higher price. On the other hand, a higher $\theta^{\dagger}$ corresponds to selling to a smaller group of high-type buyers, that is, to lower demand. Specifically, the price is linearly increasing in $\theta^\dagger$, whereas demand is linearly decreasing in $\theta^\dagger$. At the interior optimum, the seller balances these two opposing effects. Ceteris paribus, a higher reputation increases the price one-to-one but has no effect on demand. Taking the two opposing effects into account, it is profitable for the seller to trade off a fraction of the price increase for an increase in demand by choosing a lower $\theta^\dagger$. This intuition remains valid even when accounting for the effect on the continuation value. The continuation value is increasing in $\theta^\dagger$ because tomorrow's reputation is higher if the seller sells exclusively to a small group of high-type buyers. This strengthens the positive effect that an increase in today's reputation has on today's price through $\theta^\dagger$. Nevertheless, the optimum remains interior and the seller continues to trade off the two opposing effects.

Corollary 2.1. *There exists a unique SPNE.*

Proposition 2.1 yields existence of a Markov perfect equilibrium and hence of a Nash equilibrium. By Proposition 2.3 there exists a unique policy function and therefore a unique Markov perfect equilibrium. Because we only model the seller as a player, the subgame perfect Nash equilibrium is also unique.

In our representation of the reputation transition (2.6), reputation in the next period depends on the cutoff buyer $\theta^\dagger$ as a summary statistic for demand in the current period. Note that we retain the linear-quadratic structure for the seller's value and policy function if we consider a more general reputation transition where the next period's reputation depends on a linear function of demand or, equivalently, the cutoff buyer. As types are uniformly distributed, this includes, for example, the conditional expectation of the types of buyers who purchase the good. Thus, our qualitative results carry over to these alternative specifications of the reputation transition that depend more generally on demand. For computational convenience, however, we stick to specification (2.6) in the following.

Long-run Reputation. With the seller's equilibrium strategy at hand which is implicitly determined by her policy function, we analyze in detail how she dynamically manages her reputation over time. We start by showing that reputation converges to some $\hat{\lambda} \in (0,1)$ in the long run. We refer to $\hat{\lambda}$ also as *long-run reputation* in the following.

Proposition 2.4. *As $t \to \infty$, reputation converges to a unique $\hat{\lambda}$, for any $\lambda_0 \in [0,1]$.*

(i) For $\alpha < 1$,

$$\hat{\lambda} = \frac{1 + \frac{\delta(1-\alpha)}{1-\delta\alpha}}{3 + \frac{\delta(1-\alpha)}{1-\delta\alpha}}, \tag{2.8}$$

which is increasing in δ and decreasing in α.

(ii) For $\alpha = 1$, $\hat{\lambda} = \lambda_0$.

For each reputation, the seller trades off the value of reputation tomorrow against demand and, hence, profits today. In order to increase reputation, she has to sell to a more exclusive clientele which decreases her profits today. Targeting a larger group of buyers today decreases her reputation and lowers profits from future periods. From Proposition 2.4, we learn that there exists a reputation level, $\hat{\lambda}$, where these effects are in balance and the seller does not want to change her reputation. Whereas it is optimal for the seller to dynamically adjust her reputation and clientele in the short run, the seller optimally targets a fixed clientele and has a constant reputation in the long run.

Note that, in contrast to the private reputation literature, reputation is not a short-run phenomenon. Even in the long run, the seller takes into account that her current clientele determines her future reputation.

As second insight, we provide comparative statics results for the long-run reputation $\hat{\lambda}$ with respect to the seller's discount factor δ as well as the good's durability α. Intuitively, as δ increases the seller becomes more patient and values future periods more. In each period, she trades part of her current-period profits off against reputation tomorrow, recall the discussion following Proposition 2.2. When future periods become more valuable to the seller, she is willing to invest more of her profits today into higher reputation in subsequent periods in order to increase profits from future periods. As a result, the long-run reputation increases.

Furthermore, Proposition 2.4 reveals that it is optimal for the seller to decrease her reputation over time if she has a high reputation initially. On average, the seller realizes higher profits in early periods at the cost of driving down her reputation over time. Conversely, if the seller starts with a low reputation, it is optimal for her to increase reputation over time at the cost of lower profits in early periods. In the long run, however, reputation is stable. Intuitively, above $\hat{\lambda}$, the benefits of higher profits in early periods outweigh the costs of decreasing reputation in the long run; below $\hat{\lambda}$ the benefits of higher reputation in the long run outweigh the costs of lower profits in early periods. At the long-run reputation, the trade-off between realizing current-period profits and changing reputation is balanced.

As α increases, the good becomes more durable and reputation more persistent. From the seller's perspective, she has to give up more demand to increase reputation and loses less reputation when increasing demand. Therefore, the costs of increasing current-period profits go down, whereas the benefits of higher current-period profits are unaffected by the change in α. Consequently, the level where these costs and benefits are in balance is lower, that is, long-run reputation $\hat{\lambda}$ decreases in α.

Although it seems that the conditions imposed on the general model, outlined in Section 2.2, are not sufficient to guarantee existence of a long-run reputation, we can establish existence in frameworks that are more general than our linear-quadratic setup. See, for example, Scheinkman (1976), Araujo and Scheinkman (1977), McKenzie (1986), Stokey et al. (1989) for sufficient conditions.

In the following, we say that the seller has *good reputation* in period t if $\lambda_t > \hat{\lambda}$ and *bad reputation* in period t if $\lambda_t < \hat{\lambda}$. Further, we say that the seller *milks* reputation in period t if she targets a clientele which decreases her next-period reputation, that is $\theta_t^{\dagger} < \lambda_t$. Conversely, the seller *builds* reputation if she targets a clientele which increases her next-period reputation, that is $\theta_t^{\dagger} > \lambda_t$.

In the next step, we are interested how the underlying characteristics of the market affect the convergence dynamics towards the long-run reputation. The market is characterized by the good's durability α. Therefore, we analyze the effect of the good's durability on how reputation converges. The following definition turns out to be useful.

Definition 2.2. *A sequence of reputations* $\{\lambda_t\}_{t=0}^{\infty}$

(i) oscillates towards long-run reputation $\hat{\lambda}$ *if* $\lambda_t \leq (\geq)\hat{\lambda}$ *implies that* $\lambda_{t+1} \geq (\leq)\hat{\lambda}$*, for all* t*, and* $|\lambda_t - \hat{\lambda}| \downarrow 0$.

(ii) stagnates at long-run reputation $\hat{\lambda}$ *if* $\lambda_t = \hat{\lambda}$*, for all* $t \geq 1$.

(iii) converges smoothly to long-run reputation $\hat{\lambda}$ *if* $\lambda_t \leq \lambda_{t+1} \leq \hat{\lambda}$*, for all* t*, or* $\lambda_t \geq \lambda_{t+1} \geq \hat{\lambda}$*, for all* t*, and* $|\lambda_t - \hat{\lambda}| \downarrow 0$.

Intuitively, if reputation oscillates towards $\hat{\lambda}$, a period of good reputation is followed by a period of bad reputation and vice versa. The magnitude of the fluctuation between good and bad reputation decreases over time. If reputation stagnates at $\hat{\lambda}$, it jumps to the long-run reputation $\hat{\lambda}$ in the first period and remains constant from then on. There are two scenarios in which reputation converges smoothly to $\hat{\lambda}$. First, reputation is good in every period and deteriorates over time to $\hat{\lambda}$, i.e., reputation converges to $\hat{\lambda}$ monotonically from above. Second, reputation is bad in all periods and increases over time to $\hat{\lambda}$, i.e., reputation converges to $\hat{\lambda}$ monotonically from below.
For a better understanding of the dynamics, we study two benchmark cases. In the first case, the good is everlasting, i.e. $\alpha = 1$, and in the second case the good is most ephemeral, i.e. $\alpha = 0$.

Benchmark: Everlasting Good ($\alpha = 1$). In the extreme case when the good is everlasting, the good is always associated with the initial buyers' types. Current buyers do not affect reputation. Intuitively, as the good is everlasting, the reference group of buyers who determine the seller's reputation is unchanged over time. Consequently, reputation remains at its initial value λ_0 in each period. From the seller's perspective, the dynamic trade-off disappears. Her reputation tomorrow is not influenced by her clientele today and, hence, she solves the same static problem in each period.

Lemma 2.2. *Suppose $\alpha = 1$. For any t, it is optimal for the seller to set*

$$\theta_t^{\dagger} = \frac{1-\lambda_0}{2},$$

and her profits are

$$\frac{1}{1-\delta}\left(\frac{1+\lambda_0}{2}\right)^2.$$

As reputation remains constant at λ_0, this is an extreme case of smooth convergence to $\hat{\lambda}$ in the sense of Definition 2.2.

Benchmark: Most Ephemeral Good ($\alpha = 0$). Contrary to the previous case, the good is most ephemeral such that buyers use their good for one period only. Therefore, the group of buyers who possess the good consists exclusively of current-period buyers, and, consequently, the good is only associated with their types. The seller's reputation tomorrow is independent of her reputation today. Hence, her choice of today's price not only determines her current-period profits but also fully determines her reputation tomorrow. Consequently, in this benchmark, reputation reacts most sensitively to adjustments in the seller's clientele.

Lemma 2.3. *Suppose $\alpha = 0$. The seller's policy function is*

$$\theta^*(\lambda) = -\frac{\lambda}{1+\sqrt{1-\delta}} + \frac{1+\delta}{1+\delta+\sqrt{1-\delta}},$$

and, as $t \to \infty$, her reputation oscillates towards

$$\hat{\lambda} = \frac{1+\delta}{3+\delta}.$$

See Figure 2.1 for a graphical illustration of the resulting dynamics of reputation, prices, and demand. We observe that reputation alternates between periods of high reputation and periods of low reputation if the good is most ephemeral. Also, the corresponding sequences of prices and demand fluctuate substantially. Over time, the amplitude of the fluctuations decreases and reputation oscillates towards the long-run reputation.

Intermediate Durability: $\alpha \in (0,1)$. For goods of intermediate durability, the set of buyers who possess the good, and whose types are thus associated with the good, consists of current-period buyers ($\alpha < 1$) and buyers from previous periods ($\alpha > 0$). Therefore, both influence the seller's next-period reputation. In this case, we obtain the following result.

Proposition 2.5. *There exists a cutoff $\bar{\alpha} \in (0,1)$ such that reputation*

(i) *converges smoothly to $\hat{\lambda}$ if $\alpha > \bar{\alpha}$.*

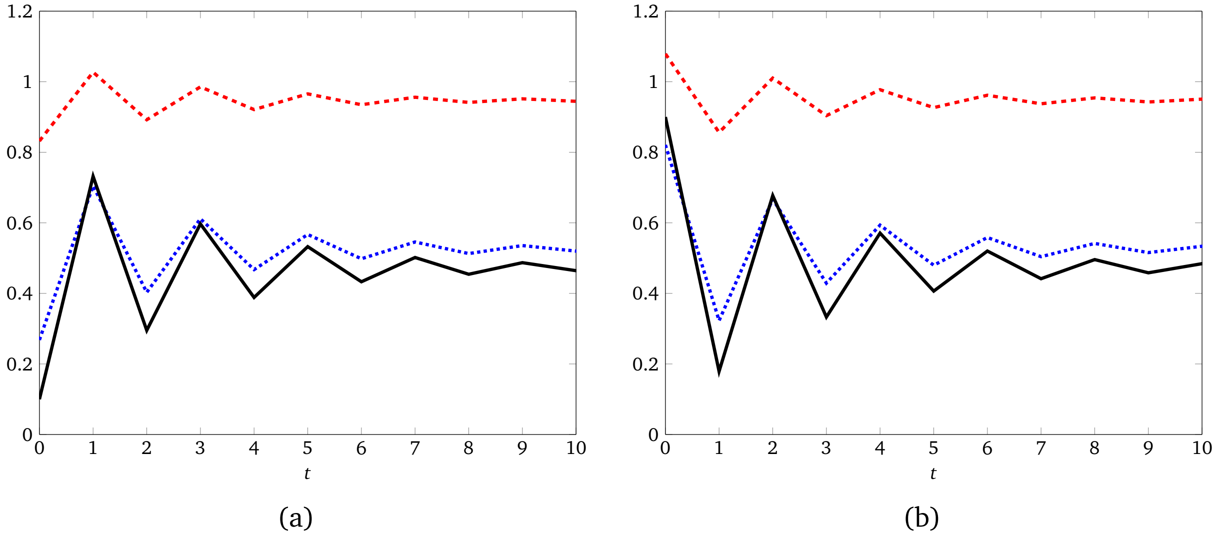

Figure 2.1: Reputation (solid black line), prices (dashed red line), and demand (dotted blue line) for $\alpha = 0$ and $\delta = 0.8$ in $t \in \{0, \dots, 10\}$ with $\lambda_0 = 0.1$ (a), and $\lambda_0 = 0.9$ (b).

(ii) stagnates at $\hat{\lambda}$ if $\alpha = \bar{\alpha}$.

(iii) oscillates towards $\hat{\lambda}$ if $\alpha < \bar{\alpha}$.

Proposition 2.5 describes how the seller manages her reputation over time, depending on the durability of the good. From Proposition 2.4 we know that the seller's reputation converges to $\hat{\lambda}$, starting from any initial reputation λ_0. The seller manages her reputation by setting prices and thereby targeting a particular clientele which in turn influences her reputation in subsequent periods. Proposition 2.5 shows that, depending on the durability of the good, there are three ways in which the seller's reputation converges to the long-run reputation. These differ substantially in the convergence behavior of reputation and prices.

First, consider the case where the good is relatively durable, $\alpha > \bar{\alpha}$, and the seller's initial reputation is good, $\lambda_0 > \hat{\lambda}$. Initially, the seller sells to a large group of buyers with a diverse set of types by setting a price which is high in absolute terms but comparatively low given her high reputation. As a result, reputation decreases. Because durability is high, it takes time until the low-type buyers constitute a significant proportion of the set of buyers who possess the good and affect the reputation negatively. Therefore, the decline in reputation is relatively small. As time passes, the seller gradually targets less buyers overall, however, she continues to sell to a large group of buyers. Consequently, reputation declines at a decreasing rate. The effect on prices is twofold. On the one hand decreasing demand corresponds to increasing prices. On the other hand, as reputation decreases, the seller has to lower prices because the buyers' willingness to pay decreases. Overall, we observe that prices decrease smoothly over time which corresponds to demand decreasing slower than reputation. See Figure 2.2b for a graphical representation. Intuitively, if the good's durability is high and the initial reputation of the seller is good, the seller gradually milks her reputation in every period. That is, she

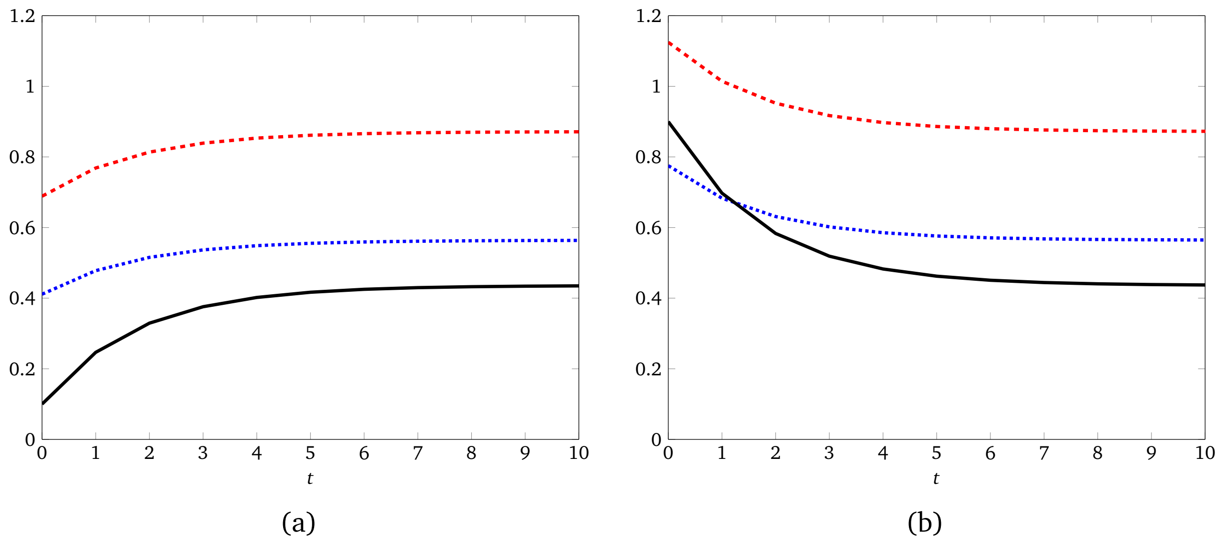

Figure 2.2: Reputation (solid black line), prices (dashed red line), and demand (dotted blue line) for $\alpha = 0.7$ and $\delta = 0.8$ in $t \in \{0, \ldots, 10\}$ with $\lambda_0 = 0.1$ (a), and $\lambda_0 = 0.9$ (b).

exploits that the good is mostly associated with her initially exclusive clientele so that selling to a broader, less exclusive, clientele only gradually deteriorates her reputation. As a result she realizes high profits in each period at the expense of slowly driving down next period's reputation.

If the good is relatively durable, $\alpha > \hat{\alpha}$, but the seller's reputation is bad initially $\lambda_0 < \hat{\lambda}$, we observe the mirror image of this behavior. See Figure 2.2a for a graphical representation. At first, the seller targets a small group of high-type buyers, and demand is low. The corresponding price is comparatively high given her low reputation, and, thus, reputation increases. As the good is durable, however, it takes time until the good is associated with these high-type buyers. Therefore, reputation only increases slowly. Over time, the seller gradually expands the set of buyers to which she sells but overall she continues to target a small, exclusive clientele. As a result, reputation increases at a decreasing rate. Analogously to the last case, prices increase over time because demand increases slower than reputation. Intuitively, even though the seller targets a larger group of buyers which would require setting a lower price, she can charge a higher price because the buyer's willingness to pay increases as her reputation improves. In this case, the seller builds reputation over time, i.e, she sacrifices profits in each period by selling only to an exclusive clientele in order to build reputation for future periods.

Altogether, the dynamics are relatively smooth and monotone if the good's durability is high. If the seller's initial reputation is high, reputation deteriorates monotonically to the long-run reputation because the seller milks her reputation in every period. If her initial reputation is low, reputation improves monotonically to the long-run reputation as the seller builds reputation in every period.

Second, we analyze the case when the good has low durability, $\alpha < \hat{\alpha}$. Consider a period in which the seller has good reputation. She sets a comparatively low price given her high repu-

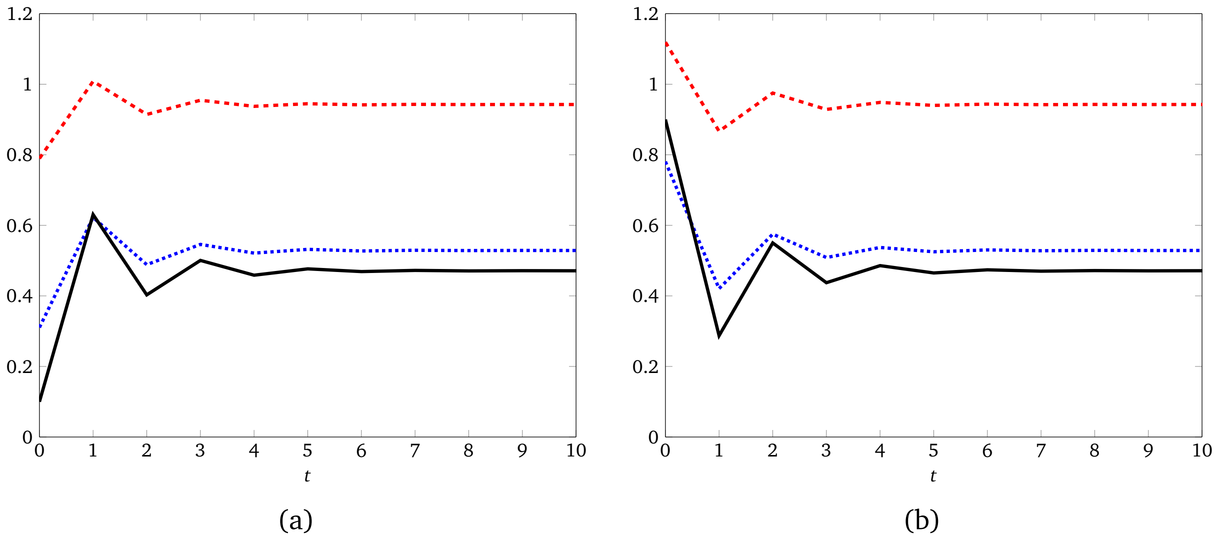

Figure 2.3: Reputation (solid black line), prices (dashed red line), and demand (dotted blue line) for $\alpha = 0.1$ and $\delta = 0.8$ in $t \in \{0, \dots, 10\}$ with $\lambda_0 = 0.1$ (a), and $\lambda_0 = 0.9$ (b).

tation such that demand is high, i.e., a large group of buyers with a diverse set of types purchases the good. In other words, the seller milks her high reputation and consequently reputation deteriorates. This is reminiscent of the behavior of a seller with good reputation when the good's durability is high. As the good's durability is low, however, current-period buyers represent a significant proportion of the buyers who possess the good in this case. Thus, the types of recent buyers, in particular of current-period buyers, have a significant impact on the seller's reputation in the next period. Consequently, reputation does not decrease gradually as before but drops below the long-run reputation. As a result, the seller has a bad reputation in the next period. Now consider a period in which the seller has bad reputation. The seller sets a comparatively high price and sells exclusively to a small group of high-type buyers, i.e., she builds reputation. As durability is low, these high-type buyers make up a significant part of the buyers who possess the good. As a result, reputation jumps above the long-run reputation. These observations imply that a period of good reputation is followed by a period of bad reputation and vice versa. As reputation converges to the long-run reputation, the amplitude of the alternations in reputation decreases over time. Analogously, prices alternate with decreasing amplitude. See Figure 2.3 for a graphical representation.

All in all, if the good's durability is low, reputation, demand, and prices fluctuate substantially. Periods of good reputation in which the seller milks her reputation alternate with periods of bad reputation in which the seller builds reputation. Intuitively, milking good reputation is profitable because the seller increases demand at a comparatively high price. Therefore, her current-period profits are particularly high. Conversely, building reputation in periods of bad reputation is less costly because, in any case, the seller's price is comparatively low due to her bad reputation. Hence, the seller has to give up relatively little in terms of profits today to increase reputation tomorrow.

Third, in the knife-edge case where the durability of the good coincides with $\bar{\alpha}$, reputation jumps immediately to the long-run reputation which balances the trade-off between current profits and the future benefits from a good reputation in the long run. Consequently, reputation remains constant in all subsequent periods.

We further investigate the effect of the good's durability on the reputation dynamics. Proposition 2.5 suggests that there is a monotone relationship between the durability of the good and the volatility of reputation dynamics. Intuitively, the lower the durability of the good, the more reputation fluctuates around the long-run reputation. This conjecture raises the issue of finding an appropriate measure for volatility to compare the volatility of two reputation sequences. In the following, we propose such a measure and formalize this intuitive conjecture.

It is instructive to make the dependence of reputation on α explicit. Fix a distance $h \in \mathbb{R}$ from long-run reputation $\hat{\lambda}(\alpha)$ and consider the reputation

$$\lambda(\alpha) \;:=\; \hat{\lambda}(\alpha) + h.$$

Let $\lambda'(\alpha)$ be the next-period reputation when following the optimal policy if the current-period reputation is $\lambda(\alpha)$, i.e.,

$$\lambda'(\alpha) \;:=\; \alpha\lambda(\alpha) + (1-\alpha)\theta^*(\lambda(\alpha)).$$

Intuitively, we define reputation relative to the long-run reputation $\hat{\lambda}(\alpha)$ to account for changes in $\hat{\lambda}(\alpha)$ as the durability of the good α changes.[11] Our measure for volatility is the distance between between $\lambda'(\alpha)$ and $\lambda(\alpha)$. The following result formalizes our conjecture:

Proposition 2.6. $|\lambda'(\alpha) - \lambda(\alpha)|$ *is decreasing in* α.

We show that the higher the durability of the good, the smaller the volatility of the seller's reputation. In other words, the higher α, the smaller the distance between two subsequent reputation levels. In this sense, there is a continuous transition between the three patterns in which reputation converges to the long-run reputation $\hat{\lambda}$, outlined in Proposition 2.5. This continuous relationship, gives rise to the question where the substantially different convergence behavior of the three cases comes from. Technically, it is the consequence of the fact that the distance between two subsequent reputation levels is increasing in α combined with the fact that the policy function is decreasing in λ and has a fixed point at $\hat{\lambda}$. For illustration, consider a seller who has a good reputation initially. If α is high, the distance between two subsequent reputation levels is relatively small. Thus, if the current-period reputation is larger than $\hat{\lambda}$, the next-period's reputation is larger than $\hat{\lambda}$ as well. Therefore, the seller milks

[11] See also the discussion following Proposition 2.6.

reputation in both periods. As α decreases the distance between two subsequent reputation levels increases. If α is sufficiently low, the next-period reputation is below $\hat{\lambda}$. Because the policy function is decreasing and has a fixed point at $\hat{\lambda}$, the seller switches from milking to building reputation in the next period.
To account for the fact that $\hat{\lambda}$ depends on α, we use the distance of two subsequent reputation levels defined relative to $\hat{\lambda}$ as our volatility measure. This is crucial because if we do not measure reputation relative to $\hat{\lambda}$, we ignore the dependence of $\hat{\lambda}$ on α. Then, it can happen for reputations close to $\hat{\lambda}$ that the decrease in $\hat{\lambda}$ outweighs the decrease in fluctuations around $\hat{\lambda}$ such that the distance between two subsequent reputation levels is increasing in α. As we are interested in the way reputation transitions to and fluctuates around $\hat{\lambda}$, however, we measure distances relative to $\hat{\lambda}$ in (the proof of) Proposition 2.6.

Intuitively, Proposition 2.5 and 2.6 predict that for goods with high durability, one should observe less fluctuations in reputation, clientele, and hence prices. The good is longer associated with its buyers' types and the seller manages her reputation such that it only gradually approaches the long-run reputation. In other words, in markets for relatively durable goods in which buyers care about the seller's past clientele, we should rarely observe price discounts. As mentioned in the introduction, this finding is consistent with the observation that the prices and reputation of the Swiss watchmaker *Rolex* have gradually increased over the last sixty years.
In contrast, for producers of goods with low durability our model predicts significant fluctuations in reputation, clientele, and prices. For these goods, reputation is mostly determined by current-period buyers. Hence, if a seller with a good reputation milks reputation by selling to a broader clientele, she incurs a severe negative effect on her next-period reputation. As outlined in the introduction, this provides a rationale for substantial changes in clientele and reputation of fashion labels such as *Abercrombie & Fitch* and *Ed Hardy*.

2.5 Planned Obsolescence

After studying the effect of the good's durability on the seller's reputation dynamics, we analyze how durability affects the profit of the seller. This analysis sheds light on planned obsolescence. Does a good with low durability yield higher profits for the seller than a good with high durability? The classical explanation of planned obsolescence is that it is optimal for a monopolist to produce a good with uneconomically short life to increase demand in future periods, see Bulow (1986). In our model, durability does not affect the size of the market in future periods. Nevertheless, we argue that a seller with low initial reputation can obtain higher profits if the good is less durable. Intuitively, if the seller has a low reputation, she can only charge comparatively low prices and hence, wants to build reputation over time. She improves her reputation to generate higher profits in the future by selling to an exclusive clien-

tele at the cost of current-period profits. If the durability of the good is comparatively low, the seller has to sacrifice less current-period profits to improve reputation because the new high types quickly replace the initial low types in the seller's clientele. As a result, the seller can obtain higher profits if the good is less durable. Put differently, reputation is more sensitive to changes in the seller's clientele if the durability of the good is low. While this makes it more difficult to maintain a high initial reputation, it is desirable for a seller with low initial reputation. Thus, our model reveals additional incentives for planned obsolescence from a reputational perspective which complements the classical explanation.
For the following result, it is convenient to make the dependence of $V(\lambda)$ on α explicit and write $V(\lambda, \alpha)$.

Proposition 2.7. *For any $\alpha \in (0,1)$ there exists a cutoff $\bar{\lambda}_0(\alpha) \in (0,1)$ such that*

(i) for all $\lambda_0 \geq \bar{\lambda}_0(\alpha)$ it holds that $V(\lambda_0, 1) \geq V(\lambda_0, \alpha)$,

(ii) for all $\lambda_0 \leq \bar{\lambda}_0(\alpha)$ it holds that $V(\lambda_0, 1) \leq V(\lambda_0, \alpha)$.

In Proposition 2.7 we compare the seller's profits from an everlasting good and a good with durability $\alpha \in (0,1)$. We show that there exists a cutoff $\bar{\lambda}_0(\alpha)$ such that a seller with initial reputation below $\bar{\lambda}_0(\alpha)$ can obtain higher profits if her good is not everlasting. Conversely, an everlasting good yields higher profits than the less durable good for a seller with initial reputation above $\bar{\lambda}_0$.

2.6 Concluding Remarks

In a setting where buyers care about who else consumes the good, we study how a seller dynamically manages reputation which is determined by her clientele. In contrast to the seller's private reputation which is the market's belief about the seller's private type, this public reputation captures the consumption externalities that buyers impose on each other. In a general setup, we show that public reputation is valuable for the seller and that there exists a trade-off between current-period profits and a higher reputation in subsequent periods. In a linear-uniform specification of the model, we find that reputation converges to a stable long-run reputation. In line with anecdotal evidence from various industries, the convergence path depends strongly on the good's characteristics. As the examples of *Ed Hardy* and *Abercrombie & Fitch* illustrate, reputation of less durable goods such as fashionable apparel fluctuates substantially. In contrast, the reputation of goods with higher durability such as luxury watches evolves monotonically as the case of the Swiss watchmaker *Rolex* exemplifies.

In future research, we think it is interesting to analyze reputation dynamics when other market characteristics, in addition to the good's durability, influence the reputation transition.

In a similar vein, studying the implications on reputation dynamics of competition between multiple sellers could shed light on how companies with very different reputation coexist in one market. For example, in the automobile industry it seems that the Italian car manufacturer *Ferrari* and the Romanian car manufacturer *Dacia* not only differ with respect to quality-related private reputation but also with respect to what their cars symbolize to others.
Although our results about price and reputation dynamics seem to match observations in several markets, our model considers public reputation as the only determinant of prices. Testing empirically how strongly pricing is influenced by public, clientele driven reputation as opposed to private, quality driven reputation would provide interesting insights.
Finally, we would like to highlight that the public reputation analyzed in this paper is present in several other applications as well, such as in the job market. A professor who considers whether to join a university, for example, takes into account the university's reputation, which depends on the academic standing of its past and current faculty members. In turn, she also influences the university's future reputation. Alternatively, young academics oftentimes work as an intern for a small salary in prestigious companies to improve their resume. Then again, the company's prestige is influenced by the quality of their employees. We believe that analyzing the implications of public reputation in this and further applications are promising avenues for future research.

2.A Appendix

Proof of Proposition 2.1.

We apply Theorems 4.2 - 4.6 from Stokey et al. (1989). Theorems 4.2 - 4.5 establish that we can study the seller's problem of choosing an infinite sequence of prices by means of the Bellman equation (2.4). Theorem 4.6 yields existence of a value function, $V(\lambda)$, which solves (2.4) as well as non-emptiness of the corresponding optimal policy correspondence. Taken together, these results yield existence of equilibrium. To employ Theorems 4.2 - 4.6, we verify that Assumptions 4.3 - 4.4 from Stokey et al. (1989) are satisfied in our model. We start by rewriting (2.4) to match their notational convention:

$$V(\lambda) = \sup_{\lambda' \in [\phi(\lambda,0),\phi(\lambda,1)]} \left\{ \left(1 - F(\phi^{-1}(\lambda,\lambda'))\right) u(\phi^{-1}(\lambda,\lambda'),\lambda) + \delta V(\lambda') \right\}, \tag{2.9}$$

where $\phi^{-1}(\lambda,\lambda')$ denotes the inverse of ϕ with respect to its second argument with domain $[0,1] \times [\phi(\lambda,0),\phi(\lambda,1)]$. Consider Assumption 4.3. First, note that the set of all possible states, $[0,1]$, is convex. Second, the correspondence which maps the initial state λ into the set of feasible next-period states, $[\phi(\lambda,0),\phi(\lambda,1)]$, is non-empty, compact-valued, and continuous because ϕ is continuous and strictly increasing in its second argument. To verify Assumption 4.4, we consider the per-period return function

$$\left(1 - F(\phi^{-1}(\lambda,\lambda'))\right) u(\phi^{-1}(\lambda,\lambda'),\lambda) \tag{2.10}$$

with domain $\{(\lambda,\lambda') \mid (\lambda,\lambda') \in [0,1] \times [\phi(\lambda,0),\phi(\lambda,1)]\}$. Because $[0,1]$ is a compact set and F, ϕ, and u are continuous, (2.10) is continuous as well. Together with the compactness of the domain, this observation also implies that (2.10) is bounded. Thus, Assumptions 4.3 - 4.4 hold, and we can apply Theorems 4.2 - 4.6. □

Proof of Proposition 2.2.

Consider $\tilde{\lambda}$, $\lambda \in [0,1]$ such that $\tilde{\lambda} \geq \lambda$. We argue that $V(\tilde{\lambda}) \geq V(\lambda)$ by showing that for any strategy of the seller with initial reputation λ (henceforth seller$_\lambda$) we can find a strategy of the seller with initial reputation $\tilde{\lambda}$ (henceforth seller$_{\tilde{\lambda}}$) which gives her weakly higher profits. Any strategy of the seller induces a sequence of prices. Because u is strictly increasing, there is a one-to-one relationship between prices and cutoffs. For this proof, it is instructive to consider sequences of cutoffs rather than sequences of prices. Let $\{\theta_s^\dagger\}_{s=0}^\infty$ be an arbitrary sequence of cutoffs chosen by seller$_\lambda$. Note that seller$_{\tilde{\lambda}}$ can choose prices that induce the same sequence of cutoffs because buyers' utility is quasilinear in the price. If seller$_{\tilde{\lambda}}$ chooses the same sequence of cutoffs $\{\tilde{\theta}_s\}_{s=0}^\infty$ such that $\tilde{\theta}_s^\dagger = \theta_s^\dagger$:

(i) Seller$_{\tilde{\lambda}}$ sells to the same (mass of) buyers as seller$_\lambda$ in each period.

(ii) The price that seller$_{\tilde{\lambda}}$ charges in the first period, $u(\theta_0^\dagger, \tilde{\lambda})$, is higher than the price that seller$_\lambda$ charges in the first period, $u(\theta_0^\dagger, \lambda)$, because u is increasing in the second argument.

(iii) As ϕ is increasing in its first argument, we have $\phi(\tilde{\lambda}, \theta_0^\dagger) \geq \phi(\lambda, \theta_0^\dagger)$. Therefore, seller$_{\tilde{\lambda}}$'s next-period reputation is higher than seller$_\lambda$'s next-period reputation. Using the same argument as in (ii), we deduce that seller$_{\tilde{\lambda}}$ also sets higher prices in the future.

Taking (i) to (iii) together, we see that $\{\tilde{\theta}_s^\dagger\}_{s=0}^\infty$ guarantees seller$_{\tilde{\lambda}}$ a higher profit than $\{\theta_s^\dagger\}_{s=0}^\infty$ guarantees seller$_\lambda$. Furthermore, by Proposition 2.1, there exists a Markovian strategy that yields seller$_{\tilde{\lambda}}$ a weakly higher profit than $\{\tilde{\theta}_s^\dagger\}_{s=0}^\infty$. We conclude that $V(\tilde{\lambda}) \geq V(\lambda)$. □

Proof of Lemma 2.1.

First, we adjust the Bellman equation, (2.4), to account for the uniform-linear model specification. Also, as in the proof of Proposition 2.1, it is convenient to let the seller choose cutoffs directly instead of prices. With linear utility, setting a price p_t at reputation λ_t implies that the cutoff is

$$\theta^\dagger(\lambda_t, p_t) = p_t - \lambda_t.$$

Thus, we can express the price as a function of the cutoff,

$$p_t\left(\lambda_t, \theta^\dagger\right) = \theta^\dagger + \lambda_t.$$

For notational convenience, we abbreviate $\theta^\dagger(\lambda, p)$ with $\theta^\dagger$ in the following. The uniform distribution of types implies that $F\left(\theta^\dagger\right) = \theta^\dagger$. Therefore, we can rewrite (2.4) to

$$V(\lambda) = \max_{\theta^\dagger \in [0,1]} \left\{\left(\theta^\dagger + \lambda\right)\left(1 - \theta^\dagger\right) + \delta V\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right)\right\}, \tag{2.11}$$

where the seller maximizes with respect to $\theta^\dagger$ instead of p.

By Proposition 2.1, $V(\lambda)$ is the fixed point of the Bellman operator. We now prove that the Bellman operator maps the (closed) set of functions which are continuous in α, δ, and λ (henceforth denoted by C) into itself. Thus, the fixed point of the Bellman operator must lie within that closed set and, hence, $V(\lambda)$ is continuous in α, δ, and λ.
To show that the Bellman operator maps C into itself, consider an arbitrary $g \in C$. Applying the Bellman operator, we obtain

$$\max_{\theta^\dagger \in [0,1]} \left\{\left(\theta^\dagger + \lambda\right)\left(1 - \theta^\dagger\right) + \delta g\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right)\right\} \tag{2.12}$$

We use Berge's maximum theorem to deduce that (2.12) is a continuous function of α, δ, and λ. First, note that for any λ the set of feasible choices is $[0, 1]$ which is compact and continuous as a correspondence of α, δ, and λ because it is independent of α, δ, and λ. Second, note that

$$\left(\theta^\dagger + \lambda\right)\left(1 - \theta^\dagger\right) + \delta g\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right)$$

is continuous in α, δ, λ, and $\theta^\dagger$. Thus, by the maximum theorem, (2.12) is an element of C. Consequently, the Bellman operator maps C into itself and $V \in C$. □

Proof of Proposition 2.3.

We consider the case where $\alpha < 1$. For $\alpha = 1$ see Lemma 2.2. Recall that the Bellman equation is given by (2.11).

(i) *Step 1: We solve the Bellman equation with a guess and verify approach.*

Consider the first-order derivative of the right-hand side of the Bellman equation with respect to $\theta^\dagger$

$$1 - \theta^\dagger - \left(\theta^\dagger + \lambda\right) + \delta(1-\alpha)V'\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right), \tag{2.13}$$

and the first-order derivative of $V(\lambda)$ with respect to λ

$$V'(\lambda) \; = \; 1 - \theta^\dagger + \delta\alpha V'\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right). \tag{2.14}$$

We guess that the value function is quadratic, that is,

$$V(\lambda) \; = \; v_2\lambda^2 + v_1\lambda + v_0.$$

This implies for the first-order derivative

$$V'(\lambda) \; = \; 2v_2\lambda + v_1. \tag{2.15}$$

In order to verify the guess, we consider the first-order condition of the Bellman equation by setting (2.13) equal to zero,

$$1 - \theta^\dagger - \left(\theta^\dagger + \lambda\right) + \delta(1-\alpha)V'\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right) \; = \; 0. \tag{2.16}$$

Using (2.15), we can rewrite (2.16) to

$$1 - \theta^\dagger - \left(\theta^\dagger + \lambda\right) + \delta(1-\alpha)\left(2v_2\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right) + v_1\right) \; = \; 0.$$

Solving for $\theta^\dagger$ yields

$$\theta^\dagger = \frac{1 + \delta(1-\alpha)v_1 - \lambda(1 - \delta(1-\alpha)(2v_2\alpha))}{2(1-\delta(1-\alpha)^2 v_2)}. \tag{2.17}$$

Now, we consider (2.14) where we plug in (2.15), the guess for $V'(\lambda)$, and (2.17). After some algebra, we obtain

$$\begin{aligned} V'(\lambda) =& \frac{\lambda(1 + 4\delta\alpha v_2(2\alpha - 1))}{2(1-\delta(1-\alpha)^2 v_2)} \\ &+ \frac{2(1+\delta\alpha v_1)\big(1 - \delta(1-\alpha)^2 v_2\big) - 1 + \delta(1-\alpha)(2\alpha v_2(1+\delta(1-\alpha)v_1) - v_1)}{2(1-\delta(1-\alpha)^2 v_2)}. \end{aligned} \tag{2.18}$$

We use the guess for the left-hand side in order to match and thereby solve for coefficients. First, we require that the coefficient of λ on the right-hand side of (2.18) equals $2v_2$, the coefficient of λ in (2.15), i.e.,

$$2v_2 = \frac{1 + 4\delta\alpha v_2(2\alpha - 1)}{2(1-\delta(1-\alpha)^2 v_2)}.$$

Solving for v_2 gives

$$v_2 = \frac{1 - \delta\alpha(2\alpha-1) \pm \sqrt{(1-\delta\alpha(2\alpha-1))^2 - \delta(1-\alpha)^2}}{2\delta(1-\alpha)^2}. \tag{2.19}$$

We see that there are two possible solutions for v_2. Second, we require that the constant on the right-hand side of (2.18) equals v_1 which is the constant in the guess (2.15), i.e.,

$$v_1 = \frac{2(1+\delta\alpha v_1)\big(1 - \delta(1-\alpha)^2 v_2\big) - 1 + \delta(1-\alpha)(2\alpha v_2(1+\delta(1-\alpha)v_1) - v_1)}{2(1-\delta(1-\alpha)^2 v_2)}.$$

Solving for v_1 yields

$$v_1 = \frac{2\big(1 - \delta(1-\alpha)^2 v_2\big) - 1 + 2\delta\alpha(1-\alpha)v_2}{2(1-\delta(1-\alpha)^2 v_2)(1-\delta\alpha) + \delta(1-\alpha)(1 - 2\delta\alpha(1-\alpha)v_2)} \tag{2.20}$$

which depends on v_2.

Step 2: We argue that only one of the two solutions for v_2 is valid.

Denote the two candidates for v_2 by

$$v_2^+ = \frac{1 - \delta\alpha(2\alpha - 1) + \sqrt{(1 - \delta\alpha(2\alpha - 1))^2 - \delta(1 - \alpha)^2}}{2\delta(1 - \alpha)^2}, \tag{2.21}$$

$$v_2^- = \frac{1 - \delta\alpha(2\alpha - 1) - \sqrt{(1 - \delta\alpha(2\alpha - 1))^2 - \delta(1 - \alpha)^2}}{2\delta(1 - \alpha)^2}. \tag{2.22}$$

First, recall from Lemma 2.1 that the value function $V(\lambda)$ is continuous in α and δ. Rewriting the square root in the numerator gives

$$\sqrt{(1 - \delta\alpha^2)(1 - \delta + 4\alpha\delta(1 - \alpha))}.$$

Observe that both solutions coincide if and only if $\alpha = \delta = 1$. As we consider $\delta \in (0, 1)$ we conclude that $v_2^+ \neq v_2^-$. Hence, only one of the two candidates can be the global solution for any pair (α, δ). If both were solutions, $V(\lambda)$ would not be continuous.

Second, we argue that value function induced by v_2^+ is not increasing for all pairs (α, δ) which contradicts Proposition 2.2. Let v_1^+ be the solution to (2.20) when we plug in v_2^+ for v_2. Some algebra yields

$$v_1^+ = \frac{\alpha + (1 - 2\alpha)\left(\delta\alpha(2\alpha - 1) - \sqrt{(1 - \delta\alpha^2)(1 - \delta + 4\alpha\delta(1 - \alpha))}\right)}{(1 - \alpha)\left(1 - \delta\alpha + \delta(1 - \alpha)(1 - 2\alpha) - \sqrt{(1 - \delta\alpha^2)(1 - \delta + 4\alpha\delta(1 - \alpha))}\right)}. \tag{2.23}$$

Next, we show that there exists a pair of (α, δ) for which v_1^+ is strictly negative which contradicts that the value function $V(\lambda)$ is strictly increasing in λ for any pair (δ, α). For $\alpha = 0$ we get

$$\frac{-\sqrt{1 - \delta}}{1 + \delta - \sqrt{1 - \delta}}, \tag{2.24}$$

which is strictly negative for all $\delta \in (0, 1)$. By continuity, we conclude that there exists an $\alpha > 0$ such that $v_1^+ < 0$. Consequently, v_2^+ is not the solution. Instead, the solution is given by v_2^- and the corresponding solution for v_1, that is,

$$v_1^- = \frac{\alpha + (1 - 2\alpha)\left(\delta\alpha(2\alpha - 1) + \sqrt{(1 - \delta\alpha^2)(1 - \delta + 4\alpha\delta(1 - \alpha))}\right)}{(1 - \alpha)\left(1 - \delta\alpha + \delta(1 - \alpha)(1 - 2\alpha) + \sqrt{(1 - \delta\alpha^2)(1 - \delta + 4\alpha\delta(1 - \alpha))}\right)}. \tag{2.25}$$

From now on, we identify v_1 with v_1^- and v_2 with v_2^-.

Step 3: We verify that the second-order condition is satisfied.

Consider the second-order derivative of the right-hand side of the Bellman equation with respect to $\theta^\dagger$

$$-2\left(1-\delta(1-\alpha)^2 v_2\right). \tag{2.26}$$

We start by showing that $v_2 \geq 0$ which is the case if the numerator of v_2^- is positive, i.e.,

$$1-\delta\alpha(2\alpha-1) \;\geq\; \sqrt{(1-\delta\alpha(2\alpha-1))^2-\delta(1-\alpha)^2}.$$

By squaring both sides of the equation and subtracting $(1-\delta\alpha(2\alpha-1))^2$, the condition simplifies to

$$\delta(1-\alpha)^2 \;\geq\; 0$$

which is always satisfied. Therefore, the second-order condition holds if and only if $\delta(1-\alpha)^2 v_2 < 1$. We show that $2(1-\alpha)v_2 \leq 1$ which implies that $\delta(1-\alpha)^2 v_2 < 1$. After plugging in (2.22), we can rewrite $2(1-\alpha)v_2 \leq 1$ to

$$1-\delta\alpha(2\alpha-1)-\sqrt{(1-\delta\alpha(2\alpha-1))^2-\delta(1-\alpha)^2} \;\leq\; \delta(1-\alpha).$$

Rearranging terms and squaring both sides yields

$$(1-\delta\alpha(2\alpha-1))^2-\delta(1-\alpha)^2 \;\geq\; (1-\delta+2\delta\alpha(1-\alpha))^2.$$

After some algebra, this condition can be reformulated to

$$-1+\alpha^2+\delta\left(1-4\alpha+7\alpha^2-4\alpha^3\right) \;\leq\; 0. \tag{2.27}$$

For $\delta = 0$, it simplifies to

$$\alpha^2-1 \;\leq\; 0$$

which is always true, and for $\delta = 1$ we get

$$-4(1-\alpha)^2 \;\leq\; 0$$

which is satisfied as well. As (2.27) is linear in δ, the condition holds for all $\delta \in (0,1)$ and we can conclude that $\delta(1-\alpha)^2 v_2 < 1$. Hence, the second-order condition is satisfied.

(ii) The shape of the policy function follows now immediately from (2.17), where we plug in v_1 and v_2 in order to determine the solutions for θ_0 and θ_1. We are left to show that

$-1 < \theta_1 < 0$. θ_1 is defined as the coefficient of λ in (2.17), i.e.,

$$\theta_1 = -\frac{1 - \delta(1-\alpha)(2v_2\alpha)}{2(1-\delta(1-\alpha)^2 v_2)}. \tag{2.28}$$

Recall from Step 3 of the previous part of the proof that $v_2 \geq 0$ and $2(1-\alpha)v_2 \leq 1$ which implies that $2\delta(1-\alpha)\alpha v_2 < 1$ and $\delta(1-\alpha)^2 v_2 < 1$. Therefore, numerator and denominator are positive and θ_1 is strictly negative. Moreover, $\theta_1 > -1$ if the denominator is larger than the numerator. This is the case if $2\delta(1-\alpha)^2 v_2 < 1$ which follows from $2(1-\alpha)v_2 \leq 1$. Recall also that the second-order derivative of the right-hand side of the Bellman equation, (2.26), is strictly negative. Hence, the objective is strictly concave and its maximizer is unique. Therefore, the policy function is also unique. This concludes the proof. □

Proof of Proposition 2.4.

(i) Consider $\alpha < 1$. In the first step, we define

$$T : [0,1] \longrightarrow [0,1], \quad \lambda_t \longmapsto \lambda_{t+1},$$

as the mapping of today's into tomorrow's reputation. Using the seller's policy function (2.7) and the reputation transition (2.6), we obtain

$$T(\lambda_t) = \alpha\lambda_t + (1-\alpha)(\theta_0 + \theta_1\lambda_t).$$

Note that the mapping is linear in λ_t with slope $\alpha + (1-\alpha)\theta_1$. As $\theta_1 \in (-1, 0)$, it follows that $\alpha + (1-\alpha)\theta_1 \in (0,1)$ for any $\alpha \in [0,1)$ and, hence, T is a contraction mapping. By Banach's fixed-point theorem, the contraction mapping T has a unique fixed point, that is, there is a unique $\hat{\lambda}$ such that $T(\hat{\lambda}) = \hat{\lambda}$, and, moreover, the sequence $\lambda_{t+1} = T(\lambda_t)$ converges to $\hat{\lambda}$ as $t \to \infty$.
Second, we explicitly compute $\hat{\lambda}$. To this end, we revisit the seller's problem from the proof of Proposition 2.3, in particular, the seller's first-order condition, as in (2.16),

$$1 - 2\theta^\dagger(\lambda) - \lambda + \delta(1-\alpha)V'\big(\alpha\lambda + (1-\alpha)\theta^\dagger(\lambda)\big) = 0,$$

and the first-order derivative with respect to λ, as in (2.18),

$$V'(\lambda) = 1 - \theta^\dagger(\lambda) + \delta\alpha V'\big(\alpha\lambda + (1-\alpha)\theta^\dagger(\lambda)\big).$$

Note that $T(\hat{\lambda}) = \hat{\lambda}$ implies $\theta^{\dagger}(\hat{\lambda}) = \hat{\lambda}$. By evaluating both conditions at the fixed point $\lambda = \hat{\lambda}$, we get

$$1 - 3\hat{\lambda} + \delta(1-\alpha)V'(\hat{\lambda}) \;=\; 0, \tag{2.29}$$

and

$$V'(\hat{\lambda}) \;=\; 1 - \hat{\lambda} + \delta\alpha V'(\hat{\lambda}). \tag{2.30}$$

Solving (2.30) for $V'(\hat{\lambda})$ yields

$$V'(\hat{\lambda}) \;=\; \frac{1-\hat{\lambda}}{1-\delta\alpha} \tag{2.31}$$

which we plug into (2.29) in order to solve for $\hat{\lambda}$. The solution is

$$\hat{\lambda} \;=\; \frac{1 + \frac{\delta(1-\alpha)}{1-\delta\alpha}}{3 + \frac{\delta(1-\alpha)}{1-\delta\alpha}},$$

as claimed.

Finally, we consider the derivatives of $\hat{\lambda}$ with respect to δ and α. For convenience, denote $\frac{\delta(1-\alpha)}{1-\delta\alpha}$ by $x(\alpha,\delta)$. As

$$\frac{\partial\hat{\lambda}}{\partial\delta} \;=\; \frac{2\frac{\partial x(\alpha,\delta)}{\partial\delta}}{(3 + x(\alpha,\delta))^2}, \quad \text{and} \quad \frac{\partial\hat{\lambda}}{\partial\alpha} \;=\; \frac{2\frac{\partial x(\alpha,\delta)}{\partial\alpha}}{(3 + x(\alpha,\delta))^2},$$

the signs of the derivatives of $\hat{\lambda}$ are determined by the signs of the derivatives of $x(\alpha,\delta)$. The latter are given by

$$\frac{\partial x(\alpha,\delta)}{\partial\delta} \;=\; \frac{1-\alpha}{(1-\delta\alpha)^2} \;>\; 0, \quad \text{and} \quad \frac{\partial x(\alpha,\delta)}{\partial\alpha} \;=\; -\frac{\delta(1-\delta)}{(1-\delta\alpha)^2} \;<\; 0.$$

Hence, $\hat{\lambda}$ is increasing in δ and decreasing in α.

(ii) For $\alpha = 1$, (2.6) implies that $\lambda_{t+1} = \lambda_t$ and, hence, $\lambda_t = \lambda_0$, for all $t \geq 0$, which completes the proof. □

Proof of Lemma 2.2.

Note that for $\alpha = 1$, (2.6) implies that $\lambda_{t+1} = \lambda_t$. When starting with initial reputation λ_0, it follows that $\lambda_t = \lambda_0$, for all $t \in \{1, \ldots\}$. Thus, the seller's decision about $\theta^{\dagger}$ in period t does not influence her reputation in the next period. Her problem reduces to maximize stage game payoffs in each period.

Therefore, the seller faces the following problem, in each period,

$$\max_{\theta^\dagger}\left\{\left(\lambda_0+\theta^\dagger\right)\left(1-\theta^\dagger\right)\right\}.$$

The objective is strictly concave in $\theta^\dagger$. The first-order condition is given by

$$1-\lambda_0-2\theta^\dagger \;=\; 0,$$

and solving for $\theta^\dagger$ yields

$$\theta^\dagger \;=\; \frac{1-\lambda_0}{2}.$$

The seller's profits are given by the discounted infinite sum of period profits. In each period, the seller receives

$$\left(\lambda_0+\frac{1-\lambda_0}{2}\right)\left(1-\frac{1-\lambda_0}{2}\right) \;=\; \left(\frac{1+\lambda_0}{2}\right)^2.$$

Hence, her profits are given by

$$\sum_{t=0}^{\infty}\delta^t\left(\frac{1+\lambda_0}{2}\right)^2 \;=\; \frac{1}{1-\delta}\left(\frac{1+\lambda_0}{2}\right)^2.$$

□

Proof of Lemma 2.3.

This result follows from Propositions 2.3 and 2.4. First, consider v_1 as in (2.25) and v_2 as in (2.22) for $\alpha=0$,

$$v_1 \;=\; \frac{\sqrt{1-\delta}}{1+\delta+\sqrt{1-\delta}}, \tag{2.32}$$

$$v_2 \;=\; \frac{1-\sqrt{1-\delta}}{2\delta}. \tag{2.33}$$

We determine the seller's policy function by evaluating (2.13) for $\alpha=0$,

$$\theta^*(\lambda) \;=\; -\frac{\lambda}{2(1-\delta v_2)}+\frac{1+\delta v_1}{2(1-\delta v_2)}. \tag{2.34}$$

With (2.32) and (2.33), we can rewrite (2.34) after some algebra to

$$\theta^*(\lambda) \;=\; -\frac{\lambda}{1+\sqrt{1-\delta}}+\frac{1+\delta}{1+\delta+\sqrt{1-\delta}}. \tag{2.35}$$

In order to determine the long-run reputation, we evaluate (2.8) at $\alpha = 0$ which yields

$$\hat{\lambda} = \frac{1+\delta}{3+\delta}.$$

Further, from (2.6) for $\alpha = 0$ follows $\lambda_{t+1} = \theta_t^{\dagger}$. From the proof of Proposition 2.4, we know that $\theta^*(\hat{\lambda}) = \hat{\lambda}$. Moreover, $\theta^*(\lambda)$ is decreasing by Proposition 2.3. Therefore, $\lambda_t > \hat{\lambda}$ implies $\lambda_{t+1} < \hat{\lambda}$ and vice versa, that is, reputation oscillates towards $\hat{\lambda}$. This completes the proof. □

Proof of Proposition 2.5.

We proceed in several steps.
Step 1: We derive θ_1 in closed form.
It is convenient to define

$$\Delta := (1-\alpha^2\delta)(1-\delta(1-2\alpha)^2).$$

Recall that we obtained θ_1 from the first-order condition as

$$\theta_1 = \frac{\delta(1-\alpha)\alpha 2v_2 - 1}{2(1-\delta(1-\alpha)^2 v_2)},$$

where

$$v_2 = \frac{1-\delta\alpha(2\alpha-1)-\sqrt{\Delta}}{2\delta(1-\alpha)^2}.$$

Inserting v_2 into θ_1 gives

$$\frac{\frac{\alpha}{1-\alpha}\left(1-\frac{1-\alpha}{\alpha}-\delta\alpha(2\alpha-1)-\sqrt{\Delta}\right)}{1+\delta\alpha(2\alpha-1)+\sqrt{\Delta}}. \tag{2.36}$$

Next, we multiply the numerator and the denominator of (2.36) by

$$1-\frac{1-\alpha}{\alpha}-\delta\alpha(2\alpha-1)+\sqrt{\Delta}.$$

Consider first the numerator of (2.36):

$$\frac{\alpha}{1-\alpha}\left(\left(1-\frac{1-\alpha}{\alpha}-\delta\alpha(2\alpha-1)\right)^2-\Delta\right).$$

Some algebra shows that the numerator can be written as

$$-3+\frac{1}{\alpha}+\alpha\delta(3\alpha-1). \tag{2.37}$$

Now, consider the denominator of (2.36),

$$\left(1 + \delta\alpha(2\alpha - 1) + \sqrt{\Delta}\right)\left(1 - \frac{1-\alpha}{\alpha} - \delta\alpha(2\alpha - 1) + \sqrt{\Delta}\right).$$

Simplifying the latter expression yields

$$\left(\sqrt{\Delta} + (1 - \alpha^2\delta)\right)\left(\frac{3\alpha - 1}{\alpha}\right). \tag{2.38}$$

Putting (2.37) and (2.38) together, we obtain for θ_1:

$$\theta_1 = \frac{-3 + \frac{1}{\alpha} + \alpha\delta(3\alpha - 1)}{\left(\sqrt{\Delta} + (1 - \alpha^2\delta)\right)\left(\frac{3\alpha-1}{\alpha}\right)},$$

which we can further simplify to

$$\theta_1 = \frac{-(1 - \alpha^2\delta)}{1 - \alpha^2\delta + \sqrt{\Delta}}. \tag{2.39}$$

Step 2: We derive $\frac{\partial\theta_1}{\partial\alpha}$ *in closed form.*
We start by taking the derivative of θ_1, see (2.39) from Step 1, with respect to α:

$$\frac{2\alpha\delta(1 - \alpha^2\delta + \sqrt{\Delta}) - (\alpha^2\delta - 1)\left(-2\alpha\delta + \frac{1}{2}\Delta^{-\frac{1}{2}}\left(-2\delta\alpha(1 - \delta(1 - 2\alpha)^2) + (1 - \alpha^2\delta)(4\delta(1 - 2\alpha))\right)\right)}{\left(1 - \alpha^2\delta + \sqrt{\Delta}\right)^2}.$$

Next, we multiply numerator and denominator by $\sqrt{\Delta}$. The numerator becomes

$$2\alpha\delta\Delta + (1 - \alpha^2\delta)\frac{1}{2}\left(-2\delta\alpha(1 - \delta(1 - 2\alpha)^2) + (1 - \alpha^2\delta)(4\delta(1 - 2\alpha))\right),$$

and simplifying yields

$$\delta(1 - \alpha^2\delta)(2 - 3\alpha - \alpha\delta + 2\alpha^2\delta).$$

Putting numerator and denominator together gives

$$\frac{\partial\theta_1}{\partial\alpha} = \frac{\delta(1 - \alpha^2\delta)(2 + 2\alpha^2\delta - 3\alpha - \alpha\delta)}{(1 - \alpha^2\delta + \sqrt{\Delta})^2\sqrt{\Delta}}.$$

Inserting the definition of Δ yields

$$\frac{\partial\theta_1}{\partial\alpha} = \frac{\delta(1 - \alpha^2\delta)(2 + 2\alpha^2\delta - 3\alpha - \alpha\delta)}{(1 - \alpha^2\delta + \sqrt{(1 - \alpha^2\delta)(1 - \delta(1 - 2\alpha)^2)})^2\sqrt{(1 - \alpha^2\delta)(1 - \delta(1 - 2\alpha)^2)}}. \tag{2.40}$$

Step 3: $(1 - \alpha)\theta_1$ *is increasing in* α.

The statement is equivalent to showing that

$$-\theta_1 + (1-\alpha)\frac{\partial \theta_1}{\partial \alpha} \;\geq\; 0. \tag{2.41}$$

We plug in the explicit formulas for θ_1 and $\frac{\partial \theta_1}{\partial \alpha}$ which we derived in Step 1 and 2 and rewrite the left side of inequality (2.41) as a single fraction over the denominator

$$(1-\alpha^2\delta + \sqrt{\Delta})^2\sqrt{\Delta}.$$

As the denominator is positive, it suffices to argue that the numerator

$$(1-\alpha)\delta(1-\alpha^2\delta)(2+2\alpha^2\delta-3\alpha-\alpha\delta) + (1-\alpha^2\delta)(1-\alpha^2\delta+\sqrt{\Delta})(\sqrt{\Delta})$$

is also positive. The latter expression is bounded from below by

$$(1-\alpha)\delta(1-\alpha^2\delta)(2+2\alpha^2\delta-3\alpha-\alpha\delta) + (1-\alpha^2\delta)\Delta.$$

The definition of Δ and some algebra yield

$$(1-\alpha^2\delta)\big((1-\delta)(1-\alpha^2\delta) + \delta(1-\alpha)(2+2\alpha^2\delta+\alpha-\alpha\delta-4\alpha^3\delta)\big).$$

It remains to be shown that

$$2+2\alpha^2\delta+\alpha-\alpha\delta-4\alpha^3\delta \tag{2.42}$$

is positive. Observe that (2.42) is linear in δ. Therefore, it attains its minimum either at $\delta = 1$ or $\delta = 0$. For $\delta = 0$ we obtain

$$2+\alpha \;>\; 0,$$

and for $\delta = 1$, (2.42) becomes

$$2+2\alpha^2-4\alpha^3 \;\geq\; 0.$$

Consequently, (2.42) is positive for all $\delta \in (0,1)$ which implies that (2.41) holds.

Step 4: We prove (i)-(iii).

Note that $\alpha + (1-\alpha)\theta_1$ is continuous in α, strictly increasing in α by Step 3, less than zero at $\alpha = 0$, and larger than zero at $\alpha = 1$. Therefore, we can uniquely define $\bar{\alpha}$ through

$$\bar{\alpha} + (1-\bar{\alpha})\theta_1 \;=\; 0.$$

(i) Consider $\alpha > \bar{\alpha}$. Fix $\lambda_0 \geq \hat{\lambda}$. Denote by λ the current-period reputation and by λ' the next-period reputation. First, we argue that $\lambda \geq \hat{\lambda}$ implies $\lambda' \geq \hat{\lambda}$:

$$
\begin{aligned}
\lambda' &= \alpha\lambda + (1-\alpha)\theta^*(\lambda) \\
&= \alpha\lambda + (1-\alpha)(\theta_1\lambda + \theta_0) \\
&= \lambda \underbrace{(\alpha + (1-\alpha)\theta_1)}_{>0 \text{ as } \alpha>\bar{\alpha}} + (1-\alpha)\theta_0 \\
&\geq \hat{\lambda}(\alpha + (1-\alpha)\theta_1) + (1-\alpha)\theta_0 \\
&= \alpha\hat{\lambda} + (1-\alpha)\theta^*(\hat{\lambda}) \\
&= \hat{\lambda}.
\end{aligned}
\tag{2.43}
$$

Thus, if $\lambda_0 \geq \hat{\lambda}$ then $\lambda_t \geq \hat{\lambda}$, for all t. Second, we show that $\lambda \geq \hat{\lambda}$ implies $\lambda \geq \lambda'$. Recall that $\theta^*(\hat{\lambda}) = \hat{\lambda}$, and that θ^* is decreasing in λ. Thus, we have

$$\theta^*(\lambda) \;<\; \hat{\lambda} \;<\; \lambda,$$

and consequently

$$\lambda' \;=\; \alpha\lambda + (1-\alpha)\theta^*(\lambda) \;\leq\; \alpha\lambda + (1-\alpha)\lambda \;=\; \lambda.$$

We deduce that if $\lambda_0 \geq \hat{\lambda}$ then $\lambda_t \geq \lambda_{t+1}$, $\forall t$. From Proposition 2.4 we know that $\lim_{t\to\infty} \lambda_t = \hat{\lambda}$, hence, $|\lambda_t - \hat{\lambda}| \downarrow 0$. The proof for $\lambda_0 \leq \hat{\lambda}$ is analogous.

(ii) Let $\alpha = \bar{\alpha}$. Fix any $\lambda_0 \in [0,1]$. Carefully inspecting the arguments from (i) shows that (2.43) holds with equality if $\alpha = \hat{\alpha}$. Thus, we have $\lambda_t = \hat{\lambda}$, for all $t \geq 1$.

(iii) Consider $\alpha < \bar{\alpha}$. Fix any $\lambda_0 \in [0,1]$. In this case the inequality in (2.43) is reversed, i.e., if $\lambda \geq \hat{\lambda}$ then $\lambda' \leq \hat{\lambda}$. The proof that if $\lambda \leq \hat{\lambda}$ then $\lambda' \geq \hat{\lambda}$ follows from an analogous computation. Next, we show that

$$|\lambda - \hat{\lambda}| \;\geq\; |\lambda' - \hat{\lambda}|. \tag{2.44}$$

Again, we focus on $\lambda \geq \hat{\lambda}$, the proof for $\lambda \leq \hat{\lambda}$ is analogous. If $\lambda \geq \hat{\lambda}$, (2.44) becomes

$$\lambda - \hat{\lambda} \;\geq\; \hat{\lambda} - \lambda' \quad\Leftrightarrow\quad \lambda + \lambda' \;\geq\; 2\hat{\lambda}. \tag{2.45}$$

Observe that

$$
\begin{aligned}
\lambda + \lambda' &= \lambda \underbrace{(1 + \alpha + (1-\alpha)\theta_1)}_{>0} + (1-\alpha)\theta_0 \\
&\geq \hat{\lambda} + \hat{\lambda}(\alpha + (1-\alpha)\theta_1) + (1-\alpha)\theta_0 \\
&= 2\hat{\lambda},
\end{aligned}
$$

which yields (2.45). Together with the convergence established in Proposition 2.4, we obtain $|\lambda_t - \hat{\lambda}| \downarrow 0$. This concludes the proof. □

Proof of Proposition 2.6.

We argue that

$$|\lambda'(\alpha) - \lambda(\alpha)| \tag{2.46}$$

is decreasing in α. Inserting the definition of $\lambda'(\alpha)$ into (2.46), we obtain

$$(1-\alpha)|\theta^*(\lambda(\alpha)) - \lambda(\alpha)|.$$

For $h < 0$, we have $\lambda(\alpha) < \hat{\lambda}(\alpha)$ and thus $\theta^*(\lambda(\alpha)) > \hat{\lambda}(\alpha) > \lambda(\alpha)$. Therefore, we obtain

$$(1-\alpha)(\theta^*(\lambda(\alpha)) - \lambda(\alpha)). \tag{2.47}$$

Recall that

$$\theta^*(\lambda(\alpha)) \; = \; \theta_1\lambda(\alpha) + \theta_0. \tag{2.48}$$

Taking the derivative of (2.47) with respect to α we obtain

$$-\big(\theta_1(\hat{\lambda}(\alpha)+h) + \theta_0 - (\hat{\lambda}(\alpha)+h)\big) + (1-\alpha)\left(\frac{\partial\theta_1}{\partial\alpha}(\hat{\lambda}(\alpha)+h) + \theta_1\frac{\partial\hat{\lambda}}{\partial\alpha} + \frac{\partial\theta_0}{\partial\alpha} - \frac{\partial\hat{\lambda}}{\partial\alpha}\right), \tag{2.49}$$

where we used (2.48) and the definition of $\lambda(\alpha)$. By definition of the long-run reputation $\hat{\lambda}(\alpha)$, it holds

$$\hat{\lambda}(\alpha) \; = \; \theta_1\hat{\lambda}(\alpha) + \theta_0. \tag{2.50}$$

Taking the derivative of the fixed-point condition with respect to α, we get

$$\frac{\partial\hat{\lambda}}{\partial\alpha} \; = \; \frac{\partial\theta_1}{\partial\alpha}\hat{\lambda}(\alpha) + \theta_1\frac{\partial\hat{\lambda}}{\partial\alpha} + \frac{\partial\theta_0}{\partial\alpha}. \tag{2.51}$$

Substituting (2.50) and (2.51), simplifies (2.49) to

$$h\left(1 - \theta_1 + (1-\alpha)\frac{\partial\theta_1}{\partial\alpha}\right). \tag{2.52}$$

We want to argue that (2.52) is negative. This is equivalent to showing that the term in brackets is positive as $h < 0$. From Step 3 of the proof of Proposition 2.5 we know that $(1-\alpha)\theta_1$ is increasing in α, i.e.,

$$-\theta_1 + (1-\alpha)\frac{\partial\theta_1}{\partial\alpha} \; \geq \; 0. \tag{2.53}$$

Thus (2.52) is negative and the distance between $\lambda'(\alpha)$ and $\lambda(\alpha)$ is decreasing in α for $h < 0$. The proof for $h > 0$ is analogous. □

Proof of Proposition 2.7.

We know from Lemma 2.2 that

$$V(\lambda_0, 1) \;=\; \frac{1}{1-\delta}\left(\frac{1+\lambda_0}{2}\right)^2,$$

and from Proposition 2.3 that

$$V(\lambda_0, \alpha) \;=\; v_2\lambda_0^2 + v_1\lambda_0 + v_0,$$

for $\alpha < 1$, where v_2 is given by (2.22) and v_1 by (2.25). Hence, $V(\lambda_0, 1) - V(\lambda_0, \alpha)$ is given by

$$\frac{1}{1-\delta}\left(\frac{1+\lambda_0}{2}\right)^2 - \left(v_2\lambda_0^2 + v_1\lambda_0 + v_0\right). \tag{2.54}$$

From here, we proceed in four steps.

Step 1: The difference $V(\lambda_0, 1) - V(\lambda_0, \alpha)$ is strictly convex, for all $\alpha < 1$.

In order to determine the curvature of (2.54), we consider the coefficient of λ_0^2, which is

$$\frac{1}{4(1-\delta)} - \frac{1 - \delta\alpha(2\alpha-1) - \sqrt{(1-\delta\alpha(2\alpha-1))^2 - \delta(1-\alpha)^2}}{2\delta(1-\alpha)^2}.$$

Rewriting fractions to have the same denominator yields

$$\frac{(1-\alpha)^2\delta - 2(1-\delta)(1-\delta\alpha(2\alpha-1)) + 2(1-\delta)\sqrt{(1-\delta\alpha(2\alpha-1))^2 - \delta(1-\alpha)^2}}{4\delta(1-\delta)(1-\alpha)^2}. \tag{2.55}$$

Note that the denominator of (2.55) is strictly positive for any $\alpha < 1$. Let

$$\begin{aligned} a \;&:=\; (1-\alpha)^2\delta - 2(1-\delta)(1-\delta\alpha(2\alpha-1)), \quad \text{and} \\ b \;&:=\; 2(1-\delta)\sqrt{(1-\delta\alpha(2\alpha-1))^2 - \delta(1-\alpha)^2}. \end{aligned}$$

The difference, (2.54), is strictly convex if and only if the numerator of (2.55) is strictly positive, that is, if $a + b > 0$. We start by considering the product $(a+b)(a-b) = a^2 - b^2$, that is

$$\left((1-\alpha)^2\delta - 2(1-\delta)(1-\delta\alpha(2\alpha-1))\right)^2 - 4(1-\delta)^2\left((1-\delta\alpha(2\alpha-1))^2 - \delta(1-\alpha)^2\right).$$

After some algebra, we obtain

$$\delta^2(1-\alpha)^2\Big[(1-\alpha)^2-(1-\delta)\big(4+4\alpha-8\alpha^2\big)\Big]. \tag{2.56}$$

The sign of (2.56) is determined by the sign of

$$(1-\alpha)^2-(1-\delta)\big(4+4\alpha-8\alpha^2\big). \tag{2.57}$$

As $4+4\alpha-8\alpha^2=4(1-\alpha^2)+4\alpha(1-\alpha)>0$, (2.57) is strictly increasing in δ. Consider (2.57) for $\delta=0$,

$$(1-\alpha)^2-\big(4+4\alpha-8\alpha^2\big) \;=\; -\big[3(1-\alpha^2)+6\alpha(1-\alpha)\big] \;<\; 0.$$

Showing that (2.57) is strictly negative for $\delta=0$ implies that $(a+b)(a-b)=a^2-b^2$ is strictly negative for all $\delta\in(0,1)$. Hence, either $a-b$ or $a+b$ must be negative but not both. Since $a+b>a-b$, we know that only $a-b$ can be negative and, hence, $a+b$ is positive which implies that (2.54) is strictly convex, as claimed.

Step 2: $V(1,1)$ is an upper bound for $V(\lambda_0,\alpha)$.

In period t, the seller's period profits from choosing $\theta_t^\dagger$ are given by

$$\big(1-\theta_t^\dagger\big)\big(\lambda_t+\theta_t^\dagger\big). \tag{2.58}$$

Note that (2.58) is strictly increasing in λ_t. Hence, obtaining the maximum of (2.58) for $\lambda_t=1$ in each period is an upper bound for $V(\lambda_0,\alpha)$. For the case where $\alpha=1$, we have $\lambda_t=\lambda_0=1$ for all $t\in\{0,1,\ldots\}$, and, therefore, the seller attains this upper bound by Lemma 2.2. For $\alpha<1$, the seller's value $V(\lambda_0,\alpha)$ is smaller than the upper bound which implies that (2.54) is positive for $\lambda_0=1$.

Step 3: For any $\alpha<1$, $V(0,1)$ is strictly smaller than $V(0,\alpha)$.

From Lemma 2.2, we know that $\theta_t^\dagger=1/2$, for all $t\in\{0,1,\ldots\}$, if $\alpha=1$ and $\lambda_0=0$. In each period, the seller thus realizes demand of $1/2$ at a constant price of $1/2$ which yields a value of

$$\frac{1}{4(1-\delta)}.$$

Now, consider the seller's profit for the case where $\alpha<1$ when she equally sets $\theta_t^\dagger=1/2$ for all $t\in\{0,1,\ldots\}$. Again, she realizes demand of $1/2$ in each period but at higher prices for $t\geq 1$ because $\alpha<1$ implies that $\lambda_t>0$ for $t\geq 1$. Consequently, $V(0,1)$ is strictly smaller than $V(0,\alpha)$ which implies that (2.54) is negative for $\lambda_0=0$.

Step 4: Existence of $\bar{\lambda}_0(\alpha)\in(0,1)$.

We know from the previous steps that (2.54) is strictly convex, strictly positive at $\lambda_0=1$, and strictly negative at $\lambda_0 = 0$. This implies that for any $\alpha < 1$, there exists a unique cutoff in $(0, 1)$. For any $\lambda_0 > \bar{\lambda}_0(\alpha)$, (2.54) is strictly positive, and hence $V(\lambda_0, 1) > V(\lambda_0, \alpha)$. Conversely, for any $\lambda_0 < \bar{\lambda}_0(\alpha)$, (2.54) is strictly negative, and hence $V(\lambda_0, 1) < V(\lambda_0, \alpha)$. For $\lambda_0 = \bar{\lambda}_0(\alpha)$, (2.54) equals zero which concludes the proof. □

3

On Reputation and Delivery Information in Competition between Online and Offline Retailers.

3.1 Introduction

In the last decade, e-commerce has developed into a significant factor in retail sales. In the United States, the fraction of e-commerce sales has more than doubled from 2.8 percent of all retail sales in March 2006 to 8.1 percent in June 2016[1]. E-commerce sales are growing by more than 20 percent per year and are expected to generate a revenue of almost 400 billion US-dollars in 2016[2]. Inevitably, this development has increased competition between online and offline retailers.

As one of the main differences between these competitors, retail stores and online shops engage with their customers differently. Buyers can physically purchase a good in the classical retail store whereas an online shop has to rely on a shipping service to deliver the goods to her buyers. This creates uncertainty for customers of an online shop as to their overall waiting time for the good, which does not only depend on the online shop's logistic capabilities but also on the ability of the assigned shipper to deliver the parcel quickly. On the one hand, the growing importance of e-commerce indicates that consumers accept waiting times for a good when the online shop compensates them with a lower price or spares them the need to travel to a retail store. In a study for the European Commission of the state of play of EU parcel markets, Okholm et al. (2013) are also concerned with the preferences of customers of online shops. They find that one of "*the most important delivery aspects for e-shoppers are (...) access to electronic delivery notifications and track and trace (...)*." On the other hand, an online seller depends on the reliability of a shipping service to finalize her transactions in

1 *YCharts*, "US E-Commerce Sales as Percent of Retail Sales'".

2 *Statista*, "Annual retail e-commerce sales growth worldwide from 2014 to 2019", and "Retail e-commerce sales in the United States from 2013 to 2019".

a time frame that is satisfying to consumers. As a result, in online transactions, the shipper has a significant impact on the buyer's shopping experience and, in turn, on the seller's reputation. The US digital media company *Business.com* highlights the importance of managing shipments in an article about shipping services advising small businesses: "*The quality and reliability of your shipping service can make or break your customer's service experience, reflecting not just on the shipper, but directly on your company.*[3]". Similarly, while the online community *Fit Small Business* recommends one specific US shipper as "*they have by far the best rates*", they also recognize the importance of a swift delivery for customer satisfaction: "*Many online complaints focus on packages that are either delivered late or go missing. If you are using insured services, you will be reimbursed but your customers may not be happy when deliveries are late or do not arrive at all.*".[4]
It is, thus, interesting to examine from an theoretical economic perspective how an online seller who competes with an offline retailer manages the assignment of shipping services optimally, and which information about the delivery process is in fact beneficial to consumers.

In this paper, we analyze the dynamic reputation problem of a long-lived online retailer. She competes with an offline retailer in an adapted two-period Hotelling competition model with fixed prices. Short-lived buyers are located on a unit interval and have to travel to the offline retailer to purchase a good. The online seller, however, can only ship goods to her buyers by choosing a shipper which is available in two qualities at different rates. The buyers' waiting time depends on two components. In the first step, the seller provides the good to a shipper, and in the second step, the shipper delivers the good. The seller's type, which can be either high or low, and is unknown to buyers, determines stochastically the speed of the first step, and the shipper's quality determines stochastically the speed of the second step. Thus, the buyers' overall waiting time depends on the type of the seller as well as the quality of the chosen shipping service. In our paper, reputation represents the buyers' belief about the seller's type.
We consider four specifications of the model in which the information available to buyers about the delivery process changes. In particular, we differentiate cases where buyers can or cannot *track and trace* the delivery of their ordered good, as well as cases where buyers can or cannot distinguish which shipper quality the seller chose when the good is delivered.
We start by analyzing a benchmark model in which buyers only observe their overall waiting time and, hence, draw conclusions about the seller's type solely from this information. We show that the seller types pool on the cheaper low quality shipper if reputation is relatively low or relatively high, that is, if buyers are sufficiently convinced that the seller is of the low type or of the high type, respectively. If her reputation is in an intermediate range, however, that is, buyers are relatively uncertain about the seller's type, a separating equilibrium exists,

[3] See Wood (2015).
[4] See Marsan (2015).

where only the high type seller chooses the more expensive, high quality shipper. When buyers are uncertain, reputation reacts most sensitively to their realized waiting time. Increasing the probability of a quick delivery and decreasing the probability of a slow delivery by choosing the high quality shipper is hence most profitable. In addition to higher expected increase in reputation from a quick delivery, the high type seller protects herself from being falsely identified as a low type from a slow delivery.

Subsequently, we change the setup with respect to which additional pieces of information buyers observe and derive the equilibria for each setup. When buyers can observe the shipper's quality upon delivery of the good, we find that pooling on the low quality shipper is always an equilibrium. Pooling on the high quality shipper, however, exists as an equilibrium if and only if the seller's reputation is sufficiently high. Deviating from the pooling equilibrium on the high quality shipper saves a seller additional costs but buyers assign an off-path belief that she is of the high type after detecting a deviation. If the seller's reputation is already low, a low off-path belief cannot reduce her profits enough in relation to the saved costs, and hence, a deviation can be profitable. With observable shipper qualities, there exist no separating equilibria. In any separating equilibrium, buyers learn the seller's type perfectly from observing the shipper's quality upon delivery of the good. If the benefits from being identified as one shipper type outweigh the additional costs, both types would choose the same respective shipper.

The ability to track the delivery process enables buyers to disentangle how the seller and the shipper contributed to their total waiting time. Observing the speed of the first delivery step is then a signal about the seller's type only whereas the speed of the second delivery step is a signal about the shipper's quality which, in turn, may be an indirect signal about the seller's type as well, depending on the equilibrium strategy. When such a tracking technology is available to buyers, pooling on the low quality shipper remains an equilibrium. In a pooling equilibrium with tracking, the signal about the shipper's step is non-informative and the only signal that buyers learn from is the speed of the seller's step. As the speed of this step depends on the seller's fixed type only, a seller cannot influence the buyers' information with her shipper choice. Thus, seller's always prefer the cheaper, low quality shipper which implies that pooling on the high quality shipper can never be an equilibrium when buyers can track their delivery but are not able to observe the shipper's quality upon. The separating equilibrium can only exists in this setup if the seller's reputation is low, and for a specific range of intermediate cost parameters.

Comparing the equilibria of these setups, we find that the seller's optimal strategy differs significantly with respect to the information that buyers obtain about the delivery process. The incentives for the seller to assign a shipper of higher or lower quality depend on which pieces of information buyers take into account when they draw conclusions about the seller's type. In particular, incentives for the high type seller to choose the high quality shipper decrease when buyers can track the delivery and, hence, learn about her high type through the tracking technology instead of only through the noisier overall waiting time. As a result, the high type seller might choose the high quality shipper only if reputation is very low such that from the consumers' perspective, the ability to track parcels decreases their expected welfare. The results are ambiguous, however, when we consider the effect on the buyers' welfare from their ability to identify the shipper's quality upon delivery of the ordered good. It harms their welfare if the seller has comparatively low reputation, but can improve their welfare if reputation is sufficiently high.

In an online transaction, we assume that not the buyer but the seller pays the shipping fee. In recent years, larger online shops have typically absorbed these costs from their customers which provides a justification for this assumption[5]. It is also worth noting, that we only consider the effect of a tracking technology on reputation. The ability to track a parcel can simplify taking delivery of the shipper, as some services provide, for example, an estimated arrival time of the parcel. Online sellers might also have a preference to track their delivery for insurance and warranty related aspects. These additional channels though which tracking might have positive effect on the consumers' welfare are not captured in our model.

The growing importance of e-commerce has fostered empirical literature on online sellers' strategies, on consumer behavior in online markets, and on the importance of a seller's reputation on her sales. Dinlersoz and Li (2006) find that internet book retailers strategically adjust their strategies for shipping prices and average delivery time to their market position. In online auction environments, Melnik and Alm (2002) and Livingston (2005) show empirically that sellers with better reputation are able to generate higher revenues. Recently, Jolivet et al. (2016) confirm that this empirical relation can be found on online e-commerce platforms as well. Hence, the empirical literature suggests that reputation is an important determinant for success in online markets and sellers are, to a certain extent, aware of the strategic implications of shipping speed on their reputation.

In these papers, however, reputation relates to the consumers' belief that that the seller behaves trustworthy and does in fact deliver a good after receiving payment. In contrast, in our paper reputation refers to an online seller's unknown logistical ability. This is reminiscent of the classical dynamic reputation literature where buyers observe noisy signals and update their belief about the seller's unknown type accordingly. See, for example, Bar-Isaac and Tadelis (2008) for an overview. To see similarities, it is instructive to reframe the model slightly. We can interpret the seller's logistical ability as her type and the chosen shipper quality as her effort. The realized waiting time corresponds to a signal generated by the seller's type and effort.

[5] According to an article from *Business Insider* from May 27, 2014, the fraction of transactions with free shipping has increased from 48 percent in the first quarter of 2013 to 58 percent in the first quarter of 2014, see Smith (2014).

In Kreps et al. (1982), Mailath and Samuelson (2001), Cripps et al. (2004), and Cripps et al. (2007), reputation reflects the market's belief that the seller of a competent type instead of a behavioral type. Whereas the latter always chooses the same effort, a competent type chooses her effort strategically. Buyers observe noisy signals whose distributions depend on the seller's effort, and derive utility not from the seller's type directly but from the signal realization. Kreps et al. (1982) introduce behavioral types to explain cooperation in finitely repeated prisoner's dilemma through reputation effects. Mailath and Samuelson (2001) find that there is no equilibrium in which a competent seller always exerts high effort. When firms are able leave the market and new firms can acquire their name (and reputation) unobserved, competent types prefer to acquire average reputations whereas good reputations are most valuable to behavioral types. Cripps et al. (2004), and Cripps et al. (2007) show that reputation is a short-run phenomenon when types are fixed as buyers learn about the seller's type over time.

In Holmström (1999), Board and Meyer-ter Vehn (2013), and Dilmé (2016), the seller's type is either high or low and reputation corresponds to the market's belief that the seller is of the high type. In Holmström (1999), the seller's type is unknown to herself as well as to the market. Buyers gain utility from a noisy signal whose distribution depends on the seller's type and her exerted effort. He shows that the concern for reputation can increase effort in the first periods but as the market learns about the seller's type, incentives to exert effort diminish. In Board and Meyer-ter Vehn (2013), and Dilmé (2016), buyers gain utility from the type of a seller at the time of purchasing her good. The seller's type is her private information and changes dynamically over time. Board and Meyer-ter Vehn (2013) examine a model in which buyers observe and learn from signals whose realization depends on the seller's type. They show that depending on whether the market learns through good or bad news signals, the seller exerts effort if reputation is low or high, respectively. In Dilmé (2016), the seller can strategically choose her type in each period but suffers costs from switching. He shows that a seller changes her type only if reputation is either very high or very low.

Our setup is related more closely to Holmström (1999). Buyers gain utility from the signal whose realization depends on both the seller's type as well as her effort. The seller, however, is aware of her type and exerts effort strategically to manage her reputation. With our focus on the information structure, the research question of this paper is more related to Board and Meyer-ter Vehn (2013). We put emphasis on the effects of changing information that is available to buyers by explicitly modeling how the signal is generated from two independent signals about the seller's type, and about her effort. We use this structure to examine how transparency about the generation of the signal influences the seller's strategic problem, and, in turn, buyers' welfare.

The paper is organized as follows. We start in Section 3.2 by introducing the adapted two-period Hotelling competition model, and solve the buyer's problem for each period as well as the seller's problem in the final second period. In Section 3.3, we analyze the seller's problem in the first period for a benchmark case where buyers can only observe their realized waiting time. In Section 3.4, we allow buyers to additionally observe the shipper's quality upon delivery of the good. In Section 3.5, buyers track their delivery and are thus able to observe the contribution of the seller and the shipper to their waiting time but are no longer able to observe the shipper's quality. In Section 3.6, we combine the two previous cases and allow for both tracking as well as observable shipper quality. Finally, in Section 3.7 we compare the derived equilibria of all four setups from the buyers' perspective. Section 3.8 concludes. All proofs are relegated to the Appendix.

3.2 Model

Time is discrete and there are two periods $t \in \{1, 2\}$. In each period a unit mass of short-lived buyers can purchase a good at a fixed price of 1 which provides them with net utility of u.[6] Two long-lived sellers provide the good at zero production costs. Buyers are located on a unit interval and their position $\eta \in [0, 1]$ is uniformly distributed. One of the two sellers, retailer R, is located at $\eta = 0$. If a buyer from position η chooses to purchase the good from R, he suffers costs η from the distance he has to travel to R. Formally, a buyer's utility from buying from R and traveling distance η is

$$U^R(\eta) \; = \; u - \eta. \tag{3.1}$$

The second seller, online seller S (in the following also referred to as *seller*), is of an unknown type $\theta \in \{L, H\}$. If a buyer purchases from S in period t, the seller needs to ship the good to the buyer. In order to do so, she assigns a shipper which is available in a low and a high quality, $d_t \in \{l, h\}$. Shipping costs are $c(d_t)$, where $c(h) > c(l)$, and we denote the cost difference by $\Delta \equiv c(h) - c(l)$. Buyers who purchase the good from S in period t suffer losses from their experienced waiting time $w_t \in \{0, 1\}$ which is drawn from a probability distribution depending on both the seller's type and the shipper's quality. Formally, a buyer's utility from buying from S and waiting w_t is given by

$$U^S(w_t) \; = \; u - w_t. \tag{3.2}$$

Note that a buyer who is located at $\eta = 0$ is always better off by purchasing the good from R whereas a buyer located at $\eta = 1$ always prefers to order from S. We assume

$$u \; > \; 1 \tag{3.3}$$

[6] As the seller provides the good at zero costs, the price only serves as a multiplier in her profits. Normalizing the fixed price to 1 simplifies notation and is without loss of generality.

so that not buying the good is a dominated strategy for all consumers and we only have to take care of how the total demand is split between R and S.

In each period $t \in \{1, 2\}$, shipping proceeds in two steps. First, the seller processes the order and provides the good to a shipper. In the second step, that shipper delivers the good to the buyers. Both steps can be processed quickly or slowly. Denote the speed of the first step in period t by $w_t^1 \in \{0, 1\}$ and of the second step by $w_t^2 \in \{0, 1\}$. We interpret a realization of 0 as a quickly processed step and a realization of 1 as a delayed step. We assume that the overall waiting time in period t is long ($w_t = 1$) whenever there is delay in at least one step. Only if both steps proceeded quickly, the overall waiting time is short ($w_t = 0$)[7]. Formally,

$$w_t = \max\left\{w_t^1, w_t^2\right\}. \tag{3.4}$$

We denote the probability that the high and the low type seller process the first step in period t quickly by p and r respectively, i.e.

$$\Pr\left[w_t^1 = 0|\theta = H\right] = p,$$
$$\Pr\left[w_t^1 = 0|\theta = L\right] = r,$$

where $p > r$. Note that the seller's ability to manage the first step quickly depends only on her type but not on the chosen shipper quality. For the shipper, the probability of managing the second step quickly depends on its quality but is independent of the seller's type[8]. We denote the probability for no delay in the second step in period t from choosing the high and the low shipper quality by q^h and q^l, respectively, i.e.

$$\Pr\left[w_t^2 = 0|d_t = h\right] = q^h,$$
$$\Pr\left[w_t^2 = 0|d_t = l\right] = q^l,$$

where $q^h > q^l$. Given (3.4), we can derive the probability of $w_t = 0$ for both seller types from

---

7 One can also think of the opposite case and even intermediate cases. It turns out, however, that this assumption does not have an effect on the existence of equilibria, but only affects demand in equilibrium.

8 The ability may not only depend on how quickly a shipper can bring the good to a buyer's home. One can also include its infrastructure allowing buyers to pick up a parcel from a nearby store if they were not at home at the time of delivery. As buyers care about the waiting time until they actually receive the good, net shipping time is important but not the only factor that determines the total waiting time.

Buyers' information	only w_t observable	w_t^1 and w_t^2 observable
d_t not observable	Section 3.3	Section 3.5
d_t observable	Section 3.4	Section 3.6

Table 3.1: Overview of the differences in the buyers' information in Sections 3.3 through 3.6.

choosing a shipper quality as

$$\begin{aligned}
\Pr[w_t &= 0|\theta = H, d_t = h] = p \cdot q^h, \\
\Pr[w_t &= 0|\theta = H, d_t = l] = p \cdot q^l, \\
\Pr[w_t &= 0|\theta = H, d_t = h] = p \cdot q^h, \\
\Pr[w_t &= 0|\theta = L, d_t = l] = r \cdot q^l,
\end{aligned}$$

with the corresponding counter-probabilities for $w_t = 1$. We assume that

$$pq^h > \frac{1}{2} > rq^l. \tag{3.5}$$

In each period t, buyers form a belief that S is of type $\theta = H$ which we denote by μ_t. Their common prior is given by $\mu_1 \in (0, 1)$. We will also refer to μ_t as the seller's *reputation* in period t in the following.

For tractability, we assume that the realization of w_t is perfectly correlated for all buyers who purchase the good from S and perfectly observable even for R's buyers. This implies that all buyers observe the same history of play and, hence, update their common prior belief identically.

We analyze the model in four specifications in the following. Buyers always observe w_t but may or may not observe w_t^1, w_t^2, and d_t. In each section, we change the setup with respect to the information that buyers observe about the previous periods in addition to w_t, as depicted in Table 3.1. We denote the information that buyers can observe at the beginning of each period $t \in \{1, 2\}$ about the previous history of play by the subset $h^t \subseteq \left(w_s, w_s^1, w_s^2, d_s\right)_{s=1}^{t-1}$, where $h^1 = \emptyset$.

The timing of the stage game in period t is as follows. First, buyers observe h^t and assign probability μ_t that the seller is of the high type. Then, they decide whether to buy the good from S, R or not at all. Buyers at η who buy from R receive their utility $U^R(\eta)$ immediately. Second, seller S realizes her demand, chooses shipper quality d_t, and pays the according fee $c(d_t)$. Third, the random variables w_t^1 and w_t^2 realize, and w_t is determined. Buyers who bought from S receive utility $U^S(w_t)$, and, finally, all buyers leave the market.

3.2.1 Strategies and Equilibrium

For each respective h^t and any prior belief $\mu_1 \in (0,1)$, a buyer's strategy specifies for each period $t \in \{1,2\}$ whether to buy the good from R, S, or not at all (N) as a function of his location η on the unit interval as well as the seller's reputation μ_t. We denote the buyers' strategy by $\beta = \left(\beta_1^\eta, \beta_2^\eta\right)_{\eta \in [0,1]}$, where $\beta_t^\eta \in \{S, R, N\}$.
The seller's strategy specifies whether to choose the high quality shipper or the low quality shipper as a function of her type and her reputation μ_t for each period $t \in \{1,2\}$. In order to simplify writing down profits, we denote a strategy in period t in terms of its implied probability of realizing $w_t^2 = 0$. A seller's strategy is hence given by $\sigma = \left(\sigma_1^\theta, \sigma_2^\theta\right)_{\theta \in \{H,L\}}$, where $\sigma_t^\theta \in \left\{q^l, q^h\right\}$. For example, $\sigma_t^\theta = q^h$ denotes a strategy where seller type θ plays $d = h$ in period t. Analogously, we rewrite the shipping fees for the high and the low quality shipper to $c(q^h)$ and $c(q^l)$, respectively.

We consider Markov Perfect Equilibria (MPE) in pure strategies. In the following, we refer to a profile of strategies and beliefs that constitutes a MPE in pure strategies as *equilibrium*.

3.2.2 Consumers' Behavior and Demand

Before we analyze the seller's problem, we derive the equilibrium behavior of the buyers for both periods. Fix some reputation μ_t and a seller's strategy σ. Consider the problem of a buyer who is located at position $\eta \in [0,1]$. Following a markov strategy, he conditions his decision on μ_t only and buys from S only if

$$\begin{aligned} \mathbb{E}\left[U^S(w_t)|\mu_t, \sigma\right] \;&\geq\; U^R(\eta) \\ \Leftrightarrow\; \eta \;&\geq\; \mathbb{E}[w_t|\mu_t, \sigma] \;\equiv\; \overline{\eta}(\mu_t), \end{aligned}$$

and from R otherwise, where the cutoff is

$$\begin{aligned} \overline{\eta}(\mu_t) &= \mathbb{E}[w_t|\mu_t, \sigma] \\ &= \mu_t\left(p\sigma_t^H \cdot 0 + (1 - p\sigma_t^H) \cdot 1\right) + (1-\mu_t)\left(r\sigma_t^L \cdot 0 + (1 - r\sigma_t^L) \cdot 1\right) \\ &= 1 - r\sigma_t^L - \mu_t\left(p\sigma_t^H - r\sigma_t^L\right). \end{aligned} \tag{3.6}$$

The following result summarizes the buyers' equilibrium strategies derived above.

Proposition 3.1. *For any μ_t and a seller's equilibrium strategy $\sigma^* = \left(\sigma_1^\theta, \sigma_2^\theta\right)_{\theta \in \{H,L\}}$, buyers follow a cutoff strategy and buy from S in period t only if $\eta \geq \overline{\eta}(\mu_t)$, where*

$$\overline{\eta}(\mu_t) \;=\; 1 - r\sigma_t^L - \mu_t\left(p\sigma_t^H - r\sigma_t^L\right),$$

and buy from R otherwise.

Given the buyers' equilibrium strategy β, a seller's equilibrium strategy σ^* and the uniform distribution of η on the unit interval, in period t a seller with reputation μ_t faces demand of

$$\begin{aligned} D_t(\mu_t, \sigma^*) &= 1 - \overline{\eta}(\mu_t) \\ &= r\sigma_t^L + \mu_t \left(p\sigma_t^H - r\sigma_t^L\right). \end{aligned} \tag{3.7}$$

The seller's objective is to maximize the sum of expected profits with respect to her actually played strategy σ,

$$\max_{\sigma} \left\{ \sum_{t=1}^{2} \mathbb{E}\left[D_t(\mu_t, \sigma^*) - c(\sigma_t)\right] \right\}.$$

In the following, we focus on the seller's problem.

3.2.3 Equilibrium behavior in $t = 2$

We solve this game of finite horizon by backward induction and, thus, start by analyzing the second period. In $t = 2$, the seller always chooses the low quality shipper and buyers act according to their cutoff equilibrium strategy. This result holds independently of which pieces of information buyers receive, and is hence valid for all of the following sections.

Lemma 3.1. *In any equilibrium,* $\sigma_2^H = \sigma_2^L = q^l$.

Consequently, for any $\mu_2 \in (0, 1)$, the seller's demand in $t = 2$ in any equilibrium is given by

$$D_2(\mu_2, \sigma_2) = rq^l + \mu_2 q^l (p - r). \tag{3.8}$$

Intuitively, in the final period of the game, the seller cannot credibly commit to the high quality shipper. Suppose there exists an equilibrium in which she plays $d_2 = h$. After buyers order the good from S, she can deviate to the low quality shipper and save the cost difference Δ. As the game terminates after the delivery of the goods at the end of the second period, there are neither negative effects on reputation nor are the buyers able to punish her in a following period. Consequently, that deviation is always profitable and playing $d_2 = h$ cannot be part of any equilibrium strategy.

In the following, we determine the seller's behavior in the first period in equilibrium. Focusing on $t = 1$, we change notation slightly and denote a seller's equilibrium strategy for $t = 1$ by $\sigma^* = \left(\sigma_H^*, \sigma_L^*\right)$ and a sellers actually played strategy by $\sigma = (\sigma_H, \sigma_L)$, where we continue to denote strategies by the induced probability of realizing $w_1^2 = 0$. Additionally, we drop the time indices from the buyers waiting times, and, hence, denote w_1 by w, w_1^1 by w^1 and w_1^2 by w^2 in the following.
With the above changes in notation, we rewrite the seller's profits in $t = 2$ in any equilibrium

to

$$\pi_2(\sigma^*;\mu_2) = rq^l + \mu_2 q^l(p-r) - c(q^l). \tag{3.9}$$

We can now turn to analyze the seller's problem in the first period.

3.3 Benchmark: Buyers observe w only

In order to set a benchmark we first consider the case where buyers have no access to any information except for the realization of w. Thus, buyers observe $h^t = (w_s)_{s=1}^{t-1}$ at the beginning of each period t. Neither can they distinguish the shipper's type nor observe w^1 and w^2, but only draw conclusions about the seller's type from the realization of w. This implies for the seller that deviations from her equilibrium strategy cannot be perfectly detected.
Unlike in the second period, the seller might have an incentive to choose the higher shipper quality in $t=1$. She trades off the value of a higher probability of realizing $w=0$ against the additional costs of Δ. Buyers update their beliefs based on the realization of w only. Depending on the equilibrium strategy, the initial reputation μ_1 as well as the parameter values of p and r, a realization of $w=0$ can be a strong signal for $\theta=H$ and, thus, may increase demand and profits in $t=2$ considerably. If in such a case the cost difference between both shipper qualities is comparatively small, the seller can choose the high quality shipper in equilibrium in the first period.

We denote by $\mu_2(w|\sigma^*)$ the buyers' updated beliefs according to Bayes' rule as a function of the realization $w \in \{0,1\}$ as well as the seller's equilibrium strategy in $t=1$, σ^*,

$$\begin{aligned}\mu_2(w = 0|\mu_1,\sigma^*) &= \frac{\mu_1 p\sigma_H^*}{\mu_1 p\sigma_H^* + (1-\mu_1)r\sigma_L^*},\\ \mu_2(w = 1|\mu_1,\sigma^*) &= \frac{\mu_1(1-p\sigma_H^*)}{\mu_1(1-p\sigma_H^*)+(1-\mu_1)(1-r\sigma_L^*)}.\end{aligned} \tag{3.10}$$

With (3.9), we can formulate the seller's total profits as a function of her equilibrium strategy in $t=1$, σ^*, her played strategy in $t=1$, σ, and her prior reputation, μ_1. For the high type seller, we have

$$\begin{aligned}\pi^H(\sigma,\sigma^*;\mu_1) = & rq^l + \mu\left(p\sigma_H^* - r\sigma_L^*\right) - c(\sigma_H) + p\sigma_H \cdot \pi_2(\mu_2(w = 0|\mu_1,\sigma^*))\\ & + (1-p\sigma_H)\cdot\pi_2(\mu_2(w = 1|\mu_1,\sigma^*)),\end{aligned} \tag{3.11}$$

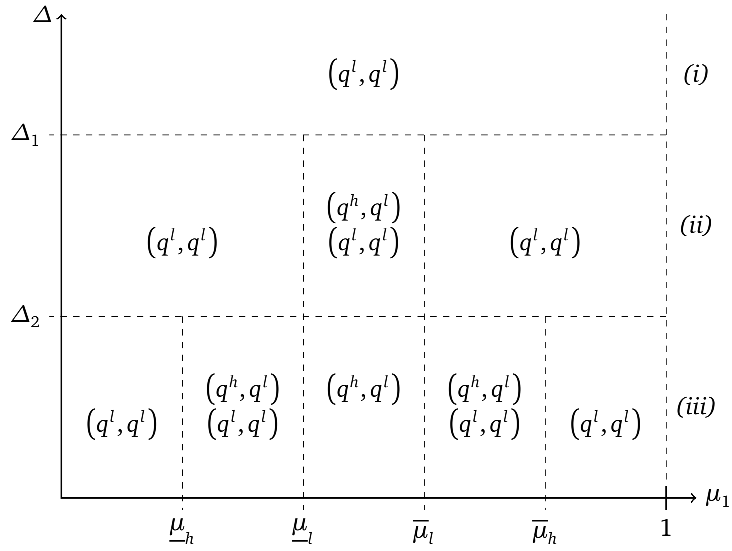

Figure 3.1: Overview of the benchmark equilibria for the cases distinguished in Proposition 3.3.

and for the low type seller, we get

$$\pi^L(\sigma,\sigma^*;\mu_1) = rq^l + \mu\left(p\sigma_H^* - r\sigma_L^*\right) - c(\sigma_L) + r\sigma_L \cdot \pi_2(\mu_2(w = 0|\mu_1,\sigma^*)) + (1 - r\sigma_L) \cdot \pi_2(\mu_2(w = 1|\mu_1,\sigma^*)). \tag{3.12}$$

We are now in a position to derive the equilibrium strategies for $t = 1$. First, we show that $\theta = L$ always chooses the low quality shipper in the first period when r is sufficiently small.

Proposition 3.2. *There exists a $\hat{r} \in (0,1)$ such that for $r < \hat{r}$, $\sigma_L = q^h$ is a strictly dominated strategy in $t = 1$.*

Intuitively, the low type seller must decide for each value of μ_1 whether investing Δ into the high quality shipper is worth the expected increase in reputation for the next period. If r is small enough, a realization of $w = 0$ is a relatively precise signal about the seller's type. The probability of realizing $w = 0$ for the low type seller, however, is small, so that in expectation reputation only improves very little compared to the constant additional costs. Hence, the investment is not worth its costs. From now on, we assume $r < \hat{r}$ such that Proposition 3.2 holds. This allows us to fix $\sigma_L^* = q^l$ for any equilibrium in the benchmark.
Second, we characterize the solution to the problem of the high type seller.

Proposition 3.3. *There are two cutoffs on Δ, $\Delta_2 < \Delta_1$, and four cutoffs on μ_1, $\underline{\mu}_h < \underline{\mu}_l < \overline{\mu}_l < \overline{\mu}_h$, such that*

(i) for $\Delta > \Delta_1$, $\sigma^* = \left(q^l, q^l\right)$ is the unique equilibrium for any $\mu_1 \in (0,1)$.

(ii) for $\Delta_1 \geq \Delta \geq \Delta_2$, $\sigma^* = \left(q^l, q^l\right)$ is an equilibrium for any $\mu_1 \in (0,1)$, and $\sigma^* = \left(q^h, q^l\right)$ is an equilibrium for $\mu_1 \in [\underline{\mu}_l, \overline{\mu}_l]$.

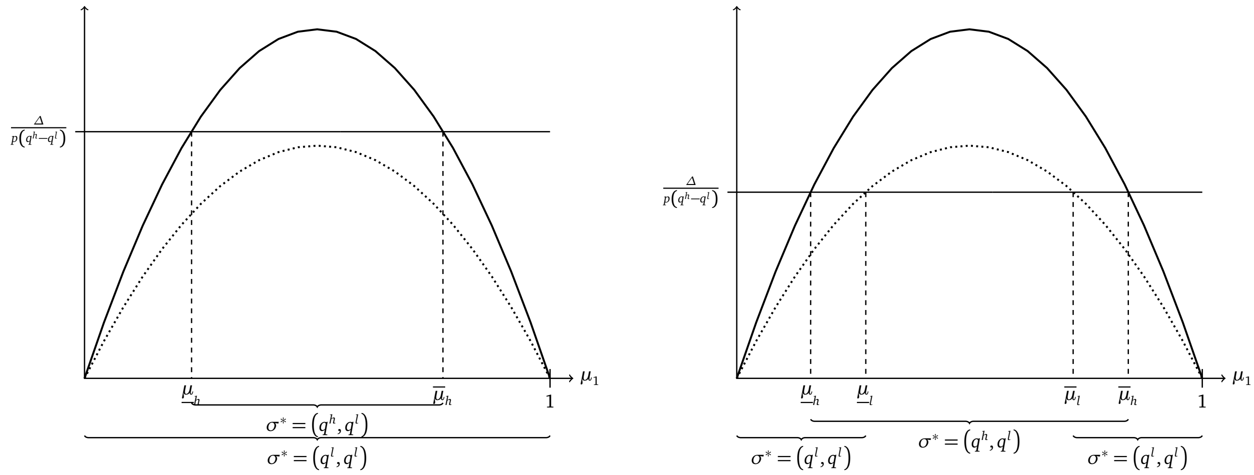

(a) (b)

Figure 3.2: The graphs display expected benefits from $d_1 = h$ if the equilibrium is $\sigma^* = (q^h, q^l)$ (solid black line) and $\sigma^* = (q^l, q^l)$ (dotted black line). See the case where $\Delta_1 > \Delta > \Delta_2$ on the left (a), and the case where $\Delta < \Delta_2$ on the right (b).

(iii) for $\Delta_2 > \Delta$,

$$\sigma^* \in \begin{cases} \{(q^l, q^l)\} & \textit{for } \mu_1 \in [0, \underline{\mu}_h) \cup (\overline{\mu}_h, 1] \\ \{(q^h, q^l)\} & \textit{for } \mu_1 \in (\underline{\mu}_l, \overline{\mu}_l) \\ \{(q^h, q^l), (q^l, q^l)\} & \textit{for } \mu_1 \in [\underline{\mu}_h, \underline{\mu}_l] \cup [\overline{\mu}_l, \overline{\mu}_h]. \end{cases}$$

Figure 3.2 shows a graphical representation of the seller's problem. She trades off the additional costs against the expected additional profits from hiring the high quality seller instead of the low quality seller. While the costs are constant in μ_1, the gains are strictly concave and higher if the seller plays $d = h$ in her equilibrium strategy. This is because the realization of w is more indicative of the seller's type if both types separate in their shipper choice, and hence, the expected additional profits from increasing the probability of $w = 0$ through the high quality shipper are higher.

When Δ is high enough, i.e., $\Delta > \Delta_1$, the additional costs for the high quality shipper cannot be compensated by the expected additional profits in the next period and both seller types always choose the low quality shipper in the first period.

When $\Delta_1 \geq \Delta \geq \Delta_2$ and the equilibrium strategy is $\sigma^* = (q^l, q^l)$, neither type has an incentive to deviate because the additional expected profits from the high quality shipper are relatively small whereas the additional costs are relatively high. As the costs are smaller than Δ_1, however, $\sigma^* = (q^h, q^l)$ can be an equilibrium in an intermediate range of μ_1 where the expected additional profits from the high quality shipper are highest. That is the case because reputation changes slowly if it is relatively small or relatively high. Buyer are already convinced that the seller is of the low or the high type, respectively, and the realized waiting

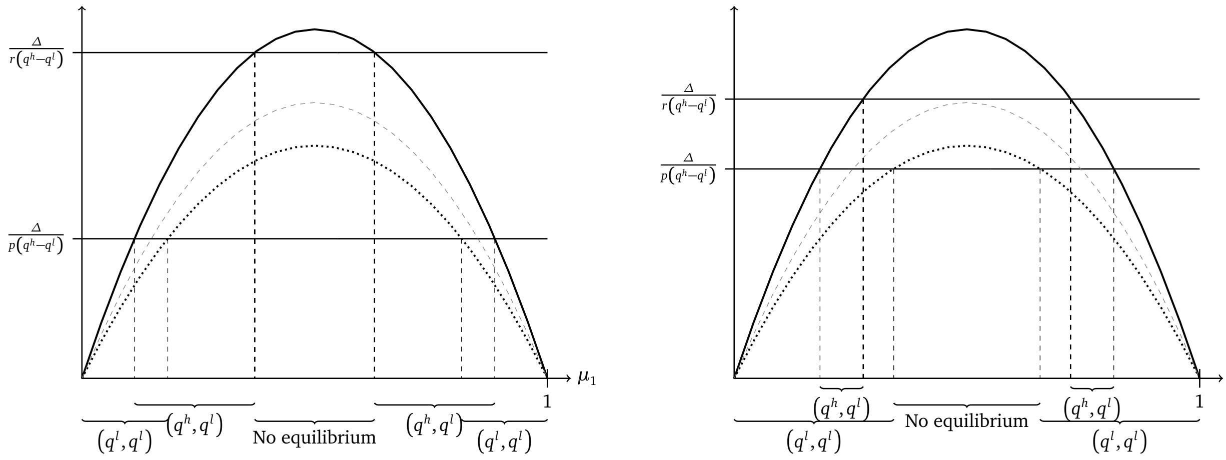

(a) (b)

Figure 3.3: The graphs display expected benefits from $d_1 = h$ if the equilibrium is $\sigma^* = (q^h, q^l)$ (solid black line), $\sigma^* = (q^h, q^h)$ (dashed grey line), and $\sigma^* = (q^l, q^l)$ (dotted black line). On the left (a), there is no equilibrium (in pure strategies) for $\mu_1 \in I$. On the right (b), there is no equilibrium (in pure strategies) for $\mu_1 \in (\underline{\mu}_l, \overline{\mu}_l)$.

time has little impact on their beliefs. If buyers are uncertain about the seller's type, however, that is if reputation is in an intermediate range, the buyers' waiting time has a higher impact on the reputation of the next period. As a result, reducing waiting time in expectation by hiring the high quality shipper is most valuable. See Figure 3.2a for a graphical representation of this case.

Finally, when Δ is small enough, i.e., $\Delta < \Delta_2$, $\sigma^* = (q^l, q^l)$ does not constitute an equilibrium for all $\mu_1 \in (0,1)$ anymore. For intermediate levels of reputation, the additional expected profits cover its additional costs of the high quality shipper even if the equilibrium is $\sigma^* = (q^l, q^l)$ and w is a less precise indicator for the seller's type. See Figure 3.2b for a graphical representation of this case.

Proposition 3.3 suggests that the seller's problem becomes trivial if $\Delta > \Delta_1$. The cost difference is so high that hiring the high quality shipper never pays off in expectation. In order to focus on those cases where the high type seller faces a meaningful problem in the benchmark, we assume that $\Delta < \Delta_1$ in the following.

The restriction to $r < \hat{r}$ rules out that the low type seller chooses the high quality shipper in equilibrium which makes the description of the seller's equilibrium strategy more convenient. Moreover, this assumption guarantees that the equilibrium characterization is complete in the sense that there always exists an equilibrium (in pure strategies) for all $\mu_1 \in (0,1)$. To see this, recall that $r < \hat{r}$ implies that for any $\mu_1 \in (0,1)$ the low type seller prefers $d = l$ over $d = h$, even in the equilibrium $\sigma^* = (q^h, q^l)$ where $d_1 = h$ is most profitable. Now, suppose to the contrary that the low type seller prefers $d_1 = h$ for μ_1 in some intermediate interval

$I \subset (0,1)$ but always prefers $d_1 = l$ if the equilibrium is $\sigma^* = \left(q^h, q^l\right)$. In other words, $\sigma^* = \left(q^h, q^l\right)$ cannot be an equilibrium for $\mu_1 \in I$, because $\theta = L$ deviates, and $\sigma^* = \left(q^h, q^h\right)$ can never be an equilibrium because $\theta = L$ deviates for any $\mu_1 \in (0,1)$. If, additionally, $\Delta < \Delta_2$ for example, such that $\sigma^* = \left(q^l, q^l\right)$ is not an equilibrium for all $\mu_1 \in (0,1)$, there exists an intermediate interval where no equilibrium (in pure strategies) exists. See Figures 3.3a and 3.3b for a graphical representation of two such cases. These cases are ruled out by assuming $r < \hat{r}$. Hence, for any $\mu_1 \in (0,1)$ either $\sigma^* = \left(q^h, q^l\right)$ or $\sigma^* = \left(q^l, q^l\right)$ constitutes an equilibrium.

3.4 Labeled Shippers: Buyers observe d and w

In the benchmark, only the seller has information about the performances of both shippers with respect to their expected speed and is able to rank them accordingly. In this section, in contrast, the chosen shipper's quality d is perfectly observable for buyers. Hence, buyers observe $h^t = (w_s, d_s)_{s=1}^{t-1}$ at the beginning of each period t. We can interpret this case as one where shippers carry observable labels and buyers correctly infer the shipper's quality from them upon delivery. As a result, information asymmetries about the quality of the hired shipper between the seller and the buyers vanish. Upon delivery of the good, a buyer now realizes whether the seller assigned the high or the low quality shipper. Consequently, buyers can detect deviations of the seller from her equilibrium strategy perfectly. A newly released product test is an example for information that allows buyers to connect shipper labels with qualities. In the following, we will refer to the shippers whose quality buyers are able to observe as *labeled shippers*.

In contrast to the benchmark, the game dynamic becomes a signaling game with observable actions. In this setup we have to keep track of two beliefs instead of one because buyers can now distinguish between the subgame after observing $d_1 = h$ and the subgame after $d_1 = l$. Accordingly, we denote the buyers' beliefs that the seller is of the high type after observing $d_1 \in \{l, h\}$ by $\mu_2(d_1)$.

First, we investigate potential separating equilibria and show that these cannot exist. If the seller plays a separating strategy, her observable choice of d_1 is a perfect signal for her type. Buyers ignore the noisy signal w and learn the seller's type perfectly from observing d_1. For each separating equilibrium strategy, $\sigma^* = \left(q^h, q^l\right)$ and $\sigma^* = \left(q^l, q^h\right)$, both types, in fact, face the same trade-off which they cannot optimally resolve in different ways.

Proposition 3.4. *There exists no equilibrium where the seller types separate.*

Intuitively, in both potential separating equilibria, one shipper quality is associated with the high seller type and the other one with the low seller type. Suppose the high type plays

$d_1 = h$ in equilibrium and does not want to deviate. In that case, she prefers to be identified as the high type instead of the low type in $t = 2$ at the expense of additional costs Δ. The other seller type who plays $d_1 = l$ in equilibrium, however, faces the same trade-off. As the observed shipper quality is the only determinant for tomorrows reputation, she would prefer to deviate to $d_1 = h$ for the same reason why the other type does not want to deviate from $d_1 = h$. In other words, if one seller prefers to be identified as the high type tomorrow at the cost of Δ, the other type has the same preference. Thus, a deviation is always profitable for one of the seller types and a separating equilibrium cannot exist. The same argument applies to rule out that the second separating equilibrium exists.

In the second step, we verify the existence of pooling equilibria.

Proposition 3.5. *For any Δ,*

(i) pooling on $d_1 = l$ is an equilibrium for any $\mu_1 \in (0, 1)$.

(ii) there exist a cutoff $\hat{\mu}$ such that pooling on $d_1 = h$ is an equilibrium if and only if $\mu_1 \leq \hat{\mu}$. Moreover, $\hat{\mu}$ is increasing in Δ.

When the seller pools on one shipper quality she does not reveal additional information to the buyers about her type with her shipper choice. On path, buyers thus update according to the realization of w. After a deviation, buyers update their off-path belief based on the realization of w.

If the seller pools on the low quality shipper, she saves on costs Δ. A deviation to the high quality shipper can be profitable only if the off-path belief would reveal her with sufficiently high probability as the high type, that is only if $\mu_2(h)$ is relatively high. Conversely, if the off-path belief is small enough, the additional costs for the high quality shipper outweigh the potential expected gains in tomorrow's reputation from deviating and pooling on $d_1 = l$ is an equilibrium. The first part of Proposition 3.5 shows that we can indeed always identify an off-path belief small enough such that there is no profitable deviation from $d_1 = l$ for any $\mu_1 \in (0, 1)$.

Pooling on the high quality shipper, however, is more difficult to sustain. In contrast to pooling on $d_1 = l$, a deviation now saves on additional costs Δ. The expected loss in reputation from a deviation is bounded by the worst off-path belief $\mu_2(l) = 0$ and, moreover, is smaller the lower μ_1 is. In other words, when her reputation is bad already, a seller has little to lose in terms of reputation from deviating to $d_1 = l$ but saves on costs of Δ. If reputation is low enough, a deviation is indeed profitable for any off-path belief $\mu_2(l)$. As a result, we can identify a cutoff on reputation, $\hat{\mu}$, below which pooling on $d_1 = h$ cannot be an equilibrium. For $\mu_1 \geq \hat{\mu}$, however, the pooling equilibrium exists and the intuition from above applies. When reputation is high and the off-path belief $\mu_2(l)$ is small enough, the expected losses in tomorrow's reputation from deviating outweigh the saved costs. For a graphical representation see also Figure 3.4.

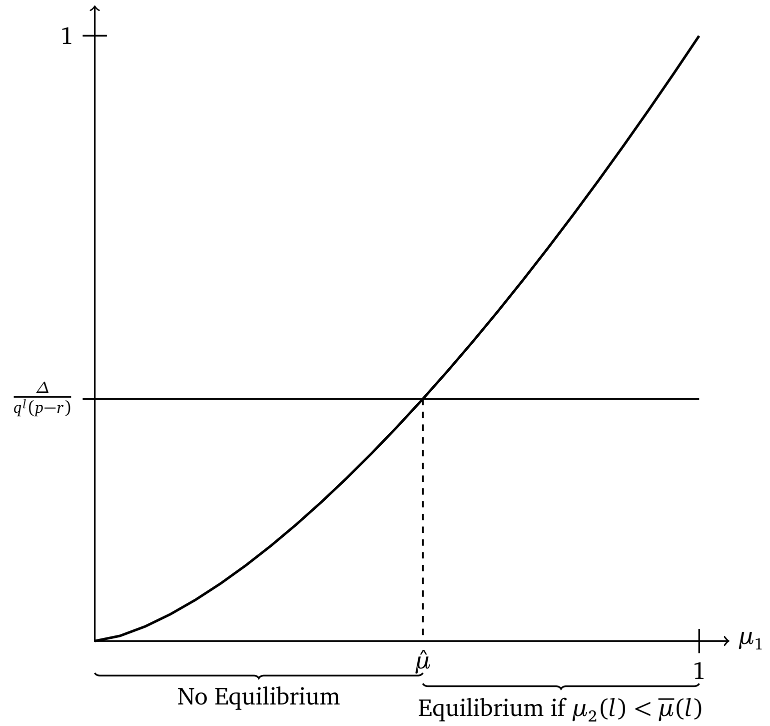

Figure 3.4: Graphical derivation of $\hat{\mu}$ for the pooling equilibrium on $d_1 = h$. For $\mu_1 < \hat{\mu}$, the additional costs for the high quality shipper outweigh the expected losses from deviating. For $\mu_1 \geq \hat{\mu}$, the expected losses from deviating are higher than the additional costs for $d_1 = h$.

3.5 Anonymous Tracking: Buyers observe w^1 and w^2

Nowadays, buyers can usually follow their order online and *track and trace* their parcel at any time. This information impacts how buyers update their belief about the seller's type. In order to see how updating changes and which equilibria result, we consider the case where buyers can track the delivery process in this setup. Instead of observing the total waiting time w only, buyers can now observe both parts of the delivery, w^1 and w^2. In contrast to Section 3.4, however, shippers are not labeled for now. Hence, buyers observe $h^t = \left(w_s, w_s^1, w_s^2\right)_{s=1}^{t-1}$ at the beginning of each period t.

In this setup, there are more cases to distinguish than in the benchmark because there is more information available to the buyers. Instead of observing w only, they observe the tuple (w^1, w^2). Therefore, we have to consider updating according to Bayes' rule for four cases. We denote the updated belief for the second period after observing (w^1, w^2) for a given equilibrium strategy profile σ^* and prior belief μ_1 as $\mu_2\left(w^1, w^2|\mu_1, \sigma^*\right)$,

$$
\begin{aligned}
\mu_2(0,0|\mu_1,\sigma^*) &= \frac{\mu_1 p\sigma_H^*}{\mu_1 p\sigma_H^* + (1-\mu_1)r\sigma_L^*},\\
\mu_2(0,1|\mu_1,\sigma^*) &= \frac{\mu_1 p(1-\sigma_H^*)}{\mu_1 p(1-\sigma_H^*) + (1-\mu_1)r(1-\sigma_L^*)},\\
\mu_2(1,0|\mu_1,\sigma^*) &= \frac{\mu_1(1-p)\sigma_H^*}{\mu_1(1-p)\sigma_H^* + (1-\mu_1)(1-r)\sigma_L^*},\\
\mu_2(1,1|\mu_1,\sigma^*) &= \frac{\mu_1(1-p)(1-\sigma_H^*)}{\mu_1(1-p)(1-\sigma_H^*) + (1-\mu_1)(1-r)(1-\sigma_L^*)}.
\end{aligned}
\tag{3.13}
$$

The high type seller's total payoffs from a strategy σ_H, given equilibrium strategy profile $\sigma^* = \left(\sigma_L^*, \sigma_H^*\right)$ and prior belief μ_1, is then given by

$$
\begin{aligned}
\pi_1^H(\sigma_H;\mu_1,\sigma^*) =& r\sigma_L^* + \mu_1(p\sigma_H^* - r\sigma_L^*) - c(\sigma_H) + rq^l - c(l)\\
&+ q^l(p-r)\Big(p\sigma_H\mu_2(0,0|\mu_1,\sigma^*) + p(1-\sigma_H)\mu_2(0,1|\mu_1,\sigma^*)\\
&+ (1-p)\sigma_H\mu_2(1,0|\mu_1,\sigma^*) + (1-p)(1-\sigma_H)\mu_2(1,1|\mu_1,\sigma^*)\Big).
\end{aligned}
\tag{3.14}
$$

We derive the total profits of a low type seller analogously,

$$
\begin{aligned}
\pi_1^L(\sigma_L;\mu_1,\sigma^*) =& r\sigma_L^* + \mu_1(p\sigma_H^* - r\sigma_L^*) - c(\sigma_L) + rq^l - c(l)\\
&+ q^l(p-r)\Big(r\sigma_L\mu_2(0,0|\mu_1,\sigma^*) + r(1-\sigma_L)\mu_2(0,1|\mu_1,\sigma^*)\\
&+ (1-r)\sigma_L\mu_2(1,0|\mu_1,\sigma^*) + (1-r)(1-\sigma_L)\mu_2(1,1|\mu_1,\sigma^*)\Big).
\end{aligned}
\tag{3.15}
$$

We proceed to characterize the equilibria of this setup. As in the benchmark, the low type seller never chooses the high quality shipper in any equilibrium. In contrast to Proposition 3.2, however, we do not need any assumption on r.

Lemma 3.2. *There exist no equilibrium where the low type seller plays $d_1 = h$.*

The seller only has an influence on the probability distribution of w^2 with her choice of the shipper quality, whereas w^1 is solely determined by her type. Suppose the low type seller chooses the high quality shipper. If the high type seller chooses the low quality shipper, a realization of $w^2 = 0$ is an indicator that the seller is of the low type, and otherwise, if the high type chooses the high quality shipper as well, w^2 does not convey any information about the seller's type. Thus, the low type seller is better off by deviating to the low quality shipper. On the one hand, she saves costs Δ. On the other hand she reduces the probability of realizing $w^2 = 0$ which would (weakly) decrease her reputation of being a high type seller. As both effects go in the same direction, the low type seller never chooses the high quality shipper in any equilibrium.

This result allows us to focus on equilibria where the low type seller plays $d_1 = l$, as in the benchmark. In the following, we are thus concerned with the equilibrium strategy of the high type seller.

Proposition 3.6. *For any $\mu_1 \in (0,1)$, $\sigma^* = \left(q^l, q^l\right)$ is an equilibrium.*

If both seller types pool on the low quality shipper, there is no information available for the buyers from the shipper choice. Buyers only learn from the realization of w^1 and update accordingly. This information, however, cannot be influenced by the sellers directly but is generated passively depending on her type. As a result, both types are better off by saving costs, as the realization of w^2 is ignored by the buyers and investing into a better shipper quality does not pay off.

Proposition 3.7. *There exists a $\overline{\mu} \in (0,1)$ and an interval $D(\mu_1) \equiv \left[\underline{\Delta}(\mu_1), \overline{\Delta}(\mu_1)\right]$ such that*

(i) $\left(\sigma_H^*, \sigma_L^*\right) = \left(q^h, q^l\right)$ *is an equilibrium if and only if $\Delta \in D(\mu_1)$,*

(ii) $D(\mu_1)$ *is non-empty only if $\mu_1 \leq \overline{\mu}$.*

Pooling on the bad shipper quality is not always the unique equilibrium. This result shows that an additional separating equilibrium exists, where the high type chooses the high quality shipper and the low type chooses the low quality shipper if the cost difference Δ is in an intermediate range. This range, however, is non-empty only if μ_1 is sufficiently small. In fact, if we take a closer look at the necessary condition

$$\mu_1^2 p(1-p) q^h (1-q^h) \;\leq\; (1-\mu_1)^2 r(1-r) q^l (1-q^l) \tag{3.16}$$

we can see that only the right-hand side carries the factor r which is smaller than $\hat{r}$ from Proposition 3.2. As the right-hand side needs to be larger than the left-hand side, the restriction on μ_1 to be small is indeed relatively strong. Hence, we can conclude that the separating equilibrium only exists for a relatively narrow set of parameters μ_1 and Δ.
For intuition for the necessary condition (3.16), recall that in this separating equilibrium both types face a trade-off when they choose the shipper quality between the cost difference Δ and the expected change in reputation for $t = 2$. The high type, $\theta = H$, does not deviate if a deviation to the low quality shipper decreases her profits from a lower reputation more than the saved costs of Δ. In contrast, the low type does not deviate if the increase in profits from an expected higher reputation is smaller than the additional costs of Δ. Consequently, for any $\mu_1 \in (0,1)$, this equilibrium can exist only if the reputation changes more from a deviation of the high type than from a deviation of the low type.
Note that the presence of tracking implies that buyers update twice, once after w^1 and the second time after w^2 is observed. Sellers, however, can only affect the realization of w^2 with their choice of d_1. Hence, the high type seller passively increases her reputation in expectation through w^1 whereas the reputation of a low type seller passively decrease in expectation.

When reputation is already high, that is buyers are relatively convinced that the seller is of the high type, reputation increases relatively little after $w^1 = 0$ or $w^2 = 0$ are observed. Moreover, if the seller is indeed of the high type, reputation increases naturally through w^1 in expectation and investing into a good realization of w^2 in addition is not very lucrative. If the seller is of the low type, however, reputation decreases already from the realization of w^1 in expectation. She then has stronger incentives to deviate to $d_1 = h$ in order to counteract the expected losses in reputation from her low type and protect the high reputation from a unfavorable realization of w^2. Hence, $\theta = L$ has more to gain from a deviation than $\theta = H$ has to lose. In such a case, the separation equilibrium cannot exist.
In contrast, when reputation is low and buyers believe with high probability that the seller is of the low type, reputation increases relatively strongly after $w^2 = 0$ or $w^1 = 0$ is observed. If the seller is of the high type, investing in a good realization of w^2 is very lucrative. She can expect reputation to increase strongly after $w^2 = 0$ is observed and her probability for $w^2 = 0$ is relatively high when she does not deviate. This is not the case for the low type, however. As $w^1 = 0$ is already very unlikely to realize, investing into $d_1 = h$ does not payoff as much as for the high type. In such a scenario the separating equilibrium can exist if, additionally, the cost difference is neither to low nor to high to provide the right incentives to each seller type.

3.6 Labeled Tracking: Buyers observe d, w^1, and w^2

To complete our characterization, we are interested in the setup where not only w_1 and w^2 but also the chosen shipper d_t can be observed. Thus, buyers observe $h^t = \left(w_s, w_s^1, w_s^2, d_s\right)_{s=1}^{t-1}$, at the beginning of each period. In other words, buyers can track their deliveries and shippers are labeled.
For this setup, we obtain qualitatively the same equilibria as in chapter 3.4. We start by investigating possible separating equilibria and argue along the lines of the proof of Proposition 3.4 that these cannot exist, even with tracking. Afterwards we consider potential pooling equilibria and show that pooling on the low quality shipper can always be supported as an equilibrium by a non-empty range of off-path beliefs. As before, however, a pooling equilibrium where both seller types choose the high quality shipper exists if and only if the seller's reputation is high enough.

When we consider potential separating equilibria where one type assigns a high quality shipper and the other type a low quality shipper, recall thatobserving the shipper quality perfectly reveals the seller's type to buyers in equilibrium. Hence, reputation in the second period is either 1 or 0, depending on the respective equilibrium strategy and the actually chosen shipper quality. In comparison to chapter 3.4, this setup only differs in the buyers' ability to observe w^1 and w^2 instead of only w. In a separating equilibrium, however, this information is ignored because the observation of d perfectly informs the buyers of the seller's type already. Consequently, we obtain the same result as Proposition 3.4 to which we refer concerning the proof and economic intuition of the following result.

Lemma 3.3. *There exists no equilibrium where the seller types separate.*

As in chapter 3.4, pooling on the low quality shipper is always an equilibrium and pooling on the high quality shipper is an equilibrium if and only if reputation high enough. Recall that in a pooling equilibrium, observing d_1 does not reveal any information and, in contrast to separating strategies, buyers consider their noisy signals, their waiting times, in order to update beliefs. Because updating differs when buyers' can track the shipment and observe w^1 and w^2 instead of w only, the equilibrium conditions on μ_1 for pooling on $d = h$ differ compared to Proposition 3.5.

Proposition 3.8. *For any Δ,*

(i) pooling on $d_1 = l$ is an equilibrium for any $\mu_1 \in (0, 1)$.

(ii) there exists a cutoff $\tilde{\mu}$ such that pooling on $d_1 = h$ is an equilibrium if and only if $\mu_1 \leq \tilde{\mu}$. Moreover, $\tilde{\mu}$ is increasing in Δ.

We are now in a position to analyze how information about the shipper and the shipment affects the welfare of buyers. In the following section, we summarize the insights from the previous sections and compare their equilibria from the buyers' perspective.

3.7 Welfare Comparison

We compare the equilibria of the previous sections from the perspective of a regulator who has the same information as the buyers before the first period starts. We want to analyze which equilibrium benefits buyers the most. The seller's type is her private information in this consideration and the regulator shares the buyers' prior belief μ_1 about her type. We show that pooling on the high quality shipper outperforms the separating strategy $\sigma = \left(q^h, q^l\right)$ as well as pooling on the low quality shipper in terms of buyers' expected welfare. To this end, it is helpful to introduce the following notation. Consider two equilibrium strategies for $t = 1$, σ and σ'. We write $\sigma \succ \sigma'$ if the the buyer' expected welfare is higher in σ than in σ'.

Lemma 3.4. *From the buyers' perspective, the three equilibria*

$$\sigma^* \in \left\{\left(q^l, q^l\right), \left(q^h, q^l\right), \left(q^h, q^h\right)\right\}$$

rank in the following order

$$\left(q^h, q^h\right) \succ \left(q^h, q^l\right) \succ \left(q^l, q^l\right).$$

This ranking is intuitive. Buyers care about their waiting time and a high quality shipper is more likely to reduce the waiting time than a low quality shipper, independent of the seller's type. Consequently, equilibria where the high quality shipper is chosen more often benefit the buyers in expectation. Note, however, that this result depends on the fact that the seller's type is private information and might be different if the regulator knew about the seller's type. As the separating strategy reveals more information about the type to buyers, a regulator might prefer the additional information for a better decision on where to buy in $t = 2$ compared to higher probability of a quick delivery in $t = 1$, especially if the seller is known to be of the low type.

We evaluate the analyzed setups in two steps. First, we compare the setups where buyers cannot track the delivery and investigate whether buyers benefit from labeled shippers, that is, from the ability to identify the shipper's quality upon delivery of the good. Second, we look into potential benefits for buyers from a tracking technology. We analyze the differences between setups with and without tracking for situations where the shipper is or is not labeled. In this comparison, we follow the buyers' ex-ante perspective. When we compare two setups with each other, we focus for any μ_1 on the equilibrium that is most beneficial for buyers if more than one equilibrium exists. Consider some interval $J \subset (0, 1)$ and two setups, A and B. We say that *A dominates B* on J if, for any $\mu_1 \in J$, buyers are at least as well off in the best equilibrium for buyers in A as in the best equilibrium for buyers in B, and strictly better off in the best equilibrium for buyers in A for some $\mu_1 \in J$. Intuitively, when A dominates B in the above sense, and we evaluate the setups before μ_1 is drawn from some continuous distribution with support J, the expected welfare of buyers is strictly higher in A than in B, comparing the best equilibrium for buyers for any μ_1. If A dominates B on $(0, 1)$, we say that *A strictly dominates B*.

We start by focusing on the cases where buyers cannot track the delivery, that is, buyers only observe w but not w^1 and w^2. In particular, we compare the equilibria of the benchmark in Section 3.3 to the equilibria of Section 3.4 in which shippers are labeled. If shippers are labeled, recall from Proposition 3.5 that pooling on $d_1 = l$ is the unique equilibrium for $\mu_1 < \hat{\mu}$. If $\mu_1 \geq \hat{\mu}$, pooling on both $d_1 = l$ and $d_1 = h$ exist. In the benchmark, the high type seller plays $d_1 = h$ only for intermediate reputation levels while the low type seller always plays $d_1 = l$. Thus, if the seller's reputation is low, that is, if $\mu_1 < \hat{\mu}$, labeled shippers never dominates the benchmark. In fact, if r is sufficiently small, the benchmark dominates labeled shippers for $\mu_1 < \hat{\mu}$.

Lemma 3.5. *Suppose r is sufficiently small. Then $\underline{\mu}_h < \hat{\mu}$.*

Intuitively, as r gets smaller, the high type seller's incentives for playing $d_1 = h$ in the benchmark increase through two channels. First, the realization of w gets more informative about

the seller's type. Thus, the seller has higher incentives to increase the probability of $w = 0$ by choosing $d_1 = h$. Second, reputation becomes more valuable for $t = 2$. As both types always play $d_2 = l$, the difference in expected waiting times between the high and the low type seller is $q^l(p - r)$. Consequently, the buyers' belief about the seller's type has a higher impact on profits when this difference increases. Both effects reinforce themselves and increase incentives to hire the high quality shipper in the first period in the benchmark. Investing in a higher reputation is more profitable and reputation is more sensitive towards the realization of w when r gets smaller. Consequently, choosing $d_1 = h$ can be an equilibrium in the benchmark for lower reputation levels the lower r.
In contrast, pooling on $d_1 = h$ is not an equilibrium for small μ_1 even if signals are more informative when r is small. The seller trades off the additional costs of Δ against the loss in reputation from a detected deviation to $d_1 = l$. If reputation is low already, the off-path belief cannot be small enough to deter this deviation even if updating based on the realization of w after a deviation becomes more accurate. As a result, if r is small enough and reputation is relatively low, a high type seller can choose the high quality shipper in the benchmark in equilibrium but when shippers are labeled, pooling on the high quality shipper is not an equilibrium. These insights are captured in Lemma 3.5.
Recall that in the benchmark, the equilibrium where the high type seller plays $d_1 = h$ exists if $\mu_1 \in (\underline{\mu}_h, \overline{\mu}_h)$. With labeled shippers, pooling on $d_1 = l$ is the unique equilibrium for $\mu_1 < \hat{\mu}$. If r is small enough, the interval $(\underline{\mu}_h, \hat{\mu})$ is non-empty by the above Lemma and labeled shippers are dominated by the benchmark for $\mu_1 \in (0, \hat{\mu})$.
Pooling on $d_1 = h$ is an equilibrium with labeled shippers for $\mu_1 \geq \hat{\mu}$, and by Lemma 3.4, is the best equilibrium for buyers. As this equilibrium never exists in the benchmark, the best equilibrium for buyers with labeled shippers is always strictly better than any equilibrium in the benchmark. Thus, labeled shippers dominate the benchmark for $\mu_1 \geq \hat{\mu}$.
We can now summarize the insights from Sections 3.3 and 3.4 as well as Lemma 3.5 without further proof.

Proposition 3.9. *Suppose buyers cannot track the delivery.*
On $(0, \hat{\mu})$*, labeled shippers never dominate the benchmark but are dominated by the benchmark if* r *is sufficiency small.*
On $(\hat{\mu}, 1)$*, labeled shippers dominate the benchmark.*

Next, we are concerned with the role of a tracking technology for buyers' welfare. In the first step, we compare the benchmark with the setup in Section 3.5. Shippers are not labeled in either case but in contrast to the benchmark, buyers can track their delivery in the setup of Section 3.5. With tracking, in the only equilibrium that always exists, the seller always chooses the low quality shipper. If reputation is low and the cost difference is in an intermediate range, there exists an additional equilibrium where the high type seller chooses the high quality shipper instead. Hence, in comparison to the benchmark, buyers are always worse off unless

reputation is low. In that case they are better off with tracking only if the cost difference is in an intermediate range, as described in Proposition 3.7.
Second, we look into the role of tracking in the setups where shippers are labeled and buyers can observe their shipper's quality upon delivery of the good. Hence, we compare Section 3.4, where buyers cannot track their delivery, and Section 3.6, where a tracking technology is available. On the one hand, we can see that in both setups pooling on $d_1 = l$ is always an equilibrium. On the other hand, the cutoff on μ_1 above which pooling on $d_1 = h$ is an equilibrium differs and, in particular, is higher with tracking. In other words, with tracking the seller must have a higher reputation such that pooling on $d_1 = h$ is an equilibrium. The latter insight is proven in the following Lemma.

Lemma 3.6. *For any Δ, $\hat{\mu} < \tilde{\mu}$.*

Pooling on $d_1 = h$ is always an equilibrium without tracking if it is an equilibrium with tracking, but not vice versa. In other words, the set of beliefs μ_1 for which the best equilibrium for buyers exists with tracking is a strict subset of the set of beliefs μ_1 for which this equilibrium exists without tracking. Consequently, tracking is strictly dominated when shippers are labeled. For a better intuition for this result, recall that with tracking buyers ignore the realization of w^2 in a pooling equilibrium and update only through the realization of w^1. The expected gains in reputation from not deviating are then smaller compared to the case where buyers cannot track their delivery and update based on the realization of w.
For the case where shippers are not labeled, pooling on $d_1 = l$ is always an equilibrium. Recall from Proposition 3.7 that $\sigma^* = (q^h, q^l)$ exists if and only if $\mu_1 < \overline{\mu}$, and $\Delta \in D(\mu_1)$. In contrast, $\sigma^* = (q^h, q^l)$ is an equilibrium in the benchmark for $\mu_1 \in (\underline{\mu}_l, \overline{\mu}_l)$.
We can summarize the insights from sections 3.3, 3.6 and Lemma 3.6 without further proof.

Proposition 3.10.

(i) If shippers are labeled, tracking is strictly dominated.

(ii) If shippers are not labeled, tracking is dominated by the benchmark on $(\overline{\mu}, 1)$. On $(0, \overline{\mu})$, tracking dominates the benchmark if and only if $\Delta \in D(\mu_1)$, and is dominated by the benchmark otherwise.

Intuitively, when shippers are not labeled, tracking harms buyers because it removes incentives from the high type seller to protect her reputation with the high quality shipper. When a tracking technology is available to buyers, they update through two channels, using w^1 to draw conclusions about the seller's type, and w^2 to draw conclusions about the shipper's quality. In expectation, a high quality seller passively generates good signals about her type through w^1 and, hence, does not necessarily need to invest into the high quality shipper in addition. This is particularly the case, if the seller types pool on the same shipper quality in equilibrium and w^2 is an uninformative signal. In the benchmark case, buyers update their beliefs only based on the realization of w. The seller has to protect her reputation by decreasing the probability of realizing $w = 1$ by choosing the high quality shipper if reputation is intermediate and, hence, most sensitive. In other words, when buyers cannot observe additional signals about the seller's type through w^1 via a tracking technology, the overall waiting time w is the only piece of information that buyers draw conclusions from about the seller's type. When reputation is sensitive to new information, the high type seller has to invest into the shipper's quality in order to generate information that reflects her high type.

Tracking is beneficial to buyers only if reputation is low and the cost difference is in an intermediate range, as outlined in Proposition 3.7. Then, the high type seller chooses the high quality shipper in an equilibrium which is not the case in the benchmark.

3.8 Concluding Remarks

This paper contributes to the literature on dynamic reputation games. We develop a model to analyze competition between an online seller and an offline retailer which is governed by the online seller's reputation. Consumers are unsure of their waiting time for an ordered good which, is determined by the seller's logistical ability as well as the quality of an assigned shipment service. Our model of reputation captures this specific feature in online markets, and we are able to answer how changing information for buyers about the delivery process affects the online seller's strategy as well as consumer's welfare.

Evaluating the derived setup from the seller's perspective is a natural extension. Although we argue that the ability to *track and trace* is generally harmful for buyers' welfare, we do not explain the emergence of comprehensive tracking abilities in recent years. To this end, it might be insightful to analyze which information the seller prefers buyers to observe about the delivery process.

Moreover, the model offers the opportunity to also consider strategic shippers. As their revenues come from shipping fees, higher demand for the online seller benefits shipping services. Hence, the online seller's reputation is valuable to shippers as well which might create incentives to disguise a seller's low type or to reveal a seller's high type. This could serve as an approach for tracking technologies from the strategic perspective of the shippers.

3.A Appendix

Proof of Proposition 3.1

In chapter 3.2.2 we argue that a buyer is better off buying from S than from R only if his position η is to the right of the cutoff type $\overline{\eta}(\mu_t)$. We complete the proof by showing that a buyer does not benefit from not buying the good at all. Recall that offline retailer R is located at $\eta = 0$. Moreover, the realization of w when a buyer buys from S is independent of his position η. Therefore, a buyer at $\eta = 1$ has the lowest utility from buying the good in any equilibrium. He is furthest from R, always buys from S and is only indifferent between his options if the shipment from S is delayed with probability one. His expected utility in equilibrium is then bounded from above by

$$\mathbb{E}\left[U^S(w)\right] = u - \mathbb{E}[w|\sigma] > u - 1 > 0.$$

The first inequality follows from the fact the probability of realizing $w = 0$ is always strictly positive, the second inequality follows from the assumption (3.3) that $u > 1$. The buyer who is worst off in any equilibrium in expectation always prefers to buy the good and, hence, all buyers never choose to not buy the good in any equilibrium. This completes the proof. □

Proof of Lemma 3.1

Consider the high type seller first. Suppose there exists an equilibrium in which she plays $d_2 = h$. Then her profits are given by

$$r\sigma_2^L + \mu_2\left(p\sigma_2^H - r\sigma_2^L\right) - c(q^h).$$

Note that the game ends after the buyers receive the good in $t = 2$, and they already paid the seller. Even if buyers can detect a deviation from the seller, they would only realize after all interactions with the seller finalized. Consequently, deviating to $d_2 = l$ does not affect the seller's demand but saves on costs. Her profits from that deviation are given by

$$r\sigma_2^L + \mu_2\left(p\sigma_2^H - r\sigma_2^L\right) - c(q^l).$$

The deviation saves her Δ and is strictly profitable. Hence, the high type seller cannot play $d_2 = h$ in any equilibrium. The proof for the low type seller follows from the same argument and completes the proof. □

Proof of Proposition 3.2

We start by determining the low type seller's profits from her played strategy σ_L given the equilibrium strategy σ^*, and her reputation μ_1. We rewrite (3.12) by plugging in (3.10) which yields

$$\begin{aligned}\pi^L(\sigma,\sigma^*;\mu_1) \;&=\; 2rq^l + \mu_1\left(p\sigma_H^* - r\sigma_L^*\right) - c(\sigma_L) - c(q^l)\\ &+ q^l(p-r)\left[r\sigma_L\frac{\mu_1 p\sigma_H^*}{\mu_1 p\sigma_H^* + (1-\mu_1)r\sigma_L^*} + (1-r\sigma_L)\frac{\mu_1(1-p\sigma_H^*)}{\mu_1(1-p\sigma_H^*)+(1-\mu_1)(1-r\sigma_L^*)}\right].\end{aligned} \tag{3.17}$$

Her strategy $\sigma_L = q^h$ is strictly dominated if for any equilibrium strategy σ^*, she prefers $d_1 = l$ over $d_1 = h$, i.e.,

$$\pi^L\left(\left(\sigma_H^*, q^h\right),\sigma^*;\mu_1\right) \;<\; \pi^L\left(\left(\sigma_H^*, q^l\right),\sigma^*;\mu_1\right). \tag{3.18}$$

Using (3.17), we can rewrite this condition to

$$\frac{\Delta}{r(q^h-q^l)} \;>\; q^l(p-r)\left(\frac{\mu_1 p\sigma_H^*}{\mu_1 p\sigma_H^* + (1-\mu_1)r\sigma_L^*} - \frac{\mu_1(1-p\sigma_H^*)}{\mu_1(1-p\sigma_H^*)+(1-\mu_1)(1-r\sigma_L^*)}\right). \tag{3.19}$$

It is instructive to investigate the right-hand-side of (3.19) in more detail. First, it is equal to zero for $\mu_1 = 0$ and $\mu_1 = 1$, and strictly positive for $\mu_1 \in (0,1)$. The first-order derivative with respect to μ_1 is given by

$$\frac{\partial}{\partial\mu_1} \;=\; q^l(p-r)\left(\frac{p\sigma_H^* r\sigma_L^*}{\left(\mu_1 p\sigma_H^* + (1-\mu_1)r\sigma_L^*\right)^2} - \frac{\left(1-p\sigma_H^*\right)\left(1-r\sigma_L^*\right)}{\left(\mu_1(1-p\sigma_H^*)+(1-\mu_1)(1-r\sigma_L^*)\right)^2}\right), \tag{3.20}$$

and the second-order derivative is

$$\begin{aligned}\frac{\partial}{\partial\mu_1} =&q^l(p-r)\Bigg(\frac{p\sigma_H^* r\sigma_L^*\left[-2\left(r\sigma_L^* + \mu_1\left(p\sigma_H^* - r\sigma_L^*\right)\right)\left(p\sigma_H^* - r\sigma_L^*\right)\right]}{\left(\mu_1 p\sigma_H^* + (1-\mu_1)r\sigma_L^*\right)^4}\\ &- \frac{\left(1-p\sigma_H^*\right)\left(1-r\sigma_L^*\right)\left[2\left(1-r\sigma_L^* - \mu_1\left(p\sigma_H^* - r\sigma_L^*\right)\right)\left(p\sigma_H^* - r\sigma_L^*\right)\right]}{\left(\mu_1(1-p\sigma_H^*)+(1-\mu_1)(1-r\sigma_L^*)\right)^4}\Bigg)\\ =&-2q^l(p-r)\left(p\sigma_H^* - r\sigma_L^*\right)\left(\frac{p\sigma_H^* r\sigma_L^*}{\left(\mu_1 p\sigma_H^* + (1-\mu_1)r\sigma_L^*\right)^3} + \frac{\left(1-p\sigma_H^*\right)\left(1-r\sigma_L^*\right)}{\left(\mu_1(1-p\sigma_H^*)+(1-\mu_1)(1-r\sigma_L^*)\right)^3}\right) \;<\; 0.\end{aligned} \tag{3.21}$$

As a result, the right-hand side of (3.19) is strictly concave when $p\sigma_H^* > r\sigma_L^*$, equal to zero for $\mu_1 \in \{0,1\}$, and strictly positive for $\mu_1 \in (0,1)$. We can focus on the case where $p\sigma_H^* > r\sigma_L^*$, because this inequality does not hold for $\sigma^* = \left(q^l, q^h\right)$ only. For $\sigma^* = \left(q^l, q^h\right)$, we can rewrite the right-hand side of (3.19) to

$$\frac{\mu_1(1-\mu_1)q^l(p-r)\left(pq^l-rq^h\right)}{(\mu_1 pq^l+(1-\mu_1)rq^h)(\mu_1(1-pq^l)+(1-\mu_1)(1-rq^h))}.$$

If $pq^l \leq rq^h$, the right-hand side is (weakly) negative and, hence, (3.18) always holds for $\sigma^* = \left(q^l, q^h\right)$. Thus, restricting to $p\sigma_H^* > r\sigma_L^*$ in the following does not exclude any case where (3.18) is not satisfied. Second, the right-hand side is increasing in σ_H^*, as

$$\frac{\partial}{\partial \sigma_H^*} = q^l(p-r)\left(\frac{\mu_1(1-\mu_1)prq^l}{\left(\mu_1 p\sigma_H^*+(1-\mu_1)r\sigma_L^*\right)^2}+\frac{\mu_1(1-\mu_1)p(1-r\sigma_L^*)}{\left(\mu_1(1-p\sigma_H^*)+(1-\mu_1)(1-r\sigma_L^*)\right)^2}\right) > 0,$$

and, decreasing in σ_L^*, as

$$\frac{\partial}{\partial \sigma_L^*} = -q^l(p-r)\left(\frac{\mu_1(1-\mu_1)prq^h}{\left(\mu_1 p\sigma_H^*+(1-\mu_1)r\sigma_L^*\right)^2}+\frac{\mu_1(1-\mu_1)p(1-p\sigma_H^*)}{\left(\mu_1(1-p\sigma_H^*)+(1-\mu_1)(1-r\sigma_L^*)\right)^2}\right) < 0.$$

Consequently, the right-hand side is largest for $\sigma^* = \left(q^h, q^l\right)$. In fact, if (3.19) holds for this equilibrium strategy, the low type seller strictly prefers $d_1 = l$ in any equilibrium and $\sigma_L = q^h$ is thus strictly dominated. Therefore, consider (3.19) for $\sigma^* = \left(q^h, q^l\right)$ in the following. The left-hand side is decreasing in r, because its denominator increases in r. On the right-hand side, both factors are decreasing in r, because

$$\begin{aligned}&\frac{\partial}{\partial r}\left(\frac{\mu_1 pq^h}{\mu_1 pq^h+(1-\mu_1)rq^l}-\frac{\mu_1(1-pq^h)}{\mu_1(1-pq^h)+(1-\mu_1)(1-rq^l)}\right)\\&= -\frac{\mu_1 pq^h(1-pq^h)\left(\mu_1\left(2pq^h-1\right)+2(1-\mu_1)\left(1-2rq^l\right)\right)+(1-\mu_1)^2\left(pq^h(1-rq^l)^2-(1-q^h)(rq^l)^2\right)}{(\mu_1 pq^h+(1-\mu_1)rq^l)^2(\mu_1(1-pq^h)+(1-\mu_1)(1-rq^l))^2} < 0\end{aligned}$$

as $2pq^h > 1$ and $1 > 2rq^l$, implying that the right-hand side is decreasing in r as well. Conversely, both sides are increasing as r decreases. When r goes to zero, the left-hand side diverges to infinity whereas the right-hand-side is bounded from above by

$$\frac{(1-\mu_1)pq^l}{1-\mu_1 pq^h} < 1.$$

Hence, there exists a cutoff $\hat{r} \in (0,1)$ such that for any $r < \hat{r}$, (3.19) holds for any equilibrium strategy σ^*, which establishes the claimed dominance of $\sigma_L = q^h$ for $r \leq \hat{r}$ and completes the proof. □

Proof of Proposition 3.3

Analogously to the proof of Proposition (3.2), we start by determining the high type seller's profits from her played strategy σ_H given the equilibrium strategy σ^*, and her reputation μ_1.

We rewrite (3.11) by plugging in (3.10), and $\sigma_L^* = q^l$ which yields

$$\pi^H(\sigma,\sigma^*;\mu_1) = 2rq^l + \mu_1\left(p\sigma_H^* - rq^l\right) - c(\sigma_H) - c\left(q^l\right) + q^l(p-r)\left[p\sigma_H\frac{\mu_1 p\sigma_H^*}{\mu_1 p\sigma_H^* + (1-\mu_1)rq^l} + (1-p\sigma_H)\frac{\mu_1(1-p\sigma_H^*)}{\mu_1(1-p\sigma_H^*) + (1-\mu_1)(1-rq^l)}\right]. \tag{3.22}$$

It is instructive to derive a conditions for the high type seller's choice in $t = 1$. Fix some equilibrium strategy σ_H^*. She prefers the high quality shipper or the low quality shipper if

$$\frac{\Delta}{p(q^h - q^l)} \leq q^l(p-r)\left(\frac{\mu_1 p\sigma_H^*}{\mu_1 p\sigma_H^* + (1-\mu_1)rq^l} - \frac{\mu_1(1-p\sigma_H^*)}{\mu_1(1-p\sigma_H^*) + (1-\mu_1)(1-rq^l)}\right), \tag{3.23}$$

or

$$\frac{\Delta}{p(q^h - q^l)} \geq q^l(p-r)\left(\frac{\mu_1 p\sigma_H^*}{\mu_1 p\sigma_H^* + (1-\mu_1)rq^l} - \frac{\mu_1(1-p\sigma_H^*)}{\mu_1(1-p\sigma_H^*) + (1-\mu_1)(1-rq^l)}\right), \tag{3.24}$$

respectively. Note that the right-hand sides of (3.23) and (3.24) are identical to the right-hand side of (3.18) from the proof of Proposition 3.2. Hence, they are strictly positive for $\mu_1 \in (0,1)$, equal to zero for $\mu_1 \in \{0,1\}$, strictly concave, and, strictly increasing in σ_H^*.
In order to define the cutoffs Δ_1 and Δ_2, we derive the maximum value of the right-hand side of (3.23) and (3.24). To this end, we solve the first-order condition with respect to μ_1 by setting the first-order derivative (3.20), from the proof of Proposition 3.2, for $\sigma_L^* = q^l$ equal to zero. After some algebra, rewriting yields

$$\mu_1^2 p\sigma_H^*\left(1 - p\sigma_H^*\right) = (1-\mu_1)^2 rq^l\left(1 - rq^l\right),$$

and solving for μ_1 gives

$$\mu_1 = \frac{1}{1 + \sqrt{a\left(\sigma_H^*\right)}} \equiv \mu_1^*\left(\sigma_H^*\right),$$

where

$$a\left(\sigma_H^*\right) = \frac{p\sigma_H^*\left(1 - p\sigma_H^*\right)}{rq^l(1 - rq^l)}.$$

Thus, the maximum of the right-hand side is attained at $\mu_1^*(\sigma_H^*)$. After plugging $\mu_1^*(\sigma_H^*)$ into the right-hand-side, we obtain

$$\frac{\sqrt{a(\sigma_H^*)}(p\sigma_H^* - rq^l)}{\left(p\sigma_H^* + \sqrt{a(\sigma_H^*)}rq^l\right)\left(1 - p\sigma_H^* + \sqrt{a(\sigma_H^*)}(1 - rq^l)\right)}. \tag{3.25}$$

We define Δ_1 as the value of Δ for which (3.23) and (3.24) with $\sigma_H^* = q^h$ evaluated at $\mu_1^*(q^h)$ hold with equality, that is,

$$\frac{\Delta_1}{p(q^h - q^l)} = \frac{\sqrt{a(q^h)}(pq^h - rq^l)}{\left(pq^h + \sqrt{a(q^h)}rq^l\right)\left(1 - pq^h + \sqrt{a(q^h)}(1 - rq^l)\right)},$$

and Δ_2 analogously for $\sigma_H^* = q^l$, that is

$$\frac{\Delta_2}{p(q^h - q^l)} = \frac{\sqrt{a(q^l)}(p - r)}{\left(p + \sqrt{a(q^l)}r\right)\left(1 - pq^l + \sqrt{a(q^l)}(1 - rq^l)\right)}.$$

Recall that the monotonicity of the right-hand side in σ_H^* implies that (3.25) increases in $\sigma_H^* = q^h$ as well. Consequently, it holds that $\Delta_1 > \Delta_2$. We use these cutoffs to distinguish three cases in the following.

(i) $\Delta > \Delta_1$. In this case, (3.24) is satisfied for any $\mu_1 \in (0,1)$ because Δ is high enough that the left-hand-side exceeds the maximum value of the right-hand side. Conversely, (3.23) is never satisfied. As a result, for any equilibrium strategy σ_H^* and any reputation $\mu_1 \in (0,1)$, the seller always prefers $d = l$ over $d = h$ and, hence $\sigma^* = (q^l, q^l)$ is the unique equilibrium for all $\mu_1 \in (0,1)$.

(ii) $\Delta_1 \geq \Delta \geq \Delta_2$. Consider the equilibrium strategy $\sigma^* = (q^l, q^l)$ first. As $\Delta \geq \Delta_2$, (3.24) with $\sigma^* = (q^l, q^l)$ is satisfied for any $\mu_1 \in (0,1)$ and is thus an equilibrium for any $\mu_1 \in (0,1)$. In contrast, as $\Delta_1 \geq \Delta$, (3.23) with $\sigma^* = (q^h, q^l)$ is satisfied at $\mu_1^*(q^h)$ but not for $\mu_1 \in \{0,1\}$, where the right-hand side is equal to zero and the left-hand side is strictly positive. While the left-hand side is constant in μ_1, recall that the right-hand side is strictly concave in μ_1. Therefore, there exist two cutoffs, $\underline{\mu}_h < \overline{\mu}_h$, such that (3.23) with $\sigma^* = (q^h, q^l)$ is satisfied for $\mu_1 \in [\underline{\mu}_h, \overline{\mu}_h]$. Hence, $\sigma^* = (q^h, q^l)$ is an equilibrium for $\mu_1 \in [\underline{\mu}_h, \overline{\mu}_h]$. If $\Delta = \Delta_1$, this equilibrium exists only at $\mu_1 = \mu_1^*(q^h) = \underline{\mu}_h = \overline{\mu}_h$.

(iii) $\Delta_2 > \Delta$. As $\Delta_1 > \Delta$, $\sigma^* = (q^h, q^l)$ is an equilibrium for $\mu_1 \in [\underline{\mu}_h, \overline{\mu}_h]$ as before. From $\Delta_2 > \Delta$ follows that (3.24) with $\sigma^* = (q^l, q^l)$ is not satisfied for $\mu_1^*(q^l)$ anymore. Analogously to the other equilibrium, there exist two cutoffs, $\underline{\mu}_l < \overline{\mu}_l$, such that (3.24) with $\sigma^* = (q^l, q^l)$ is satisfied for $\mu_1 \in [\underline{\mu}_l, \overline{\mu}_l]$. Hence, $\sigma^* = (q^l, q^l)$ is an equilibrium for $\mu_1 \in [\underline{\mu}_l, \overline{\mu}_l]$. Recall that the right-hand side of (3.23) and (3.24) is increasing in σ_H^* which implies that $\underline{\mu}_l < \underline{\mu}_h$ and $\overline{\mu}_h > \overline{\mu}_l$ which completes the proof. □

Proof of Proposition 3.4

There are two potential separating equilibria, $\sigma^* = \left(q^h, q^l\right)$ and $\sigma^* = \left(q^l, q^h\right)$. We show that the first equilibrium can never exist and rule out the second equilibrium with the same argumentation.

Note that in a separating equilibrium with observable actions, the buyers' beliefs are either 1 or 0 after the first period. In the first potential equilibrium $\sigma^* = \left(q^h, q^l\right)$, buyers belief with probability one that the seller is of the high type after $d_1 = h$ is observed and of the low type after $d_1 = l$ is observed, i.e. $\mu_2(h) = 1$ and $\mu_2(l) = 0$. Consider the high type seller. Her profits from playing $d_1 = h$ and $d_1 = l$ given the equilibrium strategy $\sigma^* = \left(q^h, q^l\right)$ and prior belief μ_1 are

$$\begin{aligned}
\pi^H\left(\sigma_H = q^h, \sigma^*; \mu_1\right) &= rq^l + \mu_1\left(pq^h - rq^l\right) - c(q^h) + \pi_2(\sigma^*; \mu_2(h)) \\
&= 2rq^l + \mu_1\left(pq^h - rq^l\right) + q^l(p - r) - c(q^h) - c(q^l), \quad \text{and,} \\
\pi^H\left(\sigma_H = q^l, \sigma^*; \mu_1\right) &= rq^l + \mu_1\left(pq^h - rq^l\right) - c(q^h) + \pi_2(\sigma^*; \mu_2(l)) \\
&= 2rq^l + \mu_1\left(pq^h - rq^l\right) - 2c(q^l),
\end{aligned}$$

respectively. After canceling terms, we can show that she does not want to deviate if and only if

$$\frac{\Delta}{q^l(p - r)} \leq 1. \tag{3.26}$$

If we look at the low type seller in this equilibrium, we see that her profits from playing $d_1 = h$ and $d_1 = l$ are identical to those from the high type seller,

$$\begin{aligned}
\pi^L\left(\sigma_L = q^h, \sigma^*; \mu_1\right) &= rq^l + \mu_1\left(pq^h - rq^l\right) - c(q^h) + \pi_2(\sigma^*; \mu_2(h)) \\
&= 2rq^l + \mu_1\left(pq^h - rq^l\right) + q^l(p - r) - c(q^h) - c(q^l), \quad \text{and,} \\
\pi^L\left(\sigma_L = q^l, \sigma^*; \mu_1\right) &= rq^l + \mu_1\left(pq^h - rq^l\right) - c(q^h) + \pi_2(\sigma^*; \mu_2(l)) \\
&= 2rq^l + \mu_1\left(pq^h - rq^l\right) - 2c(q^l),
\end{aligned}$$

respectively. She does not have a profitable deviation if and only if

$$\frac{\Delta}{q^l(p - r)} \geq 1 \tag{3.27}$$

which directly contradicts (3.26), unless $\frac{\Delta}{q^l(p-r)} = 1$. Recall that $\Delta < \Delta_1 < q^(p - r)$. Therefore, (3.26) and (3.27) always contradict and this separating equilibrium cannot exist.

For the other potential separating equilibrium $\sigma^* = \left(q^l, q^h\right)$, the same argument applies. The only thing that changes is that the conditions for not having a profitable deviation switch. We can hence rule out the second potential separating equilibrium similarly and complete the proof. □

Proof of Proposition 3.5

Before we prove both parts of the Proposition, we abbreviate notation for convenience and write in the first part, where the seller types pool of $d_1 = l$, μ_h for the off-path belief $\mu_2(h)$ after $d_1 = h$ is observed. In the second part, where the seller types pool on $d_1 = h$, we analogously denote the off-path belief $\mu_2(l)$ in brief by μ_l.

(i) Fix $\sigma^* = \left(q^l, q^l\right)$. We need to show that for any $\mu_1 \in (0,1)$, there is a $\overline{\mu}(h) \in [0,1]$ such that for any $\mu_h < \overline{\mu}(h)$ neither seller type has an inventive to deviate from playing $d_1 = l$. We start out by defining how players update after observing d_1 and a realization of w,

$$\begin{aligned}
\mu_2(w = 0, d = l|\mu_1, \sigma^*) &= \frac{\mu_1 p q^l}{\mu_1 p q^l + (1-\mu_1) r q^l} = \frac{\mu_1 p}{\mu_1 p + (1-\mu_1) r}, \\
\mu_2(w = 1, d = l|\mu_1, \sigma^*) &= \frac{\mu_1(1-pq^l)}{\mu_1(1-pq^l) + (1-\mu_1)(1-rq^l)}, \\
\mu_2(w = 0, d = h|\mu_1, \sigma^*) &= \frac{\mu_h p q^h}{\mu_h p q^h + (1-\mu_h) r q^h} = \frac{\mu_h p}{\mu_h p + (1\mu_h) r}, \\
\mu_2(w = 1, d = h|\mu_1, \sigma^*) &= \frac{\mu_h(1-pq^h)}{\mu_h(1-pq^h) + (1-\mu_h)(1-rq^h)}.
\end{aligned} \tag{3.28}$$

Consider the high type seller first. Analogously to the derivation in the benchmark case and by using (3.28), her profits from playing $d_1 = l$ and deviating to $d_1 = h$ are

$$\begin{aligned}
\pi^H\left(\sigma_H = q^l, \sigma^*; \mu_1\right) =& 2rq^l + \mu_1 q^l (p-r) - 2c(q^l) \\
&+ q^l(p-r)\left[pq^l \frac{\mu_1 p}{\mu_1 p + (1-\mu_1) r} + (1-pq^l)\frac{\mu_1(1-pq^l)}{\mu_1(1-pq^l) + (1-\mu_1)(1-rq^l)}\right], \\
\pi^H\left(\sigma_H = q^h, \sigma^*; \mu_1\right) =& 2rq^l + \mu_1 q^l (p-r) - c(q^l) - c(q^h) \\
&+ q^l(p-r)\left[pq^h \frac{\mu_h p}{\mu_h p + (1-\mu_h) r} + (1-pq^h)\frac{\mu_h(1-pq^h)}{\mu_h(1-pq^h) + (1-\mu_h)(1-rq^h)}\right],
\end{aligned}$$

respectively. Then, for any $\mu_1 \in (0,1)$ and some $\mu_h \in [0,1]$, she does not have a profitable deviation to $d_1 = h$ if and only if

$$\begin{aligned}
\frac{\Delta}{q^l(p-r)} \geq& pq^h \frac{\mu_h p}{\mu_h p + (1-\mu_h) r} - pq^l \frac{\mu_1 p}{\mu_1 p + (1-\mu_1) r} \\
&+ (1-pq^h)\frac{\mu_h(1-pq^h)}{\mu_h(1-pq^h) + (1-\mu_h)(1-rq^h)} - (1-pq^l)\frac{\mu_1(1-pq^h)}{\mu_1(1-pq^h) + (1-\mu_1)(1-rq^h)}.
\end{aligned} \tag{3.29}$$

The right-hand side of (3.29) is weakly negative for $\mu_h = 0$ and strictly positive for $\mu_h = 1$. Moreover, it is increasing in μ_h for any $\mu_1 \in (0,1)$, because

$$\frac{\partial}{\partial \mu_h} = pq^h \frac{pr}{(\mu_h p + (1-\mu_h) r)^2} + (1-pq^h)\frac{(1-pq^h)(1-rq^h)}{(\mu_h(1-pq^h) + (1-\mu_h)(1-rq^h))^2} \geq 0.$$

As the left-hand side of (3.29) is constant, this implies that for any μ_1 we can indeed find a cutoff $\overline{\mu}_h^H \in (0,1)$ such that for there is no profitable deviation for $\theta = H$ for any $\mu_h < \overline{\mu}_h^H$.
We proceed analogously for the low type seller. From her profits from both shipper qualities, we formulate the sufficient and necessary condition that she does not have a profitable deviation, which is

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} \geq & rq^h \frac{\mu_h p}{\mu_h p + (1-\mu_h)r} - rq^l \frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} \\ & + (1-rq^h)\frac{\mu_h(1-pq^h)}{\mu_h(1-pq^h)+(1-\mu_h)(1-rq^h)} - (1-rq^l)\frac{\mu_1(1-pq^h)}{\mu_1(1-pq^h)+(1-\mu_1)(1-rq^h)}.\end{aligned} \tag{3.30}$$

As for the high type, the right-hand side of (3.30) is weakly negative for $\mu_h = 0$, strictly positive for $\mu_h = 1$, and increasing in μ_h for any $\mu_1 \in (0,1)$. Consequently, there exist a cutoff $\overline{\mu}_h^L \in (0,1)$ such that for any $\mu_h < \overline{\mu}_h^L$ there is no profitable deviation for $\theta = L$. Whenever $\mu_h < \min\{\overline{\mu}_h^L, \overline{\mu}_h^H\} \equiv \overline{\mu}(h)$, neither type has a profitable deviation and, hence, pooling on $d = l$ is indeed an equilibrium for any $\mu_h < \overline{\mu}(h)$.

(ii) Fix $\sigma^* = (q^h, q^h)$. We proceed similarly to the previous case. Buyers update as follows,

$$\begin{aligned}\mu_2(w=0, d=h|\mu_1,\sigma^*) &= \frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r}, \\ \mu_2(w=1, d=h|\mu_1,\sigma^*) &= \frac{\mu_1(1-pq^h)}{\mu_1(1-pq^h)+(1-\mu_1)(1-rq^h)}, \\ \mu_2(w=0, d=l|\mu_1,\sigma^*) &= \frac{\mu_l p}{\mu_l p + (1-\mu_l)r}, \\ \mu_2(w=1, d=l|\mu_1,\sigma^*) &= \frac{\mu_h(1-pq^l)}{\mu_h(1-pq^l)+(1-\mu_h)(1-rq^l)}.\end{aligned} \tag{3.31}$$

Next, we derive the seller types' profits from either choice of d proceeding similarly as in the previous case using (3.31). Consider $\theta = H$ first,

$$\begin{aligned}\pi^H(\sigma_H = q^h, \sigma^*; \mu_1) =& 2rq^l + \mu_1 q^l(p-r) - c(q^l) - c(q^h) \\ & + q^l(p-r)\left[pq^h \frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} + (1-pq^h)\frac{\mu_1(1-pq^h)}{\mu_1(1-pq^h)+(1-\mu_1)(1-rq^h)}\right], \\ \pi^H(\sigma_H = q^l, \sigma^*; \mu_1) =& 2rq^l + \mu_1 q^l(p-r) - 2c(q^l) \\ & + q^l(p-r)\left[pq^l \frac{\mu_l p}{\mu_l p + (1-\mu_l)r} + (1-pq^l)\frac{\mu_l(1-pq^h)}{\mu_l(1-pq^h)+(1-\mu_l)(1-rq^h)}\right].\end{aligned}$$

She does not have a profitable deviation if and only if

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} \leq & pq^h \frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} - pq^l \frac{\mu_l p}{\mu_l p + (1-\mu_l)r} \\ & + (1-pq^h)\frac{\mu_1(1-pq^h)}{\mu_1(1-pq^h)+(1-\mu_1)(1-rq^h)} - (1-pq^l)\frac{\mu_l(1-pq^l)}{\mu_l(1-pq^l)+(1-\mu_l)(1-rq^l)}.\end{aligned}$$

For $\theta = L$, we proceed analogously. She does not deviate to $d_1 = l$ if and only if

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} \leq & rq^h \frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} - rq^l \frac{\mu_l p}{\mu_l p + (1-\mu_l)r} \\ & + (1-rq^h)\frac{\mu_1(1-pq^h)}{\mu_1(1-pq^h) + (1-\mu_1)(1-rq^h)} - (1-rq^l)\frac{\mu_l(1-pq^l)}{\mu_l(1-pq^l) + (1-\mu_l)(1-rq^l)}.\end{aligned}$$

Summing up, pooling on $d_1 = h$ is an equilibrium if and only if neither type has an incentive to deviate, that is if and only if

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} \leq & xq^h \frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} - xq^l \frac{\mu_l p}{\mu_l p + (1-\mu_l)r} \\ & + (1-xq^h)\frac{\mu_1(1-pq^h)}{\mu_1(1-pq^h) + (1-\mu_1)(1-rq^h)} - (1-xq^l)\frac{\mu_l(1-pq^l)}{\mu_l(1-pq^l) + (1-\mu_l)(1-rq^l)}\end{aligned} \tag{3.32}$$

for $x \in \{r, p\}$. The right-hand side of (3.32) is strictly negative for $\mu_l = 0$ and strictly positive for $\mu_l = 1$. From the first-order derivative with respect to μ_l, we can see that it is strictly decreasing in μ_l,

$$\frac{\partial}{\partial \mu_l} = -\frac{xq^l pr}{(\mu_l p + (1-\mu_l)r)^2} - \frac{(1-xq^l)(1-pq^l)(1-rq^l)}{(\mu_l(1-pq^l) + (1-\mu_l)(1-rq^l))^2} < 0.$$

As a result, a deviation is least profitable if $\mu_l = 0$. If for $\mu_l = 0$, however, one of the types has an incentive to deviate, that is the right-hand side of (3.32) is smaller than $\frac{\Delta}{q^l(p-r)}$, then pooling on $d = h$ cannot be an equilibrium for any $\mu_l \in [0, 1]$.
In the next step, we show that there exists a cutoff $\hat{\mu}$ such that this situation occurs if $\mu_1 < \hat{\mu}$. Consider the right-hand side of (3.32) at $\mu_l = 0$,

$$xq^h \frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} + (1-xq^h)\frac{\mu_1(1-pq^h)}{\mu_1(1-pq^h) + (1-\mu_1)(1-rq^h)}. \tag{3.33}$$

From the first-order derivative with respect to x, we can see that (3.33) is smaller for $\theta = L$, because

$$\frac{\partial}{\partial x} = \frac{\mu_1(1-\mu_1)q^h(p-r)}{(\mu_1 p + (1-\mu_1)r)(\mu_1(1-pq^h) + (1-\mu_1)(1-rq^h))} > 0 \tag{3.34}$$

which implies that $\theta = L$ has higher incentives to deviate if incentives for deviations are highest, i.e., if $\mu_l = 0$. Consequently, if deviating is profitable for $\theta = L$ at $\mu_l = 0$, pooling on $d_1 = h$ cannot be an equilibrium. Conversely, if $\theta = L$ cannot profitably deviate at $\mu_l = 0$, $\theta = H$ can neither, and pooling on the high quality shipper can be an equilibrium for μ_l sufficiently small.

Hence, for any Δ, pooling on $d_1 = h$ is an equilibrium for some $\mu_l \in [0,1]$ if and only if

$$\frac{\Delta}{q^l(p-r)} \leq rq^h \frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} + (1-rq^h)\frac{\mu_1(1-pq^h)}{\mu_1(1-pq^h)+(1-\mu_1)(1-rq^h)}. \tag{3.35}$$

Note that the right-hand side is equal to 0 if $\mu_1 = 0$, is equal to 1 if $\mu_1 = 1$, and is strictly increasing in μ_1, because

$$\frac{\partial}{\partial \mu_1} = rq^h \frac{pr}{(\mu_1 p + (1-\mu_1)r)^2} + (1-rq^h)\frac{(1-pq^h)(1-rq^h)}{(\mu_1(1-pq^h)+(1-\mu_1)(1-rq^h))^2} > 0.$$

Recall that $\frac{\Delta}{q^l(p-r)} \in (0,1)$. We define $\hat{\mu}$ as that value of μ_1 which solves (3.35) with equality, i.e., $\hat{\mu}$ is implicitly defined by

$$\frac{\Delta}{q^l(p-r)} = rq^h \frac{\hat{\mu} p}{\hat{\mu} p + (1-\hat{\mu})r} + (1-rq^h)\frac{\hat{\mu}(1-pq^h)}{\hat{\mu}(1-pq^h)+(1-\hat{\mu})(1-rq^h)}. \tag{3.36}$$

As the right-hand side of (3.35) is strictly increasing in μ_1, we can conclude that pooling on $d_1 = h$ cannot be an equilibrium for $\mu_1 < \hat{\mu}$, but there exist $\mu_l \in [0,1]$ supporting that equilibrium for $\mu_1 \in [\hat{\mu}, 1)$. Moreover, as the left-hand side is increasing in Δ, the cutoff $\hat{\mu}$ is increasing in Δ, which completes the proof. □

Proof of Lemma 3.2

We need to show that whenever the low type seller plays $d_1 = h$ in equilibrium, she has a profitable deviation to $d_1 = l$ regardless of the high type seller's strategy.
Fix an equilibrium strategy $\sigma^* = \left(q^h, \sigma_H^*\right)$ and consider the difference in profits for $\theta = L$ from playing $d = h$ and $d = l$,

$$\pi_1^L\left(\sigma_L = q^h; \mu_1, \sigma^*\right) - \pi_1^L\left(\sigma_L = q^l; \mu_1, \sigma^*\right). \tag{3.37}$$

As we want to show that there is always a profitable deviation, we need to show that (3.37) is strictly negative for any $\sigma_H^* \in \left\{q^l, q^h\right\}$. Using (3.15), we rewrite this condition to

$$\frac{\Delta}{q^l(p-r)} > \left(q^h - q^l\right)\Big[r\left(\mu_2(0,0|\mu_1,\sigma^*) - \mu_2(0,1|\mu_1,\sigma^*)\right) + (1-r)\left(\mu_2(1,0|\mu_1,\sigma^*) - \mu_2(1,1|\mu_1,\sigma^*)\right)\Big]. \tag{3.38}$$

We plug (3.13) into (3.38), and rewrite the inequality after some algebra to

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} >&\mu_1(1-\mu_1)\left(q^h-q^l\right)\left(\sigma_H^*-q^h\right)\Big[\frac{pr^2}{\left(\mu_1 p\sigma_H^*+(1-\mu_1)rq^h\right)\left(\mu_1 p(1-\sigma_H^*)+(1-\mu_1)r(1-q^h)\right)}\\&+\frac{(1-p)(1-r)^2}{\left(\mu_1(1-p)\sigma_H^*+(1-\mu_1)(1-r)q^h\right)\left(\mu_1(1-p)(1-\sigma_H^*)+(1-\mu_1)(1-r)(1-q^h)\right)}\Big].\end{aligned}$$

This inequality is always fulfilled because $\sigma_H^*-q^h\leq 0$. Hence, $\theta=L$ always has a profitable deviation if her equilibrium strategy is $\sigma_L^*=q^h$, which proves the claim. □

Proof of Proposition 3.6

We have to show that neither seller type has a profitable deviation if the equilibrium strategy is $\sigma^*=\left(q^l,q^l\right)$.

The low type seller does not have a profitable deviation if and only if

$$\pi_1^L\left(\sigma_L = q^l;\mu_1,\sigma^*\right)-\pi_1^L\left(\sigma_L = q^h;\mu_1,\sigma^*\right) \geq 0. \tag{3.39}$$

As in the proof of Lemma 3.2, we plug (3.15) into this condition. We can rewrite (3.39) to

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} \geq& \left(q^h-q^l\right)\Big[r\left(\mu_2(0,0|\mu_1,\sigma^*)-\mu_2(0,1|\mu_1,\sigma^*)\right)\\&+(1-r)\left(\mu_2(1,0|\mu_1,\sigma^*)-\mu_2(1,1|\mu_1,\sigma^*)\right)\Big].\end{aligned} \tag{3.40}$$

As $\sigma^*=\left(q^l,q^l\right)$, the realization of w^2 is uninformative and buyers have the same updated belief after observing w^2. Hence,

$$\begin{aligned}\mu_2(0,0|\mu_1,\sigma^*)-\mu_2(0,1|\mu_1,\sigma^*) &= 0, \quad \text{and}\\ \mu_2(1,0|\mu_1,\sigma^*)-\mu_2(1,1|\mu_1,\sigma^*) &= 0.\end{aligned}$$

This implies that the right-hand side of (3.40) is zero while its left-hand side is strictly positive. As a result, the condition is always satisfied, that is, the low type seller does not have a profitable deviation.

The same argument applies to the high type seller. Her profit difference

$$\pi_1^H\left(\sigma_L = q^l;\mu_1,\sigma^*\right)-\pi_1^H\left(\sigma_L = q^h;\mu_1,\sigma^*\right)$$

is non-negative if and only if

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} \geq& \left(q^h-q^l\right)\Big[p\left(\mu_2(0,0|\mu_1,\sigma^*)-\mu_2(0,1|\mu_1,\sigma^*)\right)\\&+(1-p)\left(\mu_2(1,0|\mu_1,\sigma^*)-\mu_2(1,1|\mu_1,\sigma^*)\right)\Big].\end{aligned} \tag{3.41}$$

By the same argument as for $\theta = L$, the right-hand side of (3.41) is zero and, hence, $\theta = H$ does not have a profitable deviation either, which completes the proof. □

Proof of Proposition 3.7

We first show that $\sigma^* = \left(q^h, q^l\right)$ is an equilibrium if Δ is in an intermediate interval of values. The second step is then to show that this interval is non-empty only if μ_1 is sufficiently small.

(i) Fix $\left(\sigma_H^*, \sigma_L^*\right) = \left(q^h, q^l\right)$. We first derive the conditions on Δ that neither player has a profitable deviation. Analogous to the proofs of Lemma 3.2 and Proposition 3.6, the low type seller does not want to deviate to $d_1 = h$ if and only if

$$\frac{\Delta}{q^l(p-r)} \geq \left(q^h - q^l\right)\Big[r\left(\mu_2(0,0|\mu_1,\sigma^*) - \mu_2(0,1|\mu_1,\sigma^*)\right) + (1-r)\left(\mu_2(1,0|\mu_1,\sigma^*) - \mu_2(1,1|\mu_1,\sigma^*)\right)\Big].$$

After using (3.13), we rewrite this condition to

$$\frac{\Delta}{q^l(p-r)} \geq \mu_1(1-\mu_1)\left(q^h-q^l\right)^2\Bigg[\frac{pr^2}{(\mu_1 pq^h + (1-\mu_1)rq^l)(\mu_1 p(1-q^h) + (1-\mu_1)r(1-q^l))} + \frac{(1-p)(1-r)^2}{(\mu_1(1-p)q^h + (1-\mu_1)(1-r)q^l)(\mu_1(1-p)(1-q^h) + (1-\mu_1)(1-r)(1-q^l))}\Bigg] \tag{3.42}$$

We proceed for $\theta = H$ analogously. She does not want to deviate to $d_1 = l$ if and only if

$$\frac{\Delta}{q^l(p-r)} \leq \mu_1(1-\mu_1)\left(q^h-q^l\right)^2\Bigg[\frac{p^2 r}{(\mu_1 pq^h + (1-\mu_1)rq^l)(\mu_1 p(1-q^h) + (1-\mu_1)r(1-q^l))} + \frac{(1-p)^2(1-r)}{(\mu_1(1-p)q^h + (1-\mu_1)(1-r)q^l)(\mu_1(1-p)(1-q^h) + (1-\mu_1)(1-r)(1-q^l))}\Bigg] \tag{3.43}$$

Define the right-hand side of (3.42) as $\frac{\underline{\Delta}(\mu_1)}{q^l(p-r)}$, the right-hand side of (3.43) as $\frac{\overline{\Delta}(\mu_1)}{q^l(p-r)}$, and $D(\mu_1) \equiv \left[\underline{\Delta}(\mu_1), \overline{\Delta}(\mu_1)\right]$. Now the first claim follows immediately. The strategy σ^* is an equilibrium if and only if neither player has an incentive to deviate, that is, if and only if (3.42) and (3.43) hold at the same time. This is equivalent to requiring that $\Delta \in D(\mu_1)$.

(ii) The interval $D(\mu_1)$ is non-empty only if $\underline{\Delta}(\mu_1) \leq \overline{\Delta}(\mu_1)$. By using (3.42) and (3.43), this condition can be expressed equivalently by

$$\frac{pr(p-r)}{(\mu_1 pq^h + (1-\mu_1)rq^l)(\mu_1 p(1-q^h) + (1-\mu_1)r(1-q^l))} \geq \frac{(1-p)(1-r)(p-r)}{(\mu_1(1-p)q^h + (1-\mu_1)(1-r)q^l)(\mu_1(1-p)(1-q^h) + (1-\mu_1)(1-r)(1-q^l))}. \tag{3.44}$$

After multiplying both sides with the other side's denominator and canceling terms, (3.44) reduces to

$$\mu_1^2 p(1-p)q^h(1-q^h) \;\le\; (1-\mu_1)^2 r(1-r)q^l(1-q^l). \tag{3.45}$$

This inequality is satisfied for $\mu_1 = 0$ but not for $\mu_1 = 1$. Moreover, the left-hand side of (3.45) is strictly increasing in μ_1 while its right-hand side is strictly decreasing in μ_1. Therefore, there exists a cutoff $\overline{\mu} \in (0,1)$ such that (3.45) holds if and only if $\mu_1 < \overline{\mu}$, as claimed. □

Proof of Proposition 3.8

As in the proof of Proposition 3.5, we abbreviate notation for convenience and write in the first part, where the seller types pool on $d_1 = l$, μ_h for the off-path belief $\mu_2(h)$ after $d_1 = h$ is observed. In the second part, where the seller types pool on $d_1 = h$, we denote the off-path belief $\mu_2(l)$ by μ_l.

(i) Fix $\sigma^* = \left(q^l, q^l\right)$. As in the proof of Proposition 3.5, we show that for any $\mu_1 \in (0,1)$ there is a $\overline{\mu}(h) \in [0,1]$ such that no seller type has an inventive to deviate. In this setup, buyers can observe both w^1 and w^2. In a pooling equilibrium, however, the realization of w^2 is uninformative and buyers update based on w^1 as well as the observed shipper quality,

$$\begin{aligned}
\mu_2\left(w^1 = 0, w^2, d = l|\mu_1, \sigma^*\right) &= \frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r},\\
\mu_2\left(w^1 = 1, w^2, d = l|\mu_1, \sigma^*\right) &= \frac{\mu_1(1-p)}{\mu_1(1-p) + (1-\mu_1)},\\
\mu_2\left(w^1 = 0, w^2, d = h|\mu_1, \sigma^*\right) &= \frac{\mu_h p}{\mu_h p + (1-\mu_h)r},\\
\mu_2\left(w^1 = 1, w^2, d = h|\mu_1, \sigma^*\right) &= \frac{\mu_h(1-p)}{\mu_h(1-p) + (1-\mu_h)}.
\end{aligned} \tag{3.46}$$

Consider the high type seller first. Analogously to the derivation in the benchmark case and by using (3.46), her profits from playing $d_1 = l$ and $d_1 = h$ are

$$\begin{aligned}
\pi^H\left(\sigma_H = q^l, \sigma^*; \mu_1\right) =& 2rq^l + \mu_1 q^l(p-r) - 2c(q^l)\\
&+ q^l(p-r)\left[p\frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} + (1-p)\frac{\mu_1(1-p)}{\mu_1(1-p) + (1-\mu_1)(1-r)}\right],\\
\pi^H\left(\sigma_H = q^h, \sigma^*; \mu_1\right) =& 2rq^l + \mu_1 q^l(p-r) - c(q^l) - c(q^h)\\
&+ q^l(p-r)\left[p\frac{\mu_h p}{\mu_h p + (1-\mu_h)r} + (1-p)\frac{\mu_h(1-p)}{\mu_h(1-p) + (1-\mu_h)(1-r)}\right].
\end{aligned}$$

respectively. For any $\mu_1 \in (0,1)$ and some $\mu_h \in [0,1]$ there is hence no profitable deviation if and only if

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} \geq & p\left(\frac{\mu_h p}{\mu_h p+(1-\mu_h)r}-\frac{\mu_1 p}{\mu_1 p+(1-\mu_1)r}\right)\\ &+(1-p)\left(\frac{\mu_h(1-p)}{\mu_h(1-p)+(1-\mu_h)(1-r)}-\frac{\mu_1(1-p)}{\mu_1(1-p)+(1-\mu_1)(1-r)}\right).\end{aligned} \tag{3.47}$$

Just as in the proof of Proposition 3.5, the right-hand side of (3.47) is weakly negative for $\mu_h = 0$, strictly positive for $\mu_h = 1$, and increasing in μ_h for any $\mu_1 \in (0,1)$. As the left-hand side of (3.47) is constant, this implies that for any μ_1 we can find a cutoff $\overline{\mu}_h^H \in (0,1)$ such that for there is no profitable deviation for $\theta = H$ for any $\mu_h < \overline{\mu}_h^H$. We proceed analogously for the low type seller. From her profits from both shipper qualities, we setup the sufficient and necessary condition that there is no profitable deviation,

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} \geq & r\left(\frac{\mu_h p}{\mu_h p+(1-\mu_h)r}-\frac{\mu_1 p}{\mu_1 p+(1-\mu_1)r}\right)\\ &+(1-r)\left(\frac{\mu_h(1-p)}{\mu_h(1-p)+(1-\mu_h)(1-r)}-\frac{\mu_1(1-p)}{\mu_1(1-p)+(1-\mu_1)(1-r)}\right).\end{aligned} \tag{3.48}$$

As for the high type, the right-hand side of (3.48) is weakly negative for $\mu_h = 0$, strictly positive for $\mu_h = 1$, and increasing in μ_h for any $\mu_1 \in (0,1)$. Consequently, there exist a cutoff $\overline{\mu}_h^L \in (0,1)$ such that for any $\mu_h < \overline{\mu}_h^L$ there is no profitable deviation. Whenever $\mu_h < \min\left\{\overline{\mu}_h^L, \overline{\mu}_h^H\right\} \equiv \overline{\mu}(h)$, neither type has a profitable deviation and, hence, pooling on $d_1 = l$ is an equilibrium for any $\mu_h < \overline{\mu}(h)$.

(ii) Now fix $\sigma^* = \left(q^h, q^h\right)$. Analogously to the previous case, buyers update as follows,

$$\begin{aligned}\mu_2\left(w^1 = 0, w^2, d = h|\mu_1, \sigma^*\right) &= \frac{\mu_1 p}{\mu_1 p+(1-\mu_1)r},\\ \mu_2\left(w^1 = 1, w^2, d = h|\mu_1, \sigma^*\right) &= \frac{\mu_1(1-p)}{\mu_1(1-p)+(1-\mu_1)},\\ \mu_2\left(w^1 = 0, w^2, d = l|\mu_1, \sigma^*\right) &= \frac{\mu_l p}{\mu_l p+(1-\mu_l)r},\\ \mu_2\left(w^1 = 1, w^2, d = l|\mu_1, \sigma^*\right) &= \frac{\mu_l(1-p)}{\mu_l(1-p)+(1-\mu_l)}.\end{aligned} \tag{3.49}$$

We derive the seller's profits from either choice of d_1 using (3.49), analogously to the previous case. Consider $\theta = H$ first,

$$\begin{aligned}\pi^H\left(\sigma_H = q^h, \sigma^*; \mu_1\right) =& 2rq^l + \mu_1 q^l(p-r) - c(q^l) - c(q^h)\\ &+q^l(p-r)\left[p\frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} + (1-p)\frac{\mu_1(1-p)}{\mu_1(1-p)+(1-\mu_1)(1-r)}\right],\\ \pi^H\left(\sigma_H = q^l, \sigma^*; \mu_1\right) =& 2rq^l + \mu_1 q^l(p-r) - 2c(q^l)\\ &+q^l(p-r)\left[p\frac{\mu_l p}{\mu_l p + (1-\mu_l)r} + (1-p)\frac{\mu_l(1-p)}{\mu_l(1-p)+(1-\mu_l)(1-r)}\right].\end{aligned}$$

There is no profitable deviation if and only if

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} \le& p\left(\frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} - \frac{\mu_l p}{\mu_l p + (1-\mu_l)r}\right)\\ &+(1-p)\left(\frac{\mu_1(1-p)}{\mu_1(1-p)+(1-\mu_1)(1-r)} - \frac{\mu_l(1-p)}{\mu_l(1-p)+(1-\mu_l)(1-r)}\right).\end{aligned}$$

For $\theta = L$, we proceed analogously. She does not deviate to $d_1 = l$ if and only if

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} \le& r\left(\frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} - \frac{\mu_l p}{\mu_l p + (1-\mu_l)r}\right)\\ &+(1-r)\left(\frac{\mu_1(1-p)}{\mu_1(1-p)+(1-\mu_1)(1-r)} - \frac{\mu_l(1-p)}{\mu_l(1-p)+(1-\mu_l)(1-r)}\right).\end{aligned}$$

Summing up, pooling on $d_1 = h$ is an equilibrium if and only if neither type has an incentive to deviate, that is, if and only if

$$\begin{aligned}\frac{\Delta}{q^l(p-r)} \le& x\left(\frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} - \frac{\mu_l p}{\mu_l p + (1-\mu_l)r}\right)\\ &+(1-x)\left(\frac{\mu_1(1-p)}{\mu_1(1-p)+(1-\mu_1)(1-r)} - \frac{\mu_l(1-p)}{\mu_l(1-p)+(1-\mu_l)(1-r)}\right).\end{aligned} \tag{3.50}$$

for $x \in \{r, p\}$. The right-hand side of (3.50) is strictly negative for $\mu_l = 0$, strictly positive for $\mu_l = 1$, and decreasing in μ_l for $x \in \{r, p\}$, because

$$\frac{\partial}{\partial \mu_l} = -\frac{xpr}{(\mu_l p + (1-\mu_l)r)^2} - \frac{(1-x)(1-p)(1-r)}{(\mu_l(1-p)+(1-\mu_l)(1-r))^2} < 0. \tag{3.51}$$

Hence, a deviation is least profitable if $\mu_l = 0$. If, however, for $\mu_l = 0$ one of the types has an incentive to deviate, that is, the right-hand side of (3.50) is smaller than $\frac{\Delta}{q^l(p-r)}$, then pooling on $d_1 = h$ is not an equilibrium.

Analogously to the proof of Proposition 3.5, for any Δ, pooling on $d = h$ is an equilibrium for some $\mu_l \in [0,1]$ if and only if

$$\frac{\Delta}{q^l(p-r)} \leq r\frac{\mu_1 p}{\mu_1 p + (1-\mu_1)r} + (1-r)\frac{\mu_1(1-p)}{\mu_1(1-p)+(1-\mu_1)(1-r)}. \tag{3.52}$$

The right-hand side is equal to zero for $\mu_1 = 0$, is equal to one for $\mu_1 = 1$, and is strictly increasing in μ_1, because

$$\frac{\partial}{\partial \mu_1} = rq^h\frac{pr}{(\mu_1 p + (1-\mu_1)r)^2} + (1-rq^h)\frac{(1-pq^h)(1-rq^h)}{(\mu_1(1-pq^h)+(1-\mu_1)(1-rq^h))^2} > 0.$$

We define $\tilde{\mu}$ as that value of μ_1 which solves (3.52) with equality, i.e., $\tilde{\mu}$ is implicitly defined by

$$\frac{\Delta}{q^l(p-r)} = r\frac{\tilde{\mu} p}{\tilde{\mu} p + (1-\tilde{\mu})r} + (1-r)\frac{\tilde{\mu}(1-p)}{\tilde{\mu}(1-p)+(1-\tilde{\mu})(1-r)}.$$

We conclude that pooling on $d_1 = h$ cannot be an equilibrium for $\mu_1 < \tilde{\mu}$, but there exist $\mu_l \in [0,1]$ supporting that equilibrium otherwise. Moreover, as the left-hand side is increasing in Δ, the cutoff $\tilde{\mu}$ is increasing in Δ, which completes the proof. □

Proof of Lemma 3.4

Recall that $\sigma_2^H = \sigma_2^L = q^l$ in any equilibrium. When we compute the buyers' expected welfare, we compute their expected utilities from both periods separately.
In $t = 2$, the cutoff type $\overline{\eta}(\mu_2)$ determines the buyers' equilibrium behavior. Note that welfare does not depend on where the indifferent buyer buys. As he receives the same utility from buying from either seller, we can w.l.o.g. assume he buys from R. Thus, all types $\eta \in [0, \overline{\eta}(\mu_2)]$ buy from R and receive utility of

$$U^R(\eta) = u - \eta,$$

and all types $(\overline{\eta}(\mu_2), 1]$ buy from S and receive expected utility of

$$U^S(\mu_2) = u - \mathbb{E}[w|\sigma, \mu_2].$$

In any equilibrium, all buyers buy the good. As a result, we only have to consider the buyers' overall expected loss in $t = 2$ from their travel costs and their expected waiting time for a given μ_2. We denote this loss by Σ_2.

$$\Sigma_2 = -\int_0^{\overline{\eta}(\mu_2)} \eta \, \mathrm{d}\eta - \int_{\overline{\eta}(\mu_2)}^1 \mathbb{E}[w|\sigma, \mu_2] \, \mathrm{d}\eta$$

Using (3.6), that is $\overline{\eta}(\mu_2) = \mathbb{E}[w|\sigma, \mu_2]$, we get

$$\begin{aligned}\Sigma_2 &= -\int_0^{\overline{\eta}(\mu_2)} \eta \,\mathrm{d}\eta - \int_{\overline{\eta}(\mu_2)}^{1} \overline{\eta}(\mu_2)\,\mathrm{d}\eta \\ &= -\frac{1}{2}\overline{\eta}(\mu_2)^2 - (1-\overline{\eta}(\mu_2))\overline{\eta}(\mu_2) \\ &= -\overline{\eta}(\mu_2)\left(1 - \frac{1}{2}\overline{\eta}(\mu_2)\right).\end{aligned}$$

From here, it is easy to show that the buyers' welfare in $t = 2$ is decreasing in $\overline{\eta}(\mu_2)$, as

$$\frac{\partial \Sigma_2}{\partial \overline{\eta}(\mu_2)} = -(1 - \overline{\eta}(\mu_2)) < 0. \tag{3.53}$$

The seller always plays $d_2 = l$ so that

$$\overline{\eta}(\mu_2) = 1 - rq^l - \mu_2 q^l (p - r)$$

which is linearly decreasing in μ_2. Together with (3.53) this implies that the buyers' welfare in $t = 2$ is increasing in μ_2.

In the next step, we argue that in $t = 1$, the expected value of μ_2 is independent of the equilibrium strategy in $t = 1$ by using the martingale property of Bayes' rule. Consider any prior belief $\mu_1 \in (0, 1)$ and any equilibrium strategy σ^*. The expected updated belief μ_2 is then equal to the prior belief μ_1, that is

$$\mathbb{E}[\mu_2|\sigma^*] = \mu_1. \tag{3.54}$$

Therefore, in expectation, the equilibrium strategy in $t = 1$ does not affect the buyers' welfare in $t = 2$. Consequently, we can focus on the buyers' welfare in the first period only.

For the same reasons as above, in $t = 1$ the buyers' welfare decreases in $\overline{\eta}(\mu_1)$. For any equilibrium strategy σ^*, $\overline{\eta}(\mu_1)$ is given by

$$1 - r\sigma_L^* - \mu_1\left(p\sigma_H^* - r\sigma_L^*\right).$$

To proof our claim, we are left to show that

$$\overline{\eta}\left(\mu_1|\sigma^* = \left(q^l, q^l\right)\right) > \overline{\eta}\left(\mu_1|\sigma^* = \left(q^h, q^l\right)\right) > \overline{\eta}\left(\mu_1|\sigma^* = \left(q^h, q^h\right)\right). \tag{3.55}$$

For the first first inequality, we plug the equilibrium strategies into (3.55). The claim holds if

$$\begin{aligned}&1 - rq^l - \mu_1 q^l (p - r) > 1 - rq^l - \mu_1\left(pq^h - rq^l\right) \\ \Leftrightarrow &q^h > q^l\end{aligned}$$

which is a true statement. For the second inequality, we get

$$\begin{aligned} 1 - rq^l - \mu_1\left(pq^h - rq^l\right) &> 1 - rq^h - \mu_1 q^h (p - r) \\ \Leftrightarrow \mu_1 &< 1 \end{aligned}$$

which, again, is a true statement and completes the proof. □

Proof of Lemma 3.5

We start by establishing some comparative statics results for the benchmark equilibrium. Recall, that the benchmark equilibria are composed by four cutoffs $\underline{\mu}_h < \underline{\mu}_l < \overline{\mu}_l < \overline{\mu}_h$. The equilibrium $\sigma^* = (q^h, q^l)$ exists for any $\mu_1 \in \left(\underline{\mu}_h, \overline{\mu}_h\right)$.
We show that $\underline{\mu}_h$ decreases when r is decreasing. Moreover, as $r \to 0$, $\underline{\mu}_h$ converges to 0 as well. In the limit, the equilibrium $\sigma^* = (q^h, q^l)$ exists for any $\mu_1 \in \left(0, \overline{\mu}_h\right]$. Recall from the proof of Proposition 3.3 that $\underline{\mu}_h$ and $\overline{\mu}_h$ are implicitly defined by inequality (**??**),

$$\frac{\Delta}{q^l(p-r)} = p\left(q^h - q^l\right)\left(\frac{\mu_1 pq^h}{\mu_1 pq^h + (1-\mu_1)rq^l} - \frac{\mu_1(1-pq^h)}{\mu_1(1-pq^h) + (1-\mu_1)(1-rq^l)}\right), \tag{3.56}$$

which they solve with equality. Note that both left-hand side is decreasing as r decreases. At the same time, the right-hand side increases, because the first-order derivative with respect to r is

$$\frac{\partial}{\partial r} = -\mu_1(1-\mu_1)q^l\left(\frac{pq^h}{(\mu_1 pq^h + (1-\mu_1)rq^l)^2} - \frac{1-pq^h}{(\mu_1(1-pq^h) + (1-\mu_1)(1-rq^l))^2}\right) < 0.$$

Consequently, the cutoff shift to the extremes of the unit interval as r decreases and, in particular, $\underline{\mu}_h$ decreases. For the limit case where $r = 0$, consider (3.56) for $r = 0$, that is,

$$\frac{\Delta}{pq^l} = \frac{(1-\mu_1)p\left(q^h - q^l\right)}{1 - \mu_1 pq^h}. \tag{3.57}$$

In contrast to (3.56), the right-hand side is strictly positive for $\mu_1 = 0$ and strictly decreasing in μ_1. As a result, only two cutoffs, $\overline{\mu}_l < \overline{\mu}_h$, characterize the equilibrium in the limit such that $\underline{\mu}_l = \underline{\mu}_h = 0$. We can conclude that when r gets smaller, the high type seller can play $d_1 = h$ in equilibrium for lower values of μ_1.
Next, we consider the pooling equilibrium on $d_1 = h$ with labeled shippers which exists only for $\mu_1 \geq \hat{\mu}$. Unlike in the benchmark, the comparative statics of $\hat{\mu}$ with respect to r are ambiguous. Recall, that $\hat{\mu}$ is implicitly defined as the solution of (3.36). The left-hand side is identical to that of (3.56) and, hence, decreases as r decreases. In contrast to the benchmark,

however, the derivative of the right-hand side with respect to r is positive,

$$\frac{\partial}{\partial r} = \frac{\mu_1^2 q^h \left(2\mu_1(1-\mu_1)p\left(1-pq^h\right)(p-r) + (1-\mu_1)^2\left(p^2(1-rq^h)^2 - r^2(1-pq^h)^2\right)\right)}{\left(\mu_1 p + (1-\mu_1)r\right)^2\left(\mu_1(1-pq^h) + (1-\mu_1)(1-rq^h)\right)^2} > 0.$$

Thus, as r decreases, both sides of (3.36) decrease and the overall effect on $\hat{\mu}$ is unclear. Still, for any r, the right-hand side is strictly increasing, equals zero for $\mu_1 = 0$, and equals one for $\mu_1 = 1$. As the left-hand side is bounded below by $\frac{\Delta}{pq^l} > 0$, we can conclude $\hat{\mu} > 0$ for any $r \geq 0$.

We can now combine these insights. When r gets small, $\underline{\mu}_h$ gets small as well. As $\hat{\mu}$ is strictly positive for any $r \geq 0$, we can deduce that $\underline{\mu}_h < \hat{\mu}$ for r sufficiently small. This completes the proof. □

Proof of Lemma 3.6

We prove the claim by showing that the right-hand side of (3.35) is strictly higher than the right-hand side of (3.52). In particular, we show that

$$\frac{\mu_1\left(\mu_1 p(1-pq^h) + (1-\mu_1)r(1-rq^h)\right)}{(\mu_1 p + (1-\mu_1)r)(\mu_1(1-pq^h) + (1-\mu_1)(1-rq^h))} > \frac{\mu_1(\mu_1 p(1-p) + (1-\mu_1)r(1-r))}{(\mu_1 p + (1-\mu_1)r)(\mu_1(1-p) + (1-\mu_1)(1-r))}.$$

Multiplying both side with the denominators and canceling terms yields

$$p(1-r)\left(1-pq^h\right) + r(1-p)\left(1-rq^h\right) > p(1-p)\left(1-rq^h\right) + r(1-r)\left(1-pq^h\right). \tag{3.58}$$

After regrouping and canceling terms, (3.58) reduces to

$$1 > q^h,$$

which is always true and, hence, completes the proof. □

Bibliography

Alabi, M. (September 30, 2013). Ed Hardy: From art to infamy and back again. *CNN*, Retrieved from http://edition.cnn.com.

Amaldoss, W. and Jain, S. (2005a). Conspicuous consumption and sophisticated thinking. *Management Science*, 51(10):1449–1466.

Amaldoss, W. and Jain, S. (2005b). Pricing of conspicuous goods: A competitive analysis of social effects. *Journal of Marketing Research*, 42(1):30–42.

Araujo, A. and Scheinkman, J. A. (1977). Smoothness, comparative dynamics, and the turnpike property. *Econometrica*, pages 601–620.

Austen-Smith, D. (1990). Information transmission in debate. *American Journal of Political Science*, 34(1):pp. 124–152.

Austen-Smith, D. and Banks, J. S. (1996). Information aggregation, rationality, and the condorcet jury theorem. *The American Political Science Review*, 90(1):34–45.

Austen-Smith, D. and Feddersen, T. J. (2006). Deliberation, preference uncertainty, and voting rules. *American Political Science Review*, null:209–217.

Bagwell, L. S. and Bernheim, B. D. (1996). Veblen effects in a theory of conspicuous consumption. *The American Economic Review*, pages 349–373.

Bar-Isaac, H. and Tadelis, S. (2008). Seller reputation. *Foundations and Trends in Microeconomics*, 4(4):273–351.

Becker, G. S. (1991). A note on restaurant pricing and other examples of social influences on price. *Journal of Political Economy*, pages 1109–1116.

Becker, G. S. and Murphy, K. M. (1993). A simple theory of advertising as a good or bad. *The Quarterly Journal of Economics*, pages 941–964.

Bensaid, B. and Lesne, J.-P. (1996). Dynamic monopoly pricing with network externalities. *International Journal of Industrial Organization*, 14(6):837–855.

Board, S. (2009). Monopolistic group design with peer effects. *Theoretical Economics*, 4(1):89–125.

Board, S. and Meyer-ter Vehn, M. (2013). Reputation for quality. *Econometrica*, 81(6):2381–2462.

Bognar, K., ter Vehn, M. M., and Smith, L. (2013). A conversational war of attrition.

Bredan, D. (January 1, 2015). How and why Rolex prices have increased over time. *Business Insider*, Retrieved from http://www.businessinsider.com.

Bulow, J. (1986). An economic theory of planned obsolescence. *The Quarterly Journal of Economics*, pages 729–750.

Condorcet, M. d. (1785). Essay on the application of analysis to the probability of majority decisions. *Paris: Imprimerie Royale*.

Coughlan, P. J. (2000). In defense of unanimous jury verdicts: Mistrials, communication, and strategic voting. *The American Political Science Review*, 94(2):375–393.

Cripps, M. W., Mailath, G. J., and Samuelson, L. (2004). Imperfect monitoring and impermanent reputations. *Econometrica*, 72(2):407–432.

Cripps, M. W., Mailath, G. J., and Samuelson, L. (2007). Disappearing private reputations in long-run relationships. *Journal of Economic Theory*, 134(1):287–316.

Damiano, E., Li, H., and Suen, W. (2010). Delay in strategic information aggregation.

Damiano, E., Li, H., and Suen, W. (2012a). Optimal deadlines for agreements. *Theoretical Economics*, 7(2):357–393.

Damiano, E., Li, H., and Suen, W. (2012b). Optimal delay in committees.

Deimen, I., Ketelaar, F., and Thordal-Le Quement, M. (2015). Consistency and communication in committees. *Journal of Economic Theory*, 160, Part A:24 – 35.

Denizet-Lewis, B. (January 24, 2006). The man behind Abercrombie & Fitch. *Salon*, Retrieved from http://www.salon.com.

Dhebar, A. and Oren, S. S. (1985). Optimal dynamic pricing for expanding networks. *Marketing Science*, 4(4):336–351.

Dilmé, F. (2016). Reputation building through costly adjustment. *mimeo*.

Dinlersoz, E. M. and Li, H. (2006). The shipping strategies of internet retailers: Evidence from internet book retailing. *Quantitative marketing and Economics*, 4(4):407–438.

Feddersen, T. J. and Pesendorfer, W. (1998). Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. *The American Political Science Review*, 92(1):23–35.

Gabszewicz, J. J. and Garcia, F. (2008). A note on expanding networks and monopoly pricing. *Economics Letters*, 98(1):9–15.

Goeree, J. K. and Yariv, L. (2011). An experimental study of collective deliberation. *Econometrica*, 79(3):893–921.

Grilo, I., Shy, O., and Thisse, J.-F. (2001). Price competition when consumer behavior is characterized by conformity or vanity. *Journal of Public Economics*, 80(3):385–408.

Hashimoto, K. and Matsubayashi, N. (2014). A note on dynamic monopoly pricing under consumption externalities. *Economics Letters*, 124(1):1–8.

Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *The Review of Economic Studies*, 66(1):169–182.

Hummel, P. (2012). Deliberation in large juries with diverse preferences. *Public Choice*, 150(3-4):595–608.

Jolivet, G., Jullien, B., and Postel-Vinay, F. (2016). Reputation and prices on the e-market: Evidence from a major french platform. *International Journal of Industrial Organization*, 45:59–75.

Katz, M. L. and Shapiro, C. (1985). Network externalities, competition, and compatibility. *The American Economic Review*, 75(3):424–440.

Kreps, D. M., Milgrom, P., Roberts, J., and Wilson, R. (1982). Rational cooperation in the finitely-repeated prisoners' dilemma. Technical report, DTIC Document.

Ladha, K. K. (1992). The condorcet jury theorem, free speech, and correlated votes. *American Journal of Political Science*, 36(3):pp. 617–634.

Lee, I. H. and Mason, R. (2001). Market structure in congestible markets. *European Economic Review*, 45(4–6):809 – 818. 15th Annual Congress of the European Economic Association.

Lee, J. and Liu, Q. (2013). Gambling reputation: Repeated bargaining with outside options. *Econometrica*, 81(4):1601–1672.

Leibenstein, H. (1950). Bandwagon, snob, and veblen effects in the theory of consumers' demand. *The Quarterly Journal of Economics*, pages 183–207.

Lindner, R. and Löhr, J. (August 31, 2015). Der Niedergang der amerikanischen Modeketten. *Frankfurter Allgemeine Zeitung*, Retrieved from http://www.faz.net.

Liu, Q. (2011). Information acquisition and reputation dynamics. *The Review of Economic Studies*, 78(4):1400–1425.

Livingston, J. A. (2005). How valuable is a good reputation? A sample selection model of internet auctions. *Review of Economics and Statistics*, 87(3):453–465.

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

Marsan, J. (June 26, 2015). Fedex vs. UPS vs. USPS: Who's best for shipping? *Fit Small Business*, Retrieved from http://fitsmallbusiness.com.

McKenzie, L. W. (1986). Optimal economic growth, turnpike theorems and comparative dynamics. *Handbook of Mathematical Economics*, 3:1281–1355.

Meirowitz, A. (2007). In defense of exclusionary deliberation: Communication and voting with private beliefs and values. *Journal of Theoretical Politics*, 19(3):301–327.

Melnik, M. I. and Alm, J. (2002). Does a seller's ecommerce reputation matter? evidence from eBay auctions. *The journal of industrial economics*, 50(3):337–349.

Miller, N. R. (1986). Information, electorates, and democracy: some extensions and interpretations of the condorcet jury theorem. *Information pooling and group decision making*, 2:173–192.

Morgan, J. and Stocken, P. C. (2008). Information aggregation in polls. *The American Economic Review*, 98(3):864–896.

Okholm, H. B., Thelle, M. H., Möller, A., Basalisco, B., and Rølmer, S. (2013). E-commerce and delivery. a study of the state of play of EU parcel markets with particular emphasis on e-commerce. *European Commission DG Internal Market and Services, July 15th*.

Pesendorfer, W. (1995). Design innovation and fashion cycles. *The American Economic Review*, pages 771–792.

Piketty, T. (1999). The information-aggregation approach to political institutions. *European Economic Review*, 43(4–6):791 – 800.

Piketty, T. (2000). Voting as communicating. *The Review of Economic Studies*, 67(1):169–191.

Rayo, L. (2013). Monopolistic signal provision. *The BE Journal of Theoretical Economics*, 13(1):27–58.

Robert III., H. M., Evans, W. J., Honemann, D. D., and Balch, T. J. (2000). *Robert's Rules of Order Newley Revised*. Robert's Rules of Order. Perseus Publishing, Cambridge, Massachusetts, 10th edition.

Scheinkman, J. A. (1976). On optimal steady states of n-sector growth models when utility is discounted. *Journal of Economic Theory*, 12(1):11–30.

Smith, C. (My 27, 2014). Chart: Retailers are offering free shipping with more e-commerce transactions. *Business Insider*, Retrieved from http://www.businessinsider.com.

Stock, A. and Balachander, S. (2005). The making of a "hot product": A signaling explanation of marketers' scarcity strategy. *Management Science*, 51(8):1181–1192.

Stokey, N. L., Lucas Jr, R. E., and Prescott, E. (1989). *Recursive Methods in Economic Dynamics*.

Tadelis, S. (1999). What's in a name? Reputation as a tradeable asset. *The American Economic Review*, 89(3):548–563.

Thordal-Le Quement, M. (2013). Communication compatible voting rules. *Theory and Decision*, 74(4):479–507.

Thordal-Le Quement, M. and Yokeeswaran, V. (2015). Subgroup deliberation and voting. *Social Choice and Welfare*, 45(1):155–186.

Veblen, T. (1899). *The theory of the leisure class: An economic study of institutions*. London: Unwin Books.

Wood, M. (June 17, 2015). FedEx vs. UPS vs. USPS: Head-to-head comparison of shipping services. *Business.com*, Retrieved from http://www.business.com.

Young, H. P. (1988). Condorcet's theory of voting. *American Political Science Review*, 82:1231–1244.